

ASDM Book 1: Cisco ASA Series General Operations ASDM Configuration Guide, 7.23

First Published: 2025-01-27

Last Modified: 2025-06-03

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

PREFACE

About This Guide li

Document Objectives li

Related Documentation li

Document Conventions li

Communications, Services, and Additional Information liii

PART I

Getting Started with the ASA 55

CHAPTER 1

Introduction to the Secure Firewall ASA 1

ASDM Requirements 1

ASDM Java Requirements 1

ASDM Compatibility Notes 2

Hardware and Software Compatibility 7

VPN Compatibility 7

New Features 7

New Features in ASA 9.23(1)/ASDM 7.23(1) 8

Firewall Functional Overview 10

Security Policy Overview 11

Permitting or Denying Traffic with Access Rules 11

Applying NAT 11

Protecting from IP Fragments 11

Applying HTTP, HTTPS, or FTP Filtering 11

Applying Application Inspection 11

Applying QoS Policies 12

Applying Connection Limits and TCP Normalization 12

Enabling Threat Detection 12

Firewall Mode Overview 12
Stateful Inspection Overview 13
VPN Functional Overview 14
Security Context Overview 14
ASA Clustering Overview 15
Special, Deprecated, and Legacy Services 15

CHAPTER 2 Getting Started 17
Access the Console for the Command-Line Interface 17
Access the ISA 3000 Console 17
Access the Firepower 1000, and Secure Firewall 1200/3100/4200 Console 18
Access the ASA Console on the Firepower 4100/9300 Chassis 20
Configure ASDM Access 21
Use the Factory Default Configuration for ASDM Access 21
Customize ASDM Access 22
Start ASDM 25
Customize ASDM Operation 26
Install an Identity Certificate for ASDM 26
Increase the ASDM Configuration Memory 26
Increase the ASDM Configuration Memory in Windows 26
Increase the ASDM Configuration Memory in Mac OS 27
Factory Default Configurations 28
Restore the Factory Default Configuration 29
Restore the ASA Virtual Deployment Configuration 31
Firepower 1010 Default Configuration 32
Firepower 1100 Default Configuration 34
Secure Firewall 1210/1220 Default Configuration 35
Secure Firewall 1230/1240/1250 Default Configuration 37
Secure Firewall 3100 Default Configuration 38
Secure Firewall 4200 Default Configuration 39
Firepower 4100/9300 Chassis Default Configuration 40
ISA 3000 Default Configuration 40
ASA Virtual Deployment Configuration 42
Get Started with the Configuration 44

Use the Command Line Interface Tool in ASDM 44
Use the Command Line Interface Tool 44
Show Commands Ignored by ASDM on the Device 45
Apply Configuration Changes to Connections 46

CHAPTER 3 **ASDM Graphical User Interface** 47
About the ASDM User Interface 47
Navigate the ASDM User Interface 50
Menus 51
File Menu 51
View Menu 52
Tools Menu 53
Wizards Menu 54
Window Menu 55
Help Menu 55
Toolbar 56
ASDM Assistant 57
Status Bar 57
Connection to Device 58
Device List 58
Common Buttons 59
Keyboard Shortcuts 59
Find Function in ASDM Panes 61
Find Function in Rule Lists 62
Enable Extended Screen Reader Support 62
Organizational Folder 63
Home Pane (Single Mode and Context) 63
Device Dashboard Tab 63
Device Information Pane 64
Interface Status Pane 66
VPN Sessions Pane 66
Failover Status Pane 66
System Resources Status Pane 66
Traffic Status Pane 66

Latest ASDM Syslog Messages Pane 66
Firewall Dashboard Tab 67
Traffic Overview Pane 68
Top 10 Access Rules Pane 69
Top Usage Status Pane 69
Top Ten Protected Servers Under SYN Attack Pane 69
Top 200 Hosts Pane 70
Top Botnet Traffic Filter Hits Pane 70
Cluster Dashboard Tab 70
Cluster Firewall Dashboard Tab 72
Content Security Tab 73
Intrusion Prevention Tab 74
ASA CX Status Tab 76
ASA FirePower Status Tabs 76
Home Pane (System) 77
Define ASDM Preferences 78
Search with the ASDM Assistant 80
Enable History Metrics 81
Unsupported Commands 81
Ignored and View-Only Commands 81
Effects of Unsupported Commands 82
Discontinuous Subnet Masks Not Supported 82
Interactive User Commands Not Supported by the ASDM CLI Tool 83

CHAPTER 4 **Licenses: Product Authorization Key Licensing for the ISA 3000** 85
About PAK Licenses 85
Preinstalled License 85
Permanent License 85
Time-Based Licenses 86
Time-Based License Activation Guidelines 86
How the Time-Based License Timer Works 86
How Permanent and Time-Based Licenses Combine 86
Stacking Time-Based Licenses 87
Time-Based License Expiration 87

License Notes 88
Secure Client Advantage, Secure Client Premier, and Secure Client VPN Only Licenses 88
Other VPN License 88
Total VPN Sessions Combined, All Types 88
VPN Load Balancing 89
Legacy VPN Licenses 89
Encryption License 89
Total TLS Proxy Sessions 89
VLANs, Maximum 90
Shared Secure Client Premium Licenses (AnyConnect 3 and Earlier) 90
Failover 90
Failover License Requirements and Exceptions 90
How Failover Licenses Combine 91
Loss of Communication Between Failover Units 91
Upgrading Failover Pairs 91
No Payload Encryption Models 91
Licenses FAQ 92
Guidelines for PAK Licenses 93
Configure PAK Licenses 94
Order License PAKs and Obtain an Activation Key 94
Obtain a Strong Encryption License 95
Activate or Deactivate Keys 98
Configure a Shared License (Secure Client 3 and Earlier) 99
About Shared Licenses 99
About the Shared Licensing Server and Participants 99
Communication Issues Between Participant and Server 100
About the Shared Licensing Backup Server 100
Failover and Shared Licenses 101
Maximum Number of Participants 102
Configure the Shared Licensing Server 103
Configure the Shared Licensing Participant and the Optional Backup Server 103
Supported Feature Licenses Per Model 104
Licenses Per Model 104
ISA 3000 License Features 104

Monitoring PAK Licenses 105
Viewing Your Current License 105
Monitoring the Shared License 106
History for PAK Licenses 106

CHAPTER 5 **Licenses: Smart Software Licensing 113**
About Smart Software Licensing 113
Smart Software Licensing for the ASA on the Firepower 4100/9300 Chassis 114
Smart Software Manager and Accounts 114
Offline Management 114
Permanent License Reservation 114
Smart Software Manager On-Prem 119
Licenses and Devices Managed per Virtual Account 120
Evaluation License 120
About Licenses by Type 121
Secure Client Advantage, Secure Client Premier, and Secure Client VPN Only Licenses 121
Other VPN Peers 121
Total VPN Peers Combined, All Types 122
Encryption License 122
Carrier License 123
Total TLS Proxy Sessions 124
VLANs, Maximum 125
Botnet Traffic Filter License 125
Failover or ASA Cluster Licenses 125
Failover Licenses for the ASAv 125
Failover Licenses for the Firepower 1010 125
Failover Licenses for the Firepower 1100 125
Failover Licenses for the Secure Firewall 1210/1220 127
Failover Licenses for the Secure Firewall 1230/1240/1250 128
Failover Licenses for the Secure Firewall 3100 129
Failover Licenses for the Secure Firewall 4200 130
Failover Licenses for the Firepower 4100/9300 131
ASA Cluster Licenses for the Secure Firewall 3100 133
ASA Cluster Licenses for the Secure Firewall 4200 134

ASA Cluster Licenses for the ASAv 135
ASA Cluster Licenses for the Firepower 4100/9300 136
Prerequisites for Smart Software Licensing 137
Smart Software Manager Regular and On-Prem Prerequisites 137
Permanent License Reservation Prerequisites 137
Guidelines for Smart Software Licensing 138
Defaults for Smart Software Licensing 138
ASA Virtual: Configure Smart Software Licensing 139
ASA Virtual: Configure Regular Smart Software Licensing 139
ASA Virtual: Configure Smart Software Manager On-Prem Licensing 142
ASA Virtual: Configure Utility (MSLA) Smart Software Licensing 145
ASA Virtual: Configure Permanent License Reservation 148
Install the ASA Virtual Permanent License 149
(Optional) Return the ASA Virtual Permanent License 153
Switch from Nonflexible PLR Mode to Flexible PLR Mode 154
Switch from Flexible PLR Mode to Nonflexible PLR Mode 155
(Optional) Deregister the ASA Virtual (Regular and On-Prem) 155
(Optional) Renew the ASA Virtual ID Certificate or License Entitlement (Regular and On-Prem) 156
1000/1200/3100/4200: Configure Smart Software Licensing 156
1000/1200/3100/4200: Configure Regular Smart Software Licensing 157
1000/1200/3100/4200: Configure Smart Software Manager On-Prem Licensing 161
1000/1200/3100/4200: Configure Permanent License Reservation 165
Install the 1000/1200/3100/4200 Permanent License 165
(Optional) Return the 1000/1200/3100/4200 Permanent License 168
(Optional) Deregister the 1000/1200/3100/4200 (Regular and On-Prem) 169
(Optional) Renew the 1000/1200/3100/4200 ID Certificate or License Entitlement (Regular and On-Prem) 169
Firepower 4100/9300: Configure Smart Software Licensing 170
Licenses Per Model 171
ASA Virtual 171
Firepower 1010 176
Firepower 1100 Series 176
Secure Firewall 1210 and 1220 178
Secure Firewall 1230, 1240, 1250 179

Secure Firewall 3100 Series 180
Firepower 4100 181
Secure Firewall 4200 Series 182
Firepower 9300 183
License PIDs Per Model 184
Monitoring Smart Software Licensing 189
Viewing Your Current License 189
Viewing Smart License Status 189
Viewing the UDI 189
Smart Software Manager Communication 189
Device Registration and Tokens 190
Periodic Communication with the Smart Software Manager 190
Out-of-Compliance State 190
Smart Call Home Infrastructure 191
Smart License Certificate Management 191
History for Smart Software Licensing 191

CHAPTER 6 **Logical Devices for the Firepower 4100/9300** 197
About Interfaces 197
Chassis Management Interface 197
Interface Types 198
FXOS Interfaces vs. Application Interfaces 199
About Logical Devices 200
Standalone and Clustered Logical Devices 200
Requirements and Prerequisites for Hardware and Software Combinations 200
Guidelines and Limitations for Logical Devices 201
Guidelines and Limitations for Interfaces 201
General Guidelines and Limitations 202
Requirements and Prerequisites for High Availability 202
Configure Interfaces 202
Enable or Disable an Interface 203
Configure a Physical Interface 203
Add an EtherChannel (Port Channel) 204
Configure Logical Devices 206

Add a Standalone ASA 206
Add a High Availability Pair 209
Change an Interface on an ASA Logical Device 210
Connect to the Console of the Application 211
History for Logical Devices 212

CHAPTER 7 **Transparent or Routed Firewall Mode 215**
About the Firewall Mode 215
About Routed Firewall Mode 215
About Transparent Firewall Mode 215
Using the Transparent Firewall in Your Network 216
Management Interface 216
Passing Traffic For Routed-Mode Features 216
About Bridge Groups 217
Bridge Virtual Interface (BVI) 217
Bridge Groups in Transparent Firewall Mode 217
Bridge Groups in Routed Firewall Mode 218
Passing Traffic Not Allowed in Routed Mode 219
Allowing Layer 3 Traffic 219
Allowed MAC Addresses 220
BPDU Handling 220
MAC Address vs. Route Lookups 220
Unsupported Features for Bridge Groups in Transparent Mode 222
Unsupported Features for Bridge Groups in Routed Mode 222
Default Settings 223
Guidelines for Firewall Mode 223
Set the Firewall Mode (Single Mode) 225
Examples for Firewall Mode 226
How Data Moves Through the ASA in Routed Firewall Mode 226
An Inside User Visits a Web Server 226
An Outside User Visits a Web Server on the DMZ 227
An Inside User Visits a Web Server on the DMZ 228
An Outside User Attempts to Access an Inside Host 229
A DMZ User Attempts to Access an Inside Host 230

How Data Moves Through the Transparent Firewall 230
An Inside User Visits a Web Server 231
An Inside User Visits a Web Server Using NAT 232
An Outside User Visits a Web Server on the Inside Network 234
An Outside User Attempts to Access an Inside Host 235
History for the Firewall Mode 236

CHAPTER 8 **Startup Wizard** 239
Access the Startup Wizard 239
Guidelines for the Startup Wizard 239
Startup Wizard Screens 239
Starting Point or Welcome 239
Basic Configuration 240
Interface Screens 240
Outside Interface Configuration (Routed Mode) 240
Outside Interface Configuration - PPPoE (Routed Mode, Single Mode) 240
Management IP Address Configuration (Transparent Mode) 240
Other Interfaces Configuration 240
Static Routes 240
DHCP Server 240
Address Translation (NAT/PAT) 241
Administrative Access 241
IPS Basic Configuration 241
ASA CX Basic Configuration (ASA 5585-X) 241
ASA FirePOWER Basic Configuration 241
Time Zone and Clock Configuration 241
Auto Update Server (Single Mode) 241
Startup Wizard Summary 242
History for the Startup Wizard 242

PART II **High Availability and Scalability** 245

CHAPTER 9 **Multiple Context Mode** 247
About Security Contexts 247

Common Uses for Security Contexts 247
Context Configuration Files 248
Context Configurations 248
System Configuration 248
Admin Context Configuration 248
How the ASA Classifies Packets 248
Valid Classifier Criteria 248
Classification Examples 249
Cascading Security Contexts 251
Management Access to Security Contexts 252
System Administrator Access 252
Context Administrator Access 252
Management Interface Usage 252
About Resource Management 253
Resource Classes 253
Resource Limits 253
Default Class 254
Use Oversubscribed Resources 255
Use Unlimited Resources 255
About MAC Addresses 256
MAC Addresses in Multiple Context Mode 256
Automatic MAC Addresses 256
VPN Support 257
Licensing for Multiple Context Mode 257
Prerequisites for Multiple Context Mode 258
Guidelines for Multiple Context Mode 259
Defaults for Multiple Context Mode 260
Configure Multiple Contexts 260
Enable or Disable Multiple Context Mode 261
Enable Multiple Context Mode 261
Restore Single Context Mode 262
Configure a Class for Resource Management 263
Configure a Security Context 266
Assign MAC Addresses to Context Interfaces Automatically 268

Change Between Contexts and the System Execution Space 269
Manage Security Contexts 269
Remove a Security Context 270
Change the Admin Context 270
Change the Security Context URL 271
Reload a Security Context 272
Reload by Clearing the Configuration 272
Reload by Removing and Re-adding the Context 273
Monitoring Security Contexts 273
Monitor Context Resource Usage 273
View Assigned MAC Addresses 275
View MAC Addresses in the System Configuration 275
View MAC Addresses Within a Context 275
History for Multiple Context Mode 276

CHAPTER 10 **Failover for High Availability** 281
About Failover 281
Failover Modes 281
Failover System Requirements 282
Hardware Requirements 282
Software Requirements 282
License Requirements 283
Failover and Stateful Failover Links 283
Failover Link 283
Stateful Failover Link 285
Avoiding Interrupted Failover and Data Links 285
MAC Addresses and IP Addresses in Failover 287
Stateless and Stateful Failover 288
Stateless Failover 289
Stateful Failover 289
Bridge Group Requirements for Failover 291
Bridge Group Requirements for Appliances, ASAv 291
Failover Health Monitoring 291
Unit Health Monitoring 292

Heartbeat Module Redundancy 292
Interface Monitoring 292
Failover Times 294
Configuration Synchronization 295
Running Configuration Replication 295
File Replication 296
Command Replication 296
Config-Sync Optimization 297
About Active/Standby Failover 298
Primary/Secondary Roles and Active/Standby Status 298
Active Unit Determination at Startup 298
Failover Events 298
About Active/Active Failover 299
Active/Active Failover Overview 299
Primary/Secondary Roles and Active/Standby Status for a Failover Group 300
Active Unit Determination for Failover Groups at Startup 300
Failover Events 300
Licensing for Failover 301
Guidelines for Failover 302
Defaults for Failover 305
Configure Active/Standby Failover 305
Configure Active/Active Failover 306
Configure Optional Failover Parameters 308
Configure Failover Criteria and Other Settings 308
Configure Interface Monitoring and Standby Addresses 310
Configure Support for Asymmetrically Routed Packets (Active/Active Mode) 311
Manage Failover 313
Modify the Failover Setup 314
Force Failover 316
Disable Failover 316
Restore a Failed Unit 317
Re-Sync the Configuration 318
Monitoring Failover 318
Failover Messages 318

Failover Syslog Messages 318
Failover Debug Messages 318
SNMP Failover Traps 319
Monitoring Failover Status 319
System 319
Failover Group 1 and Failover Group 2 320
History for Failover 320

CHAPTER 11 **Failover for High Availability in the Public Cloud 325**
About Failover in the Public Cloud 325
About Active/Backup Failover 326
Primary/Secondary Roles and Active/Backup Status 326
Failover Connection 326
Polling and Hello Messages 326
Active Unit Determination at Startup 327
Failover Events 327
Guidelines and Limitations 328
Licensing for Failover in the Public Cloud 329
Defaults for Failover in the Public Cloud 329
About ASA Virtual High Availability in Microsoft Azure 330
About the Azure Service Principal 331
Configuration Requirements for ASA Virtual High Availability in Azure 332
Configure Active/Backup Failover 333
Configure Optional Failover Parameters 335
Configure Azure Route Tables 335
Manage Failover in the Public Cloud 336
Force Failover 336
Update Routes 336
Validate Azure Authentication 337
Monitor Failover in the Public Cloud 338
Failover Status 338
Failover Messages 338
History for Failover in the Public Cloud 339

CHAPTER 12 **ASA Cluster for the Secure Firewall 3100/4200** 341

About ASA Clustering 341

How the Cluster Fits into Your Network 341

Cluster Members 342

Bootstrap Configuration 342

Control and Data Node Roles 342

Cluster Interfaces 342

Cluster Control Link 342

Configuration Replication 343

ASA Cluster Management 343

Management Network 343

Management Interface 343

Control Unit Management Vs. Data Unit Management 344

Crypto Key Replication 344

ASDM Connection Certificate IP Address Mismatch 344

Inter-Site Clustering 344

Licenses for ASA Clustering 345

Requirements and Prerequisites for ASA Clustering 346

Guidelines for Clustering 348

Configure ASA Clustering 354

Back Up Your Configurations (Recommended) 354

Cable the Units and Configure Interfaces 354

About Cluster Interfaces 354

Cable the Cluster Units and Configure Upstream and Downstream Equipment 361

Configure the Cluster Interface Mode on the Control Unit 361

(Recommended; Required in Multiple Context Mode) Configure Interfaces on the Control Node 364

Create or Join a Cluster Using the High Availability Wizard 368

Customize the Clustering Operation 372

Configure Basic ASA Cluster Parameters 372

Configure Interface Health Monitoring and Auto-Rejoin Settings 378

Configure the Cluster TCP Replication Delay 379

Configure Inter-Site Features 379

Configure Distributed Site-to-Site VPN 382

Manage Cluster Nodes 387
Add a New Data Node from the Control Node 387
Become an Inactive Node 388
Deactivate a Data Node from the Control Node 389
Rejoin the Cluster 390
Leave the Cluster 391
Change the Control Node 392
Execute a Command Cluster-Wide 393
Monitoring the ASA Cluster 394
Monitoring Cluster Status 394
Capturing Packets Cluster-Wide 394
Monitoring Cluster Resources 394
Monitoring Cluster Traffic 394
Monitoring the Cluster Control Link 395
Monitoring Cluster Routing 395
Monitoring Distributed Site-to-Site VPN 395
Configuring Logging for Clustering 395
Troubleshooting Distributed Site-to-Site VPN 396
Examples for ASA Clustering 397
Sample ASA and Switch Configuration 397
ASA Configuration 397
Cisco IOS Switch Configuration 399
Firewall on a Stick 400
Traffic Segregation 402
OTV Configuration for Routed Mode Inter-Site Clustering 404
Examples for Inter-Site Clustering 407
Individual Interface Routed Mode North-South Inter-Site Example 407
Spanned EtherChannel Routed Mode Example with Site-Specific MAC and IP Addresses 408
Spanned EtherChannel Transparent Mode North-South Inter-Site Example 409
Spanned EtherChannel Transparent Mode East-West Inter-Site Example 410
Reference for Clustering 411
ASA Features and Clustering 411
Unsupported Features with Clustering 411
Centralized Features for Clustering 412

Features Applied to Individual Nodes 413
AAA for Network Access and Clustering 413
Connection Settings and Clustering 414
FTP and Clustering 414
ICMP Inspection and Clustering 414
Multicast Routing and Clustering 414
NAT and Clustering 414
Dynamic Routing and Clustering 416
SCTP and Clustering 418
SIP Inspection and Clustering 419
SNMP and Clustering 419
STUN and Clustering 419
Syslog and NetFlow and Clustering 419
Cisco TrustSec and Clustering 419
VPN and Clustering 419
Performance Scaling Factor 420
Control Node Election 420
High Availability Within the Cluster 420
Node Health Monitoring 421
Interface Monitoring 421
Status After Failure 421
Rejoining the Cluster 422
Data Path Connection State Replication 422
How the Cluster Manages Connections 423
Connection Roles 423
New Connection Ownership 425
Sample Data Flow for TCP 425
Sample Data Flow for ICMP and UDP 426
Rebalancing New TCP Connections Across the Cluster 427
History for ASA Clustering for the Secure Firewall 3100/4200 427

CHAPTER 13 **ASA Cluster for the Firepower 4100/9300** 429
About Clustering on the Firepower 4100/9300 Chassis 429
Bootstrap Configuration 430

Cluster Members 430
Cluster Control Link 430
Size the Cluster Control Link 431
Cluster Control Link Redundancy 431
Cluster Control Link Reliability 432
Cluster Control Link Network 432
Cluster Interfaces 432
Connecting to a Redundant Switch System 432
Configuration Replication 433
Secure Firewall ASA Cluster Management 433
Management Network 433
Management Interface 433
Control Unit Management Vs. Data Unit Management 433
Crypto Key Replication 434
ASDM Connection Certificate IP Address Mismatch 434
Spanned EtherChannels (Recommended) 434
Inter-Site Clustering 435
Requirements and Prerequisites for Clustering on the Firepower 4100/9300 Chassis 436
Licenses for Clustering on the Firepower 4100/9300 Chassis 437
Clustering Guidelines and Limitations 439
Configure Clustering on the Firepower 4100/9300 Chassis 443
FXOS: Add an ASA Cluster 444
Create an ASA Cluster 444
Add More Cluster Members 450
ASA: Change the Firewall Mode and Context Mode 452
ASA: Configure Data Interfaces 452
ASA: Customize the Cluster Configuration 454
Configure Basic ASA Cluster Parameters 454
Configure Interface Health Monitoring and Auto-Rejoin Settings 457
Configure the Cluster TCP Replication Delay 459
Configure Inter-Site Features 459
Configure Distributed Site-to-Site VPN 462
FXOS: Remove a Cluster Node 468
ASA: Manage Cluster Members 469

Become an Inactive Member 469
Deactivate a Data Unit from the Control Unit 470
Rejoin the Cluster 471
Change the Control Unit 472
Execute a Command Cluster-Wide 472
ASA: Monitoring the ASA Cluster on the Firepower 4100/9300 chassis 474
Monitoring Cluster Status 474
Capturing Packets Cluster-Wide 474
Monitoring Cluster Resources 474
Monitoring Cluster Traffic 474
Monitoring the Cluster Control Link 475
Monitoring Cluster Routing 475
Monitoring Distributed S2S VPN 475
Configuring Logging for Clustering 475
Troubleshooting Distributed S2S VPN 476
Examples for ASA Clustering 477
Firewall on a Stick 478
Traffic Segregation 479
OTV Configuration for Routed Mode Inter-Site Clustering 479
Examples for Inter-Site Clustering 482
Spanned EtherChannel Routed Mode Example with Site-Specific MAC and IP Addresses 482
Spanned EtherChannel Transparent Mode North-South Inter-Site Example 483
Spanned EtherChannel Transparent Mode East-West Inter-Site Example 484
Reference for Clustering 485
ASA Features and Clustering 485
Unsupported Features with Clustering 485
Centralized Features for Clustering 486
Features Applied to Individual Units 487
AAA for Network Access and Clustering 487
Connection Settings 488
FTP and Clustering 488
ICMP Inspection 488
Multicast Routing and Clustering 488
NAT and Clustering 488

Dynamic Routing and Clustering 490
SCTP and Clustering 491
SIP Inspection and Clustering 491
SNMP and Clustering 491
STUN and Clustering 491
Syslog and NetFlow and Clustering 491
Cisco TrustSec and Clustering 491
VPN and Clustering on the Secure Firewall eXtensible Operating System (FXOS) Chassis 491
Performance Scaling Factor 492
Control Unit Election 492
High Availability Within the Cluster 493
Chassis-Application Monitoring 493
Unit Health Monitoring 493
Interface Monitoring 493
Decorator Application Monitoring 494
Status After Failure 494
Rejoining the Cluster 494
Data Path Connection State Replication 495
How the Cluster Manages Connections 495
Connection Roles 495
New Connection Ownership 497
Sample Data Flow for TCP 497
Sample Data Flow for ICMP and UDP 498
Rebalancing New TCP Connections Across the Cluster 499
History for ASA Clustering on the Firepower 4100/9300 500

CHAPTER 14 **ASA Cluster for the ASA Virtual for the Private Cloud** 507
About ASA Virtual Clustering 507
How the Cluster Fits into Your Network 508
Cluster Nodes 508
Bootstrap Configuration 508
Control and Data Node Roles 508
Individual Interfaces 509
Policy-Based Routing 509

Equal-Cost Multi-Path Routing 510
Cluster Control Link 510
Cluster Control Link Traffic Overview 511
Cluster Control Link Failure 511
Configuration Replication 511
ASA Virtual Cluster Management 511
Management Network 511
Management Interface 511
Control Node Management Vs. Data Node Management 512
Crypto Key Replication 512
ASDM Connection Certificate IP Address Mismatch 512
Inter-Site Clustering 513
Licenses for ASA Virtual Clustering 513
Requirements and Prerequisites for ASA Virtual Clustering 513
Guidelines for ASA Virtual Clustering 514
Configure the ASA Virtual Clustering Using a Day0 Configuration 515
Configure ASA Virtual Clustering after Deployment 518
Back Up Your Configurations (Recommended) 518
Configure Interface Settings 518
Configure the Cluster Interface Mode on the Control Node 518
Configure the Cluster Control Link on the Control Node 520
Configure Individual Interfaces 522
Create or Join a Cluster Using the High Availability Wizard 524
Customize the Clustering Operation 527
Configure Basic ASA Cluster Parameters 527
Configure Interface Health Monitoring and Auto-Rejoin Settings 530
Configure the Cluster TCP Replication Delay 531
Configure Inter-Site Features 532
Configure Cluster Flow Mobility 532
Manage Cluster Nodes 535
Add a New Data Node from the Control Node 535
Become an Inactive Node 536
Deactivate a Data Node from the Control Node 537
Rejoin the Cluster 537

Leave the Cluster 538
Change the Control Node 539
Execute a Command Cluster-Wide 539
Monitoring the ASA Virtual Cluster 540
Monitoring Cluster Status 540
Capturing Packets Cluster-Wide 540
Monitoring Cluster Resources 541
Monitoring Cluster Traffic 541
Monitoring the Cluster Control Link 541
Monitoring Cluster Routing 541
Configuring Logging for Clustering 542
Examples for ASA Virtual Clustering 542
Individual Interface Routed Mode North-South Inter-Site Example 542
Reference for Clustering 543
ASA Features and Clustering 543
Unsupported Features with Clustering 543
Centralized Features for Clustering 544
Features Applied to Individual Nodes 545
AAA for Network Access and Clustering 545
Connection Settings and Clustering 546
Dynamic Routing and Clustering 546
FTP and Clustering 547
ICMP Inspection and Clustering 547
Multicast Routing and Clustering 547
NAT and Clustering 547
SCTP and Clustering 549
SIP Inspection and Clustering 549
SNMP and Clustering 549
STUN and Clustering 549
Syslog and NetFlow and Clustering 549
Cisco TrustSec and Clustering 550
VPN and Clustering 550
Performance Scaling Factor 550
Control Node Election 550

High Availability Within the ASA Virtual Cluster 551
Node Health Monitoring 551
Interface Monitoring 551
Status After Failure 552
Rejoining the Cluster 552
Data Path Connection State Replication 552
How the ASA Virtual Cluster Manages Connections 553
Connection Roles 553
New Connection Ownership 555
Sample Data Flow for TCP 555
Sample Data Flow for ICMP and UDP 556
Rebalancing New TCP Connections Across the Cluster 557
History for ASA Virtual Clustering 557

PART III **Interfaces** 559

CHAPTER 15 **Basic Interface Configuration** 561
About Basic Interface Configuration 561
Auto-MDI/MDIX Feature 561
Management Interface 562
Management Interface Overview 562
Management Slot/Port Interface 562
Use Any Interface for Management-Only Traffic 562
Management Interface for Transparent Mode 563
Guidelines for Basic Interface Configuration 563
Default Settings for Basic Interface Configuration 564
Enable the Physical Interface and Configure Ethernet Parameters 565
Enable Jumbo Frame Support (ASA Virtual, ISA 3000) 567
Manage the Network Module for the Secure Firewall 3100/4200 567
Configure Breakout Ports 568
Add a Network Module 568
Hot Swap the Network Module 569
Replace the Network Module with a Different Type 570
Remove the Network Module 571

Examples for Basic Interfaces 571
Physical Interface Parameters Example 571
Multiple Context Mode Example 572
History for Basic Interface Configuration 572

CHAPTER 16 **Basic Interface Configuration for Firepower 1010 and Secure Firewall 1210/1220 Switch Ports** 575
About Switch Ports 575
Understanding Switch Ports and Interfaces 575
Auto-MDI/MDIX Feature 576
Guidelines and Limitations for Switch Ports 576
Configure Switch Ports and Power Over Ethernet 578
Configure a VLAN Interface 578
Configure Switch Ports as Access Ports 579
Configure Switch Ports as Trunk Ports 580
Configure Power Over Ethernet 581
Monitoring Switch Ports 582
History for Switch Ports 583

CHAPTER 17 **EtherChannel Interfaces** 585
About EtherChannels 585
About EtherChannels 585
Channel Group Interfaces 586
Connecting to an EtherChannel on Another Device 586
Link Aggregation Control Protocol 587
Load Balancing 587
EtherChannel MAC Address 588
Guidelines for EtherChannels 588
Default Settings for EtherChannels Interfaces 590
Configure an EtherChannel 590
Add Interfaces to the EtherChannel 590
Customize the EtherChannel 592
Examples for EtherChannels 594
History for EtherChannels 595

CHAPTER 18 **Loopback Interfaces** 597

About Loopback Interfaces 597

Guidelines for Loopback Interfaces 598

Configure a Loopback Interface 598

Rate-Limit Traffic to the Loopback Interface 599

History for Loopback Interfaces 603

CHAPTER 19 **VLAN Subinterfaces** 605

About VLAN Subinterfaces 605

Licensing for VLAN Subinterfaces 605

Guidelines and Limitations for VLAN Subinterfaces 606

Default Settings for VLAN Subinterfaces 607

Configure VLAN Subinterfaces and 802.1Q Trunking 607

Examples for VLAN Subinterfaces 609

History for VLAN Subinterfaces 610

CHAPTER 20 **VXLAN Interfaces** 611

About VXLAN Interfaces 611

Encapsulation 611

VXLAN Tunnel Endpoint 612

VTEP Source Interface 612

VNI Interfaces 612

VXLAN Packet Processing 613

Peer VTEP 614

VXLAN Use Cases 614

VXLAN Bridge or Gateway Overview 615

VXLAN Bridge 615

VXLAN Gateway (Routed Mode) 615

Router Between VXLAN Domains 616

AWS Gateway Load Balancer and Geneve Single-Arm Proxy 617

AWS Gateway Load Balancer and Geneve Dual-Arm Proxy 618

Azure Gateway Load Balancer and Paired Proxy 619

Requirements and Prerequisites for VXLAN Interfaces 620

Guidelines for VXLAN Interfaces 620
Default Settings for VXLAN Interfaces 621
Configure VXLAN Interfaces 622
Configure the VTEP Source Interface 622
Configure the VNI Interface 623
Configure Geneve Interfaces 624
Configure the VTEP Source Interface for Geneve 624
Configure the VNI Interface for Geneve 625
Allow Gateway Load Balancer Health Checks 626
Examples for VXLAN Interfaces 626
Transparent VXLAN Gateway Example 627
VXLAN Routing Example 629
History for VXLAN Interfaces 630

CHAPTER 21 **Routed and Transparent Mode Interfaces** 633
About Routed and Transparent Mode Interfaces 633
Security Levels 633
Dual IP Stack (IPv4 and IPv6) 634
31-Bit Subnet Mask 634
31-Bit Subnet and Clustering 634
31-Bit Subnet and Failover 634
31-Bit Subnet and Management 635
31-Bit Subnet Unsupported Features 635
Guidelines and Limitations for Routed and Transparent Mode Interfaces 635
Configure Routed Mode Interfaces 637
Configure General Routed Mode Interface Parameters 637
Configure PPPoE 640
Configure Bridge Group Interfaces 641
Configure the Bridge Virtual Interface (BVI) 641
Configure General Bridge Group Member Interface Parameters 642
Configure a Management Interface for Transparent Mode 643
Configure IPv6 Addressing 645
About IPv6 645
IPv6 Addressing 645

Modified EUI-64 Interface IDs 646
Configure the IPv6 Prefix Delegation Client 646
About IPv6 Prefix Delegation 646
Enable the IPv6 Prefix Delegation Client 648
Configure a Global IPv6 Address 649
(Optional) Configure the Link-Local Addresses Automatically 651
(Optional) Configure the Link-Local Addresses Manually 652
Configure IPv6 Neighbor Discovery 653
View and Clear Dynamically Discovered Neighbors 656
Monitoring Routed and Transparent Mode Interfaces 656
Interface Statistics and Information 656
DHCP Information 657
Static Route Tracking 657
PPPoE 657
Dynamic ACLs 657
Examples for Routed and Transparent Mode Interfaces 658
Transparent Mode Example with 2 Bridge Groups 658
Switched LAN Segment Example with 2 Bridge Groups 659
History for Routed and Transparent Mode Interfaces 661

CHAPTER 22 **Advanced Interface Configuration** 665
About Advanced Interface Configuration 665
About MAC Addresses 665
Default MAC Addresses 665
Automatic MAC Addresses 666
About the MTU 667
Path MTU Discovery 667
Default MTU 667
MTU and Fragmentation 667
MTU and Jumbo Frames 668
About the TCP MSS 668
Default TCP MSS 668
Suggested Maximum TCP MSS Setting 668
Inter-Interface Communication 669

Intra-Interface Communication (Routed Firewall Mode) 669
Automatically Assign MAC Addresses 669
Configure the Manual MAC Address, MTU, and TCP MSS 670
Allow Same Security Level Communication 672
Monitoring the ARP and MAC Address Table 672
History for Advanced Interface Configuration 672

CHAPTER 23 Traffic Zones 675
About Traffic Zones 675
Non-Zoned Behavior 675
Why Use Zones? 675
Asymmetric Routing 676
Lost Route 676
Load Balancing 677
Per-Zone Connection and Routing Tables 678
ECMP Routing 678
Non-Zoned ECMP Support 678
Zoned ECMP Support 679
How Connections Are Load-Balanced 679
Falling Back to a Route in Another Zone 679
Interface-Based Security Policy 679
Supported Services for Traffic Zones 679
Security Levels 680
Primary and Current Interface for the Flow 680
Joining or Leaving a Zone 680
Intra-Zone Traffic 680
To- and From-the-Box Traffic 681
Overlapping IP Addresses Within a Zone 681
Prerequisites for Traffic Zones 681
Guidelines for Traffic Zones 682
Configure a Traffic Zone 684
Monitoring Traffic Zones 684
Zone Information 684
Zone Connections 685

Zone Routing 685
Example for Traffic Zones 686
History for Traffic Zones 689

PART IV **Basic Settings** 691

CHAPTER 24 **Basic Settings** 693
Set the Hostname, Domain Name, and the Enable and Telnet Passwords 693
Set the Date and Time 695
Set the Date and Time Using an NTP Server 695
Set the Date and Time Manually 696
Configure Precision Time Protocol (ISA 3000) 697
Configure the Master Passphrase 698
Add or Change the Master Passphrase 699
Disable the Master Passphrase 700
Configure the DNS Servers 701
Configure the Hardware Bypass and Dual Power Supply (Cisco ISA 3000) 704
Adjust ASP (Accelerated Security Path) Performance and Behavior 705
Choose a Rule Engine Transactional Commit Model 706
Enable ASP Load Balancing 707
Monitoring the DNS Cache 707
History for Basic Settings 708

CHAPTER 25 **DHCP and DDNS Services** 713
About DHCP and DDNS Services 713
About the DHCPv4 Server 713
DHCP Options 713
About the DHCPv6 Stateless Server 714
About the DHCP Relay Agent 714
DHCP Relay Server Support on VTI 714
Guidelines for DHCP and DDNS Services 715
Configure the DHCP Server 717
Enable the DHCPv4 Server 717
Configure Advanced DHCPv4 Options 719

Configure the DHCPv6 Stateless Server 720
Configure the DHCP Relay Agent 721
Configure Dynamic DNS 722
Monitoring DHCP and DDNS Services 726
Monitoring DHCP Services 726
Monitoring DDNS Status 727
History for DHCP and DDNS Services 728

CHAPTER 26 **Digital Certificates** 731
About Digital Certificates 731
Public Key Cryptography 732
Certificate Scalability 732
Key Pairs 733
Trustpoints 733
Certificate Enrollment 734
Proxy for SCEP Requests 734
Revocation Checking 734
Supported CA Servers 734
CRLs 735
OCSP 736
Certificates and User Login Credentials 737
User Login Credentials 737
Certificates 737
Guidelines for Digital Certificates 738
Configure Digital Certificates 741
Configure Reference Identities 741
How to Set Up Specific Certificate Types 742
Identity Certificates 743
Add or Import an Identity Certificate 743
Export an Identity Certificate 748
Generate a Certificate Signing Request 749
Install Identity Certificates 750
CA Certificates 751
Add or Install a CA Certificate 751

Configure CA Certificates for Revocation 752
Configure CRL Retrieval Policy 753
Configure CRL Retrieval Methods 753
Configure OCSP Rules 754
Configure Advanced CRL and OCSP Settings 755
CA Server Management 756
Permit Weak Crypto for CA Certificates 756
Code Signer Certificate 757
Import a Code Signer Certificate 757
Export a Code Signer Certificate 757
Set a Certificate Expiration Alert (for Identity or CA Certificates) 758
Monitoring Digital Certificates 758
History for Certificate Management 759

CHAPTER 27 **ARP Inspection and the MAC Address Table** 763
About ARP Inspection and the MAC Address Table 763
ARP Inspection for Bridge Group Traffic 763
MAC Address Table 764
Default Settings 764
Guidelines for ARP Inspection and the MAC Address Table 764
Configure ARP Inspection and Other ARP Parameters 765
Add a Static ARP Entry and Customize Other ARP Parameters 765
Enable ARP Inspection 766
Customize the MAC Address Table for Bridge Groups 767
Add a Static MAC Address for Bridge Groups 767
Configure MAC Address Learning 767
History for ARP Inspection and the MAC Address Table 768

PART V **IP Routing** 771

CHAPTER 28 **Routing Overview** 773
Path Determination 773
Supported Route Types 774
Static Versus Dynamic 774

Single-Path Versus Multipath 774
Flat Versus Hierarchical 774
Link-State Versus Distance Vector 775
Supported Internet Protocols for Routing 775
Routing Table 776
How the Routing Table Is Populated 776
Administrative Distances for Routes 776
Backup Dynamic and Floating Static Routes 778
How Forwarding Decisions Are Made 778
Dynamic Routing and Failover 778
Dynamic Routing and Clustering 779
Dynamic Routing in Spanned EtherChannel Mode 779
Dynamic Routing in Individual Interface Mode 780
Dynamic Routing in Multiple Context Mode 781
Route Resource Management 781
Routing Table for Management Traffic 781
Management Interface Identification 782
Equal-Cost Multi-Path (ECMP) Routing 783
Disable Proxy ARP Requests 783
Display the Routing Table 784
History for Route Overview 784

CHAPTER 29 **Static and Default Routes** 785
About Static and Default Routes 785
Default Route 785
Static Routes 785
Route to null0 Interface to Drop Unwanted Traffic 786
Route Priorities 786
Transparent Firewall Mode and Bridge Group Routes 786
Static Route Tracking 786
Guidelines for Static and Default Routes 787
Configure Default and Static Routes 788
Configure a Default Route 788
Configure a Static Route 789

Configure Static Route Tracking 790
Monitoring a Static or Default Route 791
Examples for Static or Default Routes 791
History for Static and Default Routes 792

CHAPTER 30 **Policy Based Routing** 793
About Policy Based Routing 793
Why Use Policy Based Routing? 793
Equal-Access and Source-Sensitive Routing 794
Quality of Service 794
Cost Saving 794
Load Sharing 795
Implementation of PBR 795
Guidelines for Policy Based Routing 795
Path Monitoring 797
Configure Path Monitoring 798
Configure Policy Based Routing 798
History for Policy Based Routing 801

CHAPTER 31 **Route Maps** 803
About Route Maps 803
Permit and Deny Clauses 804
Match and Set Clause Values 804
Guidelines for Route Maps 805
Define a Route Map 805
Customize a Route Map 808
Define a Route to Match a Specific Destination Address 808
Configure Prefix Rules 809
Configure Prefix Lists 809
Configure the Metric Values for a Route Action 810
Example for Route Maps 810
History for Route Maps 811

CHAPTER 32 **Bidirectional Forwarding Detection Routing** 813

About BFD Routing 813
BFD Asynchronous Mode and Echo Function 813
BFD Session Establishment 814
BFD Timer Negotiation 815
BFD Failure Detection 816
BFD Deployment Scenarios 816
Guidelines for BFD Routing 816
Configure BFD 817
Create the BFD Template 817
Configure BFD Interfaces 819
Configure BFD Maps 820
History for BFD Routing 820

CHAPTER 33 BGP 821
About BGP 821
When to Use BGP 821
Routing Table Changes 821
BGP Path Selection 823
BGP Multipath 823
Guidelines for BGP 824
Configure BGP 825
Enable BGP 825
Define the Best Path for a BGP Routing Process 827
Configure Policy Lists 827
Configure AS Path Filters 829
Configure Community Rules 829
Configure IPv4 Address Family Settings 830
Configure IPv4 Family General Settings 830
Configure IPv4 Family Aggregate Address Settings 831
Configure IPv4 Family Filtering Settings 831
Configure IPv4 Family BGP Neighbor Settings 832
Configure IPv4 Network Settings 835
Configure IPv4 Redistribution Settings 836
Configure IPv4 Route Injection Settings 836

Configure IPv6 Address Family Settings 837
Configure IPv6 Family General Settings 837
Configure IPv6 Family Aggregate Address Settings 838
Configure IPv6 Family BGP Neighbor Settings 839
Configure IPv6 Network Settings 842
Configure IPv6 Redistribution Settings 842
Configure IPv6 Route Injection Settings 843
Monitoring BGP 843
History for BGP 844

CHAPTER 34 OSPF 847
About OSPF 847
OSPF Support for Fast Hello Packets 849
Prerequisites for OSPF Support for Fast Hello Packets 849
About OSPF Support for Fast Hello Packets 849
Implementation Differences Between OSPFv2 and OSPFv3 850
Guidelines for OSPF 850
Configure OSPFv2 853
Configure a Key Chain for Authentication 854
Configure OSPFv2 Router ID 855
Manually Configure OSPF Router-ID 856
Router ID Behaviour while Migrating 856
Customize OSPFv2 857
Redistribute Routes Into OSPFv2 857
Configure Route Summarization When Redistributing Routes Into OSPFv2 859
Add a Route Summary Address 859
Add or Edit an OSPF Summary Address 860
Configure Route Summarization Between OSPFv2 Areas 860
Configure OSPFv2 Interface Parameters 861
Configure OSPFv2 Area Parameters 864
Configure OSPFv2 Filter Rules 865
Configure an OSPFv2 NSSA 866
Configure an IP Address Pool for Clustering (OSPFv2 and OSPFv3) 867
Define Static OSPFv2 Neighbors 869

Configure Route Calculation Timers 869
Log Neighbors Going Up or Down 870
Configure a Key Chain for Authentication 870
Configure Filtering in OSPF 872
Configure a Virtual Link in OSPF 873
Configure OSPFv3 874
Enable OSPFv3 874
Configure OSPFv3 Interface Parameters 875
Configure OSPFv3 Area Parameters 876
Configure a Virtual Link Neighbor 877
Configure OSPFv3 Passive Interfaces 878
Configure OSPFv3 Administrative Distance 879
Configure OSPFv3 Timers 879
Define Static OSPFv3 Neighbors 880
Send Syslog Messages 881
Suppress Syslog Messages 882
Calculate Summary Route Costs 882
Generate a Default External Route into an OSPFv3 Routing Domain 882
Configure an IPv6 Summary Prefix 883
Redistribute IPv6 Routes 883
Configure Graceful Restart 884
Configuring Graceful Restart for OSPFv2 885
Configure Cisco NSF Graceful Restart for OSPFv2 885
Configure IETF NSF Graceful Restart for OSPFv2 886
Configuring Graceful Restart for OSPFv3 886
Configuring Graceful Restart Wait Timer for OSPF 887
Remove the OSPFv2 Configuration 888
Remove the OSPFv3 Configuration 888
Example for OSPFv2 888
Examples for OSPFv3 890
Monitoring OSPF 891
History for OSPF 892

CHAPTER 35 IS-IS 895

About IS-IS 895
About NET 895
IS-IS Dynamic Hostname 896
IS-IS PDU Types 896
Operation of IS-IS on Multiaccess Circuits 897
IS-IS Election of the Designated IS 898
IS-IS LSPDB Synchronization 899
IS-IS Shortest Path Calculation 900
IS-IS Shutdown Protocol 901
Prerequisites for IS-IS 901
Guidelines for IS-IS 901
Configure IS-IS 902
Enable IS-IS Routing Globally 902
Enable IS-IS Authentication 903
Configure IS-IS LSP 904
Configure IS-IS Summary Addresses 906
Configure IS-IS NET 907
Configure IS-IS Passive Interfaces 907
Configure IS-IS Interfaces 908
Configure IS-IS IPv4 Address Family 912
Configure IS-IS IPv6 Address Family 915
Monitoring IS-IS 917
History for IS-IS 917

CHAPTER 36 EIGRP 919
About EIGRP 919
Guidelines for EIGRP 921
Configure an EIGRP Process 922
Configure EIGRP 922
Enable EIGRP 923
Enable EIGRP Stub Routing 924
Customize EIGRP 925
Define a Network for an EIGRP Routing Process 925
Configure Interfaces for EIGRP 926

Configure Passive Interfaces 927
Configure the Summary Aggregate Addresses on Interfaces 927
Change the Interface Delay Value 928
Enable EIGRP Authentication on an Interface 928
Define an EIGRP Neighbor 930
Redistribute Routes Into EIGRP 930
Filter Networks in EIGRP 932
Customize the EIGRP Hello Interval and Hold Time 933
Disable Automatic Route Summarization 934
Configure Default Information in EIGRP 935
Disable EIGRP Split Horizon 936
Restart the EIGRP Process 936
Configure an EIGRPv6 Process 937
Enable EIGRPv6 937
Filter Rules in EIGRPv6 937
Configure Interfaces for EIGRPv6 938
Configure Passive Interfaces for EIGRPv6 939
Redistribute Routes Into EIGRPv6 940
Define an EIGRPv6 Neighbor 941
Monitoring for EIGRP 942
History for EIGRP 943

CHAPTER 37 **Multicast Routing** 945
About Multicast Routing 945
Stub Multicast Routing 945
PIM Multicast Routing 946
PIM Source Specific Multicast Support 946
PIM Bootstrap Router (BSR) 946
PIM Bootstrap Router (BSR) Terminology 947
Multicast Group Concept 947
Multicast Addresses 947
Clustering 948
Guidelines for Multicast Routing 948
Enable Multicast Routing 949

Customize Multicast Routing 949
Configure Stub Multicast Routing and Forward IGMP Messages 949
Configure a Static Multicast Route 950
Configure IGMP Features 951
Disable IGMP on an Interface 951
Configure IGMP Group Membership 952
Configure a Statically Joined IGMP Group 952
Control Access to Multicast Groups 953
Limit the Number of IGMP States on an Interface 953
Modify the Query Messages to Multicast Groups 954
Change the IGMP Version 955
Configure PIM Features 956
Enable and Disable PIM on an Interface 956
Configure a Static Rendezvous Point Address 956
Configure the Designated Router Priority 957
Configure and Filter PIM Register Messages 958
Configure PIM Message Intervals 958
Configure a Route Tree 959
Configure a Multicast Group 959
Filter PIM Neighbors 960
Configure a Bidirectional Neighbor Filter 961
Configure the ASA as a Candidate BSR 962
Configure a Multicast Boundary 962
Monitoring for PIM 963
Example for Multicast Routing 964
History for Multicast Routing 965

PART VI **AAA Servers and the Local Database** 967

CHAPTER 38 **AAA and the Local Database** 969
About AAA and the Local Database 969
Authentication 969
Authorization 970
Accounting 970

Interaction Between Authentication, Authorization, and Accounting 970
AAA Servers and Server Groups 970
About the Local Database 972
Fallback Support 973
How Fallback Works with Multiple Servers in a Group 973
Guidelines for the Local Database 973
Add a User Account to the Local Database 974
Test Local Database Authentication and Authorization 975
Monitoring the Local Database 976
History for the Local Database 976

CHAPTER 39 RADIUS Servers for AAA 981
About RADIUS Servers for AAA 981
Supported Authentication Methods 981
User Authorization of VPN Connections 982
Supported Sets of RADIUS Attributes 982
Supported RADIUS Authorization Attributes 982
Supported IETF RADIUS Authorization Attributes 990
RADIUS Accounting Disconnect Reason Codes 991
Guidelines for RADIUS Servers for AAA 992
Configure RADIUS Servers for AAA 992
Configure RADIUS Server Groups 993
Add a RADIUS Server to a Group 995
Add an Authentication Prompt 997
Test RADIUS Server Authentication and Authorization 998
Monitoring RADIUS Servers for AAA 998
History for RADIUS Servers for AAA 999

CHAPTER 40 TACACS+ Servers for AAA 1001
About TACACS+ Servers for AAA 1001
TACACS+ Attributes 1001
Guidelines for TACACS+ Servers for AAA 1002
Configure TACACS+ Servers 1003
Configure TACACS+ Server Groups 1003

Add a TACACS+ Server to a Group 1004
Add an Authentication Prompt 1005
Test TACACS+ Server Authentication and Authorization 1006
Monitoring TACACS+ Servers for AAA 1006
History for TACACS+ Servers for AAA 1007

CHAPTER 41 LDAP Servers for AAA 1009
About LDAP and the ASA 1009
How Authentication Works with LDAP 1009
LDAP Hierarchy 1010
Search the LDAP Hierarchy 1010
Bind to an LDAP Server 1011
LDAP Attribute Maps 1011
Guidelines for LDAP Servers for AAA 1012
Configure LDAP Servers for AAA 1013
Configure LDAP Attribute Maps 1013
Configure LDAP Server Groups 1014
Add an LDAP Server to a Server Group 1015
Test LDAP Server Authentication and Authorization 1017
Monitoring LDAP Servers for AAA 1017
History for LDAP Servers for AAA 1018

CHAPTER 42 RSA SecurID Servers for AAA 1019
About RSA SecurID Servers 1019
Guidelines for RSA SecurID Servers for AAA 1019
Configure RSA SecurID Servers for AAA 1020
Configure RSA SecurID AAA Server Groups 1020
Add RSA SecurID Servers to an SDI Server Group 1020
Import the SDI Node Secret File 1021
Monitor RSA SecurID Servers for AAA 1022
History for RSA SecurID Servers for AAA 1022

PART VII System Administration 1025

CHAPTER 43 **Management Access** 1027

Configure Management Remote Access 1027

Configure ASA Access for HTTPS, Telnet, or SSH 1027

Configure HTTPS Access for ASDM, Other Clients 1028

Configure SSH Access 1029

Configure Telnet Access 1037

Configure HTTP Redirect for ASDM Access or Clientless SSL VPN 1037

Configure Management Access Over a VPN Tunnel 1038

Change the Console Timeout 1039

Customize a CLI Prompt 1039

Configure a Login Banner 1041

Set a Management Session Quota 1042

Configure AAA for System Administrators 1043

Configure Management Authentication 1043

About Management Authentication 1043

Configure Authentication for CLI, ASDM, and enable command Access 1045

Configure ASDM Certificate Authentication 1046

Control CLI and ASDM Access with Management Authorization 1046

Configure Command Authorization 1048

About Command Authorization 1048

Configure Local Command Authorization 1050

Configure Commands on the TACACS+ Server 1051

Configure TACACS+ Command Authorization 1054

Configure a Password Policy for Local Database Users 1055

Change Your Password 1056

Enable and View the Login History 1056

Configure Management Access Accounting 1057

Recover from a Lockout 1058

Monitoring Device Access 1059

History for Management Access 1060

CHAPTER 44 **Software and Configurations** 1069

Upgrade the Software 1069

Load an Image Using ROMMON (ISA 3000) 1069
Upgrade the ROMMON Image: ISA 3000 1071
Downgrade Your Software 1072
Guidelines and Limitations for Downgrading 1072
Incompatible Configuration Removed After Downgrading 1074
Downgrade the ASA Appliance 1075
Downgrade the Firepower 4100/9300 1076
Downgrade the ISA 3000 1076
Manage Files 1077
Configure File Access 1077
Configure the FTP Client Mode 1078
Configure the ASA Secure Copy Client 1078
Configure the ASA TFTP Client Path 1079
Add Mount Points 1080
Access the File Management Tool 1081
Transfer Files 1082
Transfer Files Between Local PC and Flash 1082
Transfer Files Between Remote Server and Flash 1082
Set the ASA Image, ASDM, and Startup Configuration 1084
Back Up and Restore Configurations or Other Files 1086
Perform a Complete System Backup or Restoration 1086
Before You Begin Backup or Restore 1086
Back Up the System 1087
Restore the Backup 1089
Configure Automatic Backup and Restore (ISA 3000) 1089
Configure Automatic Backup (ISA 3000) 1090
Configure Automatic Restore (ISA 3000) 1090
Save the Running Configuration to a TFTP Server 1091
Schedule a System Restart 1092
Hot Swap an SSD on the Secure Firewall 3100/4200 1093
Disable the USB Port 1095
History for Software and Configurations 1096

CHAPTER 45 **Response Automation for System Events** 1099

About the EEM 1099
Supported Events 1099
Actions on Event Manager Applets 1100
Output Destinations 1100
Guidelines for the EEM 1100
Configure the EEM 1101
Create an Event Manager Applet and Configure Events 1101
Configure an Action and Destinations for Output from an Action 1102
Run an Event Manager Applet 1103
Track Memory Allocation and Memory Usage 1103
Monitoring the EEM 1104
History for the EEM 1104

CHAPTER 46 **Testing and Troubleshooting** 1105
Recover Enable and Telnet Passwords 1105
Recover Passwords on the ISA 3000 1105
Recover Passwords or Images on the ASA Virtual 1107
Disable Password Recovery for ISA 3000 Hardware 1108
Configure and Run Captures with the Packet Capture Wizard 1109
Guidelines for Packet Capture 1111
Ingress Traffic Selector 1112
Egress Traffic Selector 1113
Buffers 1114
Summary 1114
Run Captures 1114
Save Captures 1115
CPU Usage and Reporting 1115
vCPU Usage in the ASA Virtual 1115
CPU Usage Example 1115
VMware CPU Usage Reporting 1116
ASA Virtual and vCenter Graphs 1116
Amazon CloudWatch CPU Usage Reporting 1117
ASA Virtual and Amazon CloudWatch Graphs 1117
Azure CPU Usage Reporting 1117

ASA Virtual and Azure Graphs 1118
Hyper-V CPU Usage Reporting 1118
ASA Virtual and Hyper-V Graphs 1119
OCI CPU Usage Reporting 1119
ASA Virtual and OCI Graphs 1120
Test Your Configuration 1120
Test Basic Connectivity: Pinging Addresses 1120
What You Can Test Using Ping 1120
Choosing Between ICMP and TCP Ping 1121
Enable ICMP 1121
Ping Hosts 1122
Test ASA Connectivity Systematically 1122
Trace Routes to Hosts 1125
Make the ASA Visible on Trace Routes 1125
Determine Packet Routes 1126
Using the Packet Tracer to Test Policy Configuration 1127
Monitoring Performance and System Resources 1128
Monitoring Performance 1128
Monitoring Memory Blocks 1128
Monitoring CPU 1129
Monitoring Memory 1129
Monitoring Per-Process CPU Usage 1130
Monitoring Connections 1130
History for Testing and Troubleshooting 1130

PART VIII **Monitoring** 1133

CHAPTER 47 **Logging** 1135
About Logging 1135
Logging in Multiple Context Mode 1136
Syslog Message Analysis 1136
Syslog Message Format 1136
Severity Levels 1138
Syslog Message Filtering 1139

Syslog Message Classes 1139
Sort Messages in the Log Viewers 1142
Custom Message Lists 1142
Clustering 1143
Guidelines for Logging 1143
Configure Logging 1145
Enable Logging 1145
Configure an Output Destination 1145
Send Syslog Messages to an External Syslog Server 1145
Send Syslog Messages to the Internal Log Buffer 1149
Send Syslog Messages to an E-mail Address 1151
Send Syslog Messages to the Console Port 1153
Send Syslog Messages to a Telnet or SSH Session 1153
Configure Syslog Messages 1154
Configure Syslog Messaging 1154
Edit Syslog ID Settings 1155
Include a Device ID in Non-EMBLEM Formatted Syslog Messages 1155
Include the Date and Time in Syslog Messages 1156
Disable a Syslog Message 1156
Change the Severity Level of a Syslog Message 1156
Block Syslog Messages on a Standby Unit 1157
Include the Device ID in Non-EMBLEM Format Syslog Messages 1157
Create a Custom Event List 1158
Configure Logging Filters 1159
Apply Message Filters to a Logging Destination 1159
Apply Logging Filters 1159
Add or Edit a Syslog Message ID Filter 1160
Add or Edit a Message Class and Severity Filter 1160
Send All Syslog Messages in a Class to a Specified Output Destination 1161
Limit the Rate of Syslog Message Generation 1161
Assign or Change Rate Limits for Individual Syslog Messages 1162
Add or Edit the Rate Limit for a Syslog Message 1162
Edit the Rate Limit for a Syslog Severity Level 1163
Assign or Change Rate Limits for Dynamic Logging 1163

Monitoring the Logs 1164
Filter Syslog Messages Through the Log Viewers 1164
Edit Filtering Settings 1166
Issue Certain Commands Using the Log Viewers 1166
History for Logging 1167

CHAPTER 48 SNMP 1171
About SNMP 1171
SNMP Terminology 1172
SNMP Version 3 Overview 1172
Security Models 1172
SNMP Groups 1173
SNMP Users 1173
SNMP Hosts 1173
Implementation Differences Between the ASA and Cisco IOS Software 1173
SNMP Syslog Messaging 1174
Application Services and Third-Party Tools 1174
Guidelines for SNMP 1174
Configure SNMP 1176
Configure an SNMP Management Station 1176
Configure SNMP Traps 1177
Configure Parameters for SNMP Version 1 or 2c 1179
Configure Parameters for SNMP Version 3 1180
Configure a Group of Users 1181
Monitoring SNMP 1182
History for SNMP 1183

CHAPTER 49 Cisco Success Network and Telemetry Data 1189
About Cisco Success Network 1189
Supported Platforms and Required Configurations 1189
How Does ASA Telemetry Data Reach the SSE Cloud 1190
Enable or Disable Cisco Success Network 1190
View ASA Telemetry Data 1191
Cisco Success Network - Telemetry Data 1191

CHAPTER 50 **Anonymous Reporting and Smart Call Home** 1199

About Anonymous Reporting 1199

DNS Requirement 1200

About Smart Call Home 1200

Guidelines for Anonymous Reporting and Smart Call Home 1201

Configure Anonymous Reporting and Smart Call Home 1202

Configure Anonymous Reporting 1202

Configure Smart Call Home 1202

Configure Auto Import of Trustpool Certificates 1206

Monitoring Anonymous Reporting and Smart Call Home 1206

History for Anonymous Reporting and Smart Call Home 1207

PART IX **Reference** 1209

CHAPTER 51 **Addresses, Protocols, and Ports** 1211

IPv4 Addresses and Subnet Masks 1211

Classes 1211

Private Networks 1212

Subnet Masks 1212

Determine the Subnet Mask 1212

Determine the Address to Use with the Subnet Mask 1213

IPv6 Addresses 1215

IPv6 Address Format 1215

IPv6 Address Types 1216

Unicast Addresses 1216

Multicast Address 1218

Anycast Address 1219

Required Addresses 1219

IPv6 Address Prefixes 1220

Protocols and Applications 1220

TCP and UDP Ports 1221

Local Ports and Protocols 1224

ICMP Types 1226



About This Guide

The following topics explain how to use this guide.

- Document Objectives, on page li
- Related Documentation, on page li
- Document Conventions, on page li
- Communications, Services, and Additional Information, on page liii

Document Objectives

The purpose of this guide is to help you configure general operations for the Secure Firewall ASA series using the Adaptive Security Device Manager (ASDM). This guide does not cover every feature, but describes only the most common configuration scenarios.

Throughout this guide, the term "ASA" applies generically to supported models, unless specified otherwise.

Note ASDM supports many ASA versions. The ASDM documentation and online help includes all of the latest features supported by the ASA. If you are running an older version of ASA software, the documentation might include features that are not supported in your version. Please refer to the feature history table for each chapter to determine when features were added. For the minimum supported version of ASDM for each ASA version, see *Cisco ASA Series Compatibility*.

Related Documentation

For more information, see *Navigating the Cisco ASA Series Documentation* at http://www.cisco.com/go/asadocs.

Document Conventions

This document adheres to the following text, display, and alert conventions.

Text Conventions

Convention	Indication
boldface	Commands, keywords, button labels, field names, and user-entered text appear in **boldface**. For menu-based commands, the full path to the command is shown.
italic	Variables, for which you supply values, are presented in an *italic* typeface. Italic type is also used for document titles, and for general emphasis.
`monospace`	Terminal sessions and information that the system displays appear in `monospace` type.
{x | y | z}	Required alternative keywords are grouped in braces and separated by vertical bars.
[ ]	Elements in square brackets are optional.
[x | y | z]	Optional alternative keywords are grouped in square brackets and separated by vertical bars.
[ ]	Default responses to system prompts are also in square brackets.
< >	Non-printing characters such as passwords are in angle brackets.
!, #	An exclamation point (!) or a number sign (#) at the beginning of a line of code indicates a comment line.

Reader Alerts

This document uses the following for reader alerts:



Note Means *reader take note*. Notes contain helpful suggestions or references to material not covered in the manual.



Tip Means *the following information will help you solve a problem.*



Caution Means *reader be careful*. In this situation, you might do something that could result in equipment damage or loss of data.



Timesaver Means *the described action saves time.* You can save time by performing the action described in the paragraph.

Warning **Means *reader be warned.* In this situation, you might perform an action that could result in bodily injury.**

Communications, Services, and Additional Information

- To receive timely, relevant information from Cisco, sign up at Cisco Profile Manager.
- To get the business results you're looking for with the technologies that matter, visit Cisco Services.
- To submit a service request, visit Cisco Support.
- To discover and browse secure, validated enterprise-class apps, products, solutions and services, visit Cisco DevNet.
- To obtain general networking, training, and certification titles, visit Cisco Press.
- To find warranty information for a specific product or product family, access Cisco Warranty Finder.

Cisco Bug Search Tool

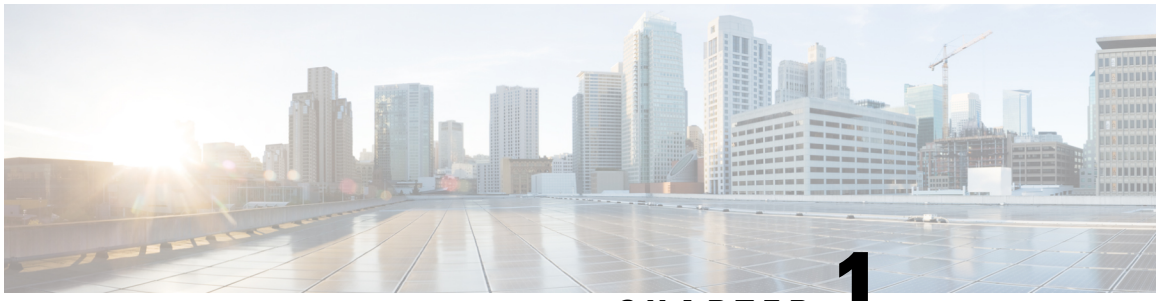
Cisco Bug Search Tool (BST) is a web-based tool that acts as a gateway to the Cisco bug tracking system that maintains a comprehensive list of defects and vulnerabilities in Cisco products and software. BST provides you with detailed defect information about your products and software.



PART I

Getting Started with the ASA

- Introduction to the Secure Firewall ASA, on page 1
- Getting Started, on page 17
- ASDM Graphical User Interface, on page 47
- Licenses: Product Authorization Key Licensing for the ISA 3000, on page 85
- Licenses: Smart Software Licensing, on page 113
- Logical Devices for the Firepower 4100/9300, on page 197
- Transparent or Routed Firewall Mode, on page 215
- Startup Wizard, on page 239



CHAPTER 1

Introduction to the Secure Firewall ASA

The Secure Firewall ASA provides advanced stateful firewall and VPN concentrator functionality in one device. The ASA includes many advanced features, such as multiple security contexts (similar to virtualized firewalls), clustering (combining multiple firewalls into a single firewall), transparent (Layer 2) firewall or routed (Layer 3) firewall operation, advanced inspection engines, IPsec VPN, SSL VPN, and clientless SSL VPN support, and many more features.

Note ASDM supports many ASA versions. The ASDM documentation and online help includes all of the latest features supported by the ASA. If you are running an older version of ASA software, the documentation might include features that are not supported in your version. Please refer to the feature history table for each chapter to determine when features were added. For the minimum supported version of ASDM for each ASA version, see Cisco ASA Compatibility. See also Special, Deprecated, and Legacy Services, on page 15.

- ASDM Requirements, on page 1
- Hardware and Software Compatibility, on page 7
- VPN Compatibility, on page 7
- New Features, on page 7
- Firewall Functional Overview, on page 10
- VPN Functional Overview, on page 14
- Security Context Overview, on page 14
- ASA Clustering Overview, on page 15
- Special, Deprecated, and Legacy Services, on page 15

ASDM Requirements

ASDM Java Requirements

You can install ASDM using Oracle JRE 8.0 (**asdm-*version*.bin**) or OpenJRE 1.8.x (**asdm-openjre-*version*.bin**).

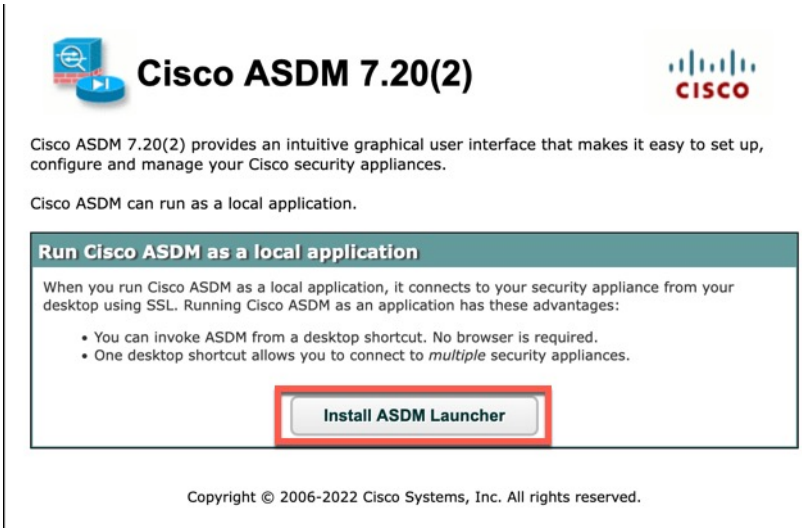
Note ASDM is not supported on Linux.

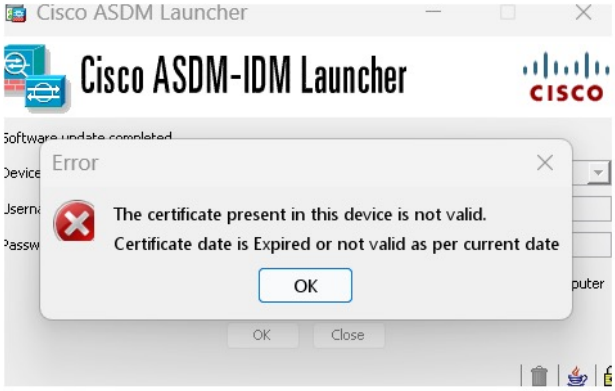
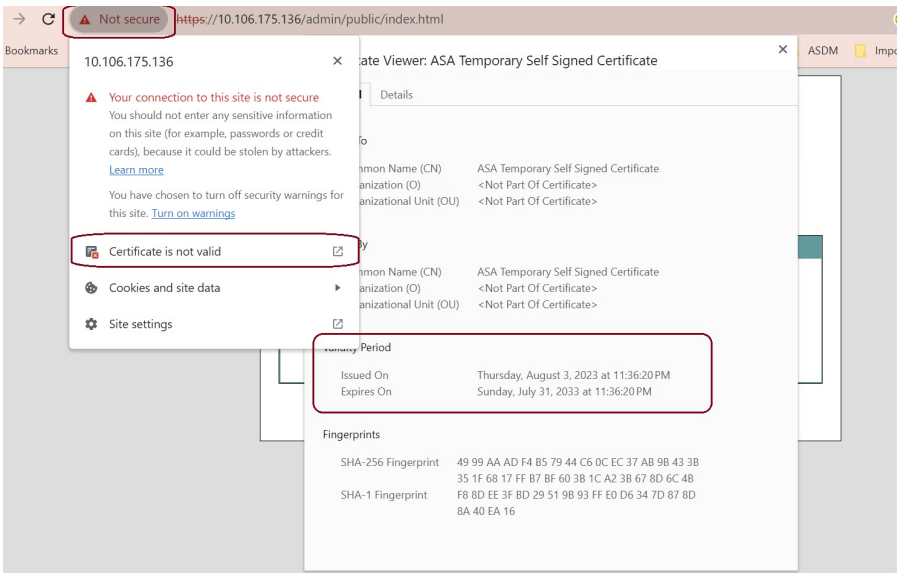
Table 1: ASDM Operating System and Browser Requirements

Operating System	Browser			Oracle JRE	OpenJRE
	Firefox	Safari	Chrome		
Microsoft Windows (English and Japanese): • 11 • 10 **Note** See Windows 10 in ASDM Compatibility Notes, on page 2 if you have problems with the ASDM shortcut. • 8 • 7 • Server 2016 and Server 2019 • Server 2012 R2 • Server 2012 • Server 2008	Yes	No support	Yes	8.0 version 8u261 or later	1.8 **Note** No support for Windows 7 or 10 32-bit
Apple OS X 10.4 and later	Yes	Yes	Yes (64-bit version only)	8.0 version 8u261 or later	1.8

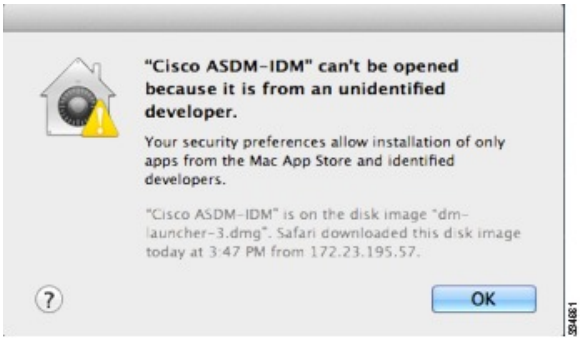
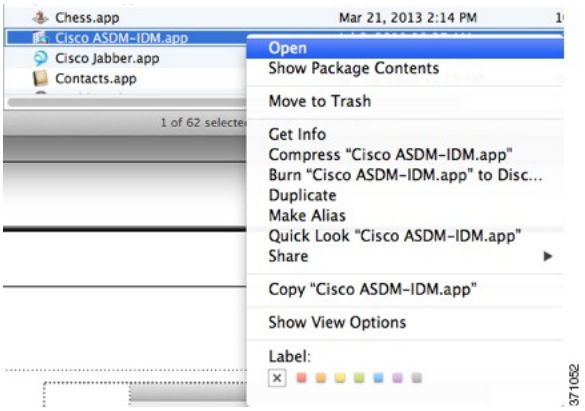
ASDM Compatibility Notes

The following table lists compatibility caveats for ASDM.

Conditions	Notes
ASDM Launcher compatibility with ASDM version	"**Unable to Launch Device Manager**" error message. If you upgrade to a new ASDM version and then get this error, you may need to re-install the latest Launcher. **1.** Open the ASDM web page on the ASA: https://<asa_ip_address>. **2.** Click **Install ASDM Launcher**. ***Figure 1: Install ASDM Launcher***  **3.** Leave the username and password fields empty (for a new installation), and click **OK**. With no HTTPS authentication configured, you can gain access to ASDM with no username and the **enable** password, which is blank by default. When you enter the **enable** command at the CLI for the first time, you are prompted to change the password; this behavior is not enforced when you log into ASDM. We suggest that you change the enable password as soon as possible so that it does not remain blank. **Note**: If you enabled HTTPS authentication, enter your username and associated password. Even without authentication, if you enter a username and password at the login screen (instead of leaving the username blank), ASDM checks the local database for a match.

Conditions	Notes
Self-signed certificate not valid due to a time and date mismatch with ASA	ASDM validates the self-signed SSL certificate, and if the ASA's date is not within the certificate's **Issued On** and **Expires On** date, ASDM will not launch. If there is a time and date mismatch, you will see the following error: ***Figure 2: Certificate Not Valid***  **To fix the issue:** Set the correct time on the ASA and reload. To check the certificate dates, (example shown is Chrome): **1.** Go to https://*device_ip*. **2.** Click the **Not secure** text in the menu bar. **3.** Click **Certificate is not valid** to open the Certificate Viewer. **4.** Check the Validity Period. ***Figure 3: Certificate Viewer*** 

Conditions	Notes
Windows Active Directory directory access	In some cases, Active Directory settings for Windows users may restrict access to program file locations needed to successfully launch ASDM on Windows. Access is needed to the following directories: • Desktop folder • C:\Windows\System32C:\Users\<username>\.asdm • C:\Program Files (x86)\Cisco Systems If your Active Directory is restricting directory access, you need to request access from your Active Directory administrator.
Windows 10	"**This app can't run on your PC**" error message. When you install the ASDM Launcher, Windows 10 might replace the ASDM shortcut target with the Windows Scripting Host path, which causes this error. To fix the shortcut target: **1.** Choose **Start** > **Cisco ASDM-IDM Launcher**, and right-click the **Cisco ASDM-IDM Launcher** application. **2.** Choose **More** > **Open file location**. Windows opens the directory with the shortcut icon. **3.** Right click the shortcut icon, and choose **Properties**. **4.** Change the **Target** to: **C:\Windows\System32\wscript.exe invisible.vbs run.bat** **5.** Click **OK**.
OS X	On OS X, you may be prompted to install Java the first time you run ASDM; follow the prompts as necessary. ASDM will launch after the installation completes.

Conditions	Notes
OS X 10.8 and later	You need to allow ASDM to run because it is not signed with an Apple Developer ID. If you do not change your security preferences, you see an error screen.  1. To allow ASDM to run, right-click (or Ctrl-Click) the Cisco ASDM-IDM Launcher icon, and choose **Open**.  2. You see a similar error screen; however, you can open ASDM from this screen. Click **Open**. The ASDM-IDM Launcher opens. 

Conditions	Notes
(ASA 5500 and ISA 3000) Requires Strong Encryption license (3DES/AES) on ASA **Note** Smart licensing models allow access with ASDM without the Strong Encryption license.	ASDM requires an SSL connection to the ASA. You can request a 3DES PAK license from Cisco: 1. Go to https://www.cisco.com/go/license. 2. Under **Traditional Licenses**, click **Access LRP**. 3. Click **Get Licenses** and then choose **IPS, Crypto, Other...** from the drop-down list. 4. Type **ASA** in to the **Search by Keyword** field. 5. Select **Cisco ASA 3DES/AES License** in the **Product** list, and click **Next**. 6. Enter the serial number of the ASA, and follow the prompts to request a 3DES/AES license for the ASA.
• Self-signed certificate or an untrusted certificate • IPv6 • Firefox and Safari	When the ASA uses a self-signed certificate or an untrusted certificate, Firefox and Safari are unable to add security exceptions when browsing using HTTPS over IPv6. See https://bugzilla.mozilla.org/show_bug.cgi?id=633001. This caveat affects all SSL connections originating from Firefox or Safari to the ASA (including ASDM connections). To avoid this caveat, configure a proper certificate for the ASA that is issued by a trusted certificate authority.
• SSL encryption on the ASA must include both RC4-MD5 and RC4-SHA1 or disable SSL false start in Chrome. • Chrome	If you change the SSL encryption on the ASA to exclude both RC4-MD5 and RC4-SHA1 algorithms (these algorithms are enabled by default), then Chrome cannot launch ASDM due to the Chrome "SSL false start" feature. We suggest re-enabling one of these algorithms (see the **Configuration** > **Device Management** > **Advanced** > **SSL Settings** pane); or you can disable SSL false start in Chrome using the **--disable-ssl-false-start** flag according to Run Chromium with flags.

Hardware and Software Compatibility

For a complete list of supported hardware and software, see Cisco ASA Compatibility.

VPN Compatibility

See Supported VPN Platforms, Cisco ASA Series.

New Features

This section lists new features for each release.

Note New, changed, and deprecated syslog messages are listed in the syslog message guide.

New Features in ASA 9.23(1)/ASDM 7.23(1)

Released: March 5, 2025

Feature	Description
Platform Features	
Secure Firewall 1230/1240/1250	The Secure Firewall 1230/1240/1250 is a 1RU rackmountable firewall.
Increased connection limits for the Secure Firewall 4200	Connection limits have been increased: • 4225: 80M → **90M** • 4245: 80M → **180M**
Firewall Features	
Support for the RADIUS Message-Authenticator attribute.	The Message-Authenticator attribute is used to protect against Blast-RADIUS attacks. If you have upgraded your RADIUS server so it supports the message authenticator, you can enable this option to help protect against these attacks. When enabled, all requests and responses must have the message authenticator, or authentication will fail. We added the **Require message authenticator from RADIUS Server** option to the Add/Edit RADIUS Server dialog box.
New Umbrella API.	You can now configure Umbrella using the Umbrella Open API, which uses an API key with a Secret key. We updated the following screen: **Configuration** > **Firewall** > **Objects** > **Umbrella**.
Flow offload is enabled by default for the Secure Firewall 3100/4200	Flow offload is now enabled by default. Added/modified screens: **Configuration** > **Firewall** > **Advanced** > **Offload Engine**
High Availability and Scalability Features	
Multiple context support for all Secure Firewall 1200 models	We added support for multiple context mode for the Secure Firewall 1210/1220: • Secure Firewall 1210CE—5 contexts. • Secure Firewall 1210CP—5 contexts. • Secure Firewall 1220CX—10 contexts. Switchports are not supported in multiple context mode, and you must convert all interfaces to router interfaces before you can convert to multiple context mode. The Secure Firewall 1230/1240/1250 also supports multiple context mode in its initial release: • Secure Firewall 1230—25 contexts. • Secure Firewall 1240—25 contexts. • Secure Firewall 1250—25 contexts.

Feature	Description
Cluster redirect: flow offload support for the Secure Firewall 4200 asymmetric cluster traffic	For asymmetric flows, cluster redirect lets the forwarding node offload flows to hardware. This feature is enabled by default. When traffic for an existing flow is sent to a different node, then that traffic is redirected to the owner node over the cluster control link. Because asymmetric flows can create a lot of traffic on the cluster control link, letting the forwarder offload these flows can improve performance. Added/modified screens: **Configuration** > **Firewall** > **Advanced** > **Offload Engine** > **Cluster Redirect Offload**
Improved role-switch time during failover	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.
Interface Features	
Secure Firewall 1210CP IEEE 802.3bt support (PoE++ and Hi-PoE)	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: • **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces** > **Edit** > **Power Over Ethernet** • **Monitoring** > **Interfaces** > **Power Over Ethernet**
License Features	
Flexible Permanent License Reservation for ASA virtual	For an ASA virtual, you can configure any model-specific license for permanent license reservation irrespective of the RAM and vCPUs. You can switch between the permanent license reservation licenses irrespective of the memory allocated to the ASA virtual. You can also change the memory and vCPUs assigned to the ASA virtual without changing the model license. If you downgrade the ASA virtual to versions earlier than 9.23.1, the license status becomes Unregistered. We recommend that you do not downgrade an ASA virtual with flexible permanent license reservation.
Administrative, Monitoring, and Troubleshooting Features	

Feature	Description
Automated Certificate Management Environment (ACME) protocol for TLS device certificates.	You can configure Automated Certificate Management Environment (ACME) protocol to ASA trustpoint to manage the TLS device certificates. ACME enables simplified certificate management through auto renewal, domain validation, and easy enrolling and revoking of certificates. You can choose to use the Let's Encrypt CA server or use any other ACME server for the authentication. ACME uses http01 method for authentication. New or modified screens: **Add Identity Certificate > Add a new Identity Certificate > Advanced > Request from a CA**
VPN Features	
Distributed site-to-site VPN with clustering on the Secure Firewall 4200	An ASA cluster on the Secure Firewall 4200 supports site-to-site VPN in distributed mode. Distributed mode provides the ability to have many site-to-site IPsec IKEv2 VPN connections distributed across members of an ASA cluster, not just on the control node (as in centralized mode). This significantly scales VPN support beyond centralized VPN capabilities and provides high availability. New or modified screens: **Monitoring > ASA Cluster > ASA Cluster > VPN Cluster Summary** **Monitoring > VPN > VPN Statistics > Sessions** **Configuration > Device Management > High Availablility and Scalability > ASA Cluster** **Wizards > Site-to-Site** **Monitoring > VPN > VPN Statistics > Sessions** **Monitoring > ASA Cluster > ASA Cluster > VPN Cluster Summary** **Monitoring > ASA Cluster > ASA Cluster > System Resource Graphs > CPU/Memory** **Monitoring > Logging > Real-Time Log Viewer**
IPsec flow offload for traffic on the cluster control link on the Secure Firewall 4200 in distributed site-to-site VPN mode	For asymmetric flows in distributed site-to-site VPN mode, IPsec flow offload now lets the flow owner decrypt IPsec traffic in hardware that was forwarded over the cluster control link. This feature is not configurable and is always available when you enable IPsec flow offload. Added/modified screens: **Configuration** > **Firewall** > **Advanced** > **IPsec Offload**

Firewall Functional Overview

Firewalls protect inside networks from unauthorized access by users on an outside network. A firewall can also protect inside networks from each other, for example, by keeping a human resources network separate from a user network. If you have network resources that need to be available to an outside user, such as a web or FTP server, you can place these resources on a separate network behind the firewall, called a *demilitarized zone* (DMZ). The firewall allows limited access to the DMZ, but because the DMZ only includes the public servers, an attack there only affects the servers and does not affect the other inside networks. You can also control when inside users access outside networks (for example, access to the Internet), by allowing

only certain addresses out, by requiring authentication or authorization, or by coordinating with an external URL filtering server.

When discussing networks connected to a firewall, the *outside* network is in front of the firewall, the *inside* network is protected and behind the firewall, and a *DMZ*, while behind the firewall, allows limited access to outside users. Because the ASA lets you configure many interfaces with varied security policies, including many inside interfaces, many DMZs, and even many outside interfaces if desired, these terms are used in a general sense only.

Security Policy Overview

A security policy determines which traffic is allowed to pass through the firewall to access another network. By default, the ASA allows traffic to flow freely from an inside network (higher security level) to an outside network (lower security level). You can apply actions to traffic to customize the security policy.

Permitting or Denying Traffic with Access Rules

You can apply access rules to limit traffic from inside to outside, or allow traffic from outside to inside. For bridge group interfaces, you can also apply an EtherType access rule to allow non-IP traffic.

Applying NAT

Some of the benefits of NAT include the following:

- You can use private addresses on your inside networks. Private addresses are not routable on the Internet.
- NAT hides the local addresses from other networks, so attackers cannot learn the real address of a host.
- NAT can resolve IP routing problems by supporting overlapping IP addresses.

Protecting from IP Fragments

The ASA provides IP fragment protection. This feature performs full reassembly of all ICMP error messages and virtual reassembly of the remaining IP fragments that are routed through the ASA. Fragments that fail the security check are dropped and logged. Virtual reassembly cannot be disabled.

Applying HTTP, HTTPS, or FTP Filtering

Although you can use access lists to prevent outbound access to specific websites or FTP servers, configuring and managing web usage this way is not practical because of the size and dynamic nature of the Internet.

You can configure Cloud Web Security on the ASA. You can also use the ASA in conjunction with an external product such as the Cisco Web Security Appliance (WSA).

Applying Application Inspection

Inspection engines are required for services that embed IP addressing information in the user data packet or that open secondary channels on dynamically assigned ports. These protocols require the ASA to do a deep packet inspection.

Applying QoS Policies

Some network traffic, such as voice and streaming video, cannot tolerate long latency times. QoS is a network feature that lets you give priority to these types of traffic. QoS refers to the capability of a network to provide better service to selected network traffic.

Applying Connection Limits and TCP Normalization

You can limit TCP and UDP connections and embryonic connections. Limiting the number of connections and embryonic connections protects you from a DoS attack. The ASA uses the embryonic limit to trigger TCP Intercept, which protects inside systems from a DoS attack perpetrated by flooding an interface with TCP SYN packets. An embryonic connection is a connection request that has not finished the necessary handshake between source and destination.

TCP normalization is a feature consisting of advanced TCP connection settings designed to drop packets that do not appear normal.

Enabling Threat Detection

You can configure scanning threat detection and basic threat detection, and also how to use statistics to analyze threats.

Basic threat detection detects activity that might be related to an attack, such as a DoS attack, and automatically sends a system log message.

A typical scanning attack consists of a host that tests the accessibility of every IP address in a subnet (by scanning through many hosts in the subnet or sweeping through many ports in a host or subnet). The scanning threat detection feature determines when a host is performing a scan. Unlike IPS scan detection that is based on traffic signatures, the ASA scanning threat detection feature maintains an extensive database that contains host statistics that can be analyzed for scanning activity.

The host database tracks suspicious activity such as connections with no return activity, access of closed service ports, vulnerable TCP behaviors such as non-random IPID, and many more behaviors.

You can configure the ASA to send system log messages about an attacker or you can automatically shun the host.

Firewall Mode Overview

The ASA runs in two different firewall modes:

- Routed
- Transparent

In routed mode, the ASA is considered to be a router hop in the network.

In transparent mode, the ASA acts like a “bump in the wire,” or a “stealth firewall,” and is not considered a router hop. The ASA connects to the same network on its inside and outside interfaces in a "bridge group".

You might use a transparent firewall to simplify your network configuration. Transparent mode is also useful if you want the firewall to be invisible to attackers. You can also use a transparent firewall for traffic that would otherwise be blocked in routed mode. For example, a transparent firewall can allow multicast streams using an EtherType access list.

Routed mode supports Integrated Routing and Bridging, so you can also configure bridge groups in routed mode, and route between bridge groups and regular interfaces. In routed mode, you can replicate transparent mode functionality; if you do not need multiple context mode or clustering, you might consider using routed mode instead.

Stateful Inspection Overview

All traffic that goes through the ASA is inspected using the Adaptive Security Algorithm and either allowed through or dropped. A simple packet filter can check for the correct source address, destination address, and ports, but it does not check that the packet sequence or flags are correct. A filter also checks *every* packet against the filter, which can be a slow process.



Note The TCP state bypass feature allows you to customize the packet flow.

A stateful firewall like the ASA, however, takes into consideration the state of a packet:

- Is this a new connection?

  If it is a new connection, the ASA has to check the packet against access lists and perform other tasks to determine if the packet is allowed or denied. To perform this check, the first packet of the session goes through the "session management path," and depending on the type of traffic, it might also pass through the "control plane path."

  The session management path is responsible for the following tasks:

  - Performing the access list checks
  - Performing route lookups
  - Allocating NAT translations (xlates)
  - Establishing sessions in the "fast path"

  The ASA creates forward and reverse flows in the fast path for TCP traffic; the ASA also creates connection state information for connectionless protocols like UDP, ICMP (when you enable ICMP inspection), so that they can also use the fast path.



Note For other IP protocols, like SCTP, the ASA does not create reverse path flows. As a result, ICMP error packets that refer to these connections are dropped.

  Some packets that require Layer 7 inspection (the packet payload must be inspected or altered) are passed on to the control plane path. Layer 7 inspection engines are required for protocols that have two or more channels: a data channel, which uses well-known port numbers, and a control channel, which uses different port numbers for each session. These protocols include FTP, H.323, and SNMP.

- Is this an established connection?

  If the connection is already established, the ASA does not need to re-check packets; most matching packets can go through the "fast" path in both directions. The fast path is responsible for the following tasks:

- IP checksum verification
- Session lookup
- TCP sequence number check
- NAT translations based on existing sessions
- Layer 3 and Layer 4 header adjustments

Data packets for protocols that require Layer 7 inspection can also go through the fast path.

Some established session packets must continue to go through the session management path or the control plane path. Packets that go through the session management path include HTTP packets that require inspection or content filtering. Packets that go through the control plane path include the control packets for protocols that require Layer 7 inspection.

VPN Functional Overview

A VPN is a secure connection across a TCP/IP network (such as the Internet) that appears as a private connection. This secure connection is called a tunnel. The ASA uses tunneling protocols to negotiate security parameters, create and manage tunnels, encapsulate packets, transmit or receive them through the tunnel, and unencapsulate them. The ASA functions as a bidirectional tunnel endpoint: it can receive plain packets, encapsulate them, 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, unencapsulate them, and send them to their final destination. The ASA invokes various standard protocols to accomplish these functions.

The ASA performs the following functions:

- Establishes tunnels
- Negotiates tunnel parameters
- Authenticates users
- Assigns user addresses
- Encrypts and decrypts data
- Manages security keys
- Manages data transfer across the tunnel
- Manages data transfer inbound and outbound as a tunnel endpoint or router

The ASA invokes various standard protocols to accomplish these functions.

Security Context Overview

You can partition a single ASA into multiple virtual devices, known as security contexts. Each context is an independent device, with its own security policy, interfaces, and administrators. Multiple contexts are similar to having multiple standalone devices. Many features are supported in multiple context mode, including routing tables, firewall features, IPS, and management; however, some features are not supported. See the feature chapters for more information.

In multiple context mode, the ASA includes a configuration for each context that identifies the security policy, interfaces, and almost all the options you can configure on a standalone device. The system administrator adds and manages contexts by configuring them in the system configuration, which, like a single mode configuration, is the startup configuration. The system configuration identifies basic settings for the ASA. The system configuration does not include any network interfaces or network settings for itself; rather, when the system needs to access network resources (such as downloading the contexts from the server), it uses one of the contexts that is designated as the admin context.

The admin context is just like any other context, except that when a user logs into the admin context, then that user has system administrator rights and can access the system and all other contexts.

ASA Clustering Overview

ASA Clustering lets you group multiple ASAs 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 perform all configuration (aside from the bootstrap configuration) on the control unit only; the configuration is then replicated to the member units.

Special, Deprecated, and Legacy Services

For some services, documentation is located outside of the main configuration guides and online help.

Special Services Guides

Special services allow the ASA to interoperate with other Cisco products; for example, by providing a security proxy for phone services (Unified Communications), or by providing Botnet traffic filtering in conjunction with the dynamic database from the Cisco update server, or by providing WCCP services for the Cisco Web Security Appliance. Some of these special services are covered in separate guides:

- Cisco ASA Botnet Traffic Filter Guide
- Cisco ASA NetFlow Implementation Guide
- Cisco ASA Unified Communications Guide
- Cisco ASA WCCP Traffic Redirection Guide
- SNMP Version 3 Tools Implementation Guide

Deprecated Services

For deprecated features, see the configuration guide for your ASA version. Similarly, for redesigned features such as NAT between Version 8.2 and 8.3 or transparent mode interfaces between Version 8.3 and 8.4, refer to the configuration guide for your version. Although ASDM is backwards compatible with previous ASA releases, the configuration guide and online help only cover the latest release.

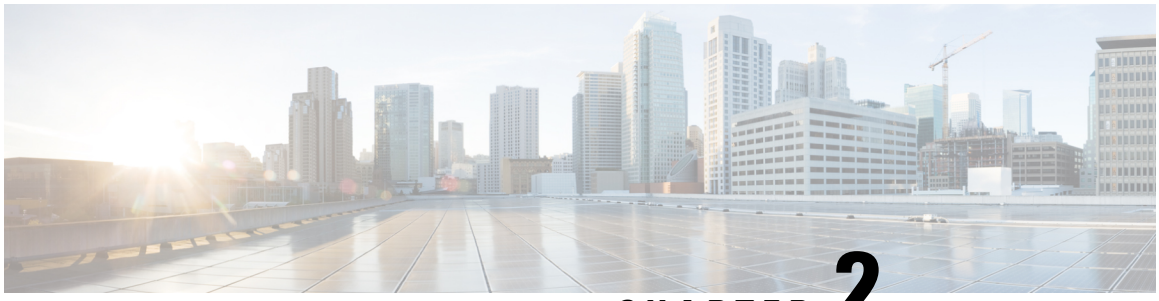
Legacy Services Guide

Legacy services are still supported on the ASA, however there may be better alternative services that you can use instead. Legacy services are covered in a separate guide:

Cisco ASA Legacy Feature Guide

This guide includes the following chapters:

- Configuring RIP
- AAA Rules for Network Access
- Using Protection Tools, which includes Preventing IP Spoofing (**ip verify reverse-path**), Configuring the Fragment Size (**fragment**), Blocking Unwanted Connections (**shun**), Configuring TCP Options (for ASDM), and Configuring IP Audit for Basic IPS Support (**ip audit**).
- Configuring Filtering Services



CHAPTER 2

Getting Started

This chapter describes how to get started with your ASA.

- Access the Console for the Command-Line Interface, on page 17
- Configure ASDM Access, on page 21
- Start ASDM, on page 25
- Customize ASDM Operation, on page 26
- Factory Default Configurations, on page 28
- Get Started with the Configuration, on page 44
- Use the Command Line Interface Tool in ASDM, on page 44
- Apply Configuration Changes to Connections, on page 46

Access the Console for the Command-Line Interface

In some cases, you may need to use the CLI to configure basic settings for ASDM access.

For initial configuration, access the CLI directly from the console port. Later, you can configure remote access using Telnet or SSH according to Management Access, on page 1027. If your system is already in multiple context mode, then accessing the console port places you in the system execution space.

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.

Note For ASA virtual console access, see the ASA virtual quick start guide.

Access the ISA 3000 Console

Follow these steps to access the appliance console.

Procedure

Step 1 Connect a computer to the console port using the provided console cable, and connect to the console using a terminal emulator set for 9600 baud, 8 data bits, no parity, 1 stop bit, no flow control.

See the hardware guide for your ASA for more information about the console cable.

Step 2 Press the **Enter** key to see the following prompt:

```
ciscoasa>
```

This prompt indicates that you are in user EXEC mode. Only basic commands are available from user EXEC mode.

Step 3 Access privileged EXEC mode.

enable

You are prompted to change the password the first time you enter the **enable** command:

Example:

```
ciscoasa> enable
Password:
The enable password is not set. Please set it now.
Enter Password: ******
Repeat Password: ******
ciscoasa#
```

All non-configuration commands are available in privileged EXEC mode. You can also enter configuration mode from privileged EXEC mode.

To exit privileged mode, enter the **disable**, **exit**, or **quit** command.

Step 4 Access global configuration mode.

configure terminal

Example:

```
ciscoasa# configure terminal
ciscoasa(config)#
```

You can begin to configure the ASA from global configuration mode. To exit global configuration mode, enter the **exit**, **quit**, or **end** command.

Access the Firepower 1000, and Secure Firewall 1200/3100/4200 Console

The Firepower 1000, and Secure Firewall 1200/3100/4200 console port connects you to the ASA CLI. From the ASA CLI, you can then connect to the FXOS CLI using Telnet for troubleshooting purposes.

Procedure

Step 1 Connect your management computer to the console port. Be sure to install any necessary serial drivers for your operating system. Use the following serial settings:

- 9600 baud
- 8 data bits
- No parity
- 1 stop bit

You connect to the ASA CLI. There are no user credentials required for console access by default.

Step 2 Access privileged EXEC mode.

enable

You are prompted to change the password the first time you enter the **enable** command.

Example:

```
ciscoasa> enable
Password:
The enable password is not set. Please set it now.
Enter Password: ******
Repeat Password: ******
ciscoasa#
```

The enable password that you set on the ASA is also the FXOS **admin** user password if the ASA fails to boot up, and you enter FXOS failsafe mode.

All non-configuration commands are available in privileged EXEC mode. You can also enter configuration mode from privileged EXEC mode.

To exit privileged EXEC mode, enter the **disable**, **exit**, or **quit** command.

Step 3 Access global configuration mode.

configure terminal

Example:

```
ciscoasa# configure terminal
ciscoasa(config)#
```

You can begin to configure the ASA from global configuration mode. To exit global configuration mode, enter the **exit**, **quit**, or **end** command.

Step 4 (Optional) Connect to the FXOS CLI.

connect fxos [**admin**]

- **admin**—Provides admin-level access. Without this option, users have read-only access. Note that no configuration commands are available even in admin mode.

You are not prompted for user credentials. The current ASA username is passed through to FXOS, and no additional login is required. To return to the ASA CLI, enter **exit** or type **Ctrl-Shift-6**, **x**.

Within FXOS, you can view user activity using the **scope security/show audit-logs** command.

Example:

```
ciscoasa# connect fxos admin
Connecting to fxos.
Connected to fxos. Escape character sequence is 'CTRL-^X'.
firepower#
firepower# exit
Connection with FXOS terminated.
Type help or '?' for a list of available commands.
ciscoasa#
```

Access the ASA Console on the Firepower 4100/9300 Chassis

For initial configuration, access the command-line interface by connecting to the Firepower 4100/9300 chassis supervisor (either to the console port or remotely using Telnet or SSH) and then connecting to the ASA security module.

Procedure

Step 1 Connect to the Firepower 4100/9300 chassis supervisor CLI (console or SSH), and then session to the ASA:

connect module *slot* {**console** | **telnet**}

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.

The first time you access the module, you access the FXOS module CLI. You must then connect to the ASA application.

connect asa

Example:

```
Firepower# connect module 1 console
Firepower-module1> connect asa

asa>
```

Step 2 Access privileged EXEC mode, which is the highest privilege level.

enable

You are prompted to change the password the first time you enter the **enable** command.

Example:

```
asa> enable
Password:
```

```
The enable password is not set. Please set it now.
Enter Password: ******
Repeat Password: ******
asa#
```

All non-configuration commands are available in privileged EXEC mode. You can also enter configuration mode from privileged EXEC mode.

To exit privileged mode, enter the **disable**, **exit**, or **quit** command.

Step 3 Enter global configuration mode.

configure terminal

Example:

```
asa# configure terminal
asa(config)#
```

To exit global configuration mode, enter the **disable**, **exit**, or **quit** command.

Step 4 Exit the application console to the FXOS module CLI by entering **Ctrl-a, d**

You might want to use the FXOS module CLI for troubleshooting purposes.

Step 5 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-], .**

Configure ASDM Access

This section describes how to access ASDM with a default configuration and how to configure access if you do not have a default configuration.

Use the Factory Default Configuration for ASDM Access

With a factory default configuration, ASDM connectivity is pre-configured with default network settings.

Procedure

Connect to ASDM using the following interface and network settings:

- The management interface depends on your model:
  - Firepower 1010, Secure Firewall 1210/1220—Management 1/1 (192.168.45.1), or inside Ethernet 1/2 through 1/8(1010 and 1210) or 1/10 (1220) (192.168.1.1). Management hosts are limited to the 192.168.45.0/24 network, and inside hosts are limited to the 192.168.1.0/24 network.
  - Firepower 1100, Secure Firewall 1230/1240/1250, 3100, 4200—Inside Ethernet 1/2 (192.168.1.1), or Management 1/1 (from DHCP). Inside hosts are limited to the 192.168.1.0/24 network. Management hosts are allowed from any network.
  - Firepower 4100/9300—The Management type interface and IP address of your choice defined when you deployed. Management hosts are allowed from any network.
  - ASA Virtual—Management 0/0 (set during deployment). Management hosts are limited to the management network.
  - ISA 3000—Management 1/1 (192.168.1.1). Management hosts are limited to the 192.168.1.0/24 network.

Note
If you change to multiple context mode, you can access ASDM from the admin context using the network settings above.

Related Topics

Factory Default Configurations, on page 28
Enable or Disable Multiple Context Mode, on page 261
Start ASDM, on page 25

Customize ASDM Access

Use this procedure if *one or more* of the following conditions applies:

- You do not have a factory default configuration
- You want to change to transparent firewall mode
- You want to change to multiple context mode

For routed, single mode, for quick and easy ASDM access, we recommend applying the factory default configuration with the option to set your own management IP address. Use the procedure in this section only if you have special needs such as setting transparent or multiple context mode, or if you have other configuration that you need to preserve.



Note For the ASAv, you can configure transparent mode when you deploy, so this procedure is primarily useful after you deploy if you need to clear your configuration, for example.

Procedure

Step 1 Access the CLI at the console port.

Step 2 (Optional) Enable transparent firewall mode:

This command clears your configuration.

firewall transparent

Step 3 Configure the management interface:

```
interface interface_id
   nameif name
   security-level level
   no shutdown
   ip address ip_address mask
```

Example:

```
ciscoasa(config)# interface management 0/0
ciscoasa(config-if)# nameif management
ciscoasa(config-if)# security-level 100
ciscoasa(config-if)# no shutdown
ciscoasa(config-if)# ip address 192.168.1.1 255.255.255.0
```

The **security-level** is a number between 1 and 100, where 100 is the most secure.

Step 4 (For directly-connected management hosts) Set the DHCP pool for the management network:

```
dhcpd address ip_address-ip_address interface_name
dhcpd enable interface_name
```

Example:

```
ciscoasa(config)# dhcpd address 192.168.1.2-192.168.1.254 management
ciscoasa(config)# dhcpd enable management
```

Make sure you do not include the interface address in the range.

Step 5 (For remote management hosts) Configure a route to the management hosts:

route *management_ifc management_host_ip mask gateway_ip* **1**

Example:

```
ciscoasa(config)# route management 10.1.1.0 255.255.255.0 192.168.1.50 1
```

Step 6 Enable the HTTP server for ASDM:

http server enable

Step 7 Allow the management host(s) to access ASDM:

http *ip_address mask interface_name*

Example:

```
ciscoasa(config)# http 192.168.1.0 255.255.255.0 management
```

Step 8 Save the configuration:

write memory

Step 9 (Optional) Set the mode to multiple mode:

mode multiple

When prompted, confirm that you want to convert the existing configuration to be the admin context. You are then prompted to reload the ASA.

Examples

The following configuration converts the firewall mode to transparent mode, configures the Management 0/0 interface, and enables ASDM for a management host:

```
firewall transparent
interface management 0/0

ip address 192.168.1.1 255.255.255.0
nameif management
security-level 100
no shutdown

dhcpd address 192.168.1.2-192.168.1.254 management
dhcpd enable management
http server enable
http 192.168.1.0 255.255.255.0 management
```

Related Topics

Restore the Factory Default Configuration, on page 29
Set the Firewall Mode (Single Mode), on page 225
Access the ISA 3000 Console, on page 17
Start ASDM, on page 25

Start ASDM

Launch ASDM using the ASDM-IDM Launcher. The Launcher is an application downloaded from the ASA using a web browser that you can use to connect to any ASA IP address. You do not need to re-download the launcher if you want to connect to other ASAs.

Within ASDM, you can choose a different ASA IP address to manage.

This section describes how to connect to ASDM initially, and then launch ASDM using the Launcher.

ASDM stores files in the local \Users\<user_id>\.asdm directory, including cache, log, and preferences, and also in the Temp directory, including Secure Client profiles.

Procedure

Step 1 On the computer that you specified as the ASDM client, enter the following URL:

https://*asa_ip_address***/admin**

Note
Be sure to specify **https://**, and not **http://** or just the IP address (which defaults to HTTP); the ASA does not automatically forward an HTTP request to HTTPS.

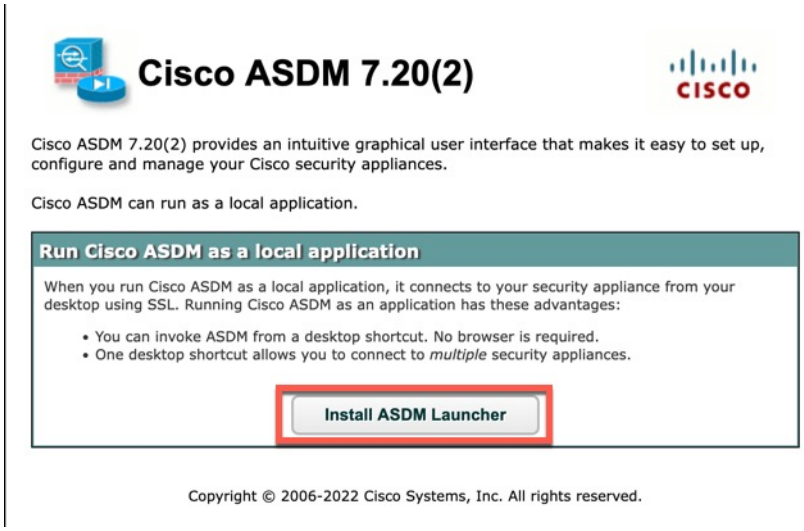
The ASDM launch page appears with the following button:

Install ASDM Launcher

Step 2 To download the Launcher and start ASDM:

a) Click **Install ASDM Launcher**.

Figure 4: Install ASDM Launcher



b) Leave the username and password fields empty (for a new installation), and click **OK**.

With no HTTPS authentication configured, you can gain access to ASDM with no username and the **enable** password, which is blank by default. When you enter the **enable** command at the CLI for the first time, you are prompted to change the password; this behavior is not enforced when you log into ASDM. We suggest that you change the enable password as soon as possible so that it does not remain blank; see Set the Hostname, Domain Name, and the Enable and Telnet Passwords, on page 693. **Note**: If you enabled HTTPS authentication, enter your username and associated password. Even without authentication, if you enter a username and password at the login screen (instead of leaving the username blank), ASDM checks the local database for a match.

c) Save the installer to your computer, and then start the installer. The ASDM-IDM Launcher opens automatically after installation is complete.

d) Enter the management IP address, the same username and password (blank for a new installation), and then click **OK**.

Customize ASDM Operation

You can install an identity certificate to successfully launch ASDM as well as increase the ASDM heap memory so it can handle larger configurations.

Install an Identity Certificate for ASDM

When using Java 7 update 51 and later, the ASDM Launcher requires a trusted certificate. An easy approach to fulfill the certificate requirements is to install a self-signed identity certificate. You can use Java Web Start to launch ASDM until you install a certificate.

See the following document to install a self-signed identity certificate on the ASA for use with ASDM, and to register the certificate with Java.

http://www.cisco.com/go/asdm-certificate

Increase the ASDM Configuration Memory

ASDM supports a maximum configuration size of 512 KB. If you exceed this amount you may experience performance issues. For example, when you load the configuration, the status dialog box shows the percentage of the configuration that is complete, yet with large configurations it stops incrementing and appears to suspend operation, even though ASDM might still be processing the configuration. If this situation occurs, we recommend that you consider increasing the ASDM system heap memory. To confirm that you are experiencing memory exhaustion, monitor the Java console for the "java.lang.OutOfMemoryError" message.

In addition, we recommend reducing your configuration size if possible, for example, by removing unused objects.

Increase the ASDM Configuration Memory in Windows

To increase the ASDM heap memory size, edit the **run.bat** file by performing the following procedure.

Procedure

Step 1 Go to the ASDM installation directory, for example C:\Program Files (x86)\Cisco Systems\ASDM.

Step 2 Edit the **run.bat** file with any text editor.

Step 3 In the line that starts with "start javaw.exe", change the argument prefixed with "-Xmx" to specify your desired heap size. For example, change it to -Xmx768M for 768 MB or -Xmx1G for 1 GB.

For very large configurations, you may need to specify a heap size up to 2 GB.

Step 4 Save the **run.bat** file.

Increase the ASDM Configuration Memory in Mac OS

To increase the ASDM heap memory size, edit the **Info.plist** file by performing the following procedure.

Procedure

Step 1 Right-click the **Cisco ASDM-IDM** icon, and choose **Show Package Contents**.

Step 2 In the **Contents** folder, double-click the **Info.plist** file. If you have Developer tools installed, it opens in the **Property List Editor**. Otherwise, it opens in **TextEdit**.

Step 3 Under **Java** > **VMOptions**, change the string prefixed with "-Xmx" to specify your desired heap size. For example, change it to -Xmx768M for 768 MB or -Xmx1G for 1 GB.

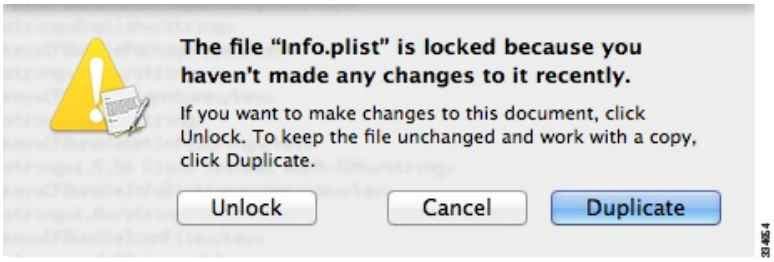
```
<key>CFBundleIconFile</key>
<string>asdm32.icns</string>

<key>VMOptions</key>
<string>-Xms64m -Xmx512m</string>


<key>CFBundleDocumentTypes</key>
  <array>
```

For very large configurations, you may need to specify a heap size up to 2 GB.

Step 4 If this file is locked, you see an error such as the following:



Step 5 Click **Unlock** and save the file.

If you do not see the **Unlock** dialog box, exit the editor, right-click the **Cisco ASDM-IDM** icon, choose **Copy Cisco ASDM-IDM**, and paste it to a location where you have write permissions, such as the Desktop. Then change the heap size from this copy.

Factory Default Configurations

The factory default configuration is the configuration applied by Cisco to new ASAs.

- Firepower 1010—The factory default configuration enables a functional inside/outside configuration. You can manage the ASA using ASDM from either the management interface or the inside switch ports.
- Firepower 1100—The factory default configuration enables a functional inside/outside configuration. You can manage the ASA using ASDM from either the management interface or the inside interface.
- Secure Firewall 1210/1220—The factory default configuration enables a functional inside/outside configuration. You can manage the ASA using ASDM from either the management interface or the inside switch ports.
- Secure Firewall 1230/1240/1250—The factory default configuration enables a functional inside/outside configuration. You can manage the ASA using ASDM from either the management interface or the inside interface.
- Secure Firewall 3100—The factory default configuration enables a functional inside/outside configuration. You can manage the ASA using ASDM from either the Management 1/1 interface or the inside interface.
- Secure Firewall 4200—The factory default configuration enables a functional inside/outside configuration. You can manage the ASA using ASDM from either the Management 1/1 interface or the inside interface.
- Firepower 4100/9300 chassis—When you deploy the standalone or cluster of ASAs, the factory default configuration configures an interface for management so that you can connect to it using ASDM, with which you can then complete your configuration.
- ASA Virtual—Depending on your hypervisor, as part of deployment, the deployment configuration (the initial virtual deployment settings) configures an interface for management so that you can connect to it using ASDM, with which you can then complete your configuration. You can also configure failover IP addresses. You can also apply a “factory default” configuration if desired.
- ISA 3000—The factory default configuration is an almost-complete transparent firewall mode configuration with all inside and outside interfaces on the same network; you can connect to the management interface with ASDM to set the IP address of your network. Hardware bypass is enabled for two interface pairs.

For appliances, the factory default configuration is available only for routed firewall mode and single context mode, except for the ISA 3000, where the factory default configuration is only available in transparent mode. For the ASA virtual and the Firepower 4100/9300 chassis, you can choose transparent or routed mode at deployment.

Note In addition to the image files and the (hidden) default configuration, the following folders and files are standard in flash memory: log/, crypto_archive/, and coredumpinfo/coredump.cfg. The date on these files may not match the date of the image files in flash memory. These files aid in potential troubleshooting; they do not indicate that a failure has occurred.

Restore the Factory Default Configuration

This section describes how to restore the factory default configuration. Both CLI and ASDM procedures are provided. For the ASA virtual, this procedure erases the deployment configuration and applies the following configuration:

```
interface management 0/0
  ip address 192.168.1.1 255.255.255.0
  nameif management
  security-level 100
  no shutdown
!
asdm logging informational
asdm history enable
!
http server enable
http 192.168.1.0 255.255.255.0 management
!
dhcpd address 192.168.1.2-192.168.1.254 management
dhcpd enable management
```

Note On the Firepower 4100/9300, restoring the factory default configuration simply erases the configuration; to restore the default configuration, you must re-deploy the ASA from the supervisor.

Before you begin

This feature is available only in routed firewall mode, except for the ISA 3000, where this command is only supported in transparent mode. In addition, this feature is available only in single context mode; an ASA with a cleared configuration does not have any defined contexts to configure automatically using this feature.

Procedure

Step 1 Restore the factory default configuration:

configure factory-default [*ip_address* [*mask*]]

Example:

```
ciscoasa(config)# configure factory-default 10.1.1.1 255.255.255.0
```

If you specify the *ip_address*, then you set the inside or management interface IP address, depending on your model, instead of using the default IP address. See the following model guidelines for which interface is set by the *ip_address* option:

- Firepower 1010—Sets the **management** interface IP address.
- Firepower 1100—Sets the **inside** interface IP address.
- Secure Firewall 1210/1220—Sets the **management** interface IP address.
- Secure Firewall 1230/1240/1250—Sets the **inside** interface IP address.
- Secure Firewall 3100—Sets the **inside** interface IP address.
- Secure Firewall 4200—Sets the **inside** interface IP address.
- Firepower 4100/9300—No effect.
- ASA Virtual—Sets the **management** interface IP address.
- ISA 3000—Sets the **management** interface IP address.

The **http** command uses the subnet you specify. Similarly, the **dhcpd address** command range consists of all available addresses higher than the IP address you specify. For example, if you specify 10.5.6.78 with a subnet mask of 255.255.255.0, then the DHCP address range will be 10.5.6.79-10.5.6.254.

For the Firepower 1000, and the Secure Firewall 1200, 3100, 4200: This command clears the **boot system** command, if present, along with the rest of the configuration. This configuration change does not affect the image at bootup: the currently-loaded image continues to be used.

For all other models: This command clears the **boot system** command, if present, along with the rest of the configuration. The **boot system** command lets you boot from a specific image. The next time you reload the ASA after restoring the factory configuration, it boots from the first image in internal flash memory; if you do not have an image in internal flash memory, the ASA does not boot.

Example:

```
docs-bxb-asa3(config)# configure factory-default 10.86.203.151 255.255.254.0
Based on the management IP address and mask, the DHCP address
pool size is reduced to 103 from the platform limit 256

WARNING: The boot system configuration will be cleared.
The first image found in disk0:/ will be used to boot the
system on the next reload.
Verify there is a valid image on disk0:/ or the system will
not boot.

Begin to apply factory-default configuration:
Clear all configuration
WARNING: The new maximum-session limit will take effect after the running-config is saved
and the system boots next time. Command accepted
WARNING: Local user database is empty and there are still 'aaa' commands for 'LOCAL'.
Executing command: interface management0/0
Executing command: nameif management
INFO: Security level for "management" set to 0 by default.
Executing command: ip address 10.86.203.151 255.255.254.0
Executing command: security-level 100
Executing command: no shutdown
Executing command: exit
Executing command: http server enable
Executing command: http 10.86.202.0 255.255.254.0 management
```

```
Executing command: dhcpd address 10.86.203.152-10.86.203.254 management
Executing command: dhcpd enable management
Executing command: logging asdm informational
Factory-default configuration is completed
ciscoasa(config)#
```

Step 2 Save the default configuration to flash memory:

write memory

This command saves the running configuration to the default location for the startup configuration, even if you previously configured the **boot** config command to set a different location; when the configuration was cleared, this path was also cleared.

Step 3 (ASDM procedure.) In the main ASDM application window, do the following:

a) Choose **File** > **Reset Device to the Factory Default Configuration**.

The **Reset Device to the Default Configuration** dialog box appears.

b) (Optional) Enter the **Management IP address** of the management or inside interface, instead of using the default address.

See the previous CLI step for details about which interface IP is set per model.

c) (Optional) Choose the **Management Subnet Mask** from the drop-down list.

d) Click **OK**.

A confirmation dialog box appears.

Note
For the Firepower 1000, and the Secure Firewall 1200, 3100, 4200: This command clears the location of the boot image, if present, along with the rest of the configuration. This configuration change does not affect the image at bootup: the currently-loaded image continues to be used.

For all other models: This action also clears the location of the boot image, if present, along with the rest of the configuration. The **Configuration > Device Management > System Image/Configuration > Boot Image/Configuration** pane lets you boot from a specific image, including an image on the external memory. The next time you reload the ASA after restoring the factory configuration, it boots from the first image in internal flash memory; if you do not have an image in *internal* flash memory, the ASA does not boot.

e) Click **Yes**.

f) After you restore the default configuration, save this configuration to internal flash memory. Choose **File** > **Save Running Configuration to Flash**.

Choosing this option saves the running configuration to the default location for the startup configuration, even if you have previously configured a different location. When the configuration was cleared, this path was also cleared.

Restore the ASA Virtual Deployment Configuration

This section describes how to restore the ASA virtual deployment (Day 0) configuration.

Procedure

Step 1 For failover, power off the standby unit.

To prevent the standby unit from becoming active, you must power it off. If you leave it on, when you erase the active unit configuration, then the standby unit becomes active. When the former active unit reloads and reconnects over the failover link, the old configuration will sync from the new active unit, wiping out the deployment configuration you wanted.

Step 2 Restore the deployment configuration after you reload. For failover, enter this command on the active unit:

write erase

Note
The ASA virtual boots the current running image, so you are not reverted to the original boot image. To use the original boot image, see the **boot image** command.

Do not save the configuration.

Step 3 Reload the ASA virtual and load the deployment configuration:

reload

Step 4 For failover, power on the standby unit.

After the active unit reloads, power on the standby unit. The deployment configuration will sync to the standby unit.

Firepower 1010 Default Configuration

The default factory configuration for the Firepower 1010 configures the following:

- **Hardware switch**—Ethernet 1/2 through 1/8 belong to VLAN 1
- **inside→outside** traffic flow—Ethernet 1/1 (outside), VLAN1 (inside)
- **management**—Management 1/1 (management), IP address 192.168.45.1
- **outside IP address** from DHCP, **inside IP address**—192.168.1.1
- **DHCP server** on inside interface, management interface
- **Default route** from outside DHCP
- **ASDM** access—Management and inside hosts allowed. Management hosts are limited to the 192.168.45.0/24 network, and inside hosts are limited to the 192.168.1.0/24 network.
- **NAT**—Interface PAT for all traffic from inside to outside.
- **DNS** servers—OpenDNS servers are pre-configured.

The configuration consists of the following commands:

```
interface Vlan1
nameif inside
security-level 100
ip address 192.168.1.1 255.255.255.0
no shutdown
!
interface Management1/1
managment-only
nameif management
no shutdown
security-level 100
ip address 192.168.45.1 255.255.255.0
!
interface Ethernet1/1
nameif outside
ip address dhcp setroute
no shutdown
!
interface Ethernet1/2
no shutdown
switchport
switchport mode access
switchport access vlan 1
!
interface Ethernet1/3
no shutdown
switchport
switchport mode access
switchport access vlan 1
!
interface Ethernet1/4
no shutdown
switchport
switchport mode access
switchport access vlan 1
!
interface Ethernet1/5
no shutdown
switchport
switchport mode access
switchport access vlan 1
!
interface Ethernet1/6
no shutdown
switchport
switchport mode access
switchport access vlan 1
!
interface Ethernet1/7
no shutdown
switchport
switchport mode access
switchport access vlan 1
!
interface Ethernet1/8
no shutdown
switchport
switchport mode access
switchport access vlan 1
!
object network obj_any
   subnet 0.0.0.0 0.0.0.0
   nat (any,outside) dynamic interface
```

```
!
dhcpd auto_config outside
dhcpd address 192.168.1.20-192.168.1.254 inside
dhcpd address 192.168.45.10-192.168.45.12 management
dhcpd enable inside
dhcpd enable management
!
http server enable
http 192.168.45.0 255.255.255.0 management
http 192.168.1.0 255.255.255.0 inside
!
dns domain-lookup outside
dns server-group DefaultDNS
   name-server 208.67.222.222 outside
   name-server 208.67.220.220 outside
!
```

Firepower 1100 Default Configuration

The default factory configuration for the Firepower 1100 configures the following:

- **inside→outside traffic flow**—Ethernet 1/1 (outside), Ethernet 1/2 (inside)
- **outside IP address** from DHCP, **inside IP address**—192.168.1.1
- **management**—Management 1/1 (management), IP address from DHCP
- **DHCP server** on inside interface
- **Default routes** from outside DHCP, management DHCP
- **ASDM** access—Management and inside hosts allowed. Inside hosts are limited to the 192.168.1.0/24 network.
- **NAT**—Interface PAT for all traffic from inside to outside.
- **DNS** servers—OpenDNS servers are pre-configured.

The configuration consists of the following commands:

```
interface Management1/1
  management-only
  nameif management
  security-level 100
  ip address dhcp setroute
  no shutdown
!
interface Ethernet1/1
  nameif outside
  security-level 0
  ip address dhcp setroute
  no shutdown
!
interface Ethernet1/2
  nameif inside
  security-level 100
  ip address 192.168.1.1 255.255.255.0
  no shutdown
!
object network obj_any
```

```
  subnet 0.0.0.0 0.0.0.0
  nat (any,outside) dynamic interface
!
http server enable
http 0.0.0.0 0.0.0.0 management
http 192.168.1.0 255.255.255.0 inside
!
dhcpd auto_config outside
dhcpd address 192.168.1.20-192.168.1.254 inside
dhcpd enable inside
!
dns domain-lookup outside
dns server-group DefaultDNS
   name-server 208.67.222.222 outside
   name-server 208.67.220.220 outside
!
```

Secure Firewall 1210/1220 Default Configuration

The default factory configuration for the Secure Firewall 1210/1220 configures the following:

- **Hardware switch**—Ethernet 1/2 through 1/8 (1210) or Ethernet 1/2 through 1/10 (1220) belong to VLAN 1
- **inside→outside** traffic flow—Ethernet 1/1 (outside), VLAN1 (inside)
- **management**—Management 1/1 (management), IP address 192.168.45.1
- **outside IP address** from DHCP, **inside IP address**—192.168.1.1
- **DHCP server** on inside interface, management interface
- **Default route** from outside DHCP
- **ASDM** access—Management and inside hosts allowed. Management hosts are limited to the 192.168.45.0/24 network, and inside hosts are limited to the 192.168.1.0/24 network.
- **NAT**—Interface PAT for all traffic from inside to outside.
- **DNS** servers—OpenDNS servers are pre-configured.

The configuration consists of the following commands:

```
interface Vlan1
nameif inside
security-level 100
ip address 192.168.1.1 255.255.255.0
no shutdown
!
interface Management1/1
managment-only
nameif management
no shutdown
security-level 100
ip address 192.168.45.1 255.255.255.0
!
interface Ethernet1/1
nameif outside
ip address dhcp setroute
no shutdown
```

```
!
interface Ethernet1/2
no shutdown
switchport
switchport mode access
switchport access vlan 1
!
interface Ethernet1/3
no shutdown
switchport
switchport mode access
switchport access vlan 1
!
interface Ethernet1/4
no shutdown
switchport
switchport mode access
switchport access vlan 1
!
interface Ethernet1/5
no shutdown
switchport
switchport mode access
switchport access vlan 1
!
interface Ethernet1/6
no shutdown
switchport
switchport mode access
switchport access vlan 1
!
interface Ethernet1/7
no shutdown
switchport
switchport mode access
switchport access vlan 1
!
interface Ethernet1/8
no shutdown
switchport
switchport mode access
switchport access vlan 1
!
! 1220
interface Ethernet1/9
no shutdown
switchport
switchport mode access
switchport access vlan 1
!
! 1220
interface Ethernet1/10
no shutdown
switchport
switchport mode access
switchport access vlan 1
!
object network obj_any
   subnet 0.0.0.0 0.0.0.0
   nat (any,outside) dynamic interface
!
dhcpd auto_config outside
dhcpd address 192.168.1.20-192.168.1.254 inside
dhcpd address 192.168.45.10-192.168.45.12 management
```

```
dhcpd enable inside
dhcpd enable management
!
http server enable
http 192.168.45.0 255.255.255.0 management
http 192.168.1.0 255.255.255.0 inside
!
dns domain-lookup outside
dns server-group DefaultDNS
   name-server 208.67.222.222 outside
   name-server 208.67.220.220 outside
!
```

Secure Firewall 1230/1240/1250 Default Configuration

The default factory configuration for the Secure Firewall 1230/1240/1250 configures the following:

- **inside→outside traffic flow**—Ethernet 1/1 (outside), Ethernet 1/2 (inside)
- **outside IP address** from DHCP, **inside IP address**—192.168.1.1
- **management**—Management 1/1 (management), IP address from DHCP
- **DHCP server** on inside interface
- **Default routes** from outside DHCP, management DHCP
- **ASDM** access—Management and inside hosts allowed. Inside hosts are limited to the 192.168.1.0/24 network.
- **NAT**—Interface PAT for all traffic from inside to outside.
- **DNS** servers—OpenDNS servers are pre-configured.

The configuration consists of the following commands:

```
interface Management1/1
  management-only
  nameif management
  security-level 100
  ip address dhcp setroute
  no shutdown
!
interface Ethernet1/1
  nameif outside
  security-level 0
  ip address dhcp setroute
  no shutdown
!
interface Ethernet1/2
  nameif inside
  security-level 100
  ip address 192.168.1.1 255.255.255.0
  no shutdown
!
object network obj_any
  subnet 0.0.0.0 0.0.0.0
  nat (any,outside) dynamic interface
!
http server enable
```

```
http 0.0.0.0 0.0.0.0 management
http 192.168.1.0 255.255.255.0 inside
!
dhcpd auto_config outside
dhcpd address 192.168.1.20-192.168.1.254 inside
dhcpd enable inside
!
dns domain-lookup outside
dns server-group DefaultDNS
   name-server 208.67.222.222 outside
   name-server 208.67.220.220 outside
!
```

Secure Firewall 3100 Default Configuration

The default factory configuration for the Secure Firewall 3100 configures the following:

- **inside→outside traffic flow**—Ethernet 1/1 (outside), Ethernet 1/2 (inside)
- **outside IP address** from DHCP, **inside IP address**—192.168.1.1
- **management**—Management 1/1 (management), IP address from DHCP
- **DHCP server** on inside interface
- **Default routes** from outside DHCP, management DHCP
- **ASDM** access—Management and inside hosts allowed. Inside hosts are limited to the 192.168.1.0/24 network.
- **NAT**—Interface PAT for all traffic from inside to outside.
- **DNS** servers—OpenDNS servers are pre-configured.

The configuration consists of the following commands:

```
interface Management1/1
  management-only
  nameif management
  security-level 100
  ip address dhcp setroute
  no shutdown
!
interface Ethernet1/1
  nameif outside
  security-level 0
  ip address dhcp setroute
  no shutdown
!
interface Ethernet1/2
  nameif inside
  security-level 100
  ip address 192.168.1.1 255.255.255.0
  no shutdown
!
object network obj_any
  subnet 0.0.0.0 0.0.0.0
  nat (any,outside) dynamic interface
!
http server enable
```

```
http 0.0.0.0 0.0.0.0 management
http 192.168.1.0 255.255.255.0 inside
!
dhcpd auto_config outside
dhcpd address 192.168.1.20-192.168.1.254 inside
dhcpd enable inside
!
dns domain-lookup outside
dns server-group DefaultDNS
   name-server 208.67.222.222 outside
   name-server 208.67.220.220 outside
!
```

Secure Firewall 4200 Default Configuration

The default factory configuration for the Secure Firewall 4200 configures the following:

- **inside→outside traffic flow**—Ethernet 1/1 (outside), Ethernet 1/2 (inside)
- **outside IP address** from DHCP, **inside IP address**—192.168.1.1
- **management**—Management 1/1 (management), IP address from DHCP
- **DHCP server** on inside interface
- **Default routes** from outside DHCP, management DHCP
- **ASDM** access—Management and inside hosts allowed. Inside hosts are limited to the 192.168.1.0/24 network.
- **NAT**—Interface PAT for all traffic from inside to outside.
- **DNS** servers—OpenDNS servers are pre-configured.

The configuration consists of the following commands:

```
interface Management1/1
  management-only
  nameif management
  security-level 100
  ip address dhcp setroute
  no shutdown
!
interface Ethernet1/1
  nameif outside
  security-level 0
  ip address dhcp setroute
  no shutdown
!
interface Ethernet1/2
  nameif inside
  security-level 100
  ip address 192.168.1.1 255.255.255.0
  no shutdown
!
object network obj_any
  subnet 0.0.0.0 0.0.0.0
  nat (any,outside) dynamic interface
!
http server enable
```

```
http 0.0.0.0 0.0.0.0 management
http 192.168.1.0 255.255.255.0 inside
!
dhcpd auto_config outside
dhcpd address 192.168.1.20-192.168.1.254 inside
dhcpd enable inside
!
dns domain-lookup outside
dns server-group DefaultDNS
   name-server 208.67.222.222 outside
   name-server 208.67.220.220 outside
!
```

Firepower 4100/9300 Chassis Default Configuration

When you deploy the ASA on the Firepower 4100/9300 chassis, you can pre-set many parameters that let you connect to the Management interface using ASDM. A typical configuration includes the following settings:

- Management interface:
  - Management type interface of your choice defined on the Firepower 4100/9300 Chassis supervisor
  - Named “management”
  - IP address of your choice
  - Security level 0
  - Management-only
- Default route through the management interface
- ASDM access—All hosts allowed.

The configuration for a standalone unit consists of the following commands. For additional configuration for clustered units, see Create an ASA Cluster, on page 444.

```
interface <management_ifc>
  management-only
  ip address <ip_address> <mask>
  ipv6 address <ipv6_address>
  ipv6 enable
  nameif management
  security-level 0
  no shutdown
!
http server enable
http 0.0.0.0 0.0.0.0 management
http ::/0 management
!
route management 0.0.0.0 0.0.0.0 <gateway_ip> 1
ipv6 route management ::/0 <gateway_ipv6>
```

ISA 3000 Default Configuration

The default factory configuration for the ISA 3000 configures the following:

- **Transparent firewall mode**—A transparent firewall 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.
- **1 Bridge Virtual Interface**—All member interfaces are in the same network (**IP address *not* pre-configured; you must set to match your network**): GigabitEthernet 1/1 (outside1), GigabitEthernet 1/2 (inside1), GigabitEthernet 1/3 (outside2), GigabitEthernet 1/4 (inside2)
- All **inside and outside** interfaces can communicate with each other.
- **Management 1/1** interface—192.168.1.1/24 for ASDM access.
- **DHCP** for clients on management.
- **ASDM** access—Management hosts allowed.
- **Hardware bypass** is enabled for the following interface pairs: GigabitEthernet 1/1 & 1/2; GigabitEthernet 1/3 & 1/4



Note When the ISA 3000 loses power and goes into hardware bypass mode, only the above interface pairs can communicate; inside1 and inside2, and outside1 and outside2 can no longer communicate. Any existing connections between these interfaces will be lost. When the power comes back on, there is a brief connection interruption as the ASA takes over the flows.

The configuration consists of the following commands:

```
firewall transparent

interface GigabitEthernet1/1
  bridge-group 1
  nameif outside1
  security-level 0
  no shutdown
interface GigabitEthernet1/2
  bridge-group 1
  nameif inside1
  security-level 100
  no shutdown
interface GigabitEthernet1/3
  bridge-group 1
  nameif outside2
  security-level 0
  no shutdown
interface GigabitEthernet1/4
  bridge-group 1
  nameif inside2
  security-level 100
  no shutdown
interface Management1/1
  management-only
  no shutdown
  nameif management
  security-level 100
  ip address 192.168.1.1 255.255.255.0
interface BVI1
  no ip address
```

```
access-list allowAll extended permit ip any any
access-group allowAll in interface outside1
access-group allowAll in interface outside2

same-security-traffic permit inter-interface

hardware-bypass GigabitEthernet 1/1-1/2
hardware-bypass GigabitEthernet 1/3-1/4

http server enable
http 192.168.1.0 255.255.255.0 management

dhcpd address 192.168.1.5-192.168.1.254 management
dhcpd enable management
```

ASA Virtual Deployment Configuration

When you deploy the ASA virtual, you can pre-set many parameters that let you connect to the Management 0/0 interface using ASDM. A typical configuration includes the following settings:

- Routed or Transparent firewall mode
- Management 0/0 interface:
  - Named “management”
  - IP address or DHCP
  - Security level 0
- Static route for the management host IP address (if it is not on the management subnet)
- HTTP server enabled or disabled
- HTTP access for the management host IP address
- (Optional) Failover link IP addresses for GigabitEthernet 0/8, and the Management 0/0 standby IP address
- DNS server
- Smart licensing ID token
- Smart licensing Throughput Level and Essentials Feature Tier
- Smart Transport URL (https://smartreceiver.cisco.com/licservice/license) and port 80.
- (Optional) Smart Transport Proxy URL and port
- (Optional) SSH management settings:
  - Client IP addresses
  - Local username and password
  - Authentication required for SSH using the LOCAL database
- (Optional) REST API enabled or disabled

Note To successfully register the ASA virtual with the Cisco Licensing Authority, the ASA virtual requires Internet access. You might need to perform additional configuration after deployment to achieve Internet access and successful license registration.

See the following sample configuration for a standalone unit:

```
interface Management0/0
  nameif management
  security-level 0
  ip address ip_address

  no shutdown
http server enable
http managemment_host_IP mask management
route management management_host_IP mask gateway_ip 1
dns server-group DefaultDNS
  name-server ip_address
call-home
  http-proxy ip_address port port
license smart
  feature tier standard
  throughput level {100M | 1G | 2G}
  license smart register idtoken id_token
aaa authentication ssh console LOCAL
username username password password
ssh source_IP_address mask management
rest-api image boot:/path
rest-api agent
```

Note The Essentials license used to be called “Standard” license.

See the following sample configuration for a primary unit in a failover pair:

```
nameif management
  security-level 0
  ip address ip_address standby standby_ip

  no shutdown
route management management_host_IP mask gateway_ip 1
http server enable
http managemment_host_IP mask management
dns server-group DefaultDNS
  name-server ip_address
call-home
  http-proxy ip_address port port
license smart
  feature tier standard
  throughput level {100M | 1G | 2G}
  license smart register idtoken id_token
aaa authentication ssh console LOCAL
username username password password
ssh source_IP_address mask management
rest-api image boot:/path
```

```
rest-api agent
failover
failover lan unit primary
failover lan interface fover gigabitethernet0/8
failover link fover gigabitethernet0/8
failover interface ip fover primary_ip mask standby standby_ip
```

Get Started with the Configuration

To configure and monitor the ASA, perform the following steps.

Note ASDM supports up to a maximum of a 512 KB configuration. If you exceed this amount, you may experience performance issues. See Increase the ASDM Configuration Memory, on page 26.

Procedure

Step 1 For initial configuration using the Startup Wizard, choose **Wizards** > **Startup Wizard**.

Step 2 To use the IPsec VPN Wizard to configure IPsec VPN connections, choose **Wizards** > **IPsec VPN Wizard** and complete each screen that appears.

Step 3 To use the SSL VPN Wizard to configure SSL VPN connections, choose **Wizards** > **SSL VPN Wizard** and complete each screen that appears.

Step 4 To configure high availability and scalability settings, choose **Wizards** > **High Availability and Scalability Wizard**.

Step 5 To use the Packet Capture Wizard to configure packet capture, choose **Wizards** > **Packet Capture Wizard**.

Step 6 To display different colors and styles available in the ASDM GUI, choose **View** > **Office Look and Feel**.

Step 7 To configure features, click the **Configuration** button on the toolbar and then click one of the feature buttons to display the associated configuration pane.

Note
If the Configuration screen is blank, click **Refresh** on the toolbar to display the screen content.

Step 8 To monitor the ASA, click the **Monitoring** button on the toolbar and then click a feature button to display the associated monitoring pane.

Use the Command Line Interface Tool in ASDM

This section tells how to enter commands using ASDM, and how to work with the CLI.

Use the Command Line Interface Tool

This feature provides a text-based tool for sending commands to the ASA and viewing the results.

The commands you can enter with the CLI tool depend on your user privileges. Review your privilege level in the status bar at the bottom of the main ASDM application window to ensure that you have the required privileges to execute privileged-level CLI commands.

Before you begin

- Commands entered via the ASDM CLI tool might function differently from those entered through a terminal connection to the ASA.
- Command errors—If an error occurs because you entered an incorrect command, the incorrect command is skipped and the remaining commands are processed. A message appears in the Response area to inform you whether or not any error occurred, as well as other related information.
- Interactive commands—Interactive commands are not supported in the CLI tool. To use these commands in ASDM, use the **noconfirm** keyword if available, as shown in the following command:

```
crypto key generate rsa modulus 1024 noconfirm
```

- Avoid conflicts with other administrators—Multiple administrative users can update the running configuration of the ASA. Before using the ASDM CLI tool to make configuration changes, check for other active administrative sessions. If more than one user is configuring the ASA at the same time, the most recent changes take effect.

  To view other administrative sessions that are currently active on the same ASA, choose **Monitoring > Properties > Device Access**.

Procedure

Step 1 In the main ASDM application window, choose **Tools** > **Command Line Interface**.

The **Command Line Interface** dialog box appears.

Step 2 Choose the type of command (single line or multiple line) that you want, and then choose the command from the drop-down list, or type it in the field provided.

Step 3 Click **Send** to execute the command.

Step 4 To enter a new command, click **Clear Response**, and then choose (or type) another command to execute.

Step 5 Check the **Enable context-sensitive help (?)** check box to provide context-sensitive help for this feature. Uncheck this check box to disable the context-sensitive help.

Step 6 After you have closed the **Command Line Interface** dialog box, if you changed the configuration, click **Refresh** to view the changes in ASDM.

Show Commands Ignored by ASDM on the Device

This feature lets you show the list of commands that ASDM does not support. Typically, ASDM ignores them. ASDM does not change or remove these commands from your running configuration. See Unsupported Commands, on page 81 for more information.

Procedure

Step 1 In the main ASDM application window, choose **Tools** > **Show Commands Ignored by ASDM on Device**.

Step 2 Click **OK** when you are done.

Apply Configuration Changes to Connections

When you make security policy changes to the configuration, all *new* connections use the new security policy. Existing connections continue to use the policy that was configured at the time of the connection establishment. **show** command output for old connections reflect the old configuration, and in some cases will not include data about the old connections.

For example, if you remove a QoS **service-policy** from an interface, then re-add a modified version, then the **show service-policy** command only displays QoS counters associated with new connections that match the new service policy; existing connections on the old policy no longer show in the command output.

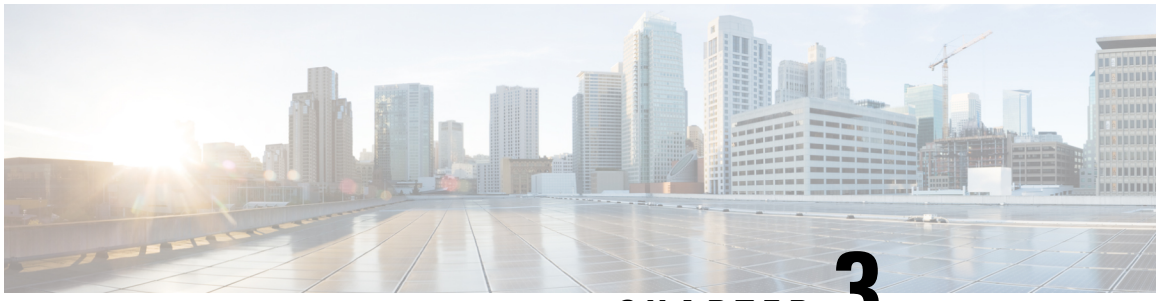
To ensure that all connections use the new policy, you need to disconnect the current connections so that they can reconnect using the new policy.

To disconnect connections, enter the following command:

- **clear conn** [**all**] [**protocol** {**tcp** | **udp**}] [**address** *src_ip* [*-src_ip*] [**netmask** *mask*]] [**port** *src_port* [*-src_port*]] [**address** *dest_ip* [*-dest_ip*] [**netmask** *mask*]] [**port** *dest_port* [*-dest_port*]]

  This command terminates connections in any state. See the **show conn** command to view all current connections.

  With no arguments, this command clears all through-the-box connections. To also clear to-the-box connections (including your current management session), use the **all** keyword. To clear specific connections based on the source IP address, destination IP address, port, and/or protocol, you can specify the desired options.



CHAPTER 3

ASDM Graphical User Interface

This chapter describes how to use the ASDM user interface.

- About the ASDM User Interface, on page 47
- Navigate the ASDM User Interface, on page 50
- Menus, on page 51
- Toolbar, on page 56
- ASDM Assistant, on page 57
- Status Bar, on page 57
- Device List, on page 58
- Common Buttons, on page 59
- Keyboard Shortcuts, on page 59
- Find Function in ASDM Panes, on page 61
- Find Function in Rule Lists, on page 62
- Enable Extended Screen Reader Support, on page 62
- Organizational Folder, on page 63
- Home Pane (Single Mode and Context), on page 63
- Home Pane (System), on page 77
- Define ASDM Preferences, on page 78
- Search with the ASDM Assistant, on page 80
- Enable History Metrics, on page 81
- Unsupported Commands, on page 81

About the ASDM User Interface

The ASDM user interface is designed to provide easy access to the many features that the ASA supports. The ASDM user interface includes the following elements:

- A menu bar that provides quick access to files, tools, wizards, and help. Many menu items also have keyboard shortcuts.
- A toolbar that enables you to navigate ASDM. From the toolbar you can access the **Home**, **Configuration**, and **Monitoring** panes. You can also get help and navigate between panes.
- A dockable left **Navigation** pane to move through the **Configuration** and **Monitoring** panes. You can click one of the three buttons in the header to maximize or restore this pane, make it a floating pane that

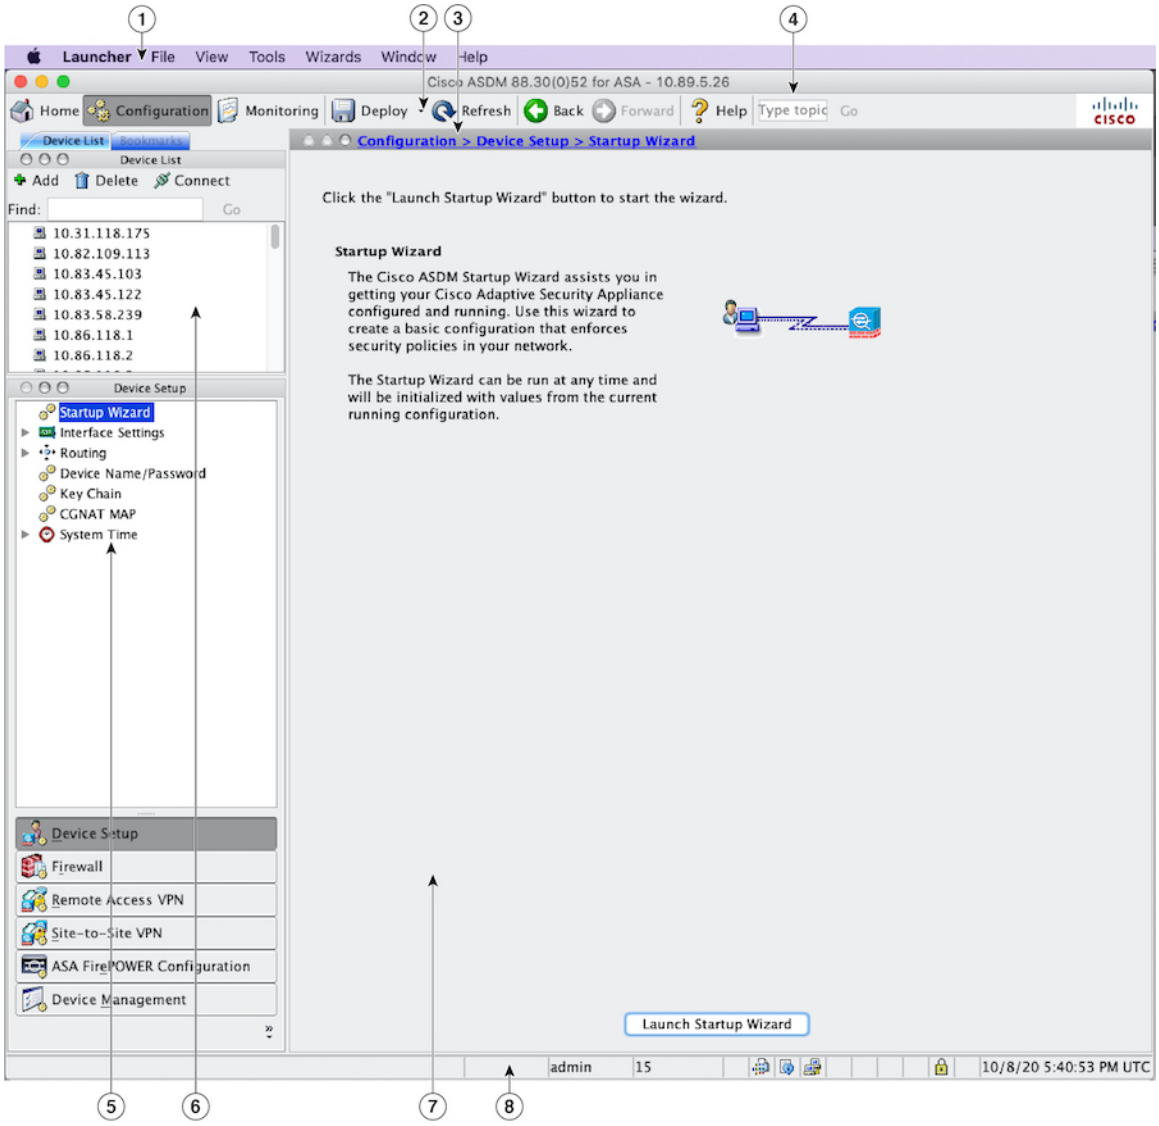
you can move, hide it, or close it. To access the **Configuration** and **Monitoring** panes, you can do one of the following:

- Click links on the left side of the application window in the left **Navigation** pane. The **Content** pane then displays the path (for example, **Configuration > Device Setup > Startup Wizard**) in the title bar of the selected pane.
- If you know the exact path, you can type it directly into the title bar of the **Content** pane on the right side of the application window, without clicking any links in the left **Navigation** pane.

- A maximize and restore button in the right corner of the **Content** pane that lets you hide and show the left **Navigation** pane.
- A dockable **Device List** pane with a list of devices that you can access through ASDM. You can click one of the three buttons in the header to maximize or restore this pane, make it a floating pane that you can move, hide it, or close it.
- A status bar that shows the time, connection status, user, memory status, running configuration status, privilege level, and SSL status at the bottom of the application window.
- A left **Navigation** pane that shows various objects that you can use in the rules tables when you create access rules, NAT rules, AAA rules, filter rules, and service rules. The tab titles within the pane change according to the feature that you are viewing. In addition, the **ASDM Assistant** appears in this pane.

The following figure shows the elements of the ASDM user interface.

Figure 5: ASDM User Interface



Legend

GUI Element	Description
1	Menu Bar
2	Toolbar
3	Navigation Path
4	Search Field
5	Left Navigation Pane
6	Device List Pane

GUI Element	Description
7	Content Pane
8	Status Bar



Note Tool tips have been added for various parts of the GUI, including **Wizards**, the **Configuration** and **Monitoring** panes, and the **Status Bar**. To view tool tips, hover your mouse over a specific user interface element, such as an icon in the status bar.

Navigate the ASDM User Interface

To move efficiently throughout the ASDM user interface, you may use a combination of menus, the toolbar, dockable panes, and the left and right **Navigation** panes, which are described in the previous section. The available functions appear in a list of buttons below the **Device List** pane. An example list could include the following function buttons:

- **Device Setup**
- **Firewall**
- **Botnet Traffic Filter**
- **Remote Access VPN**
- **Site to Site VPN**
- **Device Management**

The list of function buttons that appears is based on the licensed features that you have purchased. Click each button to access the first pane in the selected function for either the Configuration view or the Monitoring view. The function buttons are not available in the Home view.

To change the display of function buttons, perform the following steps:

Procedure

Step 1 Choose the drop-down list below the last function button to display a context menu.

Step 2 Choose one of the following options:

- Click **Show More Buttons** to show more buttons.
- Click **Show Fewer Buttons** to show fewer buttons.
- Click **Add or Remove Buttons** to add or remove buttons, then click the button to add or remove from the list that appears.
- Choose **Option** to display the **Option** dialog box, which displays a list of the buttons in their current order. Then choose one of the following:

- Click **Move Up** to move up a button in the list.
- Click **Move Down** to move down a button in the list.
- Click **Reset** to return the order of the items in the list to the default setting.

Step 3 Click **OK** to save your settings and close this dialog box.

Menus

You can access ASDM menus using the mouse or keyboard.

File Menu

The **File** menu lets you manage ASA configurations.

File Menu Item	Description
Refresh ASDM with the Running Configuration on the Device	Loads a copy of the running configuration into ASDM.
Reset Device to the Factory Default Configuration	Restores the configuration to the factory default.
Show Running Configuration in New Window	Displays the current running configuration in a new window.
Save Running Configuration to Flash	Writes a copy of the running configuration to flash memory.
Save Running Configuration to TFTP Server	Stores a copy of the current running configuration file on a TFTP server.
Save Running Configuration to Standby Unit	Sends a copy of the running configuration file on the primary unit to the running configuration of a failover standby unit.
Save Internal Log Buffer to Flash	Saves the internal log buffer to flash memory.
Deploy FirePOWER Changes	Saves configuration changes made to ASA FirePOWER module policies to the module. This option is available only if you install an ASA FirePOWER module and manage it through ASDM.
Print	Prints the current page. We recommend landscape page orientation when you print rules. When you use Internet Explorer, permission to print was already granted when you originally accepted the signed applet.
Clear ASDM Cache	Removes local ASDM images. ASDM downloads images locally when you connect to ASDM.

File Menu Item	Description
Clear ASDM Password Cache	Removes the password cache if you have defined a new password and still have a existing password that is different than the new password.
Clear Internal Log Buffer	Empties the syslog message buffer.
Exit	Closes ASDM.

View Menu

The **View** menu lets you display various parts of the ASDM user interface. Certain items are dependent on the current view. You cannot select items that cannot be displayed in the current view.

View Menu Item	Description
Home	Displays the Home view.
Configuration	Displays the Configuration view.
Monitoring	Displays the Monitoring view.
Bookmarks	Displays a list of bookmarked pages in a dockable pane.
Device List	Displays a list of devices in a dockable pane.
Navigation	Shows and hides the display of the **Navigation** pane in the Configuration and Monitoring views.
ASDM Assistant	Searches and finds useful ASDM procedural help about certain tasks.
Latest ASDM Syslog Messages	Shows and hides the display of the **Latest ASDM Syslog Messages** pane in the Home view. This pane is only available in the Home view. If you do not have sufficient memory to upgrade to the most current release, syslog message %ASA-1-211004 is generated, indicating what the installed memory is and what the required memory is. This message reappears every 24 hours until the memory is upgraded.
Addresses	Shows and hides the display of the **Addresses** pane. The **Addresses** pane is only available for the **Access Rules**, **NAT Rules**, **Service Policy Rules**, **AAA Rules**, and **Filter Rules** panes in the Configuration view.
Services	Shows and hides the display of the **Services** pane. The **Services** pane is only available for the **Access Rules**, **NAT Rules**, **Service Policy Rules**, **AAA Rules**, and **Filter Rules** panes in the Configuration view.
Time Ranges	Shows and hides the display of the **Time Ranges** pane. The **Time Ranges** pane is only available for the **Access Rules**, **Service Policy Rules**, **AAA Rules**, and **Filter Rules** panes in the Configuration view.

View Menu Item	Description
Select Next Pane	Highlights the next pane shown in a multi-pane display, for example, going from the Service Policies Rules pane to the Address pane beside it.
Select Previous Pane	Highlights the previous pane shown in multi-pane displays.
Back	Returns to the previous pane.
Forward	Goes to the next pane previously visited.
Find in ASDM	Locates an item for which you are searching, such as a feature or the **ASDM Assistant**.
Reset Layout	Returns the layout to the default configuration.
Office Look and Feel	Changes the screen fonts and colors to the Microsoft Office settings.

Tools Menu

The **Tools** menu provides you with the following series of tools to use in ASDM.

Tools Menu Item	Description
Command Line Interface	Sends commands to the ASA and view the results.
Show Commands Ignored by ASDM on Device	Displays unsupported commands that have been ignored by ASDM.
Packet Tracer	Traces a packet from a specified source address and interface to a destination. You can specify the protocol and port of any type of data and view the lifespan of a packet, with detailed information about actions taken on it. See the firewall configuration guide for more information.
Ping	Verifies the configuration and operation of the ASA and surrounding communications links, as well as performs basic testing of other network devices. See the firewall configuration guide for more information.
Traceroute	Determines the route that packets will take to their destination. See the firewall configuration guide for more information.
File Management	Views, moves, copies, and deletes files stored in flash memory. You can also create a directory in flash memory. You can also transfer files between various file systems, including TFTP, flash memory, and your local PC.
Check for ASA/ASDM Updates	Upgrades ASA software and ASDM software through a wizard.
Upgrade Software from Local Computer	Uploads an ASA image, ASDM image, or another image on your PC to flash memory.

Tools Menu Item	Description
Downgrade Software	Loads an older ASA image than the one you are currently running.
Backup Configurations	Backs up the ASA configuration, a Cisco Secure Desktop image, and SSL VPN Client images and profiles.
Restore Configurations	Restores the ASA configuration, a Cisco Secure Desktop image, and SSL VPN Client images and profiles.
System Reload	Restarts the ASDM and reload the saved configuration into memory.
Administrator's Alert to Clientless SSL VPN Users	Enables an administrator to send an alert message to clientless SSL VPN users. See the VPN configuration guide for more information.
Migrate Network Object Group Members	If you migrate to 8.3 or later, the ASA creates named network objects to replace inline IP addresses in some features. In addition to named objects, ASDM automatically creates non-named objects for any IP addresses used in the configuration. These auto-created objects are identified by the *IP address* only, do not have a name, and are not present as named objects in the platform configuration. When the ASA creates named objects as part of the migration, the matching non-named ASDM-only objects are replaced with the named objects. The only exception are non-named objects in a network object group. When the ASA creates named objects for IP addresses that are inside a network object group, ASDM retains the non-named objects as well, creating duplicate objects in ASDM. Choose **Tools > Migrate Network Object Group Members** to merge these object. See *Cisco ASA 5500 Migration to Version 8.3 and Later* for more information.
Preferences	Changes the behavior of specified ASDM functions between sessions.
ASDM Java Console	Shows the Java console.

Wizards Menu

The **Wizards** menu lets you run a wizard to configure multiple features.

Wizards Menu Item	Description
Startup Wizard	Guides you, step-by-step, through the initial configuration of the ASA.
VPN Wizards	There are separate wizards for a variety of VPN configurations. See the VPN configuration guide for more information.
High Availability and Scalability Wizard	Allows you to configure failover: VPN cluster load balancing, or ASA clustering on the ASA.

Wizards Menu Item	Description
Unified Communication Wizard	Enables you to configure unified communication features, such as an IP phone, on the ASA. See the firewall configuration guide for more information.
ASDM Identity Certificate Wizard	When using Java 7 update 51 and later, the ASDM Launcher requires a trusted certificate. An easy approach to fulfill the certificate requirements is to install a self-signed identity certificate. You can use Java Web Start to launch ASDM until you install a certificate using this wizard. See http://www.cisco.com/go/asdm-certificate for more information.
Packet Capture Wizard	Allows you to configure packet capture on the ASA. The wizard runs one packet capture on each ingress and egress interface. After you run the capture, you can save it on your computer, and then examine and analyze the capture with a packet analyzer.

Window Menu

The **Window** menu enables you to move between ASDM windows. The active window appears as the selected window.

Help Menu

The **Help** menu provides links to online help, as well as information about ASDM and the ASA.

Help Menu Items	Description
Help Topics	Opens a new browser window to show the ASDM online help. If you are managing an ASA FirePOWER module in ASDM, this item is labeled **ASDM Help Topics**.
ASA FirePOWER Help Topics	Opens a new browser window to show online help for the ASA FirePOWER module. This item is available only if you have installed the module and are managing it in ASDM.
Help for Current Screen	Opens context-sensitive help about the screen you are viewing. Alternatively, you can also click the **? Help** button in the tool bar.
Release Notes	Opens the most current version of the *ASDM release notes* on Cisco.com. The release notes contain the most current information about ASDM software and hardware requirements, and the most current information about changes in the software.
Cisco ASA Series Documentation	Opens a document on Cisco.com that includes links to all of the available product documentation.
ASDM Assistant	Opens the **ASDM Assistant**, which lets you search downloadable content from Cisco.com, with details about performing certain tasks.

Help Menu Items	Description
About Cisco Adaptive Security Appliance (ASA)	Displays information about the ASA, including the software version, hardware set, configuration file loaded at startup, and software image loaded at startup. This information is helpful in troubleshooting.
About Cisco ASDM	Displays information about ASDM such as the software version, hostname, privilege level, operating system, device type, and Java version.

Toolbar

The **Toolbar** below the menus provides access to the Home view, Configuration view, and Monitoring view. It also lets you choose between the system and security contexts in multiple context mode, and provides navigation and other commonly used features.

Toolbar Button	Description
Home	Displays the **Home** pane, which lets you view important information about your ASA such as the status of your interfaces, the version you are running, licensing information, and performance. In multiple mode, the system does not have a Home pane.
Configuration	Configures the ASA. Click a function button in the left **Navigation** pane to configure that function.
Monitoring	Monitors the ASA. Click a function button in the left **Navigation** pane to monitor various elements.
Save	Saves the running configuration to the startup configuration for write-accessible contexts only. The button is replaced by the **Deploy** button if you have an ASA FirePOWER module installed on the device and you are configuring it through ASDM.
Deploy	If you have an ASA FirePOWER module installed on the device and you are configuring it through ASDM, the **Deploy** button replaces the **Save** button and contains the following options: • **Deploy FirePOWER Changes**—Saves configuration changes made to ASA FirePOWER module policies to the module. • **Save Running Configuration to Flash**—Writes a copy of the ASA running configuration to flash memory. This is equivalent to the **Save** button for devices that do not include an ASA FirePOWER module.
Refresh	Refreshes ASDM with the current running configuration, except for graphs in any of the Monitoring panes.
Back	Returns to the last pane of ASDM that you visited.
Forward	Goes forward to the last pane of ASDM that you visited.

Toolbar Button	Description
Help	Shows context-sensitive help for the screen that is currently open.
Search	Searches for a feature in ASDM. The Search function looks through the titles of each pane and presents you with a list of matches, and gives you a hyperlink directly to that pane. Click **Back** or **Forward** to switch quickly between two different panes that you found.

ASDM Assistant

The ASDM Assistant lets you search and view useful ASDM procedural help about certain tasks. This feature is available in routed and transparent modes, and in the single and system contexts.

Choose**View** > **ASDM Assistant** > **How Do I?** or enter a search request from the **Look For** field in the menu bar to access information. Choose **How Do I?** from the **Find** drop-down list to begin the search.

To use the ASDM Assistant, perform the following steps:

Procedure

Step 1 Choose **View** > **ASDM Assistant**.

The **ASDM Assistant** pane appears.

Step 2 Enter the information that you want to find in the **Search** field, then click **Go**.

The requested information appears in the **Search Results** pane.

Step 3 Click any links that appear in the **Search Results and Features** areas to obtain more details.

Status Bar

The **Status Bar** appears at the bottom of the ASDM window. The following table lists the areas shown from left to right.

Area	Description
Status	The status of the configuration (for example, “Device configuration loaded successfully.”)
Failover	The status of the failover unit, either active or standby.
User Name	The username of the ASDM user. If you logged in without a username, the username is “admin.”
User Privilege	The privilege of the ASDM user.

Area	Description
Commands Ignored by ASDM	Click the icon to show a list of commands from your configuration that ASDM did not process. These commands will not be removed from the configuration.
Connection to Device	The ASDM connection status to the ASA.
Syslog Connection	The syslog connection is up, and the ASA is being monitored.
SSL Secure	The connection to ASDM is secure because it uses SSL.
Time	The time that is set on the ASA.

Connection to Device

ASDM maintains a constant connection to the ASA to maintain up-to-date **Monitoring** and **Home** pane data. This dialog box shows the status of the connection. When you make a configuration change, ASDM opens a second connection for the duration of the configuration, and then closes it; however, this dialog box does not represent the second connection.

Device List

The Device List is a dockable pane. You can click one of the three buttons in the header to maximize or restore this pane, make it a floating pane that you can move, hide it, or close it. This pane is available in the Home, Configuration, Monitoring, and System views. You can use this pane to switch to another device, and between the System and contexts; however, that device must run the same version of ASDM that you are currently running. To display the pane fully, you must have at least two devices listed. This pane is available in routed and transparent modes, and in the single, multiple, and system contexts.

To use this pane to connect to another device, perform the following steps:

Procedure

Step 1 Click **Add** to add another device to the list.

The **Add Device** dialog box appears.

Step 2 Enter the device name or IP address of the device, then click **OK**.

Step 3 Click **Delete** to remove a selected device from the list.

Step 4 Click **Connect** to connect to another device.

The **Enter Network Password** dialog box appears.

Step 5 Enter your username and password in the applicable fields, then click **Login**.

Common Buttons

Many ASDM panes include buttons that are listed in the following table. Click the applicable button to complete the desired task:

Button	Description
Apply	Sends changes made in ASDM to the ASA and applies them to the running configuration.
Save	Writes a copy of the running configuration to flash memory.
Reset	Discards changes and reverts to the information displayed before changes were made or the last time that you clicked Refresh or Apply. After you click **Reset**, click **Refresh** to make sure that information from the current running configuration appears.
Restore Default	Clears the selected settings and returns to the default settings.
Cancel	Discards changes and returns to the previous pane.
Enable	Displays read-only statistics for a feature.
Close	Closes an open dialog box.
Clear	Remove information from a field, or remove a check from a check box.
Back	Returns to the previous pane.
Forward	Goes to the next pane.
Help	Displays help for the selected pane or dialog box.

Keyboard Shortcuts

You can use the keyboard to navigate the ASDM user interface.

The following table lists the keyboard shortcuts you can use to move across the three main areas of the ASDM user interface.

Table 2: Keyboard Shortcuts Within the Main Window

To display the	Windows/Linux	MacOS
Home Pane	Ctrl+H	Shift+Command+H
Configuration Pane	Ctrl+G	Shift+Command+G
Monitoring Pane	Ctrl+M	Shift+Command+M
Help	F1	Command+?

To display the	Windows/Linux	MacOS
Back	Alt+Left Arrow	Command+[
Forward	Alt+Rightarrow	Command+]
Refresh the display	F5	Command+R
Cut	Ctrl+X	Command+X
Copy	Ctrl+C	Command+C
Paste	Ctrl+V	Command+V
Save the configuration	Ctrl+S	Command+S
Popup menus	Shift+F10	—
Close a secondary window	Alt+F4	Command+W
Find	Ctrl+F	Command+F
Exit	Alt+F4	Command+Q
Exit a table or text area	Ctrl_Shift or Ctrl+Shift+Tab	Ctril+Shift or Ctrl+Shift+Tab

The following table lists the keyboard shortcut you can use to navigate within a pane.

Table 3: Keyboard Shortcuts Within a Pane

To move the focus to the	Press
Next field	Tab
Previous field	Shift+Tab
Next field when the focus is in a table	Ctrl+Tab
Previous field when the focus is in a table	Shift+Ctrl+Tab
Next tab (when a tab has the focus)	Right Arrow
Previous tab (when a tab has the focus)	Left Arrow
Next cell in a table	Tab
Previous sell in a table	Shift+Tab
Next pane (when multiple panes are displayed)	F6
Previous pane (when multiple panes are displayed)	Shift+F6

The following table lists the keyboard shortcuts you can use with the Log Viewers.

Table 4: Keyboard Shortcuts for the Log Viewer

To	Windows/Linux	MacOS
Pause and Resume Real-Time Log Viewer	Ctrl+U	Command+
Refresh Log Buffer Pane	F5	Command+R
Clear Internal Log Buffer	Ctrl+Delete	Command+Delete
Copy Selected Log Entry	Ctrl+C	Command+C
Save Log	Ctrl+S	Command+S
Print	Ctrl+P	Command+P
Close a secondary window	Alt+F4	Command+W

The following table lists the keyboard shortcuts you can use to access menu items.

Table 5: Keyboard Shortcuts to Access Menu Items

To access the	Windows/Linux
Menu Bar	Alt
Next Menu	Right Arrow
Previous Menu	Left Arrow
Next Menu Option	Down Arrow
Previous Menu Option	Up Arrow
Selected Menu Option	Enter

Find Function in ASDM Panes

Some ASDM panes contain tables with many elements. To make it easier for you to search, highlight, and then edit a particular entry, several ASDM panes have a find function that allows you to search on objects within those panes.

To perform a search, you can type a phrase into the Find field to search on all columns within any given pane. The phrase can contain the wild card characters "*" and "?". The * matches one or more characters, and ? matches one character. The up and down arrows to the right of the **Find** field locate the next (up) or previous (down) occurrence of the phrase. Check the **Match Case** check box to find entries with the exact uppercase and lowercase characters that you enter.

For example, entering B*ton-L* might return the following matches:

```
Boston-LA, Boston-Lisbon, Boston-London
```

Entering Bo?ton might return the following matches:

```
Boston, Bolton
```

Find Function in Rule Lists

Because ACLs and ACEs and other rules contain many elements of different types, the find function in the any pane that displays rules allows for a more targeted search than the find function in other panes. This includes access rules, service policy rules, the ACL Manager, and any other pane that lists ACL rules, and also the NAT rules.

To find elements within the rule lists, perform the following steps:

Procedure

Step 1 Click **Find**.

Step 2 Choose one of the following options in the **Filter** field from the drop-down list.

The items you can search on differ depending on the rule type, and correspond to the columns in the table. Select **Query** if you want to create a complex search that uses more than one field.

Step 3 Unless you picked **Query**, in the second field, choose one of the following options from the drop-down list:

- **is**—Specifies an exact match to the search string. This is always the option for queries.
- **contains**—Specifies a match to any rule that includes, whether exactly or partially, your search string.

Step 4 In the third field, enter the string you want to find. Click **...** to pick an object from a list. If you are using a query, click **Define Query**.

If you search for an IP address, you can get matches to addresses that are in a network object or group, so long as that object or group was created by ASDM. That is, the group name begins with DM_INLINE. The find feature cannot find IP addresses within user-created objects.

Step 5 Click **Filter** to perform the search.

The view is updated to show only those rules that match. The rule numbers are maintained so that you can see their absolute location within the rule list.

Step 6 Click **Clear** to remove the filter and see the complete list again.

Step 7 When you are finished, click the red **x** to close the find controls.

Enable Extended Screen Reader Support

By default, labels and descriptions are not included in tab order when you press the **Tab** key to navigate a pane. Some screen readers, such as JAWS, only read screen objects that have the focus. You can include the labels and descriptions in the tab order by enabling extended screen reader support.

To enable extended screen reader support, perform the following steps:

Procedure

Step 1 Choose **Tools** > **Preferences**.

The **Preferences** dialog box appears.

Step 2 Check the **Enable screen reader support** check box on the **General** tab.

Step 3 Click **OK**.

Step 4 Restart ASDM to activate screen reader support.

Organizational Folder

Some folders in the navigation pane for the configuration and monitoring views do not have associated configuration or monitoring panes. These folders are used to organize related configuration and monitoring tasks. Clicking these folders displays a list of sub-items in the right **Navigation** pane. You can click the name of a sub-item to go to that item.

Home Pane (Single Mode and Context)

The ASDM **Home** pane lets you view important information about your ASA. Status information in the **Home** pane is updated every ten seconds. This pane usually has two tabs: **Device Dashboard** and **Firewall Dashboard**.

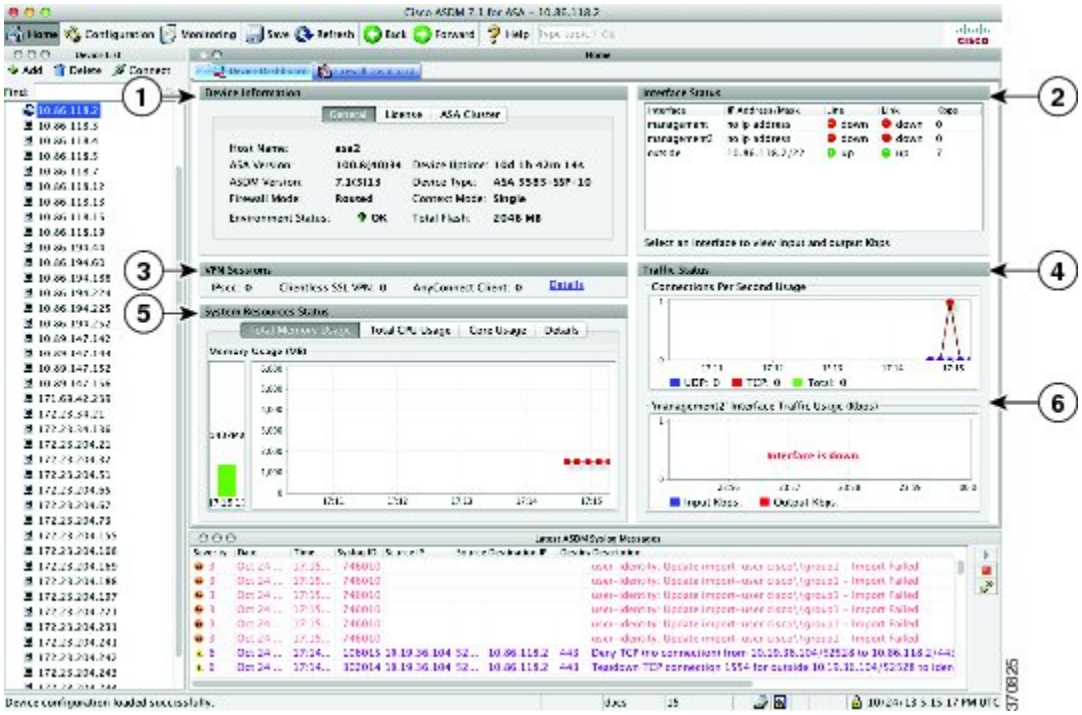
If you have hardware or software modules installed on the device, such as IPS, CX, or ASA FirePOWER modules, there are separate tabs for those modules.

Device Dashboard Tab

The **Device Dashboard** tab lets you view, at a glance, important information about your ASA, such as the status of your interfaces, the version you are running, licensing information, and performance.

The following figure shows the elements of the **Device Dashboard** tab.

Figure 6: Device Dashboard Tab



Legend

GUI Element	Description
1	Device Information Pane, on page 64
2	Interface Status Pane, on page 66
3	VPN Sessions Pane, on page 66
4	Traffic Status Pane, on page 66
5	System Resources Status Pane, on page 66
6	Traffic Status Pane, on page 66
—	Device List, on page 58
—	Latest ASDM Syslog Messages Pane, on page 66

Device Information Pane

The **Device Information** pane includes two tabs that show device information: **General** tab and **License** tab. Under the **General** tab you have access to the **Environment Status** button, which provides an at-a-glance view of the system health:

General Tab

This tab shows basic information about the ASA:

- **Host name**—Shows the hostname of the device.
- **ASA version**—Lists the version of ASA software that is running on the device.
- **ASDM version**—Lists the version of ASDM software that is running on the device.
- **Firewall mode**—Shows the firewall mode in which the device is running.
- **Total flash**—Displays the total RAM that is currently being used.
- **ASA Cluster Role**—When you enable clustering, shows the role of this unit, either Master or Slave.
- **Device uptime**—Shows the time in which the device has been operational since the latest software upload.
- **Context mode**—Shows the context mode in which the device is running.
- **Total Memory**—Shows the DRAM installed on the ASA.
- **Environment status**—Shows the system health. View hardware statistics by clicking the plus sign (+) to the right of the **Environment Status** label in the **General** tab. You can see how many power supplies are installed, track the operational status of the fan and power supply modules, and track the temperatures of the CPUs and the ambient temperature of the system.

  In general, the **Environment Status** button provides an at-a-glance view of the system health. If all monitored hardware components within the system are operating within normal ranges, the plus sign (+) button shows OK in green. Conversely, if any one component within the hardware system is operating outside of normal ranges, the plus sign (+) button turns into a red circle to show Critical status and to indicate that a hardware component requires immediate attention.

  See the hardware guide for your particular device for more information about specific hardware statistics.



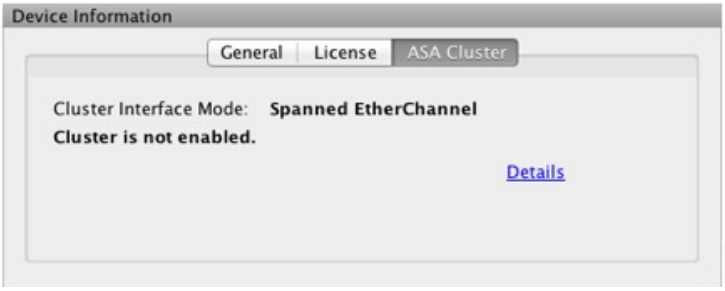
Note If you do not have enough memory to upgrade to the most current release of the ASA, the **Memory Insufficient Warning** dialog box appears. Follow the directions that appear in this dialog box to continue using the ASA and ASDM in a supported manner. Click **OK** to close this dialog box.

License Tab

This tab shows a subset of licensed features. Click **More Licenses** to view detailed license information, or to enter a new activation key; the **Configuration > Device Management > Licensing > Activation Key pane** appears.

Cluster Tab

This tab shows the cluster interface mode, as well as the cluster status



Virtual Resources Tab (ASAv)

This tab shows the virtual resources used by the ASA virtual, including the number of vCPUs, RAM, and whether the ASA virtual is over- or under-provisioned.

Interface Status Pane

This pane shows the status of each interface. If you select an interface row, the input and output throughput in Kbps displays below the table.

VPN Sessions Pane

This pane shows the VPN tunnel status. Click **Details** to go to the **Monitoring > VPN > VPN Statistics > Sessions** pane.

Failover Status Pane

This pane shows the failover status.

Click **Configure** to start the High Availability and Scalability Wizard. After you have completed the wizard, the failover configuration status (either Active/Active or Active/Standby) appears.

If failover is configured, click **Details** to open the **Monitoring > Properties > Failover > Status** pane.

System Resources Status Pane

This pane shows CPU and memory usage statistics.

Traffic Status Pane

This pane shows graphs for connections per second for all interfaces and for the traffic throughput of the lowest security interface.

When your configuration contains multiple lowest security level interfaces, and any one of them is named "outside," then that interface is used for the traffic throughput graphs. Otherwise, ASDM picks the first interface from the alphabetical list of lowest security level interfaces.

Latest ASDM Syslog Messages Pane

This pane shows the most recent system messages generated by the ASA, up to a maximum of 100 messages. Click **Enable Logging** to enable logging if it is disabled.

The following figure shows the elements of the **Latest ASDM Syslog Messages** pane.

Figure 7: Latest ASDM Syslog Messages Pane



Legend

GUI Element	Description
1	Drag the **divider** up or down to resize the pane.
2	Expands the pane. Click the **double-square** icon to return the pane to the default size.
3	Makes a floating pane. Click the **docked pane** icon to dock the pane.
4	Enables or disables Auto-hide. When Auto-hide is enabled, move your cursor over the **Latest ASDM Syslog Messages** button in the left, bottom corner and the pane displays. Move your cursor away from the pane, and it disappears.
5	Closes the pane. Choose **View Latest ASDM Syslog Messages** to show the pane.
6	Click the **green** icon on the right-hand side to continue updating the display of syslog messages.
7	Click the **red** icon on the right-hand side To stop updating the display of syslog messages.
8	Click the **filters** icon on the right-hand side to open the **Logging Filters** pane.

- Right-click an event and choose **Clear Content** to clear the current messages.
- Right-click an event and click **Save Content** to save the current messages to a file on your PC.
- Right-click an event and choose **Copy** to copy the current content.
- Right-click an event and choose **Color Settings** to change the background and foreground colors of syslog messages according to their severity.

Firewall Dashboard Tab

The **Firewall Dashboard** tab lets you view important information about the traffic passing through your ASA. This dashboard differs depending on whether you are in single context mode or multiple context mode. In multiple context mode, the **Firewall Dashboard** is viewable within each context.

The following figure shows some of the elements of the **Firewall Dashboard** tab.

Figure 8: Firewall Dashboard Tab



Legend

GUI Element	Description
1	Traffic Overview Pane, on page 68
2	Top 10 Access Rules Pane, on page 69
3	Top Usage Status Pane, on page 69
(not shown)	Top Ten Protected Servers Under SYN Attack Pane, on page 69
(not shown)	Top 200 Hosts Pane, on page 70
(not shown)	Top Botnet Traffic Filter Hits Pane, on page 70

Traffic Overview Pane

Enabled by default. If you disable basic threat detection (see the firewall configuration guide), then this area includes an **Enable** button that lets you enable basic threat detection. The runtime statistics include the following information, which is *display-only*:

- The number of connections and NAT translations.
- The rate of dropped packets per second caused by access list denials and application inspections.
- The rate of dropped packets per second that are identified as part of a scanning attack, or that are incomplete sessions detected, such as TCP SYN attack detected or no data UDP session attack detected.

Top 10 Access Rules Pane

Enabled by default. If you disable threat detection statistics for access rules (see the firewall configuration guide), then this area includes an **Enable** button that lets you enable statistics for access rules.

In the Table view, you can select a rule in the list and right-click the rule to display a popup menu item, **Show Rule**. Choose this item to go to the Access Rules table and select that rule in this table.

Top Usage Status Pane

Disabled by default. This pane include the following four tabs:

- **Top 10 Services**—Threat Detection service
- **Top 10 Sources**—Threat Detection service
- **Top 10 Destinations**—Threat Detection service
- **Top 10 Users**—Identity Firewall service

The first three tabs—**Top 10 Services**, **Top 10 Sources**, and **Top 10 Destinations**—provide statistics for threat detection services. Each tab includes an **Enable** button that let you enable each threat detection service. You can enable them according to the firewall configuration guide.

The **Top 10 Services Enable** button enables statistics for both ports and protocols (both must be enabled for the display). The **Top 10 Sources** and **Top 10 Destinations Enable** buttons enable statistics for hosts. The top usage status statistics for hosts (sources and destinations), and ports and protocols are displayed.

The fourth tab for **Top 10 Users** provides statistics for the Identity Firewall service. The Identity Firewall service provides access control based on users' identities. You can configure access rules and security policies based on user names and user groups name rather than through source IP addresses. The ASA provides this service by accessing an IP-user mapping database.

The **Top 10 Users** tab displays data only when you have configured one of the following features:

- Identity Firewall service configuration, which includes configuring these additional components: Microsoft Active Directory and Cisco Active Directory (AD) Agent. The Identity Firewall service is enabled using the **user-identity enable** command (enabled by default) and the **user-accounting statistics** command.
- VPN configuration using a RADIUS server for authenticating, authorizing, or accounting VPN users.

Depending on which option you choose, the **Top 10 Users** tab shows statistics for received EPS packets, sent EPS packets, and sent attacks for the top 10 users. For each user (displayed as *domain\user_name*), the tab displays the average EPS packet, the current EPS packet, the trigger, and total events for that user.



Caution Enabling statistics can affect the ASA performance, depending on the type of statistics enabled. Enabling statistics for hosts affects performance in a significant way; if you have a high traffic load, you might consider enabling this type of statistics temporarily. Enabling statistics for ports, however, has a modest effect.

Top Ten Protected Servers Under SYN Attack Pane

Disabled by default. This area includes an **Enable** button that lets you enable the feature, or you can enable it according to the firewall configuration guide. Statistics for the top ten protected servers under attack are displayed.

For the average rate of attack, the ASA samples the data every 30 seconds over the rate interval (by default 30 minutes).

If there is more than one attacker, then "<various>" displays, followed by the last attacker IP address.

Click **Detail** to view statistics for all servers (up to 1000) instead of just 10 servers. You can also view history sampling data. The ASA samples the number of attacks 60 times during the rate interval, so for the default 30-minute period, statistics are collected every 60 seconds.

Top 200 Hosts Pane

Disabled by default. Shows the top 200 hosts connected through the ASA. Each entry of a host contains the IP address of the host and the number of connections initiated by the host, and is updated every 120 seconds. Enter the **hpm topnenable** command to enable this display.

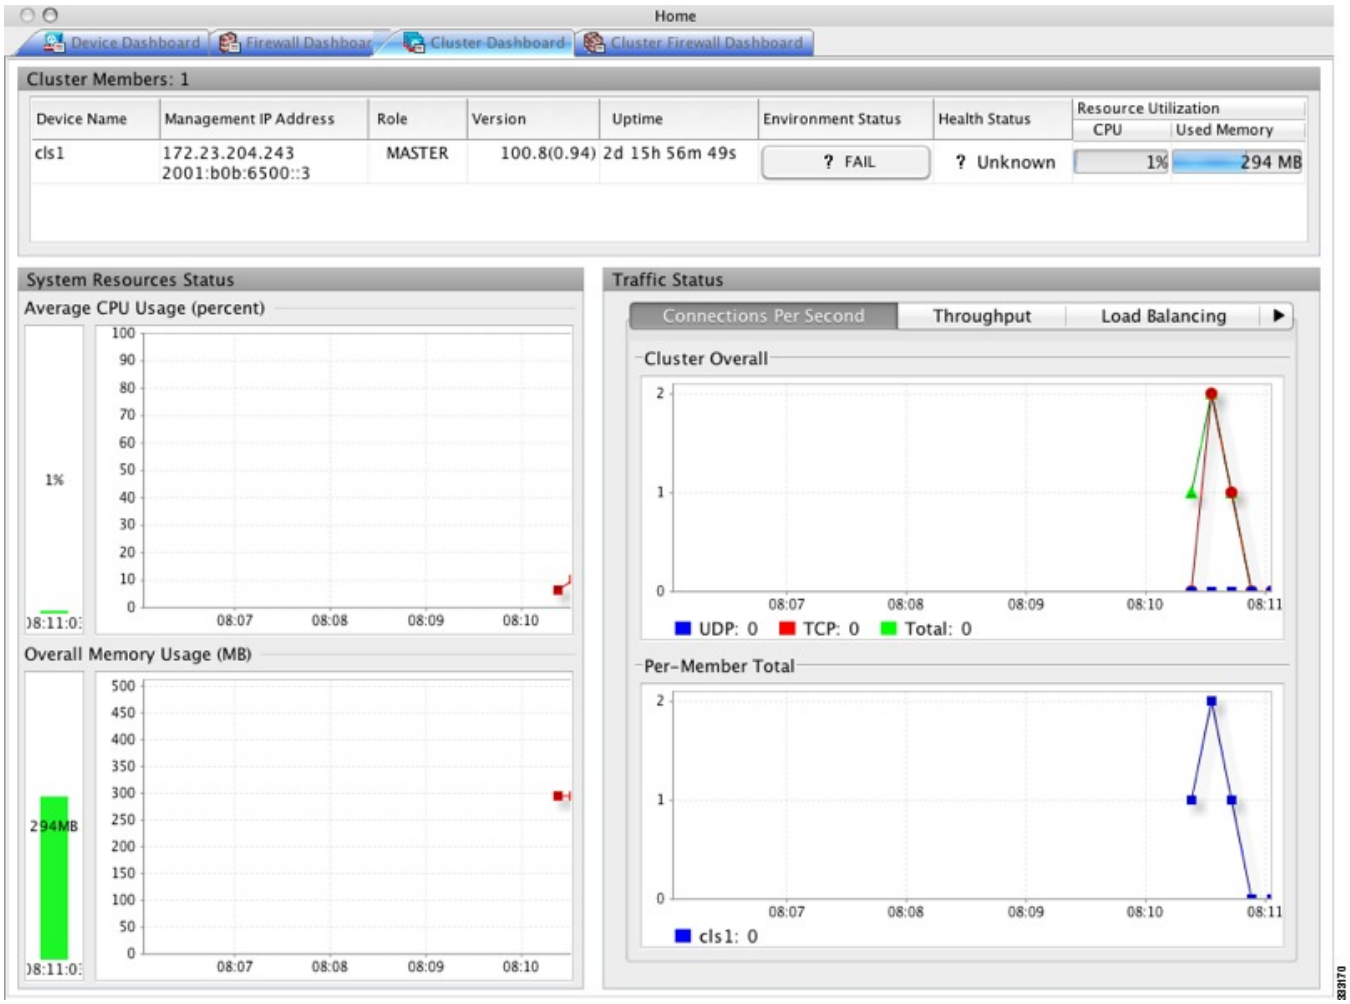
Top Botnet Traffic Filter Hits Pane

Disabled by default. This area includes links to configure the Botnet Traffic Filter. Reports of the top ten botnet sites, ports, and infected hosts provide a snapshot of the data, and may not match the top ten items since statistics started to be collected. If you right-click an IP address, you can invoke the whois tool to learn more about the botnet site.

See the Botnet configuration guide for more information.

Cluster Dashboard Tab

When you enable ASA clustering and are connected to the master unit, the **Cluster Dashboard** tab shows a summary of cluster membership and resource utilization.



- **Cluster Members**—Shows the names and basic information about the members comprising the cluster (their management IP address, version, role in the cluster, and so on) and their health status (environment status, health status, and resource utilization).

Note In multiple context mode, if you connect ASDM to the admin context, and then change to a different context, the management IP address listed does not change to show the current context management IP addresses; it continues to show the admin context management IP addresses, including the main cluster IP address to which ASDM is currently connected.

- **System Resource Status**—Shows resource utilization (CPU and memory) across the cluster and traffic graphs, both cluster-wide and per-device.
- **Traffic Status**—Each tab has the following graphs.
  - **Connections Per Second** tab:

    Cluster Overall—Shows the connections per second throughout the cluster.

Per-Member Total—Shows the average connections per second for each member.

- **Throughput** tab:

  Cluster Overall—Shows the aggregated egress throughput throughout the cluster.

  Per-Member Throughput—Shows the member throughput, one line per member.

- **Load Balancing** tab:

  Per-Member Percentage of Total Traffic—For each member, shows the percentage of total cluster traffic that the member receives.

  Per-Member Locally Processed Traffic—For each member, shows the percentage of traffic that was processed locally.

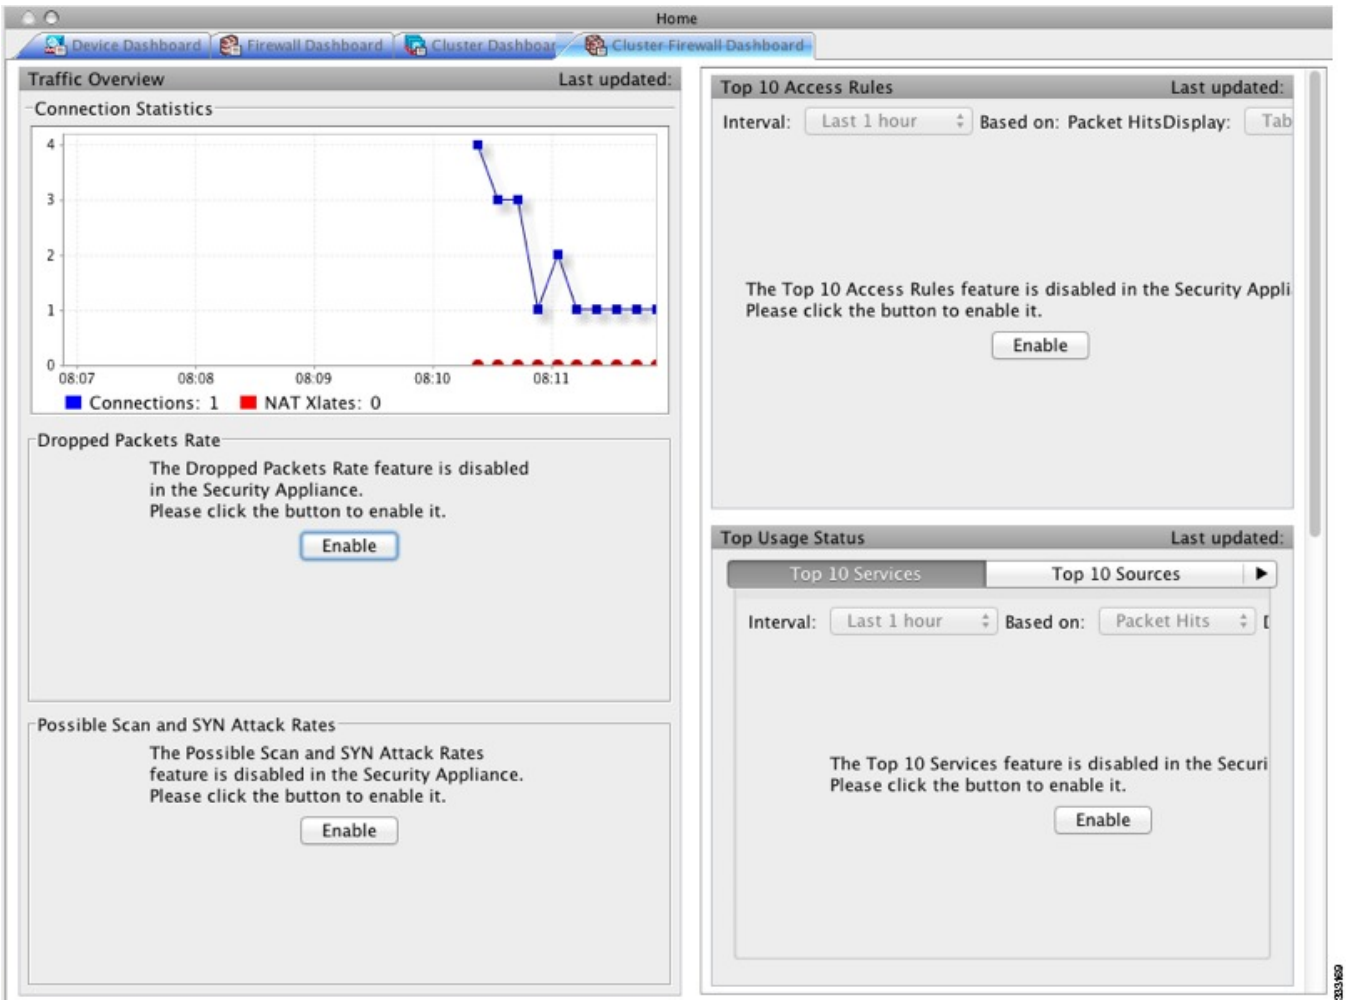
- **Control Link Usage** tab:

  Per-Member Receival Capacity Utilization—For each member, shows the usage of the transmittal capacity.

  Per-Member Transmittal Capacity Utilization—For each member, shows the usage of the receival capacity.

Cluster Firewall Dashboard Tab

The **Cluster Firewall Dashboard** tab shows traffic overview and the "top N" statistics, similar to those shown in the **Firewall Dashboard**, but aggregated across the whole cluster.



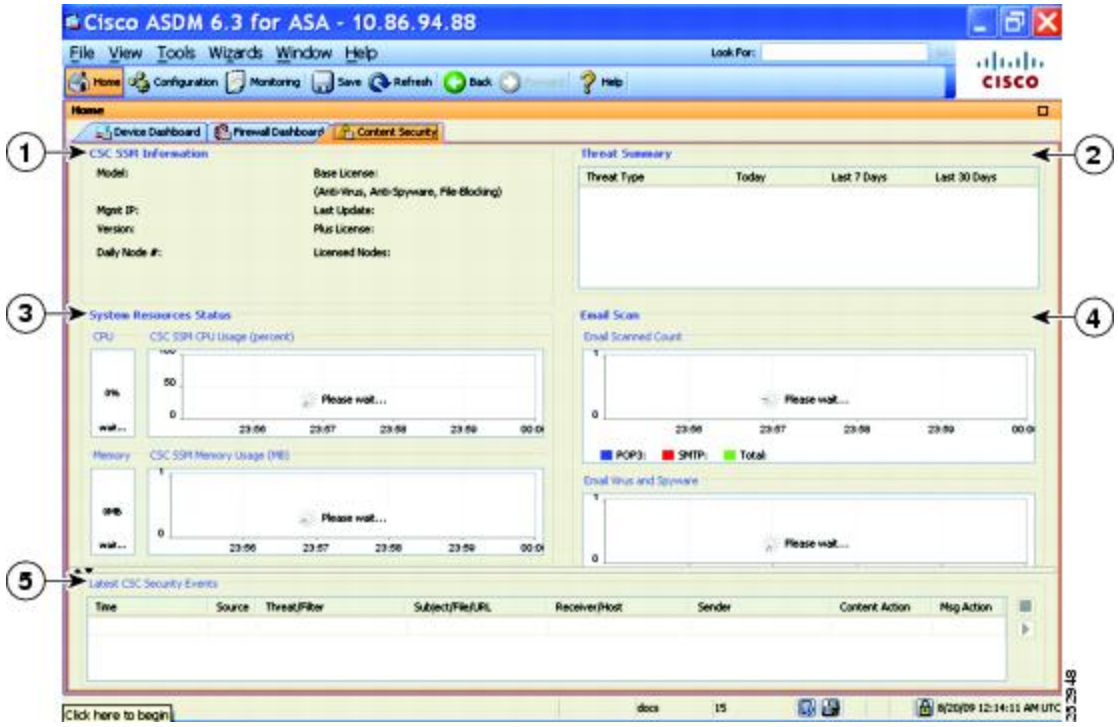
Content Security Tab

The **Content Security** tab lets you view important information about the Content Security and Control (CSC) SSM. This pane appears only if CSC software running on the CSC SSM is installed in the ASA.

Note If you have not completed the **CSC Setup Wizard** by choosing **Configuration > Trend Micro Content Security > CSC Setup**, you cannot access the panes under **Home > Content Security**. Instead, a dialog box appears and lets you access the **CSC Setup Wizard** directly from this location.

The following figure shows the elements of the **Content Security** tab.

Figure 9: Content Security Tab



Legend

GUI Element	Description
1	**CSC SSM Information** pane.
2	**Threat Summary** pane. Shows aggregated data about threats detected by the CSC SSM, including the following threat types: Virus, Spyware, URL Filtered or Blocked, Spam. Blocked, Files Blocked, and Damage Control Services.
3	**System Resources Status** pane.
4	**Email Scan** pane. The graphs display data in ten-second intervals.
5	**Latest CSC Security Events** pane.

Intrusion Prevention Tab

The **Intrusion Prevention** tab lets you view important information about IPS. This tab appears only when you have an IPS module installed on the ASA.

To connect to the IPS module, perform the following steps:

1. Click the **Intrusion Prevention** tab.

   The **Connecting to IPS** dialog box appears.

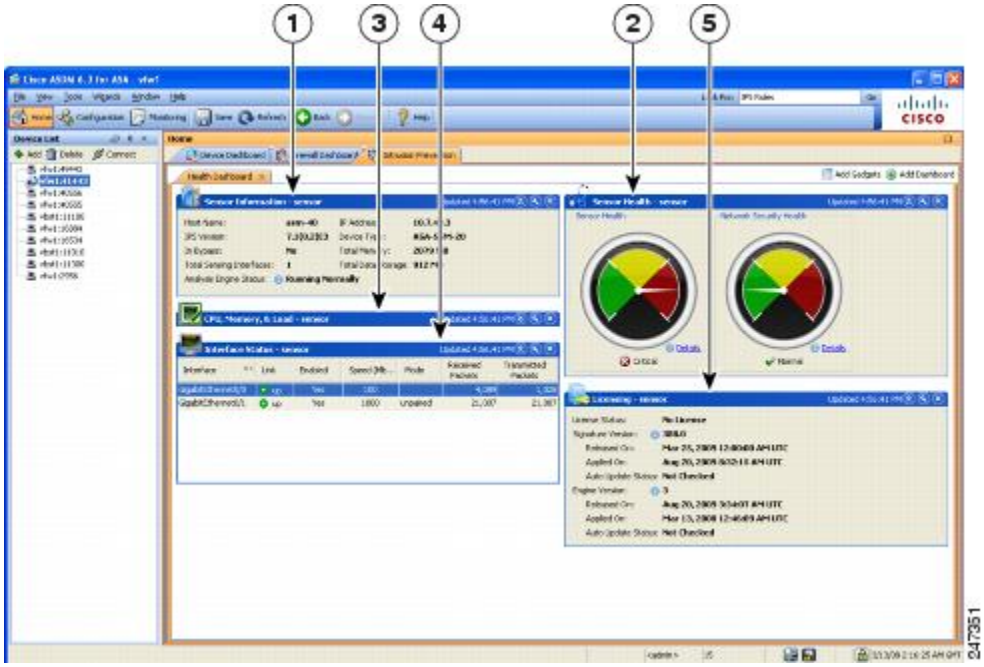


2. Enter the IP address, port, username and password. The default IP address and port is 192.168.1.2:443. The default username and password is **cisco** and **cisco**.
3. Check the **Save IPS login** information on local host check box to save the login information on your local PC.
4. Click **Continue**.

See the IPS quick start guide for more information about intrusion prevention.

The following figure shows the elements of the **Health Dashboard** tab, located on the **Intrusion Prevention** tab.

Figure 10: Intrusion Prevention Tab (Health Dashboard)

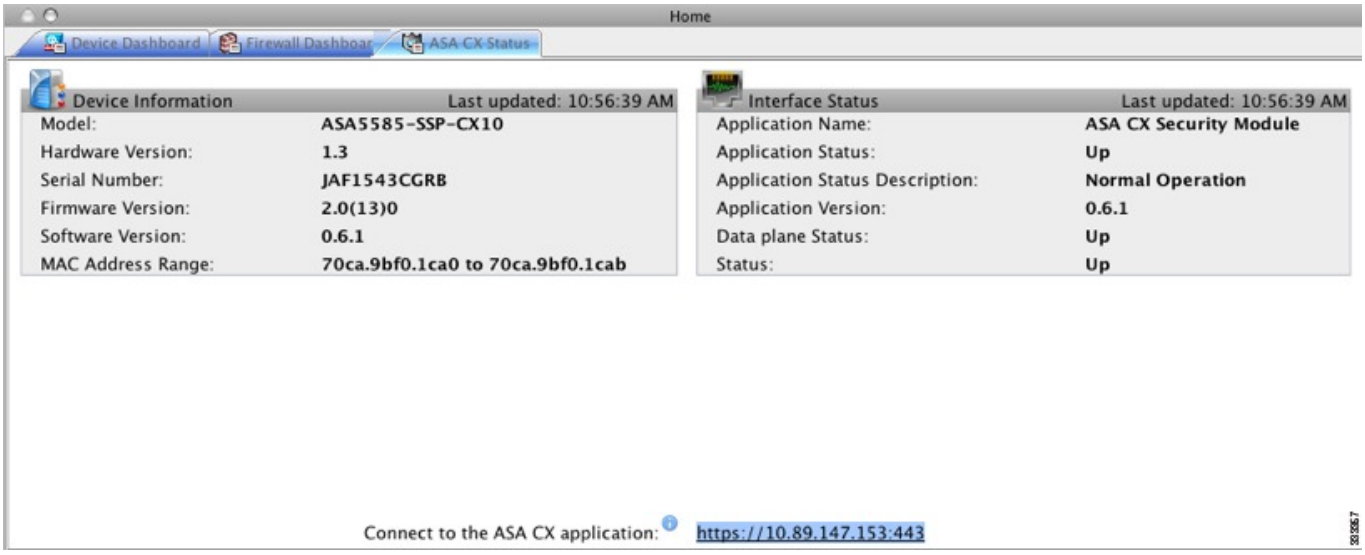


Legend

GUI Element	Description
1	**Sensor Information** pane.
2	**Sensor Health** pane.
3	**CPU**, **Memory**, and **Load** pane.
4	**Interface Status** pane.
5	**Licensing** pane.

ASA CX Status Tab

The **ASA CX Status** tab lets you view important information about the ASA CX module. This tab appears only when you have an ASA CX module installed on the ASA.



ASA FirePower Status Tabs

The **ASA FirePOWER Status** tab lets you view information about the module. This includes module information, such as the model, serial number, and software version, and module status, such as the application name and status, data plane status, and overall status. If the module is registered to a FireSIGHT Management Center, you can click the link to open the application and do further analysis and module configuration.

This tab appears only if you have an ASA FirePOWER module installed in the device.

If you are managing the ASA FirePOWER module with ASDM rather than FireSIGHT Management Center, there are additional tabs:

- **ASA FirePOWER Dashboard**—The dashboard provides summary information about the software running on the module, product updates, licensing, system load, disk usage, system time, and interface status.

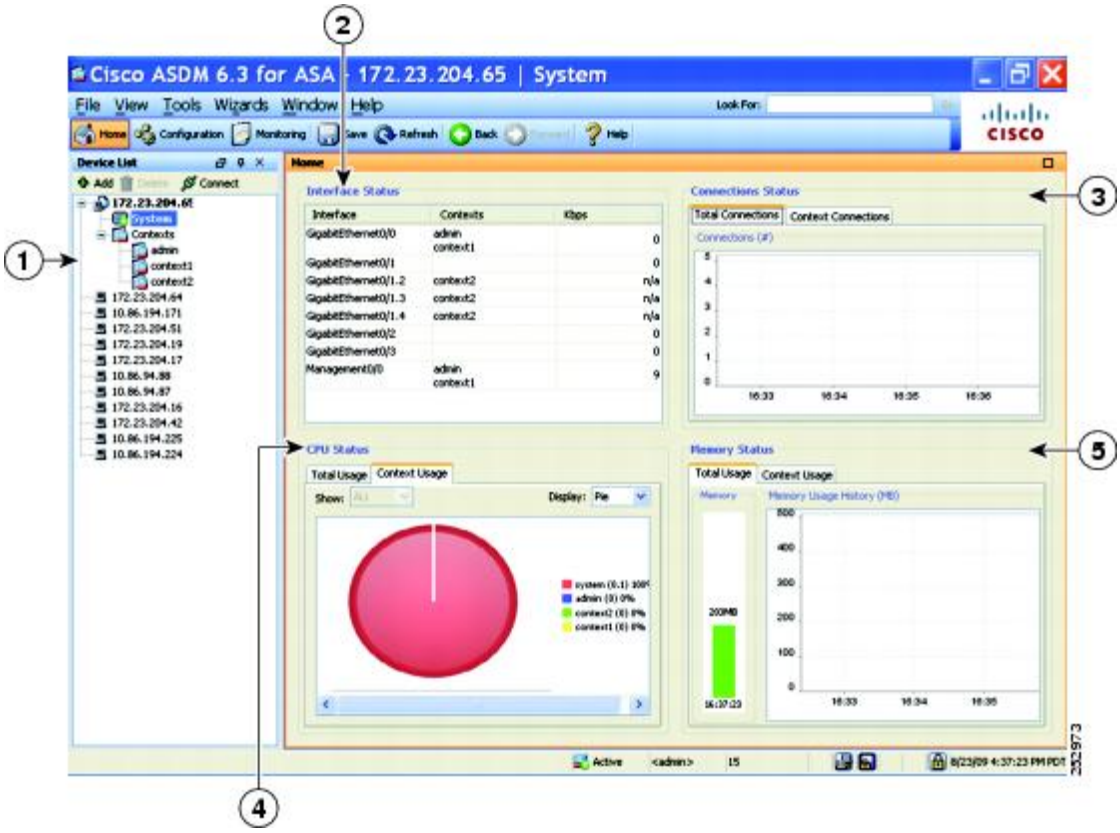
- **ASA FirePOWER Reporting**—The reporting page provides Top 10 dashboards for a wide variety of module statistics, such as web categories, users, sources, and destinations for the traffic passing through the module.

Home Pane (System)

The ASDM System **Home** pane lets you view important status information about your ASA. Many of the details available in the ASDM System **Home** pane are available elsewhere in ASDM, but this pane shows at-a-glance how your ASA is running. Status information in the System **Home** pane is updated every ten seconds.

The following figure shows the elements of the System **Home** pane.

Figure 11: System Home Pane



Legend

GUI Element	Description
1	System vs. Context selection.
2	**Interface Status** pane. Choose an interface to view the total amount of traffic through the interface.

GUI Element	Description
3	**Connection Status** pane.
4	**CPU Status** pane.
5	**Memory Status** pane.

Define ASDM Preferences

You can define the behavior of certain ASDM settings.

To change various settings in ASDM, perform the following steps:

Procedure

Step 1 Choose **Tools** > **Preferences**.

The **Preferences** dialog box appears, with three tabs: **General**, **Rules Table**, and **Syslog**.

Step 2 To define your settings, click one of these tabs: the **General** tab to specify general preferences; the **Rules Table** tab to specify preferences for the Rules table; and the **Syslog** tab to specify the appearance of syslog messages displayed in the **Home** pane and to enable the display of a warning message for NetFlow-related syslog messages.

Step 3 On the **General** tab, specify the following:

- Check the **Warn that configuration in ASDM is out of sync with the configuration in ASA** check box to be notified when the startup configuration and the running configuration are no longer in sync with each other.
- Check the **Show configuration restriction message to read-only user** check box to display the following message to a read-only user at startup. This option is checked by default.

```
"You are not allowed to modify the ASA configuration,
because you do not have sufficient privileges."
```

- Check the **Show configuration restriction message on a slave unit in an ASA cluster** check box to display a configuration restriction message to a user connected to a slave unit.
- Check the **Confirm before exiting ASDM** check box to display a prompt when you try to close ASDM to confirm that you want to exit. This option is checked by default.
- Check the **Enable screen reader support (requires ASDM restart)** check box to enable screen readers to work. You must restart ASDM to enable this option.
- Check the **Warn of insufficient ASA memory when ASDM loads** check box to receive notification when the minimum amount of ASA memory is insufficient to run complete functionality in the ASDM application. ASDM displays the memory warning in a text banner message at bootup, displays a message in the title bar text in ASDM, and sends a syslog alert once every 24 hours.
- In the **Communications** area:

- Check the **Preview commands before sending them to the device** check box to view CLI commands generated by ASDM.
- Check the **Enable cumulative (batch) CLI delivery** check box to send multiple commands in a single group to the ASA.
- In the **Minimum Configuration Sending Timeout** field, enter the minimum amount of time in seconds for a configuration to send a timeout message. The default is 60 seconds.
- For the System in multiple context mode, in the **Graph User time interval in System Context** field, enter the amount of time between updates for the graphs on the Home pane, between 1 and 40 seconds. The default is 10 seconds.

- In the **Logging** area:
  - Check the **Enable logging to the ASDM Java console** check box to configure Java logging.
  - Set the severity level by choosing a **Logging Level** from the drop-down list.
- In the **Packet Capture Wizard** area, to display captured packets, enter the name of the **Network Sniffer Application** or click **Browse** to find it in the file system.
- In the **SFR Location Wizard** area, specify the location to install ASA FirePOWER module local management files. You must have read/write privileges to the configured location.

Step 4 On the **Rules Table** tab, specify the following:

- Display settings let you change the way rules appear in the Rules table.
  - Check the **Auto-expand network and service object groups with specified prefix** check box to display the network and service object groups automatically expanded based on the Auto-Expand Prefix setting.
  - Enter the prefix of the network and service object groups to expand automatically when displayed in the **Auto-Expand Prefix** field.
  - Check the **Show members of network and service object groups** check box to display members of network and service object groups and the group name in the Rules table. If the check box is not checked, only the group name is displayed.
  - Enter the number of network and service object groups to display in the **Limit Members To** field. When the object group members are displayed, then only the first *n* members are displayed.
  - Check the **Show all actions for service policy rules** check box to display all actions in the Rules table. When unchecked, a summary appears.
- Deployment settings let you configure the behavior of the ASA when deploying changes to the Rules table.
  - Check the **Issue "clear xlate" command when deploying access lists** check box to clear the NAT table when deploying new access lists. This setting ensures the access lists that are configured on the ASA are applied to all translated addresses.
- Access Rule Hit Count Settings let you configure the frequency for which the hit counts are updated in the Access Rules table. Hit counts are applicable for explicit rules only. No hit count will be displayed for implicit rules in the Access Rules table.

- Check the **Update access rule hit counts automatically** check box to have the hit counts automatically updated in the Access Rules table.
- Specify the frequency in seconds in which the hit count column is updated in the Access Rules table. Valid values are 10 - 86400 seconds.

Step 5 On the **Syslog** tab, specify the following:

- In the **Syslog Colors** area, you can customize the message display by configuring background or foreground colors for messages at each severity level. The **Severity** column lists each severity level by name and number. To change the background color or foreground color for messages at a specified severity level, click the corresponding column. The **Pick a Color** dialog box appears. Click one of the following tabs:
  - Choose a color from the palette on the **Swatches** tab and click **OK**.
  - Specify the H, S, and B settings on the **HSB** tab, and click **OK**.
  - Specify the Red, Green, and Blue settings on the **RGB** tab, and click **OK**.
- Check the **Warn to disable redundant syslog messages when NetFlow action is first applied to the global service policy rule** check box in the **NetFlow** area to enable the display of a warning message to disable redundant syslog messages.

Step 6 After you have specified settings on these three tabs, click **OK** to save your settings and close the **Preferences** dialog box.

Note
Each time that you check or uncheck a preferences setting, the change is saved to the .conf file and becomes available to all the other ASDM sessions running on the workstation at the time. You must restart ASDM for all changes to take effect.

Search with the ASDM Assistant

The ASDM Assistant tool lets you search and view useful ASDM procedural help about certain tasks.

Choose **View** > **ASDM Assistant** > **How Do I?** to access information, or enter a search request from the **Look For** field in the menu bar. Choose **How Do I?** From the **Find** drop-down list to begin the search.

To view the ASDM Assistant, perform the following steps:

Procedure

Step 1 Choose **View** > **ASDM Assistant**.

The **ASDM Assistant** pane appears.

Step 2 Enter the information that you want to find in the **Search** field, and click **Go**.

The requested information appears in the **Search Results** pane.

Step 3 Click any links that appear in the **Search Results and Features** sections to obtain more details.

Enable History Metrics

The History Metrics pane lets you configure the ASA to keep a history of various statistics, which ASDM can display on any graph/table. If you do not enable history metrics, you can only monitor statistics in real time. Enabling history metrics lets you view statistics graphs from the last 10 minutes, 60 minutes, 12 hours, and 5 days.

To configure history metrics, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Advanced** > **History Metrics**.

The **History Metrics** pane appears.

Step 2 Check the **ASDM History Metrics** check box to enable history metrics, then click **Apply**.

Unsupported Commands

ASDM supports almost all commands available for the ASA, but ASDM ignores some commands in an existing configuration. Most of these commands can remain in your configuration; see **Tools > Show Commands Ignored by ASDM on Device** for more information.

Ignored and View-Only Commands

The following table lists commands that ASDM supports in the configuration when added through the CLI, but that cannot be added or edited in ASDM. If ASDM ignores the command, it does not appear in the ASDM GUI at all. If the command is view-only, then it appears in the GUI, but you cannot edit it.

Table 6: List of Unsupported Commands

Unsupported Commands	ASDM Behavior
capture	Ignored.
coredump	Ignored. This can be configured only using the CLI.
crypto engine large-mod-accel	Ignored.
dhcp-server (tunnel-group name general-attributes)	ASDM only allows one setting for all DHCP servers.

Unsupported Commands	ASDM Behavior
eject	Unsupported.
established	Ignored.
failover timeout	Ignored.
fips	Ignored.
nat-assigned-to-public-ip	Ignored.
pager	Ignored.
pim accept-register route-map	Ignored. You can configure only the **list** option using ASDM.

service-policy global	Ignored if it uses a **match access-list** class. For example:
	
	`access-list myacl extended permit ip any any`
	`class-map mycm`
	`  match access-list myacl`
	`policy-map mypm`
	`  class mycm`
	`    inspect ftp`
	`service-policy mypm global`

set metric	Ignored.
sysopt nodnsalias	Ignored.
sysopt uauth allow-http-cache	Ignored.
terminal	Ignored.
threat-detection rate	Ignored.

Effects of Unsupported Commands

If ASDM loads an existing running configuration and finds other unsupported commands, ASDM operation is unaffected. Choose **Tools > Show Commands Ignored by ASDM on Device** to view the unsupported commands.

Discontinuous Subnet Masks Not Supported

ASDM does not support discontinuous subnet masks such as 255.255.0.255. For example, you cannot use the following:

```
ip address inside 192.168.2.1 255.255.0.255
```

Interactive User Commands Not Supported by the ASDM CLI Tool

The ASDM CLI tool does not support interactive user commands. If you enter a CLI command that requires interactive confirmation, ASDM prompts you to enter "[yes/no]" but does not recognize your input. ASDM then times out waiting for your response.

For example:

1. Choose **Tools** > **Command Line Interface**.
2. Enter the **crypto key generate rsa** command.

   ASDM generates the default 1024-bit RSA key.
3. Enter the **crypto key generate rsa** command again.

   Instead of regenerating the RSA keys by overwriting the previous one, ASDM displays the following error:

```
Do you really want to replace them? [yes/no]:WARNING: You already have
RSA ke0000000000000$A key
Input line must be less than 16 characters in length.

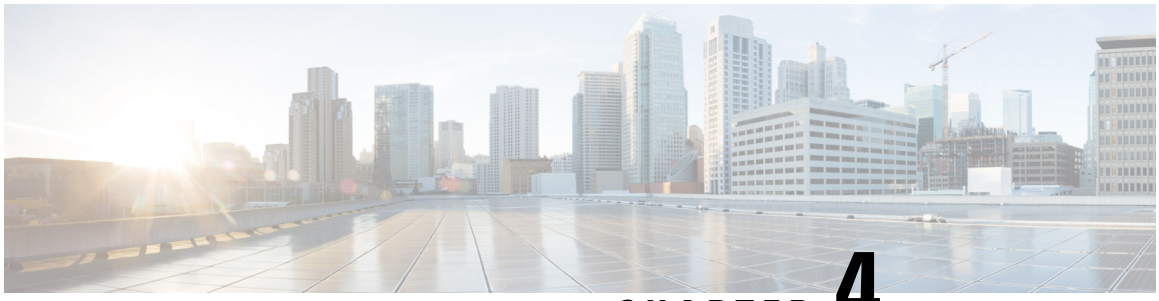
%Please answer 'yes' or 'no'.
Do you really want to replace them [yes/no]:

%ERROR: Timed out waiting for a response.
ERROR: Failed to create new RSA keys names <Default-RSA-key>
```

Workaround:

- You can configure most commands that require user interaction by means of the ASDM panes.
- For CLI commands that have a **noconfirm** option, use this option when entering the CLI command. For example:

```
crypto key generate rsa noconfirm
```



CHAPTER 4

Licenses: Product Authorization Key Licensing for the ISA 3000

A license specifies the options that are enabled on a given ASA. This document describes product authorization key (PAK) licenses for the ISA 3000. For other models, see Licenses: Smart Software Licensing, on page 113.

- About PAK Licenses, on page 85
- Guidelines for PAK Licenses, on page 93
- Configure PAK Licenses, on page 94
- Configure a Shared License (Secure Client 3 and Earlier), on page 99
- Supported Feature Licenses Per Model, on page 104
- Monitoring PAK Licenses, on page 105
- History for PAK Licenses, on page 106

About PAK Licenses

A license specifies the options that are enabled on a given ASA. It is represented by an activation key that is a 160-bit (5 32-bit words or 20 bytes) value. This value encodes the serial number (an 11 character string) and the enabled features.

Preinstalled License

By default, your ASA ships with a license already installed. This license might be the Base License, to which you want to add more licenses, or it might already have all of your licenses installed, depending on what you ordered and what your vendor installed for you.

Related Topics

Monitoring PAK Licenses, on page 105

Permanent License

You can have one permanent activation key installed. The permanent activation key includes all licensed features in a single key. If you also install time-based licenses, the ASA combines the permanent and time-based licenses into a running license.

Related Topics

How Permanent and Time-Based Licenses Combine, on page 86

Time-Based Licenses

In addition to permanent licenses, you can purchase time-based licenses or receive an evaluation license that has a time-limit. For example, you might buy a time-based Secure Client Premium license to handle short-term surges in the number of concurrent SSL VPN users.

Time-Based License Activation Guidelines

- You can install multiple time-based licenses, including multiple licenses for the same feature. However, only one time-based license per feature can be *active* at a time. The inactive license remains installed, and ready for use. For example, if you install a 1000-session Secure Client Premium license, and a 2500-session Secure Client Premium license, then only one of these licenses can be active.
- If you activate an evaluation license that has multiple features in the key, then you cannot also activate another time-based license for one of the included features.

How the Time-Based License Timer Works

- The timer for the time-based license starts counting down when you activate it on the ASA.
- If you stop using the time-based license before it times out, then the timer halts. The timer only starts again when you reactivate the time-based license.
- If the time-based license is active, and you shut down the ASA, then the timer stops counting down. The time-based license only counts down when the ASA is running. The system clock setting does not affect the license; only ASA uptime counts towards the license duration.

How Permanent and Time-Based Licenses Combine

When you activate a time-based license, then features from both permanent and time-based licenses combine to form the running license. How the permanent and time-based licenses combine depends on the type of license. The following table lists the combination rules for each feature license.

Note Even when the permanent license is used, if the time-based license is active, it continues to count down.

Table 7: Time-Based License Combination Rules

Time-Based Feature	Combined License Rule
Secure Client Premium Sessions	The higher value is used, either time-based or permanent. For example, if the permanent license is 1000 sessions, and the time-based license is 2500 sessions, then 2500 sessions are enabled. Typically, you will not install a time-based license that has less capability than the permanent license, but if you do so, then the permanent license is used.

Time-Based Feature	Combined License Rule
Unified Communications Proxy Sessions	The time-based license sessions are added to the permanent sessions, up to the platform limit. For example, if the permanent license is 2500 sessions, and the time-based license is 1000 sessions, then 3500 sessions are enabled for as long as the time-based license is active.
All Others	The higher value is used, either time-based or permanent. For licenses that have a status of enabled or disabled, then the license with the enabled status is used. For licenses with numerical tiers, the higher value is used. Typically, you will not install a time-based license that has less capability than the permanent license, but if you do so, then the permanent license is used.

Related Topics

Monitoring PAK Licenses, on page 105

Stacking Time-Based Licenses

In many cases, you might need to renew your time-based license and have a seamless transition from the old license to the new one. For features that are only available with a time-based license, it is especially important that the license not expire before you can apply the new license. The ASA allows you to *stack* time-based licenses so that you do not have to worry about the license expiring or about losing time on your licenses because you installed the new one early.

When you install an identical time-based license as one already installed, then the licenses are combined, and the duration equals the combined duration.

For example:

1. You install an 8-week 1000-session Secure Client Premium license, and use it for 2 weeks (6 weeks remain).
2. You then install another 8-week 1000-session license, and the licenses combine to be 1000-sessions for 14 weeks (8 weeks plus 6 weeks).

If the licenses are not identical (for example, a 1000-session Secure Client Premium license vs. a 2500-session license), then the licenses are *not* combined. Because only one time-based license per feature can be active, only one of the licenses can be active.

Although non-identical licenses do not combine, when the current license expires, the ASA automatically activates an installed license of the same feature if available.

Related Topics

Activate or Deactivate Keys, on page 98
Time-Based License Expiration, on page 87

Time-Based License Expiration

When the current license for a feature expires, the ASA automatically activates an installed license of the same feature if available. If there are no other time-based licenses available for the feature, then the permanent license is used.

If you have more than one additional time-based license installed for a feature, then the ASA uses the first license it finds; which license is used is not user-configurable and depends on internal operations. If you prefer to use a different time-based license than the one the ASA activated, then you must manually activate the license you prefer.

For example, you have a time-based 2500-session Secure Client Premium license (active), a time-based 1000-session Secure Client Premium license (inactive), and a permanent 500-session Secure Client Premium license. While the 2500-session license expires, the ASA activates the 1000-session license. After the 1000-session license expires, the ASA uses the 500-session permanent license.

Related Topics

Activate or Deactivate Keys, on page 98

License Notes

The following sections include additional information about licenses.

Secure Client Advantage, Secure Client Premier, and Secure Client VPN Only Licenses

The Secure Client Advantage or Premier license is a multi-use license that you can apply to multiple ASAs, all of which share a user pool as specified by the license. TheSecure Client VPN Only license applies to a specific ASA. See https://www.cisco.com/go/license, and assign the PAK separately to each ASA. When you apply the resulting activation key to an ASA, it toggles on the VPN features to the maximum allowed, but the actual number of unique users across all ASAs sharing the license should not exceed the license limit. For more information, see:

- Cisco Secure Client Ordering Guide
- Secure Client Licensing Frequently Asked Questions (FAQ)



Note The Secure Client Premier license is the only Secure Client Premier license supported for multiple context mode. Moreover, in multiple context mode, this license must be applied to each unit in a failover pair; the license is not aggregated.

Other VPN License

Other VPN peers include the following VPN types:

- IPsec remote access VPN using IKEv1
- IPsec site-to-site VPN using IKEv1
- IPsec site-to-site VPN using IKEv2

This license is included in the Base license.

Total VPN Sessions Combined, All Types

- The Total VPN Peers is the maximum VPN peers allowed of both Secure Client and Other VPN peers combined. For example, if the total is 1000, you can allow 500 Secure Client and 500 Other VPN peers

simultaneously; or 700 Secure Client and 300 Other VPN; or use all 1000 for Secure Client. If you exceed the total VPN peers, you can overload the ASA, so be sure to size your network appropriately.

VPN Load Balancing

VPN load balancing requires a Strong Encryption (3DES/AES) License.

Legacy VPN Licenses

Refer to the Supplemental end User License Agreement for Secure Client for all relevant information on licensing.



Note The Secure Client Premier license is the only Secure Client license supported for multiple context mode; you cannot use the default or legacy license.

Encryption License

If you have Strong Encryption enabled, you cannot use DES.

Total TLS Proxy Sessions

Each TLS proxy session for Encrypted Voice Inspection is counted against the TLS license limit.

Other applications that use TLS proxy sessions do not count toward the TLS limit, for example, Mobility Advantage Proxy (which does not require a license).

Some applications might use multiple sessions for a connection. For example, if you configure a phone with a primary and backup Cisco Unified Communications Manager, there are 2 TLS proxy connections.

You independently set the TLS proxy limit using the **tls-proxy maximum-sessions** command or in ASDM, using the **Configuration > Firewall > Unified Communications > TLS Proxy** pane. To view the limits of your model, enter the **tls-proxy maximum-sessions ?** command. When you apply a TLS proxy license that is higher than the default TLS proxy limit, the ASA automatically sets the TLS proxy limit to match the license. The TLS proxy limit takes precedence over the license limit; if you set the TLS proxy limit to be less than the license, then you cannot use all of the sessions in your license.



Note For license part numbers ending in “K8” (for example, licenses under 250 users), TLS proxy sessions are limited to 1000. For license part numbers ending in “K9” (for example, licenses 250 users or larger), the TLS proxy limit depends on the configuration, up to the model limit. K8 and K9 refer to whether the license is restricted for export: K8 is unrestricted, and K9 is restricted.

If you clear the configuration (using the **clear configure all** command, for example), then the TLS proxy limit is set to the default for your model; if this default is lower than the license limit, then you see an error message to use the **tls-proxy maximum-sessions** command to raise the limit again (in ASDM, use the **TLS Proxy** pane). If you use failover and enter the **write standby** command or in ASDM, use **File > Save Running Configuration to Standby Unit** on the primary unit to force a configuration synchronization, the **clear configure all** command is generated on the secondary unit automatically, so you may see the warning message on the secondary unit. Because the configuration synchronization restores the TLS proxy limit set on the primary unit, you can ignore the warning.

You might also use SRTP encryption sessions for your connections:

- For K8 licenses, SRTP sessions are limited to 250.
- For K9 licenses, there is no limit.

Note Only calls that require encryption/decryption for media are counted toward the SRTP limit; if passthrough is set for the call, even if both legs are SRTP, they do not count toward the limit.

VLANs, Maximum

For an interface to count against the VLAN limit, you must assign a VLAN to it.

Shared Secure Client Premium Licenses (AnyConnect 3 and Earlier)

Note The shared license feature on the ASA is not supported with AnyConnect 4 and later licensing. Secure Client licenses are shared and no longer require a shared server or participant license.

A shared license lets you purchase a large number of Secure Client Premium sessions and share the sessions as needed among a group of ASAs by configuring one of the ASAs as a shared licensing server, and the rest as shared licensing participants.

Failover

With some exceptions, failover units do not require the same license on each unit. For earlier versions, see the licensing document for your version.

Failover License Requirements and Exceptions

Model	License Requirement
ISA 3000	Security Plus license on both units. **Note** Each unit must have the same encryption license.

Note A valid permanent key is required; in rare instances on the ISA 3000, your PAK authentication key can be removed. If your key consists of all 0's, then you need to reinstall a valid authentication key before failover can be enabled.

How Failover Licenses Combine

For failover pairs, the licenses on each unit are combined into a single running cluster license. If you buy separate licenses for each unit, then the combined license uses the following rules:

- For licenses that have numerical tiers, such as the number of sessions, the values from each unit's licenses are combined up to the platform limit. If all licenses in use are time-based, then the licenses count down simultaneously.

  For example, for failover:

  - You have two ASAs with 10 TLS Proxy sessions installed on each; the licenses will be combined for a total of 20 TLS Proxy sessions.
  - You have an ASA with 1000 TLS Proxy sessions, and another with 2000 sessions; because the platform limit is 2000, the combined license allows 2000 TLS Proxy sessions.

- For licenses that have a status of enabled or disabled, then the license with the enabled status is used.
- For time-based licenses that are enabled or disabled (and do not have numerical tiers), the duration is the combined duration of all licenses. The primary/control unit counts down its license first, and when it expires, the secondary/data unit(s) start counting down its license, and so on.

Related Topics

Monitoring PAK Licenses, on page 105

Loss of Communication Between Failover Units

If the units lose communication for more than 30 days, then each unit reverts to the license installed locally. During the 30-day grace period, the combined running license continues to be used by all units.

If you restore communication during the 30-day grace period, then for time-based licenses, the time elapsed is subtracted from the primary/control license; if the primary/control license becomes expired, only then does the secondary/data license start to count down.

If you do not restore communication during the 30-day period, then for time-based licenses, time is subtracted from all unit licenses, if installed. They are treated as separate licenses and do not benefit from the combined license. The time elapsed includes the 30-day grace period.

Upgrading Failover Pairs

Because failover pairs do not require the same license on both units, you can apply new licenses to each unit without any downtime. If you apply a permanent license that requires a reload, then you can fail over to the other unit while you reload. If both units require reloading, then you can reload them separately so that you have no downtime.

Related Topics

Activate or Deactivate Keys, on page 98

No Payload Encryption Models

You can purchase some models with No Payload Encryption. For export to some countries, payload encryption cannot be enabled on the ASA series. The ASA software senses a No Payload Encryption model, and disables the following features:

- Unified Communications
- VPN

You can still install the Strong Encryption (3DES/AES) license for use with management connections. For example, you can use ASDM HTTPS/SSL, SSHv2, Telnet and SNMPv3.

When you view the license, VPN and Unified Communications licenses will not be listed.

Related Topics

Monitoring PAK Licenses, on page 105

Licenses FAQ

Can I activate multiple time-based licenses?

Yes. You can use one time-based license per feature at a time.

Can I "stack" time-based licenses so that when the time limit runs out, it will automatically use the next license?

Yes. For identical licenses, the time limit is combined when you install multiple time-based licenses. For non-identical licenses (for example, a 1000-session Secure Client Premium license and a 2500-session license), the ASA automatically activates the next time-based license it finds for the feature.

Can I install a new permanent license while maintaining an active time-based license?

Yes. Activating a permanent license does not affect time-based licenses.

For failover, can I use a shared licensing server as the primary unit, and the shared licensing backup server as the secondary unit?

No. The secondary unit has the same running license as the primary unit; in the case of the shared licensing server, they require a server license. The backup server requires a participant license. The backup server can be in a separate failover pair of two backup servers.

Do I need to buy the same licenses for the secondary unit in a failover pair?

No. Starting with Version 8.3(1), you do not have to have matching licenses on both units. Typically, you buy a license only for the primary unit; the secondary unit inherits the primary license when it becomes active. In the case where you also have a separate license on the secondary unit (for example, if you purchased matching licenses for pre-8.3 software), the licenses are combined into a running failover cluster license, up to the model limits.

Can I use a time-based or permanent Secure Client Premium license in addition to a shared AnyConnect Premium license?

Yes. The shared license is used only after the sessions from the locally installed license (time-based or permanent) are used up.



Note On the shared licensing server, the permanent Secure Client Premium license is not used; you can however use a time-based license at the same time as the shared licensing server license. In this case, the time-based license sessions are available for local Secure Client Premium sessions only; they cannot be added to the shared licensing pool for use by participants.

Guidelines for PAK Licenses

Context Mode Guidelines

In multiple context mode, apply the activation key in the system execution space.

Failover Guidelines

See Failover, on page 90.

Model Guidelines

- Smart Licensing is supported on the ASA virtual only.
- Shared licenses are not supported on the ASA virtual, ASA 5506-X, ASA 5508-X, and ASA 5516-X.
- The ASA 5506-X and ASA 5506W-X do not support time-based licenses.

Upgrade and Downgrade Guidelines

Your activation key remains compatible if you upgrade to the latest version from any previous version. However, you might have issues if you want to maintain downgrade capability:

- Downgrading to Version 8.1 or earlier—After you upgrade, if you activate additional feature licenses that were introduced *before 8.2*, then the activation key continues to be compatible with earlier versions if you downgrade. However if you activate feature licenses that were introduced in *8.2 or later*, then the activation key is not backwards compatible. If you have an incompatible license key, then see the following guidelines:
  - If you previously entered an activation key in an earlier version, then the ASA uses that key (without any of the new licenses you activated in Version 8.2 or later).
  - If you have a new system and do not have an earlier activation key, then you need to request a new activation key compatible with the earlier version.
- Downgrading to Version 8.2 or earlier—Version 8.3 introduced more robust time-based key usage as well as failover license changes:
  - If you have more than one time-based activation key active, when you downgrade, only the most recently activated time-based key can be active. Any other keys are made inactive. If the last time-based license is for a feature introduced in 8.3, then that license still remains the active license even though it cannot be used in earlier versions. Reenter the permanent key or a valid time-based key.
  - If you have mismatched licenses on a failover pair, then downgrading will disable failover. Even if the keys are matching, the license used will no longer be a combined license.
  - If you have one time-based license installed, but it is for a feature introduced in 8.3, then after you downgrade, that time-based license remains active. You need to reenter the permanent key to disable the time-based license.

Additional Guidelines

- The activation key is not stored in your configuration file; it is stored as a hidden file in flash memory.
- The activation key is tied to the serial number of the device. Feature licenses cannot be transferred between devices (except in the case of a hardware failure). If you have to replace your device due to a hardware failure, and it is covered by Cisco TAC, contact the Cisco Licensing Team to have your existing license transferred to the new serial number. The Cisco Licensing Team will ask for the Product Authorization Key reference number and existing serial number.
- The serial number used for licensing is the one seen on the Activation Key page. This serial number is different from the chassis serial number printed on the outside of your hardware. The chassis serial number is used for technical support, but not for licensing.
- Once purchased, you cannot return a license for a refund or for an upgraded license.
- On a single unit, you cannot add two separate licenses for the same feature together; for example, if you purchase a 25-session SSL VPN license, and later purchase a 50-session license, you cannot use 75 sessions; you can use a maximum of 50 sessions. (You may be able to purchase a larger license at an upgrade price, for example from 25 sessions to 75 sessions; this kind of upgrade should be distinguished from adding two separate licenses together).
- Although you can activate all license types, some features are incompatible with each other. In the case of the AnyConnect Essentials license, the license is incompatible with the following licenses: AnyConnect Premium license, shared AnyConnect Premium license, and Advanced Endpoint Assessment license. By default, if you install the AnyConnect Essentials license (if it is available for your model), it is used instead of the above licenses. You can disable the AnyConnect Essentials license in the configuration to restore use of the other licenses using the **Configuration > Remote Access VPN > Network (Client) Access > Advanced > AnyConnect Essentials** pane.

Configure PAK Licenses

This section describes how to obtain an activation key and how to active it. You can also deactivate a key.

Order License PAKs and Obtain an Activation Key

To install a license on the ASA, you need Product Authorization Keys, which you can then register with Cisco.com to obtain an activation key. You can then enter the activation key on the ASA. You need a separate Product Authorization Key for each feature license. The PAKs are combined to give you a single activation key. You may have received all of your license PAKs in the box with your device. The ASA has the Base or Security Plus license pre-installed, along with the Strong Encryption (3DES/AES) license if you qualify for its use. If you need to manually request the Strong Encryption license (which is free), see http://www.cisco.com/go/license.

Before you begin

When you purchase 1 or more licenses for the device, you manage them in the Cisco Smart Software Manager:

https://software.cisco.com/#module/SmartLicensing

If you do not yet have an account, set up a new account. The Smart Software Manager lets you create a master account for your organization.

Procedure

Step 1 To purchase additional licenses, see http://www.cisco.com/go/ccw. See the following Secure Client ordering guide and FAQ:

- Cisco Secure Client Ordering Guide
- Secure Client Licensing Frequently Asked Questions (FAQ)

After you order a license, you will then receive an email with a Product Authorization Key (PAK). For the Secure Client licenses, you receive a multi-use PAK that you can apply to multiple ASAs that use the same pool of user sessions. The PAK email can take several days in some cases.

Step 2 Obtain the serial number for your ASA by choosing **Configuration** > **Device Management** > **Licensing** > **Activation Key** (in multiple context mode, view the serial number in the System execution space).

The serial number used for licensing is different from the chassis serial number printed on the outside of your hardware. The chassis serial number is used for technical support, but not for licensing.

Step 3 To obtain the activation key, go to the following licensing website:

http://www.cisco.com/go/license

Step 4 Enter the following information, when prompted:

- Product Authorization Key (if you have multiple keys, enter one of the keys first. You have to enter each key as a separate process.)
- The serial number of your ASA
- Your e-mail address

An activation key is automatically generated and sent to the e-mail address that you provide. This key includes all features you have registered so far for permanent licenses. For time-based licenses, each license has a separate activation key.

Step 5 If you have additional Product Authorization Keys, repeat the process for each Product Authorization Key. After you enter all of the Product Authorization Keys, the final activation key provided includes all of the permanent features you registered.

Step 6 Install the activation key according to Activate or Deactivate Keys, on page 98.

Obtain a Strong Encryption License

To use ASDM (and many other features), you need to install the Strong Encryption (3DES/AES) license. If your ASA did not come with the Strong Encryption license pre-installed, you can request one for free. You must qualify for a Strong Encryption license based on your country.

Procedure

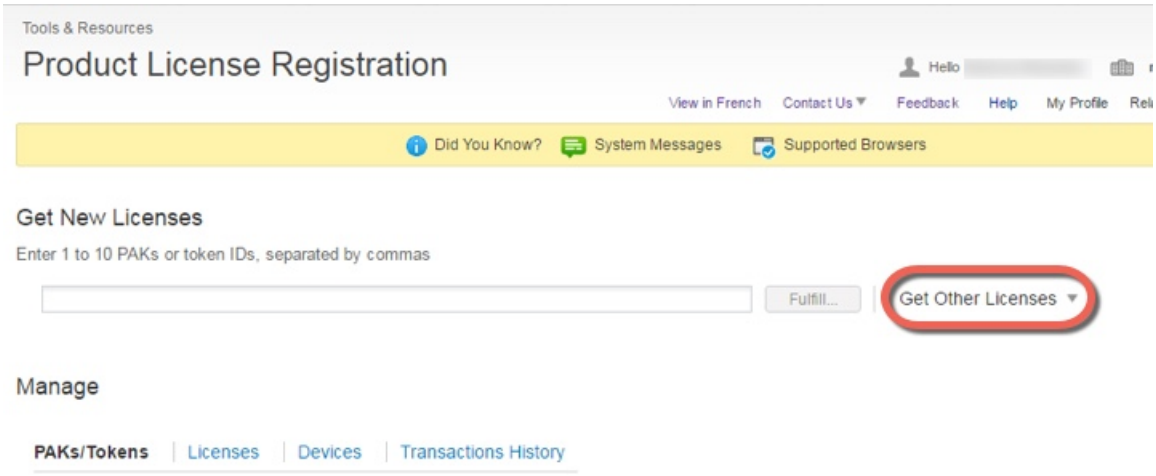
Step 1 Obtain the serial number for your ASA by entering the following command:

show version | grep Serial

This serial number is different from the chassis serial number printed on the outside of your hardware. The chassis serial number is used for technical support, but not for licensing.

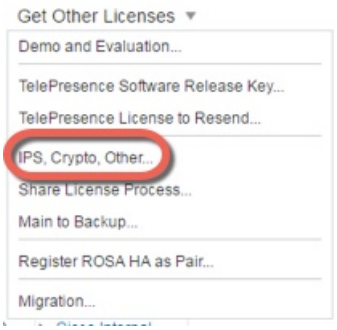
Step 2 See https://www.cisco.com/go/license, and click **Get Other Licenses**.

Figure 12: Get Other Licenses



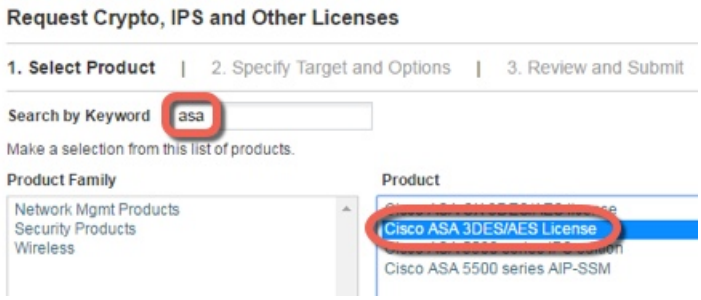
Step 3 Choose **IPS, Crypto, Other**.

Figure 13: IPS, Crypto, Other



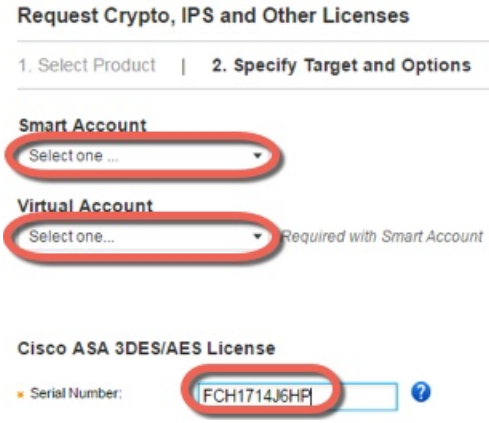
Step 4 In the **Search by Keyword** field, enter **asa**, and select **Cisco ASA 3DES/AES License**.

Figure 14: Cisco ASA 3DES/AES License



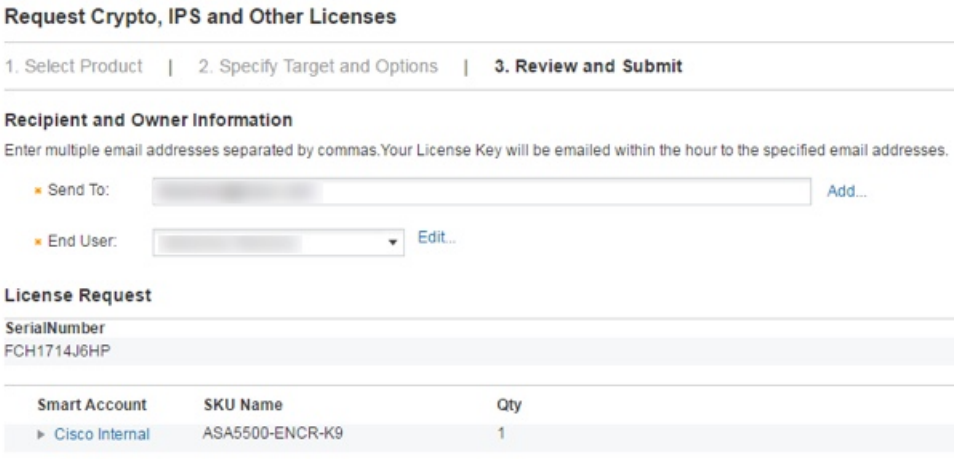
Step 5 Select your **Smart Account**, **Virtual Account**, enter the ASA **Serial Number**, and click **Next**.

Figure 15: Smart Account, Virtual Account, and Serial Number



Step 6 Your Send To email address and End User name are auto-filled; enter additional email addresses if needed. Check the **I Agree** check box, and click **Submit**.

Figure 16: Submit



Step 7 You will then receive an email with the activation key, but you can also download the key right away from the **Manage** > **Licenses** area.

Step 8 Apply the activation key according to Activate or Deactivate Keys, on page 98.

Activate or Deactivate Keys

This section describes how to enter a new activation key, and how to activate and deactivate time-based keys.

Before you begin

- If you are already in multiple context mode, enter the activation key in the system execution space.
- Some permanent licenses require you to reload the ASA after you activate them. The following table lists the licenses that require reloading.

Table 8: Permanent License Reloading Requirements

Model	License Action Requiring Reload
All models	Downgrading the Encryption license.

Procedure

Step 1 Choose **Configuration** > **Device Management**, and then choose the **Licensing** > **Activation Key** or **Licensing Activation Key** pane, depending on your model.

Step 2 To enter a new activation key, either permanent or time-based, enter the new activation key in the **New Activation Key** field.

The key is a five-element hexadecimal string with one space between each element. The leading 0x specifier is optional; all values are assumed to be hexadecimal. For example:

```
ASA0xd11b3d48 0xa80a4c0a 0x48e0fd1c 0xb0443480 0x843fc490
```

You can install one permanent key, and multiple time-based keys. If you enter a new permanent key, it overwrites the already installed one. If you enter a new time-based key, then it is active by default and displays in the Time-based License Keys Installed table. The last time-based key that you activate for a given feature is the active one.

Step 3 To activate or deactivate an installed time-based key, choose the key in the **Time-based License Keys Installed** table, and click either **Activate** or **Deactivate**.

You can only have one time-based key active for each feature.

Step 4 Click **Update Activation Key**.

Some permanent licenses require you to reload the ASA after entering the new activation key. You will be prompted to reload if it is required.

Related Topics

Time-Based Licenses, on page 86

Configure a Shared License (Secure Client 3 and Earlier)

Note The shared license feature on the ASA is not supported with Secure Client 4 and later licensing. Secure Client licenses are shared and no longer require a shared server or participant license.

This section describes how to configure the shared licensing server and participants.

About Shared Licenses

A shared license lets you purchase a large number of Secure Client Premium sessions and share the sessions as needed among a group of ASAs by configuring one of the ASAs as a shared licensing server, and the rest as shared licensing participants.

About the Shared Licensing Server and Participants

The following steps describe how shared licenses operate:

1. Decide which ASA should be the shared licensing server, and purchase the shared licensing server license using that device serial number.
2. Decide which ASAs should be shared licensing participants, including the shared licensing backup server, and obtain a shared licensing participant license for each device, using each device serial number.
3. (Optional) Designate a second ASA as a shared licensing backup server. You can only specify one backup server.

Note The shared licensing backup server only needs a participant license.

4. Configure a shared secret on the shared licensing server; any participants with the shared secret can use the shared license.
5. When you configure the ASA as a participant, it registers with the shared licensing server by sending information about itself, including the local license and model information.

Note The participant needs to be able to communicate with the server over the IP network; it does not have to be on the same subnet.

6. The shared licensing server responds with information about how often the participant should poll the server.
7. When a participant uses up the sessions of the local license, it sends a request to the shared licensing server for additional sessions in 50-session increments.
8. The shared licensing server responds with a shared license. The total sessions used by a participant cannot exceed the maximum sessions for the platform model.

Note The shared licensing server can also participate in the shared license pool. It does not need a participant license as well as the server license to participate.

a. If there are not enough sessions left in the shared license pool for the participant, then the server responds with as many sessions as available.

b. The participant continues to send refresh messages requesting more sessions until the server can adequately fulfill the request.

9. When the load is reduced on a participant, it sends a message to the server to release the shared sessions.

Note The ASA uses SSL between the server and participant to encrypt all communications.

Communication Issues Between Participant and Server

See the following guidelines for communication issues between the participant and server:

- If a participant fails to send a refresh after 3 times the refresh interval, then the server releases the sessions back into the shared license pool.
- If the participant cannot reach the license server to send the refresh, then the participant can continue to use the shared license it received from the server for up to 24 hours.
- If the participant is still not able to communicate with a license server after 24 hours, then the participant releases the shared license, even if it still needs the sessions. The participant leaves existing connections established, but cannot accept new connections beyond the license limit.
- If a participant reconnects with the server before 24 hours expires, but after the server expired the participant sessions, then the participant needs to send a new request for the sessions; the server responds with as many sessions as can be reassigned to that participant.

About the Shared Licensing Backup Server

The shared licensing backup server must register successfully with the main shared licensing server before it can take on the backup role. When it registers, the main shared licensing server syncs server settings as well as the shared license information with the backup, including a list of registered participants and the current license usage. The main server and backup server sync the data at 10 second intervals. After the initial sync, the backup server can successfully perform backup duties, even after a reload.

When the main server goes down, the backup server takes over server operation. The backup server can operate for up to 30 continuous days, after which the backup server stops issuing sessions to participants, and existing sessions time out. Be sure to reinstate the main server within that 30-day period. Critical-level syslog messages are sent at 15 days, and again at 30 days.

When the main server comes back up, it syncs with the backup server, and then takes over server operation.

When the backup server is not active, it acts as a regular participant of the main shared licensing server.

Note When you first launch the main shared licensing server, the backup server can only operate independently for 5 days. The operational limit increases day-by-day, until 30 days is reached. Also, if the main server later goes down for any length of time, the backup server operational limit decrements day-by-day. When the main server comes back up, the backup server starts to increment again day-by-day. For example, if the main server is down for 20 days, with the backup server active during that time, then the backup server will only have a 10-day limit left over. The backup server "recharges" up to the maximum 30 days after 20 more days as an inactive backup. This recharging function is implemented to discourage misuse of the shared license.

Failover and Shared Licenses

This section describes how shared licenses interact with failover.

Failover and Shared License Servers

This section describes how the main server and backup server interact with failover. Because the shared licensing server is also performing normal duties as the ASA, including performing functions such as being a VPN gateway and firewall, then you might need to configure failover for the main and backup shared licensing servers for increased reliability.

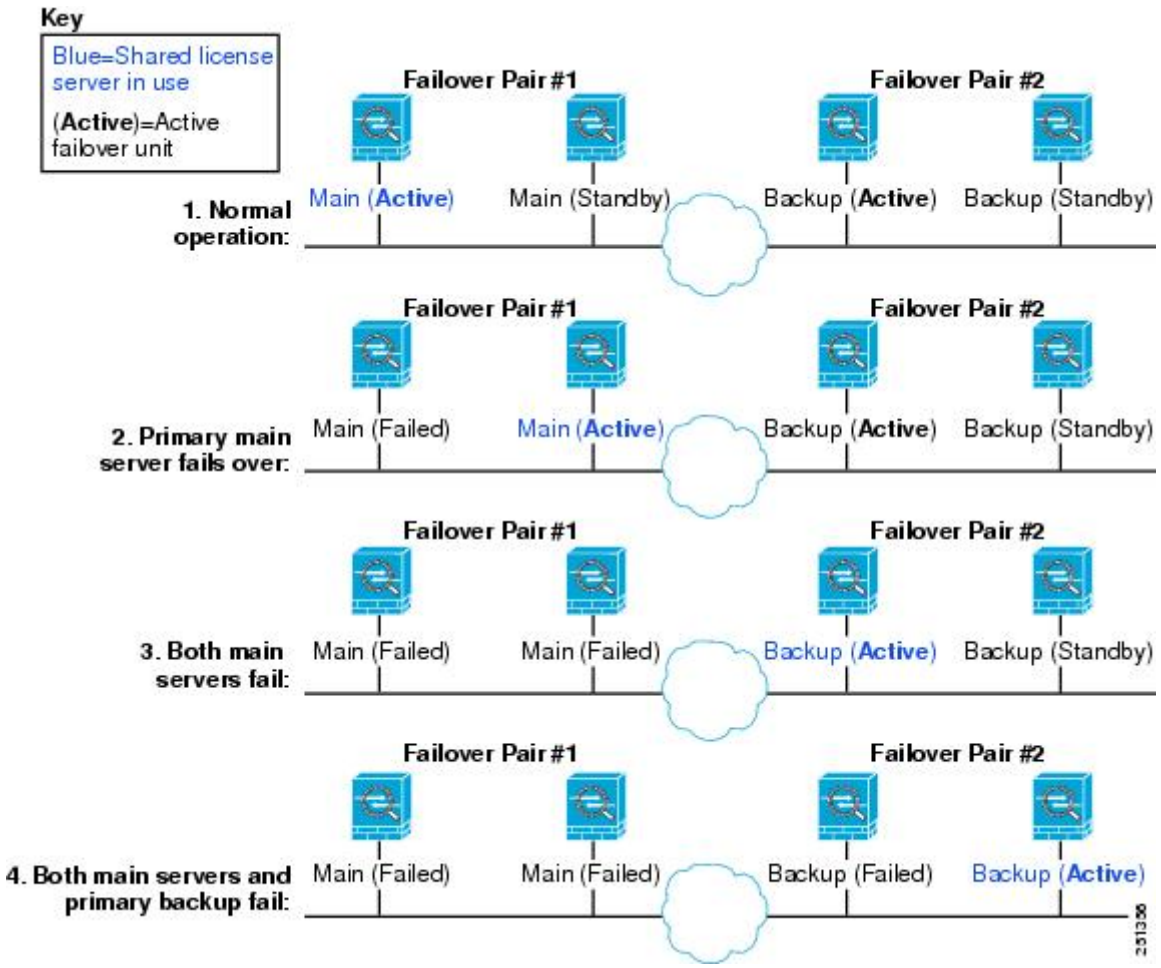
Note The backup server mechanism is separate from, but compatible with, failover.

Shared licenses are supported only in single context mode, so Active/Active failover is not supported.

For Active/Standby failover, the primary unit acts as the main shared licensing server, and the standby unit acts as the main shared licensing server after failover. The standby unit does *not* act as the backup shared licensing server. Instead, you can have a second pair of units acting as the backup server, if desired.

For example, you have a network with 2 failover pairs. Pair #1 includes the main licensing server. Pair #2 includes the backup server. When the primary unit from Pair #1 goes down, the standby unit immediately becomes the new main licensing server. The backup server from Pair #2 never gets used. Only if both units in Pair #1 go down does the backup server in Pair #2 come into use as the shared licensing server. If Pair #1 remains down, and the primary unit in Pair #2 goes down, then the standby unit in Pair #2 comes into use as the shared licensing server (see the following figure).

Figure 17: Failover and Shared License Servers



The standby backup server shares the same operating limits as the primary backup server; if the standby unit becomes active, it continues counting down where the primary unit left off.

Related Topics

About the Shared Licensing Backup Server, on page 100

Failover and Shared License Participants

For participant pairs, both units register with the shared licensing server using separate participant IDs. The active unit syncs its participant ID with the standby unit. The standby unit uses this ID to generate a transfer request when it switches to the active role. This transfer request is used to move the shared sessions from the previously active unit to the new active unit.

Maximum Number of Participants

The ASA does not limit the number of participants for the shared license; however, a very large shared network could potentially affect the performance on the licensing server. In this case, you can increase the delay between participant refreshes, or you can create two shared networks.

Configure the Shared Licensing Server

This section describes how to configure the ASA to be a shared licensing server.

Before you begin

The server must have a shared licensing server key.

Procedure

Step 1 Choose the **Configuration** > **Device Management** > **Licenses** > **Shared SSL VPN Licenses** pane.

Step 2 In the **Shared Secret** field, enter the shared secret as a string between 4 and 128 ASCII characters.

Any participant with this secret can use the license server.

Step 3 (Optional) In the **TCP IP Port** field, enter the port on which the server listens for SSL connections from participants, between 1 and 65535.

The default is TCP port 50554.

Step 4 (Optional) In the **Refresh interval** field, enter the refresh interval between 10 and 300 seconds.

This value is provided to participants to set how often they should communicate with the server. The default is 30 seconds.

Step 5 In the **Interfaces that serve shared licenses** area, check the **Shares Licenses** check box for any interfaces on which participants contact the server.

Step 6 (Optional) To identify a backup server, in the **Optional backup shared SSL VPN license server** area:

a) In the **Backup server IP address** field, enter the backup server IP address.
b) In the **Primary backup server serial number** field, enter the backup server serial number.
c) If the backup server is part of a failover pair, identify the standby unit serial number in the **Secondary backup server serial number** field.

You can only identify 1 backup server and its optional standby unit.

Step 7 Click **Apply**.

Configure the Shared Licensing Participant and the Optional Backup Server

This section configures a shared licensing participant to communicate with the shared licensing server. This section also describes how you can optionally configure the participant as the backup server.

Before you begin

The participant must have a shared licensing participant key.

Procedure

Step 1 Choose the **Configuration** > **Device Management** > **Licenses** > **Shared SSL VPN Licenses** pane.

Step 2 In the Shared Secret field, enter the shared secret as a string between 4 and 128 ASCII characters.

Step 3 (Optional) In the TCP IP Port field, enter the port on which to communicate with the server using SSL, between 1 and 65535.

The default is TCP port 50554.

Step 4 (Optional) To identify the participant as the backup server, in the Select backup role of participant area:

a) Click the **Backup Server** radio button.

b) Check the **Shares Licenses** check box for any interfaces on which participants contact the backup server.

Step 5 Click **Apply**.

Supported Feature Licenses Per Model

This section describes the licenses available for each model as well as important notes about licenses.

Licenses Per Model

This section lists the feature licenses available for each model:

Items that are in *italics* are separate, optional licenses that can replace the Base (or Security Plus, and so on) license version. You can mix and match optional licenses.

Note Some features are incompatible with each other. See the individual feature chapters for compatibility information.

If you have a No Payload Encryption model, then some of the features below are not supported. See No Payload Encryption Models, on page 91 for a list of unsupported features.

For detailed information about licenses, see License Notes, on page 88.

ISA 3000 License Features

The following table shows the licensed features for the ISA 3000.

Licenses	Base License	Security Plus License
Firewall Licenses		
Botnet Traffic Filter	No support	No Support
Firewall Conns, Concurrent	20,000	50,000

Licenses	Base License		Security Plus License	
Carrier	No Support		No Support	
Total TLS Proxy Sessions	160		160	
VPN Licenses				
Secure Client peers	Disabled	*Optional Secure Client Advantage, Secure Client Premier, or Secure Client VPN Only license: 25 maximum*	Disabled	*Optional Secure Client Advantage, Secure Client Premier, or Secure Client VPN Only license: 25 maximum*
Other VPN Peers	10		50	
Total VPN Peers, combined all types	25		50	
VPN Load Balancing	No support		No support	
General Licenses				
Encryption	Base (DES)	*Opt. lic.: Strong (3DES/AES)*	Base (DES)	*Opt. lic.: Strong (3DES/AES)*
Failover	No support		Active/Standby	
Security Contexts	No support		No Support	
Clustering	No Support		No Support	
VLANs, Maximum	5		25	

Monitoring PAK Licenses

This section describes how to view license information.

Viewing Your Current License

This section describes how to view your current license, and for time-based activation keys, how much time the license has left.

Before you begin

If you have a No Payload Encryption model, then you view the license, VPN and Unified Communications licenses will not be listed. See No Payload Encryption Models, on page 91 for more information.

Procedure

Step 1 To view the running license, which is a combination of the permanent license and any active time-based licenses, choose the **Configuration** > **Device Management** > **Licensing** > **Activation Key** pane and view the Running Licenses area.

In multiple context mode, view the activation key in the System execution space by choosing the **Configuration > Device Management > Activation Key** pane.

For a failover pair, the running license shown is the combined license from the primary and secondary units. See How Failover Licenses Combine, on page 91 for more information. For time-based licenses with numerical values (the duration is not combined), the License Duration column displays the shortest time-based license from either the primary or secondary unit; when that license expires, the license duration from the other unit displays.

Step 2 (Optional) To view time-based license details, such as the features included in the license and the duration, in the Time-Based License Keys Installed area, choose a license key, and then click **Show License Details**.

Step 3 (Optional) For a failover unit, to view the license installed on this unit (and not the combined license from both primary and secondary units), in the Running Licenses area, click **Show information of license specifically purchased for this device alone**.

Monitoring the Shared License

To monitor the shared license, choose **Monitoring > VPN > Clientless SSL VPN > Shared Licenses**.

History for PAK Licenses

Feature Name	Platform Releases	Description
Increased Connections and VLANs	7.0(5)	Increased the following limits: • ASA5510 Base license connections from 32000 to 5000; VLANs from 0 to 10. • ASA5510 Security Plus license connections from 64000 to 130000; VLANs from 10 to 25. • ASA5520 connections from 130000 to 280000; VLANs from 25 to 100. • ASA5540 connections from 280000 to 400000; VLANs from 100 to 200.
SSL VPN Licenses	7.1(1)	SSL VPN licenses were introduced.
Increased SSL VPN Licenses	7.2(1)	A 5000-user SSL VPN license was introduced for the ASA 5550 and above.

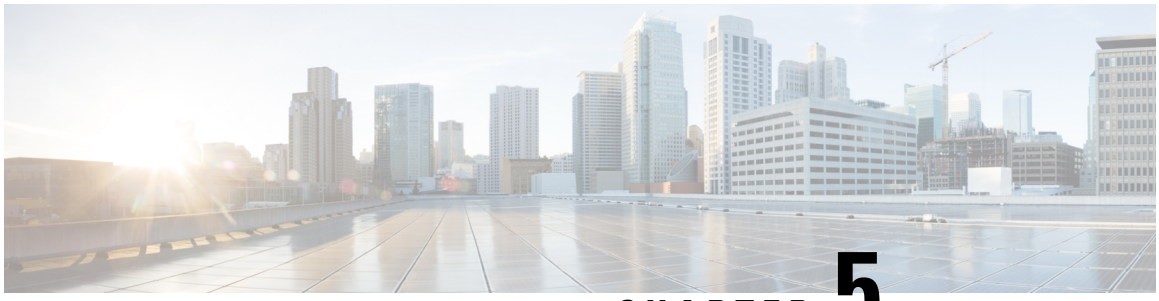
Feature Name	Platform Releases	Description
Increased interfaces for the Base license on the ASA 5510	7.2(2)	For the Base license on the ASA 5510, the maximum number of interfaces was increased from 3 plus a management interface to unlimited interfaces.
Increased VLANs	7.2(2)	The maximum number of VLANs for the Security Plus license on the ASA 5505 was increased from 5 (3 fully functional; 1 failover; one restricted to a backup interface) to 20 fully functional interfaces. In addition, the number of trunk ports was increased from 1 to 8. Now there are 20 fully functional interfaces, you do not need to use the backup interface command to cripple a backup ISP interface; you can use a fully functional interface for it. The backup interface command is still useful for an Easy VPN configuration. VLAN limits were also increased for the ASA 5510 (from 10 to 50 for the Base license, and from 25 to 100 for the Security Plus license), the ASA 5520 (from 100 to 150), the ASA 5550 (from 200 to 250).
Gigabit Ethernet Support for the ASA 5510 Security Plus License	7.2(3)	The ASA 5510 now supports Gigabit Ethernet (1000 Mbps) for the Ethernet 0/0 and 0/1 ports with the Security Plus license. In the Base license, they continue to be used as Fast Ethernet (100 Mbps) ports. Ethernet 0/2, 0/3, and 0/4 remain as Fast Ethernet ports for both licenses. **Note** The interface names remain Ethernet 0/0 and Ethernet 0/1.
Advanced Endpoint Assessment License	8.0(2)	The Advanced Endpoint Assessment license was introduced. As a condition for the completion of a Cisco AnyConnect or clientless SSL VPN connections, the remote computer scans for a greatly expanded collection of antivirus and antispyware applications, firewalls, operating systems, and associated updates. It also scans for any registry entries, filenames, and process names that you specify. It sends the scan results to the ASA. The ASA uses both the user login credentials and the computer scan results to assign a Dynamic Access Policy (DAP). With an Advanced Endpoint Assessment License, you can enhance Host Scan by configuring an attempt to update noncompliant computers to meet version requirements. Cisco can provide timely updates to the list of applications and versions that Host Scan supports in a package that is separate from Cisco Secure Desktop.
VPN Load Balancing for the ASA 5510	8.0(2)	VPN load balancing is now supported on the ASA 5510 Security Plus license.

Feature Name	Platform Releases	Description
AnyConnect for Mobile License	8.0(3)	The AnyConnect for Mobile license was introduced. It lets Windows mobile devices connect to the ASA using the Secure Client.
Time-based Licenses	8.0(4)/8.1(2)	Support for time-based licenses was introduced.
Increased VLANs for the ASA 5580	8.1(2)	The number of VLANs supported on the ASA 5580 are increased from 100 to 250.
Unified Communications Proxy Sessions license	8.0(4)	The UC Proxy sessions license was introduced. Phone Proxy, Presence Federation Proxy, and Encrypted Voice Inspection applications use TLS proxy sessions for their connections. Each TLS proxy session is counted against the UC license limit. All of these applications are licensed under the UC Proxy umbrella, and can be mixed and matched. This feature is not available in Version 8.1.
Botnet Traffic Filter License	8.2(1)	The Botnet Traffic Filter license was introduced. The Botnet Traffic Filter protects against malware network activity by tracking connections to known bad domains and IP addresses.
AnyConnect Essentials License	8.2(1)	The AnyConnect Essentials License was introduced. This license enables AnyConnect VPN client access to the ASA. This license does not support browser-based SSL VPN access or Cisco Secure Desktop. For these features, activate an AnyConnect Premium license instead of the AnyConnect Essentials license. **Note** With the AnyConnect Essentials license, VPN users can use a Web browser to log in, and download and start (WebLaunch) the Secure Client. The Secure Client software offers the same set of client features, whether it is enabled by this license or an AnyConnect Premium license. The AnyConnect Essentials license cannot be active at the same time as the following licenses on a given ASA: AnyConnect Premium license (all types) or the Advanced Endpoint Assessment license. You can, however, run AnyConnect Essentials and AnyConnect Premium licenses on different ASAs in the same network. By default, the ASA uses the AnyConnect Essentials license, but you can disable it to use other licenses by using the Configuration > Remote Access VPN > Network (Client) Access > Advanced > AnyConnect Essentials pane.
SSL VPN license changed to AnyConnect Premium SSL VPN Edition license	8.2(1)	The SSL VPN license name was changed to the AnyConnect Premium SSL VPN Edition license.

Feature Name	Platform Releases	Description
Shared Licenses for SSL VPN	8.2(1)	Shared licenses for SSL VPN were introduced. Multiple ASAs can share a pool of SSL VPN sessions on an as-needed basis.
Mobility Proxy application no longer requires Unified Communications Proxy license	8.2(2)	The Mobility Proxy no longer requires the UC Proxy license.
10 GE I/O license for the ASA 5585-X with SSP-20	8.2(3)	We introduced the 10 GE I/O license for the ASA 5585-X with SSP-20 to enable 10-Gigabit Ethernet speeds for the fiber ports. The SSP-60 supports 10-Gigabit Ethernet speeds by default. **Note** The ASA 5585-X is not supported in 8.3(x).
10 GE I/O license for the ASA 5585-X with SSP-10	8.2(4)	We introduced the 10 GE I/O license for the ASA 5585-X with SSP-10 to enable 10-Gigabit Ethernet speeds for the fiber ports. The SSP-40 supports 10-Gigabit Ethernet speeds by default. **Note** The ASA 5585-X is not supported in 8.3(x).
Non-identical failover licenses	8.3(1)	Failover licenses no longer need to be identical on each unit. The license used for both units is the combined license from the primary and secondary units. We modified the following screen: Configuration > Device Management > Licensing > Activation Key.
Stackable time-based licenses	8.3(1)	Time-based licenses are now stackable. In many cases, you might need to renew your time-based license and have a seamless transition from the old license to the new one. For features that are only available with a time-based license, it is especially important that the license not expire before you can apply the new license. The ASA allows you to *stack* time-based licenses so that you do not have to worry about the license expiring or about losing time on your licenses because you installed the new one early.
Intercompany Media Engine License	8.3(1)	The IME license was introduced.
Multiple time-based licenses active at the same time	8.3(1)	You can now install multiple time-based licenses, and have one license per feature active at a time. The following screen was modified: Configuration > Device Management > Licensing > Activation Key.
Discrete activation and deactivation of time-based licenses.	8.3(1)	You can now activate or deactivate time-based licenses using a command. We modified the following screen: Configuration > Device Management > Licensing > Activation Key.

Feature Name	Platform Releases	Description
AnyConnect Premium SSL VPN Edition license changed to AnyConnect Premium SSL VPN license	8.3(1)	The AnyConnect Premium SSL VPN Edition license name was changed to the AnyConnect Premium SSL VPN license.
No Payload Encryption image for export	8.3(2)	If you install the No Payload Encryption software on the ASA 5505 through 5550, then you disable Unified Communications, strong encryption VPN, and strong encryption management protocols. **Note** This special image is only supported in 8.3(x); for No Payload Encryption support in 8.4(1) and later, you need to purchase a special hardware version of the ASA.
Increased contexts for the ASA 5550, 5580, and 5585-X	8.4(1)	For the ASA 5550 and ASA 5585-X with SSP-10, the maximum contexts was increased from 50 to 100. For the ASA 5580 and 5585-X with SSP-20 and higher, the maximum was increased from 50 to 250.
Increased VLANs for the ASA 5580 and 5585-X	8.4(1)	For the ASA 5580 and 5585-X, the maximum VLANs was increased from 250 to 1024.
Increased connections for the ASA 5580 and 5585-X	8.4(1)	We increased the firewall connection limits: • ASA 5580-20—1,000,000 to 2,000,000. • ASA 5580-40—2,000,000 to 4,000,000. • ASA 5585-X with SSP-10: 750,000 to 1,000,000. • ASA 5585-X with SSP-20: 1,000,000 to 2,000,000. • ASA 5585-X with SSP-40: 2,000,000 to 4,000,000. • ASA 5585-X with SSP-60: 2,000,000 to 10,000,000.
AnyConnect Premium SSL VPN license changed to AnyConnect Premium license	8.4(1)	The AnyConnect Premium SSL VPN license name was changed to the AnyConnect Premium license. The license information display was changed from “SSL VPN Peers” to “AnyConnect Premium Peers.”
Increased AnyConnect VPN sessions for the ASA 5580	8.4(1)	The AnyConnect VPN session limit was increased from 5,000 to 10,000.
Increased Other VPN sessions for the ASA 5580	8.4(1)	The other VPN session limit was increased from 5,000 to 10,000.

Feature Name	Platform Releases	Description
IPsec remote access VPN using IKEv2	8.4(1)	IPsec remote access VPN using IKEv2 was added to the AnyConnect Essentials and AnyConnect Premium licenses. **Note** The following limitation exists in our support for IKEv2 on the ASA: We currently do not support duplicate security associations. IKEv2 site-to-site sessions were added to the Other VPN license (formerly IPsec VPN). The Other VPN license is included in the Base license.
No Payload Encryption hardware for export	8.4(1)	For models available with No Payload Encryption (for example, the ASA 5585-X), the ASA software disables Unified Communications and VPN features, making the ASA available for export to certain countries.
Dual SSPs for SSP-20 and SSP-40	8.4(2)	For SSP-40 and SSP-60, you can use two SSPs of the same level in the same chassis. Mixed-level SSPs are not supported (for example, an SSP-40 with an SSP-60 is not supported). Each SSP acts as an independent device, with separate configurations and management. You can use the two SSPs as a failover pair if desired. When using two SSPs in the chassis, VPN is not supported; note, however, that VPN has not been disabled.
IPS Module license for the ASA 5512-X through ASA 5555-X	8.6(1)	The IPS SSP software module on the ASA 5512-X, ASA 5515-X, ASA 5525-X, ASA 5545-X, and ASA 5555-X requires the IPS module license.
Clustering license for the ASA 5580 and ASA 5585-X.	9.0(1)	A clustering license was added for the ASA 5580 and ASA 5585-X.
Support for VPN on the ASASM	9.0(1)	The ASASM now supports all VPN features.
Unified communications support on the ASASM	9.0(1)	The ASASM now supports all Unified Communications features.
ASA 5585-X Dual SSP support for the SSP-10 and SSP-20 (in addition to the SSP-40 and SSP-60); VPN support for Dual SSPs	9.0(1)	The ASA 5585-X now supports dual SSPs using all SSP models (you can use two SSPs of the same level in the same chassis). VPN is now supported when using dual SSPs.
ASA 5500-X support for clustering	9.1(4)	The ASA 5512-X, ASA 5515-X, ASA 5525-X, ASA 5545-X, and ASA 5555-X now support 2-unit clusters. Clustering for 2 units is enabled by default in the base license; for the ASA 5512-X, you need the Security Plus license.
Support for 16 cluster members for the ASA 5585-X	9.2(1)	The ASA 5585-X now supports 16-unit clusters.
ASAv4 and ASAv30 Standard and Premium model licenses introduced	9.2(1)	The ASAv was introduced with a simple licensing scheme: ASAv4 and ASAv30 permanent licenses in Standard or Premium levels. No add-on licenses are available.



CHAPTER 5

Licenses: Smart Software Licensing

Smart Software Licensing lets you purchase and manage a pool of licenses centrally. Unlike product authorization key (PAK) licenses, smart licenses are not tied to a specific serial number. You can easily deploy or retire ASAs without having to manage each unit's license key. Smart Software Licensing also lets you see your license usage and needs at a glance.

Note For more information about Smart Licensing features and behaviors per platform, see Smart Enabled Product Families.

- About Smart Software Licensing, on page 113
- Prerequisites for Smart Software Licensing, on page 137
- Guidelines for Smart Software Licensing, on page 138
- Defaults for Smart Software Licensing, on page 138
- ASA Virtual: Configure Smart Software Licensing, on page 139
- 1000/1200/3100/4200: Configure Smart Software Licensing, on page 156
- Firepower 4100/9300: Configure Smart Software Licensing, on page 170
- Licenses Per Model, on page 171
- License PIDs Per Model, on page 184
- Monitoring Smart Software Licensing, on page 189
- Smart Software Manager Communication, on page 189
- History for Smart Software Licensing, on page 191

About Smart Software Licensing

Cisco Smart Licensing is a flexible licensing model that provides you with an easier, faster, and more consistent way to purchase and manage software across the Cisco portfolio and across your organization. And it's secure—you control what users can access. With Smart Licensing you get:

- **Easy Activation:** Smart Licensing establishes a pool of software licenses that can be used across the entire organization—no more PAKs (Product Activation Keys).
- **Unified Management:** My Cisco Entitlements (MCE) provides a complete view into all of your Cisco products and services in an easy-to-use portal, so you always know what you have and what you are using.

- **License Flexibility:** Your software is not node-locked to your hardware, so you can easily use and transfer licenses as needed.

To use Smart Licensing, you must first set up a Smart Account on Cisco Software Central (software.cisco.com).

For a more detailed overview on Cisco Licensing, go to cisco.com/go/licensingguide

Smart Software Licensing for the ASA on the Firepower 4100/9300 Chassis

For the ASA on the Firepower 4100/9300 chassis, Smart Software Licensing configuration is split between the Firepower 4100/9300 chassis supervisor and the ASA.

- Firepower 4100/9300 chassis—Configure all Smart Software Licensing infrastructure on the chassis, including parameters for communicating with the Smart Software Manager. The Firepower 4100/9300 chassis itself does not require any licenses to operate. See the FXOS configuration guide for licensing procedures.

Note Inter-chassis clustering requires that you enable the same Smart Licensing method on each chassis in the cluster.

- ASA Application—Configure all license entitlements in the ASA.

Smart Software Manager and Accounts

When you purchase 1 or more licenses for the device, you manage them in the Cisco Smart Software Manager:

https://software.cisco.com/#module/SmartLicensing

The Smart Software Manager lets you create a master account for your organization.

Note If you do not yet have an account, click the link to set up a new account. The Smart Software Manager lets you create a master account for your organization.

By default, your licenses are assigned to the *Default Virtual Account* under your master account. As the account administrator, you can optionally create additional virtual accounts; for example, you can create accounts for regions, departments, or subsidiaries. Multiple virtual accounts let you more easily manage large numbers of licenses and devices.

Offline Management

If your devices do not have internet access, and cannot register with the Smart Software Manager, you must configure offline licensing.

Permanent License Reservation

If your devices cannot access the internet for security reasons, you can optionally request permanent licenses for each ASA. Permanent licenses do not require periodic access to the Smart Software Manager. As with

PAK licenses, you can purchase a license and install the license key for the ASA. Unlike a PAK license, you can obtain and manage the licenses with the Smart Software Manager. You can easily switch between regular smart licensing mode and permanent license reservation mode.

Note ASA does not support Specific License Reservation (SLR). In SLR, specific feature entitlements are enabled permanently. ASA supports only PLR, where all the features are enabled permanently.

ASA Virtual Permanent License Reservation

Note Permanent license reservation is supported only on VMware and KVM.

You can obtain a model-specific license that enables all of the following features:

- Maximum throughput for your model
- Essentials tier
- Strong Encryption (3DES and AES) license, if you have enabled it in your Smart Licensing account
- Secure Client capabilities enabled for the platform

Note Use of Secure Client features is contingent on your purchase of the Secure Client license that enables the right to use Secure Client (see Secure Client Advantage, Secure Client Premier, and Secure Client VPN Only Licenses, on page 121).

Cisco offers two permanent license reservation modes for ASA virtual:

- Permanent License Reservation Mode Based on Memory and vCPU
- Flexible Permanent License Reservation Mode

Permanent License Reservation Mode Based on Memory and vCPU

You can configure a model-specific license based on the RAM and vCPUs assigned to your ASA virtual. For example, an ASA virtual with 8 GB RAM and 4 vCPUs always uses an ASAv30 license with 2G throughput.

The vCPU and memory-to-license relationships are as follows:

- 2 GB, 1 vCPU—ASAv5 (100 M) (You must run the **license smart set_plr5** command; otherwise, the ASAv10 license is assigned to allow 1-G throughput.)

Note In Version 9.13, the ASAv5 RAM requirements were increased to 2 GB. Because of this increase, the ASAv5 permanent license no longer worked because the ASA checked the memory assigned and determined that 2 GB of RAM was actually an ASAv10, not an ASAv5. To allow the ASAv5 permanent license to work, you must configure the ASA to recognize the extra memory for the model.

- 2 GB, 1 vCPU—ASAv10 (1G)
- 8 GB, 4 vCPUs —ASAv30 (2G)
- 16 GB, 8 vCPUs—ASAv50 (10G)
- 32 GB, 16 vCPUs—ASAv100 (20G)
- 64 GB, 32 vCPU—ASAvU
- 128 GB, 64 vCPU—ASAvU

Later, if you want to change the model level of a unit, you will have to return the current license and request a new license at the correct model level. To change the model of an already deployed ASA virtual, from the hypervisor you can change the vCPUs and DRAM settings to match the new model requirements. See the *ASA virtual Virtual Getting Started Guide* for these values.

If you stop using a license, you must return the license by generating a return code on ASA virtual and then enter that code into the Smart Software Manager. You must follow the return process correctly to ensure that you do not pay for unused licenses.

Flexible Permanent License Reservation Mode

You can configure any model-specific license irrespective of the RAM and vCPUs assigned to the ASA virtual.

The model license you specify sets the throughput level of the ASA virtual. The memory of your ASA virtual determines the maximum Secure Client sessions, TLS proxy sessions, the maximum concurrent firewall connections and VLANs, and not by the model license. However, a lower ASA virtual memory profile caps the actual number of sessions and features.

For example, if an ASA virtual has 8 GB RAM and is registered with ASAV_50_UNIVERSAL, the maximum allocated Secure Client sessions is 750. When the sessions exceed this limit, the connections are dropped.

The other benefits of the permanent license reservation license mode are:

- You can switch between permanent license reservation licenses irrespective of the memory allocated to the ASA virtual. However, if the ASA virtual has more than 16 vCPUs, you can configure only the ASAvU (Unlimited) license.
- You can change the memory and vCPUs assigned to an ASA virtual without changing the model license.

If you stop using a license, you must return the license by generating a return code on the ASA virtual and then enter that code into the Smart Software Manager. You must follow the return process correctly to ensure that you do not pay for unused licenses.

For more information about configuring permanent license reservation, see ASA Virtual: Configure Permanent License Reservation, on page 148.

Firepower 1010 Permanent License Reservation

You can obtain a license that enables all features:

- Essentials tier
- Security Plus
- Strong Encryption (3DES/AES) license if your account qualifies

- Secure Client capabilities enabled to the platform maximum.

  Use of Secure Client features is contingent on your purchase of the Secure Client license that enables the right to use Secure Client (see Secure Client Advantage, Secure Client Premier, and Secure Client VPN Only Licenses, on page 121).

Note You also need to request the entitlements in the ASA configuration so that the ASA allows their use.

If you stop using a license, you must return the license by generating a return code on the ASA, and then entering that code into the Smart Software Manager. Make sure you follow the return process correctly so you do not pay for unused licenses.

Firepower 1100 Permanent License Reservation

You can obtain a license that enables all features:

- Essentials tier
- Maximum Security Contexts
- Strong Encryption (3DES/AES) license if your account qualifies
- Secure Client capabilities enabled to the platform maximum.

  Use of Secure Client features is contingent on your purchase of the Secure Client license that enables the right to use Secure Client (see Secure Client Advantage, Secure Client Premier, and Secure Client VPN Only Licenses, on page 121).

Note You also need to request the entitlements in the ASA configuration so that the ASA allows their use.

If you stop using a license, you must return the license by generating a return code on the ASA, and then entering that code into the Smart Software Manager. Make sure you follow the return process correctly so you do not pay for unused licenses.

Secure Firewall 1210/1220 Permanent License Reservation

You can obtain a license that enables all features:

- Essentials tier
- Maximum Security Contexts
- Strong Encryption (3DES/AES) license if your account qualifies
- Secure Client capabilities enabled to the platform maximum.

  Use of Secure Client features is contingent on your purchase of the Secure Client license that enables the right to use Secure Client (see Secure Client Advantage, Secure Client Premier, and Secure Client VPN Only Licenses, on page 121).

Note You also need to request the entitlements in the ASA configuration so that the ASA allows their use.

If you stop using a license, you must return the license by generating a return code on the ASA, and then entering that code into the Smart Software Manager. Make sure you follow the return process correctly so you do not pay for unused licenses.

Secure Firewall 1230/1240/1250 Permanent License Reservation

You can obtain a license that enables all features:

- Essentials tier
- Maximum Security Contexts
- Strong Encryption (3DES/AES) license if your account qualifies
- Secure Client capabilities enabled to the platform maximum.

  Use of Secure Client features is contingent on your purchase of the Secure Client license that enables the right to use Secure Client (see Secure Client Advantage, Secure Client Premier, and Secure Client VPN Only Licenses, on page 121).

Note You also need to request the entitlements in the ASA configuration so that the ASA allows their use.

If you stop using a license, you must return the license by generating a return code on the ASA, and then entering that code into the Smart Software Manager. Make sure you follow the return process correctly so you do not pay for unused licenses.

Secure Firewall 3100/4200 Permanent License Reservation

You can obtain a license that enables all features:

- Essentials tier
- Maximum Security Contexts
- Carrier license
- Strong Encryption (3DES/AES) license if your account qualifies
- Secure Client capabilities enabled to the platform maximum.

  Use of Secure Client features is contingent on your purchase of the Secure Client license that enables the right to use Secure Client (see Secure Client Advantage, Secure Client Premier, and Secure Client VPN Only Licenses, on page 121).

Note You also need to request the entitlements in the ASA configuration so that the ASA allows their use.

If you stop using a license, you must return the license by generating a return code on the ASA, and then entering that code into the Smart Software Manager. Make sure you follow the return process correctly so you do not pay for unused licenses.

Firepower 4100/9300 Chassis Permanent License Reservation

You can obtain a license that enables all features:

- Essentials tier.
- Maximum Security Contexts
- Carrier license
- Strong Encryption (3DES/AES) license if your account qualifies
- Secure Client capabilities enabled to the platform maximum.

  Use of Secure Client features is contingent on your purchase of the Secure Client license that enables the right to use Secure Client (see Secure Client Advantage, Secure Client Premier, and Secure Client VPN Only Licenses, on page 121).



Note The license is managed on the Firepower 4100/9300 chassis, but you also need to request the entitlements in the ASA configuration so that the ASA allows their use. To manage the license, see the FXOS configuration guide.

If you stop using a license, you must return the license by generating a return code on the Firepower 4100/9300 chassis, and then entering that code into the Smart Software Manager. Make sure you follow the return process correctly so you don't pay for unused licenses.

Smart Software Manager On-Prem

If your devices cannot access the internet for security reasons, you can optionally install a local Smart Software Manager On-Prem (formerly known as "Smart Software Satellite Server") server as a virtual machine (VM). The Smart Software Manager On-Prem provides a subset of Smart Software Manager functionality, and allows you to provide essential licensing services for all your local devices. Only the Smart Software Manager On-Prem needs to connect periodically to the main Smart Software Manager to sync your license usage. You can sync on a schedule or you can sync manually.

You can perform the following functions on the Smart Software Manager On-Prem:

- Activate or register a license
- View your company's licenses
- Transfer licenses between company entities

For more information, see Cisco Smart Software Manager On-Prem Data Sheet.

Licenses and Devices Managed per Virtual Account

Licenses and devices are managed per virtual account: only that virtual account's devices can use the licenses assigned to the account. If you need additional licenses, you can transfer an unused license from another virtual account. You can also transfer devices between virtual accounts.

For the ASA on the Firepower 4100/9300 chassis—Only the chassis registers as a device, while the ASA applications in the chassis request their own licenses. For example, for a Firepower 9300 chassis with 3 security modules, the chassis counts as one device, but the modules use 3 separate licenses.

Evaluation License

ASA Virtual

The ASA virtual does not support an evaluation mode. Before the ASA virtual registers with the Smart Software Manager, it operates in a severely rate-limited state.

Firepower 1000

Before the Firepower 1000 registers with the Smart Software Manager, it operates for 90 days (total usage) in evaluation mode. Only default entitlements are enabled. When this period ends, the Firepower 1000 becomes out-of-compliance.



Note You cannot receive an evaluation license for Strong Encryption (3DES/AES); you must register with the Smart Software Manager to receive the export-compliance token that enables the Strong Encryption (3DES/AES) license.

Firepower 2100

Before the Firepower 2100 registers with the Smart Software Manager, it operates for 90 days (total usage) in evaluation mode. Only default entitlements are enabled. When this period ends, the Firepower 2100 becomes out-of-compliance.



Note You cannot receive an evaluation license for Strong Encryption (3DES/AES); you must register with the Smart Software Manager to receive the export-compliance token that enables the Strong Encryption (3DES/AES) license.

Secure Firewall 3100/4200

Before the Secure Firewall 3100/4200 registers with the Smart Software Manager, it operates for 90 days (total usage) in evaluation mode. Only default entitlements are enabled. When this period ends, the Secure Firewall 3100/4200 becomes out-of-compliance.

Note You cannot receive an evaluation license for Strong Encryption (3DES/AES); you must register with the Smart Software Manager to receive the export-compliance token that enables the Strong Encryption (3DES/AES) license.

Firepower 4100/9300 Chassis

The Firepower 4100/9300 chassis supports two types of evaluation license:

- Chassis-level evaluation mode—Before the Firepower 4100/9300 chassis registers with the Smart Software Manager, it operates for 90 days (total usage) in evaluation mode. The ASA cannot request specific entitlements in this mode; only default entitlements are enabled. When this period ends, the Firepower 4100/9300 chassis becomes out-of-compliance.
- Entitlement-based evaluation mode—After the Firepower 4100/9300 chassis registers with the Smart Software Manager, you can obtain time-based evaluation licenses that can be assigned to the ASA. In the ASA, you request entitlements as usual. When the time-based license expires, you need to either renew the time-based license or obtain a permanent license.

Note You cannot receive an evaluation license for Strong Encryption (3DES/AES); you must register with the Smart Software Manager and obtain a permanent license to receive the export-compliance token that enables the Strong Encryption (3DES/AES) license.

About Licenses by Type

The following sections include additional information about licenses by type.

Secure Client Advantage, Secure Client Premier, and Secure Client VPN Only Licenses

Secure Client licenses are not applied directly to the ASA. However, you need to purchase licenses and add them to your Smart Account to guarantee the right to use the ASA as the Secure Client headend.

- For the Secure Client Advantage and Secure Client Premier licenses, add up the number of peers you intend to use across all the ASAs in your Smart Account and purchase license(s) for that many peers.
- For the Secure Client VPN Only, purchase one license per ASA. Unlike the other licenses that provide a pool of peers that can be shared by multiple ASAs, the Secure Client VPN Only license is per headend.

For more information, see:

- Cisco Secure Client Ordering Guide
- Secure Client Licensing Frequently Asked Questions (FAQ)

Other VPN Peers

Other VPN peers include the following VPN types:

- IPsec remote access VPN using IKEv1

- IPsec site-to-site VPN using IKEv1
- IPsec site-to-site VPN using IKEv2

This license is included in the Base license.

Total VPN Peers Combined, All Types

- The Total VPN Peers is the maximum VPN peers allowed of both Secure Client and Other VPN peers combined. For example, if the total is 1000, you can allow 500 Secure Client and 500 Other VPN peers simultaneously; or 700 Secure Client and 300 Other VPN; or use all 1000 for Secure Client. If you exceed the total VPN peers, you can overload the ASA, so be sure to size your network appropriately.

Encryption License

Strong Encryption: ASA Virtual

Strong Encryption (3DES/AES) is available for management connections before you connect to the Smart Software Manager or Smart Software Manager On-Prem server, so you can launch ASDM and connect to the Smart Software Manager. For through-the-box traffic that requires strong encryption (such as VPN), throughput is severely limited until you connect to the Smart Software Manager and obtain the Strong Encryption license.

When you request the registration token for the ASA virtual from your Smart Software Licensing account, check the **Allow export-controlled functionality on the products registered with this token** check box so that the Strong Encryption (3DES/AES) license is applied (your account must be qualified for its use). If the ASA virtual becomes out-of-compliance later, as long as the export compliance token was successfully applied, the ASA virtual will retain the license and not revert to the rate-limited state. The license is removed if you re-register the ASA virtual, and export compliance is disabled, or if you restore the ASA virtual to factory default settings.

If you initially register the ASA virtual without strong encryption and later add strong encryption, then you must reload the ASA virtual for the new license to take effect.

For permanent license reservation licenses, the Strong Encryption (3DES/AES) license is enabled if your account qualifies for its use.

If your Smart Account is not authorized for strong encryption, but Cisco has determined that you are allowed to use strong encryption, you can manually add a strong encryption license to your account.

Strong Encryption: Firepower 1000, Secure Firewall 1200/3100/4200

The ASA includes 3DES capability by default for management access only, so you can connect to the Smart Software Manager and also use ASDM immediately. You can also use SSH and SCP if you later configure SSH access on the ASA. Other features that require strong encryption (such as VPN) must have Strong Encryption enabled, which requires you to first register to the Smart Software Manager.



Note If you attempt to configure any features that can use strong encryption before you register—even if you only configure weak encryption—then your HTTPS connection will be dropped on that interface, and you cannot reconnect. The exception to this rule is if you are connected to a management-only interface, such as Management 1/1. SSH is not affected. If you lose your HTTPS connection, you can connect to the console port to reconfigure the ASA, connect to a management-only interface, or connect to an interface not configured for a strong encryption feature.

When you request the registration token for the ASA from your Smart Software Licensing account, check the **Allow export-controlled functionality on the products registered with this token** check box so that the Strong Encryption (3DES/AES) license is applied (your account must be qualified for its use). If the ASA becomes out-of-compliance later, as long as the export compliance token was successfully applied, the ASA will continue to allow through the box traffic. Even if you re-register the ASA, and export compliance is disabled, the license remains enabled. The license is removed if you restore the ASA to factory default settings.

If you initially register the ASA without strong encryption and later add strong encryption, then you must reload the ASA for the new license to take effect.

For permanent license reservation licenses, the Strong Encryption (3DES/AES) license is enabled if your account qualifies for its use.

If your Smart Account is not authorized for strong encryption, but Cisco has determined that you are allowed to use strong encryption, you can manually add a strong encryption license to your account.

Strong Encryption: Firepower 4100/9300 Chassis

When the ASA is deployed as a logical device, you can launch ASDM immediately. Through the box traffic that requires strong encryption (such as VPN) is not allowed until you connect and obtain the Strong Encryption license.

When you request the registration token for the chassis from your Smart Software Licensing account, check the **Allow export-controlled functionality on the products registered with this token** check box so that the Strong Encryption (3DES/AES) license is applied (your account must be qualified for its use).

If the ASA becomes out-of-compliance later, as long as the export compliance token was successfully applied, the ASA will continue to allow through the box traffic. The license is removed if you re-register the chassis, and export compliance is disabled, or if you restore the chassis to factory default settings.

If you initially register the chassis without strong encryption and later add strong encryption, then you must reload the ASA application for the new license to take effect.

For permanent license reservation licenses, the Strong Encryption (3DES/AES) license is enabled if your account qualifies for its use.

If your Smart Account is not authorized for strong encryption, but Cisco has determined that you are allowed to use strong encryption, you can manually add a strong encryption license to your account.

DES: All Models

If you have Strong Encryption enabled, you cannot use DES.

Carrier License

The Carrier license enables the following inspection features:

- Diameter—Diameter is an Authentication, Authorization, and Accounting (AAA) protocol used in next-generation mobile and fixed telecom networks such as EPS (Evolved Packet System) for LTE (Long Term Evolution) and IMS (IP Multimedia Subsystem). It replaces RADIUS and TACACS in these networks.
- GTP/GPRS—GPRS Tunneling Protocol is used in GSM, UMTS and LTE networks for general packet radio service (GPRS) traffic. GTP provides a tunnel control and management protocol to provide GPRS network access for a mobile station by creating, modifying, and deleting tunnels. GTP also uses a tunneling mechanism for carrying user data packets.

- M3UA—MTP3 User Adaptation (M3UA) is a client/server protocol that provides a gateway to the SS7 network for IP-based applications that interface with the SS7 Message Transfer Part 3 (MTP3) layer. M3UA makes it possible to run the SS7 User Parts (such as ISUP) over an IP network. M3UA is defined in RFC 4666.
- SCTP—SCTP (Stream Control Transmission Protocol) is described in RFC 4960. The protocol supports the telephony signaling protocol SS7 over IP and is also a transport protocol for several interfaces in the 4G LTE mobile network architecture.

Total TLS Proxy Sessions

Each TLS proxy session for Encrypted Voice Inspection is counted against the TLS license limit.

Other applications that use TLS proxy sessions do not count toward the TLS limit, for example, Mobility Advantage Proxy (which does not require a license).

Some applications might use multiple sessions for a connection. For example, if you configure a phone with a primary and backup Cisco Unified Communications Manager, there are 2 TLS proxy connections.

You independently set the TLS proxy limit using the **tls-proxy maximum-sessions** command or in ASDM, using the **Configuration > Firewall > Unified Communications > TLS Proxy** pane. To view the limits of your model, enter the **tls-proxy maximum-sessions ?** command. When you apply a TLS proxy license that is higher than the default TLS proxy limit, the ASA automatically sets the TLS proxy limit to match the license. The TLS proxy limit takes precedence over the license limit; if you set the TLS proxy limit to be less than the license, then you cannot use all of the sessions in your license.



Note For license part numbers ending in "K8" (for example, licenses under 250 users), TLS proxy sessions are limited to 1000. For license part numbers ending in "K9" (for example, licenses 250 users or larger), the TLS proxy limit depends on the configuration, up to the model limit. K8 and K9 refer to whether the license is restricted for export: K8 is unrestricted, and K9 is restricted.

If you clear the configuration (using the **clear configure all** command, for example), then the TLS proxy limit is set to the default for your model; if this default is lower than the license limit, then you see an error message to use the **tls-proxy maximum-sessions** command to raise the limit again (in ASDM, use the **TLS Proxy** pane). If you use failover and enter the **write standby** command or in ASDM, use **File > Save Running Configuration to Standby Unit** on the primary unit to force a configuration synchronization, the **clear configure all** command is generated on the secondary unit automatically, so you may see the warning message on the secondary unit. Because the configuration synchronization restores the TLS proxy limit set on the primary unit, you can ignore the warning.

You might also use SRTP encryption sessions for your connections:

- For K8 licenses, SRTP sessions are limited to 250.
- For K9 licenses, there is no limit.



Note Only calls that require encryption/decryption for media are counted toward the SRTP limit; if passthrough is set for the call, even if both legs are SRTP, they do not count toward the limit.

VLANs, Maximum

For an interface to count against the VLAN limit, you must assign a VLAN to it.

Botnet Traffic Filter License

Requires a Strong Encryption (3DES/AES) License to download the dynamic database.

Failover or ASA Cluster Licenses

Failover Licenses for the ASAv

The standby unit requires the same model license as the primary unit.

Failover Licenses for the Firepower 1010

Smart Software Manager Regular and On-Prem

Both Firepower 1010 units must be registered with the Smart Software Manager or Smart Software Manager On-Prem server. Both units require you to enable the Essentials license and the Security Plus license *before* you can configure failover.

Typically, you do not also need to enable the Strong Encryption (3DES/AES) feature license in the ASA, because both units should have obtained the Strong Encryption token when you registered the units. When using the registration token, both units must have the same encryption level.

If your Smart Account is not authorized for strong encryption, but Cisco has determined that you are allowed to use strong encryption, you can manually add a strong encryption license to your account. In this case, enable it on the active unit after you enable failover. The configuration is replicated to the standby unit, but the standby unit does not use the configuration; it remains in a cached state. Only the active unit requests the license from the server. The license is aggregated into a single failover license that is shared by the failover pair, and this aggregated license is also cached on the standby unit to be used if it becomes the active unit in the future. After a failover, the new active unit continues to use the aggregated license. It uses the cached license configuration to re-request the entitlement from the server. When the old active unit rejoins the pair as a standby unit, it releases the license entitlement. Before the standby unit releases the entitlement, the new active unit's license might be in a non-compliant state if there are no available licenses in the account. The failover pair can use the aggregated license for 30 days, but if it is still non-compliant after the grace period, and you are not using the Strong Encryption token, then you will not be able to make configuration changes to features requiring the Strong Encryption (3DES/AES) feature license; operation is otherwise unaffected. The new active unit sends an entitlement authorization renewal request every 35 seconds until the license is compliant. If you disband the failover pair, then the active unit releases the entitlements, and both units retain the licensing configuration in a cached state. To re-activate licensing, you need to clear the configuration on each unit, and re-configure it.

Permanent License Reservation

For permanent license reservation, you must purchase separate licenses for each chassis and enable the licenses *before* you configure failover.

Failover Licenses for the Firepower 1100

Smart Software Manager Regular and On-Prem

Only the active unit requests licenses from the server. Licenses are aggregated into a single failover license that is shared by the failover pair. There is no extra cost for secondary units.

After you enable failover for Active/Standby failover, you can only configure smart licensing on the active unit. For Active/Active failover, you can only configure smart licensing on the unit with failover group 1 as active. The configuration is replicated to the standby unit, but the standby unit does not use the configuration; it remains in a cached state. The aggregated license is also cached on the standby unit to be used if it becomes the active unit in the future.



Note Each ASA must have the same encryption license when forming a failover pair. When you register an ASA to the smart licensing server, the Strong Encryption license is automatically enabled for qualified customers when you apply the registration token. Because of this requirement, you have two choices for licensing when you use the Strong Encryption token with failover:

- Before you enable failover, register both units to the smart licensing server. In this case, both units will have strong encryption. Then, after you enable failover, continue configuring license entitlements on the active unit. If you enable encryption for the failover link, it will use AES/3DES (strong encryption).
- Before you register the active unit to the smart licensing server, enable failover. In this case, both units will not yet have strong encryption. Then configure license entitlements and register the active unit to the smart licensing server; both units will get strong encryption from the aggregated license. Note that if you enabled encryption on the failover link, it will use DES (weak encryption) because the failover link was established before the units gained strong encryption. You must reload *both* units to use AES/3DES on the link. If you only reload one unit, then that unit will try to use AES/3DES while the original unit uses DES, which will result in both units becoming active (split brain).

Each add-on license type is managed as follows:

- Essentials—Although only the active unit requests this license from the server, the standby unit has the Essentials license enabled by default; it does not need to register with the server to use it.
- Context—Only the active unit requests this license. However, the Essentials license includes 2 contexts by default and is present on both units. The value from each unit's Essentials license plus the value of the Context license on the active unit are combined up to the platform limit. For example:
  - Active/Standby: The Essentials license includes 2 contexts; for two Firepower 1120 units, these licenses add up to 4 contexts. You configure a 3-Context license on the active unit in an Active/Standby pair. Therefore, the aggregated failover license includes 7 contexts. However, because the platform limit for one unit is 5, the combined license allows a maximum of 5 contexts only. In this case, you might only configure the active Context license to be 1 context.
  - Active/Active: The Essentials license includes 2 contexts; for two Firepower 1140 units, these licenses add up to 4 contexts. You configure a 4-Context license on the primary unit in an Active/Active pair. Therefore, the aggregated failover license includes 8 contexts. One unit can use 5 contexts and the other unit can use 3 contexts, for example; but during a failure, one unit will use all 8. Because the platform limit for one unit is 10, the combined license allows a maximum of 10 contexts; the 8 contexts are within the limit.
- Strong Encryption (3DES/AES)—If your Smart Account is not authorized for strong encryption, but Cisco has determined that you are allowed to use strong encryption, you can manually add a strong encryption license to your account. Only the active unit requests this license, and both units can use it due to license aggregation.

After a failover, the new active unit continues to use the aggregated license. It uses the cached license configuration to re-request the entitlement from the server. When the old active unit rejoins the pair as a standby unit, it releases the license entitlement. Before the standby unit releases the entitlement, the new active unit's license might be in a non-compliant state if there are no available licenses in the account. The failover pair can use the aggregated license for 30 days, but if it is still non-compliant after the grace period, you will not be able to make configuration changes to features requiring special licenses (i.e. add an extra context); operation is otherwise unaffected. The new active unit sends an entitlement authorization renewal request every 35 seconds until the license is compliant. If you disband the failover pair, then the active unit releases the entitlements, and both units retain the licensing configuration in a cached state. To re-activate licensing, you need to clear the configuration on each unit, and re-configure it.

Permanent License Reservation

For permanent license reservation, you must purchase separate licenses for each chassis and enable the licenses *before* you configure failover.

Failover Licenses for the Secure Firewall 1210/1220

Smart Software Manager Regular and On-Prem

Each unit requires the Essentials license (enabled by default) and the same encryption license. We recommend licensing each unit with the licensing server *before* you enable failover to avoid licensing mismatch issues, and when using the Strong Encryption license, issues with failover link encryption.

There is no extra cost for the Context license on data units.

The Strong Encryption license is automatically enabled for qualified customers when you apply the registration token. For the optional Strong Encryption (3DES/AES) feature license enabled in the ASA configuration, see below.

In the ASA license configuration, the Essentials license is always enabled by default on both units. After you enable failover for Active/Standby failover, you can only configure smart licensing on the active unit. For Active/Active failover, you can only configure smart licensing on the unit with failover group 1 as active. The configuration is replicated to the standby unit, but the standby unit does not use the configuration; it remains in a cached state. The aggregated license is also cached on the standby unit to be used if it becomes the active unit in the future.

Each add-on license type is managed as follows:

- Essentials—Each unit requests a Essentials license from the server.
- Context—Only the active unit requests this license. However, the Essentials license includes 2 contexts by default and is present on both units. The value from each unit’s Essentials license plus the value of the Context license on the active unit are combined up to the platform limit. For example:
  - Active/Standby: The Essentials license includes 2 contexts; for two Secure Firewall 1210 units, these licenses add up to 4 contexts. You configure a 5-Context license on the active unit in an Active/Standby pair. Therefore, the aggregated failover license includes 9 contexts. However, because the platform limit for one unit is 5, the combined license allows a maximum of 5 contexts only. In this case, you might only configure the active Context license to be 1 context.
  - Active/Active: The Essentials license includes 2 contexts; for two Secure Firewall 1220 units, these licenses add up to 4 contexts. You configure a 5-Context license on the primary unit in an Active/Active pair. Therefore, the aggregated failover license includes 9 contexts. One unit can use 5 contexts and the other unit can use 4 contexts, for example; but during a failure, one unit will use all 9. Because the platform limit for one unit is 10, the combined license allows a maximum of 10 contexts; the 9 contexts are within the limit.

- Strong Encryption (3DES/AES)—If your Smart Account is not authorized for strong encryption, but Cisco has determined that you are allowed to use strong encryption, you can manually add a strong encryption license to your account. Only the active unit requests this license, and both units can use it due to license aggregation.

After a failover, the new active unit continues to use the aggregated license. It uses the cached license configuration to re-request the entitlement from the server. When the old active unit rejoins the pair as a standby unit, it releases the license entitlement. Before the standby unit releases the entitlement, the new active unit's license might be in a non-compliant state if there are no available licenses in the account. The failover pair can use the aggregated license for 30 days, but if it is still non-compliant after the grace period, you will not be able to make configuration changes to features requiring special licenses; operation is otherwise unaffected. The new active unit sends an entitlement authorization renewal request every 35 seconds until the license is compliant. If you disband the failover pair, then the active unit releases the entitlements, and both units retain the licensing configuration in a cached state. To re-activate licensing, you need to clear the configuration on each unit, and re-configure it.

Permanent License Reservation

For permanent license reservation, you must purchase separate licenses for each chassis and enable the licenses *before* you configure failover.

Failover Licenses for the Secure Firewall 1230/1240/1250

Smart Software Manager Regular and On-Prem

Each unit requires the Essentials license (enabled by default) and the same encryption license. We recommend licensing each unit with the licensing server *before* you enable failover to avoid licensing mismatch issues, and when using the Strong Encryption license, issues with failover link encryption.

The failover feature itself does not require any licenses. There is no extra cost for the Context license on data units.

The Strong Encryption license is automatically enabled for qualified customers when you apply the registration token. For the optional Strong Encryption (3DES/AES) feature license enabled in the ASA configuration, see below.

In the ASA license configuration, the Essentials license is always enabled by default on both units. After you enable failover for Active/Standby failover, you can only configure smart licensing on the active unit. For Active/Active failover, you can only configure smart licensing on the unit with failover group 1 as active. The configuration is replicated to the standby unit, but the standby unit does not use the configuration; it remains in a cached state. The aggregated license is also cached on the standby unit to be used if it becomes the active unit in the future.

Each add-on license type is managed as follows:

- Essentials—Each unit requests a Essentials license from the server.
- Context—Only the active unit requests this license. However, the Essentials license includes 2 contexts by default and is present on both units. The value from each unit’s Essentials license plus the value of the Context license on the active unit are combined up to the platform limit. For example:
  - Active/Standby: The Essentials license includes 2 contexts; for two Secure Firewall 1230 units, these licenses add up to 4 contexts. You configure a 25-Context license on the active unit in an Active/Standby pair. Therefore, the aggregated failover license includes 29 contexts. However, because the platform limit for one unit is 25, the combined license allows a maximum of 25 contexts only. In this case, you might only configure the active Context license to be 21 contexts.

- Active/Active: The Essentials license includes 2 contexts; for two Secure Firewall 1230 units, these licenses add up to 4 contexts. You configure a 10-Context license on the primary unit in an Active/Active pair. Therefore, the aggregated failover license includes 14 contexts. One unit can use 9 contexts and the other unit can use 5 contexts, for example; but during a failure, one unit will use all 14. Because the platform limit for one unit is 25, the combined license allows a maximum of 25 contexts; the 14 contexts are within the limit.

- Strong Encryption (3DES/AES)—If your Smart Account is not authorized for strong encryption, but Cisco has determined that you are allowed to use strong encryption, you can manually add a strong encryption license to your account. Only the active unit requests this license, and both units can use it due to license aggregation.

After a failover, the new active unit continues to use the aggregated license. It uses the cached license configuration to re-request the entitlement from the server. When the old active unit rejoins the pair as a standby unit, it releases the license entitlement. Before the standby unit releases the entitlement, the new active unit's license might be in a non-compliant state if there are no available licenses in the account. The failover pair can use the aggregated license for 30 days, but if it is still non-compliant after the grace period, you will not be able to make configuration changes to features requiring special licenses (i.e. add an extra context); operation is otherwise unaffected. The new active unit sends an entitlement authorization renewal request every 35 seconds until the license is compliant. If you disband the failover pair, then the active unit releases the entitlements, and both units retain the licensing configuration in a cached state. To re-activate licensing, you need to clear the configuration on each unit, and re-configure it.

Permanent License Reservation

For permanent license reservation, you must purchase separate licenses for each chassis and enable the licenses *before* you configure failover.

Failover Licenses for the Secure Firewall 3100

Smart Software Manager Regular and On-Prem

Each unit requires the Essentials license (enabled by default) and the same encryption license. We recommend licensing each unit with the licensing server *before* you enable failover to avoid licensing mismatch issues, and when using the Strong Encryption license, issues with failover link encryption.

The failover feature itself does not require any licenses. There is no extra cost for the Context license on data units.

The Strong Encryption license is automatically enabled for qualified customers when you apply the registration token. For the optional Strong Encryption (3DES/AES) feature license enabled in the ASA configuration, see below.

In the ASA license configuration, the Essentials license is always enabled by default on both units. After you enable failover for Active/Standby failover, you can only configure smart licensing on the active unit. For Active/Active failover, you can only configure smart licensing on the unit with failover group 1 as active. The configuration is replicated to the standby unit, but the standby unit does not use the configuration; it remains in a cached state. The aggregated license is also cached on the standby unit to be used if it becomes the active unit in the future.

Each add-on license type is managed as follows:

- Essentials—Each unit requests a Essentials license from the server.

- Context—Only the active unit requests this license. However, the Essentials license includes 2 contexts by default and is present on both units. The value from each unit's Essentials license plus the value of the Context license on the active unit are combined up to the platform limit. For example:
  - Active/Standby: The Essentials license includes 2 contexts; for two Secure Firewall 3130 units, these licenses add up to 4 contexts. You configure a 100-Context license on the active unit in an Active/Standby pair. Therefore, the aggregated failover license includes 104 contexts. However, because the platform limit for one unit is 100, the combined license allows a maximum of 100 contexts only. In this case, you might only configure the active Context license to be 95 contexts.
  - Active/Active: The Essentials license includes 2 contexts; for two Secure Firewall 3130 units, these licenses add up to 4 contexts. You configure a 10-Context license on the primary unit in an Active/Active pair. Therefore, the aggregated failover license includes 14 contexts. One unit can use 9 contexts and the other unit can use 5 contexts, for example; but during a failure, one unit will use all 14. Because the platform limit for one unit is 100, the combined license allows a maximum of 100 contexts; the 14 contexts are within the limit.
- Strong Encryption (3DES/AES)—If your Smart Account is not authorized for strong encryption, but Cisco has determined that you are allowed to use strong encryption, you can manually add a strong encryption license to your account. Only the active unit requests this license, and both units can use it due to license aggregation.

After a failover, the new active unit continues to use the aggregated license. It uses the cached license configuration to re-request the entitlement from the server. When the old active unit rejoins the pair as a standby unit, it releases the license entitlement. Before the standby unit releases the entitlement, the new active unit's license might be in a non-compliant state if there are no available licenses in the account. The failover pair can use the aggregated license for 30 days, but if it is still non-compliant after the grace period, you will not be able to make configuration changes to features requiring special licenses (i.e. add an extra context); operation is otherwise unaffected. The new active unit sends an entitlement authorization renewal request every 35 seconds until the license is compliant. If you disband the failover pair, then the active unit releases the entitlements, and both units retain the licensing configuration in a cached state. To re-activate licensing, you need to clear the configuration on each unit, and re-configure it.

Permanent License Reservation

For permanent license reservation, you must purchase separate licenses for each chassis and enable the licenses *before* you configure failover.

Failover Licenses for the Secure Firewall 4200

Smart Software Manager Regular and On-Prem

Each unit requires the Essentials license (enabled by default) and the same encryption license. We recommend licensing each unit with the licensing server *before* you enable failover to avoid licensing mismatch issues, and when using the Strong Encryption license, issues with failover link encryption.

The failover feature itself does not require any licenses. There is no extra cost for the Context license on data units.

The Strong Encryption license is automatically enabled for qualified customers when you apply the registration token. For the optional Strong Encryption (3DES/AES) feature license enabled in the ASA configuration, see below.

In the ASA license configuration, the Essentials license is always enabled by default on both units. After you enable failover for Active/Standby failover, you can only configure smart licensing on the active unit. For

Active/Active failover, you can only configure smart licensing on the unit with failover group 1 as active. The configuration is replicated to the standby unit, but the standby unit does not use the configuration; it remains in a cached state. The aggregated license is also cached on the standby unit to be used if it becomes the active unit in the future.

Each add-on license type is managed as follows:

- Essentials—Each unit requests a StEssentialsandard license from the server.
- Context—Only the active unit requests this license. However, the Essentials license includes 2 contexts by default and is present on both units. The value from each unit's Essentials license plus the value of the Context license on the active unit are combined up to the platform limit. For example:
  - Active/Standby: The Essentials license includes 2 contexts; for two Secure Firewall 4215 units, these licenses add up to 4 contexts. You configure a 250-Context license on the active unit in an Active/Standby pair. Therefore, the aggregated failover license includes 254 contexts. However, because the platform limit for one unit is 250, the combined license allows a maximum of 250 contexts only. In this case, you might only configure the active Context license to be 246 contexts.
  - Active/Active: The Essentials license includes 2 contexts; for two Secure Firewall 4215 units, these licenses add up to 4 contexts. You configure a 10-Context license on the primary unit in an Active/Active pair. Therefore, the aggregated failover license includes 14 contexts. One unit can use 9 contexts and the other unit can use 5 contexts, for example; but during a failure, one unit will use all 14. Because the platform limit for one unit is 250, the combined license allows a maximum of 250 contexts; the 14 contexts are within the limit.
- Strong Encryption (3DES/AES)—If your Smart Account is not authorized for strong encryption, but Cisco has determined that you are allowed to use strong encryption, you can manually add a strong encryption license to your account. Only the active unit requests this license, and both units can use it due to license aggregation.

After a failover, the new active unit continues to use the aggregated license. It uses the cached license configuration to re-request the entitlement from the server. When the old active unit rejoins the pair as a standby unit, it releases the license entitlement. Before the standby unit releases the entitlement, the new active unit's license might be in a non-compliant state if there are no available licenses in the account. The failover pair can use the aggregated license for 30 days, but if it is still non-compliant after the grace period, you will not be able to make configuration changes to features requiring special licenses (i.e. add an extra context); operation is otherwise unaffected. The new active unit sends an entitlement authorization renewal request every 35 seconds until the license is compliant. If you disband the failover pair, then the active unit releases the entitlements, and both units retain the licensing configuration in a cached state. To re-activate licensing, you need to clear the configuration on each unit, and re-configure it.

Permanent License Reservation

For permanent license reservation, you must purchase separate licenses for each chassis and enable the licenses *before* you configure failover.

Failover Licenses for the Firepower 4100/9300

Smart Software Manager Regular and On-Prem

Both Firepower 4100/9300 must be registered with the Smart Software Manager or Smart Software Manager On-Prem server before you configure failover. There is no extra cost for secondary units.

The Strong Encryption license is automatically enabled for qualified customers when you apply the registration token. When using the token, each chassis must have the same encryption license. For the optional Strong Encryption (3DES/AES) feature license enabled in the ASA configuration, see below.

After you enable failover, for the ASA license configuration for Active/Standby failover, you can only configure smart licensing on the active unit. For Active/Active failover, you can only configure smart licensing on the unit with failover group 1 as active. The configuration is replicated to the standby unit, but the standby unit does not use the configuration; it remains in a cached state. Only the active unit requests the licenses from the server. The licenses are aggregated into a single failover license that is shared by the failover pair, and this aggregated license is also cached on the standby unit to be used if it becomes the active unit in the future. Each license type is managed as follows:

- Essentials—Although only the active unit requests this license from the server, the standby unit has the Essentials license enabled by default; it does not need to register with the server to use it.
- Context—Only the active unit requests this license. However, the Essentials license includes 10 contexts by default and is present on both units. The value from each unit's Essentials license plus the value of the Context license on the active unit are combined up to the platform limit. For example:
  - Active/Standby: The Essentials license includes 10 contexts; for 2 units, these licenses add up to 20 contexts. You configure a 250-Context license on the active unit in an Active/Standby pair. Therefore, the aggregated failover license includes 270 contexts. However, because the platform limit for one unit is 250, the combined license allows a maximum of 250 contexts only. In this case, you should only configure the active Context license to be 230 contexts.
  - Active/Active: The Essentials license includes 10 contexts; for 2 units, these licenses add up to 20 contexts. You configure a 10-Context license on the primary unit in an Active/Active pair. Therefore, the aggregated failover license includes 30 contexts. One unit can use 17 contexts and the other unit can use 13 contexts, for example; but during a failure, one unit will use all 30. Because the platform limit for one unit is 250, the combined license allows a maximum of 250 contexts; the 30 contexts are within the limit.
- Carrier—Only the active requests this license, and both units can use it due to license aggregation.
- Strong Encryption (3DES)—If your Smart Account is not authorized for strong encryption, but Cisco has determined that you are allowed to use strong encryption, you can manually add a strong encryption license to your account. Only the active unit requests this license, and both units can use it due to license aggregation.

After a failover, the new active unit continues to use the aggregated license. It uses the cached license configuration to re-request the entitlement from the server. When the old active unit rejoins the pair as a standby unit, it releases the license entitlement. Before the standby unit releases the entitlement, the new active unit's license might be in a non-compliant state if there are no available licenses in the account. The failover pair can use the aggregated license for 30 days, but if it is still non-compliant after the grace period, you will not be able to make configuration changes to features requiring special licenses; operation is otherwise unaffected. The new active unit sends an entitlement authorization renewal request every 35 seconds until the license is compliant. If you disband the failover pair, then the active unit releases the entitlements, and both units retain the licensing configuration in a cached state. To re-activate licensing, you need to clear the configuration on each unit, and re-configure it.

Permanent License Reservation

For permanent license reservation, you must purchase separate licenses for each chassis and enable the licenses *before* you configure failover.

ASA Cluster Licenses for the Secure Firewall 3100

Smart Software Manager Regular and On-Prem

Each unit requires the Essentials license (enabled by default) and the same encryption license. We recommend licensing each unit with the licensing server *before* you enable clustering to avoid licensing mismatch issues, and when using the Strong Encryption license, issues with cluster control link encryption.

The clustering feature itself does not require any licenses. There is no extra cost for the Context license on data units.

The Strong Encryption license is automatically enabled for qualified customers when you apply the registration token. For the optional Strong Encryption (3DES/AES) feature license enabled in the ASA configuration, see below.

In the ASA license configuration, the Essentials license is always enabled by default on all units. You can only configure smart licensing on the control unit. The configuration is replicated to the data units, but for some licenses, they do not use the configuration; it remains in a cached state, and only the control unit requests the license. The licenses are aggregated into a single cluster license that is shared by the cluster units, and this aggregated license is also cached on the data units to be used if one of them becomes the control unit in the future. Each license type is managed as follows:

- Essentials—Each unit requests a Essentials license from the server.
- Context—Only the control unit requests the Context license from the server. The Essentials license includes 2 contexts by default and is present on all cluster members. The value from each unit's Essentials license plus the value of the Context license on the control unit are combined up to the platform limit in an aggregated cluster license. For example:
  - You have 6 Secure Firewall 3100s in the cluster. The Essentials license includes 2 contexts; for 6 units, these licenses add up to 12 contexts. You configure an additional 20-Context license on the control unit. Therefore, the aggregated cluster license includes 32 contexts. Because the platform limit for one chassis is 100, the combined license allows a maximum of 100 contexts; the 32 contexts are within the limit. Therefore, you can configure up to 32 contexts on the control unit; each data unit will also have 32 contexts through configuration replication.
  - You have 3 Secure Firewall 3100s in the cluster. The Essentials license includes 2 contexts; for 3 units, these licenses add up to 6 contexts. You configure an additional 100-Context license on the control unit. Therefore, the aggregated cluster license includes 106 contexts. Because the platform limit for one unit is 100, the combined license allows a maximum of 100 contexts; the 106 contexts are over the limit. Therefore, you can only configure up to 100 contexts on the control unit; each data unit will also have 100 contexts through configuration replication. In this case, you should only configure the control unit Context license to be 94 contexts.
- Strong Encryption (3DES)—If your Smart Account is not authorized for strong encryption, but Cisco has determined that you are allowed to use strong encryption, you can manually add a strong encryption license to your account. Only the control unit requests this license, and all units can use it due to license aggregation.

If a new control unit is elected, the new control unit continues to use the aggregated license. It also uses the cached license configuration to re-request the control unit license. When the old control unit rejoins the cluster as a data unit, it releases the control unit license entitlement. Before the data unit releases the license, the control unit's license might be in a non-compliant state if there are no available licenses in the account. The retained license is valid for 30 days, but if it is still non-compliant after the grace period, you will not be able to make configuration changes to features requiring special licenses; operation is otherwise unaffected. The

new active unit sends an entitlement authorization renewal request every 35 seconds until the license is compliant. You should refrain from making configuration changes until the license requests are completely processed. If a unit leaves the cluster, the cached control configuration is removed, while the per-unit entitlements are retained. In particular, you would need to re-request the Context license on non-cluster units.

Permanent License Reservation

For permanent license reservation, you must purchase separate licenses for each chassis and enable the licenses *before* you configure clustering.

ASA Cluster Licenses for the Secure Firewall 4200

Smart Software Manager Regular and On-Prem

Each unit requires the Essentials license (enabled by default) and the same encryption license. We recommend licensing each unit with the licensing server *before* you enable clustering to avoid licensing mismatch issues, and when using the Strong Encryption license, issues with cluster control link encryption.

The clustering feature itself does not require any licenses. There is no extra cost for the Context license on data units.

The Strong Encryption license is automatically enabled for qualified customers when you apply the registration token. For the optional Strong Encryption (3DES/AES) feature license enabled in the ASA configuration, see below.

In the ASA license configuration, the Essentials license is always enabled by default on all units. You can only configure smart licensing on the control unit. The configuration is replicated to the data units, but for some licenses, they do not use the configuration; it remains in a cached state, and only the control unit requests the license. The licenses are aggregated into a single cluster license that is shared by the cluster units, and this aggregated license is also cached on the data units to be used if one of them becomes the control unit in the future. Each license type is managed as follows:

- Essentials—Each unit requests a Essentials license from the server.
- Context—Only the control unit requests the Context license from the server. The Essentials license includes 2 contexts by default and is present on all cluster members. The value from each unit's Essentials license plus the value of the Context license on the control unit are combined up to the platform limit in an aggregated cluster license. For example:
  - You have 6 Secure Firewall 4200s in the cluster. The Essentials license includes 2 contexts; for 6 units, these licenses add up to 12 contexts. You configure an additional 20-Context license on the control unit. Therefore, the aggregated cluster license includes 32 contexts. Because the platform limit for one chassis is 250, the combined license allows a maximum of 250 contexts; the 32 contexts are within the limit. Therefore, you can configure up to 32 contexts on the control unit; each data unit will also have 32 contexts through configuration replication.
  - You have 3 Secure Firewall 4200s in the cluster. The Essentials license includes 2 contexts; for 3 units, these licenses add up to 6 contexts. You configure an additional 250-Context license on the control unit. Therefore, the aggregated cluster license includes 256 contexts. Because the platform limit for one unit is 250, the combined license allows a maximum of 250 contexts; the 256 contexts are over the limit. Therefore, you can only configure up to 250 contexts on the control unit; each data unit will also have 250 contexts through configuration replication. In this case, you should only configure the control unit Context license to be 244 contexts.
- Strong Encryption (3DES)—If your Smart Account is not authorized for strong encryption, but Cisco has determined that you are allowed to use strong encryption, you can manually add a strong encryption

license to your account. Only the control unit requests this license, and all units can use it due to license aggregation.

If a new control unit is elected, the new control unit continues to use the aggregated license. It also uses the cached license configuration to re-request the control unit license. When the old control unit rejoins the cluster as a data unit, it releases the control unit license entitlement. Before the data unit releases the license, the control unit's license might be in a non-compliant state if there are no available licenses in the account. The retained license is valid for 30 days, but if it is still non-compliant after the grace period, you will not be able to make configuration changes to features requiring special licenses; operation is otherwise unaffected. The new active unit sends an entitlement authorization renewal request every 35 seconds until the license is compliant. You should refrain from making configuration changes until the license requests are completely processed. If a unit leaves the cluster, the cached control configuration is removed, while the per-unit entitlements are retained. In particular, you would need to re-request the Context license on non-cluster units.

Permanent License Reservation

For permanent license reservation, you must purchase separate licenses for each chassis and enable the licenses *before* you configure clustering.

ASA Cluster Licenses for the ASAv

Smart Software Manager Regular and On-Prem

Each unit requires the same Throughput license and the same encryption license. We recommend licensing each unit with the licensing server *before* you enable clustering to avoid licensing mismatch issues, and when using the Strong Encryption license, issues with cluster control link encryption.

The clustering feature itself does not require any licenses.

The Strong Encryption license is automatically enabled for qualified customers when you apply the registration token. For the optional Strong Encryption (3DES/AES) feature license enabled in the ASA configuration, see below.

In the ASA license configuration, you can only configure smart licensing on the control unit. The configuration is replicated to the data units, but for some licenses, they do not use the configuration; it remains in a cached state, and only the control unit requests the license. The licenses are aggregated into a single cluster license that is shared by the cluster units, and this aggregated license is also cached on the data units to be used if one of them becomes the control unit in the future. Each license type is managed as follows:

- Essentials—Only the control unit requests the Essentials license from the server, and all units can use it due to license aggregation.
- Throughput—Each unit requests its own Throughput license from the server.
- Strong Encryption (3DES)—If your Smart Account is not authorized for strong encryption, but Cisco has determined that you are allowed to use strong encryption, you can manually add a strong encryption license to your account. Only the control unit requests this license, and all units can use it due to license aggregation.

Permanent License Reservation

For permanent license reservation, you must purchase separate licenses for each unit and enable the licenses *before* you configure clustering.

ASA Cluster Licenses for the Firepower 4100/9300

Smart Software Manager Regular and On-Prem

The clustering feature itself does not require any licenses. To use Strong Encryption and other optional licenses, each Firepower 4100/9300 chassis must be registered with the License Authority or Smart Software Manager Regular and On-Prem server. There is no extra cost for data units.

The Strong Encryption license is automatically enabled for qualified customers when you apply the registration token. When using the token, each chassis must have the same encryption license. For the optional Strong Encryption (3DES/AES) feature license enabled in the ASA configuration, see below.

In the ASA license configuration, you can only configure smart licensing on the control unit. The configuration is replicated to the data units, but for some licenses, they do not use the configuration; it remains in a cached state, and only the control unit requests the license. The licenses are aggregated into a single cluster license that is shared by the cluster units, and this aggregated license is also cached on the data units to be used if one of them becomes the control unit in the future. Each license type is managed as follows:

- Essentials—Only the control unit requests the Essentials license from the server, and both units can use it due to license aggregation.
- Context—Only the control unit requests the Context license from the server. The Essentials license includes 10 contexts by default and is present on all cluster members. The value from each unit's Essentials license plus the value of the Context license on the control unit are combined up to the platform limit in an aggregated cluster license. For example:
  - You have 6 Firepower 9300 modules in the cluster. The Essentials license includes 10 contexts; for 6 units, these licenses add up to 60 contexts. You configure an additional 20-Context license on the control unit. Therefore, the aggregated cluster license includes 80 contexts. Because the platform limit for one module is 250, the combined license allows a maximum of 250 contexts; the 80 contexts are within the limit. Therefore, you can configure up to 80 contexts on the control unit; each data unit will also have 80 contexts through configuration replication.
  - You have 3 Firepower 4112 units in the cluster. The Essentials license includes 10 contexts; for 3 units, these licenses add up to 30 contexts. You configure an additional 250-Context license on the control unit. Therefore, the aggregated cluster license includes 280 contexts. Because the platform limit for one unit is 250, the combined license allows a maximum of 250 contexts; the 280 contexts are over the limit. Therefore, you can only configure up to 250 contexts on the control unit; each data unit will also have 250 contexts through configuration replication. In this case, you should only configure the control unit Context license to be 220 contexts.
- Carrier—Required for Distributed S2S VPN. This license is a per-unit entitlement, and each unit requests its own license from the server.
- Strong Encryption (3DES)—For pre-2.3.0 Cisco Smart Software Manager On-Prem deployment; or if your Smart Account is not authorized for strong encryption, but Cisco has determined that you are allowed to use strong encryption, you can manually add a strong encryption license to your account. This license is a per-unit entitlement, and each unit requests its own license from the server.

If a new control unit is elected, the new control unit continues to use the aggregated license. It also uses the cached license configuration to re-request the control unit license. When the old control unit rejoins the cluster as a data unit, it releases the control unit license entitlement. Before the data unit releases the license, the control unit's license might be in a non-compliant state if there are no available licenses in the account. The retained license is valid for 30 days, but if it is still non-compliant after the grace period, you will not be able to make configuration changes to features requiring special licenses; operation is otherwise unaffected. The

new active unit sends an entitlement authorization renewal request every 12 hours until the license is compliant. You should refrain from making configuration changes until the license requests are completely processed. If a unit leaves the cluster, the cached control configuration is removed, while the per-unit entitlements are retained. In particular, you would need to re-request the Context license on non-cluster units.

Permanent License Reservation

For permanent license reservation, you must purchase separate licenses for each chassis and enable the licenses *before* you configure clustering.

Prerequisites for Smart Software Licensing

Smart Software Manager Regular and On-Prem Prerequisites

Firepower 4100/9300

Configure the Smart Software Licensing infrastructure on the Firepower 4100/9300 chassis before you configure the ASA licensing entitlements.

All Other Models

- Ensure internet access, or HTTP proxy access, or Smart Software Manager On-Prem server access from the device.
- Configure a DNS server so the device can resolve the name of the Smart Software Manager.
- Set the clock for the device.
- Create an account on the Cisco Smart Software Manager:

  https://software.cisco.com/#module/SmartLicensing

  If you do not yet have an account, click the link to set up a new account. The Smart Software Manager lets you create an account for your organization.

Permanent License Reservation Prerequisites

- Create a master account on the Cisco Smart Software Manager:

  https://software.cisco.com/#module/SmartLicensing

  If you do not yet have an account, click the link to set up a new account. The Smart Software Manager lets you create a master account for your organization. Even though the ASA does need internet connectivity to the Smart Licensing server for permanent license reservation, the Smart Software Manager is used to manage your permanent licenses.
- Obtain support for permanent license reservation from the licensing team. You must provide a justification for using permanent license reservation. If your account is not approved, then you cannot purchase and apply permanent licenses.
- Purchase special permanent licenses (see License PIDs Per Model, on page 184). If you do not have the correct license in your account, then when you try to reserve a license on the ASA, you will see an error

message similar to: "The licenses cannot be reserved because the Virtual Account does not contain a sufficient surplus of the following perpetual licenses: 1 - Firepower 4100 ASA PERM UNIV(perpetual)."

- The permanent license includes all available features, including the Strong Encryption (3DES/AES) license if your account qualifies. The Secure Client capabilities are also enabled to the platform maximum, contingent on your purchase of an Secure Client license that enables the right to use Secure Client (see Secure Client Advantage, Secure Client Premier, and Secure Client VPN Only Licenses, on page 121).
- ASA Virtual: Permanent license reservation is not supported for the Azure hypervisor.

Guidelines for Smart Software Licensing

- Only Smart Software Licensing is supported. For older software on the ASA virtual, if you upgrade an existing PAK-licensed ASA virtual, then the previously installed activation key will be ignored, but retained on the device. If you downgrade the ASA virtual, the activation key will be reinstated.
- For permanent license reservation, you must return the license before you decommission the device. If you do not officially return the license, the license remains in a used state and cannot be reused for a new device.
- Because the Cisco Transport Gateway uses a certificate with a non-compliant country code, you cannot use HTTPS when using the ASA in conjunction with that product. You must use HTTP with Cisco Transport Gateway.

Defaults for Smart Software Licensing

Smart Transport

By default, all device models use Smart Transport for Smart Software License communication and use the following URL:

```
https://smartreceiver.cisco.com/licservice/license
```

For the Firepower 4100/9300, you must enable the Smart Software License communication at the chassis-level.

ASA Virtual

- When you deploy the ASA virtual, you set the feature tier and throughput level. Only the Essentials level is available at this time. For permanent license reservation, you do not need to set these parameters. When you enable permanent license reservation, these commands are removed from the configuration.



Note The Essentials license used to be known as the Standard license, and the CLI still uses the "standard" terminology.

Note If you are using a 32 or 64 core deployment on VMware or KVM, only the **unlimited** option is available with the **throughput level** command.

ASA Virtual: Configure Smart Software Licensing

This section describes how to configure Smart Software Licensing for ASA Virtual.

ASA Virtual: Configure Regular Smart Software Licensing

When you deploy ASA virtual, you can pre-configure the device and include a registration token so it registers with the Smart Software Manager and enables Smart Software Licensing. If you need to change your HTTP proxy server, license entitlement, or register the ASA virtual (for example, if you did not include the ID token in the Day0 configuration), perform this task.

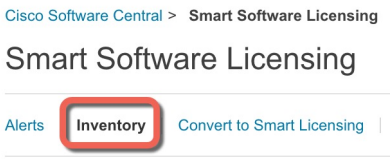
Note You may have pre-configured the HTTP proxy and license entitlements when you deployed your ASA virtual. You may also have included the registration token with your Day0 configuration when you deployed the ASA virtual; if so, you do not need to re-register using this procedure.

Procedure

Step 1 (Cisco Smart Software Manager), request and copy a registration token for the virtual account to which you want to add this device.

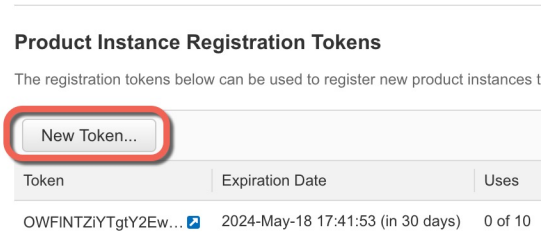
a) Click **Inventory**.

Figure 18: Inventory



b) On the **General** tab, click **New Token**.

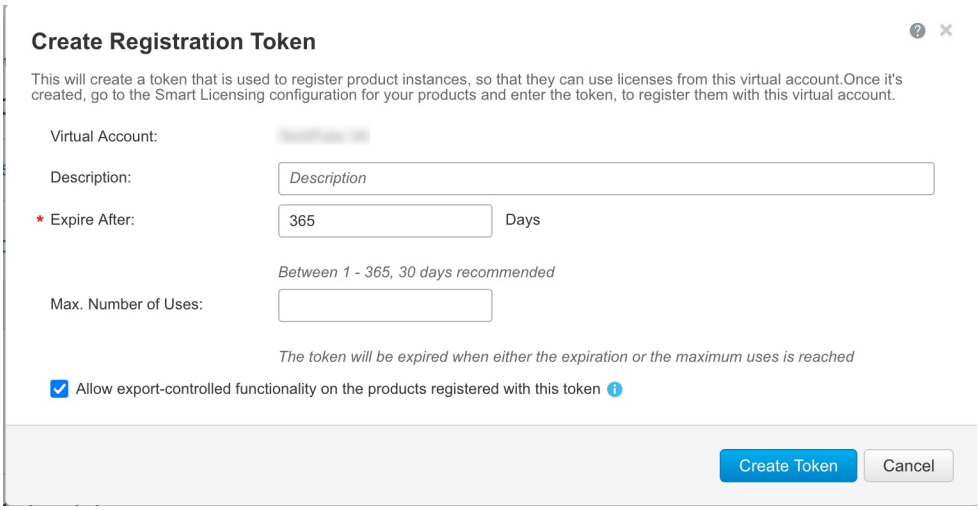
Figure 19: New Token



c) On the **Create Registration Token** dialog box enter the following settings, and then click **Create Token**:

- **Description**
- **Expire After**—Cisco recommends 30 days.
- **Max. Number of Uses**
- **Allow export-controlled functionaility on the products registered with this token**—Enables the export-compliance flag.

Figure 20: Create Registration Token



The token is added to your inventory.

d) Click the arrow icon to the right of the token to open the **Token** dialog box so you can copy the token ID to your clipboard. Keep this token ready for later in the procedure when you need to register the ASA.

Figure 21: View Token

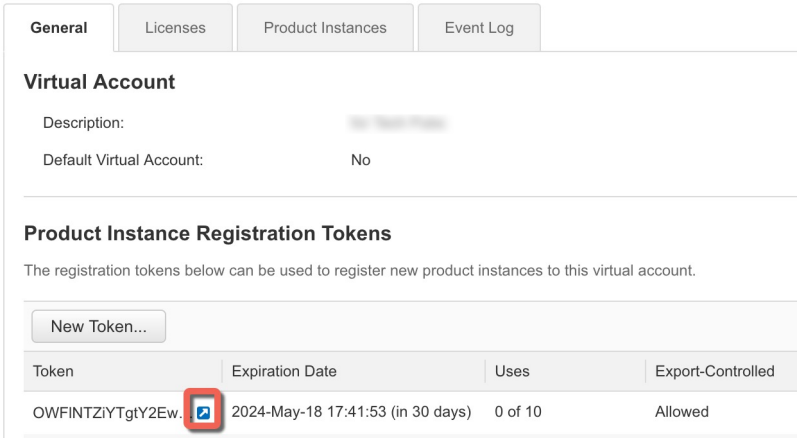
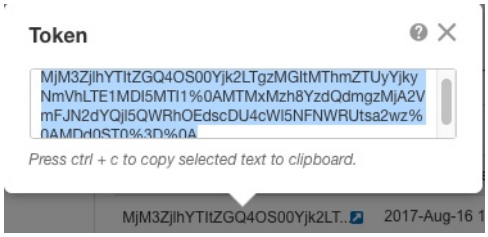


Figure 22: Copy Token



Step 2 (Optional) Use Smart Call Home instead of the default Smart Transport for communication with the Smart Licensing server.

If you know you need to use Smart Call Home instead of Smart Transport, complete these steps. Otherwise, you should use the default Smart Transport.

a) On ASA virtual, choose **Configuration** > **Device Management** > **Smart Call-Home**.
b) (Optional) Check the **Enable HTTP Proxy** check box, and enter the proxy IP address and port in the **Proxy server** and **Proxy port** fields.

For example, enter port 443 for an HTTPS server.

Note
HTTP proxy with authentication is not supported.

c) Click **Apply**.
d) Choose **Configuration** > **Device Management** > **Licensing** > **Smart Licensing**.
e) For the **Transport**, click the **Call Home** radio button.

Step 3 (Optional) Specify the HTTP Proxy for Smart Transport.

To use Smart Call Home instead of Smart Transport, see Step Step 2, on page 141.

a) For the **Transport**, click the **Smart Transport** radio button.
b) In the **Proxy URL** and **Proxy Port** fields, enter an IP address or FQDN and the proxy port number.

Step 4 Configure the license entitlements:

a) Choose **Configuration** > **Device Management** > **Licensing** > **Smart Licensing**.

b) Check the **Enable Smart license configuration** check box.
c) From the **Feature Tier** drop-down list, choose **Essentials**.

Only the Essentials tier is available, however, you need to enable it in the configuration.

d) From the **Throughput Level** drop-down list, choose one of the following: **100M**, **1G**, **2G**, **10G**, **20G**, **unlimited**.

See the following throughput/license relationship:

- 100M—ASAv5
- 1G—ASAv10
- 2G—ASAv30
- 10G—ASAv50
- 20G—ASAv100
- unlimited—ASAvU

e) (Optional) Check the **Enable strong-encryption protocol** check box.

This license is not required if you receive the strong encryption token from the Smart Software Manager. However, if your Smart Account is not authorized for strong encryption, but Cisco has determined that you are allowed to use strong encryption, you can manually add a strong encryption license to your account. Only the active unit requests this license, and both units can use it due to license aggregation.

f) Click **Apply**.

Step 5 Register the ASA virtual with the Smart Software Manager.

a) Choose **Configuration** > **Device Management** > **Licensing** > **Smart Licensing**.
b) Click **Register**.
c) Enter the registration token in the **ID Token** field.
d) (Optional) Click the **Force registration** check box to register the ASA virtual that is already registered, if it is out of sync with the Smart Software Manager.

For example, use **Force registration** if the ASA virtual was accidentally removed from the Smart Software Manager.

e) Click **Register**.

The ASA virtual attempts to register with the Smart Software Manager and request authorization for the configured license entitlements.

When you register the ASA virtual, the Smart Software Manager issues an ID certificate for communication between the ASA virtual and the Smart Software Manager. It also assigns the ASA virtual to the appropriate virtual account. Normally, this procedure is a one-time instance. However, you might need to later re-register the ASA virtual. For example, the ID certificate expires because of a communication problem.

ASA Virtual: Configure Smart Software Manager On-Prem Licensing

This procedure applies for the ASA virtual using a Smart Software Manager On-Prem.

Before you begin

- Download the Smart Software Manager On-Prem OVA file from Cisco.com and install and configure it on a VMwareESXi server. For more information, see the *Cisco Smart Software Manager On-Prem Data Sheet*.
- Smart Transport was added to the Smart Software Manager On-Prem in Version 7.0. If you are using an older version, enable Smart Call Home on the ASA virtual according to this procedure.
- Download the crypto CA trustpool before you place the device in an air-gapped network. This trustpool is normally downloaded automatically, but may be out of date in an air-gapped network:

  crypto ca trustpool import url http://www.cisco.com/security/pki/trs/ios_core.p7b

Procedure

Step 1 Request a registration token on the Smart Software Manager On-Prem.

Step 2 (Optional) Use Smart Call Home instead of the default Smart Transport for communication with the Smart Licensing server.

If you know you need to use Smart Call Home instead of Smart Transport, complete these steps. Otherwise, you should use the default Smart Transport.

a) On the ASA virtual, choose **Configuration** > **Device Management** > **Smart Call-Home**.
b) (Optional) Check the **Enable HTTP Proxy** check box, and enter the proxy IP address and port in the **Proxy server** and **Proxy port** fields.

For example, enter port 443 for an HTTPS server.

Note
HTTP proxy with authentication is not supported.

c) In the **Configure Subscription Profiles** area, edit the **License** profile to change the license server URL to go to the Smart Software Manager On-Prem.
d) In the **Deliver Subscriptions Using HTTP transport** area, select the **Subscribers** URL, and click **Edit**.
e) Change the **Subscribers** URL to the following value, and click **OK**:

https://*on-prem_ip_address*/**Transportgateway/services/DeviceRequestHandler**

f) Click **OK**.
g) Click **Apply**.
h) Choose **Configuration** > **Device Management** > **Licensing** > **Smart Licensing**.
i) For the **Transport**, click the **Call Home** radio button.

Step 3 On the ASA virtual, configure the license entitlements.

a) Choose **Configuration** > **Device Management** > **Licensing** > **Smart Licensing**.
b) Check **Enable Smart license configuration**.
c) From the **Feature Tier** drop-down menu, choose **Essentials**.

Only the **Essentials** tier is available, however, you need to enable it in the configuration.

d) To determine the license requested from the Smart Software Manager, from the **Throughput Level** drop-down list, choose one of the following: **100M**, **1G**, **2G**, **10G**, **20G**, **unlimited**.

See the following throughput/license relationship:

- 100M—ASAv5
- 1G—ASAv10
- 2G—ASAv30
- 10G—ASAv50
- 20G—ASAv100
- unlimited—ASAvU

e) (Optional) Check the **Enable strong-encryption protocol** check box.

This license is not required if you receive the strong encryption token from the Smart Software Manager. However, if your Smart Account is not authorized for strong encryption, but Cisco has determined that you are allowed to use strong encryption, you can manually add a strong encryption license to your account. Only the active unit requests this license, and both units can use it due to license aggregation.

f) Click **Apply**.

Step 4 Change the license server URL to go to the Smart Software Manager On-Prem and optionally specify the HTTP Proxy.

To use Smart Call Home instead of Smart Transport, see Step 2, on page 143.

a) For the **Transport**, click the **Smart Transport** radio button.
b) Click the **Custom URL** radio button to set the Smart Software Manager On-Prem URL in the following format:

https://*on-prem_ip_address***/SmartTransport**

c) (Optional) Configure the **Proxy URL** IP address or FQDN and **Proxy Port** number.

Step 5 Register the ASA with the Smart Software Manager.

a) Choose **Configuration** > **Device Management** > **Licensing** > **Smart Licensing**.
b) Click **Register**.
c) Enter the registration token in the **ID Token** field.
d) (Optional) Check the **Force registration** check box to register an ASA that is already registered, if it is out of sync with the Smart Software Manager.

For example, use **Force registration** if the ASA was accidentally removed from the Smart Software Manager.

e) Click **Register**.

The ASA registers with the Smart Software Manager and requests authorization for the configured license entitlements. The Smart Software Manager also applies the Strong Encryption (3DES/AES) license if your account allows. Choose **Monitoring > Properties > Smart License** to check the license status.

When you register the ASA virtual, the Smart Software Manager issues an ID certificate for communication between the ASA virtual and the Smart Software Manager. It also assigns the ASA virtual to the appropriate virtual account. Normally, this procedure is a one-time instance. However, you might need to later re-register the ASA virtual if the ID certificate expires because of a communication problem, for example.

ASA Virtual: Configure Utility (MSLA) Smart Software Licensing

Utility Licensing for a Managed Service License Agreement (MSLA) lets you pay for the amount of time a license is in use rather than paying a one time charge for a license subscription or a perpetual license. In Utility Licensing mode, the ASA virtual keeps track of license usage in units of time (15-minute intervals). The ASA virtual sends license usage reports (known as RUM reports) to the Smart Software Manager every four hours. The usage reports are then forwarded to a billing server. With Utility Licensing, Smart Call Home is not used as the transport for licensing messages. Instead the messages are sent directly via HTTP/HTTPS using *Smart Transport.*

Before you begin

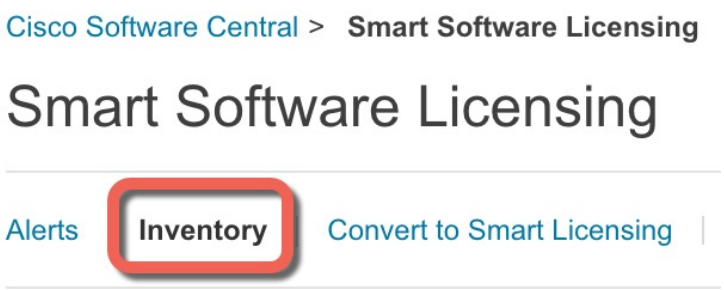
If you are using the Smart Software Manager On-Prem, download the Smart Software Manager On-Prem OVA file from Cisco.com and install and configure it on a VMware ESXi server. For more information, see the *Cisco Smart Software Manager On-Prem Data Sheet*.

Procedure

Step 1 In the Smart Software Manager (Cisco Smart Software Manager), request and copy a registration token for the virtual account to which you want to add this device.

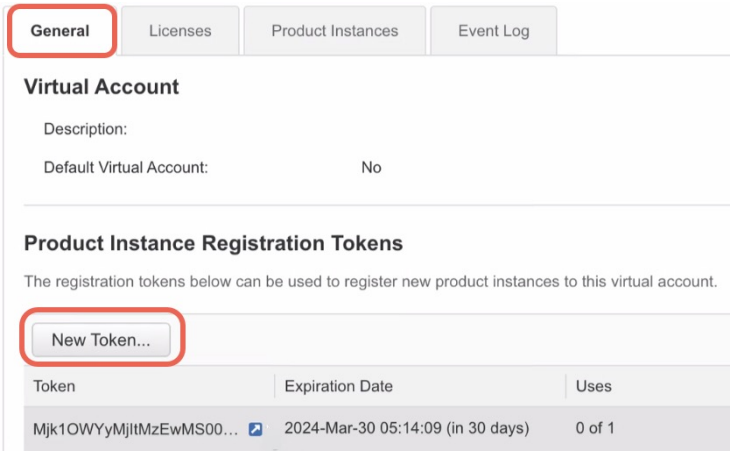
a) Click **Inventory**.

Figure 23: Inventory



b) On the **General** tab, click **New Token**.

Figure 24: New Token

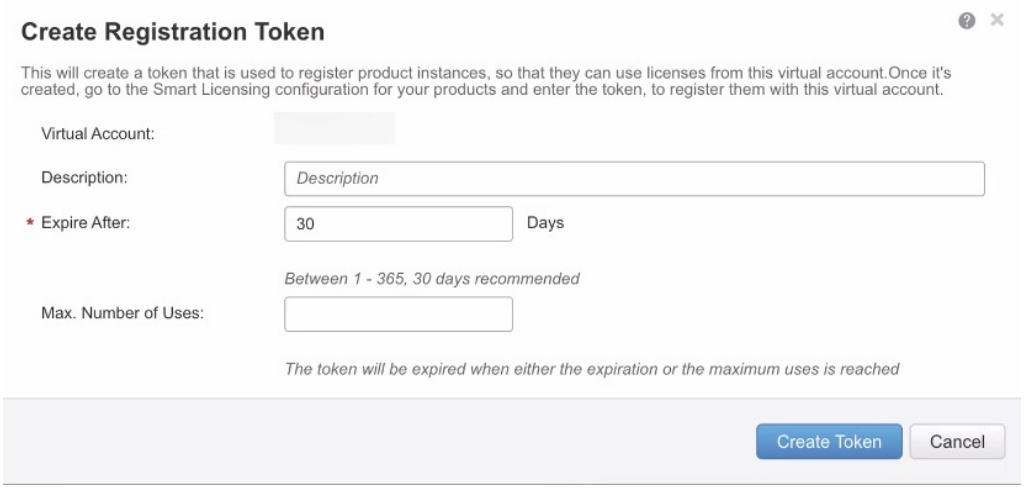


c) On the **Create Registration Token** dialog box enter the following settings, and then click **Create Token**:

- **Description**
- **Expire After**—Cisco recommends 30 days.
- **Max. Number of Uses**—The maximum number of uses of the token.

  The token expires either with the expiration date or with the maximum number of uses.

Figure 25: Create Registration Token



The token is added to your inventory.

d) Click the arrow icon to the right of the token to open the **Token** dialog box so you can copy the token ID to your clipboard. Keep this token ready for later in the procedure when you need to register the ASA.

Figure 26: View Token

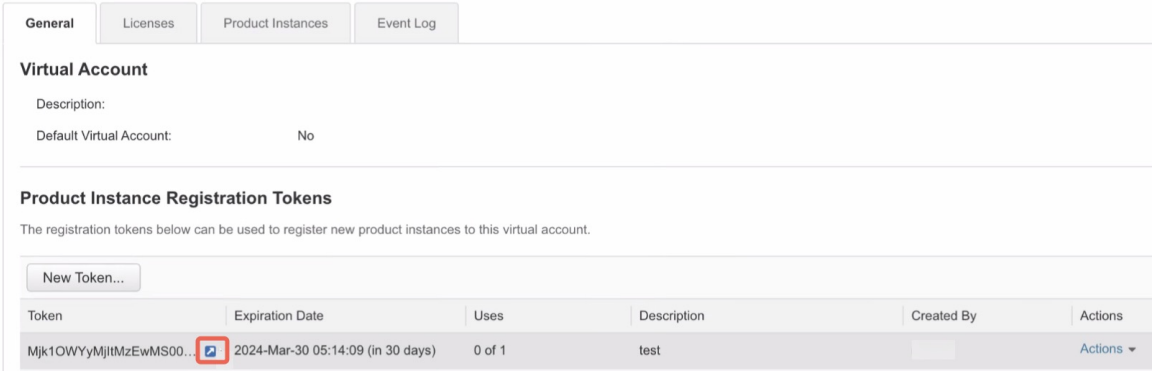
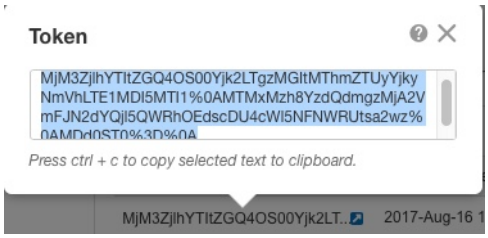


Figure 27: Copy Token



Step 2 Choose **Configuration** > **Device Management** > **Licensing** > **Smart Licensing**.

Step 3 Configure the license entitlements:

a) Check the **Enable Smart license configuration** check box.

b) From the **Feature Tier** drop-down menu, choose **Essentials**.

Only the **Essentials** tier is available, however, you need to enable it in the configuration.

c) To determine the license requested from the Smart Software Manager, from the **Throughput Level** drop-down list, choose one of the following: **100M**, **1G**, **2G**, **10G**, **20G**, **unlimited**.

See the following throughput and license relationship:

- 100M—ASAv5
- 1G—ASAv10
- 2G—ASAv30
- 10G—ASAv50
- 20G—ASAv100
- unlimited—ASAvU

d) (Optional) Check the **Enable strong-encryption protocol** check box.

This license is not required if you receive the strong encryption token from the Smart Software Manager. However, if your Smart Account is not authorized for strong encryption, but Cisco has determined that you are allowed to use strong encryption, you can manually add a strong encryption license to your account. Only the active unit requests this license, and both units can use it due to license aggregation.

Step 4 (Optional) If you want to suppress the licensing device's hostname or Smart Agent version number in the licensing messages, check the following check boxes:

a) **Host Name**.
b) **Version**.

Step 5 Click the **Smart Transport** radio button.

Step 6 Configure the URLs for Smart Transport.

a) Click the **Custom URL** radio button and enter the URL in the **URL** field.
b) In the **Registration** field, paste the Smart Software Manager Regular or On-Prem registration token.
c) In the **Utility** field, specify the URL of the Smart Software Manager Regular or On-Prem.
d) (Optional) In the **Proxy URL** field, specify the URL of the proxy if the licensing server or satellite is only reachable via a proxy.

Note
HTTP proxy with authentication is not supported.

e) (Optional) In the **Proxy Port** field, specify the proxy port number.

Step 7 Check the **Enable Standard Utility Mode** check box.

Step 8 Configure the utility licensing information, which includes customer information necessary for billing purposes.

a) In the **Custom ID** field, specify a unique customer identifier. This identifier is included in Utility Licensing usage report messages.
b) Complete the customer profile by entering the appropriate information in the remaining fields, including **Customer Company Identifier**, **Customer Company Name**, **Customer Street**, **Customer City**, **Customer State**, **Customer Country**, and **Customer Postal Code**.

Step 9 Click **Apply**.

Step 10 Click **Register** to register the ASA virtual with the Smart Software Manager Regular or On-Prem.

The ASA registers with the Smart Software Manager and requests authorization for the configured license entitlements. Choose **Monitoring > Properties > Smart License** to check the license status.

ASA Virtual: Configure Permanent License Reservation

You can assign a permanent license to the ASA virtual. This section also describes how to return a license if you retire the ASA virtual or change model tiers and need a new license.

Procedure

Step 1 Install the ASA Virtual Permanent License, on page 149

Step 2 (Optional) Return the ASA Virtual Permanent License, on page 153

Install the ASA Virtual Permanent License

For an ASA virtual that does not have internet access, you can request a permanent license from the Smart Software Manager. For more information about ASA virtual permanent license reservation, see ASA Virtual Permanent License Reservation, on page 115.

Note

- For permanent license reservation, you must return the license before you decommission the ASA virtual. If you do not officially return the license, the license remains in a used state and cannot be reused for a new ASA virtual. See (Optional) Return the ASA Virtual Permanent License, on page 153.
- If you clear your configuration after you install the permanent license, for example, by using the **write erase** command, you only have to re-enable permanent license reservation using the **license smart reservation** command without any arguments, you do not need to complete the rest of this procedure.

Before you begin

- Purchase permanent licenses so that they are available in the Smart Software Manager. Not all accounts are approved for permanent license reservation. Make sure that you have approval from Cisco for this feature before you attempt to configure it.
- Request a permanent license after the ASA virtual starts up; you cannot install a permanent license as part of the Day 0 configuration.

Procedure

Step 1 (ASAv5 only, nonflexible licensing) Run this command to allow the use of the ASAv5 permanent license when DRAM is 2 GB (the minimum that is required in 9.13 and later):

license smart set_plr5

Step 2 (Optional) Run this command to manually configure permanent license reservation irrespective of the RAM and vCPU (flexible permanent license reservation mode):

license smart flex-model {**asav5_plr** | **asav10_plr** | **asav30_plr** | **asav50_plr** | **asav100_plr** | **asavu_plr**}

Note
If the ASA virtual has more than 16 vCPUs, you can configure only the ASAvU (Unlimited) license.

To disable flexible permanent license reservation, use the **no** form of this command after you disable the permanent license reservation mode using the **no license smart reservation** command.

In the flexible permanent license reservation mode, you can change the entitlement any number of times before enabling the license reservation. To change the entitlement after you enable license reservation, you must disable permanent license reservation, enable it with the new entitlement, and then enable the license reservation.

For example, let us assume that an ASA virtual has ASAv30 entitlement and you have enabled permanent license reservation:

```
asav1# license smart flex-model asav30_plr
asav1# license smart reservation
```

Now, if you change the entitlement to ASAv50, you will get an error:

```
asav1# license smart flex-model asav50_plr
ERROR: Flexible PLR CLI is disabled when license reservation is enabled
```

To change the entitlement to ASAv50, you must run these commands:

```
asav1# no license smart reservation
asav1# license smart flex-model asav50_plr
asav1# license smart reservation
```

Step 3 From the ASA virtual CLI, enable permanent license reservation:

license smart reservation

Example:

```
ciscoasa (config)# license smart reservation
ciscoasa (config)#
```

To use regular smart licensing, use the **no** form of this command.

Step 4 Request the license code to enter in the Smart Software Manager:

license smart reservation request universal

Example:

```
ciscoasa# license smart reservation request universal
Enter this request code in the Cisco Smart Software Manager portal:
ABP:ASAv,S:9AU5ET6UQHD{A8ug5/1jRDaSp3w8uGlfeQ{53C13E
ciscoasa#
```

For flexible licensing: You can change the memory and vCPUs assigned to the ASA virtual without changing the model license.

For nonflexible licensing: If you change the memory or vCPUs, you will have to return the current license and request a new license at the new model level. The license requested must match the installed memory and vCPUs. For more information about the vCPU and memory-to-license matrix for ASA virtual permanent license reservation, see ASA Virtual Permanent License Reservation, on page 115.

To change the model of an already deployed ASA virtual, from the hypervisor you can change the vCPUs and DRAM settings to match the new model requirements; see the ASA virtual quick start guide for these values. To view your current model, use the **show vm** command.

If you re-enter this command, then the same code is displayed, even after a reload. If you have not yet entered this code into the Smart Software Manager and want to cancel the request, enter:

license smart reservation cancel

If you disable permanent license reservation, then any pending requests are canceled. If you already entered the code into the Smart Software Manager, then you must complete this procedure to apply the license to the ASA virtual, after which point you can return the license if desired. See (Optional) Return the ASA Virtual Permanent License, on page 153.

Step 5 Go to the Smart Software Manager Inventory screen, and click the **Licenses** tab:

https://software.cisco.com/#SmartLicensing-Inventory

The **Licenses** tab displays all existing licenses related to your account, both regular and permanent.

Step 6 Click **License Reservation**, and enter the ASA virtual code in the **Reservation Request Code** field.

Step 7 Click **Next**.

Verify the permanent license reservation entitlement that is displayed under **Licenses to Reserve**. It must be the entitlement that you configured in Step 2.

Step 8 Click **Next**.

Step 9 Review the information and click **Generate Authorization Code**.

The Smart Software Manager generates an authorization code. Download the code or copy it to the clipboard. The license is now in use according to the Smart Software Manager.

Step 10 Click **Close**.

Step 11 From the ASA virtual CLI, run the following command and enter the authorization code:

license smart reservation install *code*

Example:

```
ciscoasa# license smart reservation install AAu3431rGRS00Ig5HQl2vpzg{MEYCIQCBw$
INFO: ASAv platform license state is Licensed.
ciscoasa#
```

Your ASA virtual is now fully licensed.

Step 12 Run this command to save the running configuration to the startup configuration to avoid any loss of configuration during reboot.

write memory

Step 13 Verify the license status and limits of the entitlement:

- **show license all**

```
asav1# show license all

Smart Licensing Status
======================

Smart Licensing is ENABLED
License Reservation is ENABLED

Registration:
  Status: REGISTERED - UNIVERSAL LICENSE RESERVATION
  Export-Controlled Functionality: ALLOWED
  Initial Registration: SUCCEEDED on Nov 12 2024 06:40:36 UTC

License Authorization:
  Status: AUTHORIZED - RESERVED on Nov 12 2024 06:27:23 UTC

Export Authorization Key:
  Features Authorized:
    <none>

Utility:
  Status: DISABLED

Data Privacy:
  Sending Hostname: yes
    Callhome hostname privacy: DISABLED
    Smart Licensing hostname privacy: DISABLED
```

```
  Version privacy: DISABLED

Transport:
  Type: Smart
  URL: https://smartreceiver.cisco.com/licservice/license
  Proxy:
    Not Supported
  VRF:
    Not Supported

Miscellaneous:
  Custom Id: <empty>

License Usage
=============

No licenses in use

Product Information
===================
UDI: PID:ASAv,SN:9A9PJLLB849

Agent Version
=============
Smart Agent for Licensing: 5.7.25_asav/5

Reservation Info
================
License reservation: ENABLED

Overall status:
  Active: PID:ASAv,SN:9A9PJLLB849
      Reservation status: UNIVERSAL INSTALLED on Nov 12 2024 06:27:23 UTC
```

- **show license features**

```
asav1# show license features
Serial Number:  9AEJMLUP4D4
Export Compliant: YES

License mode: Smart Licensing
License reservation: Enabled
ASAv Platform License State: Licensed
Active entitlement: ASAv50_PLR, enforce mode: Authorized
Firewall throughput limited to 10 Gbps
Licensed features for this platform:
Maximum VLANs                       : 50
Inside Hosts                        : Unlimited
Failover                            : Active/Standby
Encryption-DES                      : Enabled
Encryption-3DES-AES                 : Enabled
Security Contexts                   : 0
Carrier                             : Enabled
AnyConnect Premium Peers            : 250
AnyConnect Essentials               : Disabled
Other VPN Peers                     : 250
Total VPN Peers                     : 250
AnyConnect for Mobile               : Enabled
AnyConnect for Cisco VPN Phone      : Enabled
Advanced Endpoint Assessment        : Enabled
Shared License                      : Disabled
Total TLS Proxy Sessions            : 500
```

```
Botnet Traffic Filter             : Enabled
Cluster                           : Enabled
```

(Optional) Return the ASA Virtual Permanent License

If you no longer need a permanent license, for example, you are retiring the ASA virtual or changing its model level, because of which it needs a new license, you must officially return the license to the Smart Software Manager using this procedure. If you do not follow all the steps, the license stays in a used state and cannot be freed up for use elsewhere.

Procedure

Step 1 From the ASA virtual CLI, generate a return code:

license smart reservation return

Example:

```
ciscoasa(config)# license smart reservation return
   WARNING: ASAv platform license state is Limited functionality.

   Enter this return code in Cisco Smart Software Manager portal:
   UDI: PID:ASAv,SN:9A6ST9AXQ24
       Return code: DrT1pZ-8cRT2N-xxxxx-64ABCD-xXYZeM-qD1
```

Step 2 Copy the Universal Device Identifier (UDI) so that you can find this ASA virtual instance in the Smart Software Manager.

Step 3 Go to the Smart Software Manager Inventory screen, and click the **Product Instances** tab:

https://software.cisco.com/#SmartLicensing-Inventory

The **Product Instances** tab displays all licensed products by the UDI.

Step 4 Find the ASA virtual you want to unlicense, choose **Actions** > **Remove**, and enter the ASA virtual return code in the **Reservation Return Code** field.

The permanent license is returned to the available pool.

Step 5 Click **Remove Reservation**.

Step 6 To disable permanent license reservation, do the following:

a) Run the following command to disable permanent license reservation:

no license smart reservation

b) Run the following command to disable flexible permanent license reservation:

no license smart flex-model {**asav5_plr** | **asav10_plr** | **asav30_plr** | **asav_50_plr** | **asav100_plr** | **asavu_plr**}

c) Run the following command to confirm that the license is deregistered:

show license features

The command output shows `No active entitlement`.

Step 7 Run this command to save the running configuration to the startup configuration to avoid any loss of configuration during reboot.

write memory

Switch from Nonflexible PLR Mode to Flexible PLR Mode

Procedure

Step 1 Deregister the nonflexible permanent license reservation license from Smart Account.

a) From the ASA virtual CLI, generate a return code:

license smart reservation return

Example:

```
ciscoasa(config)# license smart reservation return
   WARNING: ASAv platform license state is Limited functionality.

   Enter this return code in Cisco Smart Software Manager portal:
   UDI: PID:ASAv,SN:9A6ST9AXQ24
       Return code: DrT1pZ-8cRT2N-xxxxx-64ABCD-xXYZeM-qD1
```

b) Copy the Universal Device Identifier (UDI) so that you can find this ASA virtual instance in the Smart Software Manager.

c) Go to the Smart Software Manager Inventory screen, and click the **Product Instances** tab:

https://software.cisco.com/#SmartLicensing-Inventory

The **Product Instances** tab displays all licensed products by the UDI.

d) Find the ASA virtual you want to unlicense using the UDI.

e) Choose **Actions** > **Remove**, and enter the ASA virtual return code in the **Reservation Return Code** field.

f) Click **Remove Reservation**.

The permanent license is returned to the available pool.

Step 2 Disable permanent license reservation license from ASA virtual.

a) Run the following command to disable permanent license reservation:

no license smart reservation

Step 3 Register the ASA virtual with flexible permanent license reservation license. For more information, see Install the ASA Virtual Permanent License, on page 149.

Switch from Flexible PLR Mode to Nonflexible PLR Mode

Procedure

Step 1 Deregister the flexible permanent license reservation license from Smart Account.

a) From the ASA virtual CLI, generate a return code:

license smart reservation return

Example:

```
ciscoasa(config)# license smart reservation return
   WARNING: ASAv platform license state is Limited functionality.

   Enter this return code in Cisco Smart Software Manager portal:
   UDI: PID:ASAv,SN:9A6ST9AXQ24
       Return code: DrT1pZ-8cRT2N-xxxxx-64ABCD-xXYZeM-qD1
```

b) Copy the Universal Device Identifier (UDI) so that you can find this ASA virtual instance in the Smart Software Manager.

c) Go to the Smart Software Manager Inventory screen, and click the **Product Instances** tab:

https://software.cisco.com/#SmartLicensing-Inventory

The **Product Instances** tab displays all licensed products by the UDI.

d) Find the ASA virtual you want to unlicense using the UDI.

e) Choose **Actions** > **Remove**, and enter the ASA virtual return code in the **Reservation Return Code** field.

f) Click **Remove Reservation**.

The permanent license is returned to the available pool.

Step 2 Disable permanent license reservation license from ASA virtual.

a) Run the following command to disable permanent license reservation:

no license smart reservation

b) Run the following command to disable flexible permanent license reservation:

no license smart flex-model {**asav5_plr** | **asav10_plr** | **asav30_plr** | **asav_50_plr** | **asav100_plr** | **asavu_plr**}

Step 3 Register the ASA virtual with nonflexible permanent license reservation license. For more information, see Install the ASA Virtual Permanent License, on page 149.

(Optional) Deregister the ASA Virtual (Regular and On-Prem)

Deregistering the ASA virtual removes the ASA virtual from your account. All license entitlements and certificates on the ASA virtual are removed. You might want to deregister to free up a license for a new ASA virtual. Alternatively, you can remove the ASA virtual from the Smart Software Manager.

Note If you deregister the ASA virtual, then it will revert to a severely rate-limited state after you reload the ASA virtual.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Licensing** > **Smart Licensing**.

Step 2 Click **Unregister**.

The ASA virtual then reloads.

(Optional) Renew the ASA Virtual ID Certificate or License Entitlement (Regular and On-Prem)

By default, the ID certificate is automatically renewed every 6 months, and the license entitlement is renewed every 30 days. You might want to manually renew the registration for either of these items if you have a limited window for Internet access, or if you make any licensing changes in the Smart Software Manager, for example.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Licensing** > **Smart Licensing**.

Step 2 To renew the ID certificate, click **Renew ID Certificate**.

Step 3 To renew the license entitlement, click **Renew Authorization**.

1000/1200/3100/4200: Configure Smart Software Licensing

This section describes how to configure Smart Software Licensing for the1000/1200/3100/4200. Choose one of the following methods:

- 1000/1200/3100/4200: Configure Regular Smart Software Licensing, on page 157

  You can also (Optional) Deregister the 1000/1200/3100/4200 (Regular and On-Prem), on page 169 or (Optional) Renew the 1000/1200/3100/4200 ID Certificate or License Entitlement (Regular and On-Prem), on page 169.

- 1000/1200/3100/4200: Configure Smart Software Manager On-Prem Licensing, on page 161

  You can also (Optional) Deregister the 1000/1200/3100/4200 (Regular and On-Prem), on page 169 or (Optional) Renew the 1000/1200/3100/4200 ID Certificate or License Entitlement (Regular and On-Prem), on page 169.

• 1000/1200/3100/4200: Configure Permanent License Reservation, on page 165

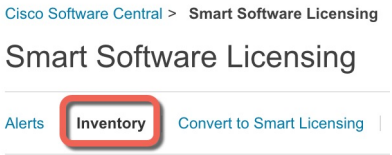
1000/1200/3100/4200: Configure Regular Smart Software Licensing

This procedure applies for an ASA using the Smart Software Manager.

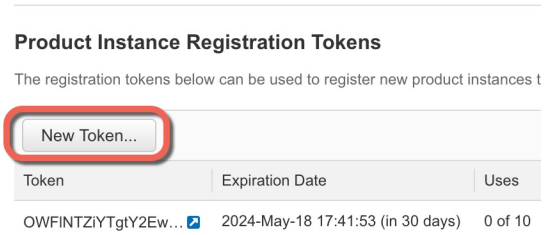
Procedure

Step 1 In the Cisco Smart Software Manager, request and copy a registration token for the virtual account to which you want to add this device.

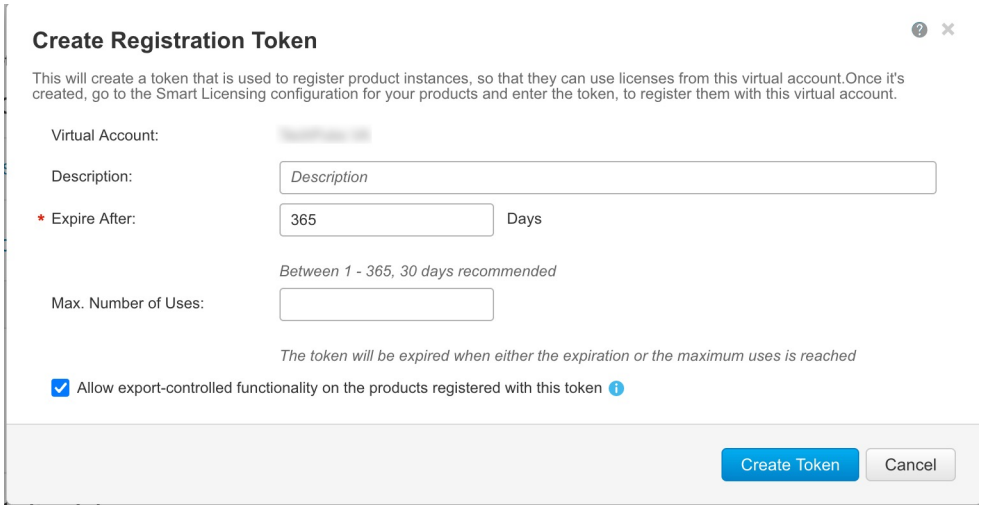
a) Click **Inventory**.



b) On the **General** tab, click **New Token**.



c) On the **Create Registration Token** dialog box enter the following settings, and then click **Create Token**:



- **Description**
- **Expire After**—Cisco recommends 30 days.

- **Max. Number of Uses**
- **Allow export-controlled functionaility on the products registered with this token**—Enables the export-compliance flag.

The token is added to your inventory.

d) Click the arrow icon to the right of the token to open the **Token** dialog box so you can copy the token ID to your clipboard. Keep this token ready for later in the procedure when you need to register the ASA.

Figure 28: View Token

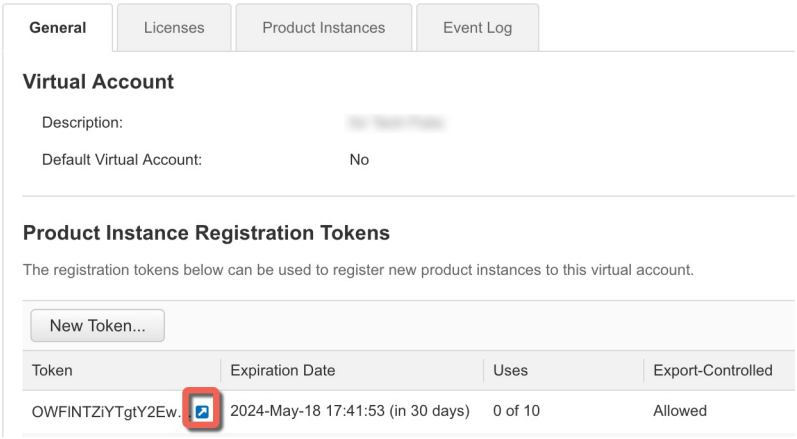
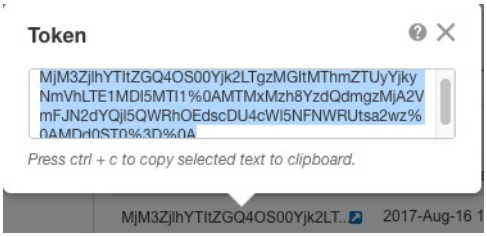


Figure 29: Copy Token



Step 2 (Optional) Use Smart Call Home instead of the default Smart Transport for communication with the Smart Licensing server.

If you know you need to use Smart Call Home instead of Smart Transport, complete these steps. Otherwise, you should use the default Smart Transport.

a) On ASA, choose **Configuration** > **Device Management** > **Smart Call-Home**.
b) (Optional) Check the **Enable HTTP Proxy** check box, and enter the proxy IP address and port in the **Proxy server** and **Proxy port** fields.

For example, enter port 443 for an HTTPS server.

Note
HTTP proxy with authentication is not supported.

c) Click **Apply**.
d) Choose **Configuration** > **Device Management** > **Licensing** > **Smart Licensing**.
e) For the **Transport**, click the **Call Home** radio button.

Step 3 (Optional) Specify the HTTP Proxy for Smart Transport.

To use Smart Call Home instead of Smart Transport, see Step Step 2, on page 158.

a) For the **Transport**, click the **Smart Transport** radio button.
b) Enter the proxy IP address or FQDN and port number in the **Proxy URL** and **Proxy Port** fields.

Step 4 Configure the license entitlements:

a) Choose **Configuration** > **Device Management** > **Licensing** > **Smart Licensing**.
b) Check the **Enable Smart license configuration** check box.
c) From the **Feature Tier** drop-down menu, choose **Essentials**.

Only the **Essentials** tier is available, however, you need to enable it in the configuration; a tier license is a prerequisite for adding other feature licenses. For Secure Firewall models, the Essentials license is always enabled and cannot be disabled.

d) (Optional) (Firepower 1010) Check the **Enable Security Plus** check box.

The Security Plus tier enables failover.

e) (Optional) For the **Context** license, enter the number of contexts.

Note
This license is not supported for the Firepower 1010.

By default, the ASA supports 2 contexts, so you should request the number of contexts you want minus the 2 default contexts. The maximum number of contexts depends on your model:

- Firepower 1120—5 contexts
- Firepower 1140—10 contexts
- Firepower 1150—25 contexts
- Secure Firewall 1210—5 contexts
- Secure Firewall 1220—10 contexts
- Secure Firewall 1230—25 contexts
- Secure Firewall 1240—25 contexts
- Secure Firewall 1250—25 contexts
- Secure Firewall 3100—100 contexts
- Secure Firewall 4200—100 contexts

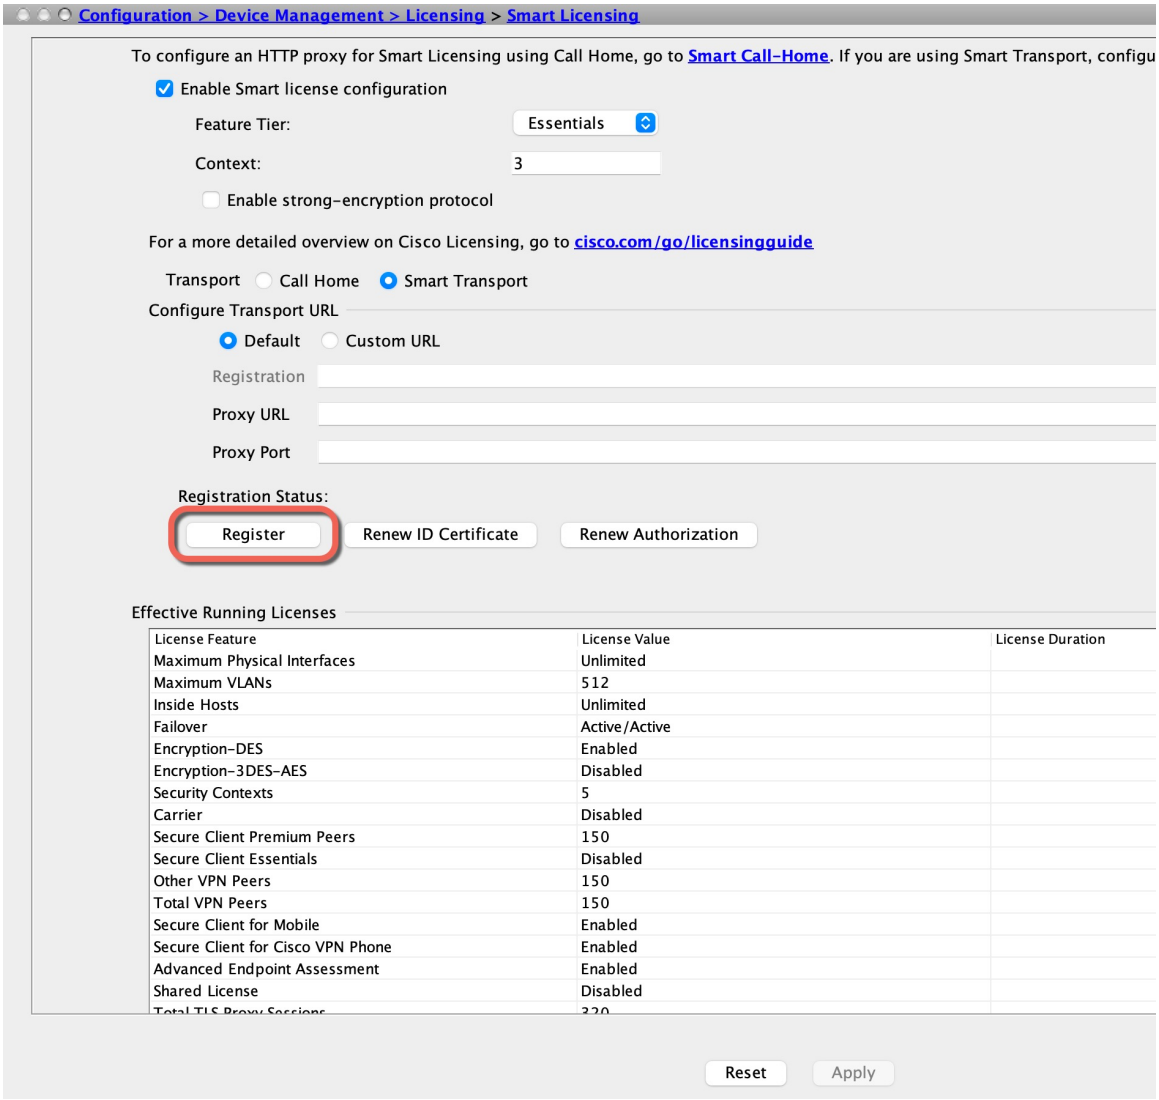
For example, to use the maximum of 25 contexts on the Firepower 1150, enter 23 for the number of contexts; this value is added to the default of 2.

f) (Optional) Check the **Enable strong-encryption protocol** check box.

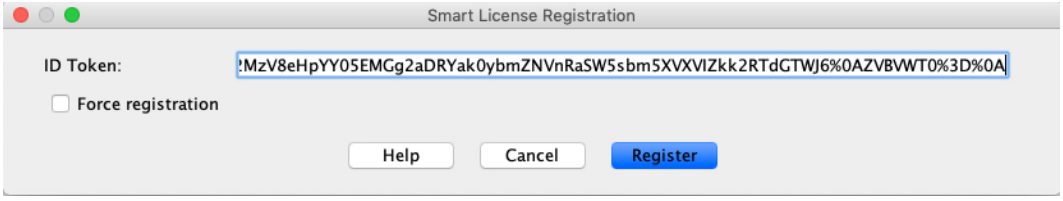
This license is not required if you receive the strong encryption token from the Smart Software Manager. However, if your Smart Account is not authorized for strong encryption, but Cisco has determined that you are allowed to use strong encryption, you can manually add a strong encryption license to your account. Only the active unit requests this license, and both units can use it due to license aggregation.

g) (Optional) (Secure Firewall 3100/4200) Check **Enable Carrier** for Diameter, GTP/GPRS, SCTP inspection.
h) Click **Apply**.

Step 5 Click **Register**.



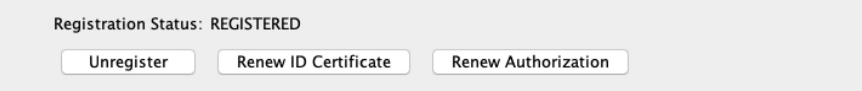
Step 6 Enter the registration token from the Cisco Smart Software Manager in the **ID Token** field.



(Optional) Click the **Force registration** check box to register an ASA that is already registered, but that might be out of sync with the Smart Software Manager. For example, use **Force registration** if the ASA was accidentally removed from the Smart Software Manager.

Step 7 Click **Register**.

ASDM refreshes the page when the license status is updated. You can also choose **Monitoring** > **Properties** > **Smart License** to check the license status, particularly if the registration fails.



Step 8 Quit ASDM and relaunch it.

When you change licenses, you need to relaunch ASDM to show updated screens.

1000/1200/3100/4200: Configure Smart Software Manager On-Prem Licensing

This procedure applies for an ASA using a Smart Software Manager On-Prem.

Before you begin

- Download the Smart Software Manager On-Prem OVA file from Cisco.com and install and configure it on a VMwareESXi server. For more information, see the *Cisco Smart Software Manager On-Prem Data Sheet*.
- Smart Transport was added to the Smart Software Manager On-Prem in Version 7.0. If you are using an older version, enable Smart Call Home on the ASA according to this procedure.
- Download the crypto CA trustpool before you place the device in an air-gapped network. This trustpool is normally downloaded automatically, but may be out of date in an air-gapped network:

  crypto ca trustpool import url http://www.cisco.com/security/pki/trs/ios_core.p7b

Procedure

Step 1 Request a registration token on the Smart Software Manager On-Prem server.

Step 2 (Optional) Use Smart Call Home instead of the default Smart Transport for communication with the Smart Licensing server.

If you know you need to use Smart Call Home instead of Smart Transport, complete these steps. Otherwise, you should use the default Smart Transport.

a) On the ASA,choose **Configuration** > **Device Management** > **Smart Call-Home**.

b) (Optional) Check the **Enable HTTP Proxy** check box, and enter the proxy IP address and port in the **Proxy server** and **Proxy port** fields.

For example, enter port 443 for an HTTPS server.

Note
HTTP proxy with authentication is not supported.

c) In the **Configure Subscription Profiles** area, edit the **License** profile to change the license server URL to go to the Smart Software Manager On-Prem.

d) In the **Deliver Subscriptions Using HTTP transport** area, select the **Subscribers** URL, and click **Edit**.
e) Change the **Subscribers** URL to the following value, and click **OK**:

https://*on-prem_ip_address***/Transportgateway/services/DeviceRequestHandler**

f) Click **OK**.
g) Click **Apply**.
h) Choose **Configuration** > **Device Management** > **Licensing** > **Smart Licensing**.
i) For the **Transport**, click the **Call Home** radio button.

Step 3 Configure the license entitlements.

a) Choose **Configuration** > **Device Management** > **Licensing** > **Smart Licensing**.
b) Check the **Enable Smart license configuration** check box.
c) From the **Feature Tier** drop-down list, choose **Essentials**.

Only the **Essentials** tier is available, however, you need to enable it in the configuration; a tier license is a prerequisite for adding other feature licenses. For Secure Firewall models, the Essentials license is always enabled and cannot be disabled.

d) (Optional) (Firepower 1010) Check **Enable Security Plus**.

The Security Plus tier enables failover.

e) (Optional) For the **Context** license, enter the number of contexts.

Note
This license is not supported for the Firepower 1010.

By default, the ASA supports 2 contexts, so you should request the number of contexts you want minus the 2 default contexts. The maximum number of contexts depends on your model:

- Firepower 1120—5 contexts
- Firepower 1140—10 contexts
- Firepower 1150—25 contexts
- Secure Firewall 1210—5 contexts
- Secure Firewall 1220—10 contexts
- Secure Firewall 1230—25 contexts
- Secure Firewall 1240—25 contexts
- Secure Firewall 1250—25 contexts
- Secure Firewall 3100—100 contexts
- Secure Firewall 4200—100 contexts

For example, to use the maximum of 25 contexts on the Firepower 1150, enter 23 for the number of contexts; this value is added to the default of 2.

f) (Optional) Check the **Enable strong-encryption protocol** check box.

This license is not required if you receive the strong encryption token from the Smart Software Manager. However, if your Smart Account is not authorized for strong encryption, but Cisco has determined that you are allowed to use strong encryption, you can manually add a strong encryption license to your account. Only the active unit requests this license, and both units can use it due to license aggregation.

g) (Optional) (Secure Firewall 3100/4200) Check **Enable Carrier** for Diameter, GTP/GPRS, SCTP inspection.
h) Click **Apply**.

Step 4 Change the license server URL to go to the Smart Software Manager On-Prem and optionally specify the HTTP Proxy.

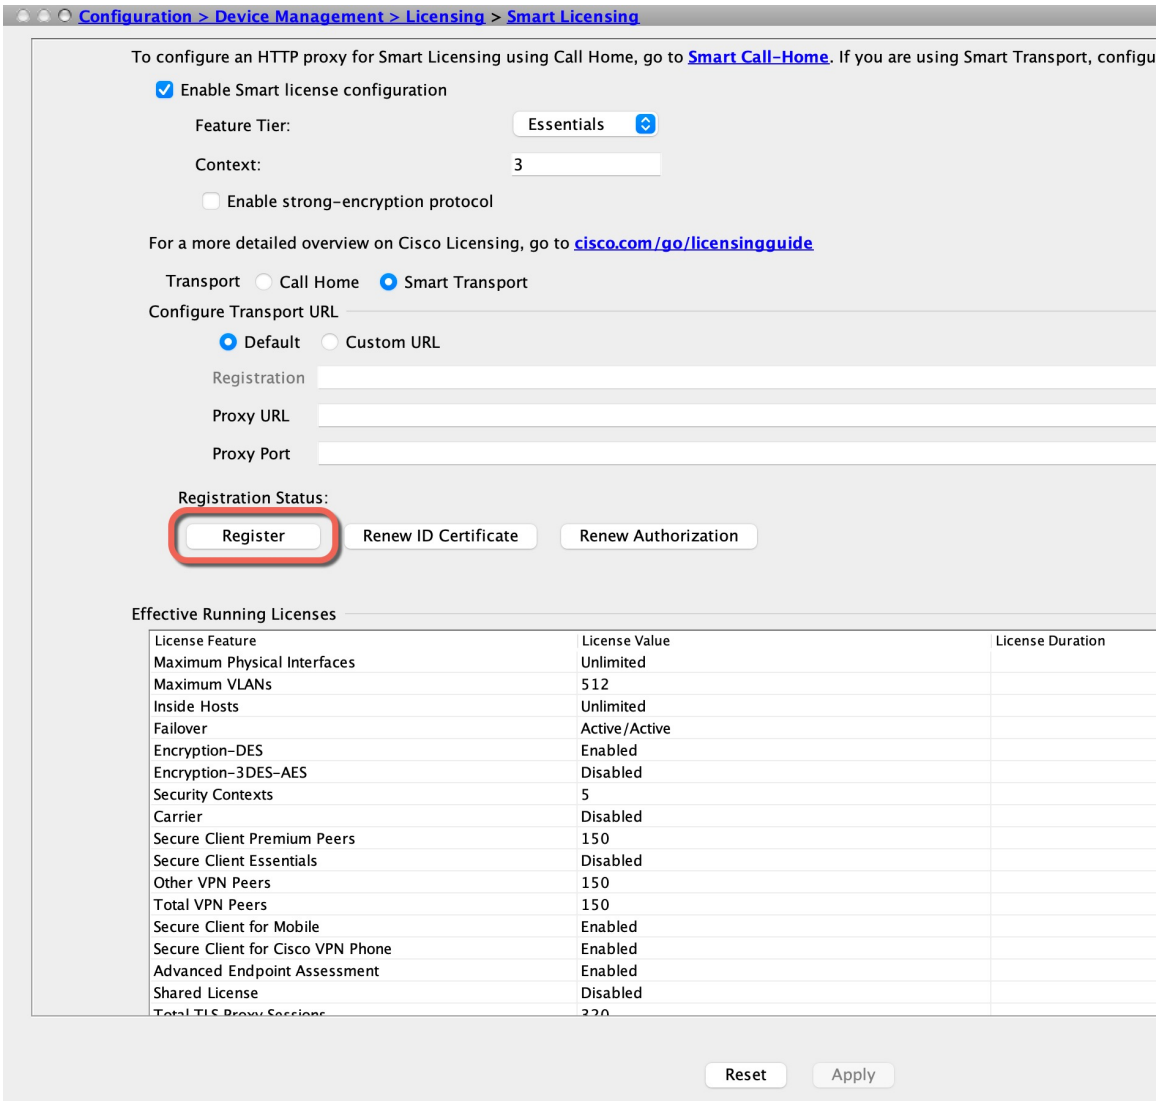
To use Smart Call Home instead of Smart Transport, see Step Step 2, on page 161.

a) For the **Transport**, click the **Smart Transport** radio button.
b) Click the **Custom URL** radio button to set the Smart Software Manager On-Prem URL in the following format:

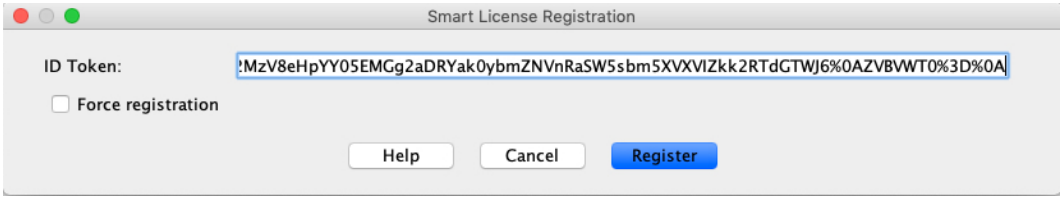
https://*on-prem_ip_address*/**SmartTransport**

c) (Optional) Configure the **Proxy URL** IP address or FQDN and **Proxy Port** number.

Step 5 Click **Register**.



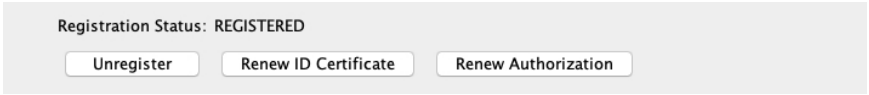
Step 6 Enter the registration token from the Cisco Smart Software Manager in the **ID Token** field.



(Optional) Click the **Force registration** check box to register an ASA that is already registered, but that might be out of sync with the Smart Software Manager. For example, use **Force registration** if the ASA was accidentally removed from the Smart Software Manager.

Step 7 Click **Register**.

ASDM refreshes the page when the license status is updated. You can also choose **Monitoring** > **Properties** > **Smart License** to check the license status, particularly if the registration fails.



Step 8 Quit ASDM and relaunch it.

When you change licenses, you need to relaunch ASDM to show updated screens.

1000/1200/3100/4200: Configure Permanent License Reservation

You can assign a permanent license to a Firepower 1000, Secure Firewall 1200/3100/4200. This section also describes how to return a license if you retire the ASA.

Procedure

Step 1 Install the 1000/1200/3100/4200 Permanent License, on page 165.

Step 2 (Optional) Return the 1000/1200/3100/4200 Permanent License, on page 168.

Install the 1000/1200/3100/4200 Permanent License

For ASAs that do not have internet access, you can request a permanent license from the Smart Software Manager. The permanent license enables all features: Essentials license with maximum Security Contexts.

Note For permanent license reservation, you must return the license before you decommission the ASA. If you do not officially return the license, the license remains in a used state and cannot be reused for a new ASA. See (Optional) Return the 1000/1200/3100/4200 Permanent License, on page 168.

Before you begin

Purchase permanent licenses so they are available in the Smart Software Manager. Not all accounts are approved for permanent license reservation. Make sure you have approval from Cisco for this feature before you attempt to configure it.

Procedure

Step 1 At the ASA CLI, enable permanent license reservation:

license smart reservation

Example:

```
ciscoasa (config)# license smart reservation
```

```
ciscoasa (config)#
```

Step 2 Request the license code to enter in the Smart Software Manager:

license smart reservation request universal

Example:

```
ciscoasa# license smart reservation request universal
Enter this request code in the Cisco Smart Software Manager portal:
BB-ZFPR-2140:JAD200802RR-AzKmHcc71-2A
ciscoasa#
```

If you re-enter this command, then the same code is displayed, even after a reload. If you have not yet entered this code into the Smart Software Manager and want to cancel the request, enter:

license smart reservation cancel

If you disable permanent license reservation, then any pending requests are canceled. If you already entered the code into the Smart Software Manager, then you must complete this procedure to apply the license to the ASA, after which point you can return the license if desired. See (Optional) Return the 1000/1200/3100/4200 Permanent License, on page 168.

Step 3 Go to the Smart Software Manager Inventory screen, and click the **Licenses** tab:

https://software.cisco.com/#SmartLicensing-Inventory

The **Licenses** tab displays all existing licenses related to your account, both regular and permanent.

Step 4 Click **License Reservation**, and type the ASA code into the box. Click **Reserve License**.

The Smart Software Manager generates an authorization code. You can download the code or copy it to the clipboard. At this point, the license is now in use according to the Smart Software Manager.

If you do not see the **License Reservation** button, then your account is not authorized for permanent license reservation. In this case, you should disable permanent license reservation and re-enter the regular smart license commands.

Step 5 On the ASA, enter the authorization code:

license smart reservation install *code*

Example:

```
ciscoasa# license smart reservation install AAu3431rGRS00Ig5HQl2vpzg{MEYCIQCBw$
ciscoasa#
```

Step 6 Request license entitlements on the ASA.

Note
Although the permanent license allows the full use of all of the licenses, you still need to turn on the entitlements in the ASA configuration so that the ASA knows it can use them.

a) Enter license smart configuration mode:

license smart

Example:

```
ciscoasa(config)# license smart
ciscoasa(config-smart-lic)#
```

b) (Firepower 1000) Set the feature tier:

feature tier standard

Only the standard (essentials) tier is available, however, you need to enable it in the configuration; a tier license is a prerequisite for adding other feature licenses.The Essentials license was formerly called the Standard license, and "standard" is still used in the CLI. For Secure Firewall models, the Essentials license is always enabled and cannot be disabled.

c) (Optional) Enable the security context license.

feature context *number*

Note
This license is not supported for the Firepower 1010.

By default, the ASA supports 2 contexts, so you should enable the number of contexts you want minus the 2 default contexts. Because the permanent license allows the maximum number, you can enable the maximum number for your model. The maximum number of contexts depends on your model:

- Firepower 1120—5 contexts
- Firepower 1140—10 contexts
- Firepower 1150—25 contexts
- Secure Firewall 1210—5 contexts
- Secure Firewall 1220—10 contexts
- Secure Firewall 1230—25 contexts
- Secure Firewall 1240—25 contexts
- Secure Firewall 1250—25 contexts
- Secure Firewall 3100—100 contexts
- Secure Firewall 4200—100 contexts

For example, to use the maximum of 25 contexts on the Firepower 1150, enter 23 for the number of contexts; this value is added to the default of 2.

Example:

```
ciscoasa(config-smart-lic)# feature context 18
```

d) (Optional) (Firepower 1010) Enable the Security Plus license to enable failover.

feature security-plus

Example:

```
ciscoasa(config-smart-lic)# feature security-plus
```

e) (Optional) (Secure Firewall 3100/4200) Enable the Carrier license for Diameter, GTP/GPRS, SCTP inspection.

feature carrier

Example:

```
ciscoasa(config-smart-lic)# feature carrier
```

f) (Optional) Enable strong encryption.

feature strong-encryption

This license is not required if you receive the strong encryption token from the Smart Software Manager. However, if your Smart Account is not authorized for strong encryption, but Cisco has determined that you are allowed to use strong encryption, you can manually add a strong encryption license to your account. Only the active unit requests this license, and both units can use it due to license aggregation.

Example:

```
ciscoasa(config-smart-lic)# feature strong-encryption
```

(Optional) Return the 1000/1200/3100/4200 Permanent License

If you no longer need a permanent license (for example, you are retiring an ASA), you must officially return the license to the Smart Software Manager using this procedure. If you do not follow all steps, then the license stays in a used state and cannot easily be freed up for use elsewhere.

Procedure

Step 1 On the ASA, generate a return code:

license smart reservation return

Example:

```
ciscoasa# license smart reservation return
Enter this return code in the Cisco Smart Software Manager portal:
Au3431rGRS00Ig5HQl2vpcg{uXiTRfVrp7M/zDpirLwYCaq8oSv60yZJuFDVBS2QliQ=
```

The ASA immediately becomes unlicensed and moves to the Evaluation state. If you need to view this code again, re-enter this command. Note that if you request a new permanent license (**license smart reservation request universal**), then you cannot re-display this code. Be sure to capture the code to complete the return. If the evaluation period has expired, then the ASA moves into an expired state. For more information about out-of-compliance states, see Out-of-Compliance State, on page 190.

Step 2 View the ASA universal device identifier (UDI) so you can find this ASA instance in the Smart Software Manager:

show license udi

Example:

```
ciscoasa# show license udi
UDI: PID:FPR-2140,SN:JAD200802RR
ciscoasa#
```

Step 3 Go to the Smart Software Manager Inventory screen, and click the **Product Instances** tab:

https://software.cisco.com/#SmartLicensing-Inventory

The **Product Instances** tab displays all licensed products by the UDI.

Step 4 Find the ASA you want to unlicense, choose **Actions** > **Remove**, and type the ASA return code into the box. Click **Remove Product Instance**.

The permanent license is returned to the available pool.

(Optional) Deregister the 1000/1200/3100/4200 (Regular and On-Prem)

Deregistering the ASA removes the ASA from your account. All license entitlements and certificates on the ASA are removed. You might want to deregister to free up a license for a new ASA. Alternatively, you can remove the ASA from the Smart Software Manager.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Licensing** > **Smart Licensing**.

Step 2 Click **Unregister**.

(Optional) Renew the 1000/1200/3100/4200 ID Certificate or License Entitlement (Regular and On-Prem)

By default, the ID certificate is automatically renewed every 6 months, and the license entitlement is renewed every 30 days. You might want to manually renew the registration for either of these items if you have a limited window for internet access, or if you make any licensing changes in the Smart Software Manager, for example.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Licensing** > **Smart Licensing**.

Step 2 To renew the ID certificate, click **Renew ID Certificate**.

Step 3 To renew the license entitlement, click **Renew Authorization**.

Firepower 4100/9300: Configure Smart Software Licensing

This procedure applies for a chassis using the Smart Software Manager, Smart Software Manager On-Prem, or for Permanent License Reservation; see the FXOS configuration guide to pre-configure licensing communication.

For Permanent License Reservation, the license enables all features: Standard tier with maximum Security Contexts and the Carrier license. However, for the ASA to "know" to use these features, you need to enable them on the ASA.

Before you begin

For an ASA cluster, you need to access the control node for configuration. Check the chassis manager to see which node is the control node.

Procedure

Step 1 In ASDM, choose **Configuration** > **Device Management** > **Licensing** > **Smart Licensing**.

Step 2 From the **Feature Tier** drop-down list, choose **Standard**.

Only the standard tier is available. A tier license is a prerequisite for adding other feature licenses. You must have sufficient tier licenses in your account. Otherwise, you cannot configure any other feature licenses or any features that require licenses.

Step 3 (Optional) Check the **Enable strong-encryption protocol** check box.

This license is not required if you receive the strong encryption token from the Smart Software Manager. However, if your Smart Account is not authorized for strong encryption, but Cisco has determined that you are allowed to use strong encryption, you can manually add a strong encryption license to your account. Only the active unit requests this license, and both units can use it due to license aggregation.

Step 4 (Optional) Check the **Carrier** check box.

Step 5 (Optional) From the **Context** drop-down list, choose the number of contexts you want.

For Permanent License Reservation, you can specify the maximum contexts (248).

Step 6 Click **Apply**.

Step 7 Quit ASDM and relaunch it.

When you change licenses, you need to relaunch ASDM to show updated screens.

Licenses Per Model

This section lists the license entitlements available for the ASAv and Firepower 4100/9300 chassis ASA security module.

ASA Virtual

When you set the throughput level in the ASA configuration or, for flexible permanent license reservation, specifying the model, it determines the license requested from the Smart Software Manager. See the following throughput level/license relationship:

- 100M—ASAv5
- 1G—ASAv10
- 2G—ASAv30
- 10G—ASAv50
- 20G—ASAv100
- Unlimited—ASAvU

The throughput level also determines the maximum Secure Client and TLS proxy sessions. However, a lower ASA virtual memory profile will cap your actual number of sessions, so to determine your sessions, you need to check both the throughput level and the memory installed.

The memory of your ASA virtual determines the maximum concurrent firewall connections and VLANs, and is not determined by the throughput level.

The following table shows the licensed features for the ASA virtual series.

Licenses	Description
Entitlement Licenses	
Throughput Level	You set the throughput level in the ASA configuration or, for flexible permanent license reservation, specifying the model. This level determines the license you need. 100M: ASAv5 1G: ASAv10 2G: ASAv30 10G: ASAv50 20G: ASAv100 Unlimited: ASAvU
Firewall Licenses	
Botnet Traffic Filter	Enabled

Licenses	Description
Firewall Connections, Concurrent	Firewall connections are determined by the ASA virtual memory. 2 GB to 7.9 GB: 100,000 8 GB to 15.9 GB: 500,000 16 GB to 31.9 GB: 2,000,000 32 GB to 64 GB: 4,000,000 ASAvU: If you are using ASAvU, which uses a minimum memory of 64 GB, the maximum number of firewall connections is 8,000,000.
Carrier	Enabled
Total TLS Proxy Sessions	TLS Proxy Sessions are determined by the throughput level and ASA virtual memory. 100M throughput + any memory: 500 1G throughput + any memory: 500 2G throughput: • 2 GB to 7.9 GB memory: 500 • 8 GB+ memory: 1000 10G throughput: • 2 GB to 7.9 GB memory: 500 • 8 GB to 15.9 GB memory: 1000 • 16 GB+ memory: 10,000 20G throughput: • 2 GB to 7.9 GB memory: 500 • 8 GB to 15.9 GB memory: 1000 • 16 GB to 31.9 GB memory: 10,000 • 32 GB+ memory: 20,000 Unlimited throughput: If you are using ASAvU, which uses a minimum memory of 64 GB, the maximum number of TLS proxy sessions is 40,000.
VPN Licenses	

Licenses	Description	
Secure Client peers	Unlicensed	Secure Client peers are determined by the throughput level and ASA virtual memory. *Optional Secure Client Advantage or Secure Client Premier license, Maximums:* *100M throughput + any memory: 50* *1G throughput + any memory: 250* *2G throughput:* • *2 GB to 7.9 GB memory: 250* • *8 GB+ memory: 750* *10G throughput:* • *2 GB to 7.9 GB memory: 250* • *8 GB to 15.9 GB memory: 750* • *16 GB+ memory: 10,000* *20G throughput:* • *2 GB to 7.9 GB memory: 250* • *8 GB to 15.9 GB memory: 750* • *16 GB to 31.9 GB: 10,000* • *32 GB+ memory: 20,000* Unlimited throughput: If you are using ASAvU, which uses a minimum memory of 64 GB, the maximum number of Secure Client peers is 40,000.

Licenses	Description
Other VPN Peers	**Note** Other VPN peers are determined by the throughput level and ASA virtual memory. 100M throughput + any memory: 50 1G throughput + any memory: 250 2G throughput: • 2 GB to 7.9 GB memory: 250 • 8 GB+ memory: 750 10G throughput: • 2 GB to 7.9 GB memory: 250 • 8 GB to 15.9 GB memory: 750 • 16 GB+ memory: 10,000 20G throughput: • 2 GB to 7.9 GB memory: 250 • 8 GB to 15.9 GB memory: 750 • 16 GB to 31.9 GB: 10,000 • 32 GB+ memory: 20,000 Unlimited throughput: If you are using ASAvU, which uses a minimum memory of 64 GB, the maximum number of other VPN peers is 40,000.

Licenses	Description
Total VPN Peers, combined all types	**Note** Total VPN peers are determined by the throughput level and ASA virtual memory. 100M throughput + any memory: 50 1G throughput + any memory: 250 2G throughput: • 2 GB to 7.9 GB memory: 250 • 8 GB+ memory: 750 10G throughput: • 2 GB to 7.9 GB memory: 250 • 8 GB to 15.9 GB memory: 750 • 16 GB+ memory: 10,000 20G throughput: • 2 GB to 7.9 GB memory: 250 • 8 GB to 15.9 GB memory: 750 • 16 GB to 31.9 GB: 10,000 • 32 GB+ memory: 20,000 Unlimited throughput: If you are using ASAvU, which uses a minimum memory of 64 GB, the maximum number of VPN peers is 40,000.
General Licenses	
Encryption	Base (DES) or Strong (3DES/AES), depending on the account's export compliance setting
Failover	Active/Standby
Security Contexts	No support
Clustering	Enabled
VLANs, Maximum	VLANs are determined by the ASA virtual memory. 2 GB to 7.9 GB—50 8 GB to 15.9 GB—200 16 GB to 31.9 GB—1024 32 GB to 64 GB—1024

Firepower 1010

The following table shows the licensed features for the Firepower 1010.

Licenses	Essentials License	
Firewall Licenses		
Botnet Traffic Filter	No Support.	
Firewall Conns, Concurrent	100,000	
Carrier	No support. Although SCTP inspection maps are not supported, SCTP stateful inspection using ACLs is supported.	
Total TLS Proxy Sessions	4,000	
VPN Licenses		
Secure Client peers	Unlicensed	*Optional Secure Client Advantage,Secure Client Premier, or Secure Client VPN Only license, Maximum: 75*
Other VPN Peers	75	
Total VPN Peers, combined all types	75	
General Licenses		
Encryption	Base (DES) or Strong (3DES/AES), depending on the account's export compliance setting	
Security Plus (failover, VPN Load Balancing)	Disabled	*Optional*
Security Contexts	No support.	
Clustering	No support.	
VLANs, Maximum	60	

Firepower 1100 Series

The following table shows the licensed features for the Firepower 1100 series.

Licenses	Essentials License
Firewall Licenses	
Botnet Traffic Filter	No Support.

Licenses	Essentials License	
Firewall Conns, Concurrent	Firepower 1120: 200,000 Firepower 1140: 400,000 Firepower 1150: 600,000	
Carrier	No support. Although SCTP inspection maps are not supported, SCTP stateful inspection using ACLs is supported.	
Total TLS Proxy Sessions	Firepower 1120: 4,000 Firepower 1140: 8,000 Firepower 1150: 8,000	
VPN Licenses		
Secure Client peers	Unlicensed	*Optional Secure Client Advantage, Secure Client Premier, or Secure Client VPN Only license, Maximum:* *Firepower 1120: 150* *Firepower 1140: 400* *Firepower 1150: 800*
Other VPN Peers	Firepower 1120: 150 Firepower 1140: 400 Firepower 1150: 800	
Total VPN Peers, combined all types	Firepower 1120: 150 Firepower 1140: 400 Firepower 1150: 800	
General Licenses		
Encryption	Base (DES) or Strong (3DES/AES), depending on the account's export compliance setting	
Security Contexts	2	*Optional License, Maximum:* *Firepower 1120: 5* *Firepower 1140: 10* *Firepower 1150: 25*
Clustering	No support.	
VLANs, Maximum	1024	

Secure Firewall 1210 and 1220

The following table shows the licensed features for the Secure Firewall 1210 and 1220.

Licenses	Essentials License	
Firewall Licenses		
Botnet Traffic Filter	No Support.	
Firewall Conns, Concurrent	Secure Firewall 1210: 200,000 Secure Firewall 1220: 300,000	
Carrier	No support. Although SCTP inspection maps are not supported, SCTP stateful inspection using ACLs is supported.	
Total TLS Proxy Sessions	Secure Firewall 1210: 320 Secure Firewall 1220: 320	
VPN Licenses		
Secure Client peers	Unlicensed	*Optional Secure Client Advantage,Secure Client Premier, or Secure Client VPN Only license, Maximum:* *Secure Firewall 1210: 200* *Secure Firewall 1220: 300*
Other VPN Peers	Secure Firewall 1210: 200 Secure Firewall 1220: 300	
Total VPN Peers, combined all types	Secure Firewall 1210: 200 Secure Firewall 1220: 300	
General Licenses		
Encryption	Base (DES) or Strong (3DES/AES), depending on the account's export compliance setting	
Security Contexts	2	*Optional License, Maximum:* *Secure Firewall 1210: 5* *Secure Firewall 1220: 10*
Clustering	No support.	
VLANs, Maximum	1024	

Secure Firewall 1230, 1240, 1250

The following table shows the licensed features for the Secure Firewall 1230, 1240, 1250.

Licenses	Essentials License	
Firewall Licenses		
Botnet Traffic Filter	No Support.	
Firewall Conns, Concurrent	Secure Firewall 1230: 400,000 Secure Firewall 1240: 600,000 Secure Firewall 1250: 1,000,000	
No support. Although SCTP inspection maps are not supported, SCTP stateful inspection using ACLs is supported.	No support. Although SCTP inspection maps are not supported, SCTP stateful inspection using ACLs is supported.	
Total TLS Proxy Sessions	Secure Firewall 1230: 1000 Secure Firewall 1240: 1000 Secure Firewall 1250: 1000	
VPN Licenses		
Secure Client peers	Unlicensed	*Optional Secure Client Advantage, Secure Client Premier, or Secure Client VPN Only license, Maximum:* *Secure Firewall 1230: 500* *Secure Firewall 1240: 1000* *Secure Firewall 1250: 1500*
Other VPN Peers	Secure Firewall 1230: 500 Secure Firewall 1240: 1000 Secure Firewall 1250: 1500	
Total VPN Peers, combined all types	Secure Firewall 1230: 500 Secure Firewall 1240: 1000 Secure Firewall 1250: 1500	
General Licenses		
Encryption	Base (DES) or Strong (3DES/AES), depending on the account's export compliance setting	

Licenses	Essentials License	
Security Contexts	2	*Optional License, Maximum:* *Secure Firewall 1230: 25* *Secure Firewall 1240: 25* *Secure Firewall 1250: 25*
Clustering	Enabled	
VLANs, Maximum	1024	

Secure Firewall 3100 Series

The following table shows the licensed features for the Secure Firewall 3100 series.

Licenses	Essentials License	
Firewall Licenses		
Botnet Traffic Filter	No Support.	
Firewall Conns, Concurrent	Secure Firewall 3105: 2,000,000 Secure Firewall 3110: 2,000,000 Secure Firewall 3120: 4,000,000 Secure Firewall 3130: 6,000,000 Secure Firewall 3140: 10,000,000	
Carrier	Disabled	*Optional License: Carrier*
Total TLS Proxy Sessions	Secure Firewall 3105: 10,000 Secure Firewall 3110: 10,000 Secure Firewall 3120: 15,000 Secure Firewall 3130: 15,000 Secure Firewall 3140: 15,000	
VPN Licenses		

Licenses	Essentials License	
Secure Client peers	Unlicensed	*Optional Secure Client Advantage, Secure Client Premier, or Secure Client VPN Only license, Maximum:* *Secure Firewall 3105: 3000* *Secure Firewall 3110: 3000* *Secure Firewall 3120: 7000* *Secure Firewall 3130: 15,000* *Secure Firewall 3140: 20,000*
Other VPN Peers	Secure Firewall 3105: 3000 Secure Firewall 3110: 3000 Secure Firewall 3120: 7000 Secure Firewall 3130: 15,000 Secure Firewall 3140: 20,000	
Total VPN Peers, combined all types	Secure Firewall 3105: 3000 Secure Firewall 3110: 3000 Secure Firewall 3120: 7000 Secure Firewall 3130: 15,000 Secure Firewall 3140: 20,000	
General Licenses		
Encryption	Base (DES) or Strong (3DES/AES), depending on the account's export compliance setting	
Security Contexts	2	*Optional License, Maximum: 100*
Clustering	Enabled	
VLANs, Maximum	1024	

Firepower 4100

The following table shows the licensed features for the Firepower 4100.

Licenses	Essentials License
Firewall Licenses	
Botnet Traffic Filter	No Support.

Licenses	Essentials License	
Firewall Conns, Concurrent	Firepower 4112: 10,000,000 Firepower 4115: 15,000,000 Firepower 4125: 25,000,000 Firepower 4145: 40,000,000	
Carrier	Disabled	*Optional License: Carrier*
Total TLS Proxy Sessions	15,000	
VPN Licenses		
Secure Client peers	Unlicensed	*Optional Secure Client Advantage, Secure Client Premier, or Secure Client VPN Only license:* *Firepower 4112: 10,000* *Firepower 4115: 15,000* *Firepower 4125: 20,000* *Firepower 4145: 20,000*
Other VPN Peers	Firepower 4112: 10,000 Firepower 4115: 15,000 Firepower 4125: 20,000 Firepower 4145: 20,000	
Total VPN Peers, combined all types	Firepower 4112: 10,000 Firepower 4115: 15,000 Firepower 4125: 20,000 Firepower 4145: 20,000	
General Licenses		
Encryption	Base (DES) or Strong (3DES/AES), depending on the account's export compliance setting	
Security Contexts	10	*Optional License: Maximum of 250*
Clustering	Enabled	
VLANs, Maximum	1024	

Secure Firewall 4200 Series

The following table shows the licensed features for the Secure Firewall 4200 series.

Licenses	Essentials License	
Firewall Licenses		
Botnet Traffic Filter	No Support.	
Firewall Conns, Concurrent	Secure Firewall 4215: 40,000,000 Secure Firewall 4225: 90,000,000 Secure Firewall 4245: 180,000,000	
Carrier	Disabled	*Optional License: Carrier*
Total TLS Proxy Sessions	15,000	
VPN Licenses		
Secure Client peers	Unlicensed	*Optional Secure Client Advantage, Secure Client Premier, or Secure Client VPN Only license, Maximum:* *Secure Firewall 4215: 20,000* *Secure Firewall 4225: 25,000* *Secure Firewall 4245: 30,000*
Other VPN Peers	Secure Firewall 4215: 20,000 Secure Firewall 4225: 25,000 Secure Firewall 4245: 30,000	
Total VPN Peers, combined all types	Secure Firewall 4215: 20,000 Secure Firewall 4225: 25,000 Secure Firewall 4245: 30,000	
General Licenses		
Encryption	Base (DES) or Strong (3DES/AES), depending on the account's export compliance setting	
Security Contexts	10	*Optional License, Maximum: 250*
Clustering	Enabled	
VLANs, Maximum	1024	

Firepower 9300

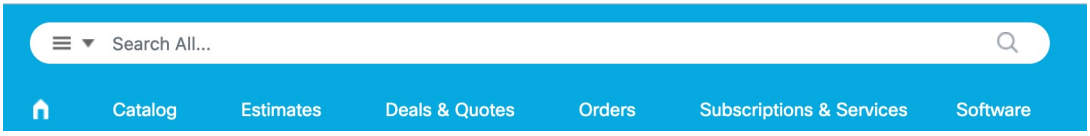
The following table shows the licensed features for the Firepower 9300.

Licenses	Essentials License	
Firewall Licenses		
Botnet Traffic Filter	No Support.	
Firewall Conns, Concurrent	Firepower 9300 SM-56: 60,000,000 Firepower 9300 SM-48: 60,000,000 Firepower 9300 SM-40: 55,000,000	
Carrier	Disabled	*Optional License: Carrier*
Total TLS Proxy Sessions	15,000	
VPN Licenses		
Secure Client peers	Unlicensed	*Optional Secure Client Advantage, Secure Client Premier, or Secure Client VPN Only license: 20,000 maximum*
Other VPN Peers	20,000	
Total VPN Peers, combined all types	20,000	
General Licenses		
Encryption	Base (DES) or Strong (3DES/AES), depending on the account's export compliance setting	
Security Contexts	10	*Optional License: Maximum of 250*
Clustering	Enabled	
VLANs, Maximum	1024	

License PIDs Per Model

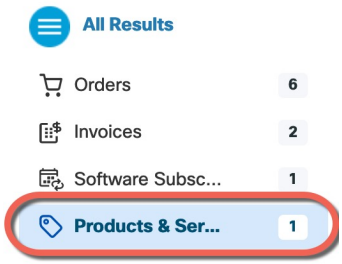
When you bought your device from Cisco or a reseller, your licenses should have been linked to your Smart Software License account. However, if you need to add licenses yourself, use the **Search All** field on the Cisco Commerce Workspace.

Figure 30: License Search



Choose **Products & Services** from the results.

Figure 31: Results



ASA Virtual PIDs

ASA Virtual Smart Software Manager Regular and On-Prem PIDs:

- ASAv5 license—L-ASAV5S-K9=
- ASAv10 license—L-ASAV10S-K9=
- ASAv30 license—L-ASAV30S-K9=
- ASAv50 license—L-ASAV50S-K9=
- ASAv100 license—L-ASAV100S-1Y=
- ASAv100 license—L-ASAV100S-3Y=
- ASAv100 license—L-ASAV100S-5Y=
- ASAvU license—L-ASA-V-U-K9=

Note The ASAv100 is a subscription-based license, available in terms of 1 year, 3 years, or 5 years.

ASA Virtual Permanent License Reservation PIDs:

The permanent license includes all available features, including the Strong Encryption (3DES/AES) license if your account qualifies. The Secure Client capabilities are also enabled to the platform maximum, contingent on your purchase of the Secure Client license that enables the right to use Secure Client (see Secure Client Advantage, Secure Client Premier, and Secure Client VPN Only Licenses, on page 121).

- ASAv5 license—L-ASAV5SR-K9=
- ASAv10 license—L-ASAV10SR-K9=
- ASAv30 license—L-ASAV30SR-K9=
- ASAv50 license—L-ASAV50SR-K9=
- ASAv100 license—L-ASAV100SR-K9=
- ASAvU license—ASA-V-U-ULR-K9=

Firepower 1010 PIDs

Firepower 1010 Smart Software Manager Regular and On-Prem PIDs:

- Essentials—L-FPR1000-ASA=. Required.
- Security Plus—L-FPR1010-SEC-PL=. The Security Plus license enables failover.
- Strong Encryption (3DES/AES)—L-FPR1K-ENC-K9=. Only required if your account is not authorized for strong encryption.

Firepower 1010 Permanent License Reservation PID:

The permanent license includes all available features, including the Strong Encryption (3DES/AES) license if your account qualifies. The Secure Client capabilities are also enabled to the platform maximum, contingent on your purchase of the Secure Client license that enables the right to use Secure Client (see Secure Client Advantage, Secure Client Premier, and Secure Client VPN Only Licenses, on page 121).

- L-FPR1K-ASA-BPU=

Firepower 1100 PIDs

Firepower 1100 Smart Software Manager Regular and On-Prem PIDs:

- Essentials—L-FPR1000-ASA=. Required.
- 5 context—L-FPR1K-ASASC-5=. Context licenses are additive; buy multiple licenses.
- 10 context—L-FPR1K-ASASC-10=. Context licenses are additive; buy multiple licenses.
- Strong Encryption (3DES/AES)—L-FPR1K-ENC-K9=. Only required if your account is not authorized for strong encryption.

Firepower 1100 Permanent License Reservation PID:

The permanent license includes all available features, including the Strong Encryption (3DES/AES) license if your account qualifies. The Secure Client capabilities are also enabled to the platform maximum, contingent on your purchase of the Secure Client license that enables the right to use Secure Client (see Secure Client Advantage, Secure Client Premier, and Secure Client VPN Only Licenses, on page 121).

- L-FPR1K-ASA-BPU=

Secure Firewall 1210/1220 PIDs

Secure Firewall 1210/1220 Smart Software Manager Regular and On-Prem PIDs:

- Essentials—*Included automatically*.
- 5 context—CSF1200-ASASC-5=. Context licenses are additive; buy multiple licenses.
- Strong Encryption (3DES/AES)—L-CSF1200-ENCK9=. Only required if your account is not authorized for strong encryption.

Secure Firewall 1210/1220 Permanent License Reservation PID:

The permanent license includes all available features, including the Strong Encryption (3DES/AES) license if your account qualifies. The Secure Client capabilities are also enabled to the platform maximum, contingent

on your purchase of the Secure Client license that enables the right to use Secure Client (see Secure Client Advantage, Secure Client Premier, and Secure Client VPN Only Licenses, on page 121).

- L-CSF1200-ASA-BPU=

Secure Firewall 1230/1240/1250 PIDs

Secure Firewall 1230/1240/1250 Smart Software Manager Regular and On-Prem PIDs:

- Essentials—*Included automatically*.
- 5 context—CSF1200-ASASC-5=. Context licenses are additive; buy multiple licenses.
- 10 context—CSF1200-ASASC-10=. Context licenses are additive; buy multiple licenses.
- Strong Encryption (3DES/AES)—L-CSF1200-ENCK9=. Only required if your account is not authorized for strong encryption.

Secure Firewall 1230/1240/1250 Permanent License Reservation PID:

The permanent license includes all available features, including the Strong Encryption (3DES/AES) license if your account qualifies. The Secure Client capabilities are also enabled to the platform maximum, contingent on your purchase of the Secure Client license that enables the right to use Secure Client (see Secure Client Advantage, Secure Client Premier, and Secure Client VPN Only Licenses, on page 121).

- L-CSF1200-ASA-BPU=

Secure Firewall 3100 PIDs

Secure Firewall 3100 Smart Software Manager Regular and On-Prem PIDs:

- Essentials—*Included automatically*.
- 5 context—L-FPR3K-ASASC-5=. Context licenses are additive; buy multiple licenses.
- 10 context—L-FPR3K-ASASC-10=. Context licenses are additive; buy multiple licenses.
- Carrier (Diameter, GTP/GPRS, M3UA, SCTP)—L-FPR3K-ASA-CAR=
- Strong Encryption (3DES/AES)—L-FPR3K-ENC-K9=. Only required if your account is not authorized for strong encryption.

Firepower 3100 Permanent License Reservation PID:

The permanent license includes all available features, including the Strong Encryption (3DES/AES) license if your account qualifies. The Secure Client capabilities are also enabled to the platform maximum, contingent on your purchase of the Secure Client license that enables the right to use Secure Client (see Secure Client Advantage, Secure Client Premier, and Secure Client VPN Only Licenses, on page 121).

- L-FPR3K-ASA-BPU=

Firepower 4100 PIDs

Firepower 4100 Smart Software Manager Regular and On-Prem PIDs:

- Essentials—L-FPR4100-ASA=. Required.
- 10 context—L-FPR4K-ASASC-10=. Context licenses are additive; buy multiple licenses.

- 230 context—L-FPR4K-ASASC-230=. Context licenses are additive; buy multiple licenses.
- 250 context—L-FPR4K-ASASC-250=. Context licenses are additive; buy multiple licenses.
- Carrier (Diameter, GTP/GPRS, M3UA, SCTP)—L-FPR4K-ASA-CAR=
- Strong Encryption (3DES/AES)—L-FPR4K-ENC-K9=. Only required if your account is not authorized for strong encryption.

Firepower 4100 Permanent License Reservation PID:

The permanent license includes all available features, including the Strong Encryption (3DES/AES) license if your account qualifies. The Secure Client capabilities are also enabled to the platform maximum, contingent on your purchase of the Secure Client license that enables the right to use Secure Client (see Secure Client Advantage, Secure Client Premier, and Secure Client VPN Only Licenses, on page 121).

- L-FPR4K-ASA-BPU=

Secure Firewall 4200 PIDs

Secure Firewall 4200 Smart Software Manager Regular and On-Prem PIDs:

- Essentials—*Included automatically*.
- 5 context—L-FPR4200-ASASC-5=. Context licenses are additive; buy multiple licenses.
- 10 context—L-FPR4200-ASASC-10=. Context licenses are additive; buy multiple licenses.
- Carrier (Diameter, GTP/GPRS, M3UA, SCTP)—L-FPR4200-ASA-CAR=
- Strong Encryption (3DES/AES)—L-FPR4200-ENC-K9=. Only required if your account is not authorized for strong encryption.

Firepower 4200 Permanent License Reservation PID:

The permanent license includes all available features, including the Strong Encryption (3DES/AES) license if your account qualifies. The Secure Client capabilities are also enabled to the platform maximum, contingent on your purchase of the Secure Client license that enables the right to use Secure Client (see Secure Client Advantage, Secure Client Premier, and Secure Client VPN Only Licenses, on page 121).

- L-FPR4200-ASA-BPU=

Firepower 9300 PIDs

Firepower 9300 Smart Software Manager Regular and On-Prem PIDs:

- Essentials—L-F9K-ASA=. Required.
- 10 context—L-F9K-ASA-SC-10=. Context licenses are additive; buy multiple licenses.
- Carrier (Diameter, GTP/GPRS, M3UA, SCTP)—L-F9K-ASA-CAR=
- Strong Encryption (3DES/AES)—L-F9K-ASA-ENCR-K9=. Only required if your account is not authorized for strong encryption.

Firepower 9300 Permanent License Reservation PIDs:

The permanent license includes all available features, including the Strong Encryption (3DES/AES) license if your account qualifies. The Secure Client capabilities are also enabled to the platform maximum, contingent on your purchase of the Secure Client license that enables the right to use Secure Client (see Secure Client Advantage, Secure Client Premier, and Secure Client VPN Only Licenses, on page 121).

- L-FPR9K-ASA-BPU=

Monitoring Smart Software Licensing

You can monitor the license features, status, and certificate, as well as enable debug messages.

Viewing Your Current License

See the followingscreen for viewing your license:

- **Configuration > Device Management > Licensing > Smart Licensing** pane and view the **Effective Running Licenses** area.

Viewing Smart License Status

See the following commands for viewing license status:

- **Monitoring > Properties > Smart License**

  Displays the state of Smart Software Licensing, Smart Agent version, UDI information, Smart Agent state, global compliance status, the entitlements status, licensing certificate information, and scheduled Smart Agent tasks.

Viewing the UDI

See the following command to view the universal product identifier (UDI):

show license udi

The following example shows the UDI for the ASAv:

```
ciscoasa# show license udi
UDI: PID:ASAv,SN:9AHV3KJBEKE
ciscoasa#
```

Smart Software Manager Communication

This section describes how your device communicates with the Smart Software Manager.

Device Registration and Tokens

For each virtual account, you can create a registration token. This token is valid for 30 days by default. Enter this token ID plus entitlement levels when you deploy each device, or when you register an existing device. You can create a new token if an existing token is expired.



Note Firepower 4100/9300 chassis—Device registration is configured in the chassis, not on the ASA logical device.

At startup after deployment, or after you manually configure these parameters on an existing device, the device registers with the Smart Software Manager. When the device registers with the token, the Smart Software Manager issues an ID certificate for communication between the device and the Smart Software Manager. This certificate is valid for 1 year, although it will be renewed every 6 months.

Periodic Communication with the Smart Software Manager

The device communicates with the Smart Software Manager every 30 days. If you make changes in the Smart Software Manager, you can refresh the authorization on the device so the change takes place immediately. Or you can wait for the device to communicate as scheduled.

You can optionally configure an HTTP proxy.

ASA Virtual

The ASA virtual must have internet access either directly or through an HTTP proxy at least every 90 days. Normal license communication occurs every 30 days, but with the grace period, your device will stay compliant for up to 90 days without calling home. After the grace period, you should contact the Smart Software Manager, or your ASA virtual will be out-of-compliance; operation is otherwise unaffected.

All Other Models

The ASA must have internet access either directly or through an HTTP proxy at least every 90 days. Normal license communication occurs every 30 days, but with the grace period, your device will operate for up to 90 days without calling home. After the grace period, you must contact the Smart Software Manager, or you will not be able to make configuration changes to features requiring special licenses; operation is otherwise unaffected.

Out-of-Compliance State

The device can become out of compliance in the following situations:

- Over-utilization—When the device uses unavailable licenses.
- License expiration—When a time-based license expires.
- Lack of communication—When the device cannot reach the Licensing Authority for re-authorization.

To verify whether your account is in, or approaching, an Out-of-Compliance state, you must compare the entitlements currently in use by your device against those in your Smart Account.

In an out-of-compliance state, the device might be limited, depending on the model:

- ASA Virtual—The ASA virtual is not affected.
- All Other Models—You will not be able to make configuration changes to features requiring special licenses, but operation is otherwise unaffected. For example, existing contexts over the Essentials license limit can continue to run, and you can modify their configuration, but you will not be able to add a *new* context. If you do not have sufficient Essentials licenses when you first register, you cannot configure any licensed features, including strong encryption features.

Smart Call Home Infrastructure

By default, a Smart Call Home profile exists in the configuration that specifies the URL for the Smart Software Manager. You cannot remove this profile. Note that the only configurable option for the License profile is the destination address URL for the Smart Software Manager. Unless directed by Cisco TAC, you should not change the Smart Software Manager URL.



Note For the Firepower 4100/9300 chassis, Smart Call Home for licensing is configured in the Firepower 4100/9300 chassis supervisor, not on the ASA.

You cannot disable Smart Call Home for Smart Software Licensing. For example, even if you disable Smart Call Home using the **no service call-home** command, Smart Software Licensing is not disabled.

Other Smart Call Home functions are not turned on unless you specifically configure them.

Smart License Certificate Management

The ASA automatically creates a trustpoint containing the certificate of the CA that issued the Smart Transport or Smart Call Home server certificate. To avoid service interruption if the issuing hierarchy of the server certificate changes, configure the **Automatic Import** area of the **Configuration** > **Remote Access VPN** > **Certificate Management** > **Trusted Certificate Pool** > **Edit Trusted Certificate Pool Policy** screen to enable the automatic update of the trustpool bundle at periodic intervals.

The server certificate received from a Smart License Server must contain "ServAuth" in the Extended Key Usage field. This check will be done on non self-signed certificates only; self-signed certificates do not provide any value in this field.

History for Smart Software Licensing

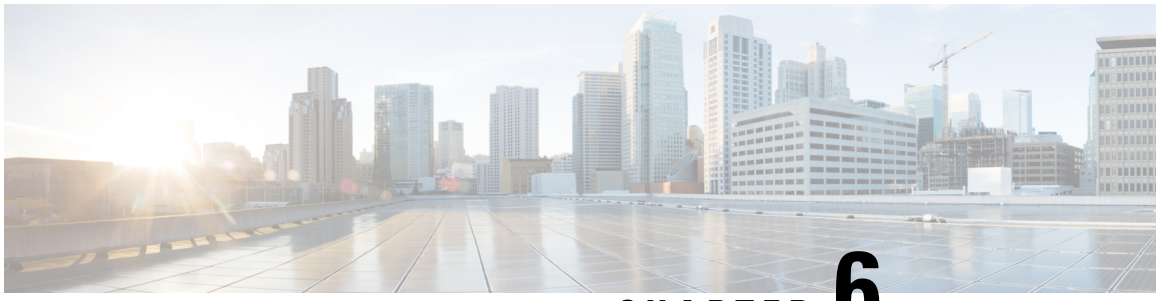
Feature Name	Platform Releases	Description
Increased connection limits for the Secure Firewall 4200	9.23(1)	Connection limits have been increased: • 4225: 80M → **90M** • 4245: 80M → **180M**

Feature Name	Platform Releases	Description
ASA Virtual Flexible Permanent License Reservation	9.23(1)	Flexible Permanent License Reservation allows you to configure any model-specific license irrespective of the RAM and vCPUs assigned to the ASA virtual. The model license sets the throughput level of the ASA virtual while the memory of ASA virtual determines the maximum Secure Client sessions, TLS proxy sessions, maximum concurrent firewall connections and VLANs. However, a lower ASA virtual memory profile caps the actual number of sessions and features.
Multiple context support for all Secure Firewall 1200 models	9.23(1)	We added support for multiple context mode for the Secure Firewall 1210/1220: • Secure Firewall 1210CE—5 contexts. • Secure Firewall 1210CP—5 contexts. • Secure Firewall 1220CX—10 contexts. Switchports are not supported in multiple context mode, and you must convert all interfaces to router interfaces before you can convert to multiple context mode. The Secure Firewall 1230/1240/1250 also supports multiple context mode in its initial release: • Secure Firewall 1230—25 contexts. • Secure Firewall 1240—25 contexts. • Secure Firewall 1250—25 contexts.
Smart Transport is the default transport mechanism to communicate with the CSSM server	9.22(1)	Smart Licensing uses Smart Transport as the default transport for all devices instead of Smart Call Home.
ASAvU license	9.22(1)	The ASAvU license achieves maximum throughput on deployments with 32-cores and 64-cores. This license is supported only on VMware and KVM. New/Modified screens: **Configuration** > **Device Management** > **Licensing** > **Smart Licensing**.
Increased connection limits for the Secure Firewall 4200	9.20(2)	Connection limits have been increased: • 4215: 15M → **40M** • 4225: 30M → **80M** • 4245: 60M → **80M**

Feature Name	Platform Releases	Description
Secure Firewall 3100 support for the Carrier license	9.18(1)	The Carrier license enables Diameter, GTP/GPRS, SCTP inspection. New/Modified screens: **Configuration** > **Device Management** > **Licensing** > **Smart Licensing**.
ASAv100 permanent license reservation	9.14(1.30)	The ASAv100 now supports permanent license reservation using product ID L-ASAV100SR-K9=. **Note:** Not all accounts are approved for permanent license reservation.
ASA Virtual MSLA Support	9.13(1)	The ASA virtual supports Cisco's Managed Service License Agreement (MSLA) program, which is a software licensing and consumption framework designed for Cisco customers and partners who offer managed software services to third parties. MSLA is a new form of Smart Licensing where the licensing Smart Agent keeps track of the usage of licensing entitlements in units of time. New/Modified screens: **Configuration** > **Device Management** > **Licensing** > **Smart Licensing**.
ASA Virtual Flexible Licensing	9.13(1)	Flexible Licensing is a new form of Smart Licensing where any ASA virtual license now can be used on any supported ASA virtual vCPU/memory configuration. Session limits for Secure Client and TLS proxy will be determined by the ASA virtual platform entitlement installed rather than a platform limit tied to a model type. New/Modified screens: **Configuration** > **Device Management** > **Licensing** > **Smart Licensing**.
Licensing changes for failover pairs on the Firepower 4100/9300 chassis	9.7(1)	Only the active unit requests the license entitlements. Previously, both units requested license entitlements. Supported with FXOS 2.1.1.
Permanent License Reservation for the ASA virtual Short String enhancement	9.6(2)	Due to an update to the Smart Agent (to 1.6.4), the request and authorization codes now use shorter strings. We did not modify any screens.
Satellite Server support for the ASA virtual	9.6(2)	If your devices cannot access the internet for security reasons, you can optionally install a local Smart Software Manager satellite server as a virtual machine (VM). We did not modify any screens.
Permanent License Reservation for the ASA on the Firepower 4100/9300 chassis	9.6(2)	For highly secure environments where communication with the Cisco Smart Software Manager is not allowed, you can request a permanent license for the ASA on the Firepower 9300 and Firepower 4100. All available license entitlements are included in the permanent license, including the Standard Tier, Strong Encryption (if qualified), Security Contexts, and Carrier licenses. Requires FXOS 2.0.1. All configuration is performed on the Firepower 4100/9300 chassis; no configuration is required on the ASA.

Feature Name	Platform Releases	Description
Permanent License Reservation for the ASA virtual	9.5(2.200) 9.6(2)	For highly secure environments where communication with the Cisco Smart Software Manager is not allowed, you can request a permanent license for the ASA virtual. In 9.6(2), we also added support for this feature for the ASA virtual on Amazon Web Services. This feature is not supported for Microsoft Azure. We introduced the following commands: **license smart reservation, license smart reservation cancel, license smart reservation install, license smart reservation request universal, license smart reservation return** No ASDM support.
Smart Agent Upgrade to v1.6	9.5(2.200) 9.6(2)	The smart agent was upgraded from Version 1.1 to Version 1.6. This upgrade supports permanent license reservation and also supports setting the Strong Encryption (3DES/AES) license entitlement according to the permission set in your license account. **Note** If you downgrade from Version 9.5(2.200), the ASA virtual does not retain the licensing registration state. You need to re-register with the **Configuration > Device Management > Licensing > Smart Licensing** page with the **Force registration** option; obtain the ID token from the Smart Software Manager. We did not change any screens.
Strong Encryption (3DES) license automatically applied for the ASA on the Firepower 9300	9.5(2.1)	For regular Cisco Smart Software Manager users, the Strong Encryption license is automatically enabled for qualified customers when you apply the registration token on the Firepower 9300. **Note** If you are using the Smart Software Manager satellite deployment, to use ASDM and other strong encryption features, after you deploy the ASA you must enable the Strong Encryption (3DES) license using the ASA CLI. This feature requires FXOS 1.1.3. We modified the following screen: **Configuration** > **Device Management** > **Licensing** > **Smart License**

Feature Name	Platform Releases	Description
Validation of the Smart Call Home/Smart Licensing certificate if the issuing hierarchy of the server certificate changes	9.5(2)	Smart licensing uses the Smart Call Home infrastructure. When the ASA first configures Smart Call Home anonymous reporting in the background, it automatically creates a trustpoint containing the certificate of the CA that issued the Smart Call Home server certificate. The ASA now supports validation of the certificate if the issuing hierarchy of the server certificate changes; you can enable the automatic update of the trustpool bundle at periodic intervals. We modified the following screen: **Configuration** > **Remote Access VPN** > **Certificate Management** > **Trusted Certificate Pool** > **Edit Trusted Certificate Pool Policy**
New Carrier license	9.5(2)	The new Carrier license replaces the existing GTP/GPRS license, and also includes support for SCTP and Diameter inspection. For the ASA on the Firepower 9300, the **feature mobile-sp** command will automatically migrate to the **feature carrier** command. We modified the following screen: **Configuration** > **Device Management** > **Licensing** > **Smart License**
Cisco Smart Software Licensing for the ASA on the Firepower 9300	9.4(1.150)	We introduced Smart Software Licensing for the ASA on the Firepower 9300. We modified the following screen: **Configuration** > **Device Management** > **Licensing** > **Smart License**
Cisco Smart Software Licensing for the ASA virtual	9.3(2)	Smart Software Licensing lets you purchase and manage a pool of licenses. Unlike PAK licenses, smart licenses are not tied to a specific serial number. You can easily deploy or retire ASA virtual's without having to manage each unit's license key. Smart Software Licensing also lets you see your license usage and needs at a glance. We introduced or modified the following screens: **Configuration > Device Management > Licensing > Smart License** **Configuration > Device Management > Smart Call-Home Monitoring > Properties > Smart License**



CHAPTER 6

Logical Devices for the Firepower 4100/9300

The Firepower 4100/9300 is a flexible security platform on which you can install one or more *logical devices*. This chapter describes basic interface configuration and how to add a standalone or High Availability logical device using the chassis manager. To add a clustered logical device, see ASA Cluster for the Firepower 4100/9300, on page 429. 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 197
- About Logical Devices, on page 200
- Requirements and Prerequisites for Hardware and Software Combinations, on page 200
- Guidelines and Limitations for Logical Devices, on page 201
- Configure Interfaces, on page 202
- Configure Logical Devices, on page 206
- History for Logical Devices, on page 212

About Interfaces

The Firepower 4100/9300 chassis supports physical interfaces 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 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 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 (threat defense-using-management center only).
- 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 197.

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 threat defense-using-management center devices.

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.

See the following table for interface type support for the threat defense and ASA applications in standalone and cluster deployments.

Table 9: Interface Type Support

Application		Data	Data: Subinterface	Data-Sharing	Data-Sharing: Subinterface	Mgmt	Eventing	Cluster (EtherChannel only)	Cluster: Subinterface
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**
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 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.

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.

About Logical Devices

A logical device lets you run one application instance (either ASA or 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 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.

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 and Container instances—When you install a container instance on a security module, that module can only support other container instances. A native instance uses all of the resources for a module, so you can only install a single native instance on a module. You can use native instances on some modules,

and container instances on the other module. For example, you can install a native instance on module 1 and module 2, but container instances on module 3.

- 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.
- 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 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 threat defense on module 3.
- ASA or 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 threat defense 6.3 on module 1, threat defense 6.4 on module 2, and 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.
- Clustering—All chassis in the cluster must be the same model.
- High Availability—High Availability is only supported between same-type models.
- ASA and threat defense application types—The Firepower 4100 can only run a single application type.

Guidelines and Limitations for Logical Devices

See the following sections for guidelines and limitations.

Guidelines and Limitations for Interfaces

Default MAC Addresses

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.

General Guidelines and Limitations

Firewall Mode

You can set the firewall mode to routed or transparent in the bootstrap configuration for the threat defense and ASA.

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.

Context Mode

- Enable multiple context mode in the ASA after you deploy.

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 other High Availability system requirements, see Failover System Requirements, on page 282.

Configure Interfaces

By default, physical interfaces are disabled. You can enable interfaces, add EtherChannels, and edit interface properties.

Note If you remove an interface in FXOS (for example, if you remove a network module, remove an EtherChannel, or reassign an interface to an EtherChannel), then the ASA configuration retains the original commands so that you can make any necessary adjustments; removing an interface from the configuration can have wide effects. You can manually remove the old interface configuration in the ASA OS.

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.

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 198 for details about interface type usage.

- **Data**
- **Mgmt**
- **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 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 198 for details about interface type usage.

- **Data**
- **Mgmt**
- **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 interfaces, choose the LACP port-channel **Mode**, **Active** or **On**.

For non-Data 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**.

Configure Logical Devices

Add a standalone logical device or a High Availability pair on the Firepower 4100/9300 chassis.

For clustering, see #unique_278.

Add a Standalone ASA

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 deploy a routed or transparent firewall mode ASA from the Firepower 4100/9300 chassis.

For multiple context mode, you must first deploy the logical device, and then enable multiple context mode in the ASA application.

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.

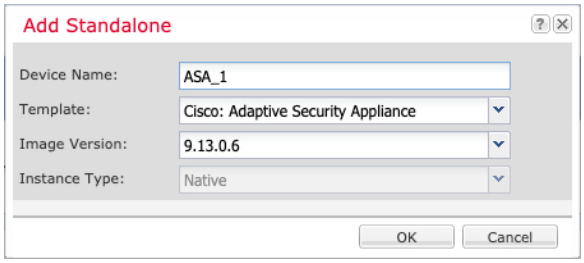
Note For the Firepower 9300, you can install different application types (ASA and threat defense) on separate modules in the chassis. You can also run different versions of an application instance type on separate modules.

- 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**).
- Gather the following information:
  - Interface IDs for this device
  - Management interface IP address and network mask
  - Gateway IP address

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: Adaptive Security Appliance**.
c) Choose the **Image Version**.
d) Click **OK**.

You see the Provisioning - *device name* window.

Step 3 Expand the **Data Ports** area, and click each port that you want to assign to the device.

You can only assign data interfaces that you previously enabled on the **Interfaces** page. You will later enable and configure these interfaces on the ASA, including setting the IP addresses.

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) Choose the **Management Interface**.

This interface is used to manage the logical device. This interface is separate from the chassis management port.

c) Choose the management interface **Address Type**: **IPv4 only**, **IPv6 only**, or **IPv4 and IPv6**.

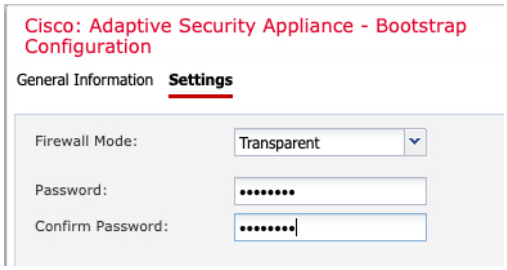
d) Configure the **Management IP** address.

Set a unique IP address for this interface.

e) Enter a **Network Mask** or **Prefix Length**.

f) Enter a **Network Gateway** address.

Step 6 Click the **Settings** tab.



Step 7 Choose the **Firewall Mode**: **Routed** or **Transparent**.

In routed mode, the ASA 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.

Step 8 Enter and confirm a **Password** for the admin user and for the enable password.

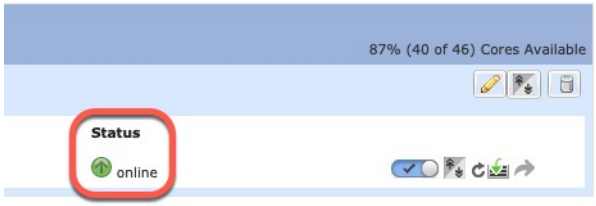
The pre-configured ASA admin user/password and enable password is useful for password recovery; if you have FXOS access, you can reset the admin user password/enable password if you forget it.

Step 9 Click **OK** to close the configuration dialog box.

Step 10 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 11 See the ASA configuration guide to start configuring your security policy.

Add a High Availability Pair

Threat DefenseASA 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 Failover System Requirements, on page 282.

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.

Step 3 Enable High Availability on the logical devices. See Failover for High Availability, on page 281.

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.

Note
For the ASA, if you remove an interface in FXOS (for example, if you remove a network module, remove an EtherChannel, or reassign an interface to an EtherChannel), then the ASA configuration retains the original

commands so that you can make any necessary adjustments; removing an interface from the configuration can have wide effects. You can manually remove the old interface configuration in the ASA OS.

Change an Interface on an ASA Logical Device

You can allocate, unallocate, or replace a management interface on an ASA logical device. ASDM discovers the new interfaces automatically.

Adding a new interface, or deleting an unused interface has minimal impact on the ASA configuration. However, if you remove an allocated interface in FXOS (for example, if you remove a network module, remove an EtherChannel, or reassign an allocated interface to an EtherChannel), and the interface is used in your security policy, removal will impact the ASA configuration. In this case, the ASA configuration retains the original commands so that you can make any necessary adjustments. You can manually remove the old interface configuration in the ASA OS.

Note You can edit the membership of an allocated EtherChannel without impacting the logical device.

Before you begin

- Configure your interfaces and add any EtherChannels according to Configure a Physical Interface, on page 203 and Add an EtherChannel (Port Channel), on page 204.
- 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 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 ASA reloads (management interface changes cause a reload), you can add the (now unallocated) management interface to the EtherChannel as well.
- For clustering or failover, make sure you add or remove the interface on all units. We recommend that you make the interface changes on the data/standby unit(s) first, and then on the control/active unit. New interfaces are added in an administratively down state, so they do not affect interface monitoring.

Procedure

Step 1 In the chassis manager, choose **Logical Devices**.

Step 2 Click the **Edit** icon at the top right to edit the logical device.

Step 3 Unallocate a data interface by de-selecting the interface in the **Data Ports** area.

Step 4 Allocate a new data interface by selecting the interface in the **Data Ports** area.

Step 5 Replace the management interface:

For this type of interface, the device reloads after you save your changes.

a) Click the device icon in the center of the page.
b) On the **General/Cluster Information** tab, choose the new **Management Interface** from the drop-down list.
c) Click **OK**.

Step 6 Click **Save**.

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 asa *name*

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

Example:

```
Firepower-module1> connect asa asa1
Connecting to asa(asa1) console... hit Ctrl + A + D  to return to bootCLI
[...]
asa>
```

Step 3 Exit the application console to the FXOS module CLI.

- ASA—Enter **Ctrl-a, d**

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-], .**

Example

The following example connects to an ASA on security module 1 and then exits back to the supervisor level of the FXOS CLI.

```
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

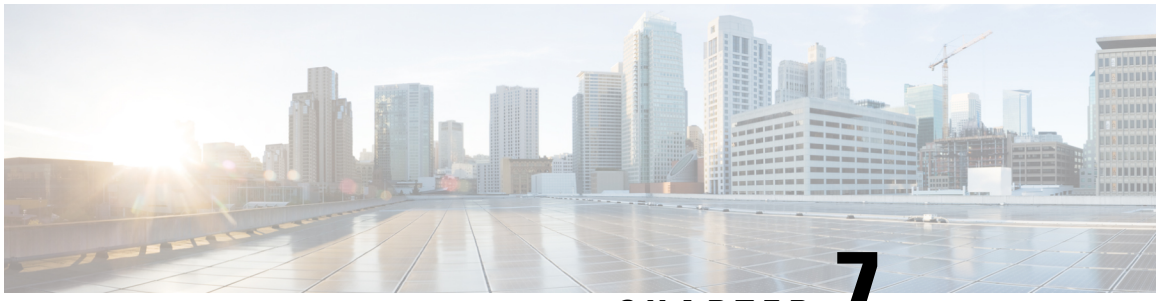
Firepower-module1>connect asa asa1
asa> ~
telnet> quit
Connection closed.
Firepower#
```

History for Logical Devices

Feature	Version	Details
ASA for the Firepower 4112	9.14(1)	We introduced the Firepower 4112. **Note** Requires FXOS 2.8.1.
Firepower 9300 SM-56 support	9.12.2	We introduced the SM-56 security module. **Note** Requires FXOS 2.6.1.157.
ASA for the Firepower 4115, 4125, and 4145	9.12(1)	We introduced the Firepower 4115, 4125, and 4145. **Note** Requires FXOS 2.6.1.

Feature	Version	Details
Firepower 9300 SM-40 and SM-48 support	9.12.1	We introduced the SM-40 and SM-48 security modules. **Note** Requires FXOS 2.6.1.
Support for ASA and threat defense on separate modules of the same Firepower 9300	9.12.1	You can now deploy ASA and threat defense logical devices on the same Firepower 9300. **Note** Requires FXOS 2.6.1.
Cluster control link customizable IP Address for the Firepower 4100/9300	9.10.1	By default, the cluster control link uses the 127.2.0.0/16 network. You can now set the network whe you deploy the cluster in FXOS. The chassis auto-generates the cluster control link interface IP addres for each unit based on the chassis ID and slot ID: 127.2.*chassis_id*.*slot_id*. However, some networkin deployments do not allow 127.2.0.0/16 traffic to pass. Therefore, you can now set a custom /16 subn for the cluster control link in FXOS except for loopback (127.0.0.0/8) and multicast (224.0.0.0/4) addresses. **Note** Requires FXOS 2.4.1. New/modified Firepower Chassis Manager screens: **Logical Devices** > **Add Device** > **Cluster Information** > **CCL Subnet IP** field
Support for data EtherChannels in On mode	9.10.1	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. **Note** Requires FXOS 2.4.1. New/Modified Firepower Chassis Manager screens: **Interfaces** > **All Interfaces** > **Edit Port Channel** > **Mode**
Inter-site clustering improvement for the ASA on the Firepower 4100/9300 chassis	9.7(1)	You can now configure the site ID for each Firepower 4100/9300 chassis when you deploy the ASA cluster. Previously, you had to configure the site ID within the ASA application; this new feature ease initial deployment. Note that you can no longer set the site ID within the ASA configuration. Also, fo best compatibility with inter-site clustering, we recommend that you upgrade to ASA 9.7(1) and FXO 2.1.1, which includes several improvements to stability and performance. We modified the following screen: **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** > **Cluster Configuration**
Support for the Firepower 4100 series	9.6(1)	With FXOS 1.1.4, the ASA supports inter-chassis clustering on the Firepower 4100 series. We did not modify any screens.
Inter-chassis clustering for 6 modules, and inter-site clustering for the Firepower 9300 ASA application	9.5(2.1)	With FXOS 1.1.3, you can now enable inter-chassis, and by extension inter-site clustering. You can include up to 6 modules in up to 6 chassis. We did not modify any screens.

Feature	Version	Details
Intra-chassis ASA Clustering for the Firepower 9300	9.4(1.150)	You can cluster up to 3 security modules within the Firepower 9300 chassis. All modules in the chassis must belong to the cluster. We introduced the following screen: **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster Replication**



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.

You can set the firewall mode independently for each context in multiple context mode.

- About the Firewall Mode, on page 215
- Default Settings, on page 223
- Guidelines for Firewall Mode, on page 223
- Set the Firewall Mode (Single Mode), on page 225
- Examples for Firewall Mode, on page 226
- History for the Firewall Mode, on page 236

About the Firewall Mode

The ASA supports two firewall modes: Routed Firewall mode and Transparent Firewall mode.

About Routed Firewall Mode

In routed mode, the ASA is considered to be a router hop in the network. Each interface that you want to route between is on a different subnet. You can share Layer 3 interfaces between contexts.

With Integrated Routing and Bridging, you can use a "bridge group" where you group together multiple interfaces on a network, and the ASA 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 ASA routes between BVIs and regular routed interfaces. If you do not need multiple context mode or clustering or EtherChannel or VNI 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.

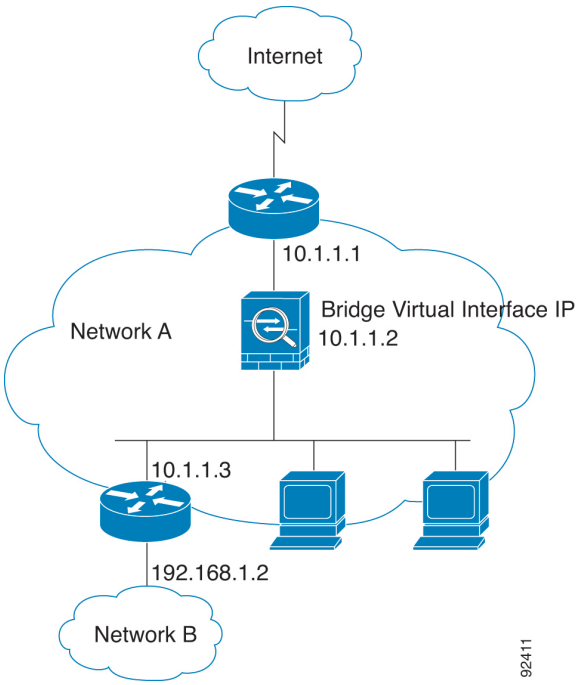
Layer 2 connectivity is achieved by using a "bridge group" where you group together the inside and outside interfaces for a network, and the ASA uses bridging techniques to pass traffic between the interfaces. Each bridge group includes a Bridge Virtual Interface (BVI) to which you assign an IP address 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 ASA 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 32: Transparent Firewall Network



Management Interface

In addition to each Bridge Virtual Interface (BVI) IP address, you can add a separate Management *slot*/*port* interface that is not part of any bridge group, and that allows only management traffic to the ASA. For more information, see Management Interface, on page 562.

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 ASA.

About Bridge Groups

A bridge group is a group of interfaces that the ASA 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 ASA 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:

- Access rules—You can configure access rules for both bridge group member interfaces and for the BVI; for inbound rules, the member interface is checked first. For outbound rules, the BVI is checked first.
- 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 ASA—When specifying a syslog server (or SNMP server, or other service where the traffic is sourced from the ASA), you can specify either the BVI or a member interface.

If you do not name the BVI in routed mode, then the ASA does not route bridge group traffic. This configuration replicates transparent firewall mode for the bridge group. If you do not need multiple context mode or clustering or EtherChannel or VNI 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.

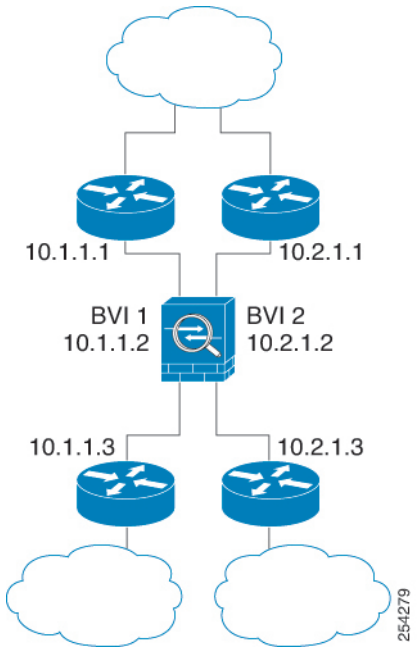
Bridge Groups in Transparent Firewall Mode

Bridge group traffic is isolated from other bridge groups; traffic is not routed to another bridge group within the ASA, and traffic must exit the ASA before it is routed by an external router back to another bridge group in the ASA. 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. For complete security policy separation, use security contexts with one bridge group in each context.

You can include multiple interfaces per bridge group. See Guidelines for Firewall Mode, on page 223 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 ASA, which has two bridge groups.

Figure 33: Transparent Firewall Network with Two Bridge Groups

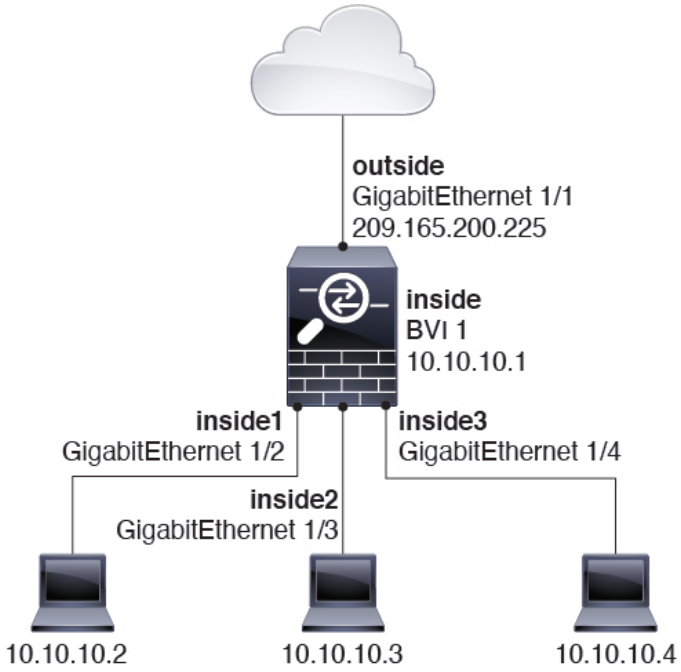


Bridge Groups in Routed Firewall Mode

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 ASA 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. For example, as in the default configuration, set all the interfaces to the same security level, and then enable same-security interface communication; no access rule is required.

Figure 34: Routed Firewall Network with an Inside Bridge Group and an Outside Routed Interface



Passing Traffic Not Allowed in Routed Mode

In routed mode, some types of traffic cannot pass through the ASA even if you allow it in an access rule. The bridge group, however, can allow almost any traffic through using either an access rule (for IP traffic) or an EtherType rule (for non-IP traffic):

- IP traffic—In routed firewall mode, broadcast and multicast traffic is blocked even if you allow it in an access rule, including unsupported dynamic routing protocols and DHCP (unless you configure DHCP relay). Within a bridge group, you can allow this traffic with an access rule.
- Non-IP traffic—AppleTalk, IPX, BPDUs, and MPLS, for example, can be configured to go through using an EtherType rule.

Note The bridge group does not pass CDP packets packets, or any packets that do not have a valid EtherType greater than or equal to 0x600. An exception is made for BPDUs and IS-IS, which are supported.

Allowing Layer 3 Traffic

- Unicast IPv4 and IPv6 traffic is allowed through the bridge group automatically from a higher security interface to a lower security interface, without an access rule.
- For Layer 3 traffic traveling from a low to a high security interface, an access rule is required on the low security interface.
- 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 219). 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
- AppleTalk multicast MAC addresses from 0900.0700.0000 to 0900.07FF.FFFF

BPDU Handling

To prevent loops using the Spanning Tree Protocol, BPDUs are passed by default. To block BPDUs, you need to configure an EtherType rule to deny them. You can also block BPDUs on the external switches. For example, you can block BPDUs on the switch if members of the same bridge group are connected to switch ports in different VLANs. In this case, BPDUs from one VLAN will be visible in the other VLAN, which can cause Spanning Tree Root Bridge election process problems.

If you are using failover, you might want to block BPDUs to prevent the switch port from going into a blocking state when the topology changes. See Bridge Group Requirements for Failover, on page 291 for more information.

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 ASA—Add a default/static route on the ASA for traffic destined for a remote network where a syslog server, for example, is located.
- Voice over IP (VoIP) and TFTP traffic with inspection enabled, and the endpoint is at least one hop away—Add a static route on the ASA for traffic destined for the remote endpoint so that secondary connections are successful. The ASA 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 ASA needs to perform a route lookup to install the pinhole on the correct interface.

  Affected applications include:

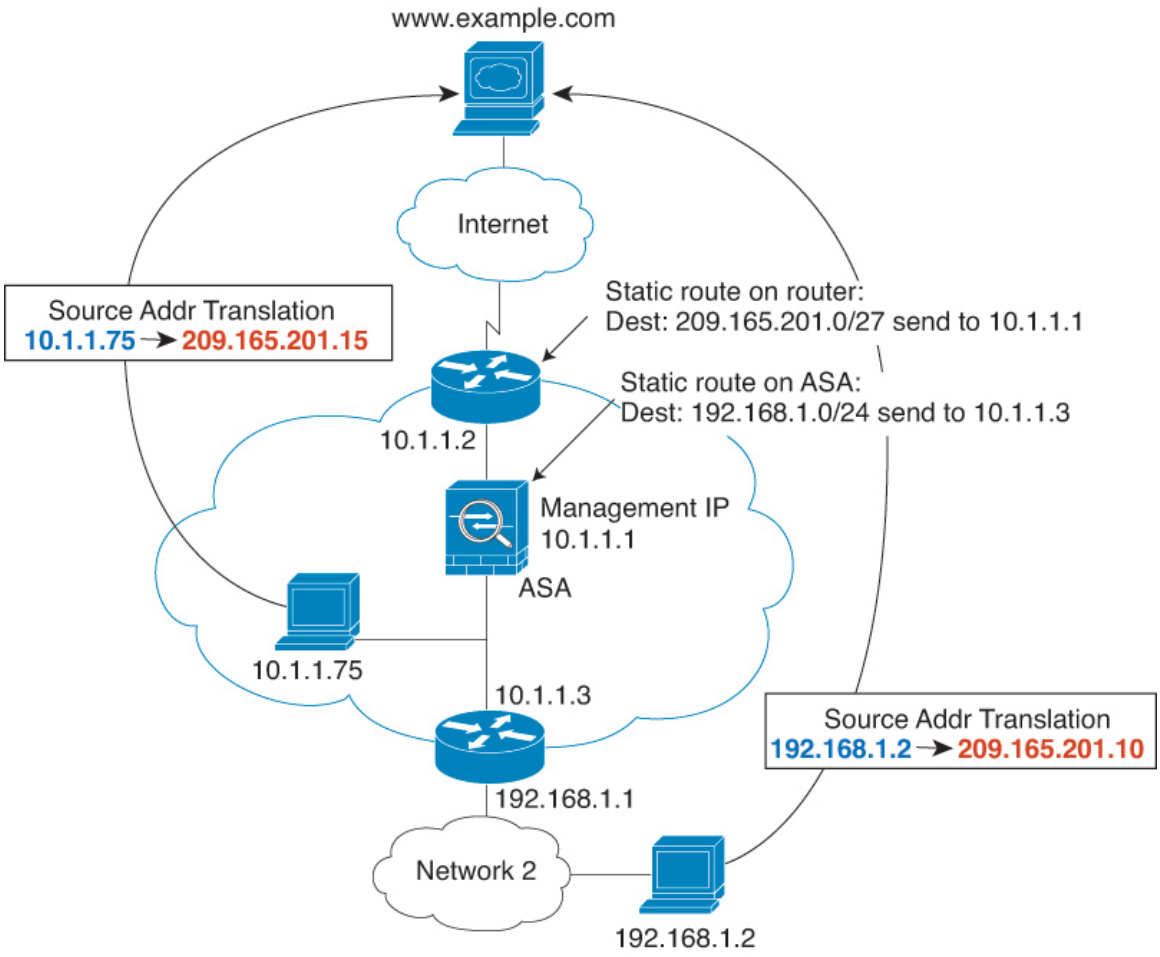
  - CTIQBE
  - GTP
  - H.323
  - MGCP
  - RTSP

- SIP
- Skinny (SCCP)
- SQL*Net
- SunRPC
- TFTP

- Traffic at least one hop away for which the ASA performs NAT—Configure a static route on the ASA 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 ASA.

  This routing requirement is also true for embedded IP addresses for VoIP and DNS with inspection and NAT enabled, and the embedded IP addresses are at least one hop away. The ASA needs to identify the correct egress interface so it can perform the translation.

Figure 35: 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 10: 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.
Dynamic routing protocols	You can, however, add static routes for traffic originating on the ASA for bridge group member interfaces. You can also allow dynamic routing protocols through the ASA using an access rule.
Multicast IP routing	You can allow multicast traffic through the ASA 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 ASA. You can pass VPN traffic through the ASA using an access rule, but it does not terminate non-management connections. Clientless SSL VPN is also not supported.
Unified Communications	—

Unsupported Features for Bridge Groups in Routed Mode

The following table lists the features are not supported in bridge groups in routed mode.

Table 11: Unsupported Features in Routed Mode

Feature	Description
EtherChannel or VNI member interfaces	Only physical 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.

Feature	Description
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 ASA using an access rule. Non-bridge group interfaces support dynamic routing.
Multicast IP routing	You can allow multicast traffic through the ASA by allowing it in an access rule. Non-bridge group interfaces support multicast routing.
Multiple Context Mode	Bridge groups are not supported in multiple context mode.
QoS	Non-bridge group interfaces support QoS.
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 ASA. You can pass VPN traffic through the bridge group using an access rule, but it does not terminate non-management connections. Clientless SSL VPN is also not supported.
Unified Communications	Non-bridge group interfaces support Unified Communications.

Default Settings

Default Mode

The default mode is routed mode.

Bridge Group Defaults

By default, all ARP packets are passed within the bridge group.

Guidelines for Firewall Mode

Context Mode Guidelines

Set the firewall mode per context.

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 ASA 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 ASA. 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 the ASAv50 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.
- 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 ASA 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.
- In transparent mode, PPPoE is not supported for the Management interface.
- In routed mode, to route between bridge groups and other routed interfaces, you must name the BVI.
- In routed mode, ASA-defined EtherChannel and VNI 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 ASA when using bridge group members. If there are two neighbors on either side of the ASA running BFD, then the ASA 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 Limitations

- When you change firewall modes, the ASA clears the running configuration because many commands are not supported for both modes. The startup configuration remains unchanged. If you reload without saving, then the startup configuration is loaded, and the mode reverts back to the original setting. See Set the Firewall Mode (Single Mode), on page 225 for information about backing up your configuration file.
- If you download a text configuration to the ASA that changes the mode with the **firewall transparent** command, be sure to put the command at the top of the configuration; the ASA changes the mode as soon as it reads the command and then continues reading the configuration you downloaded. If the command appears later in the configuration, the ASA clears all the preceding lines in the configuration.

Set the Firewall Mode (Single Mode)

This section describes how to change the firewall mode using the CLI. For single mode and for the currently connected context in multiple mode (typically the admin context), you cannot change the mode in ASDM. For other multiple mode contexts, you can set the mode in ASDM for each context; see Configure a Security Context, on page 266.

Note We recommend that you set the firewall mode before you perform any other configuration because changing the firewall mode clears the running configuration.

Before you begin

When you change modes, the ASA clears the running configuration (see Guidelines for Firewall Mode, on page 223 for more information).

- If you already have a populated configuration, be sure to back up your configuration before changing the mode; you can use this backup for reference when creating your new configuration.
- Use the CLI at the console port to change the mode. If you use any other type of session, including the ASDM Command Line Interface tool or SSH, you will be disconnected when the configuration is cleared, and you will have to reconnect to the ASA using the console port in any case.
- Set the mode within the context.

Note To set the firewall mode to transparent and also configure ASDM management access after the configuration is cleared, see Configure ASDM Access, on page 21.

Procedure

Set the firewall mode to transparent:

firewall transparent

Example:

```
ciscoasa(config)# firewall transparent
```

To change the mode to routed, enter the **no firewall transparent** command.

Note
You are not prompted to confirm the firewall mode change; the change occurs immediately.

Examples for Firewall Mode

This section includes examples of how traffic moves through the ASA in the routed and transparent firewall mode.

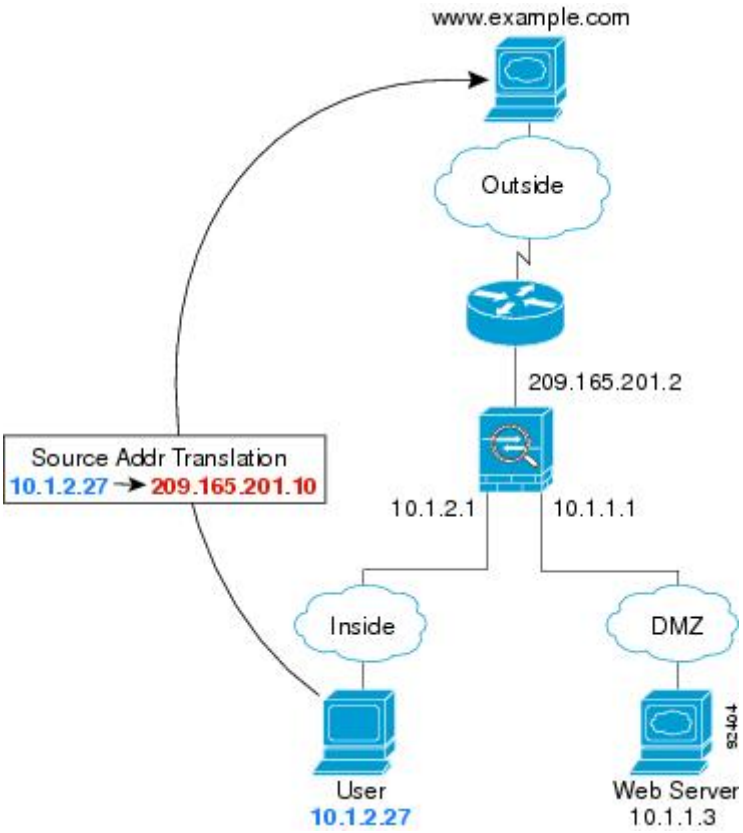
How Data Moves Through the ASA in Routed Firewall Mode

The following sections describe how data moves through the ASA in routed firewall mode in multiple scenarios.

An Inside User Visits a Web Server

The following figure shows an inside user accessing an outside web server.

Figure 36: Inside to Outside



The following steps describe how data moves through the ASA:

1. The user on the inside network requests a web page from www.example.com.
2. The ASA receives the packet and because it is a new session, it verifies that the packet is allowed according to the terms of the security policy.

   For multiple context mode, the ASA first classifies the packet to a context.

3. The ASA translates the real address (10.1.2.27) to the mapped address 209.165.201.10, which is on the outside interface subnet.

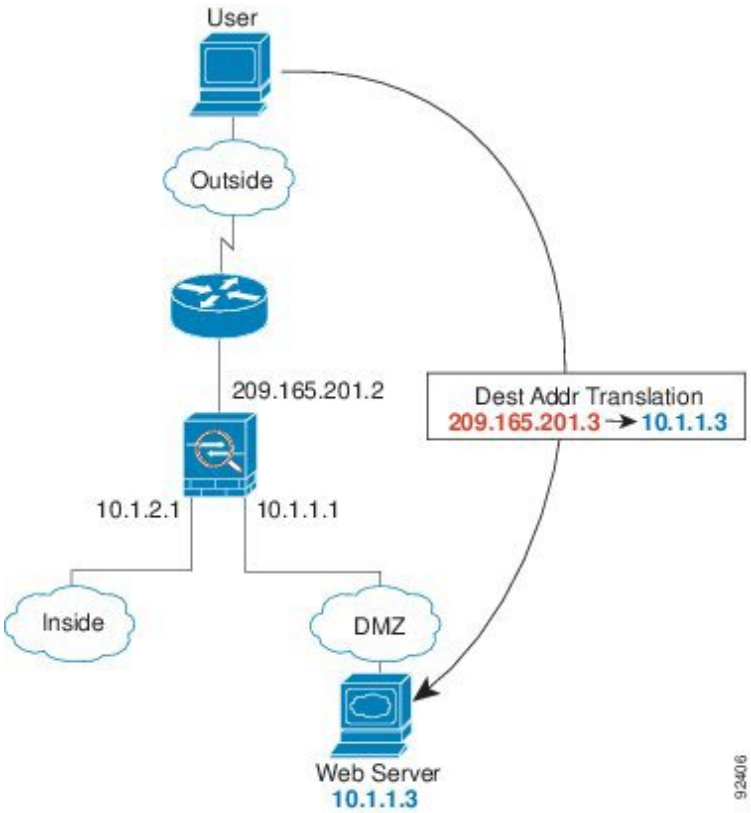
   The mapped address could be on any subnet, but routing is simplified when it is on the outside interface subnet.

4. The ASA then records that a session is established and forwards the packet from the outside interface.
5. When www.example.com responds to the request, the packet goes through the ASA, and because the session is already established, the packet bypasses the many lookups associated with a new connection. The ASA performs NAT by untranslating the global destination address to the local user address, 10.1.2.27.
6. The ASA forwards the packet to the inside user.

An Outside User Visits a Web Server on the DMZ

The following figure shows an outside user accessing the DMZ web server.

Figure 37: Outside to DMZ



The following steps describe how data moves through the ASA:

1. A user on the outside network requests a web page from the DMZ web server using the mapped address of 209.165.201.3, which is on the outside interface subnet.
2. The ASA receives the packet and untranslates the mapped address to the real address 10.1.1.3.

3. Because it is a new session, the ASA verifies that the packet is allowed according to the terms of the security policy.

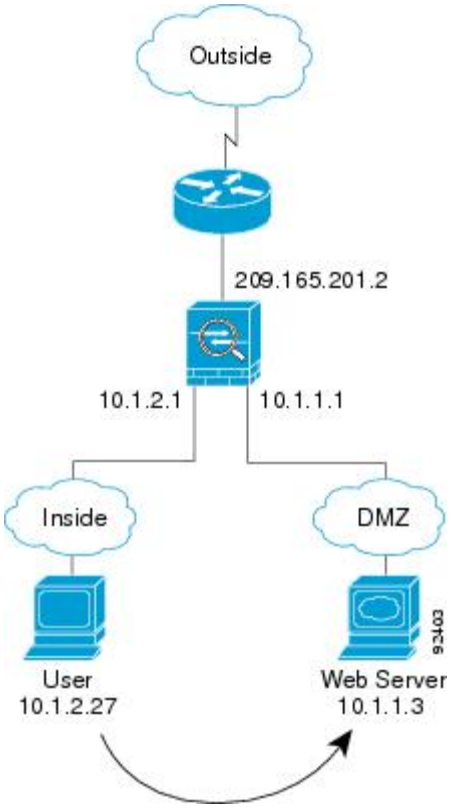
   For multiple context mode, the ASA first classifies the packet to a context.

4. The ASA then adds a session entry to the fast path and forwards the packet from the DMZ interface.
5. When the DMZ web server responds to the request, the packet goes through the ASA and because the session is already established, the packet bypasses the many lookups associated with a new connection. The ASA performs NAT by translating the real address to 209.165.201.3.
6. The ASA forwards the packet to the outside user.

An Inside User Visits a Web Server on the DMZ

The following figure shows an inside user accessing the DMZ web server.

Figure 38: Inside to DMZ



The following steps describe how data moves through the ASA:

1. A user on the inside network requests a web page from the DMZ web server using the destination address of 10.1.1.3.
2. The ASA receives the packet and because it is a new session, the ASA verifies that the packet is allowed according to the terms of the security policy.

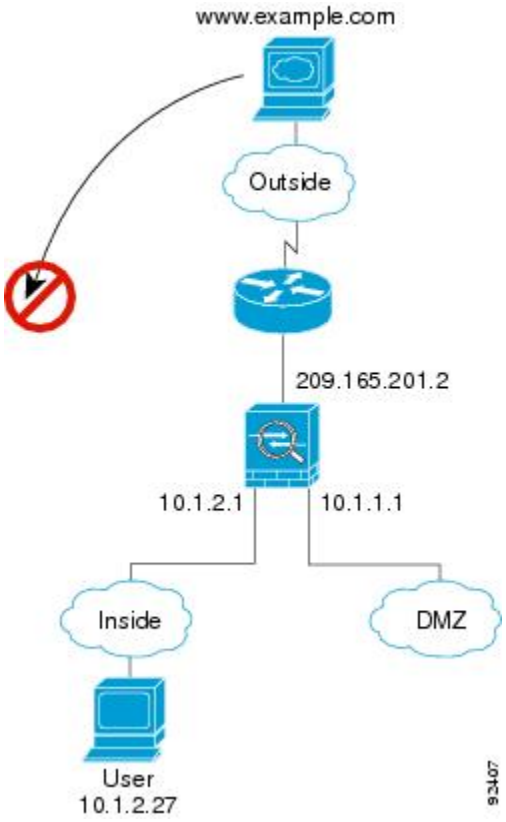
   For multiple context mode, the ASA first classifies the packet to a context.

3. The ASA then records that a session is established and forwards the packet out of the DMZ interface.
4. When the DMZ web server responds to the request, the packet goes through the fast path, which lets the packet bypass the many lookups associated with a new connection.
5. The ASA forwards the packet to the inside user.

An Outside User Attempts to Access an Inside Host

The following figure shows an outside user attempting to access the inside network.

Figure 39: Outside to Inside



The following steps describe how data moves through the ASA:

1. A user on the outside network attempts to reach an inside host (assuming the host has a routable IP address).

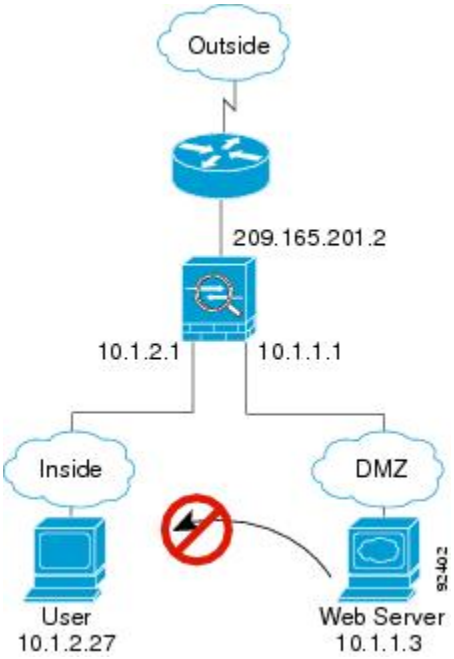
   If the inside network uses private addresses, no outside user can reach the inside network without NAT. The outside user might attempt to reach an inside user by using an existing NAT session.
2. The ASA receives the packet and because it is a new session, it verifies if the packet is allowed according to the security policy.
3. The packet is denied, and the ASA drops the packet and logs the connection attempt.

   If the outside user is attempting to attack the inside network, the ASA employs many technologies to determine if a packet is valid for an already established session.

A DMZ User Attempts to Access an Inside Host

The following figure shows a user in the DMZ attempting to access the inside network.

Figure 40: DMZ to Inside



The following steps describe how data moves through the ASA:

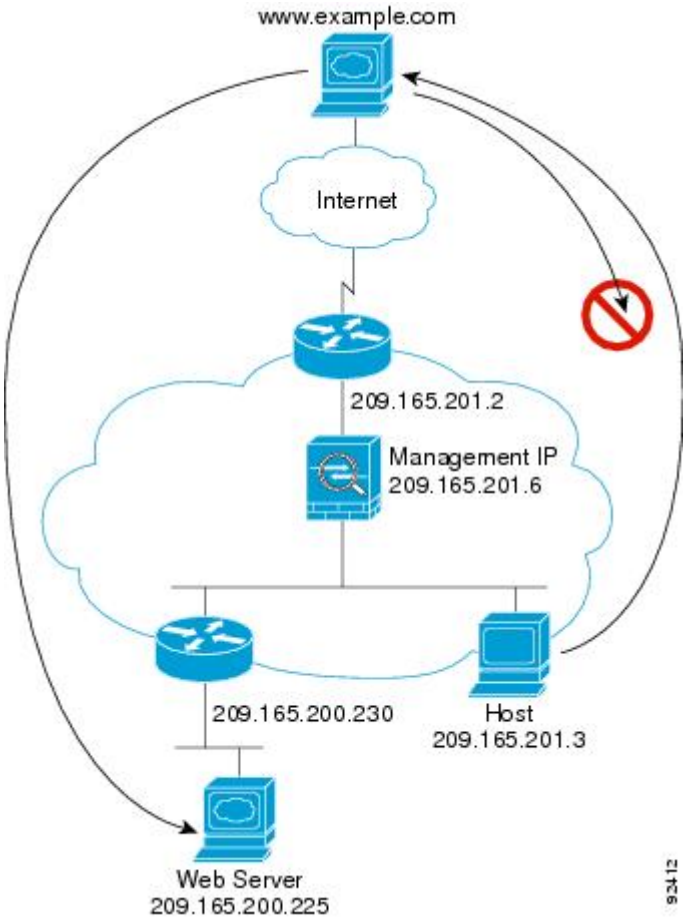
1. A user on the DMZ network attempts to reach an inside host. Because the DMZ does not have to route the traffic on the Internet, the private addressing scheme does not prevent routing.
2. The ASA receives the packet and because it is a new session, it verifies if the packet is allowed according to the security policy.

   The packet is denied, and the ASA drops the packet and logs the connection attempt.

How Data Moves Through the Transparent Firewall

The following figure shows a typical transparent firewall implementation with an inside network that contains a public web server. The ASA has an access rule so that the inside users can access Internet resources. Another access rule lets the outside users access only the web server on the inside network.

Figure 41: Typical Transparent Firewall Data Path

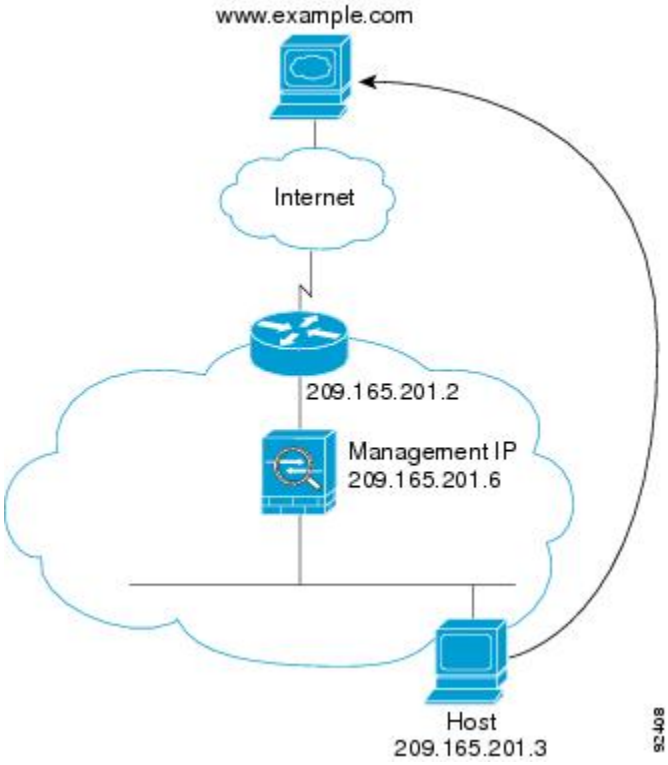


The following sections describe how data moves through the ASA.

An Inside User Visits a Web Server

The following figure shows an inside user accessing an outside web server.

Figure 42: Inside to Outside



The following steps describe how data moves through the ASA:

1. The user on the inside network requests a web page from www.example.com.
2. The ASA receives the packet and adds the source MAC address to the MAC address table, if required. Because it is a new session, it verifies that the packet is allowed according to the terms of the security policy.

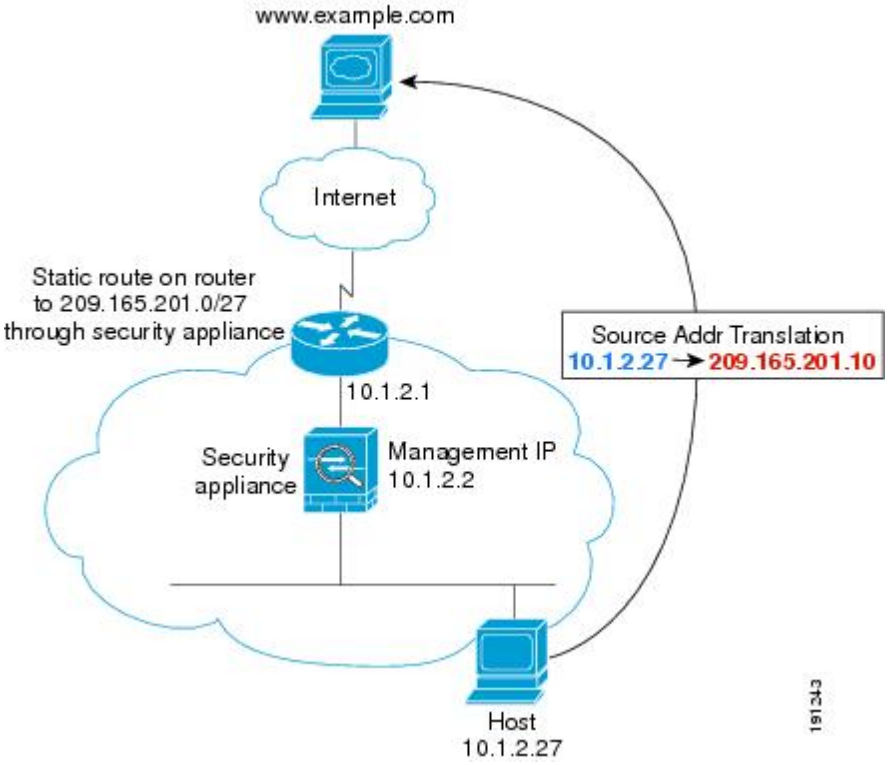
   For multiple context mode, the ASA first classifies the packet to a context.
3. The ASA records that a session is established.
4. If the destination MAC address is in its table, the ASA forwards the packet out of the outside interface. The destination MAC address is that of the upstream router, 209.165.201.2.

   If the destination MAC address is not in the ASA table, it attempts to discover the MAC address by sending an ARP request or a ping. The first packet is dropped.
5. The web server responds to the request; because the session is already established, the packet bypasses the many lookups associated with a new connection.
6. The ASA forwards the packet to the inside user.

An Inside User Visits a Web Server Using NAT

The following figure shows an inside user accessing an outside web server.

Figure 43: Inside to Outside with NAT



The following steps describe how data moves through the ASA:

1. The user on the inside network requests a web page from www.example.com.
2. The ASA receives the packet and adds the source MAC address to the MAC address table, if required. Because it is a new session, it verifies that the packet is allowed according to the terms of the security policy.

   For multiple context mode, the ASA first classifies the packet according to a unique interface.
3. The ASA translates the real address (10.1.2.27) to the mapped address 209.165.201.10.

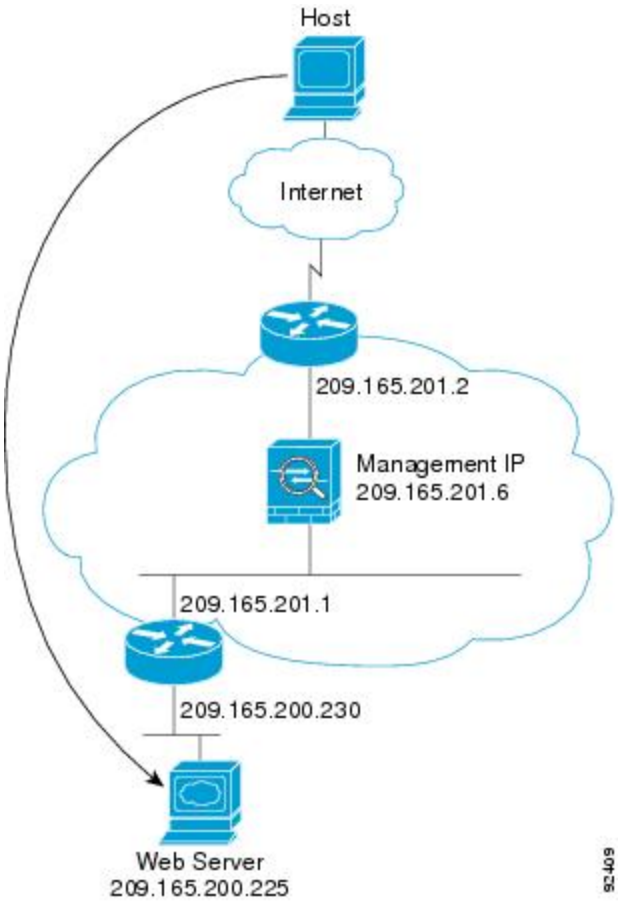
   Because the mapped address is not on the same network as the outside interface, then be sure the upstream router has a static route to the mapped network that points to the ASA.
4. The ASA then records that a session is established and forwards the packet from the outside interface.
5. If the destination MAC address is in its table, the ASA forwards the packet out of the outside interface. The destination MAC address is that of the upstream router, 10.1.2.1.

   If the destination MAC address is not in the ASA table, then it attempts to discover the MAC address by sending an ARP request and a ping. The first packet is dropped.
6. The web server responds to the request; because the session is already established, the packet bypasses the many lookups associated with a new connection.
7. The ASA performs NAT by untranslating the mapped address to the real address, 10.1.2.27.

An Outside User Visits a Web Server on the Inside Network

The following figure shows an outside user accessing the inside web server.

Figure 44: Outside to Inside



The following steps describe how data moves through the ASA:

1. A user on the outside network requests a web page from the inside web server.
2. The ASA receives the packet and adds the source MAC address to the MAC address table, if required. Because it is a new session, it verifies that the packet is allowed according to the terms of the security policy.

   For multiple context mode, the ASA first classifies the packet to a context.
3. The ASA records that a session is established.
4. If the destination MAC address is in its table, the ASA forwards the packet out of the inside interface. The destination MAC address is that of the downstream router, 209.165.201.1.

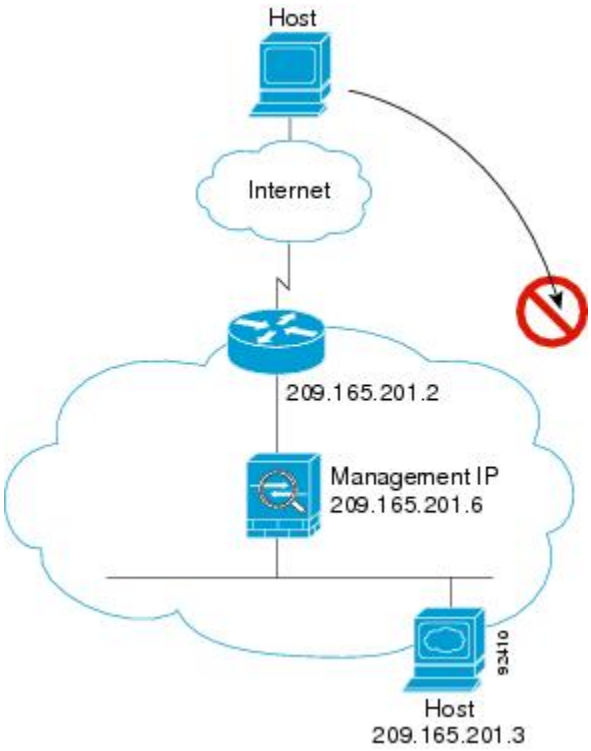
   If the destination MAC address is not in the ASA table, then it attempts to discover the MAC address by sending an ARP request and a ping. The first packet is dropped.
5. The web server responds to the request; because the session is already established, the packet bypasses the many lookups associated with a new connection.

6. The ASA forwards the packet to the outside user.

An Outside User Attempts to Access an Inside Host

The following figure shows an outside user attempting to access a host on the inside network.

Figure 45: Outside to Inside



The following steps describe how data moves through the ASA:

1. A user on the outside network attempts to reach an inside host.
2. The ASA receives the packet and adds the source MAC address to the MAC address table, if required. Because it is a new session, it verifies if the packet is allowed according to the terms of the security policy.

   For multiple context mode, the ASA first classifies the packet to a context.
3. The packet is denied because there is no access rule permitting the outside host, and the ASA drops the packet.
4. If the outside user is attempting to attack the inside network, the ASA employs many technologies to determine if a packet is valid for an already established session.

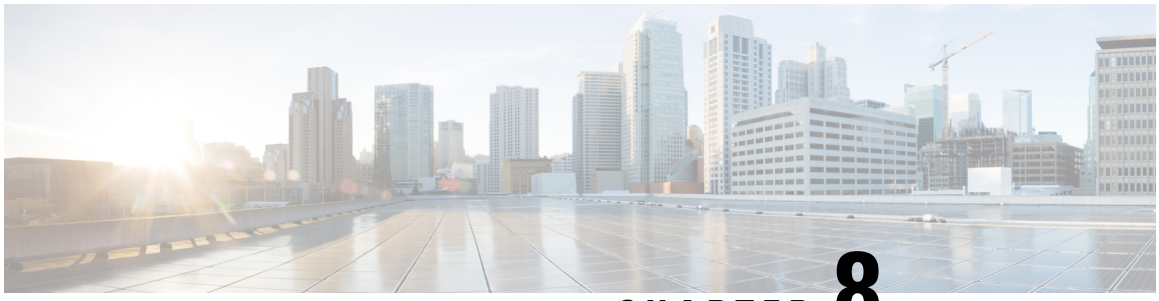
History for the Firewall Mode

Table 12: Feature History for Firewall Mode

Feature Name	Platform Releases	Feature Information
Transparent Firewall Mode	7.0(1)	A transparent firewall 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. We introduced the following commands: **firewall transparent**, **show firewall**. You cannot set the firewall mode in ASDM; you must use the command-line interface.
Transparent firewall bridge groups	8.4(1)	If you do not want the overhead of security contexts, or want to maximize your use of security contexts, you can group interfaces together in a bridge group, and then configure multiple bridge groups, one for each network. Bridge group traffic is isolated from other bridge groups. You can configure up to 8 bridge groups in single mode or per context in multiple mode, with 4 interfaces maximum per bridge group. **Note** Although you can configure multiple bridge groups on the ASA 5505, the restriction of 2 data interfaces in transparent mode on the ASA 5505 means you can only effectively use 1 bridge group. We modified or introduced the following screens: Configuration > Device Setup > Interface Settings > Interfaces Configuration > Device Setup > Interface Settings > Interfaces > Add/Edit Bridge Group Interface Configuration > Device Setup > Interface Settings > Interfaces > Add/Edit Interface
Mixed firewall mode support in multiple context mode	8.5(1)/9.0(1)	You can set the firewall mode independently for each security context in multiple context mode, so some can run in transparent mode while others run in routed mode. We modified the following command: **firewall transparent**. For single mode, you cannot set the firewall mode in ASDM; you must use the command-line interface. For multiple mode, we modified the following screen: Configuration > Context Management > Security Contexts.

Feature Name	Platform Releases	Feature Information
Transparent mode bridge group maximum increased to 250	9.3(1)	The bridge group maximum was increased from 8 to 250 bridge groups. You can configure up to 250 bridge groups in single mode or per context in multiple mode, with 4 interfaces maximum per bridge group. We modified the following screens: Configuration > Device Setup > Interface Settings > Interfaces Configuration > Device Setup > Interface Settings > Interfaces > Add/Edit Bridge Group Interface Configuration > Device Setup > Interface Settings > Interfaces > Add/Edit Interface
Transparent mode maximum interfaces per bridge group increased to 64	9.6(2)	The maximum interfaces per bridge group was increased from 4 to 64. We did not modify any screens.
Integrated Routing and Bridging	9.7(1)	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 ASA bridges instead of routes. The ASA is not a true bridge in that the ASA 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 ASA 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: multiple context mode, ASA clustering. The following features are also not supported on BVIs: dynamic routing and multicast routing. We modified the following screens: **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces** **Configuration** > **Device Setup** > **Routing** > **Static Routes** **Configuration** > **Device Management** > **DHCP** > **DHCP Server** **Configuration** > **Firewall** > **Access Rules** **Configuration** > **Firewall** > **EtherType Rules**

Feature Name	Platform Releases	Feature Information
Support for transparent mode deployment for a Firepower 4100/9300 ASA logical device	9.10(1)	You can now specify transparent or routed mode when you deploy the ASA on a Firepower 4100/9300. New/modified Firepower Chassis Manager screens: **Logical Devices** > **Add Device** > **Settings** New/Modified options: **Firewall Mode** drop-down list



CHAPTER 8

Startup Wizard

This chapter describes the ASDM Startup Wizard, which guides you through the initial configuration of the ASA and helps you define basic settings.

- Access the Startup Wizard, on page 239
- Guidelines for the Startup Wizard, on page 239
- Startup Wizard Screens, on page 239
- History for the Startup Wizard, on page 242

Access the Startup Wizard

To access the Startup Wizard, choose one of the following options:

- **Wizards > Startup Wizard**.
- **Configuration > Device Setup > Startup Wizard**, then click **Launch Startup Wizard**.

Guidelines for the Startup Wizard

Context Mode Guidelines

The Startup Wizard is not supported in the system context.

Startup Wizard Screens

The actual sequence of screens is determined by your specified configuration selections. Each screen is available for all modes or models unless otherwise noted.

Starting Point or Welcome

- Click the **Modify existing configuration** radio button to change the existing configuration.
- Click the **Reset configuration to factory defaults** radio button to set the configuration to the factory default values.

- Check the **Configure the IP address of the management interface** check box to configure the IP address and subnet mask of the Management 0/0 interface to be different from the default value (192.168.1.1).



Note If you reset the configuration to factory defaults, you cannot undo these changes by clicking **Cancel** or by closing this screen.

In multiple context mode, this screen does not include any parameters.

Basic Configuration

Set the hostname, domain name, and enable password in this screen.

Interface Screens

The interface screens depend on the mode and model selected.

Outside Interface Configuration (Routed Mode)

- Configure the IP address of the outside interface (the interface with the lowest security level).
- Configure the IPv6 address.

Outside Interface Configuration - PPPoE (Routed Mode, Single Mode)

Configure the PPPoE settings for the outside interface.

Management IP Address Configuration (Transparent Mode)

For IPv4, a management IP address is required for each bridge group for both management traffic and for traffic to pass through the ASA. This screen sets the IP address for BVI 1.

Other Interfaces Configuration

Configure parameters for other interfaces.

Static Routes

Configure static routes.

DHCP Server

Configure the DHCP server.

Address Translation (NAT/PAT)

Configures NAT or PAT for inside addresses (the interface with the highest security level) when accessing the outside (the interface with the lowest security level). See the firewall configuration guide for more information.

Administrative Access

- Configure ASDM, Telnet, or SSH access.
- Check the **Enable HTTP server for HTTPS/ASDM access** check box to enable a secure connection to an HTTP server to access ASDM.
- Check the **Enable ASDM history metrics** check box.

IPS Basic Configuration

In single context mode, use the Startup Wizard in ASDM to configure basic IPS network configuration. These settings are saved to the IPS configuration, not the ASA configuration. See the IPS quick start guide for more information.

ASA CX Basic Configuration (ASA 5585-X)

You can use the Startup Wizard in ASDM to configure the ASA CX management address and Auth Proxy Port. These settings are saved to the ASA CX configuration, not the ASA configuration. You will also need to set additional network settings at the ASA CX CLI. See the ASA CX quick start guide for information about this screen.

ASA FirePOWER Basic Configuration

You can use the Startup Wizard in ASDM to configure the ASA FirePOWER management address information and accept the end user license agreement (EULA). These settings are saved to the ASA FirePOWER configuration, not the ASA configuration. You will also need to configure some settings in the ASA FirePOWER CLI. For more information, see the chapter on the ASA FirePOWER module in the firewall configuration guide.

Time Zone and Clock Configuration

Configure the clock parameters.

Auto Update Server (Single Mode)

Follow these guidelines to configure an Auto-Update Server:

- Configure an auto update server by checking the **Enable Auto Update Server for ASA** check box.
- Check the **Enable Signature and Engine Updates from Cisco.com** check box if you have an IPS module. Set the following additional parameters:

- Enter your Cisco.com username and password, then confirm the password.
- Enter the start time in hh:mm:ss format, using a 24-hour clock.

Startup Wizard Summary

This screen summarizes all of the configuration settings that you have made for the ASA.

- Click **Back** to change any of the settings in previous screens.
- Choose one of the following:
  - If you ran the Startup Wizard directly from a browser, when you click **Finish**, the configuration settings that you created through the wizard are sent to the ASA and saved in flash memory automatically.
  - If you ran the Startup Wizard from within ASDM, you must explicitly save the configuration in flash memory by choosing **File > Save Running Configuration to Flash**.

History for the Startup Wizard

Table 13: History for the Startup Wizard

Feature Name	Platform Releases	Description
Startup Wizard	7.0(1)	This wizard was introduced. We introduced the **Wizards > Startup Wizard** screen.
ASA IPS Configuration	8.4(1)	For the ASA IPS module, the **IPS Basic Configuration** screen was added to the startup wizard. Signature updates for the IPS module were also added to the **Auto Update** screen. The **Time Zone and Clock Configuration** screen was added to ensure the clock is set on the ASA; the IPS module gets its clock from the ASA. We introduced or modified the following screens: **Wizards > Startup Wizard > IPS Basic Configuration** **Wizards > Startup Wizard > Auto Update** **Wizards > Startup Wizard > Time Zone and Clock Configuration**
ASA CX Configuration	9.1(1)	For the ASA CX module, the **ASA CX Basic Configuration** screen was added to the startup wizard. We introduced the following screens: **Wizards > Startup Wizard > ASA CX Basic Configuration**

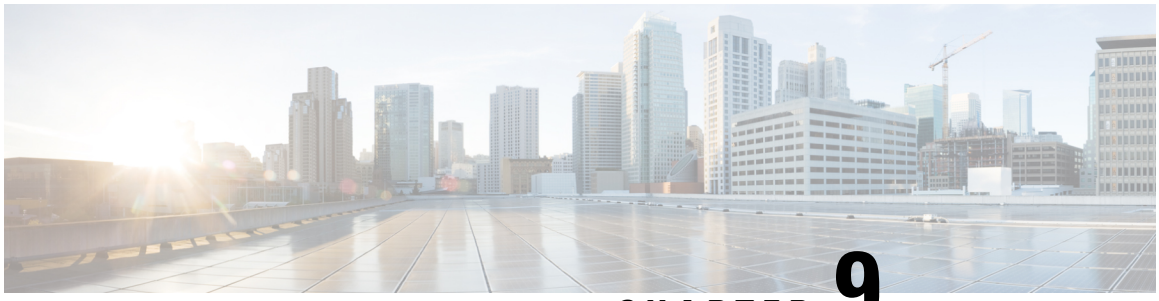
Feature Name	Platform Releases	Description
ASA FirePOWER Configuration	9.2(2.4)	For the ASA FirePOWER module, the **ASA FirePOWER Basic Configuration** screen was added to the startup wizard. We introduced the following screens: **Wizards > Startup Wizard > ASA FirePOWER Basic Configuration**



PART II

High Availability and Scalability

- Multiple Context Mode, on page 247
- Failover for High Availability, on page 281
- Failover for High Availability in the Public Cloud, on page 325
- ASA Cluster for the Secure Firewall 3100/4200, on page 341
- ASA Cluster for the Firepower 4100/9300, on page 429
- ASA Cluster for the ASA Virtual for the Private Cloud, on page 507



CHAPTER 9

Multiple Context Mode

This chapter describes how to configure multiple security contexts on the ASA.

- About Security Contexts, on page 247
- Licensing for Multiple Context Mode, on page 257
- Prerequisites for Multiple Context Mode, on page 258
- Guidelines for Multiple Context Mode, on page 259
- Defaults for Multiple Context Mode, on page 260
- Configure Multiple Contexts, on page 260
- Change Between Contexts and the System Execution Space, on page 269
- Manage Security Contexts, on page 269
- Monitoring Security Contexts, on page 273
- History for Multiple Context Mode, on page 276

About Security Contexts

You can partition a single ASA into multiple virtual devices, known as security contexts. Each context acts as an independent device, with its own security policy, interfaces, and administrators. Multiple contexts are similar to having multiple standalone devices. For unsupported features in multiple context mode, see Guidelines for Multiple Context Mode, on page 259.

This section provides an overview of security contexts.

Common Uses for Security Contexts

You might want to use multiple security contexts in the following situations:

- You are a service provider and want to sell security services to many customers. By enabling multiple security contexts on the ASA, you can implement a cost-effective, space-saving solution that keeps all customer traffic separate and secure, and also eases configuration.
- You are a large enterprise or a college campus and want to keep departments completely separate.
- You are an enterprise that wants to provide distinct security policies to different departments.
- You have any network that requires more than one ASA.

Context Configuration Files

This section describes how the ASA implements multiple context mode configurations.

Context Configurations

For each context, the ASA includes a configuration that identifies the security policy, interfaces, and all the options you can configure on a standalone device. You can store context configurations in flash memory, or you can download them from a TFTP, FTP, or HTTP(S) server.

System Configuration

The system administrator adds and manages contexts by configuring each context configuration location, allocated interfaces, and other context operating parameters in the system configuration, which, like a single mode configuration, is the startup configuration. The system configuration identifies basic settings for the ASA. The system configuration does not include any network interfaces or network settings for itself; rather, when the system needs to access network resources (such as downloading the contexts from the server), it uses one of the contexts that is designated as the *admin context*. The system configuration does include a specialized failover interface for failover traffic only.

Admin Context Configuration

The admin context is just like any other context, except that when a user logs in to the admin context, then that user has system administrator rights and can access the system and all other contexts. The admin context is not restricted in any way, and can be used as a regular context. However, because logging into the admin context grants you administrator privileges over all contexts, you might need to restrict access to the admin context to appropriate users. The admin context must reside on flash memory, and not remotely.

If your system is already in multiple context mode, or if you convert from single mode, the admin context is created automatically as a file on the internal flash memory called admin.cfg. This context is named “admin.” If you do not want to use admin.cfg as the admin context, you can change the admin context.

How the ASA Classifies Packets

Each packet that enters the ASA must be classified, so that the ASA can determine to which context to send a packet.



Note If the destination MAC address is a multicast or broadcast MAC address, the packet is duplicated and delivered to each context.

Valid Classifier Criteria

This section describes the criteria used by the classifier.



Note For management traffic destined for an interface, the interface IP address is used for classification.

The routing table is not used for packet classification.

Unique Interfaces

If only one context is associated with the ingress interface, the ASA classifies the packet into that context. In transparent firewall mode, unique interfaces for contexts are required, so this method is used to classify packets at all times.

Unique MAC Addresses

If multiple contexts share an interface, then the classifier uses unique MAC addresses assigned to the interface in each context. An upstream router cannot route directly to a context without unique MAC addresses. You can enable auto-generation of MAC addresses. You can also set the MAC addresses manually when you configure each interface.

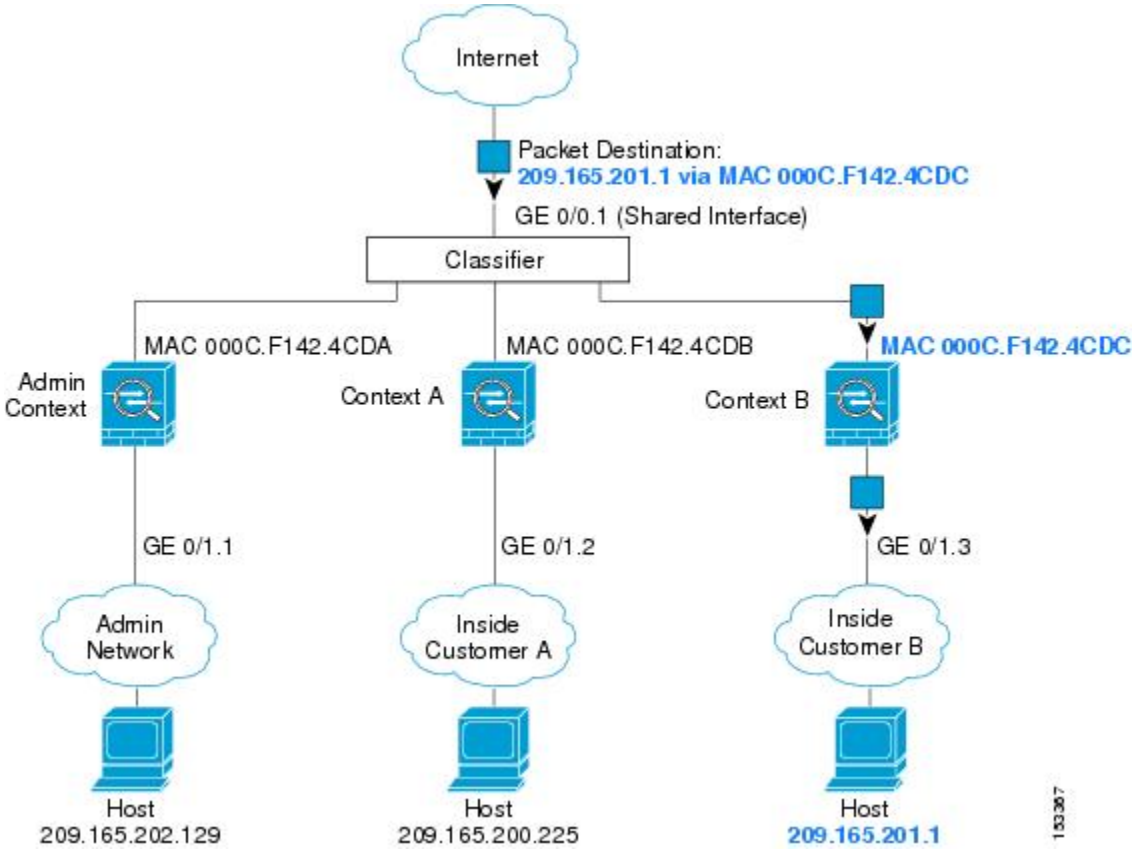
NAT Configuration

If you do not enable use of unique MAC addresses, then the ASA uses the mapped addresses in your NAT configuration to classify packets. We recommend using MAC addresses instead of NAT, so that traffic classification can occur regardless of the completeness of the NAT configuration.

Classification Examples

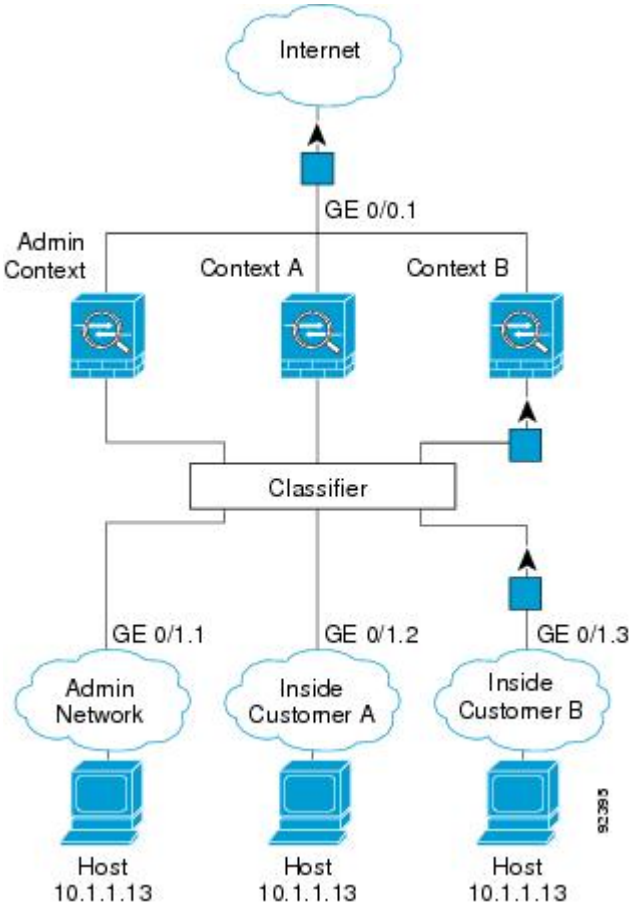
The following figure shows multiple contexts sharing an outside interface. The classifier assigns the packet to Context B because Context B includes the MAC address to which the router sends the packet.

Figure 46: Packet Classification with a Shared Interface Using MAC Addresses



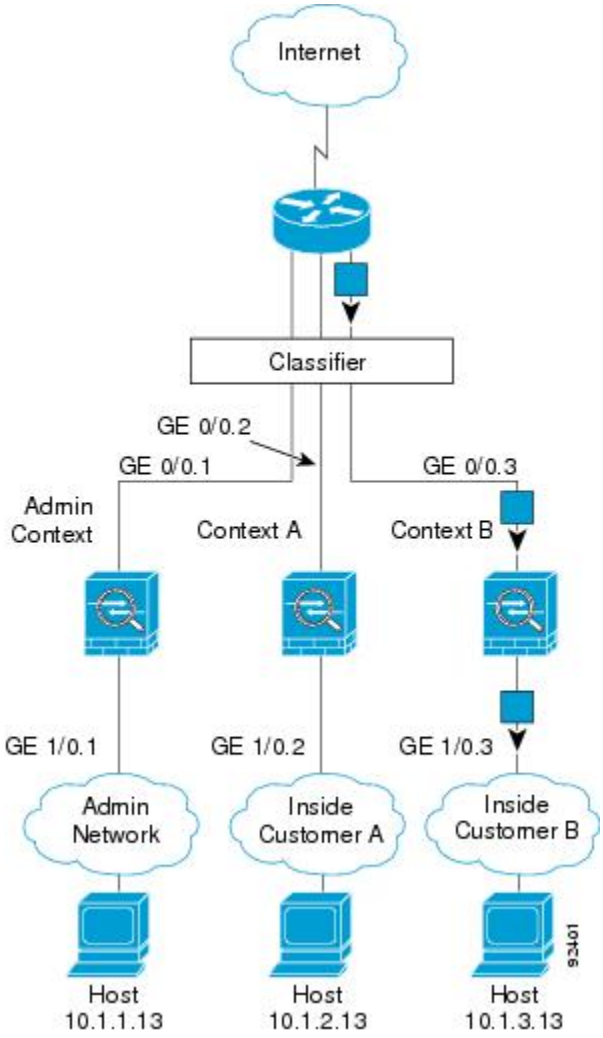
Note that all new incoming traffic must be classified, even from inside networks. The following figure shows a host on the Context B inside network accessing the Internet. The classifier assigns the packet to Context B because the ingress interface is Gigabit Ethernet 0/1.3, which is assigned to Context B.

Figure 47: Incoming Traffic from Inside Networks



For transparent firewalls, you must use unique interfaces. The following figure shows a packet destined to a host on the Context B inside network from the Internet. The classifier assigns the packet to Context B because the ingress interface is Gigabit Ethernet 1/0.3, which is assigned to Context B.

Figure 48: Transparent Firewall Contexts



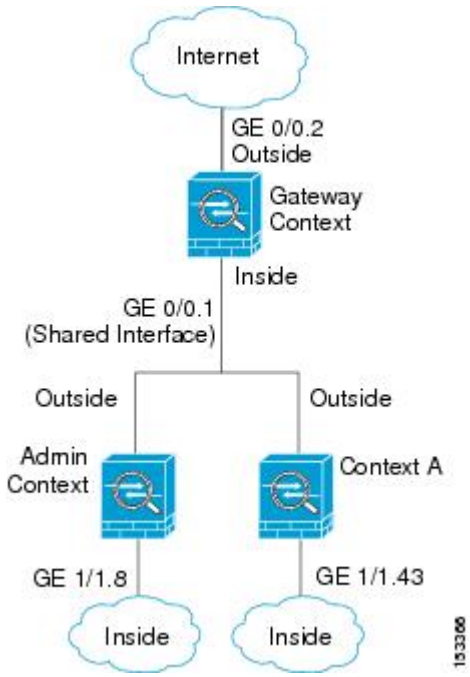
Cascading Security Contexts

Placing a context directly in front of another context is called *cascading contexts*; the outside interface of one context is the same interface as the inside interface of another context. You might want to cascade contexts if you want to simplify the configuration of some contexts by configuring shared parameters in the top context.

Note Cascading contexts requires unique MAC addresses for each context interface. Because of the limitations of classifying packets on shared interfaces without MAC addresses, we do not recommend using cascading contexts without unique MAC addresses.

The following figure shows a gateway context with two contexts behind the gateway.

Figure 49: Cascading Contexts



Management Access to Security Contexts

The ASA provides system administrator access in multiple context mode as well as access for individual context administrators.

System Administrator Access

You can access the ASA as a system administrator in two ways:

- Access the ASA console.

  From the console, you access the *system execution space*, which means that any commands you enter affect only the system configuration or the running of the system (for run-time commands).

- Access the admin context using Telnet, SSH, or ASDM.

As the system administrator, you can access all contexts.

The system execution space does not support any AAA commands, but you can configure its own enable password, as well as usernames in the local database to provide individual logins.

Context Administrator Access

You can access a context using Telnet, SSH, or ASDM. If you log in to a non-admin context, you can only access the configuration for that context. You can provide individual logins to the context.

Management Interface Usage

The Management interface is a separate interface just for management traffic.

In routed firewall mode, you can share the Management interface across all contexts.

In transparent firewall mode, the Management interface is special. In addition to the maximum allowed through-traffic interfaces, you can also use the Management interface as a separate management-only interface. However, in multiple context mode, you cannot share any interfaces across transparent contexts. You can instead use subinterfaces of the Management interface, and assign one to each context. However, only Firepower device models allow subinterfaces on the Management interface. For ASA models, you must use a data interface or a subinterface of a data interface, and add it to a bridge group within the context.

For the Firepower 4100/9300 chassis transparent context, neither the Management interface nor subinterface retains its special status. In this case, you must treat it as a data interface, and add it to a bridge group. (Note that in single context mode, the Management interface does retain its special status.)

Another consideration about transparent mode: when you enable multiple context mode, all configured interfaces are automatically assigned to the Admin context. For example, if your default configuration includes the Management interface, then that interface will be assigned to the Admin context. One option is to leave the main interface allocated to the Admin context and manage it using the native VLAN, and then use subinterfaces to manage each context. Keep in mind that if you make the Admin context transparent, its IP address will be removed; you have to assign it to a bridge group and assign the IP address to the BVI.

About Resource Management

By default, all security contexts have unlimited access to the resources of the ASA, except where maximum limits per context are enforced; the only exception is VPN resources, which are disabled by default. If you find that one or more contexts use too many resources, and they cause other contexts to be denied connections, for example, then you can configure resource management to limit the use of resources per context. For VPN resources, you must configure resource management to allow any VPN tunnels.

Resource Classes

The ASA manages resources by assigning contexts to resource classes. Each context uses the resource limits set by the class. To use the settings of a class, assign the context to the class when you define the context. All contexts belong to the default class if they are not assigned to another class; you do not have to actively assign a context to default. You can only assign a context to one resource class. The exception to this rule is that limits that are undefined in the member class are inherited from the default class; so in effect, a context could be a member of default plus another class.

Resource Limits

You can set the limit for individual resources as a percentage (if there is a hard system limit) or as an absolute value.

For most resources, the ASA does not set aside a portion of the resources for each context assigned to the class; rather, the ASA sets the maximum limit for a context. If you oversubscribe resources, or allow some resources to be unlimited, a few contexts can “use up” those resources, potentially affecting service to other contexts. The exception is VPN resource types, which you cannot oversubscribe, so the resources assigned to each context are guaranteed. To accommodate temporary bursts of VPN sessions beyond the amount assigned, the ASA supports a “burst” VPN resource type, which is equal to the remaining unassigned VPN sessions. The burst sessions *can* be oversubscribed, and are available to contexts on a first-come, first-served basis.

Default Class

All contexts belong to the default class if they are not assigned to another class; you do not have to actively assign a context to the default class.

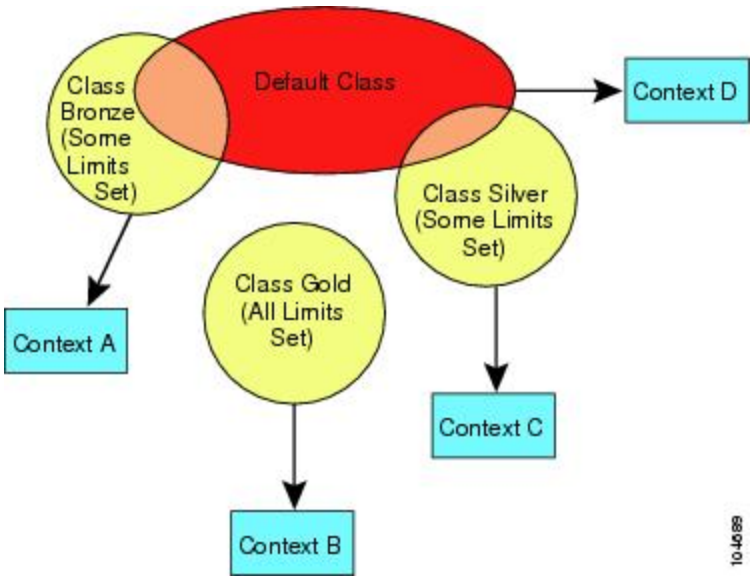
If a context belongs to a class other than the default class, those class settings always override the default class settings. However, if the other class has any settings that are not defined, then the member context uses the default class for those limits. For example, if you create a class with a 2 percent limit for all concurrent connections, but no other limits, then all other limits are inherited from the default class. Conversely, if you create a class with a limit for all resources, the class uses no settings from the default class.

For most resources, the default class provides unlimited access to resources for all contexts, except for the following limits:

- Telnet sessions—5 sessions. (The maximum per context.)
- SSH sessions—5 sessions. (The maximum per context.)
- ASDM sessions—5 sessions. (The maximum per context.)
- MAC addresses—(varies per model). (The maximum for the system.)
- Secure Client peers—0 sessions. (You must manually configure the class to allow any Secure Client peers.)
- VPN site-to-site tunnels—0 sessions. (You must manually configure the class to allow any VPN sessions.)

The following figure shows the relationship between the default class and other classes. Contexts A and C belong to classes with some limits set; other limits are inherited from the default class. Context B inherits no limits from default because all limits are set in its class, the Gold class. Context D was not assigned to a class, and is by default a member of the default class.

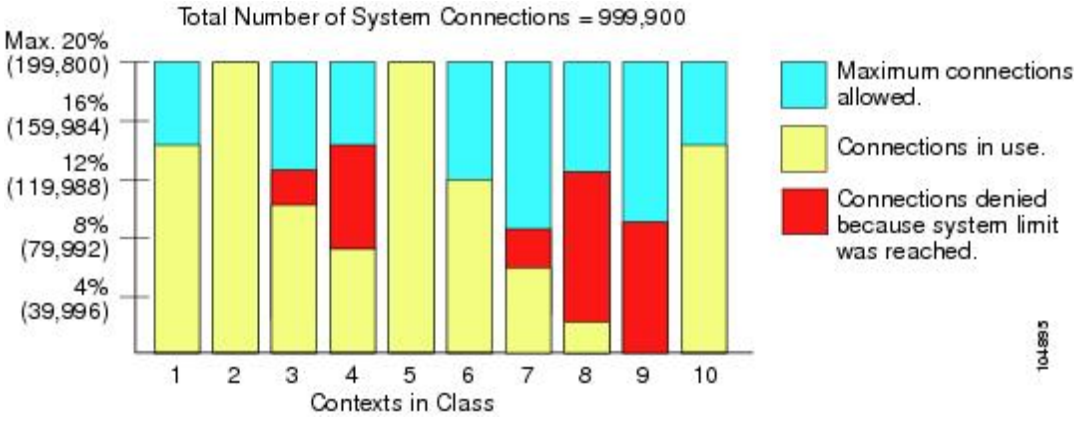
Figure 50: Resource Classes



Use Oversubscribed Resources

You can oversubscribe the ASA by assigning more than 100 percent of a resource across all contexts (with the exception of non-burst VPN resources). For example, you can set the Bronze class to limit connections to 20 percent per context, and then assign 10 contexts to the class for a total of 200 percent. If contexts concurrently use more than the system limit, then each context gets less than the 20 percent you intended.

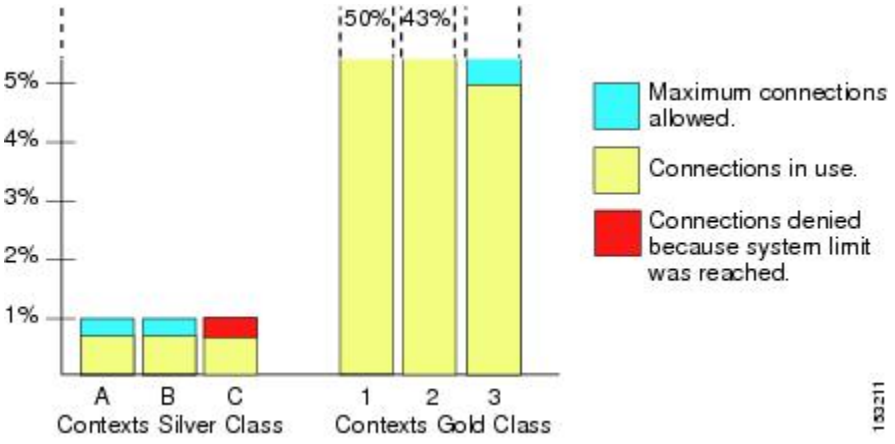
Figure 51: Resource Oversubscription



Use Unlimited Resources

The ASA lets you assign unlimited access to one or more resources in a class, instead of a percentage or absolute number. When a resource is unlimited, contexts can use as much of the resource as the system has available. For example, Context A, B, and C are in the Silver Class, which limits each class member to 1 percent of the connections, for a total of 3 percent; but the three contexts are currently only using 2 percent combined. Gold Class has unlimited access to connections. The contexts in the Gold Class can use more than the 97 percent of "unassigned" connections; they can also use the 1 percent of connections not currently in use by Context A, B, and C, even if that means that Context A, B, and C are unable to reach their 3 percent combined limit. Setting unlimited access is similar to oversubscribing the ASA, except that you have less control over how much you oversubscribe the system.

Figure 52: Unlimited Resources



About MAC Addresses

You can manually assign MAC addresses to override the default. For multiple context mode, you can automatically generate unique MAC addresses (for all interfaces assigned to a context) and single context mode (for subinterfaces)..

Note You might want to assign unique MAC addresses to subinterfaces defined on the ASA, 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 ASA device.

MAC Addresses in Multiple Context Mode

The MAC address is used to classify packets within a context. If you share an interface, but do not have unique MAC addresses for the interface in each context, then other classification methods are attempted that might not provide full coverage.

To allow contexts to share interfaces, you should enable auto-generation of virtual MAC addresses to each shared context interface.

Automatic MAC Addresses

In multiple context mode, auto-generation assigns unique MAC addresses to all interfaces assigned to a context.

If you manually assign a MAC address and also enable auto-generation, then the manually assigned MAC address is used. If you later remove the manual MAC address, the auto-generated address is used, if enabled.

In the rare circumstance that the generated MAC address conflicts with another private MAC address in your network, you can manually set the MAC address for the interface.

Because auto-generated addresses (when using a prefix) start with A2, you cannot start manual MAC addresses with A2 if you also want to use auto-generation.

The ASA generates the MAC address using the following format:

A2*xx.yyzz.zzzz*

Where *xx.yy* is a user-defined prefix or an autogenerated prefix based on the last two bytes of the interface MAC address, and *zz.zzzz* is an internal counter generated by the ASA. For the standby MAC address, the address is identical except that the internal counter is increased by 1.

For an example of how the prefix is used, if you set a prefix of 77, then the ASA converts 77 into the hexadecimal value 004D (*yyxx*). When used in the MAC address, the prefix is reversed (*xxyy*) to match the ASA native form:

A2**4D.00***zz.zzzz*

For a prefix of 1009 (03F1), the MAC address is:

A2**F1.03***zz.zzzz*

Note The MAC address format without a prefix is a legacy version. See the **mac-address auto** command in the command reference for more information about the legacy format.

VPN Support

For VPN resources, you must configure resource management to allow any VPN tunnels.

You can use site-to-site VPN in multiple context mode.

For remote access VPN, you must use AnyConnect 3.x and later for SSL VPN and IKEv2 protocol. You can customize flash storage per context for Secure Client images and customizations, as well as using shared flash memory across all contexts. For unsupported features, see Guidelines for Multiple Context Mode, on page 259. For a detailed list of supported VPN features per ASA release, see History for Multiple Context Mode, on page 276.

Note The Secure Client Premier license is required for multiple context mode; you cannot use the default or legacy license.

Licensing for Multiple Context Mode

Model	License Requirement
Firepower 1010	No support.
Firepower 1100	Essentials License: 2 contexts. *Optional License, Maximum:* *Firepower 1120: 5* *Firepower 1140: 10* *Firepower 1150: 25*
Secure Firewall 1210/1220	Essentials License: 2 contexts. *Optional License, Maximum:* *Secure Firewall 1210: 5* *Secure Firewall 1220: 10*
Secure Firewall 1230/1240/1250	Essentials License: 2 contexts. *Optional License, Maximum:* *Secure Firewall 1230: 25* *Secure Firewall 1240: 25* *Secure Firewall 1250: 25*

Model	License Requirement
Secure Firewall 3100	Essentials License: 2 contexts. *Optional License, Maximum:* *Secure Firewall 3105: 100* *Secure Firewall 3110: 100* *Secure Firewall 3120: 100* *Secure Firewall 3130: 100* *Secure Firewall 3140: 100*
Firepower 4100	Essentials License: 10 contexts. *Optional License: up to 250 contexts.*
Secure Firewall 4200	Essentials License: 2 contexts. *Optional License, Maximum:* *Secure Firewall 4215: 250* *Secure Firewall 4225: 250* *Secure Firewall 4245: 250*
Firepower 9300	Essentials License: 10 contexts. *Optional License: up to 250 contexts.*
ISA 3000	No support.
ASA Virtual	No support.

Note If the Admin context only contains management-only interfaces, and does not include any data interfaces for through traffic, then it does not count against the limit.

Note The Secure Client Premier license is required for multiple context mode; you cannot use the default or legacy license.

Prerequisites for Multiple Context Mode

After you are in multiple context mode, connect tothe admin context to access the system configuration. You cannot configure the system from a non-admin context. By default, after you enable multiple context mode, you can connect to the admin context by using the default management IP address.

Guidelines for Multiple Context Mode

Failover

Active/Active mode failover is only supported in multiple context mode.

IPv6

Cross-context IPv6 routing is not supported.

Unsupported Features

Multiple context mode does not support the following features:

- RIP
- OSPFv3. (OSPFv2 is supported.)
- Multicast routing
- Threat Detection
- Unified Communications
- QoS
- Virtual Tunnel Interfaces (VTIs)
- Static route tracking
- Switchports for the Secure Firewall 1210/1220.

Multiple context mode does not currently support the following features for remote access VPN:

- AnyConnect 2.x and earlier
- IKEv1
- SAML
- WebLaunch
- VLAN Mapping
- HostScan
- VPN load balancing
- Customization
- L2TP

Additional Guidelines

- The context mode (single or multiple) is not stored in the configuration file, even though it does endure reboots. If you need to copy your configuration to another device, set the mode on the new device to match.

- By default, Secure Firewall 1210/1220devices are in single context mode. As switchports are not supported in multiple context mode, you must convert all interfaces to router interfaces before you can convert to multiple context mode.
- If you store context configurations in the root directory of flash memory, on some models you might run out of room in that directory, even though there is available memory. In this case, create a subdirectory for your configuration files. Background: some models use the FAT 16 file system for internal flash memory, and if you do not use 8.3-compliant short names, or use uppercase characters, then fewer than 512 files and folders can be stored because the file system uses up slots to store long file names (see http://support.microsoft.com/kb/120138/en-us).
- In ACI, Policy Based Redirect (PBR) health check is executed (L2 pings) using the same MAC address on all the leaves. This results in a MAC flap. To resolve the MAC flap, you can configure the tap-mode option on the inline set. However, if the threat defense high availability is configured, you must enable MAC learning for connection handling during a failover. Thus, in an ACI environment with threat defense in HA pair using inline-set interfaces, to avoid packet losses, deploy the threat defense in standalone or in a cluster.

Defaults for Multiple Context Mode

- By default, the ASA is in single context mode.
- See Default Class, on page 254.

Configure Multiple Contexts

Procedure

Step 1 Enable or Disable Multiple Context Mode, on page 261.

Step 2 (Optional) Configure a Class for Resource Management, on page 263.

Note
For VPN support, you must configure VPN resources in a resource class; the default class does not allow VPN.

Step 3 Configure interfaces in the system execution space.

- Firepower 1100, Secure Firewall 3100/4200—Basic Interface Configuration, on page 561.
- Secure Firewall 1200—Basic Interface Configuration, on page 561 and Basic Interface Configuration for Firepower 1010 and Secure Firewall 1210/1220 Switch Ports, on page 575.
- Firepower 4100/9300—Logical Devices for the Firepower 4100/9300, on page 197

Step 4 Configure a Security Context, on page 266.

Step 5 (Optional) Assign MAC Addresses to Context Interfaces Automatically, on page 268.

Step 6 Complete interface configuration in the context. See Routed and Transparent Mode Interfaces, on page 633.

Enable or Disable Multiple Context Mode

Your ASA might already be configured for multiple security contexts depending on how you ordered it from Cisco. If you need to convert from single mode to multiple mode, follow the procedures in this section.

ASDM supports changing modes from single to multiple mode if you use the High Availability and Scalability Wizard and you enable Active/Active failover. See Failover for High Availability, on page 281 for more information. If you do not want to use Active/Active failover or want to change back to single mode, you must change modes using the CLI; because changing modes requires confirmation, you cannot use the Command Line Interface tool. This section describes changing modes at the CLI.

Enable Multiple Context Mode

When you convert from single mode to multiple mode, the ASA converts the running configuration into two files: a new startup configuration that comprises the system configuration, and admin.cfg that comprises the admin context (in the root directory of the internal flash memory). The original running configuration is saved as old_running.cfg (in the root directory of the internal flash memory). The original startup configuration is not saved. The ASA automatically adds an entry for the admin context to the system configuration with the name “admin.”

Before you begin

Back up your startup configuration if it differs from the running configuration. When you convert from single mode to multiple mode, the ASA converts the running configuration into two files. The original startup configuration is not saved. See Manage Files, on page 1077.

Procedure

Change to multiple context mode.

mode multiple

Example:

You are prompted to change the mode and convert the configuration, and then the system reloads.

Note
You will have to regenerate the RSA key pair in the Admin context before you can reestablish an SSH connection. From the console, enter the **crypto key generate rsa modulus** command. See Configure SSH Access, on page 1029 for more information.

Example:

```
ciscoasa(config)# mode multiple
WARNING: This command will change the behavior of the device
WARNING: This command will initiate a Reboot
Proceed with change mode? [confirm]
Convert the system configuration? [confirm]
!
The old running configuration file will be written to flash
```

```
Converting the configuration - this may take several minutes for a large configuration

The admin context configuration will be written to flash

The new running configuration file was written to flash
Security context mode: multiple
ciscoasa(config)#


***
*** --- START GRACEFUL SHUTDOWN ---
***
*** Message to all terminals:
***
***   change mode
Shutting down isakmp
Shutting down webvpn
Shutting down License Controller
Shutting down File system



***
*** --- SHUTDOWN NOW ---
***
*** Message to all terminals:
***
***   change mode
```

Restore Single Context Mode

To copy the old running configuration to the startup configuration and to change the mode to single mode, perform the following steps.

Before you begin

Perform this procedure in the system execution space.

Procedure

Step 1 Copy the backup version of your original running configuration to the current startup configuration:

copy disk0:old_running.cfg startup-config

Example:

```
ciscoasa(config)# copy disk0:old_running.cfg startup-config
```

Step 2 Set the mode to single mode:

mode single

Example:

```
ciscoasa(config)# mode single
```

You are prompted to reboot the ASA.

Configure a Class for Resource Management

To configure a class in the system configuration, perform the following steps. You can change the value of a particular resource limit by reentering the command with a new value.

Before you begin

- Perform this procedure in the system execution space.
- The following table lists the resource types and the limits.

Note If the System Limit is N/A, then you cannot set a percentage of the resource because there is no hard system limit for the resource.

Table 14: Resource Names and Limits

Resource Name	Rate or Concurrent	Minimum and Maximum Number per Context	System Limit	Description
ASDM Sessions	Concurrent	1 minimum 5 maximum	200	ASDM management sessions. ASDM sessions use two HTTPS connections: one for monitoring that is always present, and one for making configuration changes that is present only when you make changes. For example, the system limit of 200 ASDM sessions represents a limit of 400 HTTPS sessions.

Resource Name	Rate or Concurrent	Minimum and Maximum Number per Context	System Limit	Description
Connections Conns/sec	Concurrent or Rate	N/A	Concurrent connections: See Licenses Per Model, on page 171 for the connection limit available for your model. Rate: N/A	TCP or UDP connections between any two hosts, including connections between one host and multiple other hosts. **Note** Syslog messages are generated for whichever limit is lower, xlates or conns. For example, if you set the xlates limit to 7 and the conns to 9, then the ASA only generates syslog message 321001 ("Resource 'xlates' limit of 7 reached for context 'ctx1'") and not 321002 ("Resource 'conn rate' limit of 5 reached for context 'ctx1'").
Hosts	Concurrent	N/A	N/A	Hosts that can connect through the ASA.
Inspects/sec	Rate	N/A	N/A	Application inspections per second.
MAC Entries	Concurrent	N/A	(varies per model)	For transparent firewall mode, the number of MAC addresses allowed in the MAC address table.
Routes	Concurrent	N/A	N/A	Dynamic routes.
Secure Client Burst	Concurrent	N/A	The Secure Client Premium Peers for your model minus the sum of the sessions assigned to all contexts for Secure Client.	The number of Secure Client sessions allowed beyond the amount assigned to a context with Secure Client. For example, if your model supports 5000 peers, and you assign 4000 peers across all contexts with Secure Client, then the remaining 1000 sessions are available for Secure Client Burst. Unlike Secure Client, which guarantees the sessions to the context, Secure Client Burst can be oversubscribed; the burst pool is available to all contexts on a first-come, first-served basis.
Secure Client	Concurrent	N/A	See Licenses Per Model, on page 171 for the Secure Client Premium Peers available for your model.	Secure Client peers. You cannot oversubscribe this resource; all context assignments combined cannot exceed the model limit. The peers you assign for this resource are guaranteed to the context.

Resource Name	Rate or Concurrent	Minimum and Maximum Number per Context	System Limit	Description
Other VPN Burst	Concurrent	N/A	The Other VPN session amount for your model minus the sum of the sessions assigned to all contexts for Other VPN.	The number of site-to-site VPN sessions allowed beyond the amount assigned to a context with Other VPN. For example, if your model supports 5000 sessions, and you assign 4000 sessions across all contexts with Other VPN, then the remaining 1000 sessions are available for Other VPN Burst. Unlike Other VPN, which guarantees the sessions to the context, Other VPN Burst can be oversubscribed; the burst pool is available to all contexts on a first-come, first-served basis.
Other VPN	Concurrent	N/A	See Supported Feature Licenses Per Model, on page 104 for the Other VPN sessions available for your model.	Site-to-site VPN sessions. You cannot oversubscribe this resource; all context assignments combined cannot exceed the model limit. The sessions you assign for this resource are guaranteed to the context.
IKEv1 SAs In Negotiation	Concurrent (percentage only)	N/A	A percentage of the Other VPN sessions assigned to this context. See the Other VPN resources to assign sessions to the context.	Incoming IKEv1 SA negotiations, as a percentage of the context Other VPN limit.
SSH	Concurrent	1 minimum 5 maximum	100	SSH sessions.
Storage	MB	The maximum depends on your specified flash memory drive	The maximum depends on your specified flash memory drive	Storage limit of context directory in MB.
Syslogs/sec	Rate	N/A	N/A	Syslog messages per second.
Telnet	Concurrent	1 minimum 5 maximum	100	Telnet sessions.
Xlates	Concurrent	N/A	N/A	Network address translations.

Procedure

Step 1 If you are not already in the System configuration mode, in the **Device List** pane, double-click **System** under the active device IP address.

Step 2 Choose **Configuration** > **Context Management** > **Resource Class**, and click **Add**.

The **Add Resource Class** dialog box appears.

Step 3 Enter a class name up to 20 characters in length, in the **Resource Class** field.

Step 4 In the **Count Limited Resources** area, set the concurrent limits for resources,.

See the preceding table for a description of each resource type.

For resources that do not have a system limit, you cannot set the percentage; you can only set an absolute value. If you do not set a limit, the limit is inherited from the default class. If the default class does not set a limit, then the resource is unlimited, or the system limit if available. For most resources, 0 sets the limit to unlimited. For VPN types, 0 sets the limit none.

Note
If you also set the **Configuration** > **Device Management** > **Management Access** > **Management Session Quota** within a context to set the maximum administrative sessions (SSH, etc.), then the lower value will be used.

Step 5 In the **Rate Limited Resources** area, set the rate limit for resources.

See the preceding table for a description of each resource type.

If you do not set a limit, the limit is inherited from the default class. If the default class does not set a limit, then it is unlimited by default. 0 sets the limit to unlimited.

Step 6 Click **OK**.

Configure a Security Context

The security context definition in the system configuration identifies the context name, configuration file URL, interfaces that a context can use, and other settings.

Before you begin

- Perform this procedure in the system execution space.
- Configure interfaces. For transparent mode contexts, you cannot share interfaces between contexts, so you might want to use subinterfaces. To plan for Management interface usage, see Management Interface Usage, on page 252.
  - Firepower 1100, Secure Firewall 3100/4200—Basic Interface Configuration, on page 561.
  - Firepower 4100/9300—Logical Devices for the Firepower 4100/9300, on page 197

Procedure

Step 1 If you are not already in the System configuration mode, in the **Device List** pane, double-click **System** under the active device IP address.

Step 2 Choose **Configuration** > **Context Management** > **Security Contexts**, and click **Add**.

The **Add Context** dialog box appears.

Step 3 In the **Security Context** field, enter the context name as a string up to 32 characters long.

This name is case sensitive, so you can have two contexts named "customerA" and "CustomerA," for example. "System" or "Null" (in upper or lower case letters) are reserved names, and cannot be used.

Step 4 In the **Interface Allocation** area, click the **Add** button to assign an interface to the context.

a) From the **Interfaces** > **Physical Interface** drop-down list, choose an interface.

You can assign the main interface, in which case you leave the subinterface ID blank, or you can assign a subinterface or a range of subinterfaces associated with this interface. In transparent firewall mode, only interfaces that have not been allocated to other contexts are shown. If the main interface was already assigned to another context, then you must choose a subinterface.

b) (Optional) In the **Interfaces** > **Subinterface Range** drop-down list, choose a subinterface ID.

For a range of subinterface IDs, choose the ending ID in the second drop-down list, if available.

In transparent firewall mode, only subinterfaces that have not been allocated to other contexts are shown.

c) (Optional) In the **Aliased Names** area, check **Use Aliased Name in Context** to set an aliased name for this interface to be used in the context configuration instead of the interface ID.

- In the **Name** field, set the aliased name.

  An aliased name must start with a letter, end with a letter, and have as interior characters only letters, digits, or an underscore. This field lets you specify a name that ends with a letter or underscore; to add an optional digit after the name, set the digit in the Range field.

- (Optional) In the **Range** field, set the numeric suffix for the aliased name.

  If you have a range of subinterfaces, you can enter a range of digits to be appended to the name.

d) (Optional)Check **Show Hardware Properties in Context** to enable context users to see physical interface properties even if you set an aliased name.

e) Click **OK** to return to the **Add Context** dialog box.

Step 5 (Optional) In the **Resource Assignment** area, choose a class name from the **Resource Class** drop-down list to assign this context to a resource class.

You can add or edit a resource class directly from this area.

Step 6 From the **Config URL** drop-down list, choose a file system type. In the field, identify the URL for the context configuration location.

For example, the combined URL for FTP has the following format:

ftp://server.example.com/configs/admin.cfg

Step 7 (Optional) Click **Login** to set the username and password for external file systems.

Step 8 (Optional) From the **Failover Group** drop-down list, choose the group name to set the failover group for Active/Active failover.

Step 9 (Optional) For **Cloud Web Security**, click **Enable** to enable Web Security inspection in this context. To override the license set in the system configuration, enter a license in the **License** field.

Step 10 (Optional) In the **Description** field, add a description.

Step 11 (Optional) In the **Storage URL Assignment** area, you can allow each context to use flash memory to store VPN packages, such as Secure Client, as well as providing storage for Secure Client and clientless SSL VPN portal customizations. For example, if you are using multiple context mode to configure an Secure Client profile with Dynamic Access Policies, you must plan for context specific private and shared storage. Each context can use a private storage space as well as a shared read-only storage space. **Note:** Make sure the target directory is already present on the specified disk using **Tools** > **File Management**.

a) Check the **Configure private storage assignment** check box, and from the **Select** drop-down list, choose the private storage directory. You can specify one private storage space per context. You can read/write/delete from this directory within the context (as well as from the system execution space). Under the specified path, the ASA creates a sub-directory named after the context. For example, for contextA if you specify **disk1:/private-storage** for the path, then the ASA creates a sub-directory for this context at **disk1:/private-storage/contextA/**. You can also optionally name the path within the context by entering a name in the **is mapped to** field so that the file system is not exposed to context administrators. For example, if you specify the mapped name as **context**, then from within the context, this directory is called **context:**. To control how much disk space is allowed per context, see Configure a Class for Resource Management, on page 263.

b) Check the **Configure shared storage assignment** check box, and from the **Select** drop-down list, choose the shared storage directory. You can specify one read-only **shared** storage space per context, but you can create multiple shared directories. To reduce duplication of common large files that can be shared among all contexts, such as Secure Client packages, you can use the shared storage space. The ASA does not create context sub-directories for this storage space because it is a shared space for multiple contexts. Only the system execution space can write and delete from the shared directory.

Step 12 Click **OK** to return to the **Security Contexts** pane.

Step 13 (Optional) Select the context, and click **Change Firewall Mode** to set the firewall mode to transparent.

If this is a new context, there is no configuration to erase. Click **Change Mode** to change to transparent firewall mode.

If this is an existing context, then be sure to back up the configuration before you change the mode.

Note
You cannot change the mode of the currently connected context in ASDM (typically the admin context); see Set the Firewall Mode (Single Mode), on page 225 to set the mode at the command line.

Step 14 (Optional) To customize auto-generation of MAC addresses, see Assign MAC Addresses to Context Interfaces Automatically, on page 268.

Step 15 (Optional) Check the **Specify the maximum number of TLS Proxy sessions that the ASA needs to support** check box, to specify the maximum TLS Proxy sessions for the device. For more information about TLS proxy, see the firewall configuration guide.

Assign MAC Addresses to Context Interfaces Automatically

This section describes how to configure auto-generation of MAC addresses. The MAC address is used to classify packets within a context.

Before you begin

- When you configure a name for the interface in a context, the new MAC address is generated immediately. If you enable this feature after you configure context interfaces, then MAC addresses are generated for all interfaces immediately after you enable it. If you disable this feature, the MAC address for each interface reverts to the default MAC address. For example, subinterfaces of GigabitEthernet 0/1 revert to using the MAC address of GigabitEthernet 0/1.
- In the rare circumstance that the generated MAC address conflicts with another private MAC address in your network, you can manually set the MAC address for the interface within the context.

Procedure

Step 1 If you are not already in the System configuration mode, in the Device List pane, double-click **System** under the active device IP address.

Step 2 Choose **Configuration** > **Context Management** > **Security Contexts**, and check **Mac-Address auto**. If you do not enter a prefix, then the ASA autogenerates the prefix based on the last two bytes of the interface.

Step 3 (Optional) Check the **Prefix** check box, and in the field, enter a decimal value between 0 and 65535.

This prefix is converted to a four-digit hexadecimal number, and used as part of the MAC address.

Change Between Contexts and the System Execution Space

If you log in to the system execution space (or the admin context), you can change between contexts and perform configuration and monitoring tasks within each context. The running configuration that you edit in a configuration mode depends on your location. When you are in the system execution space, the running configuration consists only of the system configuration; when you are in a context, the running configuration consists only of that context.

Procedure

Step 1 In the Device List pane, double-click **System** under the active device IP address, to configure the System.

Step 2 In the Device List pane, double-click the context name under the active device IP address, to configure a context.

Manage Security Contexts

This section describes how to manage security contexts.

Remove a Security Context

You cannot remove the current admin context.

Note If you use failover, there is a delay between when you remove the context on the active unit and when the context is removed on the standby unit.

Before you begin

Perform this procedure in the system execution space.

Procedure

Step 1 If you are not already in the System configuration mode, in the Device List pane, double-click **System** under the active device IP address.

Step 2 Choose **Configuration** > **Context Management** > **Security Contexts**.

Step 3 Select the context you want to delete, and click **Delete**.

The Delete Context dialog box appears.

Step 4 If you might want to re-add this context later, and want to keep the configuration file for future use, uncheck the **Also delete config URL file from the disk** check box.

If you want to delete the configuration file, then leave the check box checked.

Step 5 Click **Yes**.

Change the Admin Context

The system configuration does not include any network interfaces or network settings for itself; rather, when the system needs to access network resources (such as downloading the contexts from the server), it uses one of the contexts that is designated as the admin context.

The admin context is just like any other context, except that when a user logs in to the admin context, then that user has system administrator rights and can access the system and all other contexts. The admin context is not restricted in any way, and can be used as a regular context. However, because logging into the admin context grants you administrator privileges over all contexts, you might need to restrict access to the admin context to appropriate users.

Note For ASDM, you cannot change the admin context within ASDM because your ASDM session would disconnect. You can perform this procedure using the Command Line Interface tool noting that you will have to reconnect to the new admin context.

Before you begin

- You can set any context to be the admin context, as long as the configuration file is stored in the internal flash memory.
- Perform this procedure in the system execution space.

Procedure

Step 1 If you are not already in the System configuration mode, in the Device List pane, double-click **System** under the active device IP address.

Step 2 Choose **Tools** > **Command Line Interface**.

The Command Line Interface dialog box appears.

Step 3 Enter the following command:

admin-context *context_name*

Step 4 Click **Send**.

Any remote management sessions, such as Telnet, SSH, or HTTPS (ASDM), that are connected to the admin context are terminated. You must reconnect to the new admin context.

Note
A few system configuration commands, including **ntp server**, identify an interface name that belongs to the admin context. If you change the admin context, and that interface name does not exist in the new admin context, be sure to update any system commands that refer to the interface.

Change the Security Context URL

This section describes how to change the context URL.

Before you begin

- You cannot change the security context URL without reloading the configuration from the new URL. The ASA merges the new configuration with the current running configuration.
- Reentering the same URL also merges the saved configuration with the running configuration.
- A merge adds any new commands from the new configuration to the running configuration.
  - If the configurations are the same, no changes occur.
  - If commands conflict or if commands affect the running of the context, then the effect of the merge depends on the command. You might get errors, or you might have unexpected results. If the running configuration is blank (for example, if the server was unavailable and the configuration was never downloaded), then the new configuration is used.
- If you do not want to merge the configurations, you can clear the running configuration, which disrupts any communications through the context, and then reload the configuration from the new URL.

- Perform this procedure in the system execution space.

Procedure

Step 1 If you are not already in the System configuration mode, in the Device List pane, double-click **System** under the active device IP address.

Step 2 Choose **Configuration** > **Context Management** > **Security Contexts**.

Step 3 Select the context you want to edit, and click **Edit**.

The Edit Context dialog box appears.

Step 4 Enter a new URL in the Config URL field, and click **OK**.

The system immediately loads the context so that it is running.

Reload a Security Context

You can reload the context in two ways:

- Clear the running configuration and then import the startup configuration.

  This action clears most attributes associated with the context, such as connections and NAT tables.

- Remove the context from the system configuration.

  This action clears additional attributes, such as memory allocation, which might be useful for troubleshooting. However, to add the context back to the system requires you to respecify the URL and interfaces.

Reload by Clearing the Configuration

Procedure

Step 1 In the Device List pane, double-click the context name under the active device IP address.

Step 2 Choose **Tools** > **Command Line Interface**.

The Command Line Interface dialog box appears.

Step 3 Enter the following command:

clear configure all

Step 4 Click **Send**.

The context configuration is cleared.

Step 5 Choose **Tools** > **Command Line Interface** again.

The Command Line Interface dialog box appears.

Step 6 Enter the following command:

copy startup-config running-config

Step 7 Click **Send**.

The ASA reloads the configuration. The ASA copies the configuration from the URL specified in the system configuration. You cannot change the URL from within a context.

Reload by Removing and Re-adding the Context

To reload the context by removing the context and then re-adding it, perform the steps.

Procedure

Step 1 Remove a Security Context, on page 270. Be sure to uncheck the **Also delete config URL file from the disk** check box.

Step 2 Configure a Security Context, on page 266

Monitoring Security Contexts

This section describes how to view and monitor context information.

Monitor Context Resource Usage

Procedure

Step 1 If you are not already in the System mode, in the Device List pane, double-click **System** under the active device IP address.

Step 2 Click the **Monitoring** button on the toolbar.

Step 3 Click **Context Resource Usage**.

Click each resource type to view the resource usage for all contexts:

- **ASDM/Telnet/SSH**—Shows the usage of ASDM, Telnet, and SSH connections.
  - Context—Shows the name of each context.

    For each access method, see the following usage statistics:
  - Existing Connections (#)—Shows the number of existing connections.
  - Existing Connections (%)—Shows the connections used by this context as a percentage of the total number of connections used by all contexts.

    - Peak Connections (#)—Shows the peak number of connections since the statistics were last cleared, either using the **clear resource usage** command or because the device rebooted.

- **Routes**—Shows the usage of dynamic routes.
    - Context—Shows the name of each context.
    - Existing Connections (#)—Shows the number of existing connections.
    - Existing Connections (%)—Shows the connections used by this context as a percentage of the total number of connections used by all contexts.
    - Peak Connections (#)—Shows the peak number of connections since the statistics were last cleared, either using the **clear resource usage** command or because the device rebooted.

- **Xlates**—Shows the usage of network address translations.
    - Context—Shows the name of each context.
    - Xlates (#)—Shows the number of current xlates.
    - Xlates (%)—Shows the xlates used by this context as a percentage of the total number of xlates used by all contexts.
    - Peak (#)—Shows the peak number of xlates since the statistics were last cleared, either using the **clear resource usage** command or because the device rebooted.

- **NATs**—Shows the number of NAT rules.
    - Context—Shows the name of each context.
    - NATs (#)—Shows the current number of NAT rules.
    - NATs (%)—Shows the NAT rules used by this context as a percentage of the total number of NAT rules used by all contexts.
    - Peak NATs (#)—Shows the peak number of NAT rules since the statistics were last cleared, either using the **clear resource usage** command or because the device rebooted.

- **Syslogs**—Shows the rate of system log messages.
    - Context—Shows the name of each context.
    - Syslog Rate (#/sec)—Shows the current rate of system log messages.
    - Syslog Rate (%)—Shows the system log messages generated by this context as a percentage of the total number of system log messages generated by all contexts.
    - Peak Syslog Rate (#/sec)—Shows the peak rate of system log messages since the statistics were last cleared, either using the **clear resource usage** command or because the device rebooted.

- VPN—Shows the usage of VPN site-to-site tunnels.
    - Context—Shows the name of each context.
    - VPN Connections—Shows usage of guaranteed VPN sessions.
    - VPN Burst Connections—Shows usage of burst VPN sessions.

- Existing (#)—Shows the number of existing tunnels.
- Peak (#)—Shows the peak number of tunnels since the statistics were last cleared, either using the **clear resource usage** command or because the device rebooted.

Step 4 Click **Refresh** to refresh the view.

View Assigned MAC Addresses

You can view auto-generated MAC addresses within the system configuration or within the context.

View MAC Addresses in the System Configuration

This section describes how to view MAC addresses in the system configuration.

Before you begin

If you manually assign a MAC address to an interface, but also have auto-generation enabled, the auto-generated address continues to show in the configuration even though the manual MAC address is the one that is in use. If you later remove the manual MAC address, the auto-generated one shown will be used.

Procedure

Step 1 If you are not already in the System configuration mode, in the Device List pane, double-click **System** under the active device IP address.

Step 2 Choose **Configuration** > **Context Management** > **Security Contexts**, and view the Primary MAC and Secondary MAC columns.

View MAC Addresses Within a Context

This section describes how to view MAC addresses within a context.

Procedure

Step 1 If you are not already in the System configuration mode, in the Device List pane, double-click **System** under the active device IP address.

Step 2 Choose **Configuration** > **Interfaces**, and view the MAC Address address column.

This table shows the MAC address in use; if you manually assign a MAC address and also have auto-generation enabled, then you can only view the unused auto-generated address from within the system configuration.

History for Multiple Context Mode

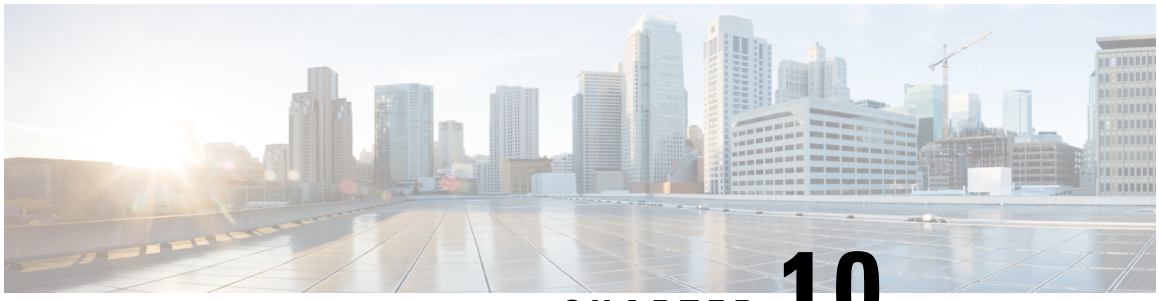
Table 15: History for Multiple Context Mode

Feature Name	Platform Releases	Feature Information
Multiple security contexts	7.0(1)	Multiple context mode was introduced. We introduced the following screens: Configuration > Context Management.
Automatic MAC address assignment	7.2(1)	Automatic assignment of MAC address to context interfaces was introduced. We modified the following screen: Configuration > Context Management > Security Contexts.
Resource management	7.2(1)	Resource management was introduced. We introduced the following screen: Configuration > Context Management > Resource Management.
Virtual sensors for IPS	8.0(2)	The AIP SSM running IPS software Version 6.0 and above can run multiple virtual sensors, which means you can configure multiple security policies on the AIP SSM. You can assign each context or single mode ASA to one or more virtual sensors, or you can assign multiple security contexts to the same virtual sensor. We modified the following screen: Configuration > Context Management > Security Contexts.
Automatic MAC address assignment enhancements	8.0(5) 8.2(2)	The MAC address format was changed to use a prefix, to use a fixed starting value (A2), and to use a different scheme for the primary and secondary unit MAC addresses in a failover pair. The MAC addresses are also now persistent across reloads. The command parser now checks if auto-generation is enabled; if you want to also manually assign a MAC address, you cannot start the manual MAC address with A2. We modified the following screen: Configuration > Context Management > Security Contexts.
Maximum contexts increased for the ASA 5550 and 5580	8.4(1)	The maximum security contexts for the ASA 5550 was increased from 50 to 100. The maximum for the ASA 5580 was increased from 50 to 250.
Automatic MAC address assignment enabled by default	8.5(1)	Automatic MAC address assignment is now enabled by default. We modified the following screen: Configuration > Context Management > Security Contexts.

Feature Name	Platform Releases	Feature Information
Automatic generation of a MAC address prefix	8.6(1)	In multiple context mode, the ASA now converts the automatic MAC address generation configuration to use a default prefix. The ASA auto-generates the prefix based on the last two bytes of the interface (ASA 5500-X) or backplane (ASASM) MAC address. This conversion happens automatically when you reload, or if you reenable MAC address generation. The prefix method of generation provides many benefits, including a better guarantee of unique MAC addresses on a segment. If you want to change the prefix, you can reconfigure the feature with a custom prefix. The legacy method of MAC address generation is no longer available. **Note** To maintain hitless upgrade for failover pairs, the ASA does *not* convert the MAC address method in an existing configuration upon a reload if failover is enabled. However, we strongly recommend that you manually change to the prefix method of generation when using failover, especially for the ASASM. Without the prefix method, ASASMs installed in different slot numbers experience a MAC address change upon failover, and can experience traffic interruption. After upgrading, to use the prefix method of MAC address generation, reenable MAC address generation to use the default prefix. We modified the following screen: Configuration > Context Management > Security Contexts
Automatic MAC address assignment disabled by default on all models except for the ASASM	9.0(1)	Automatic MAC address assignment is now disabled by default except for the ASASM. We modified the following screen: Configuration > Context Management > Security Contexts.
Dynamic routing in Security Contexts	9.0(1)	EIGRP and OSPFv2 dynamic routing protocols are now supported in multiple context mode. OSPFv3, RIP, and multicast routing are not supported.
New resource type for routing table entries	9.0(1)	A new resource type, routes, was created to set the maximum number of routing table entries in each context. We modified the following screen: Configuration > Context Management > Resource Class > Add Resource Class
Site-to-Site VPN in multiple context mode	9.0(1)	Site-to-site VPN tunnels are now supported in multiple context mode.
New resource type for site-to-site VPN tunnels	9.0(1)	New resource types, vpn other and vpn burst other, were created to set the maximum number of site-to-site VPN tunnels in each context. We modified the following screen: Configuration > Context Management > Resource Class > Add Resource Class

Feature Name	Platform Releases	Feature Information
New resource type for IKEv1 SA negotiations	9.1(2)	New resource type, ikev1 in-negotiation, was created to set the maximum percentage of IKEv1 SA negotiations in each context to prevent overwhelming the CPU and crypto engines. Under certain conditions (large certificates, CRL checking), you might want to restrict this resource. We modified the following screen: Configuration > Context Management > Resource Class > Add Resource Class
Support for Remote Access VPN in multiple context mode	9.5(2)	You can now use the following remote access features in multiple context mode: • AnyConnect 3.x and later (SSL VPN only; no IKEv2 support) • Centralized Secure Client image configuration • Secure Client image upgrade • Context Resource Management for Secure Client connections **Note** The Secure Client Premier license is required for multiple context mode; you cannot use the default or legacy license. We modified the following screen: **Configuration** > **Context Management** > **Resource Class** > **Add Resource Class**
Pre-fill/Username-from-cert feature for multiple context mode	9.6(2)	Secure Client SSL support is extended, allowing pre-fill/username-from-certificate feature CLIs, previously available only in single mode, to be enabled in multiple context mode as well. We did not modify any screens.
Flash Virtualization for Remote Access VPN	9.6(2)	Remote access VPN in multiple context mode now supports flash virtualization. Each context can have a private storage space and a shared storage place based on the total flash that is available: • Private storage—Store files associated only with that user and specific to the content that you want for that user. • Shared storage—Upload files to this space and have it accessible to any user context for read/write access once you enable it. We modified the following screens: **Configuration** > **Context Management** > **Resource Class** > **Add Resource Class** **Configuration** > **Context Management** > **Security Contexts**
Secure Clientprofiles supported in multi-context devices	9.6(2)	Secure Client profiles are supported in multi-context devices. To add a new profile using ASDM, you must have the Secure Client release 4.2.00748 or 4.3.03013 and later.
Stateful failover for Secure Client connections in multiple context mode	9.6(2)	Stateful failover is now supported for Secure Client connections in multiple context mode. We did not modify any screens.

Feature Name	Platform Releases	Feature Information
Remote Access VPN Dynamic Access Policy (DAP) is supported in multiple context mode	9.6(2)	You can now configure DAP per context in multiple context mode. We did not modify any screens.
Remote Access VPN CoA (Change of Authorization) is supported in multiple context mode	9.6(2)	You can now configure CoA per context in multiple context mode. We did not modify any screens.
Remote Access VPN localization is supported in multiple context mode	9.6(2)	Localization is supported globally. There is only one set of localization files that are shared across different contexts. We did not modify any screens.
Remote Access VPN for IKEv2 is supported in multiple context mode	9.9(2)	You can configure Remote Access VPN in multiple context mode for IKEv2.
Configurable limitation of admin sessions	9.12(1)	You can configure the maximum number of aggregate, per user, and per-protocol administrative sessions. Formerly, you could configure only the aggregate number of sessions. This feature does not affect console sessions. Note that in multiple context mode, you cannot configure the number of HTTPS sessions, where the maximum is fixed at 5 sessions. The **quota management-session** command is also no longer accepted in the system configuration, and is instead available in the context configuration. The maximum aggregate sessions is now 15; if you configured 0 (unlimited) or 16+, then when you upgrade, the value is changed to 15. New/Modified screens: **Configuration** > **Device Management** > **Management Access** > **Management Session Quota**
Firepower 1140 maximum contexts increased from 5 to 10	9.16(1)	The Firepower 1140 now supports up to 10 contexts.
Multiple context support for all Secure Firewall 1200 models	9.23(1)	We added support for multiple context mode for the Secure Firewall 1210/1220: • Secure Firewall 1210CE—5 contexts. • Secure Firewall 1210CP—5 contexts. • Secure Firewall 1220CX—10 contexts. Switchports are not supported in multiple context mode, and you must convert all interfaces to router interfaces before you can convert to multiple context mode. The Secure Firewall 1230/1240/1250 also supports multiple context mode in its initial release: • Secure Firewall 1230—25 contexts. • Secure Firewall 1240—25 contexts. • Secure Firewall 1250—25 contexts.



CHAPTER 10

Failover for High Availability

This chapter describes how to configure Active/Standby or Active/Active failover to accomplish high availability of the ASA.

- About Failover, on page 281
- Licensing for Failover, on page 301
- Guidelines for Failover, on page 302
- Defaults for Failover, on page 305
- Configure Active/Standby Failover, on page 305
- Configure Active/Active Failover, on page 306
- Configure Optional Failover Parameters, on page 308
- Manage Failover, on page 313
- Monitoring Failover, on page 318
- History for Failover, on page 320

About Failover

Configuring failover requires two identical ASAs connected to each other through a dedicated failover link and, optionally, a state link. The health of the active units and interfaces is monitored to determine whether they meet the specific failover conditions. If those conditions are met, failover occurs.

Failover Modes

The ASA supports two failover modes, Active/Active failover and Active/Standby failover. Each failover mode has its own method for determining and performing failover.

- In Active/Standby failover, one device functions as the Active unit and passes traffic. The second device, designated as the Standby unit, does not actively pass traffic. When a failover occurs, the Active unit fails over to the Standby unit, which then becomes Active. You can use Active/Standby failover for ASAs in single or multiple context mode.
- In an Active/Active failover configuration, both ASAs can pass network traffic. Active/Active failover is only available to ASAs in multiple context mode. In Active/Active failover, you divide the security contexts on the ASA into 2 *failover groups*. A failover group is simply a logical group of one or more security contexts. One group is assigned to be Active on the primary ASA, and the other group is assigned to be active on the Secondary ASA. When a failover occurs, it occurs at the failover group level.

Both failover modes support stateful or stateless failover.

Failover System Requirements

This section describes the hardware, software, and license requirements for ASAs in a Failover configuration.

Hardware Requirements

The two units in a Failover configuration must:

- Be the same model.

  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.

- 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 Failover. If you change the interfaces after you enable Failover, make the interface changes in FXOS on the Standby unit, and then make the same changes on the Active unit. If you remove an interface in FXOS (for example, if you remove a network module, remove an EtherChannel, or reassign an interface to an EtherChannel), then the ASA configuration retains the original commands so that you can make any necessary adjustments; removing an interface from the configuration can have wide effects. You can manually remove the old interface configuration in the ASA OS.

- Have the same modules installed (if any).

- Have the same RAM installed.

If you are using units with different flash memory sizes in your Failover 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 Failover configuration must:

- Be in the same context mode (single or multiple).

- For single mode: Be in the same firewall mode (routed or transparent).

  In multiple context mode, the firewall mode is set at the context-level, and you can use mixed modes.

- Have the same major (first number) and minor (second number) software version. However, you can temporarily use different versions of the software during an upgrade process; for example, you can upgrade one unit from Version 8.3(1) to Version 8.3(2) and have failover remain active. We recommend upgrading both units to the same version to ensure long-term compatibility.

- Have the same Secure Client images. If the failover pair has mismatched images when a hitless upgrade is performed, then the clientless SSL VPN connection terminates in the final reboot step of the upgrade process, the database shows an orphaned session, and the IP pool shows that the IP address assigned to the client is “in use.”

- Be in the same FIPS mode.
- (Firepower 4100/9300) Have the same flow offload mode, either both enabled or both disabled.

License Requirements

The two units in a failover configuration do not need to have identical licenses; the licenses combine to make a failover cluster license.

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.



Caution All information sent over the failover and state links is sent in clear text unless you secure the communication with an IPsec tunnel or a failover key. If the ASA is used to terminate VPN tunnels, this information includes any usernames, passwords and preshared keys used for establishing the tunnels. Transmitting this sensitive data in clear text could pose a significant security risk. We recommend securing the failover communication with an IPsec tunnel or a failover key if you are using the ASA to terminate VPN tunnels.

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, subinterface, or EtherChannel) as the failover link; however, you cannot specify an interface that is currently configured with a name. 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). For most models, you cannot use a management interface for failover unless explicitly described below.

The ASA 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.

See the following guidelines for the failover link:

- 5506-X through 5555-X—You cannot use the Management interface as the failover link; you must use a data interface. The only exception is for the 5506H-X, where you can use the management interface as the failover link.
- 5506H-X—You *can* use the Management 1/1 interface as the failover link. If you configure it for failover, you must reload the device for the change to take effect. In this case, you cannot also use the ASA Firepower module, because it requires the Management interface for management purposes.
- Firepower 4100/9300—You cannot use the management-type interface for the failover link.
- See the following guidelines for sizing the link.

Table 16: 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 (the **failover polltime unit** command).

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 ASA .

- 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.

The ASA supports Auto-MDI/MDIX on its copper Ethernet ports, so you can either use a crossover cable or a straight-through cable. If you use a straight-through cable, the interface automatically detects the cable and swaps one of the transmit/receive pairs to MDIX.

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

You can use a dedicated data interface (physical or EtherChannel) for the state link. See Interface for the Failover Link, on page 283 for requirements for a dedicated state link, and Connecting the Failover Link, on page 284 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 ASA can use the data interfaces to determine if a failover is required. Subsequently, 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 ASAs, then when a switch or inter-switch-link is down, both ASAs become active. Therefore, the following two connection methods shown in the following figures are NOT recommended.

Figure 53: Connecting with a Single Switch—Not Recommended

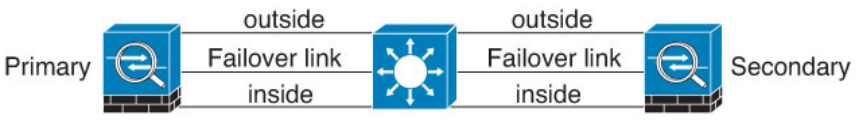
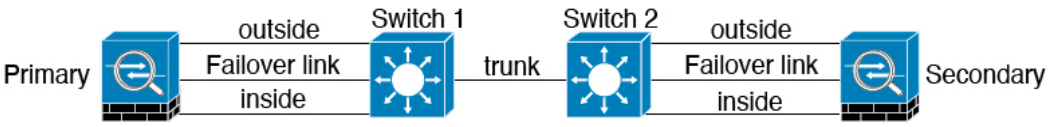


Figure 54: Connecting with a Double-Switch—Not Recommended



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 55: Connecting with a Different Switch

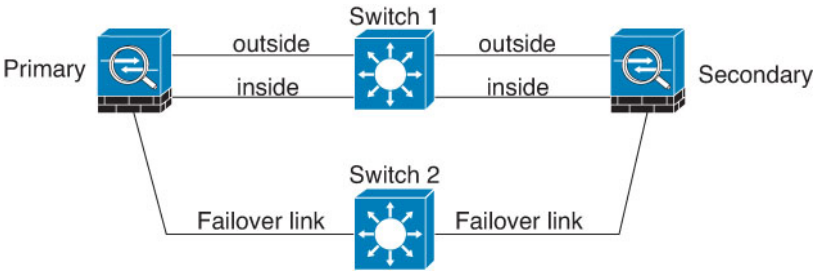
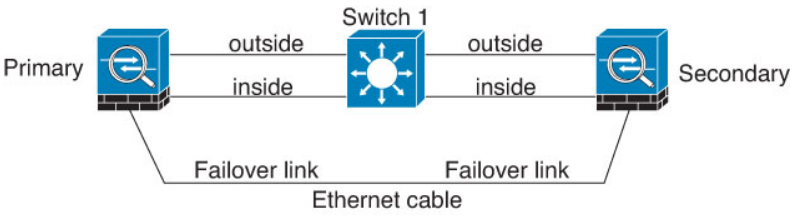


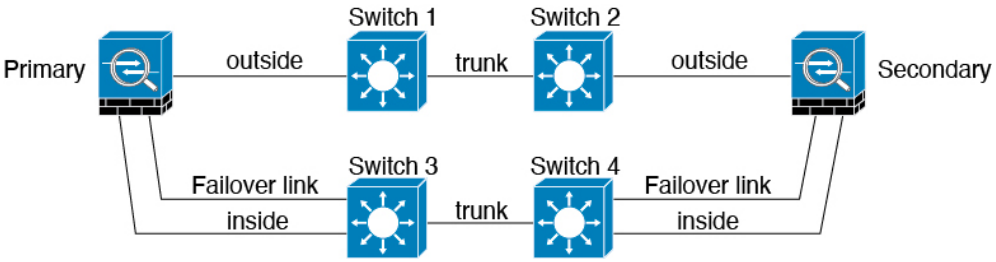
Figure 56: Connecting with a Cable



Scenario 3—Recommended

If the ASA 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 57: Connecting with a Secure Switch



MAC Addresses and IP Addresses in Failover

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 Failover, 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 failover and set the failover configurations to a disabled state, you will need to manually resume failover, or reboot the device. It is recommended to use the command **failover reset** and resume the failover 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 **failover reset** 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 ASA 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.

Active/Active IP Addresses and MAC Addresses

For Active/Active failover, see the following for IP address and MAC address usage during a failover event:

1. The primary unit autogenerates active and standby MAC addresses for all interfaces in failover group 1 and 2 contexts. You can also manually configure the MAC addresses if necessary, for example, if there are MAC address conflicts.
2. Each unit uses the active IP addresses and MAC addresses for its active failover group, and the standby addresses for its standby failover group. For example, the primary unit is active for failover group 1, so it uses the active addresses for contexts in failover group 1. It is standby for the contexts in failover group 2, where it uses the standby addresses.
3. When a unit fails over, the other unit assumes the active IP addresses and MAC addresses of the failed failover group and begins passing traffic.
4. When the failed unit comes back online, and you enabled the preempt option, it resumes the failover group.

Virtual MAC Addresses

The ASA 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. Manual methods include the interface mode **mac-address** command, the **failover mac address** command, and for Active/Active failover, the failover group mode **mac address** command, in addition to autogeneration methods described below.

In multiple context mode, you can configure the ASA to generate virtual active and standby MAC addresses automatically for shared interfaces, and these assignments are synced to the secondary unit (see the **mac-address auto** command). For non-shared interfaces, you can manually set the MAC addresses for Active/Standby mode (Active/Active mode autogenerates MAC addresses for all interfaces).

For Active/Active failover, virtual MAC addresses are always used, either with default values or with values you can set per interface.

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.

Stateless and Stateful Failover

The ASA supports two types of failover, stateless and stateful for both the Active/Standby and Active/Active modes.

Note Some configuration elements for clientless SSL VPN (such as bookmarks and customization) use the VPN failover subsystem, which is part of Stateful Failover. You must use Stateful Failover to synchronize these elements between the members of the failover pair. Stateless failover is not recommended for clientless SSL VPN.

Stateless Failover

When a failover occurs, all active connections are dropped. Clients need to reestablish connections when the new active unit takes over.

Note Some configuration elements for clientless SSL VPN (such as bookmarks and customization) use the VPN failover subsystem, which is part of Stateful Failover. You must use Stateful Failover to synchronize these elements between the members of the failover pair. Stateless (regular) failover is not recommended for clientless SSL VPN.

Stateful Failover

When Stateful Failover is enabled, the active unit continually passes per-connection state information to the standby unit, or in Active/Active failover, between the active and standby failover groups. 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 ASA:

- NAT translation table.
- TCP and UDP connections and 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.
- The HTTP connection table (unless you enable HTTP replication).
- The HTTP connection states (if HTTP replication is enabled)—By default, the ASA does not replicate HTTP session information when Stateful Failover is enabled. We suggest that you enable HTTP replication.
- SCTP connection states. However, SCTP inspection stateful failover is best effort. During failover, if any SACK packets are lost, the new active unit will drop all other out of order packets in the queue until the missing packet is received.
- 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.

- ICMP connection state—ICMP connection replication is enabled only if the respective interface is assigned to an asymmetric routing group.
- 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.
- Cisco IP SoftPhone sessions—If a failover occurs during an active Cisco IP SoftPhone session, the call remains active because the call session state information is replicated to the standby unit. When the call is terminated, the IP SoftPhone client loses connection with the Cisco Call Manager. This connection loss occurs because there is no session information for the CTIQBE hangup message on the standby unit. When the IP SoftPhone client does not receive a response back from the Call Manager within a certain time period, it considers the Call Manager unreachable and unregisters itself.
- 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 ASA:

- The user authentication (uauth) table
- TCP state bypass connections
- Multicast routing.
- Selected clientless SSL VPN features:
  - Smart Tunnels
  - Port Forwarding
  - Plugins
  - Java Applets

- IPv6 clientless or Secure Client sessions
- Citrix authentication (Citrix users must reauthenticate after failover)

Bridge Group Requirements for Failover

There are special considerations for failover when using bridge groups.

Bridge Group Requirements for Appliances, ASAv

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 configure one of the following workarounds depending on the switch port mode:

- 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.

- Trunk mode—Block BPDUs on the ASA on a bridge group's member interfaces with an EtherType access rule.

  ```
  access-list id ethertype deny bpdu
  access-group id in interface name1
  access-group id in interface name2
  ```

  Blocking BPDUs disables STP on the switch. Be sure not to have any loops involving the ASA in your network layout.

If neither of the above options are possible, then you can use one of the following less desirable workarounds that impacts failover functionality or STP stability:

- Disable interface monitoring.
- Increase interface holdtime to a high value that will allow STP to converge before the ASAs fail over.
- Decrease STP timers to allow STP to converge faster than the interface holdtime.

Failover Health Monitoring

The ASA monitors each unit for overall health and for interface health. This section includes information about how the ASA performs tests to determine the state of each unit.

Unit Health Monitoring

The ASA 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. For the Firepower 9300 and 4100 series, you can enable Bidirectional Forwarding Detection (BFD) monitoring, which is more reliable than hello messages. The action that the ASA takes depends on the response from the other unit. See the following possible actions:

- If the ASA receives a response on the failover link, then it does not fail over.
- If the ASA 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 ASA 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.
- This feature is not available in Version 7.3.

Interface Monitoring

You can monitor up to 1025 interfaces (in multiple context mode, divided between all contexts). You should monitor important interfaces. For example in multiple context mode, you might configure one context to monitor a shared interface: because the interface is shared, all contexts benefit from the monitoring.

When a unit does not receive hello messages on a monitored interface for 15 seconds (the default), it runs interface tests. (To change the period, see **Configuration** > **Device Management** > **High Availability and**

Scalability > **Failover** > **Criteria** > **Failover Poll Times**.) 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 ASA stops running tests.

If the threshold you define for the number of failed interfaces is met (see **Configuration** > **Device Management** > **High Availability and Scalability** > **Failover** > **Criteria** > **Interface Policy**), 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 ASA 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 ASA uses the IPv4 addresses to perform the health monitoring. If an interface has only IPv6 addresses configured on it, then the ASA uses IPv6 neighbor discovery instead of ARP to perform the health monitoring tests. For the broadcast ping test, the ASA uses the IPv6 all nodes address (FE02::1).



Note If a failed unit does not recover and you believe it should not be failed, you can reset the state by entering the **failover reset** command. If the failover condition persists, however, the unit will fail again.

Interface Tests

The ASA uses the following interface tests. The duration of each test is approximately 1.5 seconds by default, or 1/16 of the failover interface holdtime(see **Configuration** > **Device Management** > **High Availability and Scalability** > **Failover** > **Criteria** > **Failover Poll Times**).

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 ASA considers it failed, and testing stops. If the status is Up, then the ASA 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 ASA 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 ASA 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 ASA 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.
- Testing—Hello messages are not heard on the interface for five poll times.
- Link Down—The interface or VLAN is administratively down.
- No Link—The physical link for the interface is down.
- Failed—No traffic is received on the interface, yet traffic is heard on the peer interface.

Failover Times

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 17:

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.



Note If you manually fail over using the CLI or ASDM, or you reload the ASA, the failover starts immediately and is not subject to the timers listed below.

Table 18: ASA

Failover Condition	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 main board interface link down.	500 milliseconds	5 seconds	15 seconds
Active unit 4GE module interface link down.	2 seconds	5 seconds	15 seconds
Active unit interface up, but connection problem causes interface testing.	5 seconds	25 seconds	75 seconds

Configuration Synchronization

Failover includes various types of configuration synchronization.

Running Configuration Replication

Running configuration replication occurs when any one or both of the devices in the failover pair boot.

In Active/Standby failover, configurations are always synchronized from the active unit to the standby unit.

In Active/Active failover, whichever unit boots second obtains the running configuration from the unit that boots first, regardless of the primary or secondary designation of the booting unit. After both units are up, commands entered in the system execution space are replicated from the unit on which failover group 1 is in the active state.

When the standby/second unit completes its initial startup, it clears its running configuration (except for the **failover** commands needed to communicate with the active unit), and the active unit sends its entire configuration to the standby/second unit. When the replication starts, the ASA console on the active unit displays the message “Beginning configuration replication: Sending to mate,” and when it is complete, the ASA displays the message “End Configuration Replication to mate.” Depending on the size of the configuration, replication can take from a few seconds to several minutes.

On the unit receiving the configuration, the configuration exists only in running memory. You should save the configuration to flash memory. For example, in Active/Active failover, enter the **write memory all** command in the system execution space on the unit that has failover group 1 in the active state. The command is replicated to the peer unit, which proceeds to write its configuration to flash memory.

Note During replication, commands entered on the unit sending the configuration may not replicate properly to the peer unit, and commands entered on the unit receiving the configuration may be overwritten by the configuration being received. Avoid entering commands on either unit in the failover pair during the configuration replication process.

File Replication

Configuration syncing does not replicate the following files and configuration components, so you must copy these files manually so they match:

- Secure Client images
- CSD images
- Secure Client profiles

  The ASA uses a cached file for the Secure Client profile stored in cache:/stc/profiles, and not the file stored in the flash file system. To replicate the Secure Client profile to the standby unit, perform one of the following:

  - Enter the **write standby** command on the active unit.
  - Reapply the profile on the active unit.
  - Reload the standby unit.

- Local Certificate Authorities (CAs)
- ASA images
- ASDM images

Command Replication

After startup, commands that you enter on the active unit are immediately replicated on the standby unit.You do not have to save the active configuration to flash memory to replicate the commands.

In Active/Active failover,changes entered in the system execution space are replicated from the unit on which failover group 1 is in the active state.

Failure to enter the changes on the appropriate unit for command replication to occur causes the configurations to be out of synchronization. Those changes may be lost the next time the initial configuration synchronization occurs.

The following commands are replicated to the standby ASA:

- All configuration commands except for **mode**, **firewall**, and **failover lan unit**
- **copy running-config startup-config**
- **delete**
- **mkdir**
- **rename**
- **rmdir**
- **write memory**

The following commands are *not* replicated to the standby ASA:

- All forms of the **copy** command except for **copy running-config startup-config**
- All forms of the **write** command except for **write memory**

- **debug**
- **failover lan unit**
- **firewall**
- **show**
- **terminal pager** and **pager**

Config-Sync Optimization

When a device reboots or rejoins following a suspend or resume failover, 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 failover 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.
- ASA 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 failover 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.

Use the following commands for monitoring configuration sync optimization. You can execute these commands using **Tools > Command Line Interface**.

- **show failover config-sync checksum**

  Displays information about the device status and checksum.

- **show failover config-sync configuration**

Displays information about the device configuration and checksum.

- **show failover config-sync status**

  Displays status of configuration sync optimization functionality.

About Active/Standby Failover

Active/Standby failover lets you use a standby ASA to take over the functionality of a failed unit. When the active unit fails, the standby unit becomes the active unit. However, you must set the standby unit to primary before the failed unit is replaced, in order to retain the configuration of the secondary unit.



Note For multiple context mode, the ASA can fail over the entire unit (including all contexts) but cannot fail over individual contexts separately.

Primary/Secondary Roles and Active/Standby Status

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. Even on systems running in multiple context mode, you cannot fail over individual or groups of contexts.

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 19: 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.

About Active/Active Failover

This section describes Active/Active failover.

Active/Active Failover Overview

In an Active/Active failover configuration, both ASAs can pass network traffic. Active/Active failover is only available to ASAs in multiple context mode. In Active/Active failover, you divide the security contexts on the ASA into a maximum of 2 failover groups.

A failover group is simply a logical group of one or more security contexts. You can assign failover group to be active on the primary ASA, and failover group 2 to be active on the secondary ASA. When a failover occurs, it occurs at the failover group level. For example, depending on interface failure patterns, it is possible for failover group 1 to fail over to the secondary ASA, and subsequently failover group 2 to fail over to the

primary ASA. This event could occur if the interfaces in failover group 1 are down on the primary ASA but up on the secondary ASA, while the interfaces in failover group 2 are down on the secondary ASA but up on the primary ASA.

The admin context is always a member of failover group 1. Any unassigned security contexts are also members of failover group 1 by default. If you want Active/Active failover, but are otherwise uninterested in multiple contexts, the simplest configuration would be to add one additional context and assign it to failover group 2.



Note When configuring Active/Active failover, make sure that the combined traffic for both units is within the capacity of each unit.



Note You can assign both failover groups to one ASA if desired, but then you are not taking advantage of having two active ASAs.

Primary/Secondary Roles and Active/Standby Status for a Failover Group

As in Active/Standby failover, one unit in an Active/Active failover pair is designated the primary unit, and the other unit the secondary unit. Unlike Active/Standby failover, this designation does not indicate which unit becomes active when both units start simultaneously. Instead, the primary/secondary designation does two things:

- The primary unit provides the running configuration to the pair when they boot simultaneously.
- Each failover group in the configuration is configured with a primary or secondary unit preference. When used with preemption, this preference ensures that the failover group runs on the correct unit after it starts up. Without preemption, both groups run on the first unit to boot up.

Active Unit Determination for Failover Groups at Startup

The unit on which a failover group becomes active is determined as follows:

- When a unit boots while the peer unit is not available, both failover groups become active on the unit.
- When a unit boots while the peer unit is active (with both failover groups in the active state), the failover groups remain in the active state on the active unit regardless of the primary or secondary preference of the failover group until one of the following occurs:
  - A failover occurs.
  - A failover is manually forced.
  - A preemption for the failover group is configured, which causes the failover group to automatically become active on the preferred unit when the unit becomes available.

Failover Events

In an Active/Active failover configuration, failover occurs on a failover group basis, not a system basis. For example, if you designate both failover groups as Active on the primary unit, and failover group 1 fails, then

failover group 2 remains Active on the primary unit while failover group 1 becomes active on the secondary unit.

Because a failover group can contain multiple contexts, and each context can contain multiple interfaces, it is possible for all interfaces in a single context to fail without causing the associated failover group to fail.

The following table shows the failover action for each failure event. For each failure event, the policy (whether or not failover occurs), actions for the active failover group, and actions for the standby failover group are given.

Table 20: Failover Events

Failure Event	Policy	Active Group Action	Standby Group Action	Notes
A unit experiences a power or software failure	Failover	Become standby Mark as failed	Become active Mark active as failed	When a unit in a failover pair fails, any active failover groups on that unit are marked as failed and become active on the peer unit.
Interface failure on active failover group above threshold	Failover	Mark active group as failed	Become active	None.
Interface failure on standby failover group above threshold	No failover	No action	Mark standby group as failed	When the standby failover group is marked as failed, the active failover group does not attempt to fail over, even if the interface failure threshold is surpassed.
Formerly active failover group recovers	No failover	No action	No action	Unless failover group preemption is configured, the failover groups remain active on their current unit.
Failover link failed at startup	No failover	Become active	Become active	If the failover link is down at startup, both failover groups on 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.
Failover link failed during operation	No failover	n/a	n/a	Each unit marks the 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.

Licensing for Failover

For most models, failover units do not require the same license on each unit. If you have licenses on both units, they combine into a single running failover cluster license. There are some exceptions to this rule. See the following table for precise licensing requirements for failover.

Model	License Requirement
ASA Virtual	See Failover Licenses for the ASAv, on page 125.
Firepower 1010	Security Plus license on both units. See Failover Licenses for the Firepower 1010, on page 125.
Firepower 1100	See Failover Licenses for the Firepower 1100, on page 125.
Secure Firewall 1210/1220	See Failover Licenses for the Secure Firewall 1210/1220, on page 127.
Secure Firewall 1230/1240/1250	See Failover Licenses for the Secure Firewall 1230/1240/1250, on page 128.
Secure Firewall 3100/4200	See Failover Licenses for the Secure Firewall 3100, on page 129.
Firepower 4100/9300	See Failover Licenses for the Firepower 4100/9300, on page 131.
ISA 3000	Security Plus license on both units. **Note** Each unit must have the same encryption license.

Note A valid permanent key is required; in rare instances on the ISA 3000, your PAK authentication key can be removed. If your key consists of all 0's, then you need to reinstall a valid authentication key before failover can be enabled.

Guidelines for Failover

Context Mode

- Active/Active mode is supported only in multiple context mode.
- For multiple context mode, perform all steps in the system execution space unless otherwise noted.

Model Support

- Firepower 1010 and Secure Firewall 1210/1220:
  - You should not use the switch port functionality when using Failover. Because the switch ports operate in hardware, they continue to pass traffic on both the active *and* the standby units. Failover is designed to prevent traffic from passing through the standby unit, but this feature does not extend to switch ports. In a normal Failover 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 Failover, but a simpler setup is to use physical firewall interfaces instead.

- You can only use a firewall interface as the failover link.

- Firepower 9300—We recommend that you use inter-chassis Failover for the best redundancy.
- The ASA virtual on public cloud networks such as Microsoft Azure and Amazon Web Services are not supported with regular Failover because Layer 2 connectivity is required. Instead, see Failover for High Availability in the Public Cloud, on page 325.

ASA Virtual Failover for High Availability

When creating a failover pair with the ASA virtual, it is necessary to add the data interfaces to each ASA virtual in the same order. If the exact same interfaces are added to each ASA virtual, but in different order, errors may be presented at the ASA virtual Console. Failover functionality may also be affected

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 ASA 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.
- You can monitor up to 1025 interfaces on a unit, across all contexts.
- For Active/Standby Failover 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.
- For Active/Active failover, no two interfaces in the same context should be configured in the same ASR group.
- For Active/Active failover, you can define a maximum of two failover groups.
- For Active/Active failover, when removing failover groups, you must remove failover group 1 last. Failover group1 always contains the admin context. Any context not assigned to a failover group defaults to failover group 1. You cannot remove a failover group that has contexts explicitly assigned to it.
- 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 re-add the SNMPv3 users to the active unit to force the users to replicate to the new unit; or you can add the users directly on the new unit. Reconfigure each user by entering the **snmp-server user** *username group-name* **v3** command on the active unit or directly to the standby unit with the *priv-password* option and *auth-password* option in their unencrypted forms.
- 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 Failover 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 Failover do not dynamically synchronize the clock. Here are some examples of events when synchronization takes place:
  - A new Failover pair is created.
  - Failover is broken and re-created.
  - Communication over the failover link was disrupted and reestablished.
  - Failover status was manually changed at the CLI .
- Enabling Failover forces all routes to be deleted and are re-added after the Failover progression changes to the Active state. You could experience connection loss during this phase.
- If you enable failover on a standalone device, the data interfaces go down at negotiation state of failover, interrupting traffic.
- In the Failover 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 Failover 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 Failover 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 Failover 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 Failover 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.
- In OSPF in ASA, if a nearby switch is down and if the ASA interface is connected to the same switch, the interfaces in the firewall also go down with the switch failure. This is an expected behavior. This will trigger an High Availablity Failover, as designed.
- If the standby ASA is connected to a different switch, in this case when the interface comes up, the routing tables will be different than that of the active ASA. This will lead to an outage for a short duration (approximately 15-17 seconds) until the routes and adjacency is updated.
- 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.

Defaults for Failover

By default, the failover policy consists of the following:

- No HTTP replication in Stateful Failover.
- A single interface failure causes failover.
- The interface poll time is 5 seconds.
- The interface hold time is 25 seconds.
- The unit poll time is 1 second.
- The unit hold time is 15 seconds.
- Virtual MAC addresses are disabled in multiple context mode.
- Monitoring on all physical interfaces.

Configure Active/Standby Failover

To configure Active/Standby failover, configure basic failover settings on both the primary and secondary units. All other configuration occurs only on the primary unit, and is then synched to the secondary unit.

The **High Availability and Scalability Wizard** guides you through a step-by-step process of creating an Active/Standby failover configuration.

Procedure

Step 1 Choose **Wizards** > **High Availability and Scalability**. See select wizard guidelines in the following steps.

Step 2 On the **Failover Peer Connectivity and Compatibility** screen, enter the IP address of the peer unit. This address must be an interface that has ASDM access enabled on it.

By default, the peer address is assigned to be the standby address for the ASDM management interface.

Step 3 On the **LAN Link Configuration** screen:

- **Interface**—The interface can be a data physical interface, subinterface, or EtherChannel interface ID.On the Firepower 1010, the interface is a firewall interface ID; you cannot specify a switch port ID or VLAN ID. For the Firepower 4100/9300, you can use any data-type interface.
- **Active IP Address**—This IP address should be on an unused subnet. This subnet can be 31-bits (255.255.255.254) with only two IP addresses. 169.254.0.0/16 and fd00:0:0:*::/64 are internally used subnets, and you cannot use them for the failover or state links.
- **Standby IP Address**—This IP address must be on the same network as the active IP address.
- (Optional) **Communications Encryption**—Encrypt communications on the failover link. **Note**: Instead of a Secret Key, we recommend using an IPsec preshared key, which you can configure after you exit the wizard (see Modify the Failover Setup, on page 314).

Step 4 On the **State Link Configuration** screen, if you choose another interface for Stateful Failover:

- **Active IP Address**—This IP address should be on an unused subnet, different from the failover link. This subnet can be 31-bits (255.255.255.254) with only two IP addresses. 169.254.0.0/16 and fd00:0:0:*::/64 are internally used subnets, and you cannot use them for the failover or state links.
- **Standby IP Address**—This IP address must be on the same network as the active IP address.

Step 5 After you click **Finish**, the wizard shows the **Waiting for Config Sync** screen.

After the specified time period is over, the wizard sends the failover configuration to the secondary unit, and you see an information screen showing that failover configuration is complete.

- If you do not know if failover is already enabled on the secondary unit, then wait for the specified period.
- If you know failover is already enabled, click **Skip configuring peer**.
- If you know the secondary unit is not yet failover-enabled, click **Stop waiting** *xx* **more seconds**, and the failover bootstrap configuration is sent to the secondary unit immediately.

Configure Active/Active Failover

This section tells how to configure Active/Active failover.

The **High Availability and Scalability Wizard** guides you through a step-by-step process of creating an Active/Active failover configuration.

Procedure

Step 1 Choose **Wizards** > **High Availability and Scalability**. See select wizard guidelines in the following steps.

Step 2 In the **Failover Peer Connectivity and Compatibility Check** screen, the peer IP address must be an interface that has ASDM access enabled on it.

By default, the peer address is assigned to be the standby address for the interface to which ASDM is connected.

Step 3 In the **Security Context Configuration** screen, if you converted to multiple context mode as part of the wizard, you will only see the admin context. You can add other contexts after you exit the wizard.

Step 4 On the **LAN Link Configuration** screen:

- **Interface**—The interface can be a data physical interface, subinterface, redundant interface, or EtherChannel interface ID. For the ASA 5506H-X only, you can specify the Management 1/1 interface as the failover link. If you do so, you must save the configuration, and then reload the device. You then cannot use this interface for failover and also use the ASA Firepower module; the module requires the interface for management, and you can only use it for one function. For the Firepower 4100/9300, you can use any data-type interface.
- **Active IP Address**—This IP address should be on an unused subnet. This subnet can be 31-bits (255.255.255.254) with only two IP addresses. 169.254.0.0/16 and fd00:0:0:*::/64 are internally used subnets, and you cannot use them for the failover or state links.
- **Standby IP Address**—This IP address must be on the same network as the active IP address.
- (Optional) **Communications Encryption**—Encrypt communications on the failover link. **Note**: Instead of a Secret Key, we recommend using an IPsec preshared key, which you can configure after you exit the wizard (see Modify the Failover Setup, on page 314).

Step 5 On the **State Link Configuration** screen, if you choose another interface for Stateful Failover:

- **Active IP Address**—This IP address should be on an unused subnet, different from the failover link. This subnet can be 31-bits (255.255.255.254) with only two IP addresses. 169.254.0.0/16 and fd00:0:0:*::/64 are internally used subnets, and you cannot use them for the failover or state links.
- **Standby IP Address**—This IP address must be on the same network as the active IP address.

Step 6 After you click **Finish**, the wizard shows the **Waiting for Config Sync** screen.

After the specified time period is over, the wizard sends the failover configuration to the secondary unit, and you see an information screen showing that failover configuration is complete.

- If you do not know if failover is already enabled on the secondary unit, then wait for the specified period.
- If you know failover is already enabled, click **Skip configuring peer**.
- If you know the secondary unit is not yet failover-enabled, click **Stop waiting** *xx* **more seconds**, and the failover bootstrap configuration is sent to the secondary unit immediately.

Configure Optional Failover Parameters

You can customize failover settings as desired.

Configure Failover Criteria and Other Settings

See Defaults for Failover, on page 305 for the default settings for many parameters that you can change in this section. For Active/Active mode, you set most criteria per failover group. This section includes enabling HTTP replication per failover group for Active/Active mode; to configure HTTP replication for Active/Standby mode, see Modify the Failover Setup, on page 314.

Before you begin

- Configure these settings in the system execution space in multiple context mode.
- For Bidirectional Forwarding Detection (BFD) for unit health monitoring, see the following limitations:
  - Firepower 9300 and 4100 only.
  - Active/Standby only.
  - Routed mode only

Procedure

Step 1 Choose **Configuration** > **Device Management** > **High Availability and Scalability** > **Failover**.

Step 2 Disable the ability to make any configuration changes directly on the standby unit or context: on the **Setup** tab check the **Disable configuration changes on the standby unit** check box.

By default, configurations on the standby unit/context are allowed with a warning message.

Step 3 Under **BFD Health Check**, click **Manage** to define a BFD template to be used for failover health detection. The regular unit monitoring can cause false alarms when CPU usage is high. The BFD method is distributed, so high CPU does not affect its operation.

The **Configuration** > **Device Setup** > **Routing** > **BFD** > **Template** page opens. Click **Add** to create a single hop template; multi-hop is not supported. For the interval settings, you can specify milliseconds; microseconds are not supported. For template details, see Create the BFD Template, on page 817.

Step 4 Click the **Criteria** tab.

Step 5 Configure the unit poll times:

In the **Failover Poll Times** area:

- **Unit Failover**—The amount of time between hello messages among units. The range is between 1 and 15 seconds or between 200 and 999 milliseconds.
- **Unit Hold Time**—Sets the time during which a unit must receive a hello message on the failover link, or else the unit begins the testing process for peer failure. The range is between 1and 45 seconds or between 800 and 999 milliseconds. You cannot enter a value that is less than 3 times the polltime.

Note
Other settings on this pane apply only to Active/Standby mode. In Active/Active mode, you must configure the rest of the parameters per failover group.

Step 6 (Active/Active mode only) Click the **Active/Active** tab, then choose a failover group and click **Edit**.

Step 7 (Active/Active mode only) Change the preferred role of the failover group when used with preemption: click either **Primary** or **Secondary**.

If you used the wizard, failover group 1 is assigned to the primary unit, and failover group 2 is assigned to the secondary unit. If you want a non-standard configuration, you can specify different unit preferences if desired. These settings are only used in conjunction with the preempt setting. Both failover groups become active on the unit that boots first (even if it seems like they boot simultaneously, one unit becomes active first), despite the primary or secondary setting for the group.

Step 8 (Active/Active mode only) Configure failover group preemption: check the **Preempt after booting with optional delay of** check box.

Both failover groups become active on the unit that boots first (even if it seems like they boot simultaneously, one unit becomes active first), despite the primary or secondary setting for the group.

You can enter an optional delay value, which specifies the number of seconds the failover group remains active on the current unit before automatically becoming active on the designated unit. Valid values are from 1 to 1200.

If you manually fail over, the Peempt option is ignored.

Note
If Stateful Failover is enabled, the preemption is delayed until the connections are replicated from the unit on which the failover group is currently active.

Step 9 Configure the **Interface Policy**:

- **Number of failed interfaces that triggers failover**—Define a specific number of interfaces that must fail to trigger failover, from 1 to 250. When the number of failed monitored interfaces exceeds the value you specify, the ASA fails over.
- **Percentage of failed interfaces that triggers failover**—Define a percentage of configured interfaces that must fail to trigger failover. When the number of failed monitored interfaces exceeds the percentage you set, the ASA fails over.

Note
Do not use the **Use system failover interface policy** option. You can only set the policy per group at this time.

Step 10 (Active/Standby mode) Configure interface poll times:

In the **Failover Poll Time** area:

- **Monitored Interfaces**—Specifies the interface polltime: how long to wait between sending hello packets to the peer. The range is between 1and 15 seconds or 500 to 999 milliseconds. The default is 5 seconds.
- **Link State**—By default, each ASA in a failover pair checks the link state of its interfaces every 500 msec. You can customize the polltime; for example, if you set the polltime to 300 msec, the ASA can detect an interface failure and trigger failover faster. The range is between 300 and 799 milliseconds.

- **Interface Hold Time**—Sets the time (as a calculation) between the last-received hello message from the peer unit and the commencement of interface tests to determine the health of the interface. It also sets the duration of each interface test as *holdtime*/16. Valid values are from 5 to 75 seconds. The default is 5 times the polltime. You cannot enter a holdtime value that is less than five times the polltime.

  To calculate the time before starting interface tests (y):

  a. x = (*holdtime*/*polltime*)/2, rounded to the nearest integer. (.4 and down rounds down; .5 and up rounds up.)

  b. y = x**polltime*

  For example, if you use the default holdtime of 25 and polltime of 5, then y = 15 seconds.

For Active/Active mode, configure interface poll times on the **Add/Edit Failover Group** dialog box.

Step 11 (Active/Active mode only) Enable HTTP replication: check the **Enable HTTP replication** check box.

See Modify the Failover Setup, on page 314 section for the session replication rate.

Note
Because of a delay when deleting HTTP flows from the standby unit when using failover, the **show conn count** output might show different numbers on the active unit vs. the standby unit; if you wait several seconds and re-issue the command, you will see the same count on both units.

Step 12 Configure virtual MAC addresses:

- Active/Standby mode—click the **MAC Addresses** tab, and click **Add**.

  The **Add/Edit Interface MAC Address** dialog box appears.

- Active/Active mode—Go to the bottom of the **Active/Active** tab.

You can also set the MAC address using other methods, but 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.

a) Choose an interface from the **Physical Interface** drop-down list.
b) In the **Active MAC Address** field, type the new MAC address for the active interface.
c) In the **Standby MAC Address** field, type the new MAC address for the standby interface.
d) Click **OK**. (Active/Active mode only) Click **OK** again.

Step 13 Click **Apply**.

Configure Interface Monitoring and Standby Addresses

By default, monitoring is enabled on all physical interfaces, or for the Firepower 1010 and Secure Firewall 1210/1220, all VLAN interfaces. Firepower 1010 and Secure Firewall 1210/1220 switch ports are not elegible for interface monitoring.

You might want to exclude interfaces attached to less critical networks from affecting your failover policy.

You can monitor up to 1025 interfaces on a unit (across all contexts in multiple context mode).

If you did not configure the standby IP addresses in the wizard, you can configure them manually.

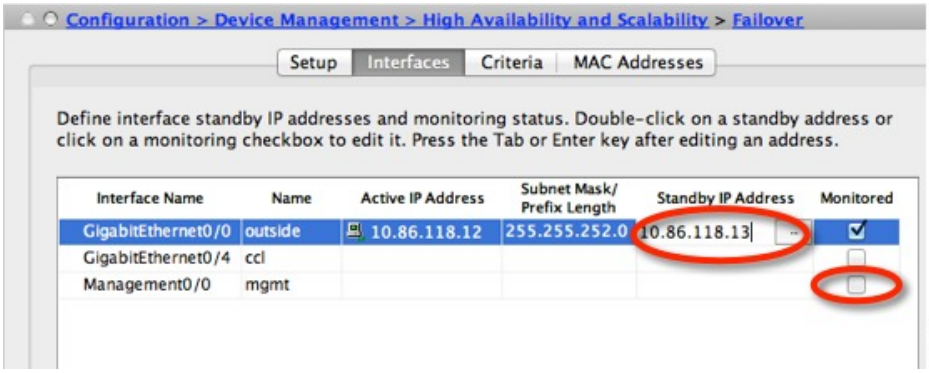
Before you begin

In multiple context mode, configure interfaces within each context.

Procedure

Step 1 In single mode, choose **Configuration** > **Device Management** > **High Availability** > **Failover** > **Interfaces**.

In multiple context mode, within a context choose **Configuration > Device Management > Failover > Interfaces**



A list of configured interfaces appears. The **Monitored** column displays whether or not an interface is monitored as part of your failover criteria. If it is monitored, a check appears in the **Monitored** check box.

The IP address for each interface appears in the **Active IP Address** column. If configured, the standby IP address for the interface appears in the **Standby IP Address** column. The failover link and state link do not display IP address; you cannot change those addresses from this tab.

Step 2 To disable monitoring of a listed interface, uncheck the **Monitored** check box for the interface.

Step 3 To enable monitoring of a listed interface, check the **Monitored** check box for the interface.

Step 4 For each interface that does not have a standby IP address, double-click the **Standby IP Address** field and enter an IP address into the field.

If you use a 31-bit subnet mask for point-to-point connections, do not configure a standby IP address.

Step 5 Click **Apply**.

Configure Support for Asymmetrically Routed Packets (Active/Active Mode)

When running in Active/Active failover, a unit might receive a return packet for a connection that originated through its peer unit. Because the ASA that receives the packet does not have any connection information for the packet, the packet is dropped. This drop most commonly occurs when the two ASAs in an Active/Active failover pair are connected to different service providers and the outbound connection does not use a NAT address.

You can prevent the return packets from being dropped by allowing asymmetrically routed packets. To do so, you assign the similar interfaces on each ASA to the same ASR group. For example, both ASAs connect

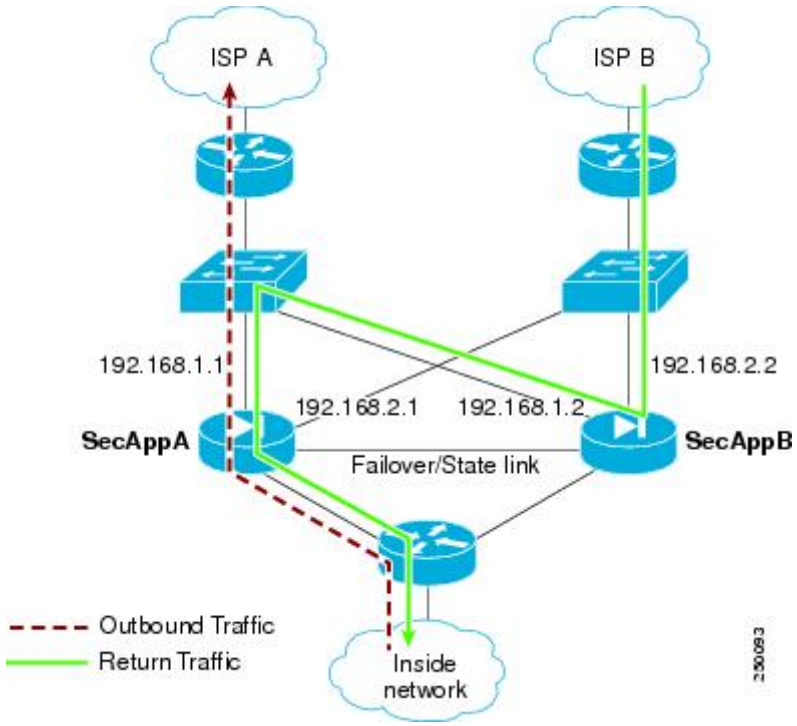
to the inside network on the inside interface, but connect to separate ISPs on the outside interface. On the primary unit, assign the active context outside interface to ASR group 1; on the secondary unit, assign the active context outside interface to the same ASR group 1. When the primary unit outside interface receives a packet for which it has no session information, it checks the session information for the other interfaces in standby contexts that are in the same group; in this case, ASR group 1. If it does not find a match, the packet is dropped. If it finds a match, then one of the following actions occurs:

- If the incoming traffic originated on a peer unit, some or all of the layer 2 header is rewritten and the packet is redirected to the other unit. This redirection continues as long as the session is active.
- If the incoming traffic originated on a different interface on the same unit, some or all of the layer 2 header is rewritten and the packet is reinjected into the stream.

Note This feature does not provide asymmetric routing; it restores asymmetrically routed packets to the correct interface.

The following figure shows an example of an asymmetrically routed packet.

Figure 58: ASR Example



1. An outbound session passes through the ASA with the active SecAppA context. It exits interface outside ISP-A (192.168.1.1).
2. Because of asymmetric routing configured somewhere upstream, the return traffic comes back through the interface outsideISP-B (192.168.2.2) on the ASA with the active SecAppB context.

3. Normally the return traffic would be dropped because there is no session information for the traffic on interface 192.168.2.2. However, the interface is configured as part of ASR group 1. The unit looks for the session on any other interface configured with the same ASR group ID.
4. The session information is found on interface outsideISP-A (192.168.1.2), which is in the standby state on the unit with SecAppB. Stateful Failover replicated the session information from SecAppA to SecAppB.
5. Instead of being dropped, the layer 2 header is rewritten with information for interface 192.168.1.1 and the traffic is redirected out of the interface 192.168.1.2, where it can then return through the interface on the unit from which it originated (192.168.1.1 on SecAppA). This forwarding continues as needed until the session ends.

Before you begin

- Stateful Failover—Passes state information for sessions on interfaces in the active failover group to the standby failover group.
- Replication HTTP—HTTP session state information is not passed to the standby failover group, and therefore is not present on the standby interface. For the ASA to be able to re-route asymmetrically routed HTTP packets, you need to replicate the HTTP state information.
- Perform this procedure within each active context on the primary and secondary units.
- You cannot configure both ASR groups and traffic zones within a context. If you configure a zone in a context, none of the context interfaces can be part of an ASR group.

Procedure

Step 1 On the primary unit active context, choose **Configuration** > **Device Setup** > **Routing** > **ASR Groups**.

Step 2 For the interface that receives asymmetrically routed packets, choose an **ASR Group ID** from the drop-down list.

Step 3 Click **Apply** to save your changes to the running configuration.

Step 4 Connect ASDM to the secondary unit, and choose the active context similar to the primary unit context.

Step 5 Choose **Configuration** > **Device Setup** > **Routing** > **ASR Groups**.

Step 6 For the similar interface on this unit, choose the same **ASR Group ID**.

Step 7 Click **Apply** to save your changes to the running configuration.

Manage Failover

This section describes how to manage Failover units after you enable Failover, including how to change the Failover setup and how to force failover from one unit to another.

Modify the Failover Setup

If you do not use the wizard, or want to change a setting, you can configure the failover setup manually. This section also includes the following options that are not included in the wizard, so you must configure them manually:

- IPsec preshared key for encrypting failover traffic
- HTTP replication rate
- HTTP replication (Active/Standby mode)

Before you begin

In multiple context mode, perform this procedure in the System execution space.

Procedure

Step 1 In single mode, choose **Configuration** > **Device Management** > **High Availability and Scalability** > **Failover** > **Setup**.

In multiple context mode, choose **Configuration > Device Management > Failover > Setup** in the System execution space.

Step 2 Check the **Enable Failover** check box.

Note
Failover is not actually enabled until you apply your changes to the device.

Step 3 To encrypt communications on the failover and state links, use one of the following options:

- **IPsec Preshared Key** (preferred)—The preshared key is used by IKEv2 to establish IPsec LAN-to-LAN tunnels on the failover links between the failover units. Note: failover LAN-to-LAN tunnels do not count against the IPsec (Other VPN) license.
- **Secret Key**—Enter the secret key used to encrypt failover communication. If you leave this field blank, failover communication, including any passwords or keys in the configuration that are sent during command replication, will be in clear text.

  Use 32 hexadecimal character key—To use a 32-hexadecimal key for the secret key, check this check box.

Step 4 In the **LAN Failover** area, set the following parameters for the failover link:

- **Interface**—Choose the interface to use for the failover link. Failover requires a dedicated interface, however you can share the interface with Stateful Failover.

  Only unconfigured interfaces or subinterfaces are displayed in this list and can be selected as the failover link. Once you specify an interface as the failover link, you cannot edit that interface in the Configuration > Interfaces pane.
- **Logical Name**—Specify the logical name of the interface used for failover communication, such as "failover". This name is informational.

- **Active IP**—Specify the active IP address for the interface. The IP address can be either an IPv4 or an IPv6 address. This IP address should be on an unused subnet.
- **Standby IP**—Specify the standby IP address for the interface, on the same subnet as the active IP address.
- **Subnet Mask**—Specify the subnet mask.
- **Preferred Role**—Select **Primary** or **Secondary** to specify whether the preferred role for this ASA is as the primary or secondary unit.

Step 5 (Optional) Configure the state link by doing the following:

- **Interface**—Choose the interface to use for the state link. You can choose an unconfigured interface or subinterface, the failover link, or the **--Use Named--** option.

  Note
  We recommend that you use two separate, dedicated interfaces for the failover link and the state link.

  If you choose an unconfigured interface or subinterface, you must supply the **Active IP**, **Subnet Mask**, **Logical Name**, and **Standby IP** for the interface.

  If you choose the failover link, you do not need to specify the **Active IP**, **Subnet Mask**, **Logical Name**, and **Standby IP** values; the values specified for the failover link are used.

  If you choose the **--Use Named--** option, the Logical Name field becomes a drop-down list of named interfaces. Choose the interface from this list. The **Active IP**, **Subnet Mask/Prefix Length**, and **Standby IP** values do not need to be specified. The values specified for the interface are used.
- **Logical Name**—Specify the logical name of the interface used for state communication, such as “state”. This name is informational.
- **Active IP**—Specify the active IP address for the interface. The IP address can be either an IPv4 or an IPv6 address. This IP address should be on an unused subnet, different from the failover link.
- **Standby IP**—Specify the standby IP address for the interface, on the same subnet as the active IP address.
- **Subnet Mask**—Specify the subnet mask.
- (Optional, Active/Standby only) **Enable HTTP Replication**—This option enables Stateful Failover to copy active HTTP sessions to the standby firewall. If you do not allow HTTP replication, then HTTP connections are disconnected in the event of a failover. In Active/Active mode, set the HTTP replication per failover group.

  Note
  Because of a delay when deleting HTTP flows from the standby unit when using failover, the **show conn count** output might show different numbers on the active unit vs. the standby unit; if you wait several seconds and re-issue the command, you will see the same count on both units.

Step 6 In the **Replication** area, set the session replication rate in connections per second. The minimum and maximum rates are determined by your model. The default is the maximum rate. To use the default, check the **Use Default check** box.

Step 7 Click **Apply**.

The configuration is saved to the device.

Step 8 If you are enabling failover, you see a dialog box to configure the failover peer.

- Click **No** if you want to connect to the failover peer later and configure the matching settings manually.
- Click **Yes** to let ASDM automatically configure the relevant failover settings on the failover peer. Provide the peer IP address in the **Peer IP Address** field.

Force Failover

To force the standby unit to become active, perform the following procedure.

Before you begin

In multiple context mode, perform this procedure in the System execution space.

Procedure

Step 1 To force failover at the unit level:

a) Choose the screen depending on your context mode:
   - In single context mode choose **Monitoring > Properties > Failover > Status**.
   - In multiple context mode, in the System choose **Monitoring > Failover > System**.

b) Click one of the following buttons:
   - Click **Make Active** to make the unit this unit.
   - Click **Make Standby** to make the other unit the active unit.

Step 2 (Active/Active mode only) To force failover at the failover group level:

a) In the System choose **Monitoring > Failover > Failover Group** #, where # is the number of the failover group you want to control.
b) Click one of the following buttons
   - Click **Make Active** to make the failover group active on this unit.
   - Click **Make Standby** to make the failover group active on the other unit.

Disable Failover

Disabling failover on one or both units causes the active and standby state of each unit to be maintained until you reload. For an Active/Active failover pair, the failover groups remain in the active state on whichever unit they are active, no matter which unit they are configured to prefer.

See the following characteristics when you disable failover:

- The standby unit/context remains in standby mode so that both units do not start passing traffic (this is called a pseudo-standby state).
- The standby unit/context continues to use its standby IP addresses even though it is no longer connected to an active unit/context.
- The standby unit/context continues to listen for a connection on the failover link. If failover is re-enabled on the active unit/context, then the standby unit/context resumes ordinary standby status after re-synchronizing the rest of its configuration.
- Do not enable failover manually on the standby unit to make it active; instead see Force Failover, on page 316. If you enable failover on the standby unit, you will see a MAC address conflict that can disrupt IPv6 traffic.
- To truly disable failover, save the no failover configuration to the startup configuration, and then reload.

Before you begin

In multiple context mode, perform this procedure in the system execution space.

Procedure

Step 1 In single mode, choose **Configuration** > **Device Management** > **High Availability and Scalability** > **Failover** > **Setup**.

In multiple context mode, choose **Configuration > Device Management > Failover > Setup** in the System execution space.

Step 2 Uncheck the **Enable Failover** check box.

Step 3 Click **Apply**.

Step 4 To completely disable failover, save the configuration and reload:

a) Click the **Save** button.

b) Choose **Tools > System Reload** and reload the ASA.

Restore a Failed Unit

To restore a failed unit to an unfailed state, perform the following procedure.

Before you begin

In multiple context mode, perform this procedure in the System execution space.

Procedure

Step 1 To restore failover at the unit level:

a) Choose the screen depending on your context mode:

- In single context mode choose **Monitoring > Properties > Failover > Status**.
- In multiple context mode, in the System choose **Monitoring > Failover > System**.

b) Click **Reset Failover**.

Step 2 (Active/Active mode only) To reset failover at the failover group level:

a) In the System choose **Monitoring > Failover > Failover Group #**, where # is the number of the failover group you want to control.

b) Click **Reset Failover**.

Re-Sync the Configuration

Replicated commands are stored in the running configuration. To save replicated commands to the flash memory on the standby unit, choose **File > Save Running Configuration to Flash**.

Monitoring Failover

This section lets you monitor the Failover status.

Failover Messages

When a failover occurs, both ASAs send out system messages.

Failover Syslog Messages

The ASA issues a number of syslog messages related to failover at priority level 2, which indicates a critical condition. To view these messages, see the syslog messages guide. The ranges of message IDs associated with failover are: 101xxx, 102xxx, 103xxx, 104xxx, 105xxx, 210xxx, 311xxx, 709xxx, 727xxx. For example, 105032 and 105043 indicate a problem with the failover link.

Note During failover, the ASA logically shuts down and then brings up interfaces, generating syslog messages 411001 and 411002. This is normal activity.

Failover Debug Messages

To see debug messages, enter the **debug fover** command. See the command reference for more information.

Note Because debugging output is assigned high priority in the CPU process, it can drastically affect system performance. For this reason, use the **debug fover** commands only to troubleshoot specific problems or during troubleshooting sessions with Cisco TAC.

SNMP Failover Traps

To receive SNMP syslog traps for failover, configure the SNMP agent to send SNMP traps to SNMP management stations, define a syslog host, and compile the Cisco syslog MIB into your SNMP management station.

Monitoring Failover Status



Note After a failover event, you should either re-launch ASDM or switch to another device in the Devices pane and then come back to the original ASA to continue monitoring the device. This action is necessary because the monitoring connection is not re-established when ASDM is disconnected from and then reconnected to the device.

Choose **Monitoring > Properties > Failover** to monitor Active/Standby failover.

Use the following screens in the **Monitoring > Properties > Failover** area to monitor Active/Active failover.

System

The System pane displays the failover state of the system. You can also control the failover state of the system by:

- Toggling the active/standby state of the device.
- Resetting a failed device.
- Reloading the standby unit.

Fields

Failover state of the system—*Display only*. Displays the failover state of the ASA. The information shown is the same output you would receive from the **show failover** command. Refer to the command reference for more information about the displayed output.

The following actions are available on the System pane:

- Make Active—Click this button to make the ASA the active unit in an active/standby configuration. In an active/active configuration, clicking this button causes both failover groups to become active on the ASA.
- Make Standby—Click this button to make the ASA the standby unit in an active/standby pair. In an active/active configuration, clicking this button causes both failover groups to go to the standby state on the ASA.
- Reset Failover—Click this button to reset a system from the failed state to the standby state. You cannot reset a system to the active state. Clicking this button on the active unit resets the standby unit.
- Reload Standby—Click this button to force the standby unit to reload.
- Refresh—Click this button to refresh the status information in the Failover state of the system field.

Failover Group 1 and Failover Group 2

The Failover Group 1 and Failover Group 2 panes display the failover state of the selected group.You can also control the failover state of the group by toggling the active/standby state of the group or by resetting a failed group.

Fields

Failover state of Group[*x*]—*Display only*. Displays the failover state of the selected failover group. The information shown is the same as the output you would receive from the **show failover group** command.

You can perform the following actions from this pane:

- Make Active—Click this button to make the failover group active unit on the ASA.
- Make Standby—Click this button to force the failover group into the standby state on the ASA.
- Reset Failover—Click this button to reset a system from the failed state to the standby state. You cannot reset a system to the active state. Clicking this button on the active unit resets the standby unit.
- Refresh—Click this button to refresh the status information in the Failover state of the system field.

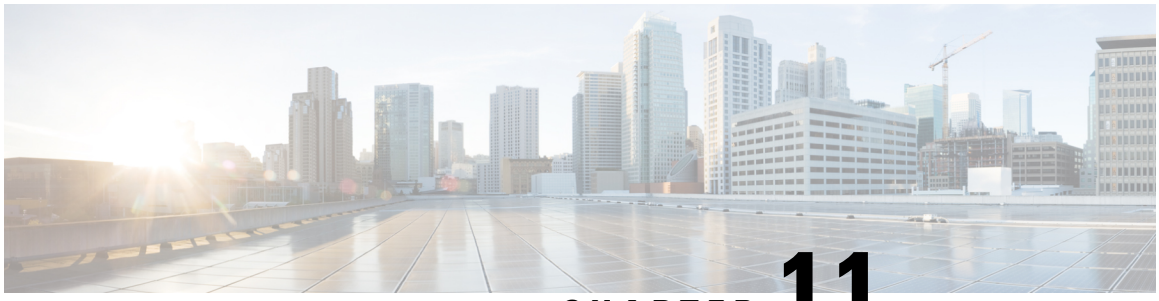
History for Failover

Feature Name	Releases	Feature Information
Active/Standby failover	7.0(1)	This feature was introduced.
Active/Active failover	7.0(1)	This feature was introduced.
Support for a hex value for the failover key	7.0(4)	You can now specify a hex value for failover link encryption. We modified the following screen: Configuration > Device Management > High Availability > Failover > Setup.
Support for the master passphrase for the failover key	8.3(1)	The failover key now supports the master passphrase, which encrypts the shared key in the running and startup configuration. If you are copying the shared secret from one ASA to another, for example from the **more system:running-config** command, you can successfully copy and paste the encrypted shared key. **Note** The **failover key** shared secret shows as ***** in **show running-config** output; this obscured key is not copyable. There were no ASDM changes.
IPv6 support for failover added.	8.2(2)	We modified the following screens: Configuration > Device Management > High Availability > Failover > Setup Configuration > Device Management > High Availability > Failover > Interfaces

Feature Name	Releases	Feature Information
Change to failover group unit preference during "simultaneous" bootup.	9.0(1)	Earlier software versions allowed “simultaneous” boot up so that the failover groups did not require the **preempt** command to become active on the preferred unit. However, this functionality has now changed so that both failover groups become active on the first unit to boot up.
Support for IPsec LAN-to-LAN tunnels to encrypt failover and state link communications	9.1(2)	Instead of using the proprietary encryption for the failover key, you can now use an IPsec LAN-to-LAN tunnel for failover and state link encryption. **Note** Failover LAN-to-LAN tunnels do not count against the IPsec (Other VPN) license. We modified the following screen: **Configuration > Device Management > High Availability > Failover > Setup**.
Disable health monitoring of a hardware module	9.3(1)	By default, the ASA monitors the health of an installed hardware module such as the ASA FirePOWER module. If you do not want a hardware module failure to trigger failover, you can disable module monitoring. We modified the following screen: **Configuration > Device Management > High Availability and Scalability > Failover > Interfaces**.
Lock configuration changes on the standby unit or standby context in a failover pair	9.3(2)	You can now lock configuration changes on the standby unit (Active/Standby failover) or the standby context (Active/Active failover) so you cannot make changes on the standby unit outside normal configuration syncing. We modified the following screen: **Configuration > Device Management > High Availability and Scalability > Failover > Setup**.
Enable use of the Management 1/1 interface as the failover link on the ASA 5506H	9.5(1)	On the ASA 5506H only, you can now configure the Management 1/1 interface as the failover link. This feature lets you use all other interfaces on the device as data interfaces. Note that if you use this feature, you cannot use the ASA Firepower module, which requires the Management 1/1 interface to remain as a regular management interface. We modified the following screen: **Configuration > Device Management > High Availability and Scalability > Failover > Setup**
Carrier Grade NAT enhancements now supported in failover and ASA clustering	9.5(2)	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). This feature is now supported in failover and ASA cluster deployments. We did not modify any screens.

Feature Name	Releases	Feature Information
Improved sync time for dynamic ACLs from Secure Client when using Active/Standby failover	9.6(2)	When you use Secure Client on a failover pair, then the sync time for the associated dynamic ACLs (dACLs) to the standby unit is now improved. Previously, with large dACLs, the sync time could take hours during which time the standby unit is busy syncing instead of providing high availability backup. We did not modify any screens.
Stateful failover for Secure Client connections in multiple context mode	9.6(2)	Stateful failover is now supported for Secure Client connections in multiple context mode. We did not modify any screens.
Interface link state monitoring polling for failover now configurable for faster detection	9.7(1)	By default, each ASA in a failover pair checks the link state of its interfaces every 500 msec. You can now configure the polling interval, between 300 msec and 799 msec; for example, if you set the polltime to 300 msec, the ASA can detect an interface failure and trigger failover faster. We modified the following screen: **Configuration** > **Device Management** > **High Availability and Scalability** > **Failover** > **Criteria**
Bidirectional Forwarding Detection (BFD) support for Active/Standby failover health monitoring on the Firepower 9300 and 4100	9.7(1)	You can enable Bidirectional Forwarding Detection (BFD) for the failover health check between two units of an Active/Standby pair on the Firepower 9300 and 4100. Using BFD for the health check is more reliable than the default health check method and uses less CPU. We modified the following screen: **Configuration** > **Device Management** > **High Availability and Scalability** > **Failover** > **Setup**
Disable failover delay	9.15(1)	When you use bridge groups or IPv6 DAD, when a failover occurs the new active unit waits up to 3000 ms for the standby unit to finish networking tasks and transition to the standby state. Then the active unit can start passing traffic. To avoid this delay, you can disable the waiting time, and the active unit will start passing traffic before the standby unit transitions. New/Modified screens: **Configuration** > **Device Management** > **High Availability and Scalability** > **Failover** > **Enable switchover waiting for peer state**
Config-Sync Optimization feature for faster HA peering	9.18(1)	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 Name	Releases	Feature Information
Heartbeat module redundancy	9.20(1)	Introduced an additional heartbeat module in the data plane of ASA high availability. This heartbeat module helps to avoid false failovers or split-brain scenarios that can happen due to traffic congestion in the control plain or CPU overload.



CHAPTER 11

Failover for High Availability in the Public Cloud

This chapter describes how to configure Active/Backup failover to accomplish high availability of the ASA virtual in a public cloud environment, such as Microsoft Azure.

- About Failover in the Public Cloud, on page 325
- Licensing for Failover in the Public Cloud, on page 329
- Defaults for Failover in the Public Cloud, on page 329
- About ASA Virtual High Availability in Microsoft Azure, on page 330
- Configure Active/Backup Failover, on page 333
- Configure Optional Failover Parameters, on page 335
- Manage Failover in the Public Cloud, on page 336
- Monitor Failover in the Public Cloud, on page 338
- History for Failover in the Public Cloud, on page 339

About Failover in the Public Cloud

To ensure redundancy, you can deploy the ASA virtual in a public cloud environment in an Active/Backup high availability (HA) configuration. HA in the public cloud implements a stateless Active/Backup solution that allows for a failure of the active ASA virtual to trigger an automatic failover of the system to the backup ASA virtual.

The following list describes the primary components in the HA public cloud solution:

- **Active ASA Virtual**—The ASA virtual in the HA pair that is set up to handle the firewall traffic for the HA peers.
- **Backup ASA Virtual**—The ASA virtual in the HA pair that is not handling firewall traffic and takes over as the active ASA virtual in the event of an active ASA virtual failure. It is referred to as a Backup rather than a Standby because it is does not take on the identify of its peer in the event of a failover.
- **HA Agent**—A lightweight process that runs on the ASA virtual and determines the HA role (active/backup) of an ASA virtual, detects failures of its HA peer, and performs actions based on its HA role.

On the physical ASA and the non-public cloud virtual ASA, the system handles failover conditions using gratuitous ARP requests where the backup ASA sends out a gratuitous ARP indicating it is now associated with the active IP and MAC addresses. Most public cloud environments do not allow broadcast traffic of this

nature. For this reason, an HA configuration in the public cloud requires ongoing connections be restarted when failover happens.

The health of the active unit is monitored by the backup unit to determine if specific failover conditions are met. If those conditions are met, failover occurs. The failover time can vary from a few seconds to over a minute depending on the responsiveness of the public cloud infrastructure.

About Active/Backup Failover

In Active/Backup failover, one unit is the active unit. It passes traffic. The backup unit does not actively pass traffic or exchange any configuration information with the active unit. Active/Backup failover lets you use a backup ASA virtual device to take over the functionality of a failed unit. When the active unit fails, it changes to the backup state while the backup unit changes to the active state.

Primary/Secondary Roles and Active/Backup Status

When setting up Active/Backup failover, you configure one unit to be primary and the other as secondary. At this point, the two units act as two separate devices for device and policy configuration, as well as for events, dashboards, reports, and health monitoring.

The main differences between the two units in a failover pair are related to which unit is active and which unit is backup, namely which unit actively passes traffic. Although both units are capable of passing traffic, only the primary unit responds to Load Balancer probes and programs any configured routes to use it as a route destination. The backup unit's primary function is to monitor the health of the primary unit. The primary unit always becomes the active unit if both units start up at the same time (and are of equal operational health).

Failover Connection

The backup ASA virtual monitors the health of the active ASA virtual using a failover connection established over TCP:

- The active ASA virtual acts as a connection server by opening a *listen port*.
- The backup ASA virtual connects to the active ASA virtual using *connect port*.
- Typically the *listen port* and the *connect port* are the same, unless your configuration requires some type of network address translation between the ASA virtual units.

The state of the failover connection detects the failure of the active ASA virtual. When the backup ASA virtual sees the failover connection come down, it considers the active ASA virtual as *failed*. Similarly, if the backup ASA virtual does not receive a response to a keepalive message sent to the active unit, it considers the active ASA virtual as *failed*

Related Topics

Polling and Hello Messages

The backup ASA virtual sends Hello messages over the failover connection to the active ASA virtual and expects a Hello Response in return. Message timing uses a polling interval, the time period between the receipt of a Hello Response by the backup ASA virtual unit and the sending of the next Hello message. The receipt

of the response is enforced by a receive timeout, called the hold time. If the receipt of the Hello Response times out, the active ASA virtual is considered to have failed.

The polling and hold time intervals are configurable parameters; see Configure Active/Backup Failover, on page 333.

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 backup 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 backup unit.

Failover Events

In Active/Backup 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 backup unit, and any special notes about the failover condition and actions.

Table 21: Failover Events

Failure Event	Policy	Active Action	Backup Action	Notes
Backup unit sees a failover connection close	Failover	n/a	Become active Mark active as failed	This is the standard failover use case.
Active unit sees a failover connection close	No failover	Mark backup as failed	n/a	Failover to an inactive unit should never occur.
Active unit sees a TCP timeout on failover link	No failover	Mark backup as failed	No action	Failover should not occur if the active unit is not getting a reponse from the backup unit.
Backup unit sees a TCP timeout on failover link	Failover	n/a	Become active Mark active as failed Try to send failover command to active unit	The backup unit assumes that the active unit is unable to continue operation and takes over. In case the active unit is still up, but fails to send a response in time, the backup unit sends the failover command to the active unit.

Failure Event	Policy	Active Action	Backup Action	Notes
Active Authentication failed	No failover	No action	No action	Because the backup unit is changing the route tables, it is the only unit that needs to be authenticated to Azure. It does not matter if the active unit is authenticated to Azure or not.
Backup Authentication failed	No failover	Mark backup as unauthenticated	No action	Failover cannot happen if the backup unit is not authenticated to Azure.
Active unit initiates intentional failover	Failover	Become backup	Become active	The active unit initiates failover by closing the Failover Link Connection. The backup unit sees the connection close and becomes the active unit.
Backup unit initiates intentional failover	Failover	Become backup	Become active	The backup unit initiates failover by sending a failover message to the active unit. When the active unit sees the message, it closes the connection and becomes the backup unit. The backup unit sees the connection close and becomes the active unit.
Formerly active unit recovers	No failover	Become backup	Mark mate as backup	Failover should not occur unless absolutely necessary.
Active unit sees failover message from backup unit	Failover	Become backup	Become active	Can occur if a manual failover was initiated by a user; or the backup unit saw the TCP timeout, but the active unit is able to receive messages from the backup unit.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

ASA Virtual Failover for High Availability in the Public Cloud

To ensure redundancy, you can deploy the ASA virtual in a public cloud environment in an Active/Backup high availability (HA) configuration.

- Supported only on the Microsoft Azure public cloud; when configuring the ASA virtual VM, the maximum supported number of vCPUs is 8; and the maximum supported memory is 64GB RAM. See the ASA virtual Getting Started Guide for comprehensive list of supported instances.
- Implements a stateless Active/Backup solution that allows for a failure of the active ASA virtual to trigger an automatic failover of the system to the backup ASA virtual.

Limitations

- Failover is on the order of seconds rather than milliseconds.
- The HA role determination and the ability to participate as an HA unit depends on TCP connectivity between HA peers and between an HA unit and the Azure infrastructure. There are several situations where an ASA virtual will not be able participate as an HA unit:
  - The inability to establish a failover connection to its HA peer.
  - The inability to retrieve an authentication token from Azure.
  - The inability to authenticate with Azure.
- There is no synching of the configuration from the Active unit to the Backup unit. Each unit must be configured individually with similar configurations for handling failover traffic.
- Failover route-table limitations

  With respect to route-tables for HA in the public cloud:
  - You can configure a maximum of 16 route-tables.
  - Within a route-table, you can configure a maximum of 64 routes.

  In each case the system alerts you when you have reached the limit, with the recommendation to remove a route-table or route and retry.
- No ASDM support.
- No IPSec Remote Access VPN support.



Note See the Cisco Adaptive Security Virtual Appliance (ASAv) Quick Start Guide for information about supported VPN topologies in the public cloud.

- ASA Virtual VM instances must be in the same availability set. If you are a current ASA virtual user in Azure, you will not be able to upgrade to HA from an existing deployment. You have to delete your instance and deploy the ASA virtual 4 NIC HA offering from the Azure Marketplace.

Licensing for Failover in the Public Cloud

The ASA virtual uses Cisco Smart Software Licensing. A smart license is required for regular operation. Each ASA virtual must be licensed independently with an ASA virtual platform license. Until you install a license, throughput is limited to 100 Kbps so you can perform preliminary connectivity tests. See the Cisco ASA Series Feature Licenses page to find precise licensing requirements for the ASA virtual.

Defaults for Failover in the Public Cloud

By default, the failover policy consists of the following:

- Stateless failover only.

- Each unit must be configured individually with similar configurations for handling failover traffic.
- The failover TCP control port number is 44442.
- The Azure Load Balancer health probe port number is 44441.
- The unit poll time is 5 seconds.
- The unit hold time is 15 seconds.
- The ASA virtual responds to health probes on the primary interface (Management 0/0).
- The ASA virtual authentication with Azure Service Principal is performed on the primary interface (Management 0/0).



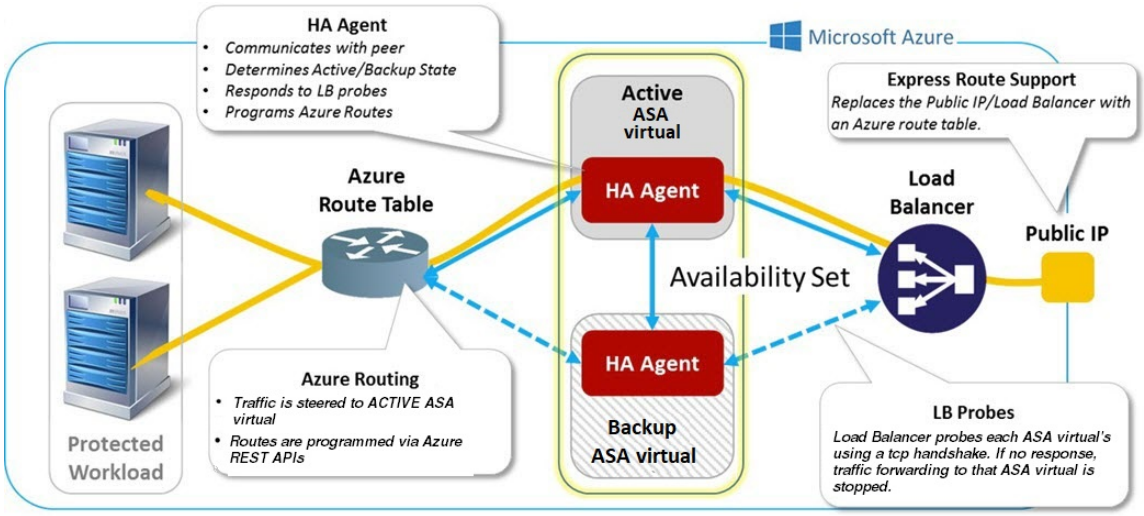
Note See Configure Optional Failover Parameters, on page 335 for options to change the failover port number, health probe port number, poll times, and primary interface.

About ASA Virtual High Availability in Microsoft Azure

The following figure shows a high-level view of an ASA virtual HA deployment in Azure. A protected workload sits behind two ASA virtual instances in an Active/Backup failover configuration. An Azure Load Balancer probes both of the ASA virtual units using a three-way TCP handshake. The active ASA virtual completes the three way handshake indicating that it is healthy, while the backup ASA virtual intentionally does not respond. By not responding to the Load Balancer, the backup ASA virtual appears unhealthy to the Load Balancer, which in turn does not send traffic to it.

On failover, the active ASA virtual stops responding to the Load Balancer probes and the backup ASA virtual starts responding, causing all new connections to be sent to the backup ASA virtual. The backup ASA virtual sends API requests to the Azure Fabric to modify the route table, redirecting traffic from the active unit to the backup unit. At this point, the backup ASA virtual becomes the active unit and the active unit becomes the backup unit or is offline, depending on the reason for the failover.

Figure 59: ASA Virtual HA Deployment in Azure



To be able to automatically make API calls to modify Azure route tables, the ASA virtual HA units need to have Azure Active Directory credentials. Azure employs the concept of a Service Principal which, in simple terms, is a service account. A Service Principal allows you to provision an account with only enough permissions and scope to run a task within a predefined set of Azure resources.

There are two steps to enable your ASA virtual HA deployment to manage your Azure subscription using a Service Principal:

1. Create an Azure Active Directory application and Service Principal; see About the Azure Service Principal, on page 331.
2. Configure the ASA virtual instances to authenticate with Azure using a Service Principal; see Configure Active/Backup Failover, on page 333.

Related Topics

See the Azure documentation for more informaion about the Load Balancer.

About the Azure Service Principal

When you have an application that needs to access or modify Azure resources, such as route tables, you must set up an Azure Active Directory (AD) application and assign the required permissions to it. This approach is preferable to running the application under your own credentials because:

- You can assign permissions to the application identity that are different than your own permissions. Typically, these permissions are restricted to exactly what the application needs to do.
- You do not have to change the application's credentials if your responsibilities change.
- You can use a certificate to automate authentication when executing an unattended script.

When you register an Azure AD application in the Azure portal, two objects are created in your Azure AD tenant: an application object, and a service principal object.

- **Application object**—An Azure AD application is defined by its one and only application object, which resides in the Azure AD tenant where the application was registered, known as the application's "home" tenant.
- **Service principal object**—The service principal object defines the policy and permissions for an application's use in a specific tenant, providing the basis for a security principal to represent the application at run-time.

Azure provides instructions on how to create an Azure AD application and service principal in the *Azure Resource Manager Documentation*. See the following topics for complete instructions:

- Use portal to create an Azure Active Directory application and service principal that can access resources
- Use Azure PowerShell to create a service principal to access resource



Note After you set up the service principal, obtain the **Directory ID**, **Application ID**, and **Secret key**. These are required to configure Azure authentication credentials; see Configure Active/Backup Failover, on page 333.

Configuration Requirements for ASA Virtual High Availability in Azure

To deploy a configuration similar to the one described in #unique_480 unique_480_Connect_42_fig_cgx_dlh_h1b you need the following :

- Azure Authentication information (see About the Azure Service Principal, on page 331):
  - Directory ID
  - Application ID
  - Secret key
- Azure route information (see Configure Azure Route Tables, on page 335):
  - Azure Subscription ID
  - Route table resource group
  - Table names
  - Address prefix
  - Next hop address
- ASA configuration (see Configure Active/Backup Failover, on page 333, Defaults for Failover in the Public Cloud, on page 329):
  - Active/Backup IP addresses
  - HA Agent communication port
  - Load Balancer probe port
  - Polling intervals

Note Configure basic failover settings on both the primary and secondary units. There is no synching of configuration from the primary unit to the secondary unit. Each unit must be configured individually with similar configurations for handling failover traffic.

Configure Active/Backup Failover

To configure Active/Backup failover, configure basic failover settings on both the primary and secondary units. There is no synching of configuration from the primary unit to the secondary unit. Each unit must be configured individually with similar configurations for handling failover traffic.

Before you begin

- Deploy your ASA virtual HA pair in an Azure Availability Set.
- Have your Azure environment information available, including your Azure Subscription ID and Azure authentication credentials for the Service Principal.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **High Availability and Scalability** > **Failover**.

Step 2 On the **Cloud** tab, check the **Unit** check box to expand the **Failover Unit** drop-down options.

Step 3 From the **Failover Unit** drop-down menu, choose **primary**.

The primary unit will assume the active HA role when both HA units come up at the same time.

Step 4 (Optional) Check the **Port** check box to expand the **Control** and **Probe** fields.

a) Enter a valid TCP control port in the **Control** field; or keep the default, port 44442.

The control port establishes the TCP failover connection established between the active ASA virtual and the backup ASA virtual.

b) Enter a valid TCP probe port in the **Probe** field; or keep the default, port 44441.

The probe port is the TCP port used as the destination port for Azure Load Balancer probes.

Step 5 (Optional) Check the **Time** check box to expand the **Poll Time** and **Hold Time** fields.

a) Enter a valid time (in seconds) in the **Poll Time** field; or keep the default, 5 seconds.

The poll time range is between 1 and 15 seconds. With a faster poll time, the ASA can detect failure and trigger failover faster. However, faster detection can cause unnecessary switchovers when the network is temporarily congested.

b) Enter a valid time (in seconds) in the **Hold Time** field; or keep the default, 15 seconds.

The hold time determines how long it takes from the time a hello packet is missed to when the unit is marked as failed. The hold time range is between 3 and 60 seconds.You cannot enter a holdtime value that is less than 3 times the unit poll time.

Step 6 Check the **Peer** check box to expand the **Peer IP-Address** and **Peer Port** fields.

a) Enter the IP address used to establish a TCP failover control connection to the HA peer in the **Peer IP-Address** field.

b) (Optional) Enter a valid TCP control port in the **Peer Port** field; or keep the default, port 44442..

The peer port establishes the TCP failover connection established between the active ASA virtual and the backup ASA virtual.

Step 7 Check the **Authentication** check box to expand the **Application-id**, **Directory-id**, and **Key** fields.

You can configure authentication credentials for an Azure Service Principal that allows your ASA virtual HA peers to access or modify Azure resources, such as route tables. Service Principals allow you to provision an Azure account that possesses the minimum permissions to perform a task within a predefined set of Azure resources. In the case of ASA virtual HA, it is limited to the permissions needed to modify user-defined routes; see About the Azure Service Principal, on page 331.

a) Enter the Azure application ID for the Azure Service Principal in the **Application-id** field.

You need this application ID when you request an access key from the Azure infrastructure.

b) Enter the Azure directory ID for the Azure Service Principal in the **Directory-id** field.

You need this directory ID when you request an access key from the Azure infrastructure.

c) Enter the Azure secret key for the Azure Service Principal in the **Key** field.

You need this secret key when requesting an access key from the Azure infrastructure. If the **Encrypt** field is checked, the secret key is encrypted in the running configuration.

Step 8 Check the **Subscription** check box to expand the **Sub-id** field.

This is the Subscription ID for the account to which the route tables that require updating belong.

Step 9 Check the **Enable Cloud Failover** check box.

Step 10 Click **Apply**.

Failover is not actually enabled until you apply your changes to the device.

Step 11 If you know the secondary unit is not yet failover-enabled, connect to the secondary ASA virtual from the **Device List**, or start a new ASDM session using the IP address of the ASA virtual: **https**://*asa_ip_address*/**admin**.

Step 12 Repeat steps 1 through 10 to configure Active/Backup failover on the secondary unit.

There is no synching of configuration from the primary unit to the secondary unit. Each unit must be configured individually with similar configurations for handling failover traffic.

Failover is not actually enabled until you apply your changes to the device.

What to do next

Configure additional parameters as needed:

- Configure Azure route information; see Configure Azure Route Tables, on page 335.

Configure Optional Failover Parameters

You can customize failover settings as necessary.

Configure Azure Route Tables

The route table configuration consists of information about Azure user-defined routes that need to be updated when the ASA virtual assumes the active role. On failover, you want to direct internal routes to the active unit, which uses the configured route table information to automatically direct the routes to itself.

Note You need to configure any Azure route table information on both the active and backup units.

Before you begin

- Configure these settings on both the primary and secondary units. There is no synching of configuration from the primary unit to the secondary unit.
- Have your Azure environment information available, including your Azure Subscription ID and Azure authentication credentials for the Service Principal.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **High Availability and Scalability** > **Failover**.

Step 2 Click the **Route-Table** tab and click **Add**.

a) In the **Route Table Name** field, enter a name for the route table.

You can configure up to 16 route tables. Alternately, you can edit or delete entries to the Route Table list.

b) (Optional) In the **Sub-id** field, enter an Azure Subscription ID.

You can update user-defined routes in more than one Azure subscription by specifying the corresponding Azure Subscription ID here. If you enter the **Route Table Name** without specifying an Azure Subscription ID, the global parameter is used.

Note
You enter the Azure Subscription ID when you configure Active/Backup failover from **Configuration** > **Device Management** > **High Availability and Scalability** > **Failover**; see Configure Active/Backup Failover, on page 333.

Step 3 Click **Route-Table-Mode**. You can add, edit, or delete route entries to the route tables.

Step 4 Click **Add**.

Enter the following values for Azure user-defined routes:

a) From the **Route Table** drop-down list, choose a route table.
b) In the **Azure Resource Group** field, enter the name of the Azure Resource Group that contains the Azure route table.

c) In the **Route Name** field, enter a unique name for the route.
d) In the **Prefix Address/Mask** field, enter the IP address prefix in CIDR notation.
e) In the **Next Hop Address** field, enter the next-hop address. This is an interface IP address on the ASA virtual.

Note
You can configure up to 64 routes.

Step 5 Click **Apply** to save your changes.

Manage Failover in the Public Cloud

This section describes how to manage Failover units in the Cloud after you enable failover, including how to change to force failover from one unit to another.

Force Failover

To force the standby unit to become active, perform the following command.

Before you begin

Use this command in the system execution space in single context mode.

Procedure

Step 1 Choose **Monitoring** > **Properties** > **Failover** > **Status**.
Step 2 To force failover at the unit level, click one of the following buttons:

- Click **Make Active** to make this unit the *active* unit.
- Click **Make Standby** to make this unit the *standby* unit.

Update Routes

If the state of the routes in Azure is inconsistent with the ASA virtual in the *active* role, you can force route updates on the ASA virtual:

Before you begin

Use this command in the system execution space in single context mode.

Procedure

Step 1 Choose **Monitoring** > **Properties** > **Failover** > **Status**.

Step 2 Click **Update Route**.

This command is only valid on the ASA virtual in the *active* role. If authentication fails the output will be `Route changes failed`.

Validate Azure Authentication

For a successful ASA virtual HA deployment in Azure, the Service Principal configuration must be complete and accurate. Without proper Azure authorization, the ASA virtual units will be unable to access resources to handle failover and to perform route updates. You can test your failover configuration to detect errors related to the following elements of your Azure Service Principal:

- Directory ID
- Application ID
- Authentication Key

Before you begin

Use this command in the system execution space in single context mode.

Procedure

Step 1 Choose **Monitoring** > **Properties** > **Failover** > **Status**.

Step 2 Click **Test Authentication**.

If authentication fails the command output will be `Authentication Failed`.

If the Directory ID or Application ID is not configured properly, Azure will not recognize the resource addressed in the REST request to obtain an authentication token. The event history for this condition entry will read:

```
Error Connection - Unexpected status in response to access token request: Bad Request
```

If the Directory ID or Application ID are correct, but the authentication key is not configured properly, Azure will not grant permission to generate the authentication token. The event history for this condition entry will read:

```
Error Connection - Unexpected status in response to access token request: Unauthorized
```

Monitor Failover in the Public Cloud

This section explains how you monitor the failover status.

Failover Status

Note After a failover event you should either re-launch ASDM or switch to another device in the Devices pane and then come back to the original ASA to continue monitoring the device. This action is necessary because the monitoring connection does not become re-established when ASDM is disconnected from and then reconnected to the device.

- Choose **Monitoring** > **Properties** > **Failover** > **Status** and click **Failover Status** to monitor Active/Backup failover status.
- Choose **Monitoring** > **Properties** > **Failover** > **History** to display failover event history with a timestamp, severity level, event type, and event text.

Failover Messages

Failover Syslog Messages

The ASA issues a number of syslog messages related to failover at priority level 2, which indicates a critical condition. To view these messages, see the syslog messages guide. Syslog messages are in the ranges of 1045xx and 1055xx.

Note During failover, the ASA logically shuts down and then brings up interfaces, generating syslog messages. This is normal activity.

The following are sample syslogs generated during a switchover:

```
%ASA-3-105509: (Primary) Error sending Hello message to peer unit 10.22.3.5, error: Unknown
 error
%ASA-1-104500: (Primary) Switching to ACTIVE - switch reason: Unable to send message to
Active unit
%ASA-5-105522: (Primary) Updating route-table wc-rt-inside
%ASA-5-105523: (Primary) Updated route-table wc-rt-inside
%ASA-5-105522: (Primary) Updating route-table wc-rt-outside
%ASA-5-105523: (Primary) Updated route-table wc-rt-outside
%ASA-5-105542: (Primary) Enabling load balancer probe responses
%ASA-5-105503: (Primary) Internal state changed from Backup to Active no peer
%ASA-5-105520: (Primary) Responding to Azure Load Balancer probes
```

Each syslog related to a Public Cloud deployment is prefaced with the unit role: (Primary) or (Secondary).

Failover Debug Messages

To see debug messages, enter the **debug fover** command. See the command reference for more information.



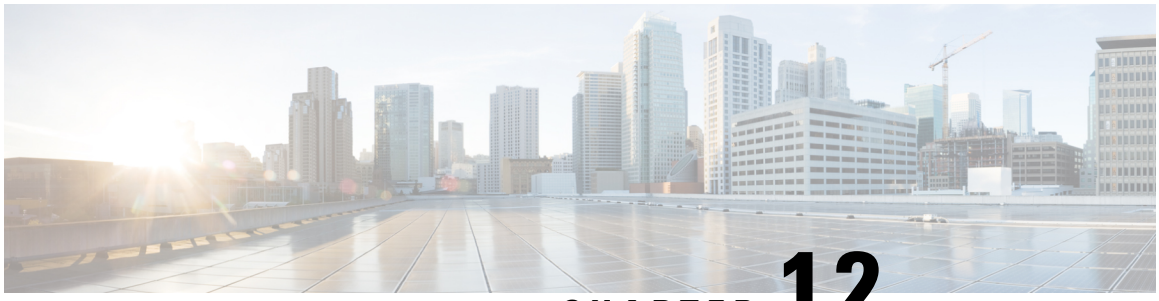
Note Because debugging output is assigned high priority in the CPU process, it can drastically affect system performance. For this reason, use the **debug fover** commands only to troubleshoot specific problems or during troubleshooting sessions with Cisco TAC.

SNMP Failover Traps

To receive SNMP syslog traps for failover, configure the SNMP agent to send SNMP traps to SNMP management stations, define a syslog host, and compile the Cisco syslog MIB into your SNMP management station.

History for Failover in the Public Cloud

Feature Name	Releases	Feature Information
Active/Backup failover on Microsoft Azure	7.9(1)	This feature was introduced.



CHAPTER 12

ASA Cluster for the Secure Firewall 3100/4200

Clustering lets you group multiple ASAs 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 411.

- About ASA Clustering, on page 341
- Licenses for ASA Clustering, on page 345
- Requirements and Prerequisites for ASA Clustering, on page 346
- Guidelines for Clustering, on page 348
- Configure ASA Clustering, on page 354
- Manage Cluster Nodes, on page 387
- Monitoring the ASA Cluster, on page 394
- Troubleshooting Distributed Site-to-Site VPN, on page 396
- Examples for ASA Clustering, on page 397
- Reference for Clustering, on page 411
- History for ASA Clustering for the Secure Firewall 3100/4200, on page 427

About ASA Clustering

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.

Cluster Members

Cluster members work together to accomplish the sharing of the security policy and traffic flows. This section describes the nature of each member role.

Bootstrap Configuration

On each device, you configure a minimal bootstrap configuration including the cluster name, cluster control link interface, and other cluster settings. The first node on which you enable clustering typically becomes the *control* node. When you enable clustering on subsequent nodes, they join the cluster as *data* nodes.

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 in the bootstrap configuration; the priority is set between 1 and 100, where 1 is the highest priority. All other members are data nodes. Typically, when you first create a cluster, the first node you add becomes the control node simply because it is the only node in the cluster so far.

You must perform all configuration (aside from the bootstrap configuration) on the control node only; the configuration is then replicated to the data nodes. In the case of physical assets, such as interfaces, the configuration of the control node is mirrored on all data nodes. For example, if you configure Ethernet 1/2 as the inside interface and Ethernet 1/1 as the outside interface, then these interfaces are also used on the data nodes as inside and outside interfaces.

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 354 for more information.

Cluster Control Link

Each unit must dedicate at least one hardware interface as the cluster control link. See Cluster Control Link, on page 354 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.

ASA Cluster Management

One of the benefits of using ASA clustering is the ease of management. This section describes how to manage the cluster.

Management Network

We recommend connecting all units to a single management network. This network is separate from the cluster control link.

Management Interface

For the management interface, we recommend using one of the dedicated management interfaces. You can configure the management interfaces as Individual interfaces (for both routed and transparent modes) or as a Spanned EtherChannel interface.

We recommend using Individual interfaces for management. Individual interfaces let you connect directly to each unit if necessary, while a Spanned EtherChannel interface only allows remote connection to the current control unit.



Note If you use Spanned EtherChannel interface mode and configure the management interface as an Individual interface, you cannot enable dynamic routing for the management interface. You must use a static route.

For an Individual interface, the Main cluster IP address is a fixed address for the cluster that always belongs to the current control unit. For each interface, you also configure a range of addresses so that each unit, including the current control unit, can use a Local address from the range. The Main cluster IP address provides consistent management access to an address; when a control unit changes, the Main cluster IP address moves to the new control unit, so management of the cluster continues seamlessly. The Local IP address is used for routing, and is also useful for troubleshooting.

For example, you can manage the cluster by connecting to the Main cluster IP address, which is always attached to the current control unit. To manage an individual member, you can connect to the Local IP address.



Note To-the-box traffic needs to be directed to the node's management IP address; to-the-box traffic is not forwarded over the cluster control link to any other node.

For outbound management traffic such as TFTP or syslog, each unit, including the control unit, uses the Local IP address to connect to the server.

For a Spanned EtherChannel interface, you can only configure one IP address, and that IP address is always attached to the control unit. You cannot connect directly to a data unit using the EtherChannel interface; we recommend configuring the management interface as an Individual interface so that you can connect to each unit. Note that you can use a device-local EtherChannel for management.

Control Unit Management Vs. Data Unit Management

All management and monitoring can take place on the control node. From the control node, you can check runtime statistics, resource usage, or other monitoring information of all nodes. You can also issue a command to all nodes in the cluster, and replicate the console messages from data nodes to the control node.

You can monitor data nodes directly if desired. Although also available from the control node, you can perform file management on data nodes (including backing up the configuration and updating images). The following functions are not available from the control node:

- Monitoring per-node cluster-specific statistics.
- Syslog monitoring per node (except for syslogs sent to the console when console replication is enabled).
- SNMP
- NetFlow

Crypto Key Replication

When you create a crypto key on the control node, the key is replicated to all data nodes. If you have an SSH session to the Main cluster IP address, you will be disconnected if the control node fails. The new control node uses the same key for SSH connections, so that you do not need to update the cached SSH host key when you reconnect to the new control node.

ASDM Connection Certificate IP Address Mismatch

By default, a self-signed certificate is used for the ASDM connection based on the Local IP address. If you connect to the Main cluster IP address using ASDM, then a warning message about a mismatched IP address might appear because the certificate uses the Local IP address, and not the Main cluster IP address. You can ignore the message and establish the ASDM connection. However, to avoid this type of warning, you can enroll a certificate that contains the Main cluster IP address and all the Local IP addresses from the IP address pool. You can then use this certificate for each cluster member. See https://www.cisco.com/c/en/us/td/docs/security/asdm/identity-cert/cert-install.html for more information.

Inter-Site Clustering

For inter-site installations, you can take advantage of ASA clustering as long as you follow the recommended guidelines.

You can configure each cluster chassis to belong to a separate site ID.

Site IDs work with site-specific MAC addresses and IP addresses. Packets egressing the cluster use a site-specific MAC address and IP address, while packets received by the cluster use a global MAC address and IP address. This feature prevents the switches from learning the same global MAC address from both sites on two different ports, which causes MAC flapping; instead, they only learn the site MAC address. Site-specific MAC addresses and IP address are supported for routed mode using Spanned EtherChannels only.

Site IDs are also used to enable flow mobility using LISP inspection, director localization to improve performance and reduce round-trip time latency for inter-site clustering for data centers, and site redundancy for connections where a backup owner of a traffic flow is always at a different site from the owner.

See the following sections for more information about inter-site clustering:

- Sizing the Data Center Interconnect—Requirements and Prerequisites for ASA Clustering, on page 346
- Inter-Site Guidelines—Guidelines for Clustering, on page 348
- Configure Cluster Flow Mobility—Configure Cluster Flow Mobility, on page 379
- Enable Director Localization—Configure Basic ASA Cluster Parameters, on page 372
- Enable Site Redundancy—Configure Basic ASA Cluster Parameters, on page 372
- Inter-Site Examples—Examples for Inter-Site Clustering, on page 407

Licenses for ASA Clustering

Smart Software Manager Regular and On-Prem

Each unit requires the Essentials license (enabled by default) and the same encryption license. We recommend licensing each unit with the licensing server *before* you enable clustering to avoid licensing mismatch issues, and when using the Strong Encryption license, issues with cluster control link encryption.

The clustering feature itself does not require any licenses. There is no extra cost for the Context license on data units.

The Strong Encryption license is automatically enabled for qualified customers when you apply the registration token. For the optional Strong Encryption (3DES/AES) feature license enabled in the ASA configuration, see below.

In the ASA license configuration, the Essentials license is always enabled by default on all units. You can only configure smart licensing on the control unit. The configuration is replicated to the data units, but for some licenses, they do not use the configuration; it remains in a cached state, and only the control unit requests the license. The licenses are aggregated into a single cluster license that is shared by the cluster units, and this aggregated license is also cached on the data units to be used if one of them becomes the control unit in the future. Each license type is managed as follows:

- Essentials—Each unit requests a Essentials license from the server.
- Context—Only the control unit requests the Context license from the server. The Essentials license includes 2 contexts by default and is present on all cluster members. The value from each unit's Essentials license plus the value of the Context license on the control unit are combined up to the platform limit in an aggregated cluster license. For example:
  - You have 6 Secure Firewall 3100s in the cluster. The Essentials license includes 2 contexts; for 6 units, these licenses add up to 12 contexts. You configure an additional 20-Context license on the control unit. Therefore, the aggregated cluster license includes 32 contexts. Because the platform limit for one chassis is 100, the combined license allows a maximum of 100 contexts; the 32 contexts are within the limit. Therefore, you can configure up to 32 contexts on the control unit; each data unit will also have 32 contexts through configuration replication.
  - You have 3 Secure Firewall 3100s in the cluster. The Essentials license includes 2 contexts; for 3 units, these licenses add up to 6 contexts. You configure an additional 100-Context license on the control unit. Therefore, the aggregated cluster license includes 106 contexts. Because the platform limit for one unit is 100, the combined license allows a maximum of 100 contexts; the 106 contexts are over the limit. Therefore, you can only configure up to 100 contexts on the control unit; each data unit will also have 100 contexts through configuration replication. In this case, you should only configure the control unit Context license to be 94 contexts.

- Strong Encryption (3DES)—If your Smart Account is not authorized for strong encryption, but Cisco has determined that you are allowed to use strong encryption, you can manually add a strong encryption license to your account. Only the control unit requests this license, and all units can use it due to license aggregation.

If a new control unit is elected, the new control unit continues to use the aggregated license. It also uses the cached license configuration to re-request the control unit license. When the old control unit rejoins the cluster as a data unit, it releases the control unit license entitlement. Before the data unit releases the license, the control unit's license might be in a non-compliant state if there are no available licenses in the account. The retained license is valid for 30 days, but if it is still non-compliant after the grace period, you will not be able to make configuration changes to features requiring special licenses; operation is otherwise unaffected. The new active unit sends an entitlement authorization renewal request every 35 seconds until the license is compliant. You should refrain from making configuration changes until the license requests are completely processed. If a unit leaves the cluster, the cached control configuration is removed, while the per-unit entitlements are retained. In particular, you would need to re-request the Context license on non-cluster units.

Permanent License Reservation

For permanent license reservation, you must purchase separate licenses for each chassis and enable the licenses *before* you configure clustering.

Requirements and Prerequisites for ASA Clustering

Model Requirements

- Secure Firewall 3100—Maximum 16 units
- Secure Firewall 4200—Maximum 16 units

ASA Hardware and Software Requirements

All units in a cluster:

- Must be the same model with the same DRAM. You do not have to have the same amount of flash memory.
- Must run the identical software except at the time of an image upgrade. Hitless upgrade is supported. Mismatched software versions can lead to poor performance, so be sure to upgrade all nodes in the same maintenance window.
- Must be in the same security context mode, single or multiple.
- (Single context mode) Must be in the same firewall mode, routed or transparent.
- New cluster members must use the same SSL encryption setting (the **ssl encryption** command) as the control unit for initial cluster control link communication before configuration replication.

Switch Requirements

- Be sure to complete the switch configuration before you configure clustering on the ASAs.
- For a list of supported switches, see Cisco ASA Compatibility.

ASA Requirements

- Provide each unit with a unique IP address before you join them to the management network.
  - See the Getting Started chapter for more information about connecting to the ASA and setting the management IP address.
  - Except for the IP address used by the control unit (typically the first unit you add to the cluster), these management IP addresses are for temporary use only.
  - After a data unit joins the cluster, its management interface configuration is replaced by the one replicated from the control unit.

Sizing the Data Center Interconnect for Inter-Site Clustering

You should reserve bandwidth on the data center interconnect (DCI) for cluster control link traffic equivalent to the following calculation:

$$\frac{\text{\# of cluster members per site}}{2} \times \text{cluster control link size per member}$$

If the number of members differs at each site, use the larger number for your calculation. The minimum bandwidth for the DCI should not be less than the size of the cluster control link for one member.

For example:

- For 4 members at 2 sites:
  - 4 cluster members total
  - 2 members at each site
  - 5 Gbps cluster control link per member

  Reserved DCI bandwidth = 5 Gbps (2/2 x 5 Gbps).
- For 6 members at 3 sites, the size increases:
  - 6 cluster members total
  - 3 members at site 1, 2 members at site 2, and 1 member at site 3
  - 10 Gbps cluster control link per member

  Reserved DCI bandwidth = 15 Gbps (3/2 x 10 Gbps).
- For 2 members at 2 sites:
  - 2 cluster members total
  - 1 member at each site
  - 10 Gbps cluster control link per member

  Reserved DCI bandwidth = 10 Gbps (1/2 x 10 Gbps = 5 Gbps; but the minimum bandwidth should not be less than the size of the cluster control link (10 Gbps)).

Other Requirements

We recommend using a terminal server to access all cluster member unit console ports. For initial setup, and ongoing management (for example, when a unit goes down), a terminal server is useful for remote management.

Guidelines for Clustering

Context Mode

The mode must match on each member unit.

Firewall Mode

For single mode, the firewall mode must match on all units.

Failover

Failover 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.
- 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. *Do not* change the load-balancing algorithm from the default on the cluster device.
- 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:

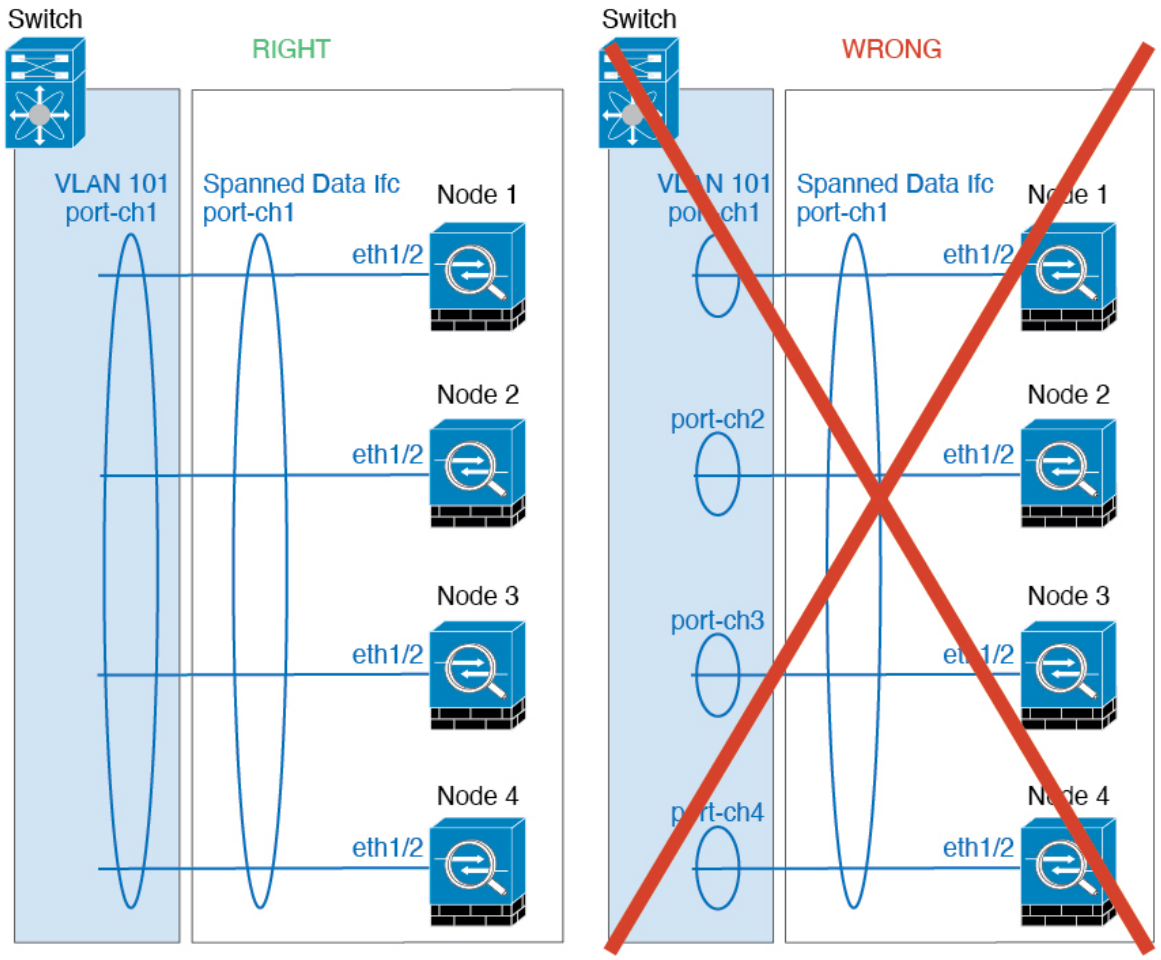
  router(config)# **port-channel** *id* **hash-distribution fixed**

  Do not change the algorithm globally; you may want to take advantage of the adaptive algorithm for the VSS peer link.

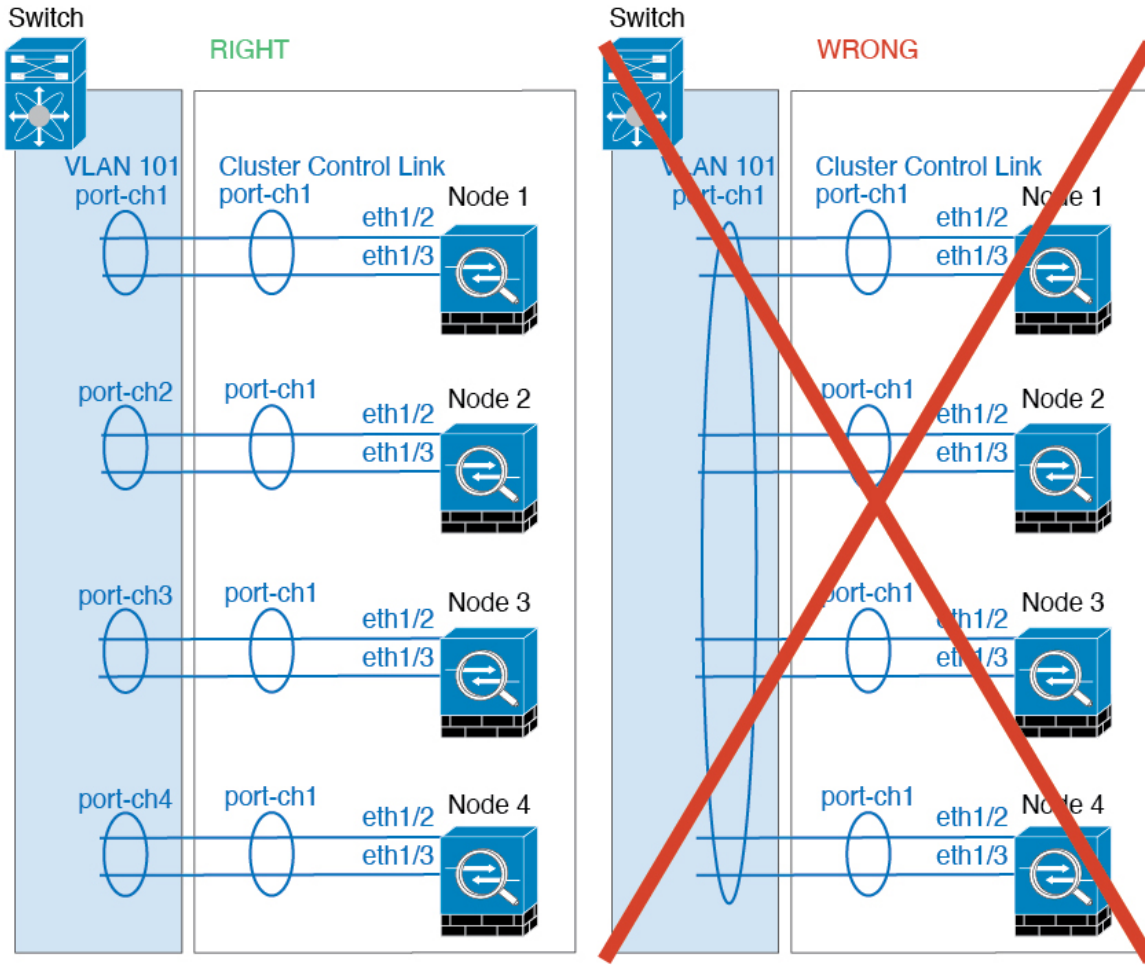
- You should disable the LACP Graceful Convergence feature on all cluster-facing EtherChannel interfaces for Cisco Nexus switches.

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.



Inter-Site Guidelines

See the following guidelines for inter-site clustering:

- Supports inter-site clustering in the following interface and firewall modes:

Interface Mode	Firewall Mode	
	Routed	**Transparent**
Individual Interface	**Yes**	N/A
Spanned EtherChannel	**Yes**	**Yes**

- For individual interface mode, when using ECMP towards a multicast Rendezvous Point (RP), we recommend that you use a static route for the RP IP address using the Main cluster IP address as the next hop. This static route prevents sending unicast PIM register packets to data units. If a data unit receives a PIM register packet, then the packet is dropped, and the multicast stream cannot be registered.

- The cluster control link latency must be less than 20 ms round-trip time (RTT).

- The cluster control link must be reliable, with no out-of-order or dropped packets; for example, you should use a dedicated link.
- Do not configure connection rebalancing; you do not want connections rebalanced to cluster members at a different site.
- The ASA does not encrypt forwarded data traffic on the cluster control link because it is a dedicated link, even when used on a Data Center Interconnect (DCI). If you use Overlay Transport Virtualization (OTV), or are otherwise extending the cluster control link outside of the local administrative domain, you can configure encryption on your border routers such as 802.1AE MacSec over OTV.
- The cluster implementation does not differentiate between members at multiple sites for incoming connections; therefore, connection roles for a given connection may span across sites. This is expected behavior. However, if you enable director localization, the local director role is always chosen from the same site as the connection owner (according to site ID). Also, the local director chooses a new owner at the same site if the original owner fails (Note: if the traffic is asymmetric across sites, and there is continuous traffic from the remote site after the original owner fails, then a node from the remote site might become the new owner if it receives a data packet within the re-hosting window.).
- For director localization, the following traffic types do not support localization: NAT or PAT traffic; SCTP-inspected traffic; Fragmentation owner query.
- For UDP long-lived flows in a North-South deployment, routing loops can occur if nodes at the original flow owner site fail and then come back up, after which the flow is directed back to the original site. If the new owner at the other site doesn't have a route to the destination, it will route the flow back to the internet, causing a loop. In this case, use the **clear conn** command on the new owner to force the flow to be reestablished.
- For transparent mode, if the cluster is placed between a pair of inside and outside routers (AKA North-South insertion), you must ensure that both inside routers share a MAC address, and also that both outside routers share a MAC address. When a cluster member at site 1 forwards a connection to a member at site 2, the destination MAC address is preserved. The packet will only reach the router at site 2 if the MAC address is the same as the router at site 1.
- For transparent mode, if the cluster is placed between data networks and the gateway router at each site for firewalling between internal networks (AKA East-West insertion), then each gateway router should use a First Hop Redundancy Protocol (FHRP) such as HSRP to provide identical virtual IP and MAC address destinations at each site. The data VLANs are extended across the sites using Overlay Transport Virtualization (OTV), or something similar. You need to create filters to prevent traffic that is destined to the local gateway router from being sent over the DCI to the other site. If the gateway router becomes unreachable at one site, you need to remove any filters so traffic can successfully reach the other site's gateway.
- For transparent mode, if the cluster is connected to an HSRP router, you must add the router HSRP MAC address as a static MAC address table entry on the ASA (see Add a Static MAC Address for Bridge Groups, on page 767). When adjacent routers use HSRP, traffic destined to the HSRP IP address will be sent to the HSRP MAC Address, but return traffic will be sourced from the MAC address of a particular router's interface in the HSRP pair. Therefore, the ASA MAC address table is typically only updated when the ASA ARP table entry for the HSRP IP address expires, and the ASA sends an ARP request and receives a reply. Because the ASA's ARP table entries expire after 14400 seconds by default, but the MAC address table entry expires after 300 seconds by default, a static MAC address entry is required to avoid MAC address table expiration traffic drops.
- For routed mode using Spanned EtherChannel, configure site-specific MAC addresses. Extend the data VLANs across the sites using OTV, or something similar. You need to create filters to prevent traffic

that is destined to the global MAC address from being sent over the DCI to the other site. If the cluster becomes unreachable at one site, you need to remove any filters so traffic can successfully reach the other site's cluster nodes. Dynamic routing is not supported when an inter-site cluster acts as the first hop router for an extended segment.

Additional Guidelines

- When significant topology changes occur (such as adding or removing an EtherChannel interface, enabling or disabling an interface on the firewall 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 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 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 do not support VXLAN in Individual Interface mode. Only Spanned EtherChannel mode supports VXLAN.
- We do not support IS-IS in Spanned EtherChannel mode. Only Individual Interface mode supports IS-IS.
- It takes time to replicate changes to all the units in a cluster. If you make a large change, for example, adding an access control rule that uses object groups (which, when deployed, are broken out into multiple rules), the time needed to complete the change can exceed the timeout for the cluster units to respond with a success message. If this happens, you might see a "failed to replicate command" message. You can ignore the message.

Defaults for Clustering

- When using Spanned EtherChannels, 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 rebalancing is disabled by default. If you enable connection rebalancing, the default time between load information exchanges is 5 seconds.
- Connection replication delay of 5 seconds is enabled by default for HTTP traffic.

Configure ASA Clustering

To configure clustering, perform the following tasks.

Note To enable or disable clustering, you must use a console connection (for CLI) or an ASDM connection.

Back Up Your Configurations (Recommended)

When you enable clustering on a data unit, the current configuration is replaced with one synced from the active unit. If you ever want to leave the cluster entirely, it may be useful to have a backup configuration with a usable management interface configuration.

Before you begin

Perform a backup on each unit.

Procedure

Step 1 Choose **Tools** > **Backup Configurations**.

Step 2 Back up at least the running configuration. See Back Up and Restore Configurations or Other Files, on page 1086 for a detailed procedure.

Cable the Units and Configure Interfaces

Before configuring clustering, cable the cluster control link network, management network, and data networks. Then configure your interfaces.

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.

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 data interface(s) for the cluster control link, with the following exceptions:

- You cannot use a VLAN subinterface as the cluster control link.
- You cannot use a Management *x*/*x* interface as the cluster control link, either alone or as an EtherChannel.

You can use an EtherChannel.

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 ASA cluster control link interfaces.

For a 2-member cluster, do not directly-connect the cluster control link from one ASA to the other ASA. 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.

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.
- AAA for network access is a centralized feature, so all traffic is forwarded to the control unit.
- 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.

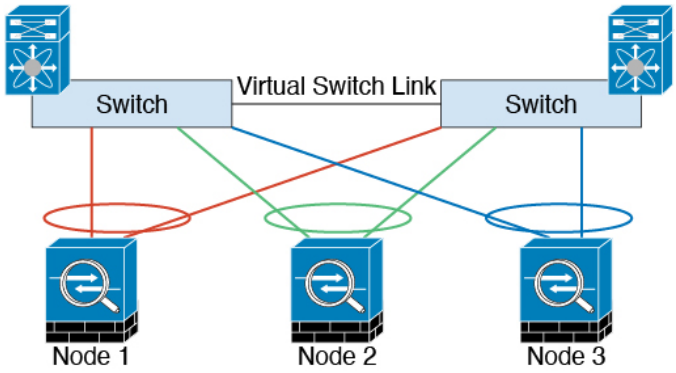


Note If your cluster has large amounts of asymmetric (rebalanced) traffic, then you should increase the cluster control link size.

Cluster Control Link Redundancy

We recommend using an EtherChannel for the cluster control link, so that you can pass traffic on multiple links in the EtherChannel while still achieving 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.

Cluster Control Link Failure

If the cluster control link line protocol goes down for a unit, then clustering is disabled; data interfaces are shut down. After you fix the cluster control link, you must manually rejoin the cluster by re-enabling clustering.

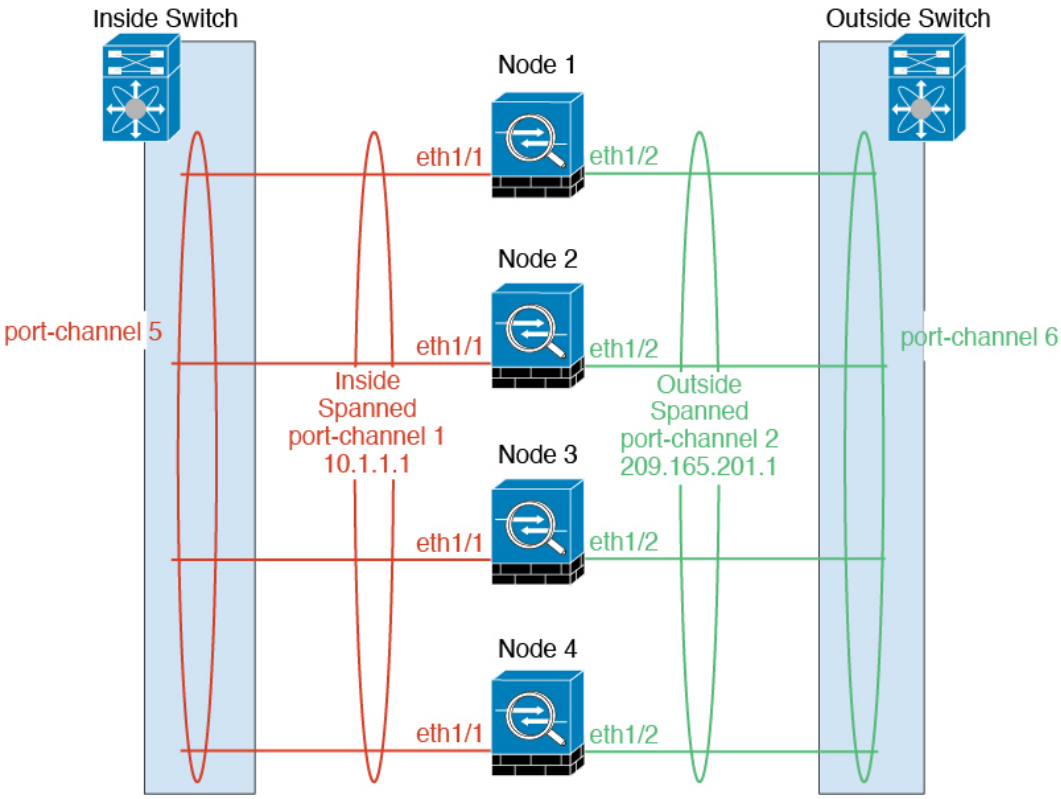
Note When the ASA becomes inactive, all data interfaces are shut down; only the management-only interface can send and receive traffic. The management interface remains up using the IP address the unit received from the cluster IP pool. However if you reload, and the unit is still inactive in the cluster, the management interface is not accessible (because it then uses the Main IP address, which is the same as the control unit). You must use the console port for any further configuration.

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.

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 ASA 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 ASAs 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.

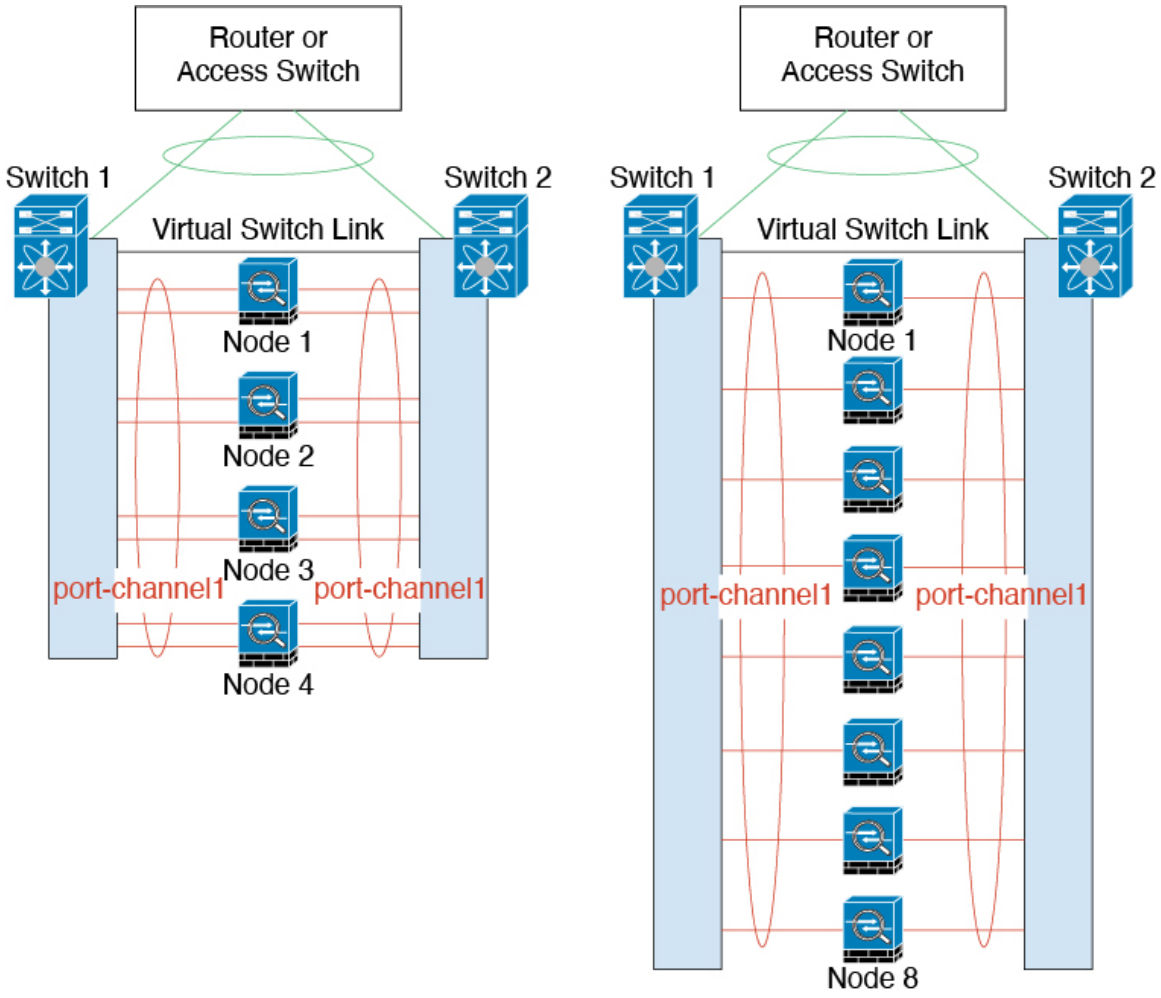
Connecting to a Redundant Switch System

You can include multiple interfaces per ASA in the Spanned EtherChannel. Multiple interfaces per ASA 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.

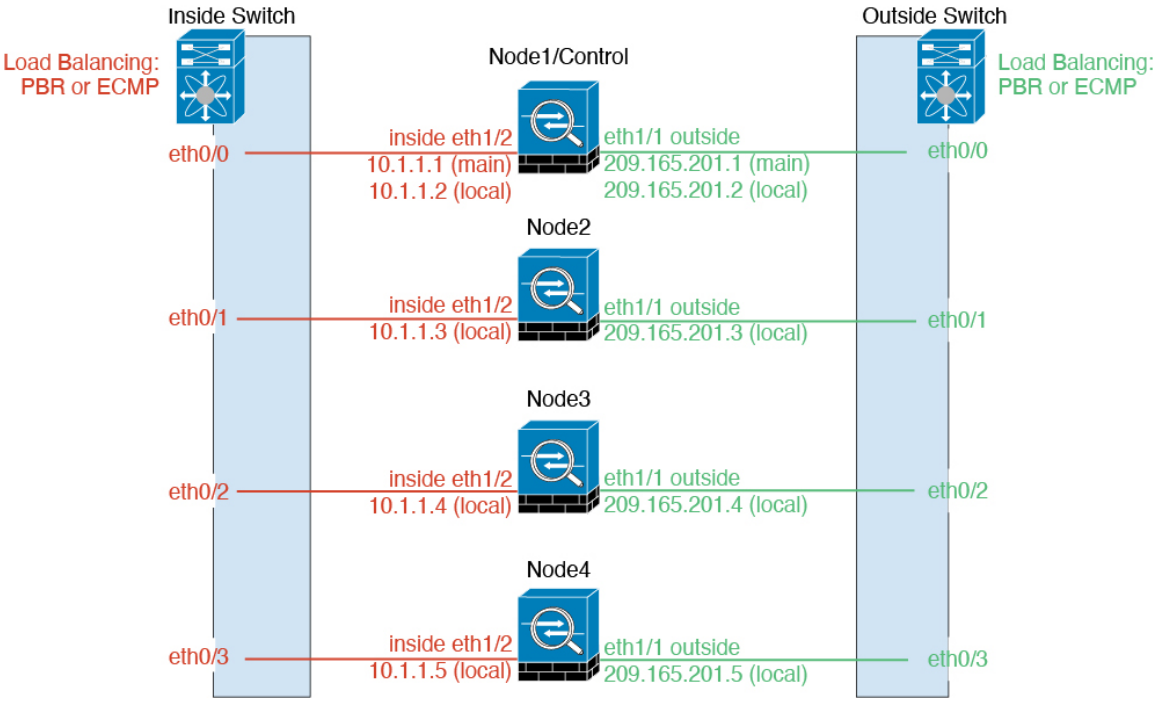


Individual Interfaces (Routed Firewall Mode Only)

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.

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 ASA 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 ASAs 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 ASA. 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 ASA using ICMP ping. PBR can then enable or disable route maps based on reachability of a particular ASA. 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 ASA 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 ASA failure can cause problems; the route continues to be used, and traffic to the failed ASA 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 ASA to participate in dynamic routing.

Cisco Intelligent Traffic Director (Routed Firewall Mode Only)

When using Individual interfaces, each ASA 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 ASA. 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 the Cluster Units and Configure Upstream and Downstream Equipment

Before configuring clustering, cable the cluster control link network, management network, and data networks.

Procedure

Step 1 Cable the cluster control link network, management network, and data networks.

Note
At a minimum, an active cluster control link network is required before you configure the nodes to join the cluster.

Step 2 You should also configure the upstream and downstream equipment. For example, if you use EtherChannels, then you should configure the upstream and downstream equipment for the EtherChannels.

Configure the Cluster Interface Mode on the Control Unit

You can only configure one type of interface for clustering: Spanned EtherChannels or Individual interfaces; you cannot mix interface types in a cluster.



Note If you do not add data units from the control unit, you must set the interface mode manually on all units according to this section, not just the control unit; if you add secondaries from the control unit, ASDM sets the interface mode automatically on the data unit.

Before you begin

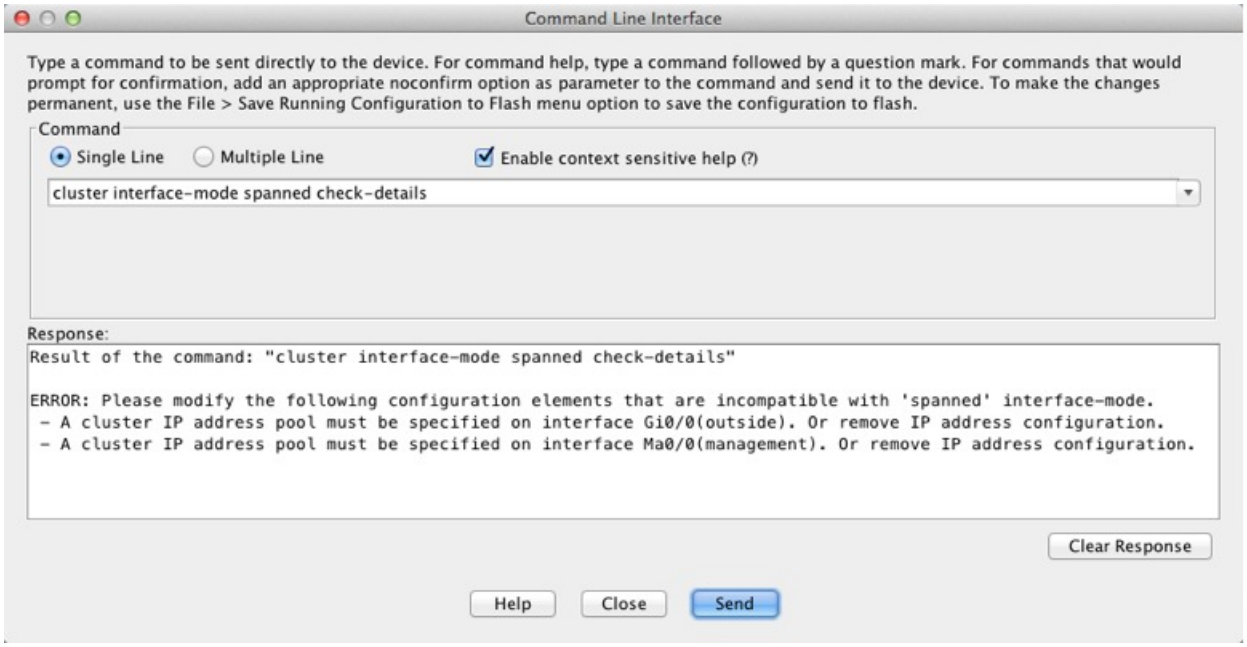
- You can always configure the management-only interface as an Individual interface (recommended), even in Spanned EtherChannel mode. The management interface can be an Individual interface even in transparent firewall mode.
- In Spanned EtherChannel mode, if you configure the management interface as an Individual interface, you cannot enable dynamic routing for the management interface. You must use a static route.
- In multiple context mode, you must choose one interface type for all contexts. For example, if you have a mix of transparent and routed mode contexts, you must use Spanned EtherChannel mode for all contexts because that is the only interface type allowed for transparent mode.

Procedure

Step 1 In ASDM on the control unit, choose **Tools** > **Command Line Interface**.Show any incompatible configuration so that you can force the interface mode and fix your configuration later; the mode is not changed with this command:

cluster interface-mode {**individual** | **spanned**} **check-details**

Example:

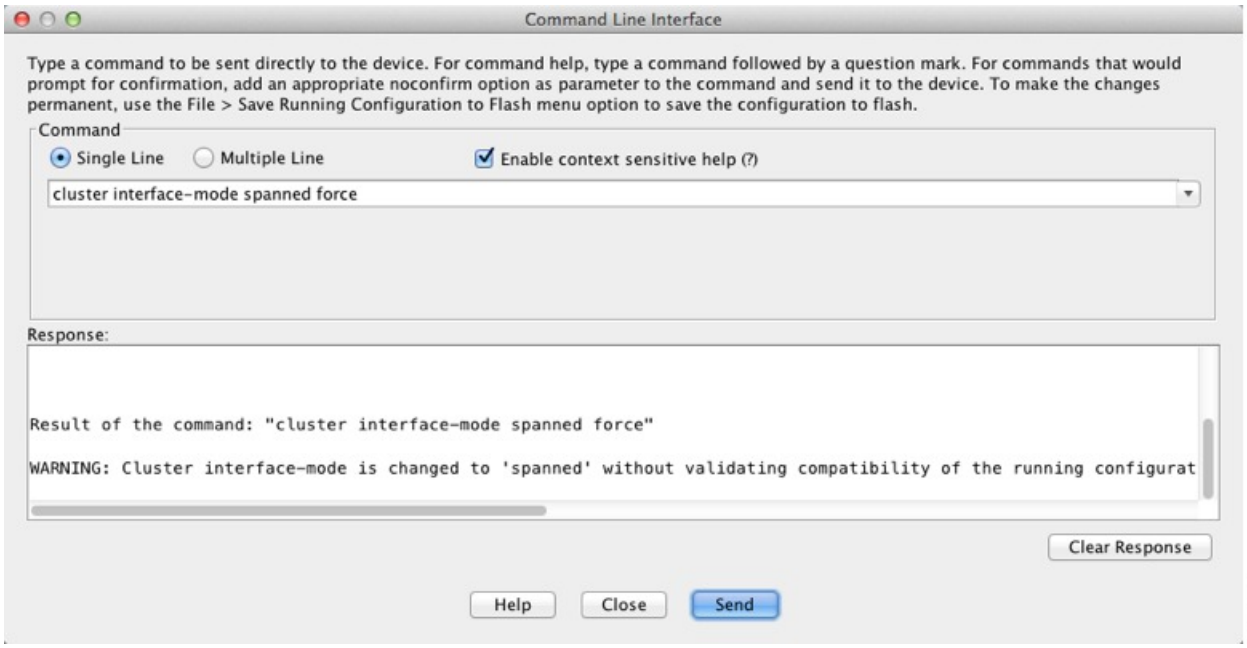


Caution
After you set the interface mode, you can continue to connect to the interface; however, if you reload the ASA before you configure your management interface to comply with clustering requirements (for example, adding a cluster IP pool), you will not be able to reconnect because cluster-incompatible interface configuration is removed. In that case, you will have to connect to the console port to fix the interface configuration.

Step 2 Set the interface mode for clustering:

cluster interface-mode {**individual** | **spanned**} **force**

Example:



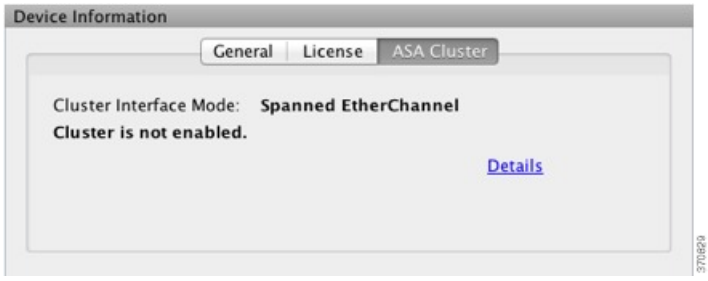
There is no default setting; you must explicitly choose the mode. If you have not set the mode, you cannot enable clustering.

The **force** option changes the mode without checking your configuration for incompatible settings. You need to manually fix any configuration issues after you change the mode. Because any interface configuration can only be fixed after you set the mode, we recommend using the **force** option so that you can at least start from the existing configuration. You can re-run the **check-details** option after you set the mode for more guidance.

Without the **force** option, if there is any incompatible configuration, you are prompted to clear your configuration and reload, thus requiring you to connect to the console port to reconfigure your management access. If your configuration is compatible (rare), the mode is changed and the configuration is preserved. If you do not want to clear your configuration, you can exit the command by typing **n**.

To remove the interface mode, enter the **no cluster interface-mode** command.

Step 3 Quit ASDM and reload. ASDM needs to be restarted to correctly account for the cluster interface mode. After you reload, you see the ASA Cluster tab on the home page:



(Recommended; Required in Multiple Context Mode) Configure Interfaces on the Control Node

You must modify any interface that is currently configured with an IP address to be cluster-ready before you enable clustering. At a minimum, you must modify the management interface to which ASDM is currently connected. For other interfaces, you can configure them before or after you enable clustering; we recommend pre-configuring all of your interfaces so that the complete configuration is synced to new cluster members. In multiple context mode, you must use the procedures in this section to fix existing interfaces or to configure new interfaces. However, in single mode, you can skip this section and configure common interface parameters within the High Availability and Scalability wizard (see Create or Join a Cluster Using the High Availability Wizard, on page 368). Note that advanced interface settings such as creating EtherChannels for Individual interfaces are not available in the wizard.

This section describes how to configure interfaces to be compatible with clustering. 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, with the exception of the management interface, which can be an Individual interface even in Spanned EtherChannel mode.

Configure Individual Interfaces (Recommended for the Management Interface)

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.

In Spanned EtherChannel mode, we recommend configuring the management interface as an Individual interface. Individual management interfaces let you connect directly to each unit if necessary, while a Spanned EtherChannel interface only allows connection to the control node.

Before you begin

- Except for the management-only interface, you must be in Individual interface mode.
- For multiple context mode, perform this procedure in each context. If you are not already in the context configuration mode in the Configuration > Device List pane, double-click the context name under the active device IP address.
- 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.
- If you are connecting remotely to the management interface using ASDM, the current IP address of prospective secondary units are for temporary use.
  - Each member will be assigned an IP address from the cluster IP pool defined on the primary unit.
  - The cluster IP pool cannot include addresses already in use on the network, including prospective secondary IP addresses.

    For example:

    1. You configure the primary unit to use 10.1.1.1.
    2. Other units use 10.1.1.2, 10.1.1.3, and 10.1.1.4.

3. When you configure the cluster IP pool on the primary unit, you cannot include the .2, .3, or .4 addresses in the pool, because they are in use.
4. Instead, you need to use other IP addresses on the network, such as .5, .6, .7, and .8.

Note The pool needs as many addresses as there are members of the cluster, including the primary unit; the original .1 address is the main cluster IP address that belongs to the current primary unit.

5. After you join the cluster, the old, temporary addresses are relinquished and can be used elsewhere.

Procedure

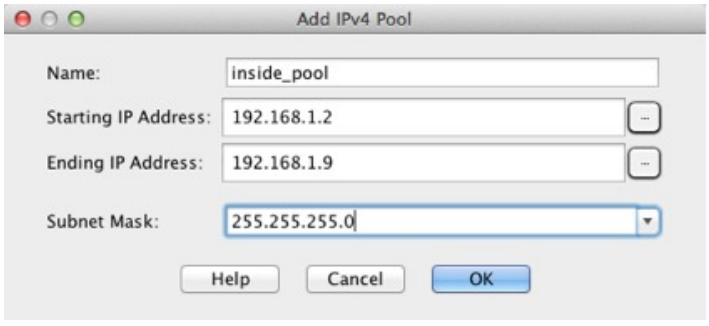
Step 1 Choose the **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces** pane.

Step 2 Choose the interface row, and click **Edit**. Set the interface parameters. See the following guidelines:

- (Required for a management interface in Spanned EtherChannel mode) **Dedicate this interface to management only**—Sets the interface to management-only mode so that it does not pass through traffic. By default, Management type interfaces are configured as management-only. In transparent mode, this command is always enabled for a Management type interface.
- **Use Static IP**—DHCP and PPPoE are not supported.

Step 3 To add the IPv4 cluster IP pool, MAC address pool, and site-specific MAC addresses, click the **Advanced** tab and set **ASA Cluster** area parameters.

a) Create a cluster IP pool by clicking the **...** button next to the **IP Address Pool** field. The valid range shown is determined by the Main IP address you set on the General tab.
b) Click **Add**.
c) Configure a range of addresses that does not include the Main cluster IP address, and that does not include any addresses currently in-use on your network. You should make the range large enough for the size of the cluster, for example, 8 addresses.



d) Click **OK** to create the new pool.
e) Select the new pool you created, and click **Assign**, and then click **OK**.

The pool name appears in the **IP Address Pool** field.

f) (Optional) (Optional) Configure a **MAC Address Pool** if you want to manually configure MAC addresses.

Step 4 To configure an IPv6 address, click the **IPv6** tab.

a) Check the **Enable IPv6** check box.

b) In the **Interface IPv6 Addresses** area, click **Add**.

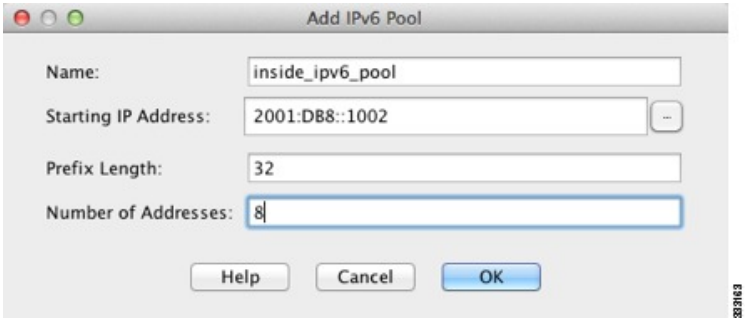
The **Enable address autoconfiguration** option is not supported. Manually configuring the link-local address is also not supported.

The **Add IPv6 Address for Interface** dialog box appears.

c) In the **Address/Prefix Length** field, enter the global IPv6 address and the IPv6 prefix length. For example, 2001:0DB8::BA98:0:3210/48.

d) Click the **...** button to configure the cluster IP pool.

e) Click **Add**.



f) Configure the starting IP address (network prefix), prefix length, and number of addresses in the pool.

g) Click **OK** to create the new pool.

h) Select the new pool you created, and click **Assign**, and then click **OK**.

The pool appears in the **ASA Cluster IP Pool** field.

i) Click **OK**.

Step 5 Click **OK** to return to the Interfaces pane.

Step 6 Click **Apply**.

Configure Spanned EtherChannels

A Spanned EtherChannel spans all ASAs in the cluster, and provides load balancing as part of the EtherChannel operation.

Before you begin

- You must be in Spanned EtherChannel interface mode.
- For multiple context mode, start this procedure in the system execution space. If you are not already in the System configuration mode in the Configuration > Device List pane, double-click **System** under the active device IP address.

- For transparent mode, configure the bridge group. See Configure the Bridge Virtual Interface (BVI), on page 641.
- When using Spanned EtherChannels, 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 Depending on your context mode:

- For single mode, choose the **Configuration > Device Setup > Interface Settings > Interfaces** pane.
- For multiple mode in the System execution space, choose the **Configuration > Context Management > Interfaces** pane.

Step 2 Choose **Add** > **EtherChannel Interface**.

The **Add EtherChannel Interface** dialog box appears.

Step 3 Enable the following:

- **Port Channel ID**
- **Enable Interface** (checked by default)
- **Members in Group**—In the **Members in Group** list, you need to add at least one interface. Multiple interfaces in the EtherChannel per unit are useful for connecting to switches in a VSS, vPC, StackWise, or StackWise Virtual.

  Make sure all interfaces are the same type and speed. The first interface you add determines the type and speed of the EtherChannel. Any non-matching interfaces you add will be put into a suspended state. ASDM does not prevent you from adding non-matching interfaces.

The rest of the fields on this screen are described later in this procedure.

Step 4 To configure the MAC address and optional parameters, click the **Advanced** tab.

- In the **MAC Address Cloning** area, set a manual global MAC address for the EtherChannel. Do not set the Standby MAC Address; it is ignored. You must configure a MAC address 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.

  In multiple context mode, if you share an interface between contexts, you should instead enable auto-generation of MAC addresses so you do not need to set the MAC address manually. Note that you must manually configure the MAC address using this command for *non-shared* interfaces.
- (Routed mode) In the **ASA Cluster** area, for inter-site clustering set **Site specific MAC Addresses** and IP addresses for a site by clicking **Add** and specifying a MAC address and IP address for the site ID (1 through 8). Repeat for up to 8 sites. The site-specific IP addresses must be on the same subnet as the global IP address. The site-specific MAC address and IP address used by a unit depends on the site ID you specify in each unit's bootstrap configuration.

Step 5 (Optional) Configure VLAN subinterfaces on this EtherChannel. The rest of this procedure applies to the subinterfaces.

Step 6 (Multiple context mode) Before you complete this procedure, you need to allocate interfaces to contexts.

a) Click **OK** to accept your changes.
b) Allocate interfaces.
c) Change to the context that you want to configure: in the **Device List** pane, double-click the context name under the active device IP address.
d) Choose the **Configuration > Device Setup > Interface Settings > Interfaces** pane, select the port-channel interface that you want to customize, and click **Edit**.

The **Edit Interface** dialog box appears.

Step 7 Click the **General** tab.

Step 8 (Transparent Mode) From the **Bridge Group** drop-down list, choose the bridge group to which you want to assign this interface.

Step 9 In the **Interface Name** field, enter a name up to 48 characters in length.

Step 10 In the **Security level** field, enter a level between 0 (lowest) and 100 (highest).

Step 11 (Routed Mode) For an IPv4 address, click the **Use Static IP** radio button and enter the IP address and mask. 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. For transparent mode, you configure the IP address for the bridge group interface, not the EtherChannel interface.

Step 12 (Routed Mode) To configure an IPv6 address, click the **IPv6** tab.

For transparent mode, you configure the IP address for the bridge group interface, not the EtherChannel interface.

a) Check the **Enable IPv6** check box.
b) In the **Interface IPv6 Addresses** area, click **Add**.

The **Add IPv6 Address for Interface** dialog box appears.

Note
The **Enable address autoconfiguration** option is not supported. Manually configuring the link-local address is also not supported.

c) In the **Address/Prefix Length** field, enter the global IPv6 address and the IPv6 prefix length. For example, 2001:DB8::BA98:0:3210/64.
d) (Optional) To use the Modified EUI-64 interface ID as the host address, check the **EUI-64** check box. In this case, just enter the prefix in the **Address/Prefix Length** field.
e) Click **OK**.

Step 13 Click **OK** to return to the **Interfaces** screen.

Step 14 Click **Apply**.

Create or Join a Cluster Using the High Availability Wizard

Each node in the cluster requires a bootstrap configuration to join the cluster. Run the High Availability and Scalability wizard on one node (that will become the control node) to create the cluster, and then add data nodes to it.

Note The wizard does not let you change the default cLACP system ID and priority, which can only be configured while clustering is disabled. So either use the wizard and then disable clustering afterwards, or see Configure Basic ASA Cluster Parameters, on page 372 to configure the cluster manually.

Before you begin

- For multiple context mode, complete this procedure in the system execution space. If you are not already in the System configuration mode, in the **Configuration > Device List** pane, double-click **System** under the active device IP address.
- The interfaces you intend to use for the cluster control link interface must be in an up state on the connected switch.
- When you add a node to a running cluster, you may see temporary, limited packet/connection drops; this is expected behavior.

Procedure

Step 1 Choose **Wizards** > **High Availability and Scalability Wizard**. See select wizard guidelines in the following steps.

Step 2 On the **Interfaces** screen, you cannot create new EtherChannels from this screen (except for the cluster control link).

Step 3 On the ASA Cluster Configuration screen, configure bootstrap settings including:

- **Member Priority**—Sets the priority of this node for control node elections, between 1 and 100, where 1 is the highest priority.
- **Site Index**—If you use inter-site clustering, set the site ID for this node so it uses a site-specific MAC address, between 1 and 8.
- (Optional) **Shared Key**—Sets an encryption 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 encryption key. This parameter does not affect datapath traffic, including connection state update and forwarded packets, which are always sent in the clear. You must configure this parameter if you also enable the password encryption service.
- (Optional) **Enable connection rebalancing for TCP traffic across all the ASAs in the cluster**—Enables connection rebalancing. This parameter is disabled by default. If enabled, ASAs in a cluster exchange load information periodically, and offload new connections from more loaded devices to less loaded devices. The frequency, between 1 and 360 seconds, specifies how often the load information is exchanged. This parameter is not part of the bootstrap configuration, and is replicated from the control node to the data nodes.

  Note
  Do not configure connection rebalancing for inter-site topologies; you do not want connections rebalanced to cluster members at a different site.

- (Optional) **Enable health monitoring of this device within the cluster**—Enables the cluster node health check feature. To determine node health, the ASA 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 holdtime period, the peer node is considered unresponsive or dead.

Note
When any topology changes occur (such as adding or removing a data interface, enabling or disabling an interface on the ASA or the switch, adding an additional switch to form a VSS or vPC) you must disable the health check 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 health check.

- **Time to Wait Before Device Considered Failed**—This value determines the amount of time between node keepalive status messages, between .3 and 45 seconds; The default is 3 seconds.
- (Optional) **Broadcast keepalive messages to all EtherChannel cluster control link ports for VSS/vPC support**—If you configure the cluster control link as an EtherChannel (recommended), and it is connected to a VSS or vPC pair, then you might need to enable this option. For some switches, when one node in the VSS/vPC is shutting down or booting up, EtherChannel member interfaces connected to that switch may appear to be Up to the ASA, but they are not passing traffic on the switch side. The ASA can be erroneously removed from the cluster if you set the ASA holdtime timeout to a low value (such as .8 seconds), and the ASA sends heartbeat messages on one of these EtherChannel interfaces. When you enable this option, the ASA floods the heartbeat messages on all EtherChannel interfaces in the cluster control link to ensure that at least one of the switches can receive them.

- (Optional) **Replicate console output**—Enables console replication from data nodes to the control node. This feature is disabled by default. The ASA may print out some messages directly to the console for certain critical events. If you enable console replication, data nodes send the console messages to the control node so that you only need to monitor one console port for the cluster. This parameter is not part of the bootstrap configuration, and is replicated from the control node to the data nodes.
- **Cluster Control Link**—Specifies the cluster control link interface.
  - **MTU**—Specifies the maximum transmission node for the cluster control link interface to be at least 100 bytes higher than the highest MTU of the data interfaces, between 1400 and 9198 bytes, but not between 2561 and 8362. Due to block pool handling, this MTU size is not optimal for system operation. The default MTU is 1500 bytes.We suggest setting the MTU to the maximum. 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. For example, because the maximum MTU is 9198 bytes, then the highest data interface MTU can be 9098, while the cluster control link can be set to 9198.

Step 4 On the **Interfaces for Health Monitoring** screen, you can exempt some interfaces from monitoring for failure. You might want to disable health monitoring of non-essential interfaces, for example, the management interface.

Note
When any topology changes occur (such as adding or removing a data interface, enabling or disabling an interface on the ASA or the switch, adding an additional switch to form a VSS or vPC) you must disable the health check 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 health check.

Step 5 On the **Interface Auto Rejoin settings** screen, customize the auto-rejoin settings in case of an interface or cluster control link failure. For each type, you can set the following:

- **Maximum Rejoin Attempts**—Define the number of attempts at rejoining the cluster by setting **Unlimited** or a value between 0 and 65535. **0** disables auto-rejoining. The default value is **Unlimited** for the cluster-interface and **3** for the data-interface.
- **Rejoin Interval**—Define the interval duration in minutes between rejoin attempts by setting the interval 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**—Define if the interval duration increases by setting the interval variation 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.

Step 6 Click **Finish**.

Step 7 The ASA scans the running configuration for incompatible commands for features that are not supported with clustering, including commands that may be present in the default configuration. Click **OK** to delete the incompatible commands. If you click **Cancel**, then clustering is not enabled.

After a period of time while ASDM enables clustering and reconnects to the ASA, the Information screen appears confirming that the ASA was added to the cluster.

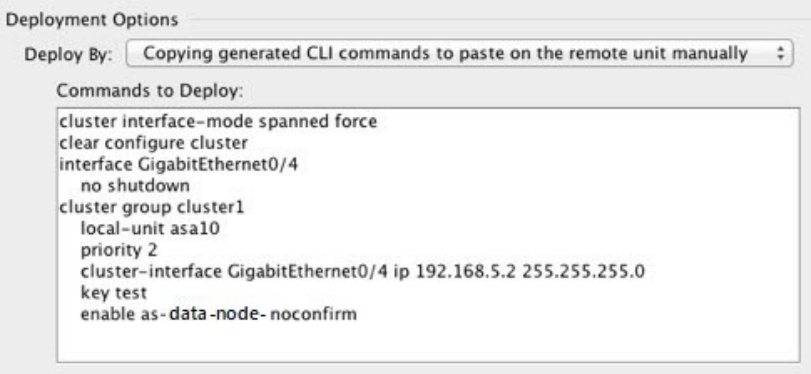
Note
In some cases, there might be an error when joining the cluster after you finish the wizard. If ASDM was disconnected, ASDM will not receive any subsequent errors from the ASA. If clustering remains disabled after you reconnect ASDM, you should connect to the ASA console port to determine the exact error condition that disabled clustering; for example, the cluster control link might be down.

Step 8 To add a data node, click **Yes**.

If you are re-running the wizard from the control node, you can add data nodes by choosing the **Add another member to the cluster** option when you first start the wizard.

Step 9 In the **Deployment Options** area, choose one of the following **Deploy By** options:

- **Sending CLI commands to the remote unit now—**Send the bootstrap configuration to the data node (temporary) management IP address. Enter the data node management IP address, username, and password.
- **Copying generated CLI commands to paste on the remote unit manually**—Generates the commands so that you can cut and paste them at the data node CLI or using the CLI tool in ASDM. In the Commands to Deploy box, select and copy the generated commands for later use.



Customize the Clustering Operation

You can customize clustering health monitoring, TCP connection replication delay, flow mobility and other optimizations.

Perform these procedures on the control node.

Configure Basic ASA Cluster Parameters

You can customize cluster settings on the control node. If you do not use the wizard to add a node to the cluster, you can configure the cluster parameters manually. If you already enabled clustering, you can edit some cluster parameters; others that cannot be edited while clustering is enabled are grayed out. This procedure also includes advanced parameters that are not included in the wizard.

For parameters that can only be changed while clustering is disabled, you need to disable clustering on all nodes and then re-enable clustering, which requires console CLI access.

Before you begin

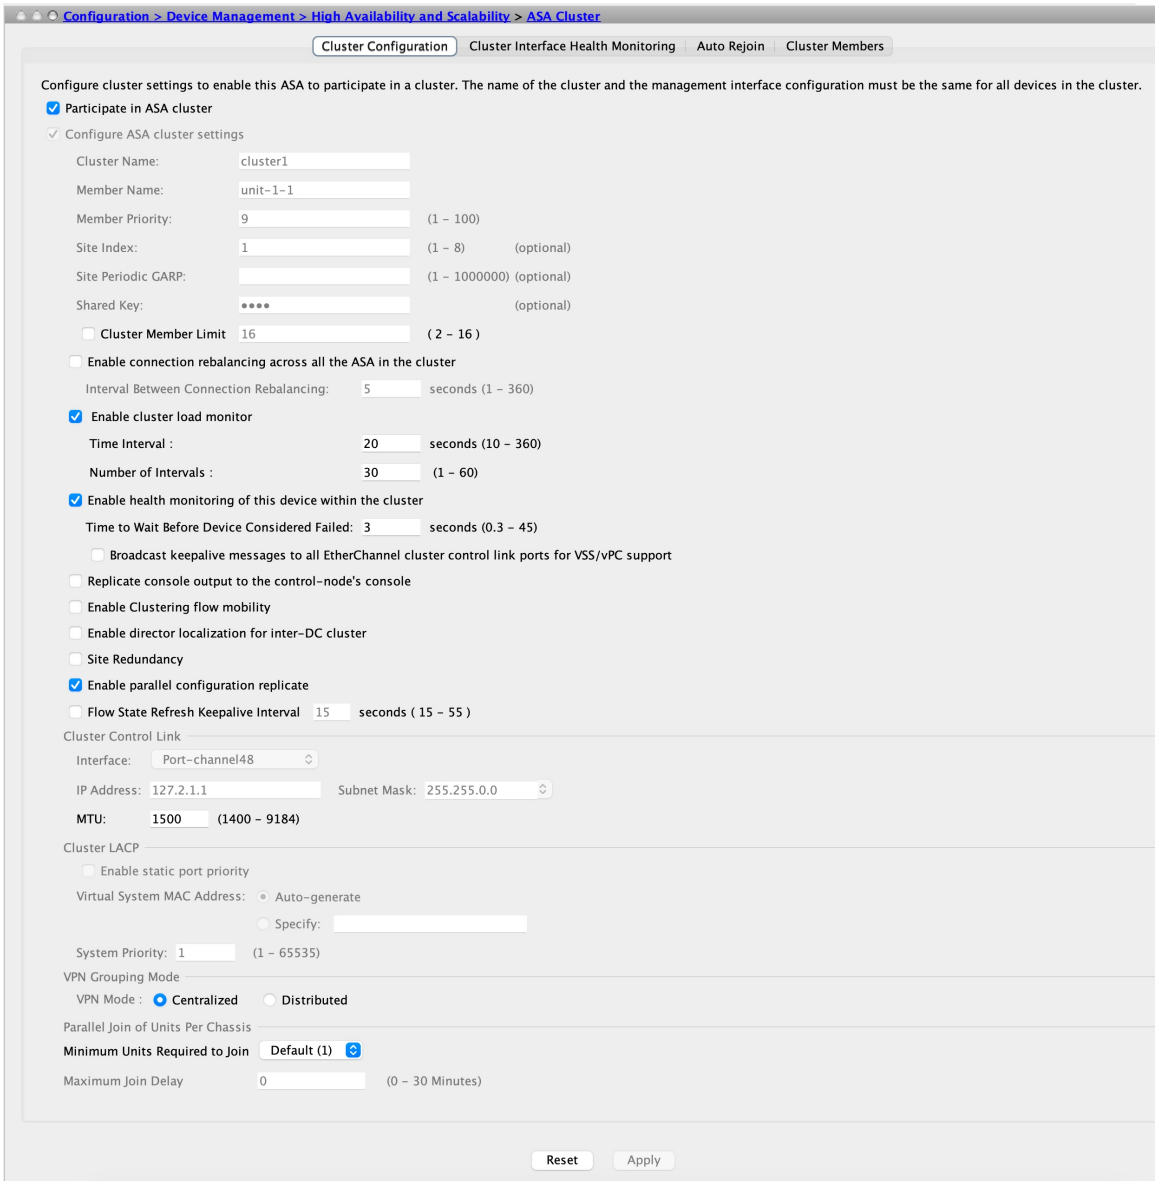
- If you did not use the wizard and want to manually join the cluster, you need to pre-configure the cluster-control-link interfaces on each node before joining the cluster. For a single interface, you must enable it; do not configure any other settings. For an EtherChannel interface, enable it and set the EtherChannel mode to On.
- For multiple context mode, complete this procedure in the system execution space on the control node. If you are not already in the System configuration mode, in the **Configuration** > **Device List** pane, double-click **System** under the active device IP address.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster**.

If your device is already in the cluster and is the control node, then this pane is on the **Cluster Configuration** tab.

Figure 60: Cluster Configuration

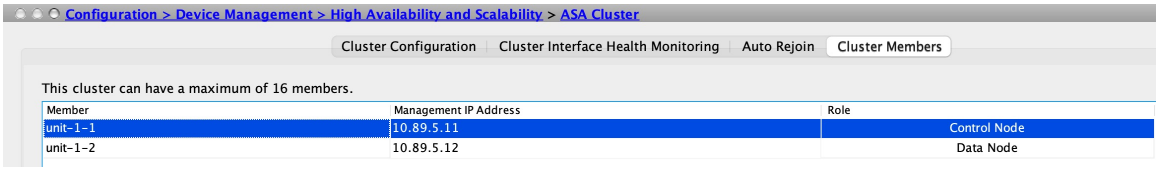


Step 2 If clustering was already enabled, and you want to configure grayed-out settings, then disable clustering on all nodes.

When an ASA becomes inactive, all data interfaces are shut down; only the management-only interface can send and receive traffic. The management interface remains up using the IP address the node received from the cluster IP pool, not the main IP address.

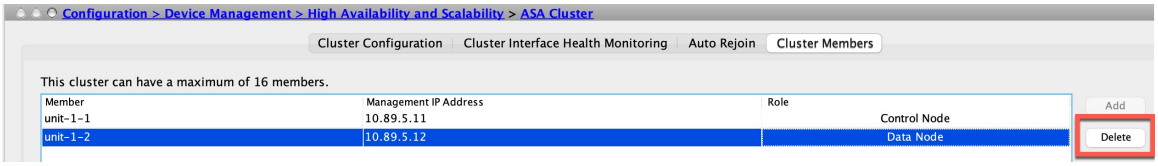
a) You can view the local IP address of the control node by clicking the **Cluster Members** tab. You will need to reconnect to this address after you disable clustering, so note the control node IP address.

Figure 61: Management IP Address of the Control Node



b) On the **Cluster Members** tab, select each data node and then click **Delete**.

Figure 62: Delete Node



The data node bootstrap configuration remains intact, so that you can later re-add the data node without losing your configuration.

c) On the **Cluster Configuration** tab, uncheck the **Participate in ASA cluster** check box, and click **Apply**.

You will be disconnected from ASDM.

d) Reconnect ASDM to the control node's local IP address.

e) Continue to make your changes according to the rest of this procedure. To set each data node to rejoin the cluster after finishing this procedure, you can either add each node as a new node, or you can go to the console port of each data node and re-enable clustering at the CLI.

Example:

```
node2(config)# cluster group cluster1
node2(cfg-cluster)# enable
```

Step 3 Check the **Configure ASA cluster settings** check box.

If you uncheck the check box, the settings are erased. Do not check **Participate in ASA cluster** until after you have set all your parameters.

Note
After you enable clustering, do not uncheck the **Configure ASA cluster settings** check box without understanding the consequences. This action clears all cluster configuration, and also shuts down all interfaces including the management interface to which ASDM is connected. To restore connectivity in this case, you need to access the CLI at the console port.

Step 4 Configure the following bootstrap parameters:

- **Cluster Name**—Names the cluster. The name must be an ASCII string from 1 to 38 characters. You can only configure one cluster per node. All members of the cluster must use the same name.
- **Member Name**—Names this member of the cluster with a unique ASCII string from 1 to 38 characters.
- **Member Priority**—Sets the priority of this node for control node elections, between 1 and 100, where 1 is the highest priority.

- (Routed mode; Spanned EtherChannel mode) **Site Index**—If you use inter-site clustering, set the site ID for this node so it uses a site-specific MAC address, between 1 and 8.
- (Optional) **Site Periodic GARP**—The ASA generates gratuitous ARP (GARP) packets to keep the switching infrastructure up to date: the highest priority member at each site periodically generates GARP traffic for the global MAC/IP addresses. GARP is enabled by default when you set the site ID for each node and the site MAC and IP address for each Spanned EtherChannel. Set the GARP interval between 1 and 1000000 seconds. The default is 290 seconds.

  When using per-site MAC and IP addresses, packets sourced from the cluster use a site-specific MAC address and IP address, while packets received by the cluster use a global MAC address and IP address. If traffic is not generated from the global MAC address periodically, you could experience a MAC address timeout on your switches for the global MAC address. After a timeout, traffic destined for the global MAC address will be flooded across the entire switching infrastructure, which can cause performance and security concerns.
- (Optional) **Shared Key**—Sets an encryption 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 encryption key. This parameter does not affect datapath traffic, including connection state update and forwarded packets, which are always sent in the clear. You must configure this parameter if you also enable the password encryption service.
- (Optional) **Enable connection rebalancing for TCP traffic across all the ASAs in the cluster**—Enables connection rebalancing. This parameter is disabled by default. This parameter is not part of the bootstrap configuration, and is replicated from the control node to the data nodes. If enabled, ASAs exchange information about the connections per second periodically, and offload new connections from devices with more connections per second to less loaded devices. Existing connections are never moved. Moreover, because this command only rebalances based on connections per second, the total number of established connections on each node is not considered, and the total number of connections may not be equal. The frequency, between 1 and 360 seconds, specifies how often the load information is exchanged. The default is 5 seconds.

  Once a connection is offloaded to a different node, it becomes an asymmetric connection.

  Do not configure connection rebalancing for inter-site topologies; you do not want new connections rebalanced to cluster members at a different site.
- **Enable cluster load monitor**—You can monitor the traffic load for cluster members, including total connection count, CPU and memory usage, and buffer drops. If the load is too high, you can choose to manually disable clustering on the node if the remaining nodes can handle the load, or adjust the load balancing on the external switch. This feature is enabled by default. You can periodically monitor the traffic load. If the load is too high, you can choose to manually disable clustering on the node.

  Set the following values:

  - **Time Interval**—Sets the time in seconds between monitoring messages, between 10 and 360 seconds. The default is 20 seconds.
  - **Number of Intervals**—Sets the number of intervals for which the ASA maintains data, between 1 and 60. The default is 30.

  See **Monitoring** > **ASA Cluster** > **Cluster Load-Monitoring** to view the traffic load.
- (Optional) **Enable health monitoring of this device within the cluster**—Enables the cluster node health check feature, and determines the amount of time between node heartbeat status messages, between .3 and 45 seconds; The default is 3 seconds. **Note**: When you are adding new nodes to the cluster, and

making topology changes on the ASA or the switch, you should disable this feature temporarily until the cluster is complete, and also disable interface monitoring for the disabled interfaces (**Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** > **Cluster Interface Health Monitoring**). You can re-enable this feature after cluster and topology changes are complete. To determine node health, the ASA 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 holdtime period, the peer node is considered unresponsive or dead.

  - (Optional) **Broadcast keepalive messages to all EtherChannel cluster control link ports for VSS/vPC support**—If you configure the cluster control link as an EtherChannel (recommended), and it is connected to a VSS or vPC pair, then you might need to enable this option. For some switches, when one node in the VSS/vPC is shutting down or booting up, EtherChannel member interfaces connected to that switch may appear to be Up to the ASA, but they are not passing traffic on the switch side. The ASA can be erroneously removed from the cluster if you set the ASA holdtime timeout to a low value (such as .8 seconds), and the ASA sends heartbeat messages on one of these EtherChannel interfaces. When you enable this option, the ASA floods the heartbeat messages on all EtherChannel interfaces in the cluster control link to ensure that at least one of the switches can receive them.

- (Optional) **Debounce Time**—Configures the debounce time before the ASA considers an interface to be failed and the node is removed from the cluster. This feature allows 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 ASA 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. The default debounce time is 500 ms, with a range of 300 ms to 9 seconds.
- (Optional) **Replicate console output**—Enables console replication from data nodes to the control node. This feature is disabled by default. The ASA may print out some messages directly to the console for certain critical events. If you enable console replication, data nodes send the console messages to the control node so that you only need to monitor one console port for the cluster. This parameter is not part of the bootstrap configuration, and is replicated from the control node to the data nodes.
- (Optional) **Enable Clustering Flow Mobility**. See Configure LISP Inspection, on page 381.
- (Optional) **Enable Director Localization for inter-DC cluster**—To improve performance and reduce round-trip time latency for inter-site clustering for data centers, you can enable director localization. New connections are typically load-balanced and owned by cluster members within a given site. However, the ASA assigns the Director role to a member at *any* site. Director localization enables additional Director roles: a Local Director at the same site as the Owner, and a Global Director that can be at any site. Keeping the Owner and Director at the same site improves performance. Also, if the original Owner fails, the Local Director will choose a new connection Owner at the same site. The Global Director is used if a cluster member receives packets for a connection that is owned on a different site.
- (Optional) **Site Redundancy**—To protect flows from a site failure, you can enable site redundancy. If the connection backup owner is at the same site as the owner, then an additional backup owner will be chosen from another site to protect flows from a site failure. Director localization and site redundancy are separate features; you can configure one or the other, or configure both.
- (Optional) **Enable config sync acceleration**—When a data node has the same configuration as the control node, it will skip syncing the configuration and will join faster. This feature is enabled by default. This feature is configured on each node, and is not replicated from the control node to the data node.

Note
Some configuration commands are not compatible with accelerated cluster joining; if these commands are present on the node, even if accelerated cluster joining is enabled, configuration syncing will always occur. You must remove the incompatible configuration for accelerated cluster joining to work. Use the **show cluster info unit-join-acceleration incompatible-config** to view incompatible configuration.

- **Enable parallel configuration replicate**—Enable the control node to sync configuration changes with data nodes in parallel. Otherwise, synching occurs sequentially, and can take more time.
- **Flow State Refresh Keepalive Interval**—Set the keepalive interval for flow state refresh messages (clu_keepalive and clu_update messages) from the flow owner to the director and backup owner, between 15 and 20 seconds. The default is 15. You may want to set the interval to be longer than the default to reduce the amount of traffic on the cluster control link.
- **Cluster Control Link**—Specifies the cluster control link interface. This interface cannot have a name configured; available interfaces are shown in the drop-down list.
  - **Interface**—Specifies the interface ID, preferably an EtherChannel. Subinterfaces and Management type interfaces are not allowed.
  - **IP Address**—Specifies an IPv4 address for the IP address; IPv6 is not supported for this interface.
  - **Subnet Mask**—Specifies the subnet mask.
  - **MTU**—Specifies the maximum transmission node for the cluster control link interface to be at least 100 bytes higher than the highest MTU of the data interfaces, between 1400 and 9198 bytes. The default MTU is 1500 bytes. 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 cluster control link MTU to the maximum. For example, because the maximum MTU is 9198 bytes, then the highest data interface MTU can be 9098, while the cluster control link can be set to 9198.
- (Optional) **Cluster LACP**—When using Spanned EtherChannels, the ASA uses cLACP to negotiate the EtherChannel with the neighbor switch. ASAs in a cluster collaborate in cLACP negotiation so that they appear as a single (virtual) device to the switch.
  - **Virtual System MAC Address**—Sets the cLACP system ID, which is in the format of a MAC address. All ASAs use the same system ID: auto-generated by the control node (the default) and replicated to all secondaries; or manually specified in the form *H.H.H*, 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. This parameter is not part of the bootstrap configuration, and is replicated from the control node to the data nodes. However, you can only change this value if you disable clustering.
  - **System Priority**—Sets the system priority, between 1 and 65535. The priority is used to decide which node is in charge of making a bundling decision. By default, the ASA uses priority 1, which is the highest priority. The priority needs to be higher than the priority on the switch. This parameter is not part of the bootstrap configuration, and is replicated from the control node to the data nodes. However, you can only change this value if you disable clustering.
- (Optional; Secure Firewall 4200) **VPN Grouping Mode**. See Configure Distributed Site-to-Site VPN, on page 382.

Step 5 Check the **Participate in ASA cluster** check box to join the cluster.

Step 6 Click **Apply**.

Configure Interface Health Monitoring and Auto-Rejoin Settings

You might want to disable health monitoring of non-essential interfaces, for example, the management interface. You can monitor any port-channel ID, redundant ID, or single physical interface ID. 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 **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** > **Cluster Interface Health Monitoring**.

Step 2 In the **Monitored Interfaces** box, select an interface, and click **Add** to move it to the **Unmonitored Interfaces** box.

Interface status messages detect link failure. 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. If a node does not receive interface status messages within the holdtime, then the amount of time before the ASA 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. Health check is enabled by default for all interfaces.

You might want to disable health monitoring of non-essential interfaces, for example, the management interface. You can specify any port-channel ID, redundant ID, or single physical interface ID. 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.

When any topology changes occur (such as adding or removing a data interface, enabling or disabling an interface on the ASA or the switch, or adding an additional switch to form a VSS, vPC, StackWise, or StackWise Virtual) you should disable the health check feature (**Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster**) 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 health check feature.

Step 3 Click the **Auto Rejoin** tab to customize the auto-rejoin settings in case of an interface, system, or cluster control link failure. For each type, click **Edit** to set the following:

- **Maximum Rejoin Attempts**—Define the number of attempts at rejoining the cluster by setting **Unlimited** or a value between 0 and 65535. **0** disables auto-rejoining. The default value is **Unlimited** for the cluster-interface and **3** for the data-interface and system.
- **Rejoin Interval**—Define the interval duration in minutes between rejoin attempts by setting the interval 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**—Define if the interval duration increases by setting the interval variation 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.

Click **Restore Defaults** to restore the default settings.

Step 4 Click **Apply**.

Configure the Cluster TCP Replication Delay

Enable the cluster replication delay for TCP connections to help eliminate the "unnecessary work" related to short-lived flows by delaying the director/backup flow creation. Note that if a unit fails before the director/backup flow is created, then those flows cannot be recovered. Similarly, if traffic is rebalanced to a different unit before the flow is created, then the flow cannot be recovered. You should not enable the TCP replication delay for traffic on which you disable TCP randomization.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster Replication**.

Step 2 Click **Add** and set the following values:

- **Replication delay**—Set the seconds between 1 and 15.
- **HTTP**—Set the delay for all HTTP traffic.
- **Source Criteria**
  - **Source**—Set the source IP address.
  - **Service**—(Optional) Set the source port. Typically you set either the source or the destination port, but not both.
- **Destination Criteria**
  - **Source**—Set the destination IP address.
  - **Service**—(Optional) Set the destination port. Typically you set either the source or the destination port, but not both.

Step 3 Click **OK**.

Step 4 Click **Apply**.

Configure Inter-Site Features

For inter-site clustering, you can customize your configuration to enhance redundancy and stability.

Configure Cluster Flow Mobility

You can inspect LISP traffic to enable flow mobility when a server moves between sites.

About LISP Inspection

You can inspect LISP traffic to enable flow mobility between sites.

About LISP

Data center virtual machine mobility such as VMware VMotion enables servers to migrate between data centers while maintaining connections to clients. To support such data center server mobility, routers need to be able to update the ingress route towards the server when it moves. Cisco Locator/ID Separation Protocol (LISP) architecture separates the device identity, or endpoint identifier (EID), from its location, or routing locator (RLOC), into two different numbering spaces, making server migration transparent to clients. For example, when a server moves to a new site and a client sends traffic to the server, the router redirects traffic to the new location.

LISP requires routers and servers in certain roles, such as the LISP egress tunnel router (ETR), ingress tunnel router (ITR), first hop routers, map resolver (MR), and map server (MS). When the first hop router for the server senses that the server is connected to a different router, it updates all of the other routers and databases so that the ITR connected to the client can intercept, encapsulate, and send traffic to the new server location.

ASA LISP Support

The ASA does not run LISP itself; it can, however, inspect LISP traffic for location changes and then use this information for seamless clustering operation. Without LISP integration, when a server moves to a new site, traffic comes to an ASA cluster member at the new site instead of to the original flow owner. The new ASA forwards traffic to the ASA at the old site, and then the old ASA has to send traffic back to the new site to reach the server. This traffic flow is sub-optimal and is known as “tromboning” or “hair-pinning.”

With LISP integration, the ASA cluster members can inspect LISP traffic passing between the first hop router and the ETR or ITR, and can then change the flow owner to be at the new site.

LISP Guidelines

- The ASA cluster members must reside between the first hop router and the ITR or ETR for the site. The ASA cluster itself cannot be the first hop router for an extended segment.
- Only fully-distributed flows are supported; centralized flows, semi-distributed flows, or flows belonging to individual nodes are not moved to new owners. Semi-distributed flows include applications, such as SIP, where all child flows are owned by the same ASA that owns the parent flow.
- The cluster only moves Layer 3 and 4 flow states; some application data might be lost.
- For short-lived flows or non-business-critical flows, moving the owner may not be worthwhile. You can control the types of traffic that are supported with this feature when you configure the inspection policy, and should limit flow mobility to essential traffic.

ASA LISP Implementation

This feature includes several inter-related configurations (all of which are described in this chapter):

1. (Optional) Limit inspected EIDs based on the host or server IP address—The first hop router might send EID-notify messages for hosts or networks the ASA cluster is not involved with, so you can limit the EIDs to only those servers or networks relevant to your cluster. For example, if the cluster is only involved with 2 sites, but LISP is running on 3 sites, you should only include EIDs for the 2 sites involved with the cluster.
2. LISP traffic inspection—The ASA inspects LISP traffic on UDP port 4342 for the EID-notify message sent between the first hop router and the ITR or ETR. The ASA maintains an EID table that correlates the EID and the site ID. For example, you should inspect LISP traffic with a source IP address of the first hop router and a destination address of the ITR or ETR. Note that LISP traffic is not assigned a director, and LISP traffic itself does not participate in cluster state sharing.

3. Service Policy to enable flow mobility on specified traffic—You should enable flow mobility on business-critical traffic. For example, you can limit flow mobility to only HTTPS traffic, and/or to traffic to specific servers.
4. Site IDs—The ASA uses the site ID for each cluster node to determine the new owner.
5. Cluster-level configuration to enable flow mobility—You must also enable flow mobility at the cluster level. This on/off toggle lets you easily enable or disable flow mobility for a particular class of traffic or applications.

Configure LISP Inspection

You can inspect LISP traffic to enable flow mobility when a server moves between sites.

Before you begin

- Assign each cluster unit to a site ID according to Configure Basic ASA Cluster Parameters, on page 372.
- LISP traffic is not included in the default-inspection-traffic class, so you must configure a separate class for LISP traffic as part of this procedure.

Procedure

Step 1 (Optional) Configure a LISP inspection map to limit inspected EIDs based on IP address, and to configure the LISP pre-shared key:

a) Choose **Configuration** > **Firewall** > **Objects** > **Inspect Maps** > **LISP**.
b) Click **Add** to add a new map.
c) Enter a name (up to 40 characters) and description.
d) For the **Allowed-EID access-list**, click **Manage**.

The **ACL Manager** opens.

The first hop router or ITR/ETR might send EID-notify messages for hosts or networks the ASA cluster is not involved with, so you can limit the EIDs to only those servers or networks relevant to your cluster. For example, if the cluster is only involved with 2 sites, but LISP is running on 3 sites, you should only include EIDs for the 2 sites involved with the cluster.

e) Add an ACL with at least one ACE according to the firewall configuration guide.
f) If necessary, enter the **Validation Key**.

If you copied an encrypted key, click the **Encrypted** radio button.

g) Click **OK**.

Step 2 Add a service policy rule to configure LISP inspection:

a) Choose **Configuration** > **Firewall** > **Service Policy Rules**.
b) Click **Add**.
c) On the **Service Policy** page, apply the rule to an interface or globally.

If you have an existing service policy you want to use, add a rule to that policy. By default, the ASA includes a global policy called **global_policy**. You can also create one service policy per interface if you do not want to apply the policy globally. LISP inspection is applied to traffic bidirectionally so you do

not need to apply the service policy on both the source and destination interfaces; all traffic that enters or exits the interface to which you apply the rule is affected if the traffic matches the class for both directions.

d) On the **Traffic Classification Criteria** page, click **Create a new traffic class**, and under **Traffic Match Criteria**, check **Source and Destination IP Address (uses ACL)**.
e) Click **Next**.
f) Specify the traffic you want to inspect. You should specify traffic between the first hop router and the ITR or ETR on UDP port 4342. Both IPv4 and IPv6 ACLs are accepted.
g) Click **Next**.
h) On the **Rule Actions** wizard page or tab, select the **Protocol Inspection** tab.
i) Check the **LISP** check box,.
j) (Optional) Click **Configure** to choose the inspection map you created.
k) Click **Finish** to save the service policy rule.

Step 3 Add a service policy rule to enable Flow Mobility for critical traffic:

a) Choose **Configuration** > **Firewall** > **Service Policy Rules**.
b) Click **Add**.
c) On the **Service Policy** page, choose the same service policy you used for LISP inspection.
d) On the **Traffic Classification Criteria** page, click **Create a new traffic class**, and under **Traffic Match Criteria**, check **Source and Destination IP Address (uses ACL)**.
e) Click **Next**.
f) Specify the business critical traffic that you want to re-assign to the most optimal site when servers change sites. For example, you can limit flow mobility to only HTTPS traffic, and/or to traffic to specific servers. Both IPv4 and IPv6 ACLs are accepted.
g) Click **Next**.
h) On the **Rule Actions** wizard page or tab, select the **Cluster** tab.
i) Check the **Enable Cluster flow-mobility triggered by LISP EID messages** check box.
j) Click **Finish** to save the service policy rule.

Step 4 Choose **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** > **Cluster Configuration**, and check the **Enable Clustering flow mobility** check box.

Step 5 Click **Apply**.

Configure Distributed Site-to-Site VPN

By default, the cluster uses centralized site-to-site VPN mode. To take advantage of the scalability of clustering, you can enable sistributed site-to-site VPN mode.

About Distributed Site-to-Site VPN

In distributed mode, site-to-site IPsec IKEv2 VPN connections are distributed across nodes of a cluster. Distributing VPN connections across the nodes of a cluster allows both the capacity and throughput of the cluster to be fully utilized, significantly scaling VPN support beyond centralized VPN capabilities.

Distributed VPN Connection Roles

When running in distributed VPN mode, the following roles are assigned to the cluster nodes:

- Active Session Owner—The node that initially receives the connection or that has transitioned a backup session to an active session. The owner maintains the state and processes packets for the complete session, including the IKE and IPsec tunnels and all traffic associated with them.
- Backup Session Owner—The node that is handling the backup session for an existing active session. If the active session owner fails, the backup session owner becomes the active session owner, and a new backup session is established on a different node.
- Forwarder—If traffic associated with a VPN session is sent to a node that does not own the VPN session, that node will use the cluster control link to forward the traffic to the node that owns the VPN session.
- Orchestrator—The orchestrator (always the control node of the cluster) is responsible for calculating which sessions will move and where they will move to when executing an Active Session Redistribution (ASR). It sends a request to the owner node X to move N sessions to node Y. Node X will respond back to the orchestrator when complete, specifying how many sessions it was able to move.

Distributed VPN Session Characteristics

Distributed site-to-site VPN Sessions have the following characteristics. Otherwise, VPN connections behave as they normally do if not on a cluster.

- VPN sessions are distributed across the cluster at the session level. Meaning the same cluster node handles the IKE and IPsec tunnels and all their traffic for a VPN connection. If VPN session traffic is sent to a cluster node that does not own that VPN session, traffic is forwarded to the cluster node that owns the VPN session.
- VPN sessions have a Session ID that is unique across the cluster. Using the session ID, traffic is validated, forwarding decisions are made, and IKE negotiation is completed.
- In a site-to-site VPN hub and spoke configuration, when clients connect through the cluster (called hair-pinning), the session traffic flowing in and the session traffic flowing out may be on different cluster nodes.

Distributed VPN Handling of Cluster Events

Event	Distributed VPN
Node failure	For all active sessions on this failed node, the backup sessions (on another node) become active, and backup sessions are reallocated on another node.
Inactivate a cluster node	For all active sessions on the cluster node being inactivated, backup sessions (on another node) become active and reallocate backup sessions on another node according to the backup strategy.
Cluster node join	If the VPN cluster mode on the new node is not set to distributed, the control node will request a mode change. After the VPN mode is compatible, the cluster node will be assigned active and backup sessions in the flow of normal operations.

IPsec IKEv2 Modifications

IKEv2 is modified while in distributed site-to-site VPN mode in the following ways:

- An identity is used in place of IP/port tuples. This allows for proper forwarding decisions on the packets, and cleanup of previous connections that may be on other cluster members.
- The (SPI) identifiers that identify a single IKEv2 session are locally generated, random 8-byte values that are unique across the cluster. An SPI embeds a time stamp and a cluster node ID. Upon receipt of an IKE negotiation packet, if the time stamp or cluster node ID check fails, the packet is dropped and a message is logged indicating the reason.
- IKEv2 processing has been modified to prevent NAT-T negotiations from failing by being split across cluster members. A new ASP classify domain, *cluster_isakmp_redirect*, and rules are added when IKEv2 is enabled on an interface.

CMPv2

The CMPv2 ID certificate and key pairs are synchronized across the cluster nodes. However, only the control node in the cluster automatically renews and rekeys the CMPv2 certificate. The control node synchronizes these new ID certificates and keys to all cluster nodes on a renewal. In this way, all nodes in the cluster utilize the CMPv2 certificates for authentication, and also any node is capable of taking over as the control node.

Licensing for Distributed Site-to-Site VPN

A Carrier license is required for distributed site-to-site VPN, on each member of the cluster.

Each VPN connection requires two *Other VPN* licensed sessions (the *Other VPN* license is part of the *Essentials* license), one for the active session and one for the backup session. The maximum VPN session capacity of the cluster can be no more than half of the licensed capacity due to using two licenses for each session.

Prerequisites for Distributed Site-to-Site VPN

Model Support

- Secure Firewall 4200

Cluster Requirements

- Spanned EtherChannel mode.
- Routed firewall mode.

Guidelines for Distributed Site-to-Site VPN

Firewall Mode

Distributed site-to-site VPN is supported in routed mode only.

Context Mode

Distributed site-to-site VPN operates in both single and multiple context modes. However, in multiple context mode, active session redistribution is done at the system level, not at the context level. This prevents an active session associated with a context from moving to a cluster member that contains active sessions associated with a different context, unknowingly creating an unsupportable load.

Unsupported Inspections

The following types of inspections are not supported or are disabled in distributed site-to-site VPN mode:

- CTIQBE
- DCERPC
- H323, H225, and RAS
- IPsec pass-through
- MGCP
- MMP
- NetBIOS
- PPTP
- RADIUS
- RSH
- RTSP
- SCCP (Skinny)
- SUNRPC
- TFTP
- WAAS
- WCCP
- XDMCP

Additional Guidelines

- Only IKEv2 IPsec site-to-site VPN is supported in distributed site-to-site VPN mode. IKEv1 is not supported. IKEv1 site-to-site is supported in centralized VPN mode.
- Inter-site clustering is not supported.
- Dynamic PAT is not available while in distributed site-to-site VPN mode.

Enable Distributed Site-to-Site VPN

Enable distributed site-to-site VPN to take advantage of the scalability of clustering for VPN sessions.



Note Changing the VPN mode between centralized and distributed requires all nodes in the cluster to reload.

Before you begin

Configure site-to-site VPN according to the VPN configuration guide.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** > **Cluster Configuration**.

Step 2 In the **VPN Cluster Mode** area, choose the **VPN Mode** for the cluster, **Centralized** or **Distributed**.

Step 3 Click **Apply**.

You will be prompted to reload. This setting is replicated to all data nodes before reloading; all nodes in the cluster will reload.

Redistribute Distributed S2S VPN Sessions

Active Session Redistribution (ASR) redistributes the active VPN session load across the cluster members. Due to the dynamic nature of beginning and ending sessions, ASR is a best effort balancing of the sessions across all cluster members. Repeated redistribution actions will optimize the balance.

Redistribution can be run at any time, should be run after any topology change in the cluster, and is recommended after a new member joins the cluster. The goal of redistribution is to creat a stable VPN cluster. A stable VPN cluster has an almost equal number of active and backup sessions across the nodes.

To move a session, the backup session becomes the active one and another node is selected to host a new backup session. Moving sessions is dependent on the location of the active session's backup and the number of active sessions already on that particular backup node. If the backup session node is unable to host the active session for some reason, the original node remains owner of the session.

In multiple-context mode, active session redistribution is done at the system level, not the individual context level. It is not done at the context level because an active session in one context could be moved a member that contains many more active sessions in a different context, creating more load on that cluster member.

Before you begin

- Enable system logs if you would like to monitor redistribution activity.
- This procedure must be carried out on the control unit of the cluster.

Procedure

Step 1 Choose **Monitoring > ASA Cluster > ASA Cluster > Cluster Summary > VPN Cluster Summary** to view how active and backup sessions are distributed across the cluster.

Depending on the number of sessions to redistribute and the load on the cluster, this may take some time. Syslogs containing the following phrases (and other system details not shown here) are provided as redistribution activity occurs:

Syslog Phrase	Notes
VPN session redistribution started	Control node only

Syslog Phrase	Notes
Sent request to move *number* sessions from *orig-member-name* to *dest-member-name*	Control node only
Failed to send session redistribution message to *member-name*	Control node only
Received request to move *number* sessions from *orig-member-name* to *dest-member-name*	Data node only
Moved *number* sessions to *member-name*	The number of active sessions moved to the named cluster.
Failed to receive session move response from *dest-member-name*	Control node only
VPN session completed	Control node only
Cluster topology change detected. VPN session redistribution aborted.	

Step 2 Click **Re-Distribute**.

Step 3 Refresh the **Monitoring > ASA Cluster > ASA Cluster > Cluster Summary > VPN Cluster Summary** to see the results of the redistribution activity.

If your redistribution was successful, and there has been no substantial system or session activity, your system will be balanced and this action is complete.

Otherwise, repeat the redistribution process to obtain a balanced and stable system.

Manage Cluster Nodes

After you deploy the cluster, you can change the configuration and manage cluster nodes.

Add a New Data Node from the Control Node

You can add additional data nodes to the cluster from the control node. You can also add data nodes using the High Availability and Scalability wizard. Adding a data node from the control node has the benefit of configuring the cluster control link and setting the cluster interface mode on each data node you add.

You can alternatively log into the data node and configure clustering directly on the node. However, after you enable clustering, your ASDM session will be disconnected, and you will have to reconnect.

Before you begin

- For multiple context mode, complete this procedure in the system execution space. If you are not already in the System configuration mode, in the **Configuration > Device List** pane, double-click **System** under the active device IP address.
- If you want to send the bootstrap configuration over the management network, be sure the data node has an accessible IP address.

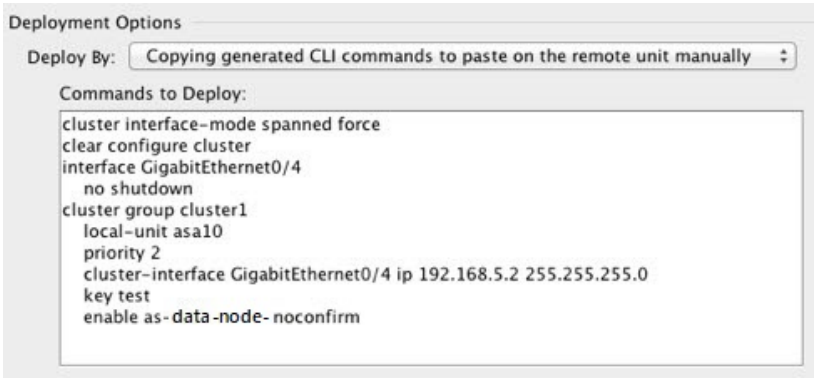
Procedure

Step 1 Choose **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** > **Cluster Members**.

Step 2 Click **Add**.

Step 3 Configure the following parameters:

- **Member Name**—Names this member of the cluster with a unique ASCII string from 1 to 38 characters.
- **Member Priority**—Sets the priority of this node for control node elections, between 1 and 100, where 1 is the highest priority.
- **Cluster Control Link > IP Address**—Specifies a unique IP address for this member for the cluster control link, on the same network as the control node cluster control link.
- In the **Deployment Options** area, choose one of the following **Deploy By** options:
  - **Sending CLI commands to the remote unit now**—Send the bootstrap configuration to the data node (temporary) management IP address. Enter the data node management IP address, username, and password.
  - **Copying generated CLI commands to paste on the remote unit manually**—Generates the commands so that you can cut and paste them at the data node CLI or using the CLI tool in ASDM. In the Commands to Deploy box, select and copy the generated commands for later use.



Step 4 Click **OK**, then **Apply**.

Become an Inactive Node

To become an inactive member of the cluster, disable clustering on the node while leaving the clustering configuration intact.

Note When an ASA becomes inactive (either manually or through a health check failure), all data interfaces are shut down; only the management-only interface can send and receive traffic. To resume traffic flow, re-enable clustering; or you can remove the node altogether from the cluster. The management interface remains up using the IP address the node received from the cluster IP pool. However if you reload, and the node is still inactive in the cluster (for example, you saved the configuration with clustering disabled), then the management interface is disabled. You must use the console port for any further configuration.

Before you begin

- For multiple context mode, perform this procedure in the system execution space. If you are not already in the System configuration mode in the Configuration > Device List pane, double-click **System** under the active device IP address.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** > **Cluster Configuration**.

Step 2 Uncheck the **Participate in ASA cluster** check box.

Note
Do not uncheck the **Configure ASA cluster settings** check box; this action clears all cluster configuration, and also shuts down all interfaces including the management interface to which ASDM is connected. To restore connectivity in this case, you need to access the CLI at the console port.

Step 3 Click **Apply**.

Deactivate a Data Node from the Control Node

To deactivate a data node, perform the following steps.

Note When an ASA becomes inactive, all data interfaces are shut down; only the management-only 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 cluster IP pool. However if you reload, and the node is still inactive in the cluster (for example, if you saved the configuration with clustering disabled), the management interface is disabled. You must use the console port for any further configuration.

Before you begin

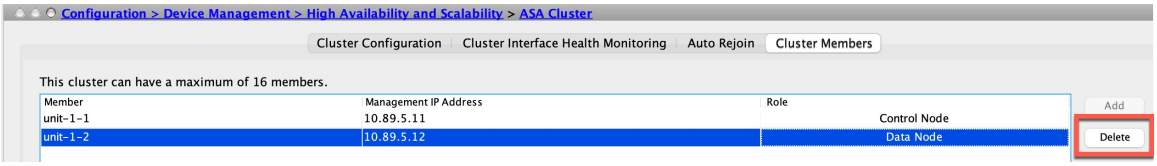
For multiple context mode, perform this procedure in the system execution space. If you are not already in the System configuration mode in the **Configuration > Device List** pane, double-click **System** under the active device IP address.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** > **Cluster Members**.

Step 2 Select the data node that you want to remove, and click **Delete**.

Figure 63: Delete Node



The data node bootstrap configuration remains intact, so that you can later re-add the data node without losing your configuration.

Step 3 Click **Apply**.

Rejoin the Cluster

If a node was removed from the cluster, for example for a failed interface or if you manually deactivated a member, you must manually rejoin the cluster.

Before you begin

- You must use the console port to reenable clustering. Other interfaces are shut down. The exception is if you manually disabled clustering in ASDM, then you can reenable clustering in ASDM if you did not save the configuration and reload. After reloading, the management interface is disabled, so console access is the only method to reenable clustering.
- For multiple context mode, perform this procedure in the system execution space. If you are not already in the System configuration mode in the **Configuration > Device List** pane, double-click **System** under the active device IP address.
- Make sure the failure is resolved before you try to rejoin the cluster.

Procedure

Step 1 If you still have ASDM access, you can reenable clustering in ASDM by connecting ASDM to the node you want to reenable.

You cannot reenable clustering for a data node from the control node unless you add it as a new member.

a) Choose **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster**.
b) Check the **Participate in ASA cluster** check box.
c) Click **Apply**.

Step 2 If you cannot use ASDM: At the console, enter cluster configuration mode:

cluster group *name*

Example:

```
ciscoasa(config)# cluster group pod1
```

Step 3 Enable clustering.

enable

Leave the Cluster

If you want to leave the cluster altogether, you need to remove the entire cluster bootstrap configuration. Because the current configuration on each node is the same (synced from the active unit), leaving the cluster also means either restoring a pre-clustering configuration from backup, or clearing your configuration and starting over to avoid IP address conflicts.

Before you begin

You must use the console port; when you remove the cluster configuration, all interfaces are shut down, including the management interface and cluster control link.

Procedure

Step 1 For a data node, disable clustering:

cluster group *cluster_name*
no enable

Example:

```
ciscoasa(config)# cluster group cluster1
ciscoasa(cfg-cluster)# no enable
```

You cannot make configuration changes while clustering is enabled on a data node.

Step 2 Clear the cluster configuration:

clear configure cluster

The ASA shuts down all interfaces including the management interface and cluster control link.

Step 3 Disable cluster interface mode:

no cluster interface-mode

The mode is not stored in the configuration and must be reset manually.

Step 4 If you have a backup configuration, copy the backup configuration to the running configuration:

copy *backup_cfg* **running-config**

Example:

```
ciscoasa(config)# copy backup_cluster.cfg running-config

Source filename [backup_cluster.cfg]?

Destination filename [running-config]?
ciscoasa(config)#
```

Step 5 Save the configuration to startup:

write memory

Step 6 If you do not have a backup configuration, reconfigure management access. Be sure to change the interface IP addresses, and restore the correct hostname, for example.

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 node you want to become the control node, use the procedure in this section. Note, however, that for centralized features, if you force a control node change using this procedure, 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.

Before you begin

For multiple context mode, perform this procedure in the system execution space. If you are not already in the System configuration mode in the Configuration > Device List pane, double-click **System** under the active device IP address.

Procedure

Step 1 Choose **Monitoring** > **ASA Cluster** > **Cluster Summary**.

Step 2 From the drop-down list, choose a data node to become control, and click the button to make it the control node.

Step 3 You are prompted to confirm the control node change. Click **Yes**.

Step 4 Quit ASDM, and reconnect using the Main cluster IP address.

Execute a Command Cluster-Wide

To send a command to all nodes in the cluster, or to a specific node, perform the following steps. Sending a **show** command to all nodes collects all output and displays it on the console of the current node. Other commands, such as **capture** and **copy**, can also take advantage of cluster-wide execution.

Before you begin

Perform this procedure at the Command Line Interface tool: choose **Tools > Command Line Interface**.

Procedure

Send a command to all nodes, or if you specify the node name, a specific node:

cluster exec [**unit** *node_name*] *command*

Example:

```
ciscoasa# cluster exec show xlate
```

To view node names, enter **cluster exec unit ?** (to see all names except the current node), or enter the **show cluster info** command.

Examples

To copy the same capture file from all nodes in the cluster at the same time to a TFTP server, enter the following command on the control node:

```
ciscoasa# cluster exec copy /pcap capture: tftp://10.1.1.56/capture1.pcap
```

Multiple PCAP files, one from each node, are copied to the TFTP server. The destination capture file name is automatically attached with the node name, such as capture1_asa1.pcap, capture1_asa2.pcap, and so on. In this example, asa1 and asa2 are cluster node names.

The following sample output for the **cluster exec show port-channel** summary command shows EtherChannel information for each node in the cluster:

```
ciscoasa# cluster exec show port-channel summary
control node(LOCAL):**********************************************************
 Number of channel-groups in use: 2
Group  Port-channel  Protocol  Span-cluster  Ports
------+-------------+-----------+-----------------------------------------------
1         Po1             LACP       Yes  Gi0/0(P)
2         Po2             LACP       Yes  Gi0/1(P)
 slave:************************************************************
 Number of channel-groups in use: 2
Group  Port-channel  Protocol  Span-cluster   Ports
------+-------------+-----------+-----------------------------------------------
1         Po1             LACP       Yes   Gi0/0(P)
```

```
2          Po2              LACP       Yes    Gi0/1(P)
```

Monitoring the ASA Cluster

You can monitor and troubleshoot cluster status and connections.

Monitoring Cluster Status

See the following screens for monitoring cluster status:

- **Monitoring > ASA Cluster > Cluster Summary**

  This pane shows cluster information about the node to which you are connected, as well as other nodes in the cluster. You can also change the primary node from this pane.

- **Cluster Dashboard**

  On the home page on the primary node, you can monitor the cluster using the Cluster Dashboard and the Cluster Firewall Dashboard.

Capturing Packets Cluster-Wide

See the following screen for capturing packets in a cluster:

Wizards > Packet Capture Wizard

To support cluster-wide troubleshooting, you can enable capture of cluster-specific traffic on the control node, which is then automatically enabled on all of the data nodes in the cluster.

Monitoring Cluster Resources

See the following screens for monitoring cluster resources:

- **Monitoring > ASA Cluster > System Resources Graphs > CPU**

  This pane lets you create graphs or tables showing the CPU utilization across the cluster nodes.

- **Monitoring > ASA Cluster > System Resources Graphs > Memory**. This pane lets you create graphs or tables showing the Free Memory and Used Memory across the cluster nodes.

Monitoring Cluster Traffic

See the following screens for monitoring cluster traffic:

- **Monitoring > ASA Cluster > Traffic Graphs > Connections**.

  This pane lets you create graphs or tables showing the Connections across the cluster members.

- **Monitoring > ASA Cluster > Traffic Graphs > Throughput**.

  This pane lets you create graphs or tables showing the traffic throughput across the cluster members.

- **Monitoring** > **ASA Cluster** > **Cluster Load-Monitoring**

  This section includes the **Load Monitor-Information** and **Load-Monitor Details** panes. **Load Monitor-Information** shows the traffic load for cluster members for the last interval and also the average over total number of intervals configured (30 by default). Use the **Load-Monitor Details** pane to view the value for each measure at each interval.

Monitoring the Cluster Control Link

See the following screen for monitoring cluster status:

Monitoring > Properties > System Resources Graphs > Cluster Control Link.

This pane lets you create graphs or tables showing the cluster control link Receival and Transmittal capacity utilization.

Monitoring Cluster Routing

See the following screen for cluster routing:

- **Monitoring** > **Routing** > **LISP-EID Table**

  Shows the ASA EID table showing EIDs and site IDs.

Monitoring Distributed Site-to-Site VPN

See the following screens for monitoring VPN cluster status:

- **Monitoring > ASA Cluster > ASA Cluster > Cluster Summary > VPN Cluster Summary**

  Displays the distribution of the session across the cluster and provides you with the ability to redistribute the sessions.

- **Monitoring > VPN > VPN Statistics > Sessions**

  Both control and data cluster nodes are listed. Click any node for details.

Configuring Logging for Clustering

See the followingscreen for configuring logging for clustering:

Configuration > Device Management > Logging > Syslog Setup

Each node in the cluster generates syslog messages independently. You can generate syslog messages with identical or different device IDs to make messages appear to come from the same or different nodes in the cluster.

Troubleshooting Distributed Site-to-Site VPN

Distributed VPN Notifications

You will be notified with messages containing the identified phrases when the following error situations occur on a cluster running distributed VPN:

Situation	Notification

If an existing or joining cluster data node is not in distributed VPN mode when attempting to join the cluster:	`New cluster member (member-name) rejected due to vpn mode mismatch.`
	and
	`Control node (control-name) rejects enrollment request from unit (unit-name) for the reason: the vpn mode capabilities are not compatible with the control node configuration`

If licensing is not properly configured on a cluster member for Distributed VPN:	`ERROR: Control node requested cluster vpn-mode change to distributed. Unable to change mode due to missing Carrier License.`

If the time stamp or member ID is invalid in the SPI of a received IKEv2 packet:	`Expired SPI received`
	or
	`Corrupted SPI detected`

If the cluster is unable to create a backup session:	`Failed to create the backup for an IKEv2 session.`
IKEv2 Initial Contact (IC) processing error:	`IKEv2 Negotiation aborted due to ERROR: Stale backup session found on backup`
Redistribution problems:	`Failed to send session redistribution message to member-name` `Failed to receive session move response from member-name (control node only)`
If the topology changes during redistribution of the sessions:	`Cluster topology change detected. VPN session redistribution aborted.`

You may be encountering one of the following situations:

- Site-to-site VPN sessions are being distributed to only one of the chassis in a cluster when the Nexus 7K switch is configured with a layer 4 port as a load-balancing algorithm using the **port-channel load-balance src-dst l4port** command. An example of the cluster session allocation looks like below:

```
SSP-Cluster/data node(cfg-cluster)# show cluster vpn-sessiondb distribution
Member 0 (unit-1-3): active: 0
Member 1 (unit-2-2): active: 13295; backups at: 0(2536), 2(2769), 3(2495), 4(2835),
5(2660)
```

```
Member 2 (unit-2-3): active: 12174; backups at: 0(2074), 1(2687), 3(2207), 4(3084),
5(2122)
Member 3 (unit-2-1): active: 13416; backups at: 0(2419), 1(3013), 2(2712), 4(2771),
5(2501)
Member 4 (unit-1-1): active: 0
Member 5 (unit-1-2): active: 0
```

Since site-to-site IKEv2 VPN uses port 500 for both source and destination ports, IKE packets are only sent to one of the links in the port channel connected between the Nexus 7K and the chassis.

Change the Nexus 7K load balancing algorithm to IP and Layer 4 port using the **port-channel load-balance src-dst ip-l4port**. Then the IKE packets are sent to all the links and thus all nodes.

For a more immediate adjustment, on the control node of the cluster, execute: **cluster redistribute vpn-sessiondb** to redistribute active VPN sessions to the cluster nodes of the other chassis.

Examples for ASA Clustering

These examples include all cluster-related ASA configuration for typical deployments.

Sample ASA and Switch Configuration

The following sample configurations connect the following interfaces between the ASA and the switch:

ASA Interface	Switch Interface
Ethernet 1/2	GigabitEthernet 1/0/15
Ethernet 1/3	GigabitEthernet 1/0/16
Ethernet 1/4	GigabitEthernet 1/0/17
Ethernet 1/5	GigabitEthernet 1/0/18

ASA Configuration

Interface Mode on Each Unit

```
cluster interface-mode spanned force
```

ASA1 Control Unit Bootstrap Configuration

```
interface Ethernet1/6
 channel-group 1 mode on
 no shutdown
!
interface Ethernet1/7
 channel-group 1 mode on
 no shutdown
!
interface Port-channel1
 description Clustering Interface
```

```
!
cluster group Moya
 local-unit A
 cluster-interface Port-channel1 ip 10.0.0.1 255.255.255.0
 priority 10
 key emphyri0
 enable noconfirm
```

ASA2 Data Unit Bootstrap Configuration

```
interface Ethernet1/6
 channel-group 1 mode on
 no shutdown
!
interface Ethernet1/7
 channel-group 1 mode on
 no shutdown
!
interface Port-channel1
 description Clustering Interface
!
cluster group Moya
 local-unit B
 cluster-interface Port-channel1 ip 10.0.0.2 255.255.255.0
 priority 11
 key emphyri0
 enable as-data-node
```

Control Unit Interface Configuration

```
ip local pool mgmt-pool 10.53.195.231-10.53.195.232

interface Ethernet1/2
 channel-group 10 mode active
 no shutdown
!
interface Ethernet1/3
 channel-group 10 mode active
 no shutdown
!
interface Ethernet1/4
 channel-group 11 mode active
 no shutdown
!
interface Ethernet1/5
 channel-group 11 mode active
 no shutdown
!
interface Management1/1
 management-only
 nameif management
 ip address 10.53.195.230 cluster-pool mgmt-pool
 security-level 100
 no shutdown
!
interface Port-channel10
 mac-address aaaa.bbbb.cccc
 nameif inside
 security-level 100
```

```
 ip address 209.165.200.225 255.255.255.224
!
interface Port-channel11
 mac-address aaaa.dddd.cccc
 nameif outside
 security-level 0
 ip address 209.165.201.1 255.255.255.224
```

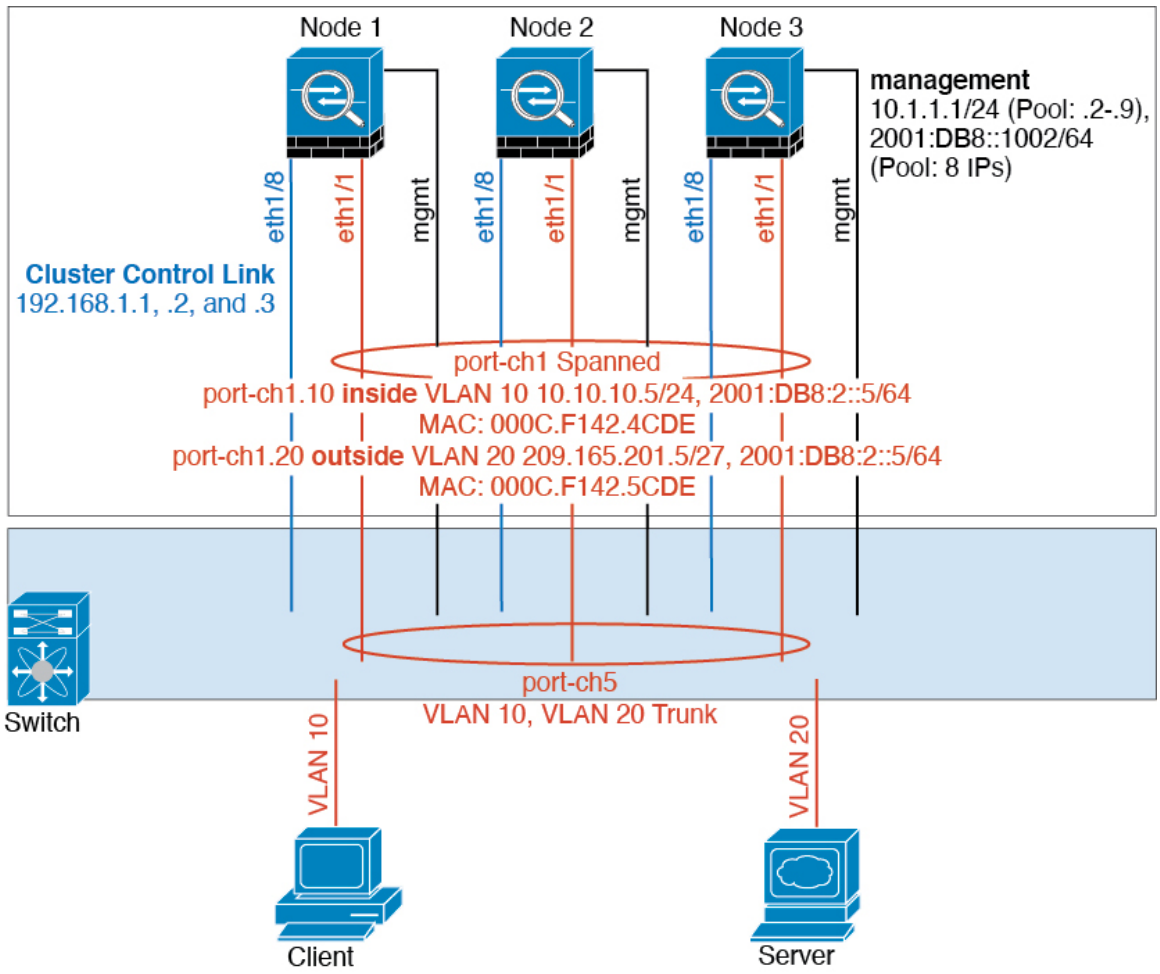
Cisco IOS Switch Configuration

```
interface GigabitEthernet1/0/15
 switchport access vlan 201
 switchport mode access
 spanning-tree portfast
 channel-group 10 mode active
!
interface GigabitEthernet1/0/16
 switchport access vlan 201
 switchport mode access
 spanning-tree portfast
 channel-group 10 mode active
!
interface GigabitEthernet1/0/17
 switchport access vlan 401
 switchport mode access
 spanning-tree portfast
 channel-group 11 mode active
!
interface GigabitEthernet1/0/18
 switchport access vlan 401
 switchport mode access
 spanning-tree portfast
 channel-group 11 mode active

interface Port-channel10
 switchport access vlan 201
 switchport mode access

interface Port-channel11
 switchport access vlan 401
 switchport mode access
```

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 ASA 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 ASA 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 ASA becomes unavailable, the switch will rebalance traffic between the remaining units.

Interface Mode on Each Unit

```
cluster interface-mode spanned force
```

Unit 1 Control Unit Bootstrap Configuration

```
interface ethernet1/8
no shutdown
description CCL
```

```
cluster group cluster1
local-unit asa1
cluster-interface ethernet1/8 ip 192.168.1.1 255.255.255.0
priority 1
key chuntheunavoidable
enable noconfirm
```

Unit 2 Data Unit Bootstrap Configuration

```
interface ethernet1/8
no shutdown
description CCL

cluster group cluster1
local-unit asa2
cluster-interface ethernet1/8 ip 192.168.1.2 255.255.255.0
priority 2
key chuntheunavoidable
enable as-data-node
```

Unit 3 Data Unit Bootstrap Configuration

```
interface ethernet1/8
no shutdown
description CCL

cluster group cluster1
local-unit asa3
cluster-interface ethernet1/8 ip 192.168.1.3 255.255.255.0
priority 3
key chuntheunavoidable
enable as-data-node
```

Control Unit Interface Configuration

```
ip local pool mgmt 10.1.1.2-10.1.1.9
ipv6 local pool mgmtipv6 2001:DB8::1002/64 8

interface management 1/1
nameif management
ip address 10.1.1.1 255.255.255.0 cluster-pool mgmt
ipv6 address 2001:DB8::1001/32 cluster-pool mgmtipv6
security-level 100
management-only
no shutdown

interface ethernet1/1
channel-group 1 mode active
no shutdown

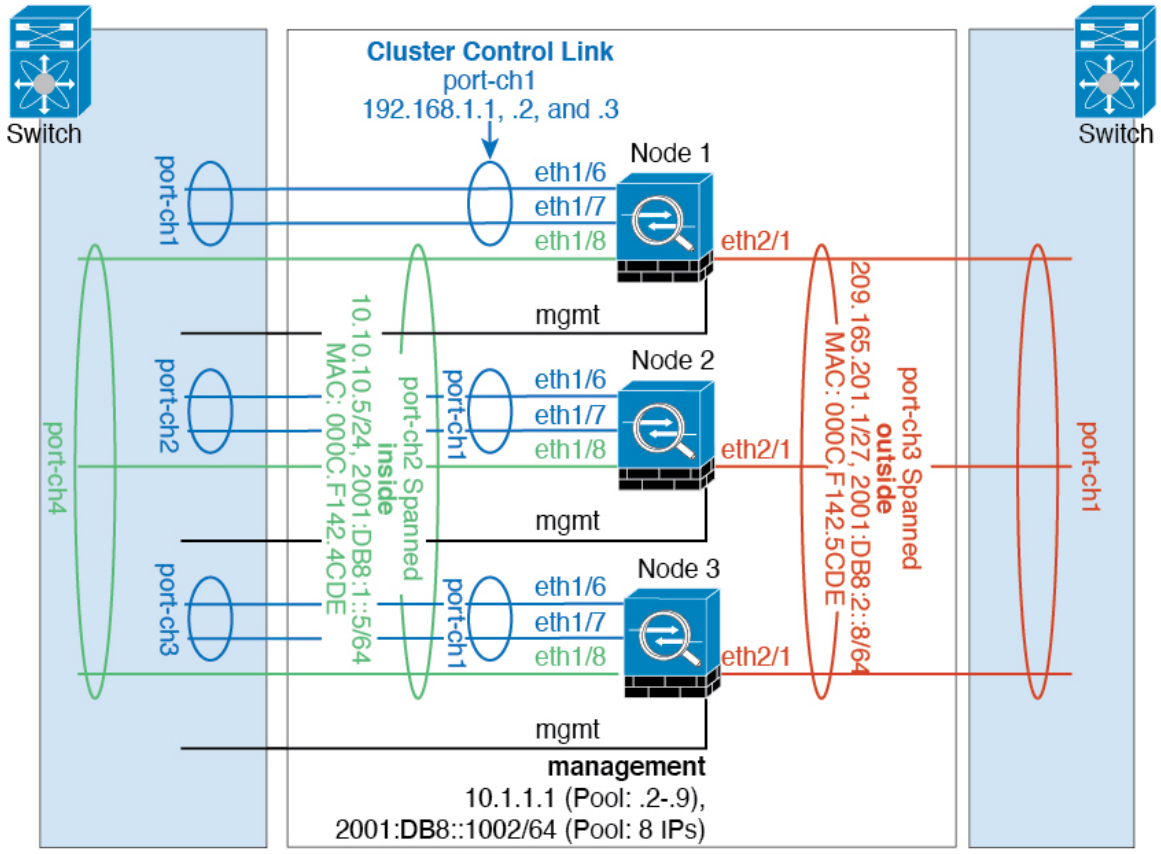
interface port-channel 1

interface port-channel 1.10
vlan 10
nameif inside
ip address 10.10.10.5 255.255.255.0
```

```
ipv6 address 2001:DB8:1::5/64
mac-address 000C.F142.4CDE

interface port-channel 1.20
vlan 20
nameif outside
ip address 209.165.201.1 255.255.255.224
ipv6 address 2001:DB8:2::8/64
mac-address 000C.F142.5CDE
```

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.

Interface Mode on Each Unit

```
cluster interface-mode spanned force
```

Unit 1 Control Unit Bootstrap Configuration

```
interface ethernet 1/6
 channel-group 1 mode on
 no shutdown

interface ethernet 1/7
 channel-group 1 mode on
 no shutdown

interface port-channel 1
 description CCL

cluster group cluster1
 local-unit asa1
 cluster-interface port-channel1 ip 192.168.1.1 255.255.255.0
 priority 1
 key chuntheunavoidable
 enable noconfirm
```

Unit 2 Data Unit Bootstrap Configuration

```
interface ethernet 1/6
 channel-group 1 mode on
 no shutdown

interface ethernet 1/7
 channel-group 1 mode on
 no shutdown

interface port-channel 1
 description CCL

cluster group cluster1
 local-unit asa2
 cluster-interface port-channel1 ip 192.168.1.2 255.255.255.0
 priority 2
 key chuntheunavoidable
 enable as-data-node
```

Unit 3 Data Unit Bootstrap Configuration

```
interface ethernet 1/6
 channel-group 1 mode on
 no shutdown

interface ethernet 1/7
 channel-group 1 mode on
 no shutdown

interface port-channel 1
 description CCL

cluster group cluster1
 local-unit asa3
 cluster-interface port-channel1 ip 192.168.1.3 255.255.255.0
 priority 3
 key chuntheunavoidable
```

```
 enable as-data-node
```

Control Unit Interface Configuration

```
ip local pool mgmt 10.1.1.2-10.1.1.9
ipv6 local pool mgmtipv6 2001:DB8::1002/64 8

interface management 1/1
 nameif management
 ip address 10.1.1.1 255.255.255.0 cluster-pool mgmt
 ipv6 address 2001:DB8::1001/32 cluster-pool mgmtipv6
 security-level 100
 management-only
 no shutdown

interface ethernet 1/8
 channel-group 2 mode active
 no shutdown

interface port-channel 2
 nameif inside
 ip address 10.10.10.5 255.255.255.0
 ipv6 address 2001:DB8:1::5/64
 mac-address 000C.F142.4CDE

interface ethernet 2/1
 channel-group 3 mode active
 no shutdown

interface port-channel 3
 nameif outside
 ip address 209.165.201.1 255.255.255.224
 ipv6 address 2001:DB8:2::8/64
 mac-address 000C.F142.5CDE
```

OTV Configuration for Routed Mode Inter-Site Clustering

The success of inter-site clustering for routed mode with Spanned EtherChannels depends on the proper configuration and monitoring of OTV. OTV plays a major role by forwarding the packets across the DCI. OTV forwards unicast packets across the DCI only when it learns the MAC address in its forwarding table. If the MAC address is not learned in the OTV forwarding table, it will drop the unicast packets.

Sample OTV Configuration

```
//Sample OTV config:
//3151 – Inside VLAN, 3152 – Outside VLAN, 202 – CCL VLAN
//aaaa.1111.1234 – ASA inside interface global vMAC
//0050.56A8.3D22 – Server MAC

feature ospf
feature otv

mac access-list ALL_MACs
  10 permit any any
mac access-list HSRP_VMAC
  10 permit aaaa.1111.1234 0000.0000.0000 any
  20 permit aaaa.2222.1234 0000.0000.0000 any
```

```
  30 permit any aaaa.1111.1234 0000.0000.0000
  40 permit any aaaa.2222.1234 0000.0000.0000
vlan access-map Local 10
  match mac address HSRP_VMAC
  action drop
vlan access-map Local 20
  match mac address ALL_MACs
  action forward
vlan filter Local vlan-list 3151-3152


//To block global MAC with ARP inspection:
arp access-list HSRP_VMAC_ARP
  10 deny aaaa.1111.1234 0000.0000.0000 any
  20 deny aaaa.2222.1234 0000.0000.0000 any
  30 deny any aaaa.1111.1234 0000.0000.0000
  40 deny any aaaa.2222.1234 0000.0000.0000
  50 permit ip any mac
ip arp inspection filter HSRP_VMAC_ARP 3151-3152

no ip igmp snooping optimise-multicast-flood
vlan 1,202,1111,2222,3151-3152

otv site-vlan 2222
mac-list GMAC_DENY seq 10 deny aaaa.aaaa.aaaa ffff.ffff.ffff
mac-list GMAC_DENY seq 20 deny aaaa.bbbb.bbbb ffff.ffff.ffff
mac-list GMAC_DENY seq 30 permit 0000.0000.0000 0000.0000.0000
route-map stop-GMAC permit 10
  match mac-list GMAC_DENY

interface Overlay1
  otv join-interface Ethernet8/1
  otv control-group 239.1.1.1
  otv data-group 232.1.1.0/28
  otv extend-vlan 202, 3151
  otv arp-nd timeout 60
  no shutdown

interface Ethernet8/1
  description uplink_to_OTV_cloud
  mtu 9198
  ip address 10.4.0.18/24
  ip igmp version 3
  no shutdown

interface Ethernet8/2

interface Ethernet8/3
  description back_to_default_vdc_e6/39
  switchport
    switchport mode trunk
    switchport trunk allowed vlan 202,2222,3151-3152
  mac packet-classify
  no shutdown

otv-isis default
  vpn Overlay1
    redistribute filter route-map stop-GMAC
otv site-identifier 0x2
//OTV flood not required for ARP inspection:
otv flood mac 0050.56A8.3D22 vlan 3151
```

OTV Filter Modifications Required Because of Site Failure

If a site goes down, the filters need to be removed from OTV because you do not want to block the global MAC address anymore. There are some additional configurations required.

You need to add a static entry for the ASA global MAC address on the OTV switch in the site that is functional. This entry will let the OTV at the other end add these entries on the overlay interface. This step is required because if the server and client already have the ARP entry for the ASA, which is the case for existing connections, then they will not send the ARP again. Therefore, OTV will not get a chance to learn the ASA global MAC address in its forwarding table. Because OTV does not have the global MAC address in its forwarding table, and per OTV design it will not flood unicast packets over the overlay interface, then it will drop the unicast packets to the global MAC address from the server, and the existing connections will break.

```
//OTV filter configs when one of the sites is down

mac-list GMAC_A seq 10 permit 0000.0000.0000 0000.0000.0000
route-map a-GMAC permit 10
  match mac-list GMAC_A

otv-isis default
  vpn Overlay1
    redistribute filter route-map a-GMAC

no vlan filter Local vlan-list 3151

//For ARP inspection, allow global MAC:
arp access-list HSRP_VMAC_ARP_Allow
  50 permit ip any mac
ip arp inspection filter HSRP_VMAC_ARP_Allow 3151-3152

mac address-table static aaaa.1111.1234 vlan 3151 interface Ethernet8/3
//Static entry required only in the OTV in the functioning Site
```

When the other site is restored, you need to add the filters back again and remove this static entry on the OTV. It is very important to clear the dynamic MAC address table on both the OTVs to clear the overlay entry for the global MAC address.

MAC Address Table Clearing

When a site goes down, and a static entry for the global MAC address is added to OTV, you need to let the other OTV learn the global MAC address on the overlay interface. After the other site comes up, these entries should be cleared. Make sure to clear the mac address table to make sure OTV does not have these entries in its forwarding table.

```
cluster-N7k6-OTV# show mac address-table
Legend:
* - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
age - seconds since last seen,+ - primary entry using vPC Peer-Link,
(T) - True, (F) - False
VLAN MAC Address Type age Secure NTFY Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
G -    d867.d900.2e42 static  -  F F sup-eth1(R)
O 202  885a.92f6.44a5 dynamic -  F F Overlay1
* 202  885a.92f6.4b8f dynamic 5  F F Eth8/3
O 3151 0050.5660.9412 dynamic -  F F Overlay1
* 3151 aaaa.1111.1234 dynamic 50 F F Eth8/3
```

OTV ARP Cache Monitoring

OTV maintains an ARP cache to proxy ARP for IP addresses that it learned across the OTV interface.

```
cluster-N7k6-OTV# show otv arp-nd-cache
OTV ARP/ND L3->L2 Address Mapping Cache

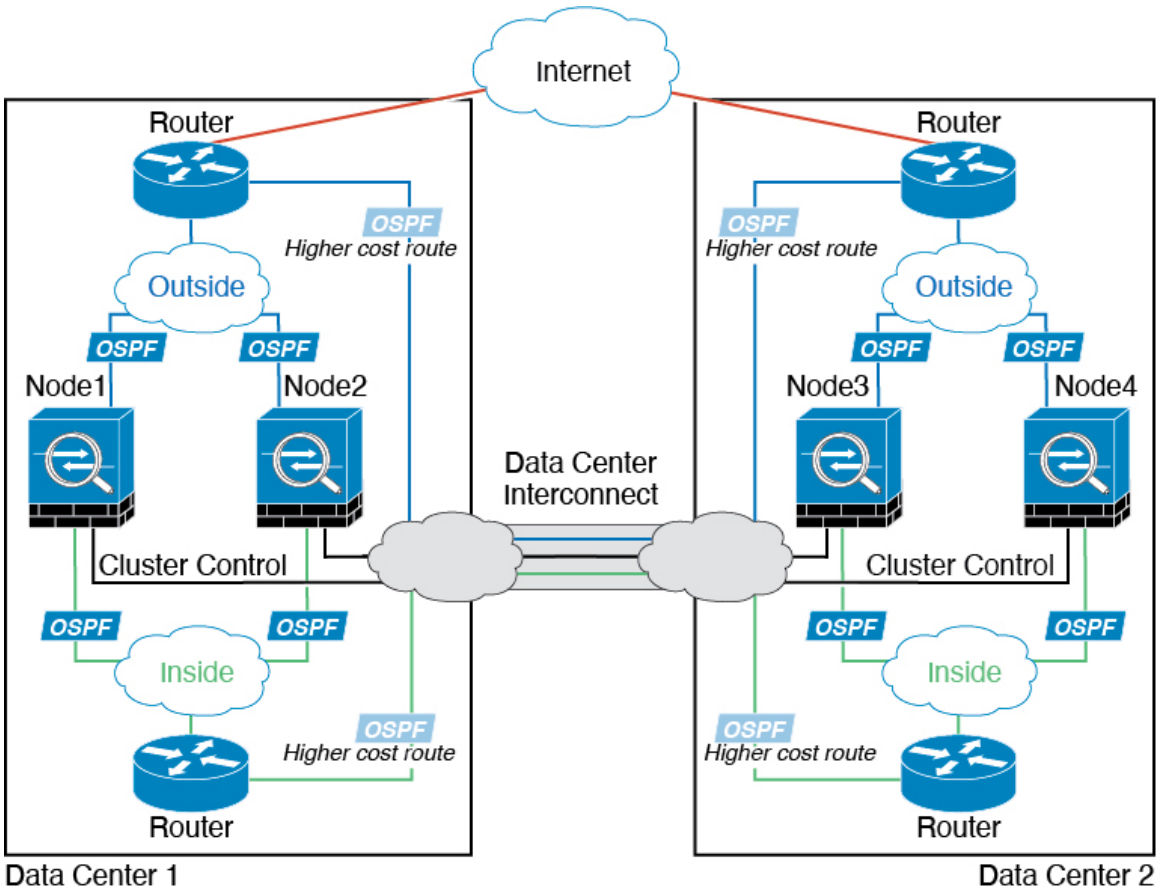
Overlay Interface Overlay1
VLAN MAC Address Layer-3 Address Age Expires In
3151 0050.5660.9412 10.0.0.2 1w0d 00:00:31
cluster-N7k6-OTV#
```

Examples for Inter-Site Clustering

The following examples show supported cluster deployments.

Individual Interface Routed Mode North-South Inter-Site Example

The following example shows 2 ASA cluster nodes at each of 2 data centers placed between inside and outside routers (North-South insertion). The cluster nodes are connected by the cluster control link over the DCI. The inside and outside routers at each data center use OSPF and PBR or ECMP to load balance the traffic between cluster members. By assigning a higher cost route across the DCI, traffic stays within each data center unless all ASA cluster nodes at a given site go down. In the event of a failure of all cluster nodes at one site, traffic goes from each router over the DCI to the ASA cluster nodes at the other site.



Spanned EtherChannel Routed Mode Example with Site-Specific MAC and IP Addresses

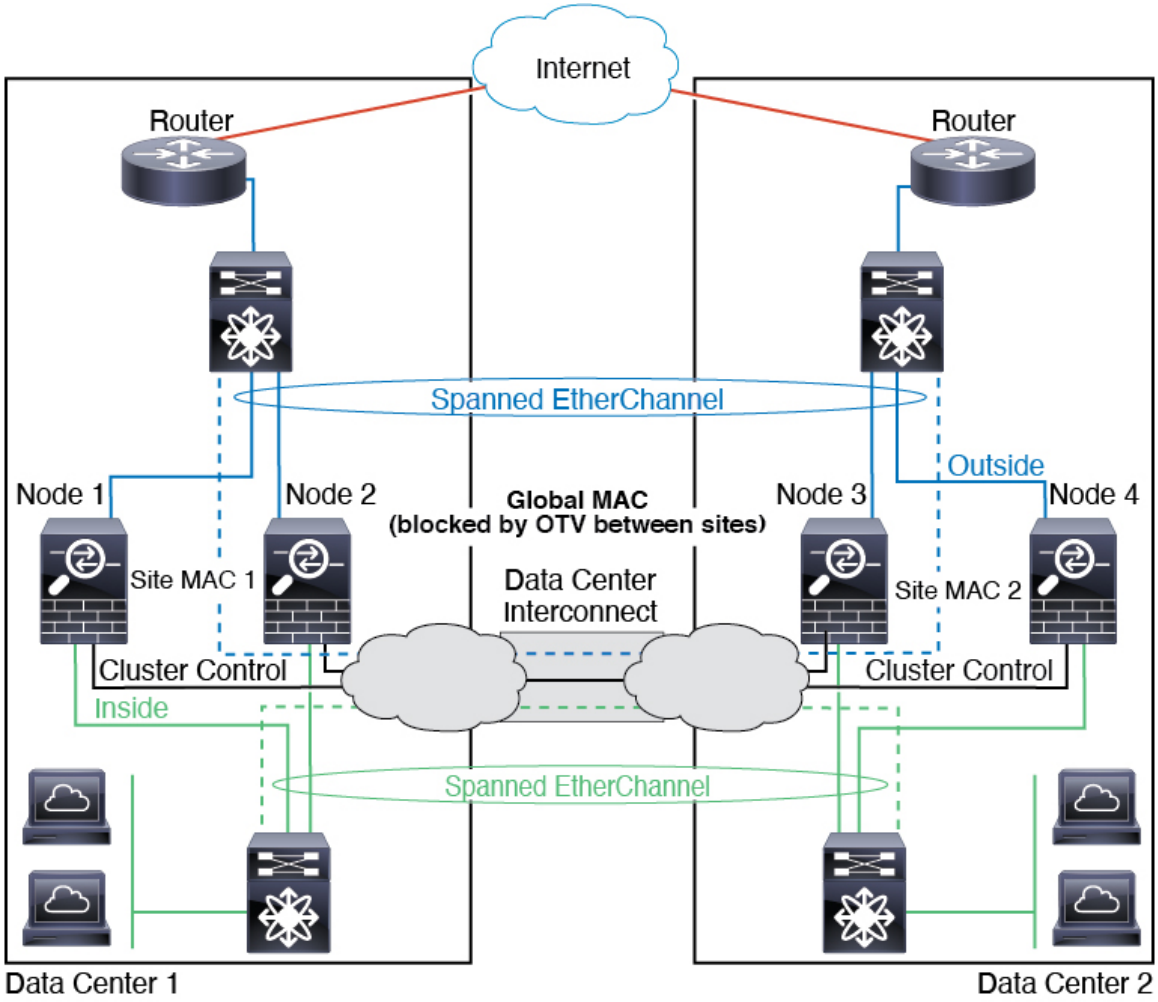
The following example shows 2 cluster members at each of 2 data centers placed between the gateway router and an inside network at each site (East-West insertion). The cluster members are connected by the cluster control link over the DCI. The cluster members at each site connect to the local switches using spanned EtherChannels for both the inside and outside networks. Each EtherChannel is spanned across all chassis in the cluster.

The data VLANs are extended between the sites using Overlay Transport Virtualization (OTV) (or something similar). You must add filters blocking the global MAC address to prevent traffic from traversing the DCI to the other site when the traffic is destined for the cluster. If the cluster nodes at one site become unreachable, you must remove the filters so traffic can be sent to the other site's cluster nodes. You should use VACLs to filter the global MAC address.For some switches, such as Nexus with the F3-series line card, you must also use ARP inspection to block ARP packets from the global MAC address. ARP inspection requires you to set both the site MAC address and the site IP address on the ASA. If you only configure the site MAC address be sure to disable ARP inspection.

The cluster acts as the gateway for the inside networks. The global virtual MAC, which is shared across all cluster nodes, is used only to receive packets. Outgoing packets use a site-specific MAC address from each DC cluster. This feature prevents the switches from learning the same global MAC address from both sites on two different ports, which causes MAC flapping; instead, they only learn the site MAC address.

In this scenario:

- All egress packets sent from the cluster use the site MAC address and are localized at the data center.
- All ingress packets to the cluster are sent using the global MAC address, so they can be received by any of the nodes at both sites; filters at the OTV localize the traffic within the data center.



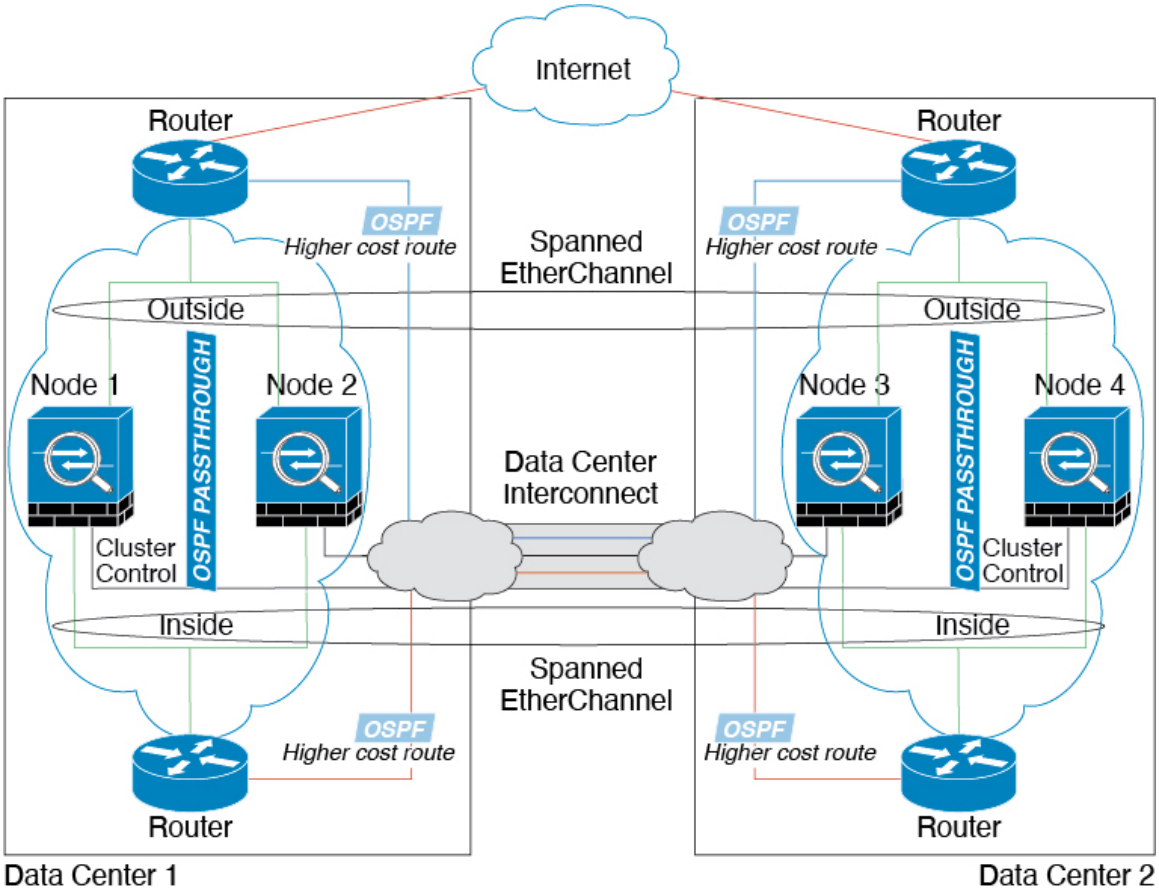
Spanned EtherChannel Transparent Mode North-South Inter-Site Example

The following example shows 2 cluster members at each of 2 data centers placed between inside and outside routers (North-South insertion). The cluster members are connected by the cluster control link over the DCI. The cluster members at each site connect to the local switches using spanned EtherChannels for the inside and outside. Each EtherChannel is spanned across all chassis in the cluster.

The inside and outside routers at each data center use OSPF, which is passed through the transparent ASAs. Unlike MACs, router IPs are unique on all routers. By assigning a higher cost route across the DCI, traffic stays within each data center unless all cluster members at a given site go down. The lower cost route through the ASAs must traverse the same bridge group at each site for the cluster to maintain asymmetric connections. In the event of a failure of all cluster members at one site, traffic goes from each router over the DCI to the cluster members at the other site.

The implementation of the switches at each site can include:

- Inter-site VSS, vPC, StackWise, or StackWise Virtual—In this scenario, you install one switch at Data Center 1, and the other at Data Center 2. One option is for the cluster nodes at each Data Center to only connect to the local switch, while the redundant switch traffic goes across the DCI. In this case, connections are for the most part kept local to each datacenter. You can optionally connect each node to both switches across the DCI if the DCI can handle the extra traffic. In this case, traffic is distributed across the data centers, so it is essential for the DCI to be very robust.
- Local VSS, vPC, StackWise, or StackWise Virtual at each site—For better switch redundancy, you can install 2 separate redundant switch pairs at each site. In this case, although the cluster nodes still have a spanned EtherChannel with Data Center 1 chassis connected only to both local switches, and Data Center 2 chassis connected to those local switches, the spanned EtherChannel is essentially "split." Each local redundant switch system sees the spanned EtherChannel as a site-local EtherChannel.

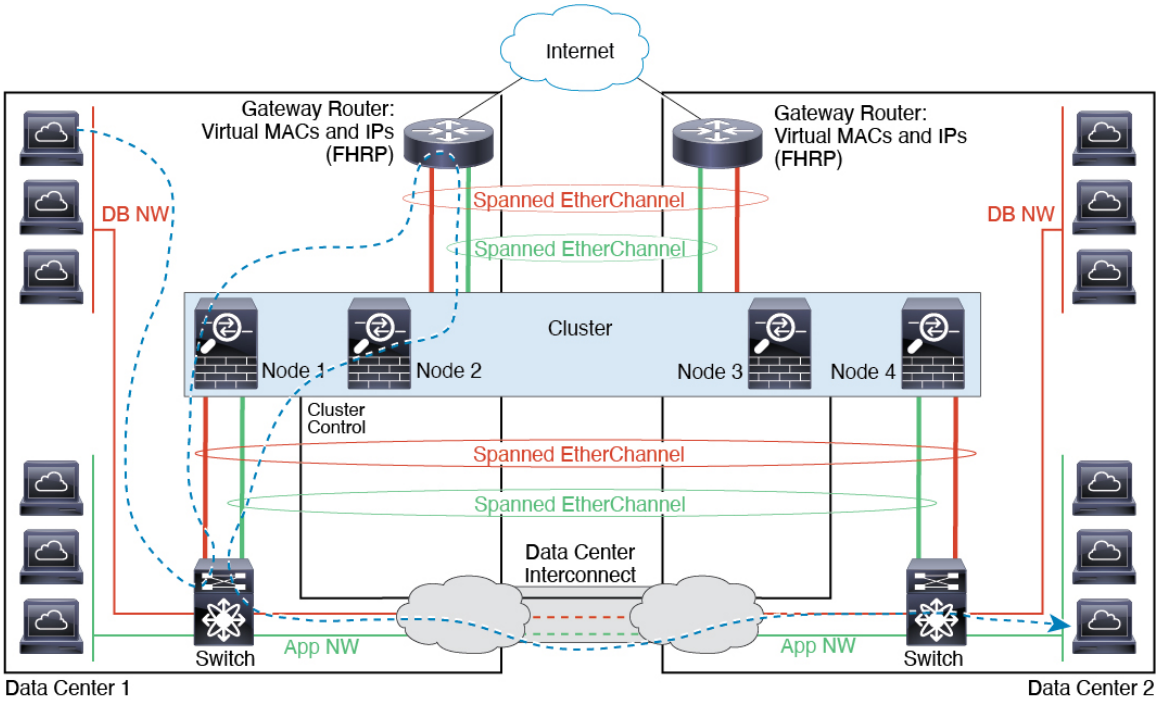


Spanned EtherChannel Transparent Mode East-West Inter-Site Example

The following example shows 2 cluster members at each of 2 data centers placed between the gateway router and two inside networks at each site, the App network and the DB network (East-West insertion). The cluster members are connected by the cluster control link over the DCI. The cluster members at each site connect to the local switches using spanned EtherChannels for both the App and DB networks on the inside and outside. Each EtherChannel is spanned across all chassis in the cluster.

The gateway router at each site uses an FHRP such as HSRP to provide the same destination virtual MAC and IP addresses at each site. A good practice to avoid unintended MAC address flapping is to statically add

the gateway routers real MAC addresses to the ASA MAC address table. Without these entries, if the gateway at site 1 communicates with the gateway at site 2, that traffic might pass through the ASA and attempt to reach site 2 from the inside interface and cause problems. The data VLANs are extended between the sites using Overlay Transport Virtualization (OTV) (or something similar). You must add filters to prevent traffic from traversing the DCI to the other site when the traffic is destined for the gateway router. If the gateway router at one site becomes unreachable, you must remove the filters so traffic can be sent to the other site's gateway router.



Reference for Clustering

This section includes more information about how clustering operates.

ASA Features and Clustering

Some ASA features are not supported with ASA clustering, and some are only supported on the control node. 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.

- Unified Communication features that rely on TLS Proxy
- Remote access VPN (SSL VPN and IPsec VPN)
- Virtual Tunnel Interfaces (VTIs)
- The following application inspections:

- CTIQBE
- H323, H225, and RAS
- IPsec passthrough
- MGCP
- MMP
- RTSP
- SCCP (Skinny)
- WAAS
- WCCP

- Botnet Traffic Filter
- Auto Update Server
- DHCP client, server, and proxy. DHCP relay is supported.
- VPN load balancing
- Failover on Azure
- Integrated Routing and Bridging
- FIPS 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.

- The following application inspections:
  - DCERPC
  - ESMTP
  - IM
  - NetBIOS
  - PPTP

  - RADIUS
  - RSH
  - SNMP
  - SQLNET
  - SUNRPC
  - TFTP
  - XDMCP
- Static route monitoring
- Authentication and Authorization for network access. Accounting is decentralized.
- Filtering Services
- 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)

Features Applied to Individual Nodes

These features are applied to each ASA node, instead of the cluster as a whole or to the control node.

- QoS—The QoS policy is synced across the cluster as part of configuration replication. However, the policy is enforced on each node independently. For example, if you configure policing on output, then the conform rate and conform burst values are enforced on traffic exiting a particular ASA. In a cluster with 3 nodes and with traffic evenly distributed, the conform rate actually becomes 3 times the rate for the cluster.
- Threat detection—Threat detection works on each node independently; for example, the top statistics is node-specific. Port scanning detection, for example, does not work because scanning traffic will be load-balanced between all nodes, and one node will not see all traffic.
- Resource management—Resource management in multiple context mode is enforced separately on each node based on local usage.
- LISP traffic—LISP traffic on UDP port 4342 is inspected by each receiving node, but is not assigned a director. Each node adds to the EID table that is shared across the cluster, but the LISP traffic itself does not participate in cluster state sharing.

AAA for Network Access and Clustering

AAA for network access consists of three components: authentication, authorization, and accounting. Authentication and authorization are implemented as centralized features on the clustering control node with replication of the data structures to the cluster data nodes. If a control node is elected, the new control node will have all the information it needs to continue uninterrupted operation of the established authenticated users

and their associated authorizations. Idle and absolute timeouts for user authentications are preserved when a control node change occurs.

Accounting is implemented as a distributed feature in a cluster. Accounting is done on a per-flow basis, so the cluster node owning a flow will send accounting start and stop messages to the AAA server when accounting is configured for a flow.

Connection Settings and Clustering

Connection limits are enforced cluster-wide (see **Configuration** > **Firewall** > **Service Policy** page). 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.
- If you use AAA for FTP access, then the control channel flow is centralized on the control node.

ICMP Inspection and Clustering

The flow of ICMP and ICMP error packets through the cluster varies depending on whether ICMP/ICMP error inspection is enabled. Without ICMP inspection, ICMP is a one-direction flow, and there is no director flow support. With ICMP inspection, the ICMP flow becomes two-directional and is backed up by a director/backup flow. One difference for an inspected ICMP flow is in the director handling of a forwarded packet: the director will forward the ICMP echo reply packet to the flow owner instead of returning the packet to the forwarder.

Multicast Routing and Clustering

Multicast routing behaves differently depending on the interface mode.

Multicast Routing in Spanned EtherChannel Mode

In Spanned EtherChannel mode: 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.

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 ASAs 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 ASA 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. 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.
- Per-session PAT feature—Although not exclusive to clustering, the per-session PAT feature improves the scalability of PAT and, for clustering, allows each data node to own PAT connections; by contrast, multi-session PAT connections have to be forwarded to and owned by the control node. By default, all TCP traffic and UDP DNS traffic use a per-session PAT xlate, whereas ICMP and all other UDP traffic uses multi-session. You can configure per-session NAT rules to change these defaults for TCP and UDP, but you cannot configure per-session PAT for ICMP. For example, with an increasing use of Quic protocol over UDP/443 as a best performance alternative compared to HTTPS TLS over TCP/443, you should enable per-session PAT for UDP/443. For traffic that benefits from multi-session PAT, such as H.323, SIP, or Skinny, you can disable per-session PAT for the associated TCP ports (the UDP ports for those H.323 and SIP are already multi-session by default). For more information about per-session PAT, see the firewall configuration guide.
- No static PAT for the following inspections—
  - FTP
  - PPTP
  - 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 and Clustering

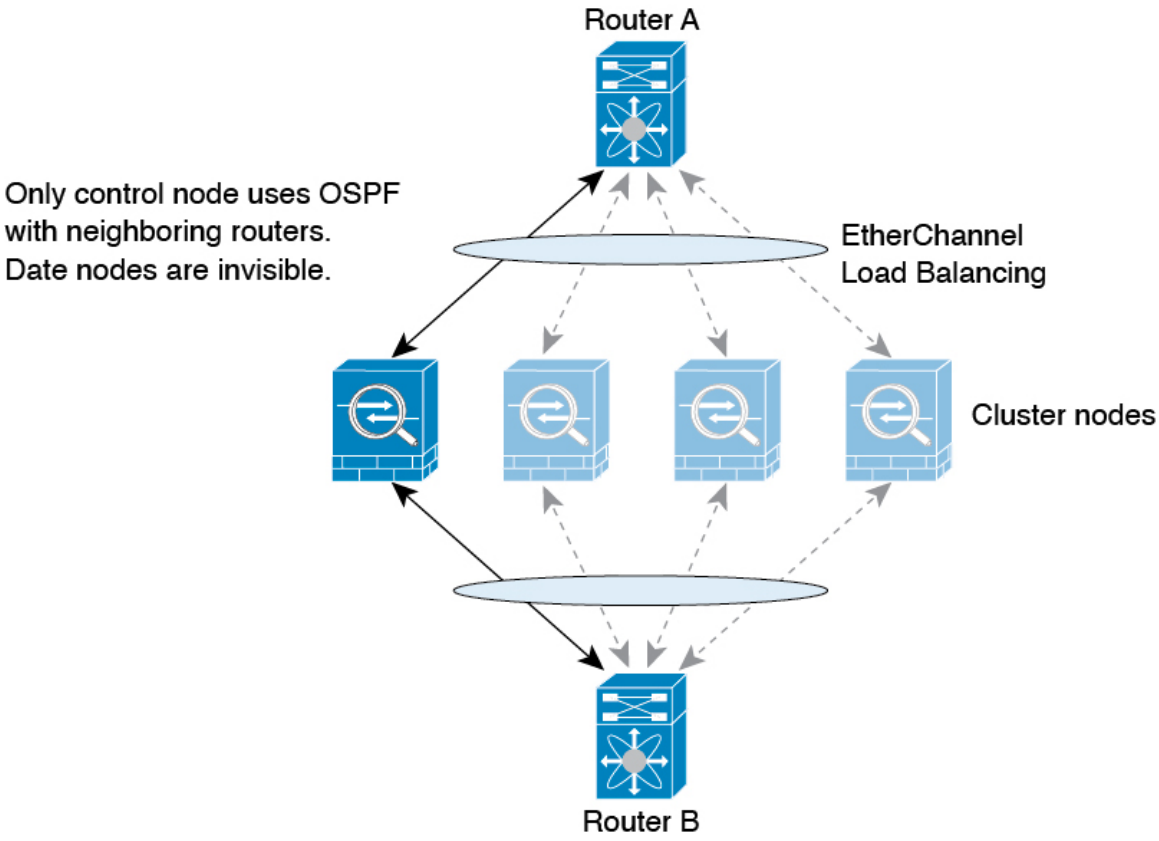
This section describes how to use dynamic routing with clustering.

Dynamic Routing in Spanned EtherChannel Mode

Note IS-IS is not supported in Spanned EtherChannel mode.

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 64: Dynamic Routing in Spanned EtherChannel Mode



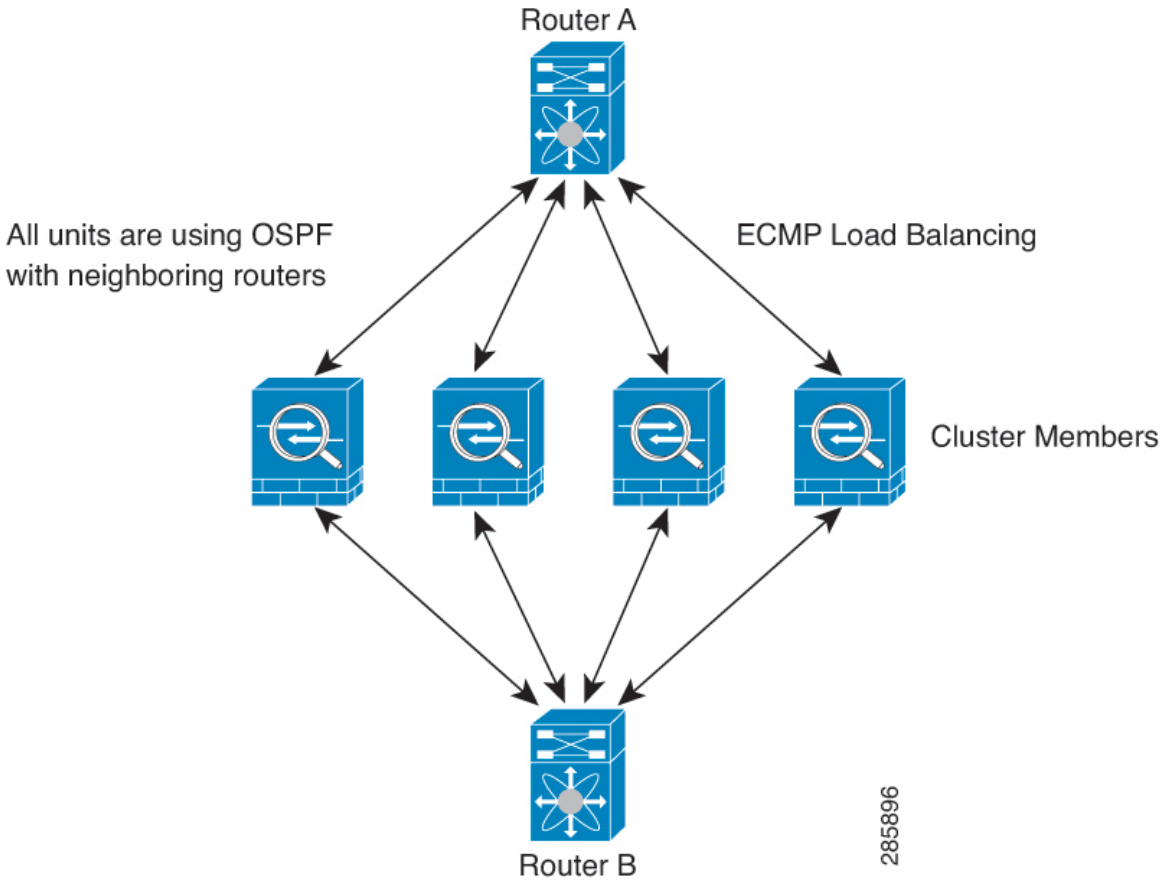
After the data node learn the routes from the control node, each node makes forwarding decisions independently.

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 65: 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 Configure a Traffic Zone, on page 684.

SCTP and Clustering

An SCTP association can be created on any node (due to load balancing); its multi-homing connections must reside on the same node.

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.

TLS Proxy configuration is not supported.

SNMP and Clustering

An SNMP agent polls each individual ASA by its Local IP address. You cannot poll consolidated data for the cluster.

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 re-add them on the control node to force the users to replicate to the new node, or directly on the data node.

STUN and Clustering

STUN inspection is supported in failover and cluster modes, as pinholes are replicated. However, the transaction ID is not replicated among nodes. In the case where a node fails after receiving a STUN Request and another node received the STUN Response, the STUN Response will be dropped.

Syslog and NetFlow and Clustering

- Syslog—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.
- NetFlow—Each node in the cluster generates its own NetFlow stream. The NetFlow collector can only treat each ASA as a separate NetFlow exporter.

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 420 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 (depending on your minimum port bundling setting).

A node is removed from the cluster if its monitored interfaces fail. The amount of time before the ASA 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 ASA removes the member after 9 seconds. The ASA 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 ASA 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 ASA automatically tries to rejoin the cluster, depending on the failure event.

Note When the ASA becomes inactive and fails to automatically rejoin the cluster, all data interfaces are shut down; only the management-only interface can send and receive traffic. The management interface remains up using the IP address the node received from the cluster IP pool. However if you reload, and the node is still inactive in the cluster, the management interface is disabled. You must use the console port for any further configuration.

Rejoining the Cluster

After a cluster node 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 ASA automatically tries to rejoin every 5 minutes, indefinitely. This behavior is configurable.
- Failed data interface—The ASA 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 ASA disables clustering. After you resolve the problem with the data interface, you have to manually enable clustering. This behavior is configurable.
- 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 and clustering is still enabled. The ASA attempts to rejoin the cluster every 5 seconds.
- Internal error—Internal failures include: application sync timeout; inconsistent application statuses; and so on. A node will attempt to rejoin the cluster automatically at the following intervals: 5 minutes, 10 minutes, and then 20 minutes. This behavior is configurable.

See Configure Basic ASA Cluster Parameters, on page 372.

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 22: 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	Includes AAA rules (uauth).

Traffic	State Support	Notes
IPv6 Neighbor database	Yes	—
Dynamic routing	Yes	—
SNMP Engine ID	**No**	—
Distributed VPN (Site-to-Site) for Firepower 4100/9300	Yes	Backup session becomes the active session, then a new backup session is created.

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.

  If you enable director localization for inter-site clustering, then there are two backup owner roles: the local backup and the global backup. The owner always chooses a local backup at the same site as itself (based on site ID). The global backup can be at any site, and might even be the same node as the local backup. The owner sends connection state information to both backups.

  If you enable site redundancy, and the backup owner is at the same site as the owner, then an additional backup owner will be chosen from another site to protect flows from a site failure. Chassis backup and site backup are independent, so in some cases a flow will have both a chassis backup and a site backup.
- 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.

If you enable director localization for inter-site clustering, then there are two director roles: the local director and the global director. The owner always chooses a local director at the same site as itself (based on site ID). The global director can be at any site, and might even be the same node as the local director. If the original owner fails, then the local director chooses a new connection owner at the same site.

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. If you enable director localization, then the forwarder always queries the local director. The forwarder only queries the global director if the local director does not know the owner, for example, if a cluster member receives packets for a connection that is owned on a different site. 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.

When a connection uses Port Address Translation (PAT), then the PAT type (per-session or multi-session) influences which member of the cluster becomes the owner of a new connection:

- Per-session PAT—The owner is the node that receives the initial packet in the connection.

  By default, TCP and DNS UDP traffic use per-session PAT.

- Multi-session PAT—The owner is always the control node. If a multi-session PAT connection is initially received by a data node, then the data node forwards the connection to the control node.

  By default, UDP (except for DNS UDP) and ICMP traffic use multi-session PAT, so these connections are always owned by the control node.

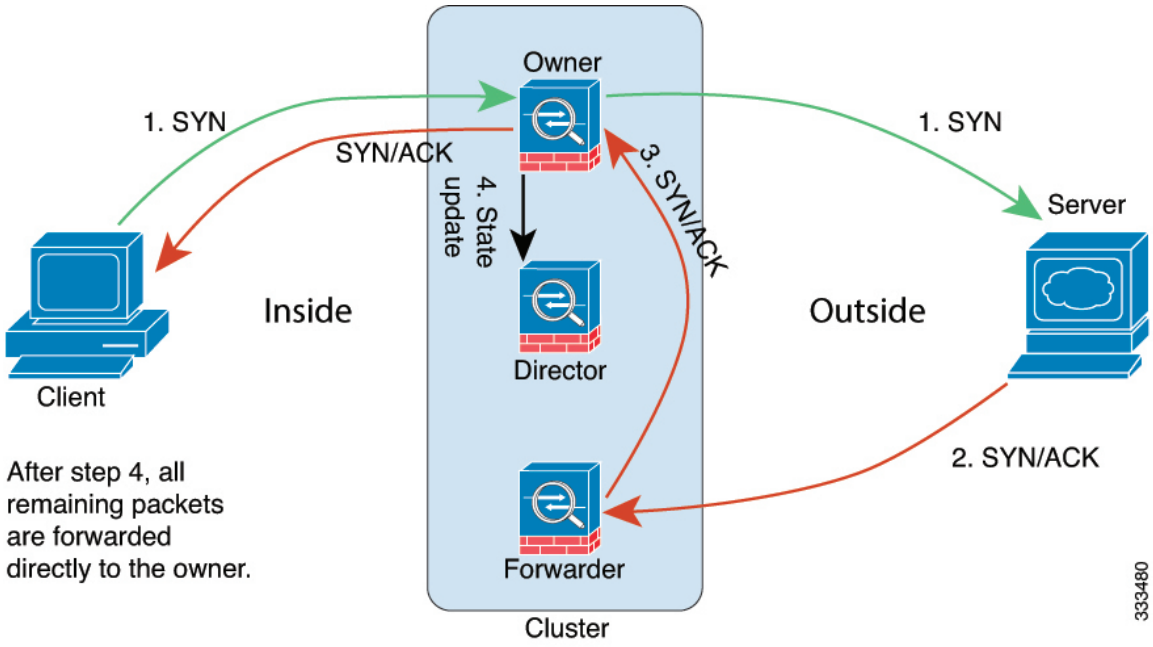
You can change the per-session PAT defaults for TCP and UDP so connections for these protocols are handled per-session or multi-session depending on the configuration. For ICMP, you cannot change from the default multi-session PAT. For more information about per-session PAT, see the firewall configuration guide.

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. For best performance, proper external load balancing is required for both directions of a flow to arrive at the same node, and for flows to be distributed evenly between nodes. 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.



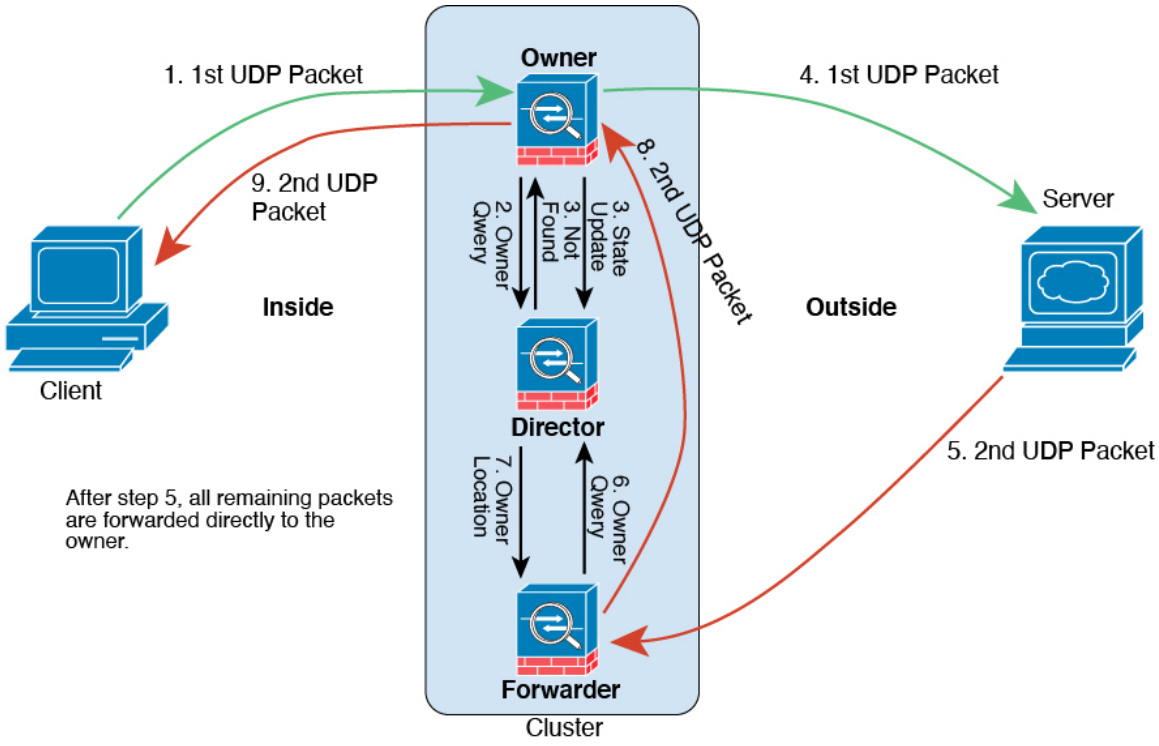
1. The SYN packet originates from the client and is delivered to one ASA (based on the load balancing method), which becomes the owner. The owner creates a flow, encodes owner information into a SYN cookie, and forwards the packet to the server.
2. The SYN-ACK packet originates from the server and is delivered to a different ASA (based on the load balancing method). This ASA 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 66: ICMP and UDP Data Flow***



The first UDP packet originates from the client and is delivered to one ASA (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.

Rebalancing New TCP Connections Across the Cluster

If the load balancing capabilities of the upstream or downstream routers result in unbalanced flow distribution, you can configure new connection rebalancing so nodes with higher new connections per second will redirect new TCP flows to other nodes. No existing flows will be moved to other nodes.

Because this command only rebalances based on connections per second, the total number of established connections on each node is not considered, and the total number of connections may not be equal.

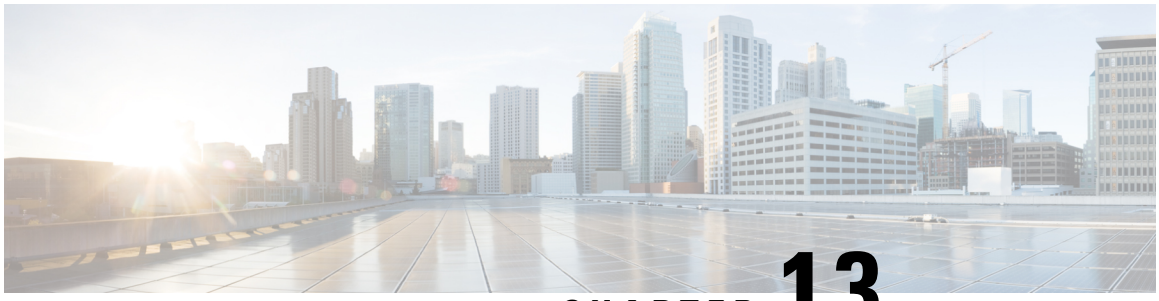
Once a connection is offloaded to a different node, it becomes an asymmetric connection.

Do not configure connection rebalancing for inter-site topologies; you do not want new connections rebalanced to cluster members at a different site.

History for ASA Clustering for the Secure Firewall 3100/4200

Feature Name	Version	Feature Information
Distributed site-to-site VPN with clustering on the Secure Firewall 4200	9.23(1)	An ASA cluster on the Secure Firewall 4200 supports site-to-site VPN in distributed mode. Distributed mode provides the ability to have many site-to-site IPsec IKEv2 VPN connections distributed across members of an ASA cluster, not just on the control node (as in centralized mode). This significantly scales VPN support beyond centralized VPN capabilities and provides high availability. New or modified screens: **Monitoring > ASA Cluster > ASA Cluster > VPN Cluster Summary** **Monitoring > VPN > VPN Statistics > Sessions** **Configuration > Device Management > High Availablility and Scalability > ASA Cluster** **Wizards > Site-to-Site** **Monitoring > VPN > VPN Statistics > Sessions** **Monitoring > ASA Cluster > ASA Cluster > VPN Cluster Summary** **Monitoring > ASA Cluster > ASA Cluster > System Resource Graphs > CPU/Memory** **Monitoring > Logging > Real-Time Log Viewer**
Cluster redirect: flow offload support for the Secure Firewall 4200 asymmetric cluster traffic	9.23(1)	For asymmetric flows, cluster redirect lets the forwarding node offload flows to hardware. This feature is enabled by default. When traffic for an existing flow is sent to a different node, then that traffic is redirected to the owner node over the cluster control link. Because asymmetric flows can create a lot of traffic on the cluster control link, letting the forwarder offload these flows can improve performance. Added/modified screens: **Configuration** > **Firewall** > **Advanced** > **Offload Engine** > **Cluster Redirect Offload**

Feature Name	Version	Feature Information
IPsec flow offload for traffic on the cluster control link on the Secure Firewall 4200 in distributed site-to-site VPN mode	9.23(1)	For asymmetric flows in distributed site-to-site VPN mode, IPsec flow offload now lets the flow owner decrypt IPsec traffic in hardware that was forwarded over the cluster control link. This feature is not configurable and is always available when you enable IPsec flow offload. Added/modified screens: **Configuration** > **Firewall** > **Advanced** > **IPsec Offload**
Maximum cluster nodes increased to 16	9.22(1)	The maximum nodes were increased from 8 to 16.
Individual interface mode	9.22(1)	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. New/Modified commands: **cluster interface-mode individual** New/Modified screens: **Wizards >** > **High Availability and Scalability Wizard**
Configurable cluster keepalive interval for flow status	9.20(1)	The flow owner sends keepalives (clu_keepalive messages) and updates (clu_update messages) to the director and backup owner to refresh the flow state. You can now set the keepalive interval. The default is 15 seconds, and you can set the interval between 15 and 55 seconds. You may want to set the interval to be longer to reduce the amount of traffic on the cluster control link. New/Modified screens: **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** > **Cluster Configuration**
Support for clustering on the Secure Firewall 4200 was introduced	9.20(1)	You can cluster up to 8 Secure Firewall 4200 nodes in Spanned EtherChannel mode.
Removal of biased language	9.19(1)	Commands, command output, and syslog messages that contained the terms "Master" and "Slave" have been changed to "Control" and "Data." New/Modified commands: **cluster control-node**, **enable as-data-node**, **prompt**, **show cluster history**, **show cluster info**
Support for clustering on the Secure Firewall 3100 was introduced	9.17(1)	You can cluster up to 8 Secure Firewall 3100 nodes in Spanned EtherChannel mode.



CHAPTER 13

ASA Cluster for the Firepower 4100/9300

Clustering lets you group multiple Firepower 4100/9300 chassis ASAs together as a single logical device. The Firepower 4100/9300 chassis series includes the Firepower 9300 and Firepower 4100 series. 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 485.

- About Clustering on the Firepower 4100/9300 Chassis, on page 429
- Requirements and Prerequisites for Clustering on the Firepower 4100/9300 Chassis, on page 436
- Licenses for Clustering on the Firepower 4100/9300 Chassis, on page 437
- Clustering Guidelines and Limitations, on page 439
- Configure Clustering on the Firepower 4100/9300 Chassis, on page 443
- FXOS: Remove a Cluster Node, on page 468
- ASA: Manage Cluster Members, on page 469
- ASA: Monitoring the ASA Cluster on the Firepower 4100/9300 chassis, on page 474
- Troubleshooting Distributed S2S VPN, on page 476
- Examples for ASA Clustering, on page 477
- Reference for Clustering, on page 485
- History for ASA Clustering on the Firepower 4100/9300, on page 500

About Clustering on the Firepower 4100/9300 Chassis

When you deploy a cluster on the Firepower 4100/9300 chassis, it does the following:

- Creates a *cluster-control link* (by default, port-channel 48) for node-to-node communication.

  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. Some parts of the bootstrap configuration may be user-configurable within the application if you want to customize your clustering environment.

- 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. Some parts of the bootstrap configuration are user-configurable if you want to customize your clustering environment.

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. See Centralized Features for Clustering, on page 486.

Cluster Control Link

The cluster-control link is an EtherChannel (port-channel 48) for unit-to-unit communication. For intra-chassis clustering, this link utilizes the Firepower 9300 backplane for cluster communications. For inter-chassis clustering, you need to manually assign physical interface(s) to this EtherChannel on the Firepower 4100/9300 chassis for communications between chassis.

For a 2-chassis inter-chassis cluster, 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.

Control traffic includes:

- Control node election.
- Configuration replication.
- Health monitoring.

Data traffic includes:

- State replication.
- Connection ownership queries and data packet forwarding.

See the following sections for more information about the cluster control link.

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.
- AAA for network access is a centralized feature, so all traffic is forwarded to the control unit.
- 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.

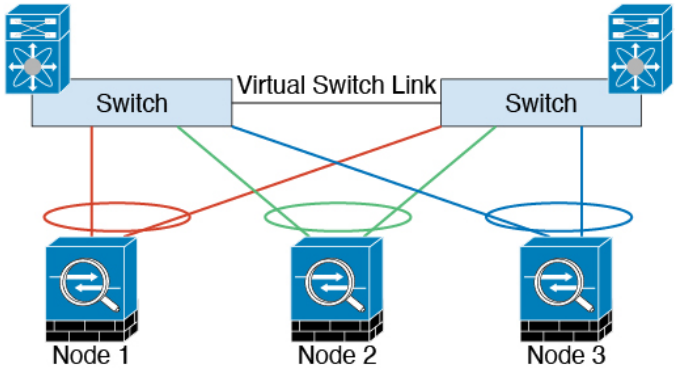


Note If your cluster has large amounts of asymmetric (rebalanced) traffic, then you should increase the cluster control link size.

Cluster Control Link Redundancy

We recommend using an EtherChannel for the cluster control link, so that you can pass traffic on multiple links in the EtherChannel while still achieving 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.

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*.You can customize this IP address when you deploy the cluster. The cluster control link network cannot include any routers between units; only Layer 2 switching is allowed. For inter-site traffic, Cisco recommends using Overlay Transport Virtualization (OTV).

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.

Individual interfaces are not supported, with the exception of a management interface.

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.

Secure Firewall ASA Cluster Management

One of the benefits of using ASA clustering is the ease of management. This section describes how to manage the cluster.

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.

The Main cluster IP address is a fixed address for the cluster that always belongs to the current control unit. You also configure a range of addresses so that each unit, including the current control unit, can use a Local address from the range. The Main cluster IP address provides consistent management access to an address; when a control unit changes, the Main cluster IP address moves to the new control unit, so management of the cluster continues seamlessly.

For example, you can manage the cluster by connecting to the Main cluster IP address, which is always attached to the current control unit. To manage an individual member, you can connect to the Local IP address.



Note To-the-box traffic needs to be directed to the node's management IP address; to-the-box traffic is not forwarded over the cluster control link to any other node.

For outbound management traffic such as TFTP or syslog, each unit, including the control unit, uses the Local IP address to connect to the server.

Control Unit Management Vs. Data Unit Management

All management and monitoring can take place on the control node. From the control node, you can check runtime statistics, resource usage, or other monitoring information of all nodes. You can also issue a command to all nodes in the cluster, and replicate the console messages from data nodes to the control node.

You can monitor data nodes directly if desired. Although also available from the control node, you can perform file management on data nodes (including backing up the configuration and updating images). The following functions are not available from the control node:

- Monitoring per-node cluster-specific statistics.
- Syslog monitoring per node (except for syslogs sent to the console when console replication is enabled).
- SNMP

- NetFlow

Crypto Key Replication

When you create a crypto key on the control node, the key is replicated to all data nodes. If you have an SSH session to the Main cluster IP address, you will be disconnected if the control node fails. The new control node uses the same key for SSH connections, so that you do not need to update the cached SSH host key when you reconnect to the new control node.

ASDM Connection Certificate IP Address Mismatch

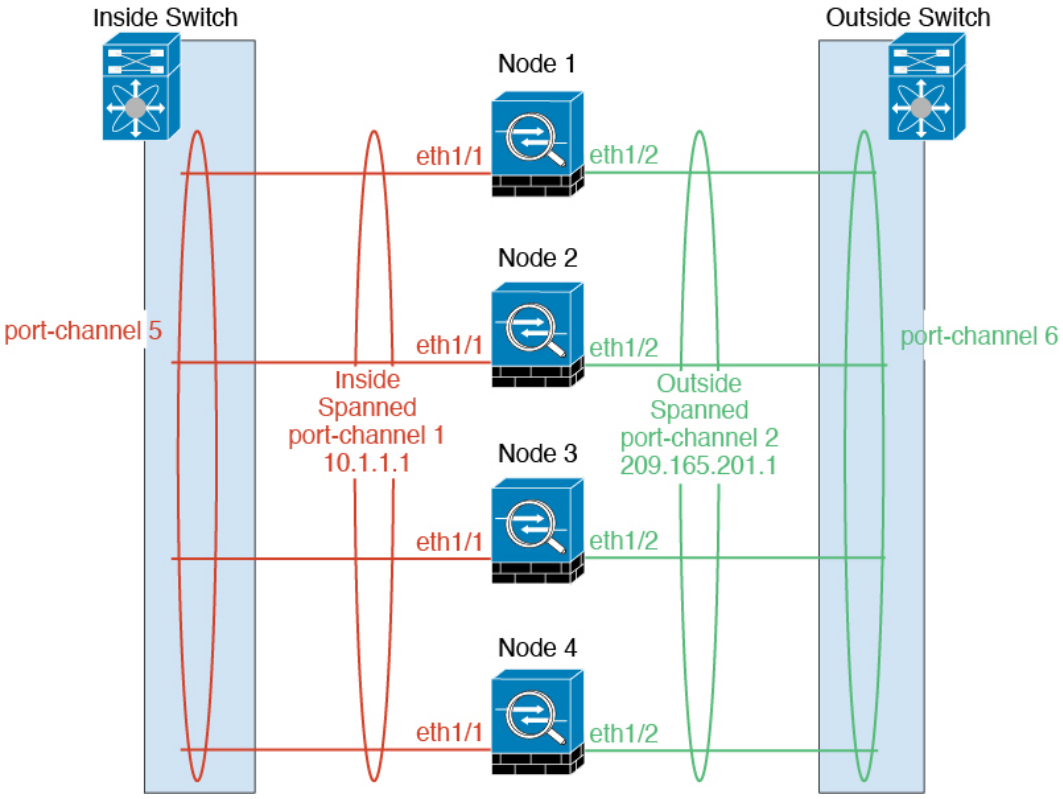
By default, a self-signed certificate is used for the ASDM connection based on the Local IP address. If you connect to the Main cluster IP address using ASDM, then a warning message about a mismatched IP address might appear because the certificate uses the Local IP address, and not the Main cluster IP address. You can ignore the message and establish the ASDM connection. However, to avoid this type of warning, you can enroll a certificate that contains the Main cluster IP address and all the Local IP addresses from the IP address pool. You can then use this certificate for each cluster member. See https://www.cisco.com/c/en/us/td/docs/security/asdm/identity-cert/cert-install.html for more information.

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.

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.



Inter-Site Clustering

For inter-site installations, you can take advantage of ASA clustering as long as you follow the recommended guidelines.

You can configure each cluster chassis to belong to a separate site ID.

Site IDs work with site-specific MAC addresses and IP addresses. Packets egressing the cluster use a site-specific MAC address and IP address, while packets received by the cluster use a global MAC address and IP address. This feature prevents the switches from learning the same global MAC address from both sites on two different ports, which causes MAC flapping; instead, they only learn the site MAC address. Site-specific MAC addresses and IP address are supported for routed mode using Spanned EtherChannels only.

Site IDs are also used to enable flow mobility using LISP inspection, director localization to improve performance and reduce round-trip time latency for inter-site clustering for data centers, and site redundancy for connections where a backup owner of a traffic flow is always at a different site from the owner.

See the following sections for more information about inter-site clustering:

- Sizing the Data Center Interconnect—Requirements and Prerequisites for Clustering on the Firepower 4100/9300 Chassis, on page 436
- Inter-Site Guidelines—Clustering Guidelines and Limitations, on page 439
- Configure Cluster Flow Mobility—Configure Cluster Flow Mobility, on page 459
- Enable Director Localization—Configure Basic ASA Cluster Parameters, on page 454

- Enable Site Redundancy—Configure Basic ASA Cluster Parameters, on page 454

Requirements and Prerequisites for Clustering on the Firepower 4100/9300 Chassis

Maximum Clustering Units Per Model

- Firepower 4100—16 chassis
- Firepower 9300—16 modules. 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.

Hardware and Software Requirements for Inter-Chassis Clustering

All chassis in a cluster:

- 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.
- 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. Note that if you remove an interface in FXOS, the ASA configuration retains the related commands so that you can make any necessary adjustments; removing an interface from the configuration can have wide effects. You can manually remove the old interface configuration.
- Must use the same NTP server. Do not set the time manually.
- ASA: Each FXOS chassis must be registered with the License Authority or satellite server. There is no extra cost for data nodes. For permanent license reservation, you must purchase separate licenses for each chassis. For threat defense, all licensing is handled by the management center.

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.

Sizing the Data Center Interconnect for Inter-Site Clustering

You should reserve bandwidth on the data center interconnect (DCI) for cluster control link traffic equivalent to the following calculation:

$$\frac{\text{\# of cluster members per site}}{2} \times \text{cluster control link size per member}$$

If the number of members differs at each site, use the larger number for your calculation. The minimum bandwidth for the DCI should not be less than the size of the cluster control link for one member.

For example:

- For 4 members at 2 sites:
  - 4 cluster members total
  - 2 members at each site
  - 5 Gbps cluster control link per member

  Reserved DCI bandwidth = 5 Gbps (2/2 x 5 Gbps).
- For 6 members at 3 sites, the size increases:
  - 6 cluster members total
  - 3 members at site 1, 2 members at site 2, and 1 member at site 3
  - 10 Gbps cluster control link per member

  Reserved DCI bandwidth = 15 Gbps (3/2 x 10 Gbps).
- For 2 members at 2 sites:
  - 2 cluster members total
  - 1 member at each site
  - 10 Gbps cluster control link per member

  Reserved DCI bandwidth = 10 Gbps (1/2 x 10 Gbps = 5 Gbps; but the minimum bandwidth should not be less than the size of the cluster control link (10 Gbps)).

Licenses for Clustering on the Firepower 4100/9300 Chassis

Smart Software Manager Regular and On-Prem

The clustering feature itself does not require any licenses. To use Strong Encryption and other optional licenses, each Firepower 4100/9300 chassis must be registered with the License Authority or Smart Software Manager Regular and On-Prem server. There is no extra cost for data units.

The Strong Encryption license is automatically enabled for qualified customers when you apply the registration token. When using the token, each chassis must have the same encryption license. For the optional Strong Encryption (3DES/AES) feature license enabled in the ASA configuration, see below.

In the ASA license configuration, you can only configure smart licensing on the control unit. The configuration is replicated to the data units, but for some licenses, they do not use the configuration; it remains in a cached state, and only the control unit requests the license. The licenses are aggregated into a single cluster license that is shared by the cluster units, and this aggregated license is also cached on the data units to be used if one of them becomes the control unit in the future. Each license type is managed as follows:

- Essentials—Only the control unit requests the Essentials license from the server, and both units can use it due to license aggregation.
- Context—Only the control unit requests the Context license from the server. The Essentials license includes 10 contexts by default and is present on all cluster members. The value from each unit's Essentials license plus the value of the Context license on the control unit are combined up to the platform limit in an aggregated cluster license. For example:
  - You have 6 Firepower 9300 modules in the cluster. The Essentials license includes 10 contexts; for 6 units, these licenses add up to 60 contexts. You configure an additional 20-Context license on the control unit. Therefore, the aggregated cluster license includes 80 contexts. Because the platform limit for one module is 250, the combined license allows a maximum of 250 contexts; the 80 contexts are within the limit. Therefore, you can configure up to 80 contexts on the control unit; each data unit will also have 80 contexts through configuration replication.
  - You have 3 Firepower 4112 units in the cluster. The Essentials license includes 10 contexts; for 3 units, these licenses add up to 30 contexts. You configure an additional 250-Context license on the control unit. Therefore, the aggregated cluster license includes 280 contexts. Because the platform limit for one unit is 250, the combined license allows a maximum of 250 contexts; the 280 contexts are over the limit. Therefore, you can only configure up to 250 contexts on the control unit; each data unit will also have 250 contexts through configuration replication. In this case, you should only configure the control unit Context license to be 220 contexts.
- Carrier—Required for Distributed S2S VPN. This license is a per-unit entitlement, and each unit requests its own license from the server.
- Strong Encryption (3DES)—For pre-2.3.0 Cisco Smart Software Manager On-Prem deployment; or if your Smart Account is not authorized for strong encryption, but Cisco has determined that you are allowed to use strong encryption, you can manually add a strong encryption license to your account. This license is a per-unit entitlement, and each unit requests its own license from the server.

If a new control unit is elected, the new control unit continues to use the aggregated license. It also uses the cached license configuration to re-request the control unit license. When the old control unit rejoins the cluster as a data unit, it releases the control unit license entitlement. Before the data unit releases the license, the control unit's license might be in a non-compliant state if there are no available licenses in the account. The retained license is valid for 30 days, but if it is still non-compliant after the grace period, you will not be able to make configuration changes to features requiring special licenses; operation is otherwise unaffected. The new active unit sends an entitlement authorization renewal request every 12 hours until the license is compliant. You should refrain from making configuration changes until the license requests are completely processed. If a unit leaves the cluster, the cached control configuration is removed, while the per-unit entitlements are retained. In particular, you would need to re-request the Context license on non-cluster units.

Permanent License Reservation

For permanent license reservation, you must purchase separate licenses for each chassis and enable the licenses *before* you configure clustering.

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.
- 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. *Do not* change the load-balancing algorithm from the default on the cluster device.
- 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:

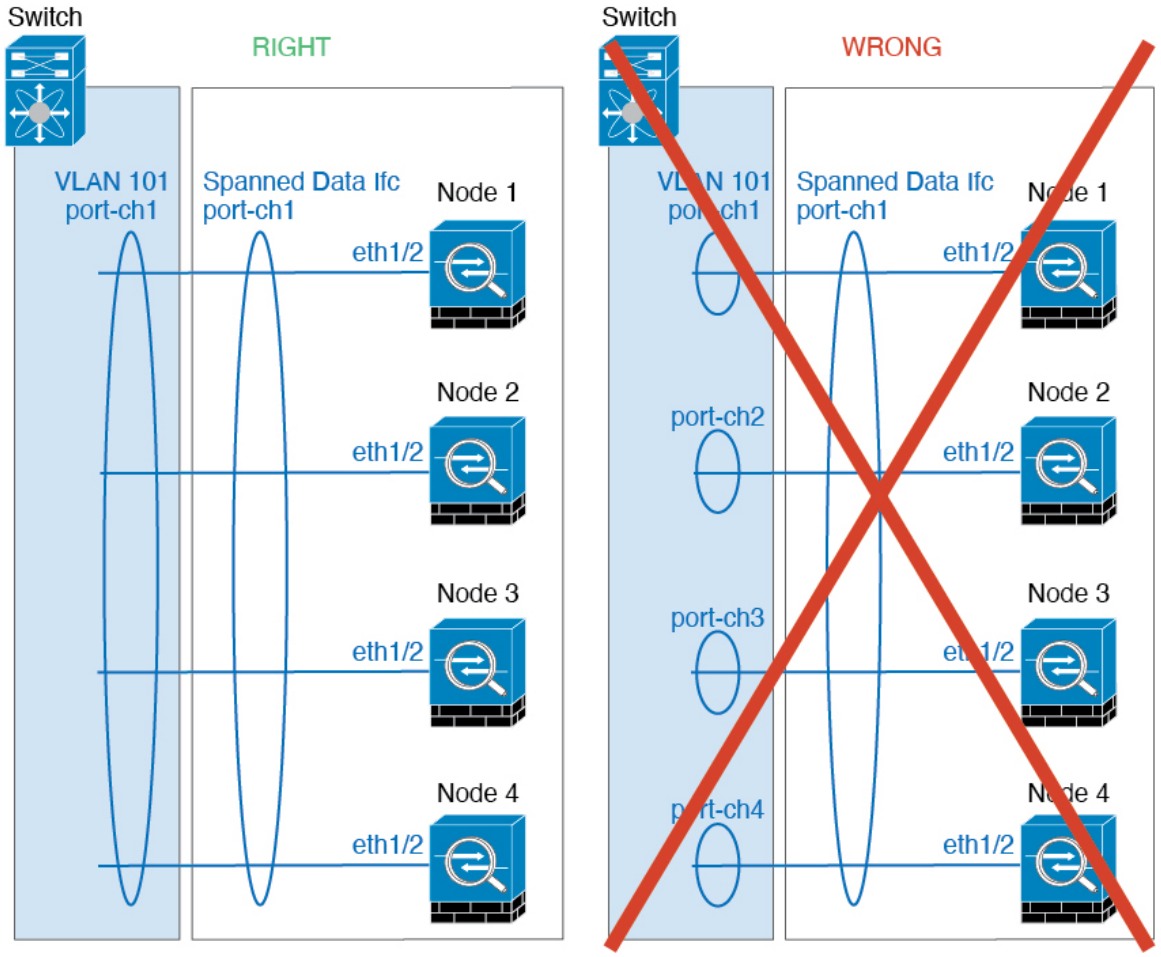
  router(config)# **port-channel** *id* **hash-distribution fixed**

  Do not change the algorithm globally; you may want to take advantage of the adaptive algorithm for the VSS peer link.

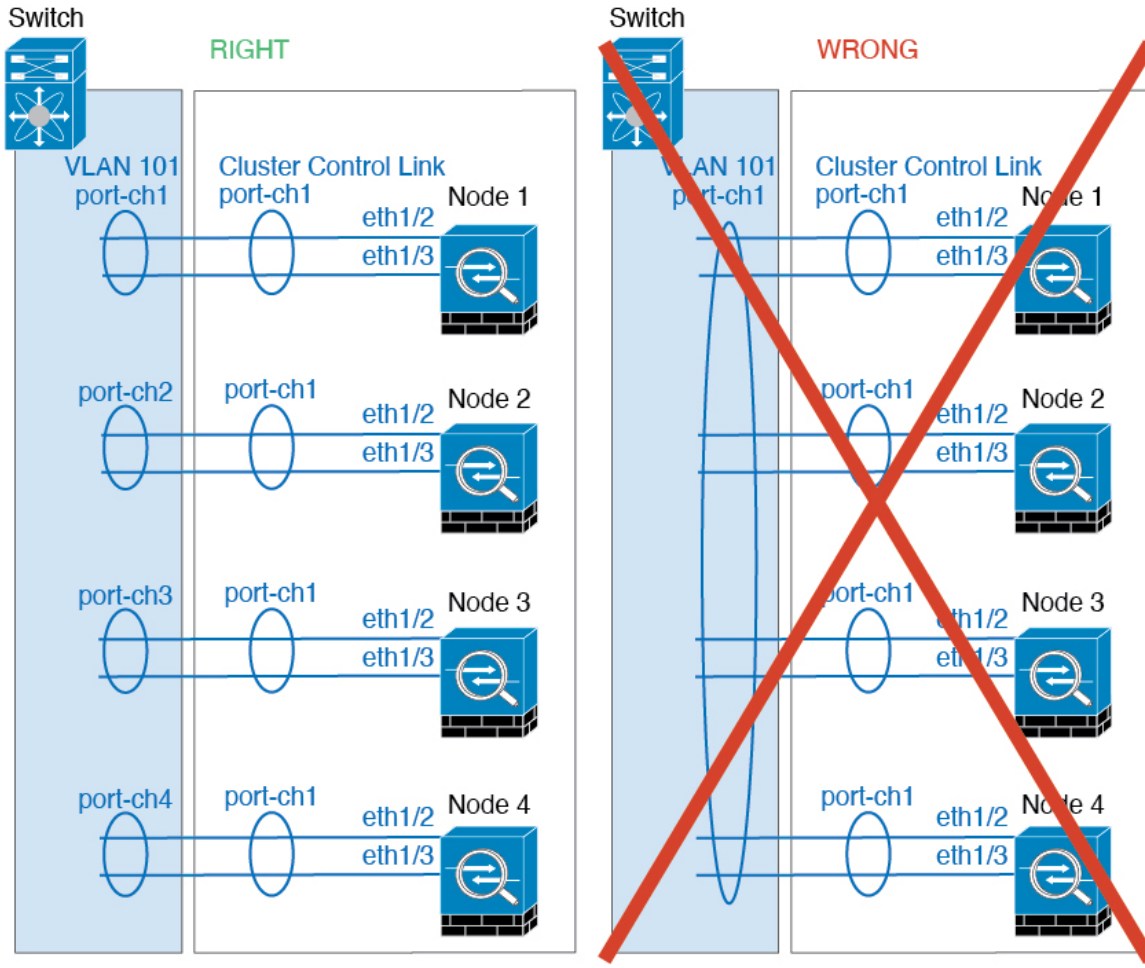
- Unlike ASA hardware clusters, Firepower 4100/9300 clusters support LACP graceful convergence. So for the platform, you can leave LACP graceful convergence enabled on connected Cisco Nexus switches.
- When you see slow bundling of a Spanned EtherChannel on the switch, you can enable LACP rate fast for 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.



Inter-Site Clustering

See the following guidelines for inter-site clustering:

- The cluster control link latency must be less than 20 ms round-trip time (RTT).
- The cluster control link must be reliable, with no out-of-order or dropped packets; for example, you should use a dedicated link.
- Do not configure connection rebalancing; you do not want connections rebalanced to cluster members at a different site.
- The ASA does not encrypt forwarded data traffic on the cluster control link because it is a dedicated link, even when used on a Data Center Interconnect (DCI). If you use Overlay Transport Virtualization (OTV), or are otherwise extending the cluster control link outside of the local administrative domain, you can configure encryption on your border routers such as 802.1AE MacSec over OTV.
- The cluster implementation does not differentiate between members at multiple sites for incoming connections; therefore, connection roles for a given connection may span across sites. This is expected behavior. However, if you enable director localization, the local director role is always chosen from the same site as the connection owner (according to site ID). Also, the local director chooses a new owner

at the same site if the original owner fails (Note: if the traffic is asymmetric across sites, and there is continuous traffic from the remote site after the original owner fails, then a node from the remote site might become the new owner if it receives a data packet within the re-hosting window.).

- For director localization, the following traffic types do not support localization: NAT or PAT traffic; SCTP-inspected traffic; Fragmentation owner query.
- For UDP long-lived flows in a North-South deployment, routing loops can occur if nodes at the original flow owner site fail and then come back up, after which the flow is directed back to the original site. If the new owner at the other site doesn't have a route to the destination, it will route the flow back to the internet, causing a loop. In this case, use the **clear conn** command on the new owner to force the flow to be reestablished.
- For transparent mode, if the cluster is placed between a pair of inside and outside routers (AKA North-South insertion), you must ensure that both inside routers share a MAC address, and also that both outside routers share a MAC address. When a cluster member at site 1 forwards a connection to a member at site 2, the destination MAC address is preserved. The packet will only reach the router at site 2 if the MAC address is the same as the router at site 1.
- For transparent mode, if the cluster is placed between data networks and the gateway router at each site for firewalling between internal networks (AKA East-West insertion), then each gateway router should use a First Hop Redundancy Protocol (FHRP) such as HSRP to provide identical virtual IP and MAC address destinations at each site. The data VLANs are extended across the sites using Overlay Transport Virtualization (OTV), or something similar. You need to create filters to prevent traffic that is destined to the local gateway router from being sent over the DCI to the other site. If the gateway router becomes unreachable at one site, you need to remove any filters so traffic can successfully reach the other site's gateway.
- For transparent mode, if the cluster is connected to an HSRP router, you must add the router HSRP MAC address as a static MAC address table entry on the ASA (see Add a Static MAC Address for Bridge Groups, on page 767). When adjacent routers use HSRP, traffic destined to the HSRP IP address will be sent to the HSRP MAC Address, but return traffic will be sourced from the MAC address of a particular router's interface in the HSRP pair. Therefore, the ASA MAC address table is typically only updated when the ASA ARP table entry for the HSRP IP address expires, and the ASA sends an ARP request and receives a reply. Because the ASA's ARP table entries expire after 14400 seconds by default, but the MAC address table entry expires after 300 seconds by default, a static MAC address entry is required to avoid MAC address table expiration traffic drops.
- For routed mode using Spanned EtherChannel, configure site-specific MAC addresses. Extend the data VLANs across the sites using OTV, or something similar. You need to create filters to prevent traffic that is destined to the global MAC address from being sent over the DCI to the other site. If the cluster becomes unreachable at one site, you need to remove any filters so traffic can successfully reach the other site's cluster nodes. Dynamic routing is not supported when an inter-site cluster acts as the first hop router for an extended segment.

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.

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.
- Connection rebalancing is disabled by default. If you enable connection rebalancing, the default time between load information exchanges is 5 seconds.
- 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 on the Firepower 4100/9300 Chassis

You can easily deploy the cluster from the Firepower 4100/9300 chassis supervisor. All initial configuration is automatically generated for each unit. This section describes the default bootstrap configuration and optional customization you can perform on the ASA. This section also describes how to manage cluster members from within the ASA. You can also manage cluster membership from the Firepower 4100/9300 chassis. See the Firepower 4100/9300 chassis documentation for more information.

Procedure

Step 1 FXOS: Add an ASA Cluster, on page 444
Step 2 ASA: Change the Firewall Mode and Context Mode, on page 452
Step 3 ASA: Configure Data Interfaces, on page 452
Step 4 ASA: Customize the Cluster Configuration, on page 454
Step 5 ASA: Manage Cluster Members, on page 469

FXOS: Add an ASA Cluster

You can add a single Firepower 9300 chassis as an intra-chassis cluster, or add multiple chassis for inter-chassis clustering. For inter-chassis clustering, 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 an ASA Cluster

Set the scope to the image version.

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, even if you do not have a module installed. If you do not configure all 3 modules, the cluster will not come up.

For multiple context mode, you must first deploy the logical device, and then enable multiple context mode in the ASA application.

When you deploy a cluster, the Firepower 4100/9300 chassis supervisor configures each ASA application with the following bootstrap configuration. You can later modify parts of the bootstrap configuration from the ASA, if desired (shown in **Bold** text).

```
interface Port-channel48
   description Clustering Interface
cluster group <service_type_name>
   key <secret>
   local-unit unit-<chassis#-module#>
   site-id <number>
   cluster-interface port-channel48 ip 127.2.<chassis#>.<module#> 255.255.255.0
   priority <auto>
   health-check holdtime 3
   health-check data-interface auto-rejoin 3 5 2
   health-check cluster-interface auto-rejoin unlimited 5 1
   enable

ip local pool cluster_ipv4_pool <ip_address>-<ip_address> mask <mask>

interface <management_ifc>
   management-only individual
   nameif management
   security-level 0
   ip address <ip_address> <mask> cluster-pool cluster_ipv4_pool
   no shutdown

http server enable
http 0.0.0.0 0.0.0.0 management
route management <management_host_ip> <mask> <gateway_ip> 1
```



Note The **local-unit** name can only be changed if you disable clustering.

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.
- Gather the following information:
  - Management interface ID, IP address, and network mask
  - Gateway IP address

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 204 or Configure a Physical Interface, on page 203.

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 439 for more information about EtherChannels.

b) Add a Management type interface or EtherChannel. See Add an EtherChannel (Port Channel), on page 204 or Configure a Physical Interface, on page 203.

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.

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 204.

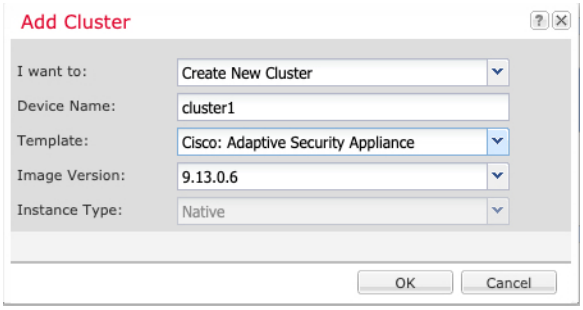
Do not add a member interface for a cluster isolated to security modules within one Firepower 9300 chassis. If you add a member, the chassis assumes this cluster will be using multiple chassis, and will only allow you to use Spanned EtherChannels, for example.

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 439 for more information about EtherChannels.

Step 2 Choose **Logical Devices**.

Step 3 Click **Add** > **Cluster**, and set the following parameters:



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 Adaptive Security Appliance**.
d) Choose the **Image Version**.
e) For the **Instance Type**, only the **Native** type is supported.
f) Click **OK**.

   You see the Provisioning - *device name* window.

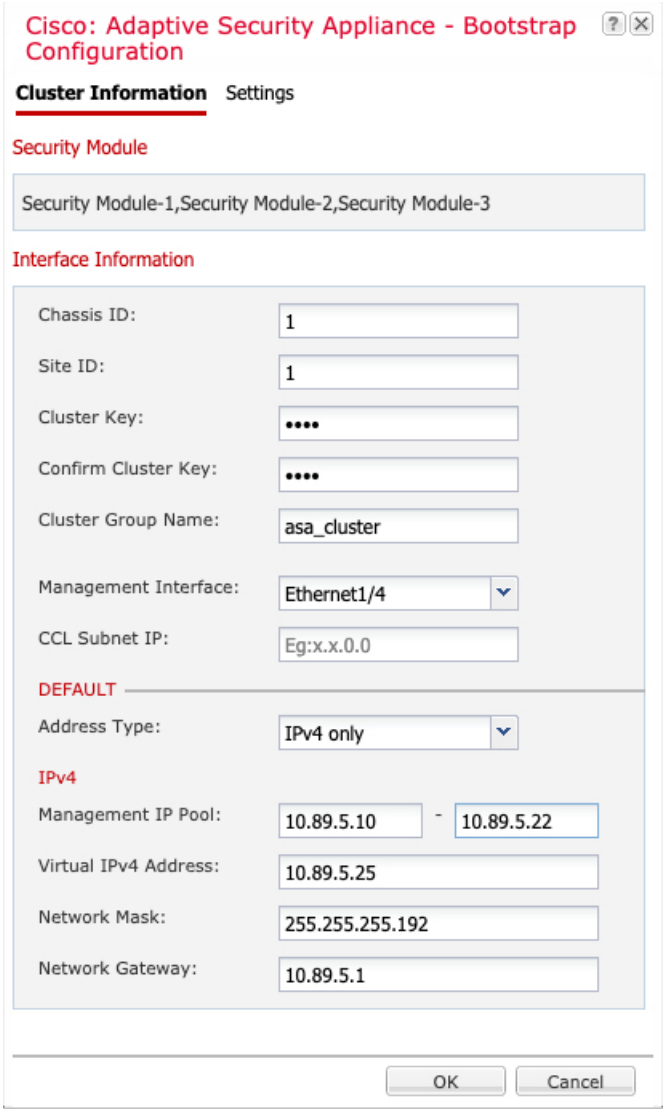
Step 4 Choose the interfaces you want to assign to this cluster.

All valid interfaces are assigned by default. If you defined multiple Cluster type interfaces, deselect all but one.

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.



a) 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.

b) For inter-site clustering, in the **Site ID** field, enter the site ID for this chassis between 1 and 8.

c) 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.

d) 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.

e) Choose the **Management Interface**.

This interface is used to manage the logical device. This interface is separate from the chassis management port.

f) Choose the **Address Type** for the management interface.

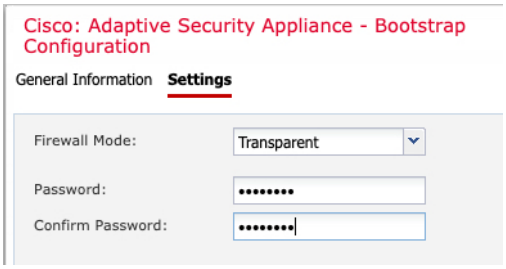
This information is used to configure a management interface in the ASA configuration. Set the following information:

- **Management IP Pool**—Configure a pool of Local IP addresses, one of which will be assigned to each cluster unit for the interface, by entering the starting and ending addresses separated by a hyphen.

  Include at least as many addresses as there are units in the cluster. Note that for the Firepower 9300, you must include 3 addresses per chassis, even if you do not have all module slots filled. If you plan to expand the cluster, include additional addresses. The Virtual IP address (known as the Main cluster IP address) that belongs to the current control unit is *not* a part of this pool; be sure to reserve an IP address on the same network for the Main cluster IP address. You can use IPv4 and/or IPv6 addresses.

- **Network Mask** or **Prefix Length**
- **Network Gateway**
- **Virtual IP address**—Set the management IP address of the current control unit. This IP address must be on the same network as the cluster pool addresses, but not be part of the pool.

Step 7 On the **Settings** page, complte the following.



a) From the **Firewall Mode** drop-down list, choose **Transparent** or **Routed**.

In routed mode, the 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.

b) Enter and confirm a **Password** for the admin user and for the enable password.

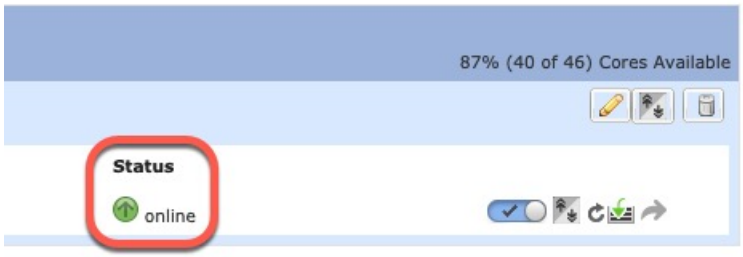
The pre-configured ASA admin user is useful for password recovery; if you have FXOS access, you can reset the admin user password if you forget it.

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 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 10 For clustering on multiple chassis, add the next chassis to the cluster:

a) On the first chassis of the chassis manager, click the **Show Configuration icon** at the top right; copy the displayed cluster configuration.

b) Connect to the 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**.

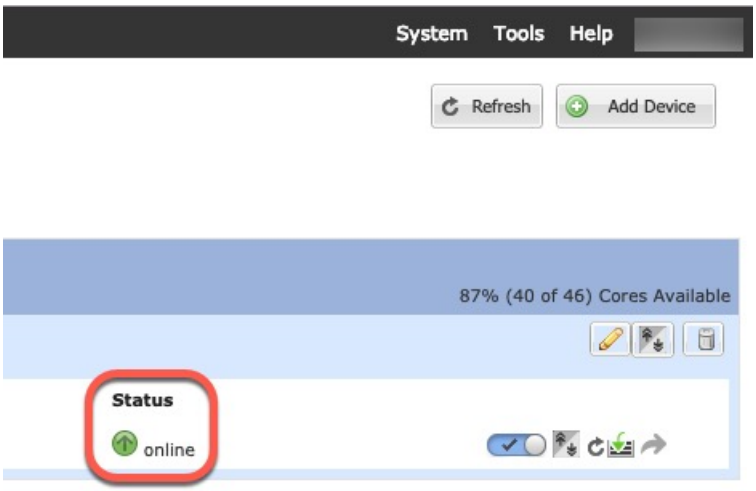
f) Click the device icon in the center of the screen. The cluster information is mostly pre-filled, but you must change the following settings:

- **Chassis ID**—Enter a unique chassis ID.
- **Site ID**—Enter the correct site ID.
- **Cluster Key**—(Not prefilled) Enter the same cluster key.

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 11 Connect to the control unit ASA to customize your clustering configuration.

Add More Cluster Members

Add or replace the ASA cluster member.

Note This procedure only applies to adding or replacing a *chassis*; if you are adding or replacing a module to a Firepower 9300 where clustering is already enabled, the module will be added automatically.

Before you begin

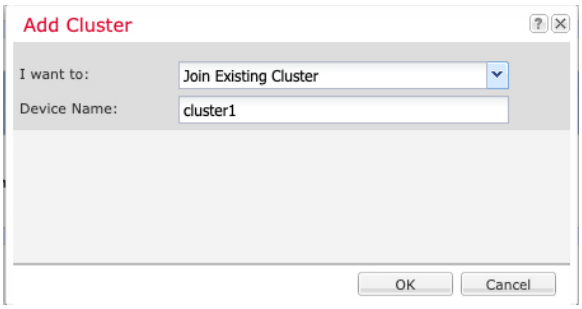
- Make sure your existing cluster has enough IP addresses in the management IP address pool for this new member. If not, you need to edit the existing cluster bootstrap configuration on each chassis before you add this new member. This change causes a restart of the logical device.
- The interface configuration must be the same on the new chassis. You can export and import FXOS chassis configuration to make this process easier.
- For multiple context mode, enable multiple context mode in the ASA application on the first cluster member; additional cluster members will inherit the multiple context mode configuration automatically.

Procedure

Step 1 On an existing cluster the chassis manager, choose **Logical Devices** to open the **Logical Devices** page.

Step 2 Click the Show Configuration icon ( ) at the top right; copy the displayed cluster configuration.

Step 3 Connect to the chassis manager on the new chassis, and click **Add** > **Cluster**.



Step 4 Choose **I want to:** > **Join Existing 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 mostly pre-filled, but you must change the following settings:

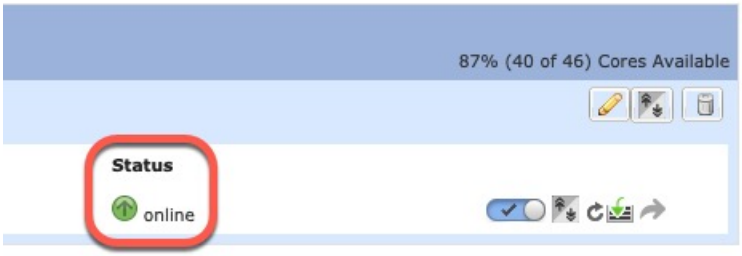
- **Chassis ID**—Enter a unique chassis ID.
- **Site ID**—Enter the correct site ID.
- **Cluster Key**—(Not prefilled) Enter the same cluster key.

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.





ASA: Change the Firewall Mode and Context Mode

By default, the FXOS chassis deploys a cluster in routed firewall mode, and single context mode.

- Change the firewall mode— To change the mode after you depoy, change the mode on the control unit; the mode is automatically changed on all data units to match. See Set the Firewall Mode (Single Mode), on page 225. In multiple context mode, you set the firewall mode per context. See Configure a Security Context, on page 266.
- Change to multiple context mode—To change to multiple context mode after you deploy, change the mode on the control unit; the mode is automatically changed on all data units to match. See Enable Multiple Context Mode, on page 261.

ASA: Configure 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.



Note The management interface was pre-configured when you deployed the cluster. You can also change the management interface parameters in ASA, but this procedure focuses on data interfaces. The management interface is an individual interface, as opposed to a Spanned interface. See Management Interface, on page 433 for more information.

Before you begin

- For multiple context mode, start this procedure in the system execution space. If you are not already in the System configuration mode in the Configuration > Device List pane, double-click **System** under the active device IP address.
- For transparent mode, configure the bridge group. See Configure the Bridge Virtual Interface (BVI), on page 641.
- When using Spanned EtherChannels for a cluster with multiple chassis, the port-channel interface will not come up until clustering is fully enabled. This requirement prevents traffic from being forwarded to a node that is not an active node in the cluster.

Procedure

Step 1 Depending on your context mode:

- For single mode, choose the **Configuration > Device Setup > Interface Settings > Interfaces** pane.
- For multiple mode in the System execution space, choose the **Configuration > Context Management > Interfaces** pane.

Step 2 Select the interface, and click **Edit**.

The **Edit Interface** dialog box appears.

Step 3 Set the following:

- (For EtherChannels) **MIO Port-channel ID**—Enter the same ID used in FXOS.
- **Enable Interface** (checked by default)

The rest of the fields on this screen are described later in this procedure.

Step 4 To configure the MAC address and optional parameters, click the **Advanced** tab.

- In the **MAC Address Cloning** area, set a manual global MAC address for the EtherChannel. Do not set the Standby MAC Address; it is ignored. You must configure a MAC address 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.

  In multiple context mode, if you share an interface between contexts, you should instead enable auto-generation of MAC addresses so you do not need to set the MAC address manually. Note that you must manually configure the MAC address using this command for *non-shared* interfaces.

- In the **ASA Cluster** area, for inter-site clustering set **Site specific MAC Addresses**, and for routed mode, the IP addresses for a site by clicking **Add** and specifying a MAC address and IP address for the site ID (1 through 8). Repeat for up to 8 sites. The site-specific IP addresses must be on the same subnet as the global IP address. The site-specific MAC address and IP address used by a unit depends on the site ID you specify in each unit's bootstrap configuration.

Step 5 (Optional) Configure VLAN subinterfaces on this EtherChannel. The rest of this procedure applies to the subinterfaces.

Step 6 (Multiple context mode) Before you complete this procedure, you need to allocate interfaces to contexts.

a) Click **OK** to accept your changes.
b) Allocate interfaces.
c) Change to the context that you want to configure: in the **Device List** pane, double-click the context name under the active device IP address.
d) Choose the **Configuration > Device Setup > Interface Settings > Interfaces** pane, select the port-channel interface that you want to customize, and click **Edit**.

   The **Edit Interface** dialog box appears.

Step 7 Click the **General** tab.

Step 8 (Transparent Mode) From the **Bridge Group** drop-down list, choose the bridge group to which you want to assign this interface.

Step 9 In the **Interface Name** field, enter a name up to 48 characters in length.

Step 10 In the **Security level** field, enter a level between 0 (lowest) and 100 (highest).

Step 11 (Routed Mode) For an IPv4 address, click the **Use Static IP** radio button and enter the IP address and mask. 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. For transparent mode, you configure the IP address for the bridge group interface, not the EtherChannel interface.

Step 12 (Routed Mode) To configure an IPv6 address, click the **IPv6** tab.

For transparent mode, you configure the IP address for the bridge group interface, not the EtherChannel interface.

a) Check the **Enable IPv6** check box.
b) In the **Interface IPv6 Addresses** area, click **Add**.

The **Add IPv6 Address for Interface** dialog box appears.

Note
The **Enable address autoconfiguration** option is not supported. Manually configuring the link-local address is also not supported.

c) In the **Address/Prefix Length** field, enter the global IPv6 address and the IPv6 prefix length. For example, 2001:DB8::BA98:0:3210/64.
d) (Optional) To use the Modified EUI-64 interface ID as the host address, check the **EUI-64** check box. In this case, just enter the prefix in the **Address/Prefix Length** field.
e) Click **OK**.

Step 13 Click **OK** to return to the **Interfaces** screen.

Step 14 Click **Apply**.

ASA: Customize the Cluster Configuration

If you want to change bootstrap settings after you deploy the cluster or configure additional options, such as clustering health monitoring, TCP connection replication delay, flow mobility, and other optimizations, you can do so on the control unit.

Configure Basic ASA Cluster Parameters

You can customize cluster settings on the control node.

Before you begin

- For multiple context mode, complete this procedure in the system execution space on the control unit. If you are not already in the System configuration mode, in the **Configuration** > **Device List** pane, double-click **System** under the active device IP address.
- The local-unit **Member Name** and several other options can only be set on the FXOS chassis, or they can only be changed on the ASA if you disable clustering, so they are not included in the following procedure.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster**.

Step 2 (Optional) Configure the following optional parameters:

- **Cluster Member Limit**—Configure the maximum number of cluster members, between 2 and 16. The default is 16. If you know that your cluster will be fewer than the maximum of 16 units, then we recommend that you set the actual planned number of units. Setting the maximum units lets the cluster

manage resources better. For example, if you use port address translation (PAT), then the control unit can allocate port blocks to the planned number of members, and it will not have to reserve ports for extra units you don't plan to use.

- **Site Periodic GARP**—The ASA generates gratuitous ARP (GARP) packets to keep the switching infrastructure up to date: the highest priority member at each site periodically generates GARP traffic for the global MAC/IP addresses. GARP is enabled by default when you set the site ID for each unit and the site MAC and IP address for each Spanned EtherChannel. Set the GARP interval between 1 and 1000000 seconds. The default is 290 seconds.

  When using per-site MAC and IP addresses, packets sourced from the cluster use a site-specific MAC address and IP address, while packets received by the cluster use a global MAC address and IP address. If traffic is not generated from the global MAC address periodically, you could experience a MAC address timeout on your switches for the global MAC address. After a timeout, traffic destined for the global MAC address will be flooded across the entire switching infrastructure, which can cause performance and security concerns.

- **Enable connection rebalancing for TCP traffic across all the ASAs in the cluster**—Enables connection rebalancing. This parameter is disabled by default. This parameter is not part of the bootstrap configuration, and is replicated from the control node to the data nodes. If enabled, ASAs exchange information about the connections per second periodically, and offload new connections from devices with more connections per second to less loaded devices. Existing connections are never moved. Moreover, because this command only rebalances based on connections per second, the total number of established connections on each node is not considered, and the total number of connections may not be equal. The frequency, between 1 and 360 seconds, specifies how often the load information is exchanged. The default is 5 seconds.

  Once a connection is offloaded to a different node, it becomes an asymmetric connection.

  Do not configure connection rebalancing for inter-site topologies; you do not want new connections rebalanced to cluster members at a different site.

- **Enable cluster load monitor**—You can monitor the traffic load for cluster members, including total connection count, CPU and memory usage, and buffer drops. If the load is too high, you can choose to manually disable clustering on the unit if the remaining units can handle the load, or adjust the load balancing on the external switch. This feature is enabled by default. For example, for inter-chassis clustering on the Firepower 9300 with 3 security modules in each chassis, if 2 security modules in a chassis leave the cluster, then the same amount of traffic to the chassis will be sent to the remaining module and potentially overwhelm it. You can periodically monitor the traffic load. If the load is too high, you can choose to manually disable clustering on the unit.

  Set the following values:

  - **Time Interval**—Sets the time in seconds between monitoring messages, between 10 and 360 seconds. The default is 20 seconds.
  - **Number of Intervals**—Sets the number of intervals for which the ASA maintains data, between 1 and 60. The default is 30.

  See **Monitoring** > **ASA Cluster** > **Cluster Load-Monitoring** to view the traffic load.

- **Enable health monitoring of this device within the cluster**—Enables the cluster unit health check feature, and determines the amount of time between unit heartbeat status messages, between .3 and 45 seconds; The default is 3 seconds. **Note**: When you are adding new units to the cluster, and making topology changes on the ASA or the switch, you should disable this feature temporarily until the cluster is complete, and also disable interface monitoring for the disabled interfaces (**Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** > **Cluster Interface Health**

Monitoring). You can re-enable this feature after cluster and topology changes are complete. To determine unit health, the ASA cluster units send heartbeat messages on the cluster control link to other units. If a unit does not receive any heartbeat messages from a peer unit within the holdtime period, the peer unit is considered unresponsive or dead.

- **Debounce Time**—Configures the debounce time before the ASA considers an interface to be failed and the unit is removed from the cluster. This feature allows 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 ASA waits the number of milliseconds specified before marking the interface as failed and the unit 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 unit just because another cluster unit was faster at bundling the ports. The default debounce time is 500 ms, with a range of 300 ms to 9 seconds.
- **Replicate console output**—Enables console replication from data units to the control unit. This feature is disabled by default. The ASA may print out some messages directly to the console for certain critical events. If you enable console replication, data units send the console messages to the control unit so that you only need to monitor one console port for the cluster. This parameter is not part of the bootstrap configuration, and is replicated from the control unit to the data units.
- **Enable Clustering Flow Mobility**. See Configure LISP Inspection, on page 461.
- **Enable Director Localization for inter-DC cluster**—To improve performance and reduce round-trip time latency for inter-site clustering for data centers, you can enable director localization. New connections are typically load-balanced and owned by cluster members within a given site. However, the ASA assigns the Director role to a member at *any* site. Director localization enables additional Director roles: a Local Director at the same site as the Owner, and a Global Director that can be at any site. Keeping the Owner and Director at the same site improves performance. Also, if the original Owner fails, the Local Director will choose a new connection Owner at the same site. The Global Director is used if a cluster member receives packets for a connection that is owned on a different site.
- **Site Redundancy**—To protect flows from a site failure, you can enable site redundancy. If the connection backup owner is at the same site as the owner, then an additional backup owner will be chosen from another site to protect flows from a site failure. Director localization and site redundancy are separate features; you can configure one or the other, or configure both.
- **Enable config sync acceleration**—When a data unit has the same configuration as the control unit, it will skip syncing the configuration and will join faster. This feature is enabled by default. This feature is configured on each unit, and is not replicated from the control unit to the data unit.

  Note
  Some configuration commands are not compatible with accelerated cluster joining; if these commands are present on the unit, even if accelerated cluster joining is enabled, configuration syncing will always occur. You must remove the incompatible configuration for accelerated cluster joining to work. Use the **show cluster info unit-join-acceleration incompatible-config** to view incompatible configuration.

- **Enable parallel configuration replicate**—Enable the control unit to sync configuration changes with data units in parallel. Otherwise, synching occurs sequentially, and can take more time.
- **Flow State Refresh Keepalive Interval**—Set the keepalive interval for flow state refresh messages (clu_keepalive and clu_update messages) from the flow owner to the director and backup owner, between 15 and 20 seconds. The default is 15. You may want to set the interval to be longer than the default to reduce the amount of traffic on the cluster control link.

Step 3 In the **Cluster Control Link** area, you can configue the cluster control link MTU. Other options in this area cannot be configured on the ASA.

- **MTU**—Specify the maximum transmission unit for the cluster control link interface to be at least 100 bytes higher than the highest MTU of the data interfaces. We suggest setting the MTU to the maximum of 9184; the minimum value is 1400 bytes. 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. 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.

  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.

Step 4 (Optional) (Firepower 9300 only) **VPN Grouping Mode**. See Configure Distributed Site-to-Site VPN, on page 462.

Step 5 (Optional) (Firepower 9300 only) In the **Parallel Join of Units Per Chassis** area, you can ensure that the security modules in a chassis join the cluster simultaneously, so that traffic is evenly distributed between the modules. If a module joins very much in advance of other modules, it can receive more traffic than desired, because the other modules cannot yet share the load.

- **Minimum Units Required to Join**—Specifies the minimum number of modules in the same chassis required to be ready before a module can join the cluster, between 1 and 3. The default is 1, meaning that a module will not wait for other modules to be ready before it joins the cluster. If you set the value to 3, for example, then each module will wait the maximum delay time or until all 3 modules are ready before joining the cluster. All 3 modules will request to join the cluster roughly simultaneously, and will all start receiving traffic around the same time.
- **Maximum Join Delay**—Specifies the maximum delay time in minutes before a module stops waiting for other modules to be ready before it joins the cluster, between 0 and 30 minutes. The default is 0, meaning the module will not wait for other modules to be ready before it joins the cluster. If you set the mimimum units to 1, then this value must be 0. If you set the mimimum units to 2 or 3, then this value must be 1 or more. This timer is per module, but when the first module joins the cluster, then all other module timers end, and the remaining modules join the cluster.

For example, you set the minimum units to 3, and the maximum delay to 5 minutes. When module 1 comes up, it starts its 5 minute timer. Module 2 comes up 2 minutes later and starts its 5 minute timer. Module 3 comes up 1 minute later, therefore all modules will now join the cluster at the 4 minute mark; they will not wait for the timers to complete. If module 3 never comes up, then Module 1 will join the cluster at the end of its 5 minute timer, and Module 2 will also join, even though its timer still has 2 minutes remaining; it will not wait for its timer to complete.

Step 6 Click **Apply**.

Configure Interface Health Monitoring and Auto-Rejoin Settings

You might want to disable health monitoring of non-essential interfaces, for example, the management interface. You can monitor any port-channel ID or single physical interface ID. 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 **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** > **Cluster Interface Health Monitoring**.

Step 2 Select an interface in the **Monitored Interfaces** box, and click **Add** to move it to the **Unmonitored Interfaces** box.

Interface status messages detect link failure. If all physical ports for a given logical interface fail on a particular unit, but there are active ports under the same logical interface on other units, then the unit is removed from the cluster. If a unit does not receive interface status messages within the holdtime, then the amount of time before the ASA removes a member from the cluster depends on the type of interface and whether the unit is an established member or is joining the cluster. Health check is enabled by default for all interfaces.

You might want to disable health monitoring of non-essential interfaces, for example, the management interface. You can specify any port-channel ID or single physical interface ID. 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.

When any topology changes occur (such as adding or removing a data interface, enabling or disabling an interface on the ASA, Firepower 4100/9300 chassis, or the switch, or adding an additional switch to form a VSS, vPC, StackWise, or StackWise Virtual) you should disable the health check feature (**Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster**) 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.

Step 3 Click the **Auto Rejoin** tab to customize the auto-rejoin settings in case of an interface, system, or cluster control link failure. For each type, click **Edit** to set the following:

- **Maximum Rejoin Attempts**—Define the number of attempts at rejoining the cluster by setting **Unlimited** or a value between 0 and 65535. **0** disables auto-rejoining. The default value is **Unlimited** for the cluster-interface and **3** for the data-interface and system.
- **Rejoin Interval**—Define the interval duration in minutes between rejoin attempts by setting the interval between 2 and 60. The default value is **5** minutes. The maximum total time that the unit attempts to rejoin the cluster is limited to 14400 minutes (10 days) from the time of last failure.
- **Interval Variation**—Define if the interval duration increases by setting the interval variation 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.

Click **Restore Defaults** to restore the default settings.

Check **Chassis Heartbeat Delay Auto-Rejoin** to set the chassis rejoin to match the **Auto Rejoin** settings for chassis heartbeat failures. By default, if the chassis heartbeat fails and then recovers, the node rejoins the cluster immediately. However, if you configure this option, it will rejoin according to the settings of the **Auto Rejoin** screen.

Step 4 Click **Apply**.

Configure the Cluster TCP Replication Delay

Enable the cluster replication delay for TCP connections to help eliminate the “unnecessary work” related to short-lived flows by delaying the director/backup flow creation. Note that if a unit fails before the director/backup flow is created, then those flows cannot be recovered. Similarly, if traffic is rebalanced to a different unit before the flow is created, then the flow cannot be recovered. You should not enable the TCP replication delay for traffic on which you disable TCP randomization.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster Replication**.

Step 2 Click **Add** and set the following values:

- **Replication delay**—Set the seconds between 1 and 15.
- **HTTP**—Set the delay for all HTTP traffic. This setting is enabled by default for 5 seconds.
- **Source Criteria**
  - **Source**—Set the source IP address.
  - Service—(Optional) Set the source port. Typically you set either the source or the destination port, but not both.
- **Destination Criteria**
  - **Source**—Set the destination IP address.
  - **Service**—(Optional) Set the destination port. Typically you set either the source or the destination port, but not both.

Step 3 Click **OK**.

Step 4 Click **Apply**.

Configure Inter-Site Features

For inter-site clustering, you can customize your configuration to enhance redundancy and stability.

Configure Cluster Flow Mobility

You can inspect LISP traffic to enable flow mobility when a server moves between sites.

About LISP Inspection

You can inspect LISP traffic to enable flow mobility between sites.

About LISP

Data center virtual machine mobility such as VMware VMotion enables servers to migrate between data centers while maintaining connections to clients. To support such data center server mobility, routers need to be able to update the ingress route towards the server when it moves. Cisco Locator/ID Separation Protocol

(LISP) architecture separates the device identity, or endpoint identifier (EID), from its location, or routing locator (RLOC), into two different numbering spaces, making server migration transparent to clients. For example, when a server moves to a new site and a client sends traffic to the server, the router redirects traffic to the new location.

LISP requires routers and servers in certain roles, such as the LISP egress tunnel router (ETR), ingress tunnel router (ITR), first hop routers, map resolver (MR), and map server (MS). When the first hop router for the server senses that the server is connected to a different router, it updates all of the other routers and databases so that the ITR connected to the client can intercept, encapsulate, and send traffic to the new server location.

Secure Firewall ASA LISP Support

The ASA does not run LISP itself; it can, however, inspect LISP traffic for location changes and then use this information for seamless clustering operation. Without LISP integration, when a server moves to a new site, traffic comes to an ASA cluster member at the new site instead of to the original flow owner. The new ASA forwards traffic to the ASA at the old site, and then the old ASA has to send traffic back to the new site to reach the server. This traffic flow is sub-optimal and is known as “tromboning” or “hair-pinning.”

With LISP integration, the ASA cluster members can inspect LISP traffic passing between the first hop router and the ETR or ITR, and can then change the flow owner to be at the new site.

LISP Guidelines

- The ASA cluster members must reside between the first hop router and the ITR or ETR for the site. The ASA cluster itself cannot be the first hop router for an extended segment.
- Only fully-distributed flows are supported; centralized flows, semi-distributed flows, or flows belonging to individual nodes are not moved to new owners. Semi-distributed flows include applications, such as SIP, where all child flows are owned by the same ASA that owns the parent flow.
- The cluster only moves Layer 3 and 4 flow states; some application data might be lost.
- For short-lived flows or non-business-critical flows, moving the owner may not be worthwhile. You can control the types of traffic that are supported with this feature when you configure the inspection policy, and should limit flow mobility to essential traffic.

ASA LISP Implementation

This feature includes several inter-related configurations (all of which are described in this chapter):

1. (Optional) Limit inspected EIDs based on the host or server IP address—The first hop router might send EID-notify messages for hosts or networks the ASA cluster is not involved with, so you can limit the EIDs to only those servers or networks relevant to your cluster. For example, if the cluster is only involved with 2 sites, but LISP is running on 3 sites, you should only include EIDs for the 2 sites involved with the cluster.
2. LISP traffic inspection—The ASA inspects LISP traffic on UDP port 4342 for the EID-notify message sent between the first hop router and the ITR or ETR. The ASA maintains an EID table that correlates the EID and the site ID. For example, you should inspect LISP traffic with a source IP address of the first hop router and a destination address of the ITR or ETR. Note that LISP traffic is not assigned a director, and LISP traffic itself does not participate in cluster state sharing.
3. Service Policy to enable flow mobility on specified traffic—You should enable flow mobility on business-critical traffic. For example, you can limit flow mobility to only HTTPS traffic, and/or to traffic to specific servers.
4. Site IDs—The ASA uses the site ID for each cluster node to determine the new owner.

5. Cluster-level configuration to enable flow mobility—You must also enable flow mobility at the cluster level. This on/off toggle lets you easily enable or disable flow mobility for a particular class of traffic or applications.

Configure LISP Inspection

You can inspect LISP traffic to enable flow mobility when a server moves between sites.

Before you begin

- Set the site ID for the chassis on Firepower 4100/9300 chassis supervisor.
- LISP traffic is not included in the default-inspection-traffic class, so you must configure a separate class for LISP traffic as part of this procedure.

Procedure

Step 1 (Optional) Configure a LISP inspection map to limit inspected EIDs based on IP address, and to configure the LISP pre-shared key:

a) Choose **Configuration** > **Firewall** > **Objects** > **Inspect Maps** > **LISP**.
b) Click **Add** to add a new map.
c) Enter a name (up to 40 characters) and description.
d) For the **Allowed-EID access-list**, click **Manage**.

The **ACL Manager** opens.

The first hop router or ITR/ETR might send EID-notify messages for hosts or networks the ASA cluster is not involved with, so you can limit the EIDs to only those servers or networks relevant to your cluster. For example, if the cluster is only involved with 2 sites, but LISP is running on 3 sites, you should only include EIDs for the 2 sites involved with the cluster.

e) Add an ACL with at least one ACE according to the firewall configuration guide.
f) If necessary, enter the **Validation Key**.

If you copied an encrypted key, click the **Encrypted** radio button.

g) Click **OK**.

Step 2 Add a service policy rule to configure LISP inspection:

a) Choose **Configuration** > **Firewall** > **Service Policy Rules**.
b) Click **Add**.
c) On the **Service Policy** page, apply the rule to an interface or globally.

If you have an existing service policy you want to use, add a rule to that policy. By default, the ASA includes a global policy called **global_policy**. You can also create one service policy per interface if you do not want to apply the policy globally. LISP inspection is applied to traffic bidirectionally so you do not need to apply the service policy on both the source and destination interfaces; all traffic that enters or exits the interface to which you apply the rule is affected if the traffic matches the class for both directions.

d) On the **Traffic Classification Criteria** page, click **Create a new traffic class**, and under **Traffic Match Criteria**, check **Source and Destination IP Address (uses ACL)**.
e) Click **Next**.

f) Specify the traffic you want to inspect. You should specify traffic between the first hop router and the ITR or ETR on UDP port 4342. Both IPv4 and IPv6 ACLs are accepted.
g) Click **Next**.
h) On the **Rule Actions** wizard page or tab, select the **Protocol Inspection** tab.
i) Check the **LISP** check box,.
j) (Optional) Click **Configure** to choose the inspection map you created.
k) Click **Finish** to save the service policy rule.

Step 3 Add a service policy rule to enable Flow Mobility for critical traffic:

a) Choose **Configuration** > **Firewall** > **Service Policy Rules**.
b) Click **Add**.
c) On the **Service Policy** page, choose the same service policy you used for LISP inspection.
d) On the **Traffic Classification Criteria** page, click **Create a new traffic class**, and under **Traffic Match Criteria**, check **Source and Destination IP Address (uses ACL)**.
e) Click **Next**.
f) Specify the business critical traffic that you want to re-assign to the most optimal site when servers change sites. For example, you can limit flow mobility to only HTTPS traffic, and/or to traffic to specific servers. Both IPv4 and IPv6 ACLs are accepted.
g) Click **Next**.
h) On the **Rule Actions** wizard page or tab, select the **Cluster** tab.
i) Check the **Enable Cluster flow-mobility triggered by LISP EID messages** check box.
j) Click **Finish** to save the service policy rule.

Step 4 Choose **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** > **Cluster Configuration**, and check the **Enable Clustering flow mobility** check box.

Step 5 Click **Apply**.

Configure Distributed Site-to-Site VPN

By default, the cluster uses centralized site-to-site VPN mode. To take advantage of the scalability of clustering, you can enable sistributed site-to-site VPN mode.

About Distributed Site-to-Site VPN

In distributed mode, site-to-site IPsec IKEv2 VPN connections are distributed across nodes of a cluster. Distributing VPN connections across the nodes of a cluster allows both the capacity and throughput of the cluster to be fully utilized, significantly scaling VPN support beyond centralized VPN capabilities.

Distributed VPN Connection Roles

When running in distributed VPN mode, the following roles are assigned to the cluster nodes:

- Active Session Owner—The node that initially receives the connection or that has transitioned a backup session to an active session. The owner maintains the state and processes packets for the complete session, including the IKE and IPsec tunnels and all traffic associated with them.
- Backup Session Owner—The node that is handling the backup session for an existing active session. If the active session owner fails, the backup session owner becomes the active session owner, and a new backup session is established on a different node.

- Forwarder—If traffic associated with a VPN session is sent to a node that does not own the VPN session, that node will use the cluster control link to forward the traffic to the node that owns the VPN session.
- Orchestrator—The orchestrator (always the control node of the cluster) is responsible for calculating which sessions will move and where they will move to when executing an Active Session Redistribution (ASR). It sends a request to the owner node X to move N sessions to node Y. Node X will respond back to the orchestrator when complete, specifying how many sessions it was able to move.

Distributed VPN Session Characteristics

Distributed site-to-site VPN Sessions have the following characteristics. Otherwise, VPN connections behave as they normally do if not on a cluster.

- VPN sessions are distributed across the cluster at the session level. Meaning the same cluster node handles the IKE and IPsec tunnels and all their traffic for a VPN connection. If VPN session traffic is sent to a cluster node that does not own that VPN session, traffic is forwarded to the cluster node that owns the VPN session.
- VPN sessions have a Session ID that is unique across the cluster. Using the session ID, traffic is validated, forwarding decisions are made, and IKE negotiation is completed.
- In a site-to-site VPN hub and spoke configuration, when clients connect through the cluster (called hair-pinning), the session traffic flowing in and the session traffic flowing out may be on different cluster nodes.
- You can require that the backup session to be allocated on a security module in another chassis; this provides protection against chassis failure. Or, you can choose to allocate backup sessions on any node in the cluster; this provides protection against node failure only. When there are two chassis in the cluster, remote-chassis backup is strongly recommended.

Distributed VPN Handling of Cluster Events

Event	Distributed VPN
Node failure	For all active sessions on this failed node, the backup sessions (on another node) become active, and backup sessions are reallocated on another node according to the backup strategy.
Chassis failure	When a remote-chassis backup strategy is being used, for all active sessions on the failed chassis, the backup sessions (on a node in the other chassis) become active. When the nodes are replaced, backup sessions for these now active sessions will be reallocated on nodes in the replaced chassis. When a flat backup strategy is being used, if both the active and backup sessions are on the failed chassis, the connection will drop. All active sessions with backup sessions on a node in the other chassis fall back to these sessions. New backup sessions will be allocated on another node in the surviving chassis.
Inactivate a cluster node	For all active sessions on the cluster node being inactivated, backup sessions (on another node) become active and reallocate backup sessions on another node according to the backup strategy.

Event	Distributed VPN
Cluster node join	If the VPN cluster mode on the new node is not set to distributed, the control node will request a mode change. After the VPN mode is compatible, the cluster node will be assigned active and backup sessions in the flow of normal operations.

IPsec IKEv2 Modifications

IKEv2 is modified while in distributed site-to-site VPN mode in the following ways:

- An identity is used in place of IP/port tuples. This allows for proper forwarding decisions on the packets, and cleanup of previous connections that may be on other cluster members.
- The (SPI) identifiers that identify a single IKEv2 session are locally generated, random 8-byte values that are unique across the cluster. An SPI embeds a time stamp and a cluster node ID. Upon receipt of an IKE negotiation packet, if the time stamp or cluster node ID check fails, the packet is dropped and a message is logged indicating the reason.
- IKEv2 processing has been modified to prevent NAT-T negotiations from failing by being split across cluster members. A new ASP classify domain, *cluster_isakmp_redirect*, and rules are added when IKEv2 is enabled on an interface.

High Availability for Distributed Site-to-Site VPN Within the Cluster

The following capabilities provide resiliency against single failure of a security module or chassis:

- VPN Sessions that are backed up on another security module in the cluster, on any chassis, withstand security module failures.
- VPN Sessions that are backed up on another chassis withstand chassis failures.
- The control node can change without losing VPN site-to-site sessions.

If an additional failure occurs before the cluster has stabilized, connections may be lost if the both active and backup sessions are on the failed nodes.

All attempts are made to ensure no sessions are lost when a node leaves the cluster in a graceful manner such as disabling the VPN cluster mode, reloading a cluster node, and other anticipated chassis changes. During these types of operations, sessions will not be lost as long as the cluster is given time to re-establish session backups between operations. If a graceful exit is triggered on the last cluster node, it will gracefully tear down existing sessions.

CMPv2

The CMPv2 ID certificate and key pairs are synchronized across the cluster nodes. However, only the control node in the cluster automatically renews and rekeys the CMPv2 certificate. The control node synchronizes these new ID certificates and keys to all cluster nodes on a renewal. In this way, all nodes in the cluster utilize the CMPv2 certificates for authentication, and also any node is capable of taking over as the control node.

Licenses for Distributed Site-to-Site VPN

A Carrier license is required for distributed site-to-site VPN, on each member of the cluster.

Each VPN connection requires two *Other VPN* licensed sessions (the *Other VPN* license is part of the *Essentials* license), one for the active session and one for the backup session. The maximum VPN session capacity of the cluster can be no more than half of the licensed capacity due to using two licenses for each session.

Prerequisites for Distributed Site-to-Site VPN

Model Support

- Firepower 9300
- Maximum of 6 modules on up to 2 chassis. You can have different quantities of installed security modules in each chassis, although we recommend an equal distribution.

Maximum VPN Sessions

Each security module supports up to 6K VPN sessions for a maximum of approximately 36K sessions across 6 nodes.

The actual number of sessions supported on a cluster node is determined by platform capacity, allocated licenses, and per context resource allocation. When utilization is close to the limit, there may be cases where session creation fails, even though the maximum capacity has not been reached on each cluster node. This is because active session allocation is determined by external switching, and backup session allocation is determined by an internal cluster algorithm. Customers are encouraged to size their utilization accordingly and allow room for uneven distribution.

Guidelines for Distributed Site-to-Site VPN

Firewall Mode

Distributed site-to-site VPN is supported in routed mode only.

Context Mode

Distributed site-to-site VPN operates in both single and multiple context modes. However, in multiple context mode, active session redistribution is done at the system level, not at the context level. This prevents an active session associated with a context from moving to a cluster member that contains active sessions associated with a different context, unknowingly creating an unsupportable load.

Unsupported Inspections

The following types of inspections are not supported or are disabled in distributed site-to-site VPN mode:

- CTIQBE
- DCERPC
- H323, H225, and RAS
- IPsec pass-through
- MGCP
- MMP
- NetBIOS

- PPTP
- RADIUS
- RSH
- RTSP
- SCCP (Skinny)
- SUNRPC
- TFTP
- WAAS
- WCCP
- XDMCP

Additional Guidelines

- Only IKEv2 IPsec site-to-site VPN is supported in distributed site-to-site VPN mode. IKEv1 is not supported. IKEv1 site-to-site is supported in centralized VPN mode.
- Inter-site clustering is not supported.
- Dynamic PAT is not available while in distributed site-to-site VPN mode.

Enable Distributed Site-to-Site VPN

Enable distributed site-to-site VPN to take advantage of the scalability of clustering for VPN sessions.



Note Changing the VPN mode between centralized and distributed requires all nodes in the cluster to reload. Changing the backup mode is dynamic and will not terminate sessions.

Before you begin

Configure site-to-site VPN according to the VPN configuration guide.

Procedure

Step 1 Choose **Configuration > Device Management > High Availability and Scalability > ASA Cluster**.

Step 2 In the **VPN Cluster Mode** area, choose the **VPN Mode** for the cluster, **Centralized** or **Distributed**.

Step 3 Choose the **Backup Distribution Mode**, **Flat** or **Remote-chassis**.

In **flat** backup mode, standby sessions are established on any other cluster node. This will protect users from module failures; however, chassis failure protection is not guaranteed.

In **remote-chassis** backup mode, standby sessions are established on a node of another chassis in the cluster. This will protect users from both module failures and chassis failures.

If **remote-chassis** is configured in a single chassis environment (intentionally configured or the result of a failure), no backups will be created until another chassis joins.

Step 4 Click **Apply**.

You will be prompted to reload. This setting is replicated to all data nodes before reloading; all nodes in the cluster will reload.

Redistribute Distributed S2S VPN Sessions

Active Session Redistribution (ASR) redistributes the active VPN session load across the cluster members. Due to the dynamic nature of beginning and ending sessions, ASR is a best effort balancing of the sessions across all cluster members. Repeated redistribution actions will optimize the balance.

Redistribution can be run at any time, should be run after any topology change in the cluster, and is recommended after a new member joins the cluster. The goal of redistribution is to creat a stable VPN cluster. A stable VPN cluster has an almost equal number of active and backup sessions across the nodes.

To move a session, the backup session becomes the active one and another node is selected to host a new backup session. Moving sessions is dependent on the location of the active session's backup and the number of active sessions already on that particular backup node. If the backup session node is unable to host the active session for some reason, the original node remains owner of the session.

In multiple-context mode, active session redistribution is done at the system level, not the individual context level. It is not done at the context level because an active session in one context could be moved a member that contains many more active sessions in a different context, creating more load on that cluster member.

Before you begin

- Enable system logs if you would like to monitor redistribution activity.
- This procedure must be carried out on the control unit of the cluster.

Procedure

Step 1 Choose **Monitoring > ASA Cluster > ASA Cluster > Cluster Summary > VPN Cluster Summary** to view how active and backup sessions are distributed across the cluster.

Depending on the number of sessions to redistribute and the load on the cluster, this may take some time. Syslogs containing the following phrases (and other system details not shown here) are provided as redistribution activity occurs:

Syslog Phrase	Notes
VPN session redistribution started	Control node only
Sent request to move *number* sessions from *orig-member-name* to *dest-member-name*	Control node only
Failed to send session redistribution message to *member-name*	Control node only
Received request to move *number* sessions from *orig-member-name* to *dest-member-name*	Data node only

Syslog Phrase	Notes
Moved *number* sessions to *member-name*	The number of active sessions moved to the named cluster.
Failed to receive session move response from *dest-member-name*	Control node only
VPN session completed	Control node only
Cluster topology change detected. VPN session redistribution aborted.	

Step 2 Click **Re-Distribute**.

Step 3 Refresh the **Monitoring > ASA Cluster > ASA Cluster > Cluster Summary > VPN Cluster Summary** to see the results of the redistribution activity.

If your redistribution was successful, and there has been no substantial system or session activity, your system will be balanced and this action is complete.

Otherwise, repeat the redistribution process to obtain a balanced and stable system.

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 chassis manager **Logical Devices** page:



- 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 (for example, you saved the configuration with clustering disabled), the Management interface is disabled.

To reenable clustering, on the ASA enter **cluster group** *name* and then **enable**.

- Disable the application instance—In the 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 chassis manager on the **Security Module/Engine** page, click the **Power Off icon**.
- Shut down the chassis—In the chassis manager on the **Overview** page, click the **Shut Down icon**.

Permanent Removal

You can permanently remove a cluster node using the following methods.

- Delete the logical device—In the 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.

ASA: Manage Cluster Members

After you deploy the cluster, you can change the configuration and manage cluster members.

Become an Inactive Member

To become an inactive member of the cluster, disable clustering on the node while leaving the clustering configuration intact.



Note When an ASA becomes inactive (either manually or through a health check failure), all data interfaces are shut down; only the management-only interface can send and receive traffic. To resume traffic flow, re-enable clustering; or you can remove the node altogether from the cluster. The management interface remains up using the IP address the node received from the cluster IP pool. However if you reload, and the node is still inactive in the cluster (for example, you saved the configuration with clustering disabled), then the management interface is disabled. You must use the console port for any further configuration.

Before you begin

- For multiple context mode, perform this procedure in the system execution space. If you are not already in the System configuration mode in the Configuration > Device List pane, double-click **System** under the active device IP address.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** > **Cluster Configuration**.

Step 2 Uncheck the **Participate in ASA cluster** check box.

Note
Do not uncheck the **Configure ASA cluster settings** check box; this action clears all cluster configuration, and also shuts down all interfaces including the management interface to which ASDM is connected. To restore connectivity in this case, you need to access the CLI at the console port.

Step 3 Click **Apply**.

Deactivate a Data Unit from the Control Unit

To deactivate a data node, perform the following steps.



Note When an ASA becomes inactive, all data interfaces are shut down; only the management-only 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 cluster IP pool. However if you reload, and the node is still inactive in the cluster (for example, if you saved the configuration with clustering disabled), the management interface is disabled. You must use the console port for any further configuration.

Before you begin

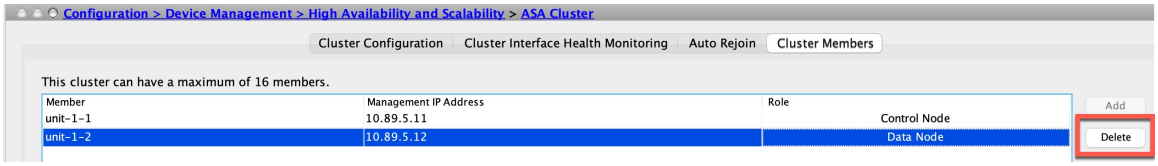
For multiple context mode, perform this procedure in the system execution space. If you are not already in the System configuration mode in the **Configuration > Device List** pane, double-click **System** under the active device IP address.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** > **Cluster Members**.

Step 2 Select the data node that you want to remove, and click **Delete**.

Figure 67: Delete Node



The data node bootstrap configuration remains intact, so that you can later re-add the data node without losing your configuration.

Step 3 Click **Apply**.

Rejoin the Cluster

If a node was removed from the cluster, for example for a failed interface or if you manually deactivated a member, you must manually rejoin the cluster.

Before you begin

- You must use the console port to reenable clustering. Other interfaces are shut down. The exception is if you manually disabled clustering in ASDM, then you can reenable clustering in ASDM if you did not save the configuration and reload. After reloading, the management interface is disabled, so console access is the only method to reenable clustering.
- For multiple context mode, perform this procedure in the system execution space. If you are not already in the System configuration mode in the **Configuration > Device List** pane, double-click **System** under the active device IP address.
- Make sure the failure is resolved before you try to rejoin the cluster.

Procedure

Step 1 If you still have ASDM access, you can reenable clustering in ASDM by connecting ASDM to the node you want to reenable.

You cannot reenable clustering for a data node from the control node unless you add it as a new member.

a) Choose **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster**.
b) Check the **Participate in ASA cluster** check box.
c) Click **Apply**.

Step 2 If you cannot use ASDM: At the console, enter cluster configuration mode:

cluster group *name*

Example:

```
ciscoasa(config)# cluster group pod1
```

Step 3 Enable clustering.

enable

Change the Control Unit



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 node you want to become the control node, use the procedure in this section. Note, however, that for centralized features, if you force a control node change using this procedure, 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.

Before you begin

For multiple context mode, perform this procedure in the system execution space. If you are not already in the System configuration mode in the Configuration > Device List pane, double-click **System** under the active device IP address.

Procedure

Step 1 Choose **Monitoring** > **ASA Cluster** > **Cluster Summary**.

Step 2 From the drop-down list, choose a data node to become control, and click the button to make it the control node.

Step 3 You are prompted to confirm the control node change. Click **Yes**.

Step 4 Quit ASDM, and reconnect using the Main cluster IP address.

Execute a Command Cluster-Wide

To send a command to all members in the cluster, or to a specific member, perform the following steps. Sending a **show** command to all members collects all output and displays it on the console of the current unit. (Note that alternatively there are show commands that you can enter on the control unit to view cluster-wide statistics.) Other commands, such as **capture** and **copy**, can also take advantage of cluster-wide execution.

Before you begin

Perform this procedure at the Command Line Interface tool: choose **Tools** > **Command Line Interface**.

Procedure

Send a command to all members, or if you specify the unit name, a specific member:

cluster exec [unit *unit_name*] *command*

Example:

```
cluster exec show xlate
```

To view member names, enter **cluster exec unit ?** (to see all names except the current unit), or enter the **show cluster info** command.

Examples

To copy the same capture file from all units in the cluster at the same time to a TFTP server, enter the following command on the control unit:

```
cluster exec copy /pcap capture: tftp://10.1.1.56/capture1.pcap
```

Multiple PCAP files, one from each unit, are copied to the TFTP server. The destination capture file name is automatically attached with the unit name, such as capture1_asa1.pcap, capture1_asa2.pcap, and so on. In this example, asa1 and asa2 are cluster unit names.

The following sample output for the **cluster exec show memory** command shows memory information for each member in the cluster:

```
cluster exec show memory
unit-1-1(LOCAL):*******************************************************
Free memory:      108724634538 bytes (92%)
Used memory:        9410087158 bytes ( 8%)
-------------     ------------------
Total memory:     118111600640 bytes (100%)


unit-1-3:*******************************************************************
Free memory:      108749922170 bytes (92%)
Used memory:        9371097334 bytes ( 8%)
-------------     ------------------
Total memory:     118111600640 bytes (100%)


unit-1-2:*******************************************************************
Free memory:      108426753537 bytes (92%)
Used memory:        9697869087 bytes ( 8%)
-------------     ------------------
Total memory:     118111600640 bytes (100%)
```

ASA: Monitoring the ASA Cluster on the Firepower 4100/9300 chassis

You can monitor and troubleshoot cluster status and connections.

Monitoring Cluster Status

See the following screens for monitoring cluster status:

- **Monitoring > ASA Cluster > Cluster Summary**

  This pane shows cluster information about the unit to which you are connected, as well as other units in the cluster. You can also change the primary unit from this pane.

- **Cluster Dashboard**

  On the home page on the primary unit, you can monitor the cluster using the Cluster Dashboard and the Cluster Firewall Dashboard.

Capturing Packets Cluster-Wide

See the following screen for capturing packets in a cluster:

Wizards > Packet Capture Wizard

To support cluster-wide troubleshooting, you can enable capture of cluster-specific traffic on the control node, which is then automatically enabled on all of the data nodes in the cluster.

Monitoring Cluster Resources

See the following screens for monitoring cluster resources:

- **Monitoring > ASA Cluster > System Resources Graphs > CPU**

  This pane lets you create graphs or tables showing the CPU utilization across the cluster members.

- **Monitoring > ASA Cluster > System Resources Graphs > Memory**.

  This pane lets you create graphs or tables showing the Free Memory and Used Memory across the cluster members.

Monitoring Cluster Traffic

See the following screens for monitoring cluster traffic:

- **Monitoring > ASA Cluster > Traffic Graphs > Connections**.

  This pane lets you create graphs or tables showing the Connections across the cluster members.

- **Monitoring > ASA Cluster > Traffic Graphs > Throughput**.

  This pane lets you create graphs or tables showing the traffic throughput across the cluster members.

- **Monitoring** > **ASA Cluster** > **Cluster Load-Monitoring**

  This section includes the **Load Monitor-Information** and **Load-Monitor Details** panes. **Load Monitor-Information** shows the traffic load for cluster members for the last interval and also the average over total number of intervals configured (30 by default). Use the **Load-Monitor Details** pane to view the value for each measure at each interval.

Monitoring the Cluster Control Link

See the following screen for monitoring cluster status:

Monitoring > Properties > System Resources Graphs > Cluster Control Link.

This pane lets you create graphs or tables showing the cluster control link Receival and Transmittal capacity utilization.

Monitoring Cluster Routing

See the following screen for cluster routing:

- **Monitoring** > **Routing** > **LISP-EID Table**

  Shows the ASA EID table showing EIDs and site IDs.

Monitoring Distributed S2S VPN

See the following screens for monitoring VPN cluster status:

- **Monitoring > ASA Cluster > ASA Cluster > Cluster Summary > VPN Cluster Summary**

  Displays the distribution of the session across the cluster and provides you with the ability to redistribute the sessions.

- **Monitoring > VPN > VPN Statistics > Sessions**

  Both control and data cluster nodes are listed. Click any node for details.

Configuring Logging for Clustering

See the followingscreen for configuring logging for clustering:

Configuration > Device Management > Logging > Syslog Setup

Each node in the cluster generates syslog messages independently. You can generate syslog messages with identical or different device IDs to make messages appear to come from the same or different nodes in the cluster.

Troubleshooting Distributed S2S VPN

Distributed VPN Notifications

You will be notified with messages containing the identified phrases when the following error situations occur on a cluster running distributed VPN:

Situation	Notification

If an existing or joining cluster data node is not in distributed VPN mode when attempting to join the cluster:	`New cluster member (member-name) rejected due to vpn mode mismatch.`
	and
	`Control node (control-name) rejects enrollment request from unit (unit-name) for the reason: the vpn mode capabilities are not compatible with the control node configuration`

If licensing is not properly configured on a cluster member for Distributed VPN:	`ERROR: Control node requested cluster vpn-mode change to distributed. Unable to change mode due to missing Carrier License.`

If the time stamp or member ID is invalid in the SPI of a received IKEv2 packet:	`Expired SPI received`
	or
	`Corrupted SPI detected`

If the cluster is unable to create a backup session:	`Failed to create the backup for an IKEv2 session.`
IKEv2 Initial Contact (IC) processing error:	`IKEv2 Negotiation aborted due to ERROR: Stale backup session found on backup`
Redistribution problems:	`Failed to send session redistribution message to member-name` `Failed to receive session move response from member-name (control node only)`
If the topology changes during redistribution of the sessions:	`Cluster topology change detected. VPN session redistribution aborted.`

You may be encountering one of the following situations:

- Site-to-site VPN sessions are being distributed to only one of the chassis in a cluster when the Nexus 7K switch is configured with a layer 4 port as a load-balancing algorithm using the **port-channel load-balance src-dst l4port** command. An example of the cluster session allocation looks like below:

```
SSP-Cluster/data node(cfg-cluster)# show cluster vpn-sessiondb distribution
Member 0 (unit-1-3): active: 0
Member 1 (unit-2-2): active: 13295; backups at: 0(2536), 2(2769), 3(2495), 4(2835),
5(2660)
```

```
Member 2 (unit-2-3): active: 12174; backups at: 0(2074), 1(2687), 3(2207), 4(3084),
5(2122)
Member 3 (unit-2-1): active: 13416; backups at: 0(2419), 1(3013), 2(2712), 4(2771),
5(2501)
Member 4 (unit-1-1): active: 0
Member 5 (unit-1-2): active: 0
```

Since site-to-site IKEv2 VPN uses port 500 for both source and destination ports, IKE packets are only sent to one of the links in the port channel connected between the Nexus 7K and the chassis.

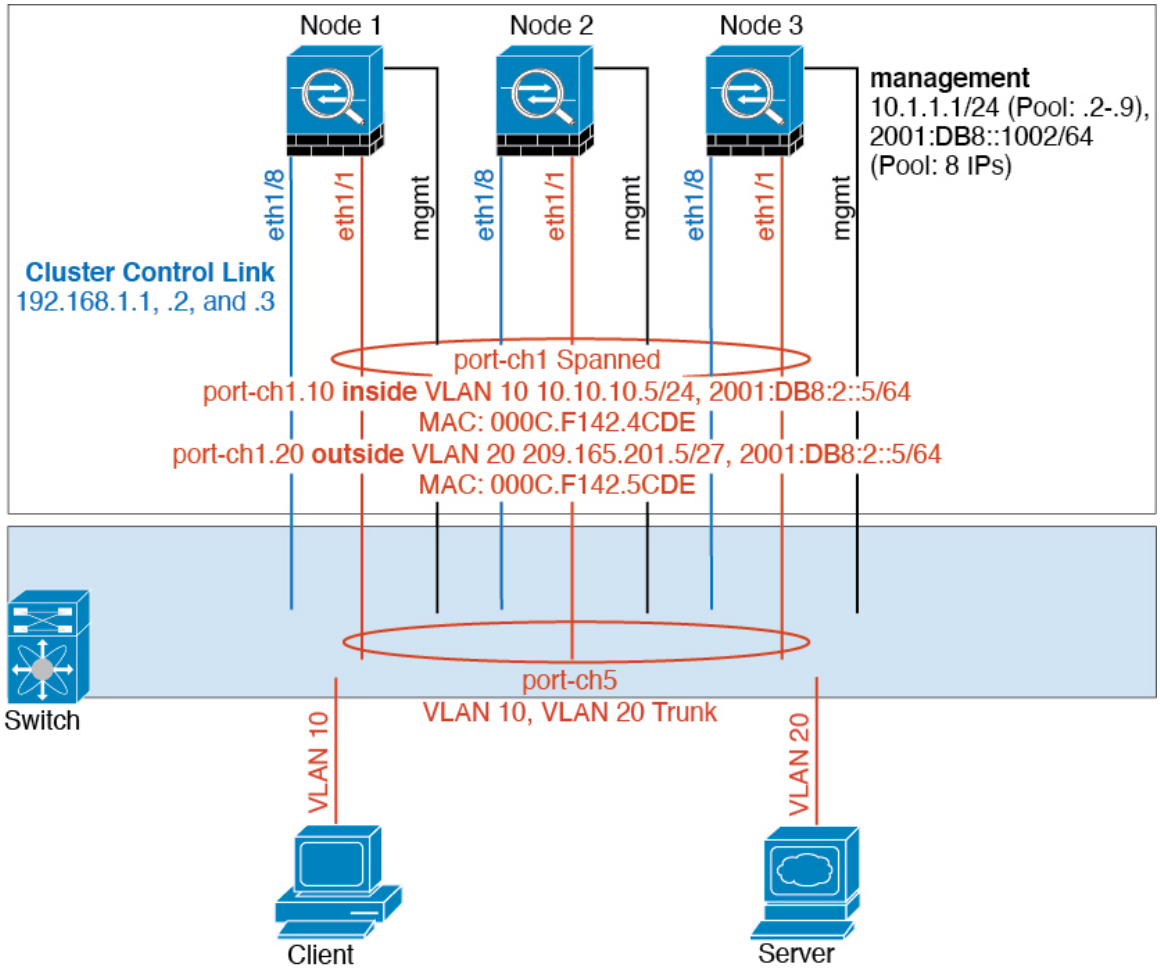
Change the Nexus 7K load balancing algorithm to IP and Layer 4 port using the **port-channel load-balance src-dst ip-l4port**. Then the IKE packets are sent to all the links and thus all nodes.

For a more immediate adjustment, on the control node of the cluster, execute: **cluster redistribute vpn-sessiondb** to redistribute active VPN sessions to the cluster nodes of the other chassis.

Examples for ASA Clustering

These examples include typical deployments.

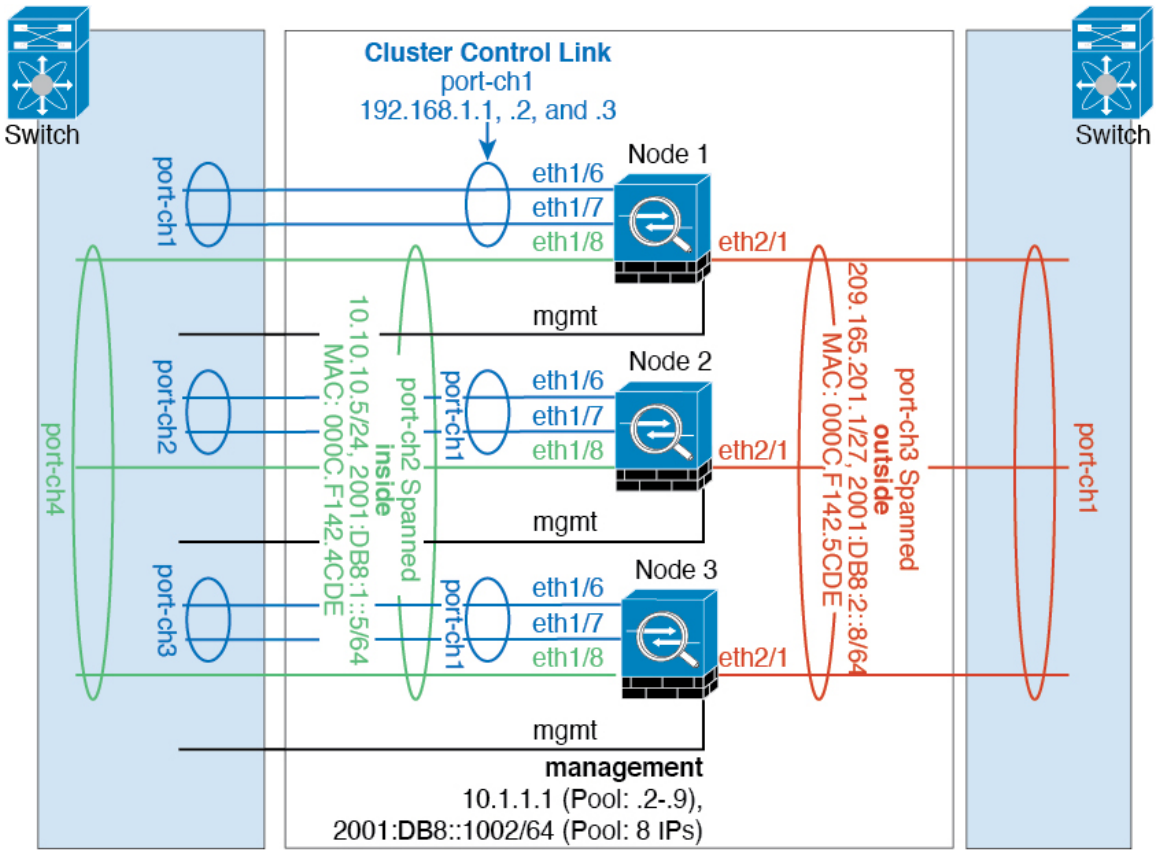
Firewall on a Stick



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 ASA 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 ASA 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 ASA 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.

OTV Configuration for Routed Mode Inter-Site Clustering

The success of inter-site clustering for routed mode with Spanned EtherChannels depends on the proper configuration and monitoring of OTV. OTV plays a major role by forwarding the packets across the DCI. OTV forwards unicast packets across the DCI only when it learns the MAC address in its forwarding table. If the MAC address is not learned in the OTV forwarding table, it will drop the unicast packets.

Sample OTV Configuration

```
//Sample OTV config:
//3151 - Inside VLAN, 3152 - Outside VLAN, 202 - CCL VLAN
//aaaa.1111.1234 - ASA inside interface global vMAC
//0050.56A8.3D22 - Server MAC

feature ospf
feature otv
```

```
mac access-list ALL_MACs
  10 permit any any
mac access-list HSRP_VMAC
  10 permit aaaa.1111.1234 0000.0000.0000 any
  20 permit aaaa.2222.1234 0000.0000.0000 any
  30 permit any aaaa.1111.1234 0000.0000.0000
  40 permit any aaaa.2222.1234 0000.0000.0000
vlan access-map Local 10
  match mac address HSRP_VMAC
  action drop
vlan access-map Local 20
  match mac address ALL_MACs
  action forward
vlan filter Local vlan-list 3151-3152


//To block global MAC with ARP inspection:
arp access-list HSRP_VMAC_ARP
  10 deny aaaa.1111.1234 0000.0000.0000 any
  20 deny aaaa.2222.1234 0000.0000.0000 any
  30 deny any aaaa.1111.1234 0000.0000.0000
  40 deny any aaaa.2222.1234 0000.0000.0000
  50 permit ip any mac
ip arp inspection filter HSRP_VMAC_ARP 3151-3152

no ip igmp snooping optimise-multicast-flood
vlan 1,202,1111,2222,3151-3152

otv site-vlan 2222
mac-list GMAC_DENY seq 10 deny aaaa.aaaa.aaaa ffff.ffff.ffff
mac-list GMAC_DENY seq 20 deny aaaa.bbbb.bbbb ffff.ffff.ffff
mac-list GMAC_DENY seq 30 permit 0000.0000.0000 0000.0000.0000
route-map stop-GMAC permit 10
  match mac-list GMAC_DENY

interface Overlay1
  otv join-interface Ethernet8/1
  otv control-group 239.1.1.1
  otv data-group 232.1.1.0/28
  otv extend-vlan 202, 3151
  otv arp-nd timeout 60
  no shutdown

interface Ethernet8/1
  description uplink_to_OTV_cloud
  mtu 9198
  ip address 10.4.0.18/24
  ip igmp version 3
  no shutdown

interface Ethernet8/2

interface Ethernet8/3
  description back_to_default_vdc_e6/39
  switchport
    switchport mode trunk
    switchport trunk allowed vlan 202,2222,3151-3152
  mac packet-classify
  no shutdown

otv-isis default
  vpn Overlay1
    redistribute filter route-map stop-GMAC
otv site-identifier 0x2
```

```
//OTV flood not required for ARP inspection:
otv flood mac 0050.56A8.3D22 vlan 3151
```

OTV Filter Modifications Required Because of Site Failure

If a site goes down, the filters need to be removed from OTV because you do not want to block the global MAC address anymore. There are some additional configurations required.

You need to add a static entry for the ASA global MAC address on the OTV switch in the site that is functional. This entry will let the OTV at the other end add these entries on the overlay interface. This step is required because if the server and client already have the ARP entry for the ASA, which is the case for existing connections, then they will not send the ARP again. Therefore, OTV will not get a chance to learn the ASA global MAC address in its forwarding table. Because OTV does not have the global MAC address in its forwarding table, and per OTV design it will not flood unicast packets over the overlay interface, then it will drop the unicast packets to the global MAC address from the server, and the existing connections will break.

```
//OTV filter configs when one of the sites is down

mac-list GMAC_A seq 10 permit 0000.0000.0000 0000.0000.0000
route-map a-GMAC permit 10
  match mac-list GMAC_A

otv-isis default
  vpn Overlay1
    redistribute filter route-map a-GMAC

no vlan filter Local vlan-list 3151

//For ARP inspection, allow global MAC:
arp access-list HSRP_VMAC_ARP_Allow
  50 permit ip any mac
ip arp inspection filter HSRP_VMAC_ARP_Allow 3151-3152

mac address-table static aaaa.1111.1234 vlan 3151 interface Ethernet8/3
//Static entry required only in the OTV in the functioning Site
```

When the other site is restored, you need to add the filters back again and remove this static entry on the OTV. It is very important to clear the dynamic MAC address table on both the OTVs to clear the overlay entry for the global MAC address.

MAC Address Table Clearing

When a site goes down, and a static entry for the global MAC address is added to OTV, you need to let the other OTV learn the global MAC address on the overlay interface. After the other site comes up, these entries should be cleared. Make sure to clear the mac address table to make sure OTV does not have these entries in its forwarding table.

```
cluster-N7k6-OTV# show mac address-table
Legend:
* - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
age - seconds since last seen,+ - primary entry using vPC Peer-Link,
(T) - True, (F) - False
VLAN MAC Address Type age Secure NTFY Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
G -    d867.d900.2e42 static  -  F F sup-eth1(R)
O 202  885a.92f6.44a5 dynamic -  F F Overlay1
```

```
* 202  885a.92f6.4b8f dynamic 5  F F Eth8/3
O 3151 0050.5660.9412 dynamic -  F F Overlay1
* 3151 aaaa.1111.1234 dynamic 50 F F Eth8/3
```

OTV ARP Cache Monitoring

OTV maintains an ARP cache to proxy ARP for IP addresses that it learned across the OTV interface.

```
cluster-N7k6-OTV# show otv arp-nd-cache
OTV ARP/ND L3->L2 Address Mapping Cache

Overlay Interface Overlay1
VLAN MAC Address Layer-3 Address Age Expires In
3151 0050.5660.9412 10.0.0.2 1w0d 00:00:31
cluster-N7k6-OTV#
```

Examples for Inter-Site Clustering

The following examples show supported cluster deployments.

Spanned EtherChannel Routed Mode Example with Site-Specific MAC and IP Addresses

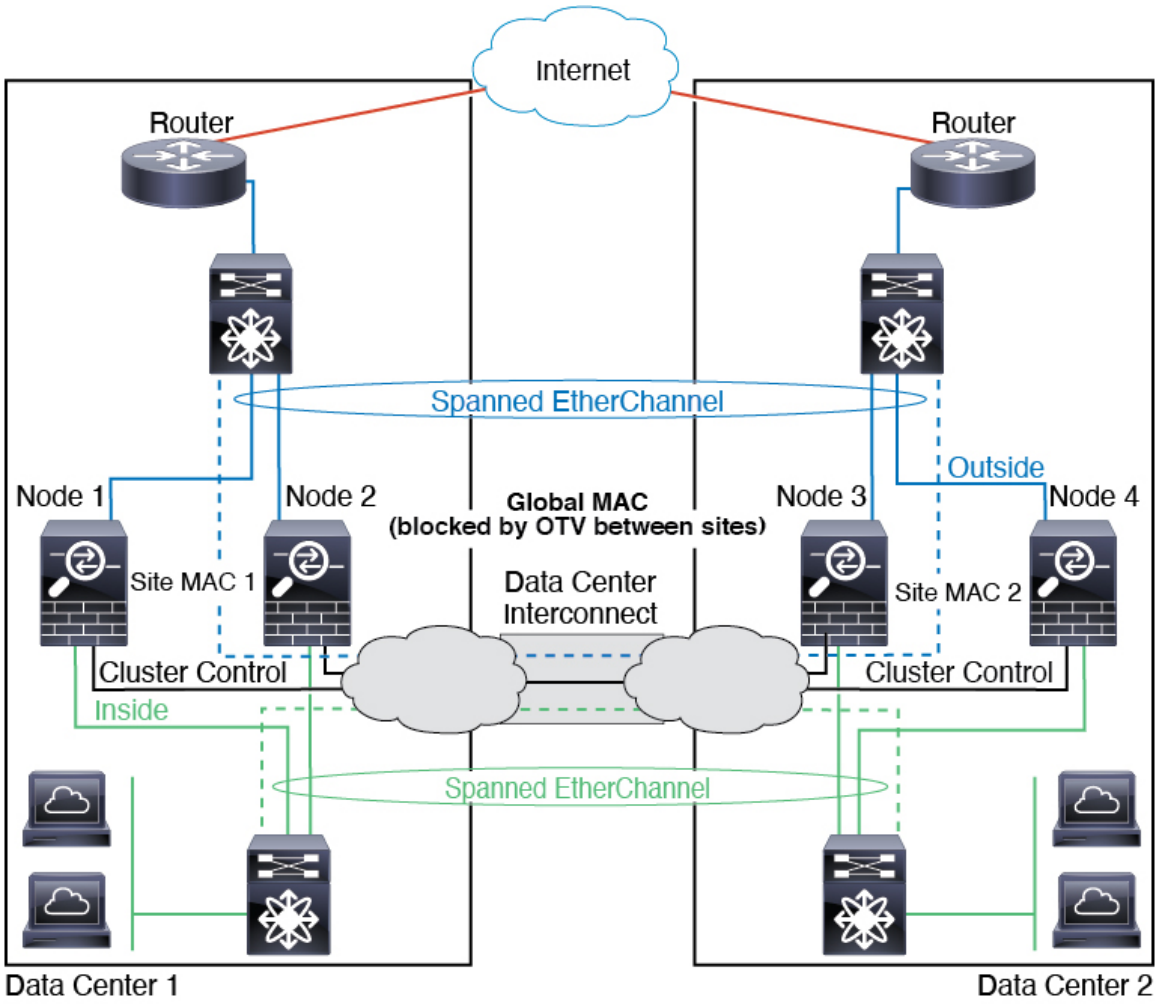
The following example shows 2 cluster members at each of 2 data centers placed between the gateway router and an inside network at each site (East-West insertion). The cluster members are connected by the cluster control link over the DCI. The cluster members at each site connect to the local switches using spanned EtherChannels for both the inside and outside networks. Each EtherChannel is spanned across all chassis in the cluster.

The data VLANs are extended between the sites using Overlay Transport Virtualization (OTV) (or something similar). You must add filters blocking the global MAC address to prevent traffic from traversing the DCI to the other site when the traffic is destined for the cluster. If the cluster nodes at one site become unreachable, you must remove the filters so traffic can be sent to the other site's cluster nodes. You should use VACLs to filter the global MAC address.For some switches, such as Nexus with the F3-series line card, you must also use ARP inspection to block ARP packets from the global MAC address. ARP inspection requires you to set both the site MAC address and the site IP address on the ASA. If you only configure the site MAC address be sure to disable ARP inspection.

The cluster acts as the gateway for the inside networks. The global virtual MAC, which is shared across all cluster nodes, is used only to receive packets. Outgoing packets use a site-specific MAC address from each DC cluster. This feature prevents the switches from learning the same global MAC address from both sites on two different ports, which causes MAC flapping; instead, they only learn the site MAC address.

In this scenario:

- All egress packets sent from the cluster use the site MAC address and are localized at the data center.
- All ingress packets to the cluster are sent using the global MAC address, so they can be received by any of the nodes at both sites; filters at the OTV localize the traffic within the data center.



Spanned EtherChannel Transparent Mode North-South Inter-Site Example

The following example shows 2 cluster members at each of 2 data centers placed between inside and outside routers (North-South insertion). The cluster members are connected by the cluster control link over the DCI. The cluster members at each site connect to the local switches using spanned EtherChannels for the inside and outside. Each EtherChannel is spanned across all chassis in the cluster.

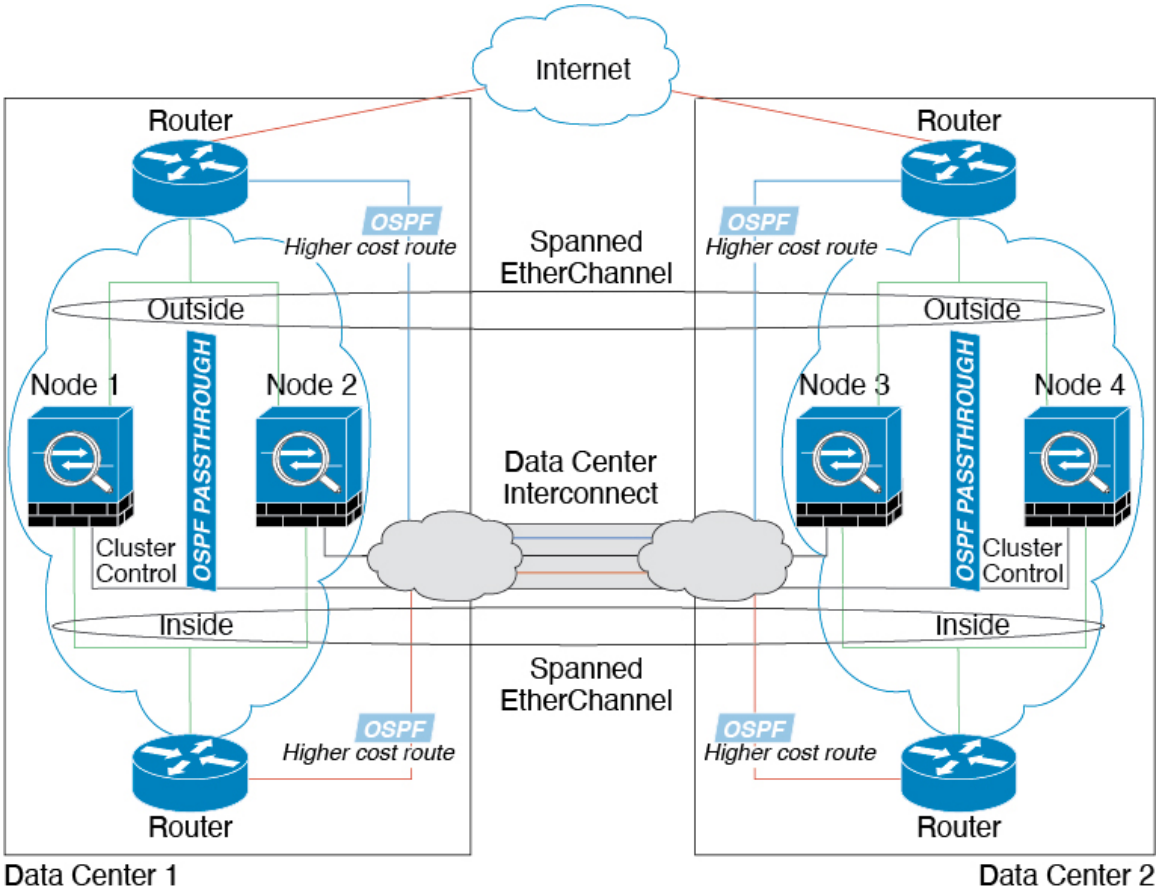
The inside and outside routers at each data center use OSPF, which is passed through the transparent ASAs. Unlike MACs, router IPs are unique on all routers. By assigning a higher cost route across the DCI, traffic stays within each data center unless all cluster members at a given site go down. The lower cost route through the ASAs must traverse the same bridge group at each site for the cluster to maintain asymmetric connections. In the event of a failure of all cluster members at one site, traffic goes from each router over the DCI to the cluster members at the other site.

The implementation of the switches at each site can include:

- Inter-site VSS, vPC, StackWise, or StackWise Virtual—In this scenario, you install one switch at Data Center 1, and the other at Data Center 2. One option is for the cluster nodes at each Data Center to only connect to the local switch, while the redundant switch traffic goes across the DCI. In this case, connections are for the most part kept local to each datacenter. You can optionally connect each node to both switches

across the DCI if the DCI can handle the extra traffic. In this case, traffic is distributed across the data centers, so it is essential for the DCI to be very robust.

- Local VSS, vPC, StackWise, or StackWise Virtual at each site—For better switch redundancy, you can install 2 separate redundant switch pairs at each site. In this case, although the cluster nodes still have a spanned EtherChannel with Data Center 1 chassis connected only to both local switches, and Data Center 2 chassis connected to those local switches, the spanned EtherChannel is essentially "split." Each local redundant switch system sees the spanned EtherChannel as a site-local EtherChannel.

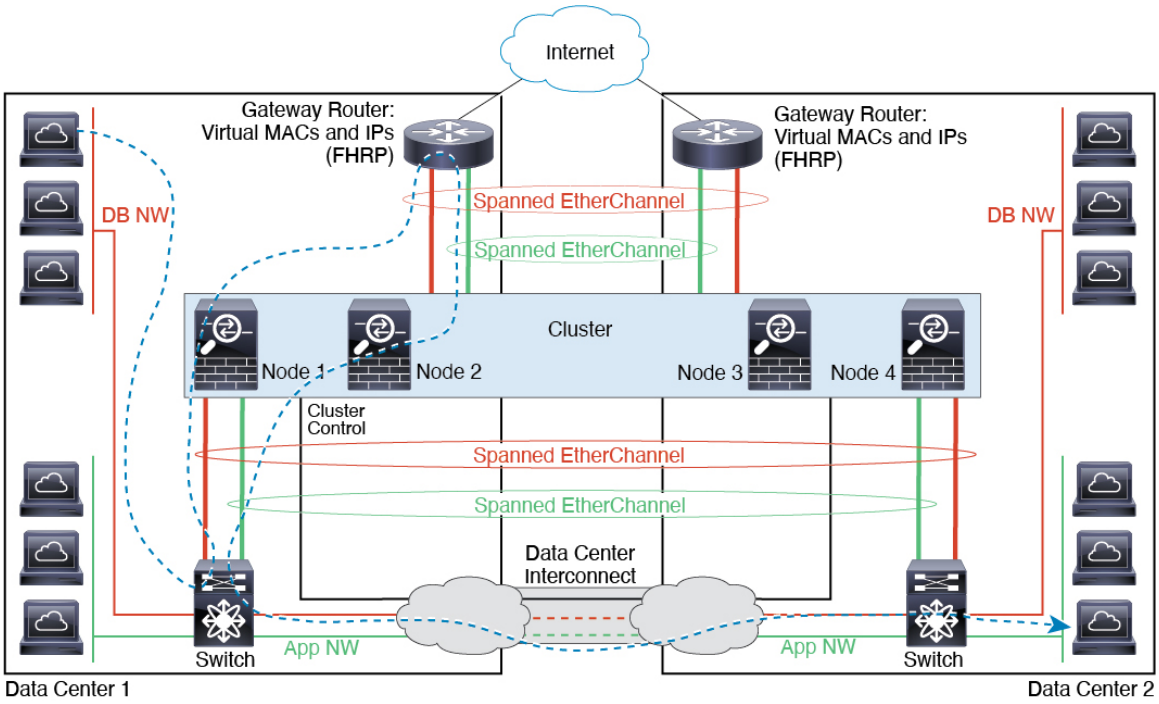


Spanned EtherChannel Transparent Mode East-West Inter-Site Example

The following example shows 2 cluster members at each of 2 data centers placed between the gateway router and two inside networks at each site, the App network and the DB network (East-West insertion). The cluster members are connected by the cluster control link over the DCI. The cluster members at each site connect to the local switches using spanned EtherChannels for both the App and DB networks on the inside and outside. Each EtherChannel is spanned across all chassis in the cluster.

The gateway router at each site uses an FHRP such as HSRP to provide the same destination virtual MAC and IP addresses at each site. A good practice to avoid unintended MAC address flapping is to statically add the gateway routers real MAC addresses to the ASA MAC address table. Without these entries, if the gateway at site 1 communicates with the gateway at site 2, that traffic might pass through the ASA and attempt to reach site 2 from the inside interface and cause problems. The data VLANs are extended between the sites using Overlay Transport Virtualization (OTV) (or something similar). You must add filters to prevent traffic from

traversing the DCI to the other site when the traffic is destined for the gateway router. If the gateway router at one site becomes unreachable, you must remove the filters so traffic can be sent to the other site's gateway router.



Reference for Clustering

This section includes more information about how clustering operates.

ASA Features and Clustering

Some ASA features are not supported with ASA clustering, and some are only supported on the control node. 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.

- Unified Communication features that rely on TLS Proxy
- Remote access VPN (SSL VPN and IPsec VPN)
- Virtual Tunnel Interfaces (VTIs)
- IS-IS routing
- The following application inspections:
  - CTIQBE

  - H323, H225, and RAS
  - IPsec passthrough
  - MGCP
  - MMP
  - RTSP
  - SCCP (Skinny)
  - WAAS
  - WCCP
- Botnet Traffic Filter
- Auto Update Server
- DHCP client, server, and proxy. DHCP relay is supported.
- VPN load balancing
- Failover
- Integrated Routing and Bridging
- Dead Connection Detection (DCD)
- FIPS 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.

- The following application inspections:
  - DCERPC
  - ESMTP
  - IM
  - NetBIOS
  - PPTP

- RADIUS
- RSH
- SNMP
- SQLNET
- SUNRPC
- TFTP
- XDMCP

- Static route monitoring
- Authentication and Authorization for network access. Accounting is decentralized.
- Filtering Services
- Site-to-site VPN

  In centralized mode, VPN connections are established with the control node of the cluster only. This is the default mode for VPN clustering. Site-to-site VPN can also be deployed in Distributed VPN Mode, where S2S IKEv2 VPN connections are distributed across nodes.
- 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

Features Applied to Individual Units

These features are applied to each ASA node, instead of the cluster as a whole or to the control node.

- QoS—The QoS policy is synced across the cluster as part of configuration replication. However, the policy is enforced on each node independently. For example, if you configure policing on output, then the conform rate and conform burst values are enforced on traffic exiting a particular ASA. In a cluster with 3 nodes and with traffic evenly distributed, the conform rate actually becomes 3 times the rate for the cluster.
- Threat detection—Threat detection works on each node independently; for example, the top statistics is node-specific. Port scanning detection, for example, does not work because scanning traffic will be load-balanced between all nodes, and one node will not see all traffic.
- Resource management—Resource management in multiple context mode is enforced separately on each node based on local usage.
- LISP traffic—LISP traffic on UDP port 4342 is inspected by each receiving node, but is not assigned a director. Each node adds to the EID table that is shared across the cluster, but the LISP traffic itself does not participate in cluster state sharing.

AAA for Network Access and Clustering

AAA for network access consists of three components: authentication, authorization, and accounting. Authentication and authorization are implemented as centralized features on the clustering control node with

replication of the data structures to the cluster data nodes. If a control node is elected, the new control node will have all the information it needs to continue uninterrupted operation of the established authenticated users and their associated authorizations. Idle and absolute timeouts for user authentications are preserved when a control node change occurs.

Accounting is implemented as a distributed feature in a cluster. Accounting is done on a per-flow basis, so the cluster node owning a flow will send accounting start and stop messages to the AAA server when accounting is configured for a flow.

Connection Settings

Connection limits are enforced cluster-wide (see **Configuration** > **Firewall** > **Service Policy** page). 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.
- If you use AAA for FTP access, then the control channel flow is centralized on the control node.

ICMP Inspection

The flow of ICMP and ICMP error packets through the cluster varies depending on whether ICMP/ICMP error inspection is enabled. Without ICMP inspection, ICMP is a one-direction flow, and there is no director flow support. With ICMP inspection, the ICMP flow becomes two-directional and is backed up by a director/backup flow. One difference for an inspected ICMP flow is in the director handling of a forwarded packet: the director will forward the ICMP echo reply packet to the flow owner instead of returning the packet to the forwarder.

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 ASAs 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 ASA 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.
- Per-session PAT feature—Although not exclusive to clustering, the per-session PAT feature improves the scalability of PAT and, for clustering, allows each data node to own PAT connections; by contrast, multi-session PAT connections have to be forwarded to and owned by the control node. By default, all TCP traffic and UDP DNS traffic use a per-session PAT xlate, whereas ICMP and all other UDP traffic uses multi-session. You can configure per-session NAT rules to change these defaults for TCP and UDP, but you cannot configure per-session PAT for ICMP. For example, with an increasing use of Quic protocol over UDP/443 as a best performance alternative compared to HTTPS TLS over TCP/443, you should enable per-session PAT for UDP/443. For traffic that benefits from multi-session PAT, such as H.323, SIP, or Skinny, you can disable per-session PAT for the associated TCP ports (the UDP ports for those

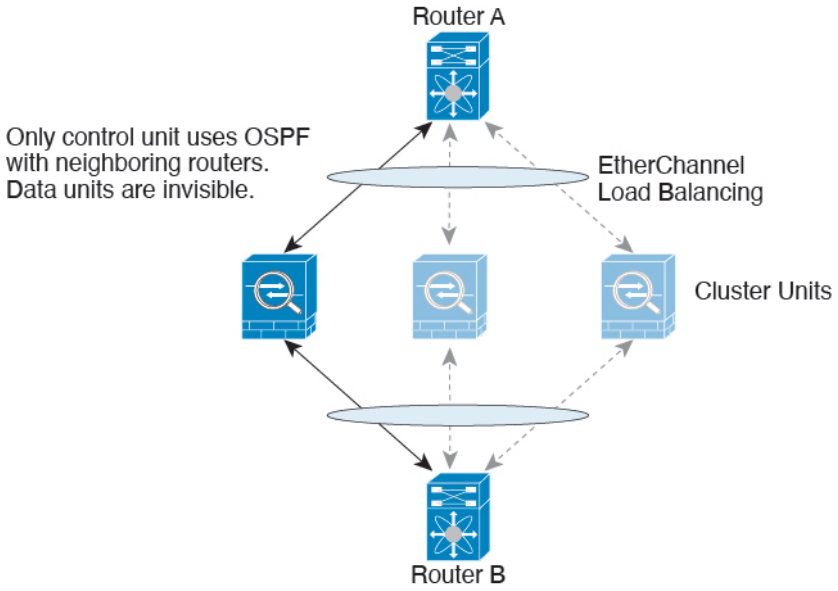
H.323 and SIP are already multi-session by default). For more information about per-session PAT, see the firewall configuration guide.

- No static PAT for the following inspections—
  - FTP
  - PPTP
  - 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 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 68: 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.

SCTP and Clustering

An SCTP association can be created on any node (due to load balancing); its multi-homing connections must reside on the same node.

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.

TLS Proxy configuration is not supported.

SNMP and Clustering

An SNMP agent polls each individual ASA by its Local IP address. You cannot poll consolidated data for the cluster.

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 re-add them on the control node to force the users to replicate to the new node, or directly on the data node.

STUN and Clustering

STUN inspection is supported in failover and cluster modes, as pinholes are replicated. However, the transaction ID is not replicated among nodes. In the case where a node fails after receiving a STUN Request and another node received the STUN Response, the STUN Response will be dropped.

Syslog and NetFlow and Clustering

- Syslog—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.
- NetFlow—Each node in the cluster generates its own NetFlow stream. The NetFlow collector can only treat each ASA as a separate NetFlow exporter.

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 on the Secure Firewall eXtensible Operating System (FXOS) Chassis

An ASA FXOS Cluster supports one of two mutually exclusive modes for S2S VPN, centralized or distributed:

- Centralized VPN Mode. The default mode. In centralized mode, VPN connections are established with the control unit of the cluster only.

  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 connected users 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.

- Distributed VPN Mode. In this mode, S2S IPsec IKEv2 VPN connections are distributed across members of an ASA cluster providing scalability. Distributing VPN connections across the members of a cluster allows both the capacity and throughput of the cluster to be fully utilized, significantly scaling VPN support beyond Centralized VPN capabilities.



Note Centralized VPN clustering mode supports S2S IKEv1 and S2S IKEv2.

Distributed VPN clustering mode supports S2S IKEv2 only.

Distributed VPN clustering mode is supported on the Firepower 9300 only.

Remote access VPN is not supported in centralized or distributed VPN clustering mode.

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 ASA application periodically (every second). If the ASA is up and cannot communicate with the Firepower 4100/9300 chassis supervisor for 3 seconds, the ASA 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 ASA. If the ASA 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 492 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 ASA 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 (depending on your minimum port bundling setting). 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 ASA removes a node from the cluster depends on whether the node is an established member or is joining the cluster. The ASA 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 ASA 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 ASA 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 ASA automatically tries to rejoin the cluster, depending on the failure event.



Note When the ASA becomes inactive and fails to automatically rejoin the cluster, all data interfaces are shut down; only the management-only interface can send and receive traffic. The management interface remains up using the IP address the node received from the cluster IP pool. However if you reload, and the node is still inactive in the cluster, the management interface is disabled. You must use the console port for any further configuration.

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 ASA automatically tries to rejoin every 5 minutes, indefinitely. This behavior is configurable.
- Failed data interface—The ASA 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 ASA disables clustering. After you resolve the problem with the data interface, you have to manually enable clustering.This behavior is configurable.
- Failed unit—If the unit was removed from the cluster because of a unit health check failure, then rejoining the cluster depends on the source of the failure. For example, a temporary power failure means the unit will rejoin the cluster when it starts up again as long as the cluster control link is up. The unit attempts to rejoin the cluster every 5 seconds.
- Failed Chassis-Application Communication—When the ASA detects that the chassis-application health has recovered, the ASA tries to rejoin the cluster immediately. Alternatively, you can configure the ASA to use the same rejoin settings as for internal errors (below).

- Failed decorator application—The ASA rejoins the cluster when it senses that the decorator application is back up.
- Internal error—Internal failures include: application sync timeout; inconsistent application statuses; and so on. A unit will attempt to rejoin the cluster automatically at the following intervals: 5 minutes, 10 minutes, and then 20 minutes. This behavior is configurable.

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 23: 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	Includes AAA rules (uauth).
IPv6 Neighbor database	Yes	—
Dynamic routing	Yes	—
SNMP Engine ID	**No**	—
Distributed VPN (Site-to-Site) for Firepower 4100/9300	Yes	Backup session becomes the active session, then a new backup session is created.

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.

If you enable director localization for inter-site clustering, then there are two backup owner roles: the local backup and the global backup. The owner always chooses a local backup at the same site as itself (based on site ID). The global backup can be at any site, and might even be the same node as the local backup. The owner sends connection state information to both backups.

If you enable site redundancy, and the backup owner is at the same site as the owner, then an additional backup owner will be chosen from another site to protect flows from a site failure. Chassis backup and site backup are independent, so in some cases a flow will have both a chassis backup and a site backup.

- 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.

  If you enable director localization for inter-site clustering, then there are two director roles: the local director and the global director. The owner always chooses a local director at the same site as itself (based on site ID). The global director can be at any site, and might even be the same node as the local director. If the original owner fails, then the local director chooses a new connection owner at the same site.

  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. If you enable director localization, then the forwarder always queries the local director. The forwarder only queries the global director if the local director does not know the owner, for example, if a cluster member receives packets for a connection that is owned on a different site. 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.

When a connection uses Port Address Translation (PAT), then the PAT type (per-session or multi-session) influences which member of the cluster becomes the owner of a new connection:

- Per-session PAT—The owner is the node that receives the initial packet in the connection.

  By default, TCP and DNS UDP traffic use per-session PAT.

- Multi-session PAT—The owner is always the control node. If a multi-session PAT connection is initially received by a data node, then the data node forwards the connection to the control node.

  By default, UDP (except for DNS UDP) and ICMP traffic use multi-session PAT, so these connections are always owned by the control node.

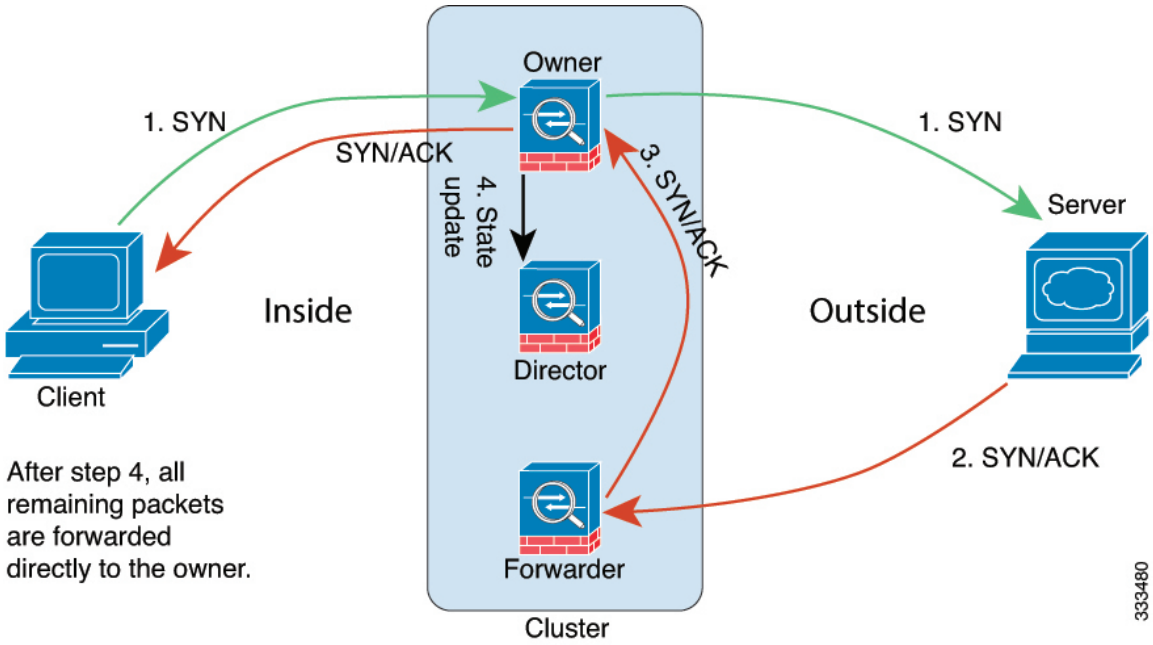
You can change the per-session PAT defaults for TCP and UDP so connections for these protocols are handled per-session or multi-session depending on the configuration. For ICMP, you cannot change from the default multi-session PAT. For more information about per-session PAT, see the firewall configuration guide.

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.

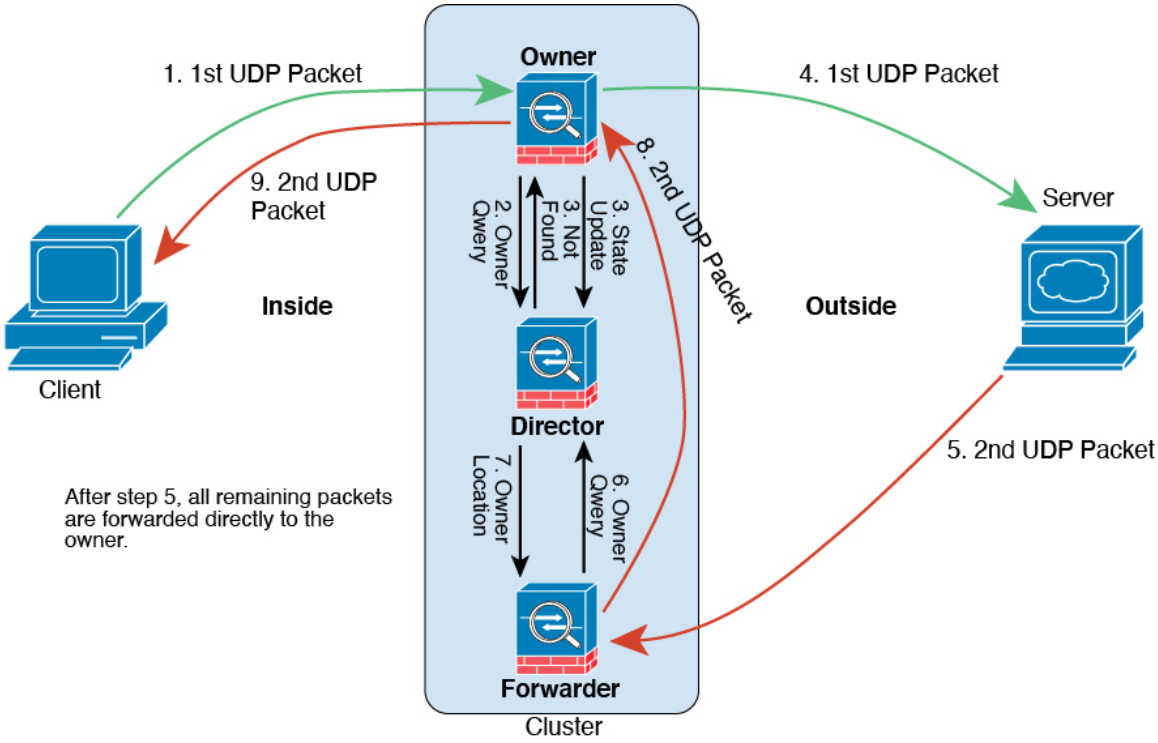


1. The SYN packet originates from the client and is delivered to one ASA (based on the load balancing method), which becomes the owner. The owner creates a flow, encodes owner information into a SYN cookie, and forwards the packet to the server.
2. The SYN-ACK packet originates from the server and is delivered to a different ASA (based on the load balancing method). This ASA 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 69: ICMP and UDP Data Flow*



The first UDP packet originates from the client and is delivered to one ASA (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.

Rebalancing New TCP Connections Across the Cluster

If the load balancing capabilities of the upstream or downstream routers result in unbalanced flow distribution, you can configure new connection rebalancing so nodes with higher new connections per second will redirect new TCP flows to other nodes. No existing flows will be moved to other nodes.

Because this command only rebalances based on connections per second, the total number of established connections on each node is not considered, and the total number of connections may not be equal.

Once a connection is offloaded to a different node, it becomes an asymmetric connection.

Do not configure connection rebalancing for inter-site topologies; you do not want new connections rebalanced to cluster members at a different site.

History for ASA Clustering on the Firepower 4100/9300

Feature Name	Version	Feature Information
Configurable delay to rejoin cluster after chassis heartbeat failure (Firepower 4100/9300)	9.20(2)	By default, if the chassis heartbeat fails and then recovers, the node rejoins the cluster immediately. However, if you configure the **health-check chassis-heartbeat-delay-rejoin** command, it will rejoin according to the settings of the **health-check system auto-rejoin** command. New/Modified screens: **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** > **Auto Rejoin**
Configurable cluster keepalive interval for flow status	9.20(1)	The flow owner sends keepalives (clu_keepalive messages) and updates (clu_update messages) to the director and backup owner to refresh the flow state. You can now set the keepalive interval. The default is 15 seconds, and you can set the interval between 15 and 55 seconds. You may want to set the interval to be longer to reduce the amount of traffic on the cluster control link. New/Modified screens: **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** > **Cluster Configuration**
Removal of biased language	9.19(1)	Commands, command output, and syslog messages that contained the terms "Master" and "Slave" have been changed to "Control" and "Data." New/Modified commands: **cluster control-node**, **enable as-data-node**, **prompt**, **show cluster history**, **show cluster info**
Improved PAT port block allocation for clustering on the Firepower 4100/9300	9.16(1)	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. 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 screens: **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** > **Cluster Configuration** > **Cluster Member Limit** field
show cluster history command improvements	9.16(1)	We have added additional outputs for the **show cluster history** command. New/Modified commands: **show cluster history brief**, **show cluster history latest**, **show cluster history reverse**, **show cluster history time**
Configuration sync to data units in parallel	9.14(1)	The control unit now syncs configuration changes with data units in parallel by default. Formerly, synching occurred sequentially. New/Modified screens: **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** > **Cluster Configuration** > **Enable parallel configuration replicate** check box

Feature Name	Version	Feature Information
Messages for cluster join failure or eviction added to **show cluster history**	9.14(1)	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.	9.13(1)	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. No modified screens.
Monitor the traffic load for a cluster	9.13(1)	You can now monitor the traffic load for cluster members, including total connection count, CPU and memory usage, and buffer drops. If the load is too high, you can choose to manually disable clustering on the unit if the remaining units can handle the load, or adjust the load balancing on the external switch. This feature is enabled by default. New/Modified screens: • **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** > **Cluster Configuration** > **Enable Cluster Load Monitor** check box • **Monitoring** > **ASA Cluster** > **Cluster Load-Monitoring**
Accelerated cluster joining	9.13(1)	When a data unit has the same configuration as the control unit, it will skip syncing the configuration and will join faster. This feature is enabled by default. This feature is configured on each unit, and is not replicated from the control unit to the data unit. **Note** Some configuration commands are not compatible with accelerated cluster joining; if these commands are present on the unit, even if accelerated cluster joining is enabled, configuration syncing will always occur. You must remove the incompatible configuration for accelerated cluster joining to work. Use the **show cluster info unit-join-acceleration incompatible-config** to view incompatible configuration. New/Modified screens: **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** > **Cluster Configuration** > **Enable config sync acceleration** check box

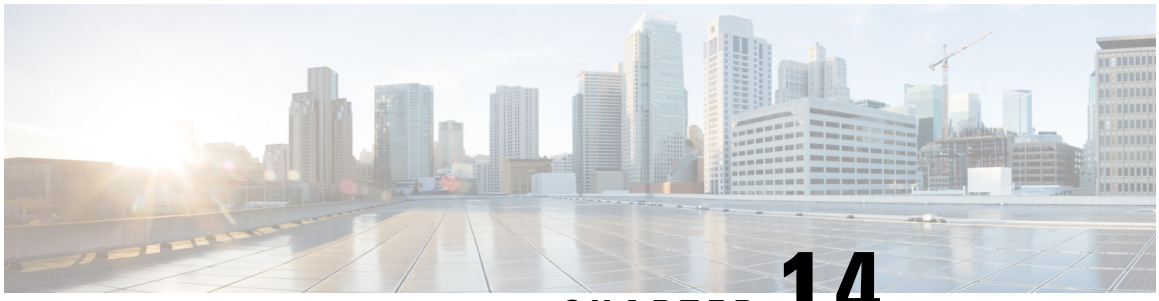
Feature Name	Version	Feature Information
Per-site gratuitous ARP for clustering	9.12(1)	The ASA now generates gratuitous ARP (GARP) packets to keep the switching infrastructure up to date: the highest priority member at each site periodically generates GARP traffic for the global MAC/IP addresses. When using per-site MAC and IP addresses, packets sourced from the cluster use a site-specific MAC address and IP address, while packets received by the cluster use a global MAC address and IP address. If traffic is not generated from the global MAC address periodically, you could experience a MAC address timeout on your switches for the global MAC address. After a timeout, traffic destined for the global MAC address will be flooded across the entire switching infrastructure, which can cause performance and security concerns. GARP is enabled by default when you set the site ID for each unit and the site MAC address for each Spanned EtherChannel. New/Modified screens: **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** > **Cluster Configuration** > **Site Periodic GARP** field
Parallel cluster joining of units per Firepower 9300 chassis	9.10(1)	For the Firepower 9300, this feature ensures that the security modules in a chassis join the cluster simultaneously, so that traffic is evenly distributed between the modules. If a module joins very much in advance of other modules, it can receive more traffic than desired, because the other modules cannot yet share the load. New/modified screens: **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** New/Modified options: **Parallel Join of Units Per Chassis** area
Cluster control link customizable IP Address for the Firepower 4100/9300	9.10(1)	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 chassis manager screens: **Logical Devices** > **Add Device** > **Cluster Information** New/Modified options: **CCL Subnet IP** field
Cluster interface debounce time now applies to interfaces changing from a down state to an up state	9.10(1)	When an interface status update occurs, the ASA waits the number of milliseconds specified in the **health-check monitor-interface debounce-time** command or the ASDM **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** screen before marking the interface as failed and the unit is removed from the cluster. This feature now applies to interfaces changing from a down state to an up state. For example, 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 unit just because another cluster unit was faster at bundling the ports. We did not modify any screens.

Feature Name	Version	Feature Information
Automatically rejoin the cluster after an internal failure	9.9(2)	Formerly, many 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 by default: 5 minutes, 10 minutes, and then 20 minutes. These values are configurable. Internal failures include: application sync timeout; inconsistent application statuses; and so on. New or modified screen: **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** > **Auto Rejoin**
Show transport related statistics for cluster reliable transport protocol messages	9.9(2)	You can now view per-unit cluster reliable transport buffer usage so you can identify packet drop issues when the buffer is full in the control plane. New or modified command: **show cluster info transport cp detail**
cluster remove unit command behavior matches **no enable** behavior	9.9(1)	The **cluster remove unit** command now removes a unit from the cluster until you manually reenable clustering or reload, similar to the **no enable** command. Previously, if you redeployed the bootstrap configuration from FXOS, clustering would be reenabled. Now, the disabled status persists even in the case of a bootstrap configuration redeployment. Reloading the ASA, however, will reenable clustering. New/Modified screen: **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster**
Improved chassis health check failure detection for the chassis	9.9(1)	You can now configure a lower holdtime for the chassis health check: 100 ms. The previous minimum was 300 ms. Note that the minimum combined time (*interval* x *retry-count*) cannot be less than 600 ms. New or modified command: **app-agent heartbeat interval** No ASDM support.
Inter-site redundancy for clustering	9.9(1)	Inter-site redundancy ensures that a backup owner for a traffic flow will always be at the other site from the owner. This feature guards against site failure. New or modified screen: **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster**

Feature Name	Version	Feature Information
Distributed Site-to-Site VPN with clustering on the Firepower 9300	9.9(1)	An ASA cluster on the Firepower 9300 supports Site-to-Site VPN in distributed mode. Distributed mode provides the ability to have many Site-to-Site IPsec IKEv2 VPN connections distributed across members of an ASA cluster, not just on the control unit (as in centralized mode). This significantly scales VPN support beyond Centralized VPN capabilities and provides high availability. Distributed S2S VPN runs on a cluster of up to two chassis, each containing up to three modules (six total cluster members), each module supporting up to 6K active sessions (12K total), for a maximum of approximately 36K active sessions (72K total). New or modified screens: **Monitoring > ASA Cluster > ASA Cluster > VPN Cluster Summary** **Monitoring > VPN > VPN Statistics > Sessions** **Configuration > Device Management > High Availablility and Scalability > ASA Cluster** **Wizards > Site-to-Site** **Monitoring > VPN > VPN Statistics > Sessions** **Monitoring > ASA Cluster > ASA Cluster > VPN Cluster Summary** **Monitoring > ASA Cluster > ASA Cluster > System Resource Graphs > CPU/Memory** **Monitoring > Logging > Real-Time Log Viewer**
Improved cluster unit health-check failure detection	9.8(1)	You can now configure a lower holdtime for the unit health check: .3 seconds minimum. The previous minimum was .8 seconds. This feature changes the unit health check messaging scheme to *heartbeats* in the data plane from *keepalives* in the control plane. Using heartbeats improves the reliability and the responsiveness of clustering by not being susceptible to control plane CPU hogging and scheduling delays. Note that configuring a lower holdtime increases cluster control link messaging activity. We suggest that you analyze your network before you configure a low holdtime; for example, make sure a ping from one unit to another over the cluster control link returns within the *holdtime*/3, because there will be three heartbeat messages during one holdtime interval. If you downgrade your ASA software after setting the hold time to .3 - .7, this setting will revert to the default of 3 seconds because the new setting is unsupported. We modified the following screen: **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster**
Configurable debounce time to mark an interface as failed for the Firepower 4100/9300 chassis	9.8(1)	You can now configure the debounce time before the ASA considers an interface to be failed, and the unit is removed from the cluster. This feature allows 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 ASA waits the number of milliseconds specified before marking the interface as failed and the unit is removed from the cluster. The default debounce time is 500 ms, with a range of 300 ms to 9 seconds. New or modified screen: **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster**

Feature Name	Version	Feature Information
Inter-site clustering improvement for the ASA on the Firepower 4100/9300 chassis	9.7(1)	You can now configure the site ID for each Firepower 4100/9300 chassis when you deploy the ASA cluster. Previously, you had to configure the site ID within the ASA application; this new feature eases initial deployment. Note that you can no longer set the site ID within the ASA configuration. Also, for best compatibility with inter-site clustering, we recommend that you upgrade to ASA 9.7(1) and FXOS 2.1.1, which includes several improvements to stability and performance. We modified the following screen: **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** > **Cluster Configuration**
Director localization: inter-site clustering improvement for data centers	9.7(1)	To improve performance and keep traffic within a site for inter-site clustering for data centers, you can enable director localization. New connections are typically load-balanced and owned by cluster members within a given site. However, the ASA assigns the director role to a member at *any* site. Director localization enables additional director roles: a local director at the same site as the owner, and a global director that can be at any site. Keeping the owner and director at the same site improves performance. Also, if the original owner fails, the local director chooses a new connection owner at the same site. The global director is used if a cluster member receives packets for a connection that is owned on a different site. We modified the following screen: **Configuration** > **Device Management** > **High Availability and Scalability** > **Cluster Configuration**
Support for 16 chassis for the Firepower 4100 series	9.6(2)	You can now add up to 16 chassis to the cluster for the Firepower 4100 series. We did not modify any screens.
Support for the Firepower 4100 series	9.6(1)	With FXOS 1.1.4, the ASA supports inter-chassis clustering on the Firepower 4100 series for up to 6 chassis. We did not modify any screens.
Support for site-specific IP addresses in Routed, Spanned EtherChannel mode	9.6(1)	For inter-site clustering in routed mode with Spanned EtherChannels, you can now configure site-specific IP addresess in addition to site-specific MAC addresses. The addition of site IP addresses allows you to use ARP inspection on the Overlay Transport Virtualization (OTV) devices to prevent ARP responses from the global MAC address from traveling over the Data Center Interconnect (DCI), which can cause routing problems. ARP inspection is required for some switches that cannot use VACLs to filter MAC addresses. We modified the following screen: **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces** > **Add/Edit EtherChannel Interface** > **Advanced**
Inter-chassis clustering for 16 modules, and inter-site clustering for the Firepower 9300 ASA application	9.5(2.1)	With FXOS 1.1.3, you can now enable inter-chassis, and by extension inter-site clustering. You can include up to 16 modules. 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. We did not modify any screens.
Site-specific MAC addresses for inter-site clustering support for Spanned EtherChannel in Routed firewall mode	9.5(2)	You can now use inter-site clustering for Spanned EtherChannels in routed mode. To avoid MAC address flapping, configure a site ID for each cluster member so that a site-specific MAC address for each interface can be shared among a site's units. We modified the following screen: **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** > **Cluster Configuration**

Feature Name	Version	Feature Information
ASA cluster customization of the auto-rejoin behavior when an interface or the cluster control link fails	9.5(2)	You can now customize the auto-rejoin behavior when an interface or the cluster control link fails. We introduced the following screen: **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** > **Auto Rejoin**
The ASA cluster supports GTPv1 and GTPv2	9.5(2)	The ASA cluster now supports GTPv1 and GTPv2 inspection. We did not modify any screens.
Cluster replication delay for TCP connections	9.5(2)	This feature helps eliminate the “unnecessary work” related to short-lived flows by delaying the director/backup flow creation. We introduced the following screen: **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster Replication**
LISP Inspection for Inter-Site Flow Mobility	9.5(2)	Cisco Locator/ID Separation Protocol (LISP) architecture separates the device identity from its location into two different numbering spaces, making server migration transparent to clients. The ASA can inspect LISP traffic for location changes and then use this information for seamless clustering operation; the ASA cluster members inspect LISP traffic passing between the first hop router and the egress tunnel router (ETR) or ingress tunnel router (ITR), and then change the flow owner to be at the new site. We introduced or modified the following screens: **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** > **Cluster Configuration** **Configuration** > **Firewall** > **Objects** > **Inspect Maps** > **LISP** **Configuration** > **Firewall** > **Service Policy Rules** > **Protocol Inspection** **Configuration** > **Firewall** > **Service Policy Rules** > **Cluster** **Monitoring** > **Routing** > **LISP-EID Table**
Carrier Grade NAT enhancements now supported in failover and ASA clustering	9.5(2)	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). This feature is now supported in failover and ASA cluster deployments. We did not modify any screens.
Configurable level for clustering trace entries	9.5(2)	By default, all levels of clustering events are included in the trace buffer, including many low level events. To limit the trace to higher level events, you can set the minimum trace level for the cluster. We did not modify any screens.
Intra-chassis ASA Clustering for the Firepower 9300	9.4(1.150)	You can cluster up to 3 security modules within the Firepower 9300 chassis. All modules in the chassis must belong to the cluster. We introduced the following screen: **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster Replication**



CHAPTER 14

ASA Cluster for the ASA Virtual for the Private Cloud

Clustering lets you group multiple ASA virtual's 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 the ASA virtual clusters using:

- KVM
- VMware

Note Only routed firewall mode is supported.

Note Some features are not supported when using clustering. See Unsupported Features with Clustering, on page 543.

- About ASA Virtual Clustering, on page 507
- Licenses for ASA Virtual Clustering, on page 513
- Requirements and Prerequisites for ASA Virtual Clustering, on page 513
- Guidelines for ASA Virtual Clustering, on page 514
- Configure the ASA Virtual Clustering Using a Day0 Configuration, on page 515
- Configure ASA Virtual Clustering after Deployment, on page 518
- Customize the Clustering Operation, on page 527
- Manage Cluster Nodes, on page 535
- Monitoring the ASA Virtual Cluster, on page 540
- Examples for ASA Virtual Clustering, on page 542
- Reference for Clustering, on page 543
- History for ASA Virtual Clustering, on page 557

About ASA Virtual Clustering

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 ASA virtual send broadcast/multicast messages over the cluster control link.
- Management access to each firewall for configuration and monitoring. The ASA 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.

Cluster Nodes

Cluster nodes work together to accomplish the sharing of the security policy and traffic flows. This section describes the nature of each node role.

Bootstrap Configuration

On each device, you configure a minimal bootstrap configuration including the cluster name, cluster control link interface, and other cluster settings. The first node on which you enable clustering typically becomes the *control* node. When you enable clustering on subsequent nodes, they join the cluster as *data* nodes.

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 in the bootstrap configuration; the priority is set between 1 and 100, where 1 is the highest priority. All other members are data nodes. Typically, when you first create a cluster, the first node you add becomes the control node simply because it is the only node in the cluster so far.

You must perform all configuration (aside from the bootstrap configuration) on the control node only; the configuration is then replicated to the data nodes. In the case of physical assets, such as interfaces, the configuration of the control node is mirrored on all data nodes. For example, if you configure Ethernet 1/2 as the inside interface and Ethernet 1/1 as the outside interface, then these interfaces are also used on the data nodes as inside and outside interfaces.

Some features do not scale in a cluster, and the control node handles all traffic for those features.

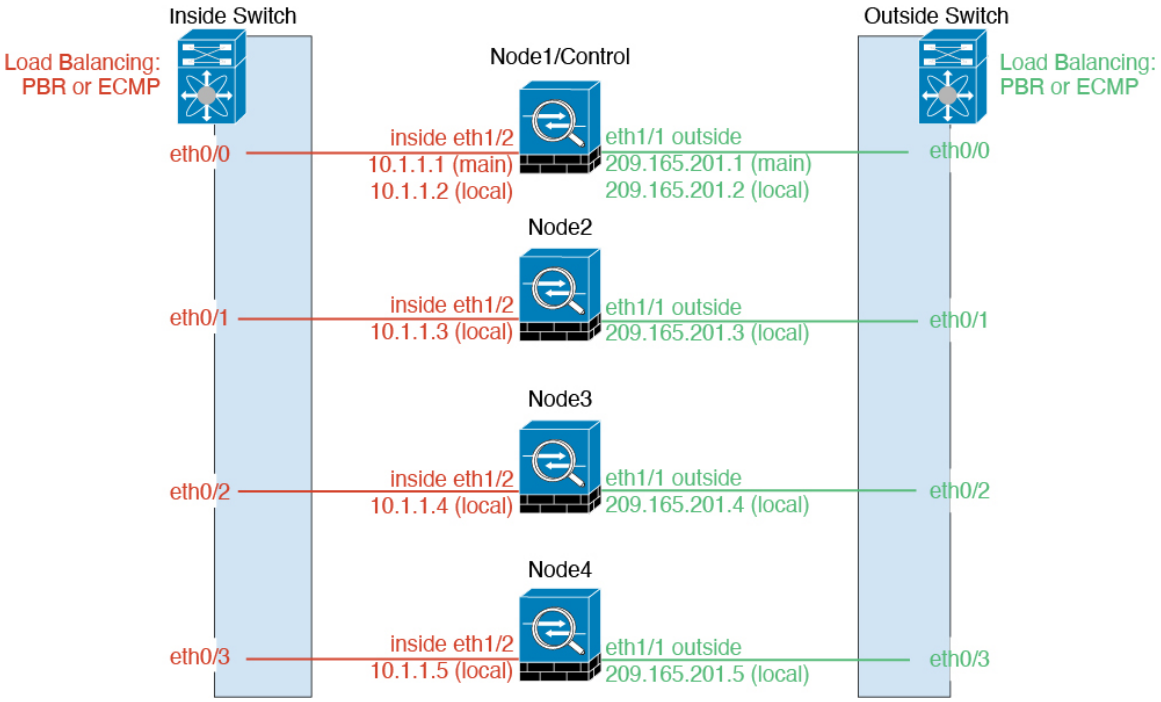
Individual Interfaces

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.

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 ASA 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 ASAs 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 ASA. 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

ASA using ICMP ping. PBR can then enable or disable route maps based on reachability of a particular ASA. 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 ASA 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 ASA failure can cause problems; the route continues to be used, and traffic to the failed ASA 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 ASA 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 VXLAN Interfaces, on page 611.

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 ASA 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 ASA 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.

Cluster Control Link Failure

If the cluster control link line protocol goes down for a unit, then clustering is disabled; data interfaces are shut down. After you fix the cluster control link, you must manually rejoin the cluster by re-enabling clustering.



Note When the ASA virtual becomes inactive, all data interfaces are shut down; only the management-only interface can send and receive traffic. The management interface remains up using the IP address the unit received from DHCP or the cluster IP pool. If you use a cluster IP pool, if you reload and the unit is still inactive in the cluster, then the management interface is not accessible (because it then uses the Main IP address, which is the same as the control node). You must use the console port (if available) for any further configuration.

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.

ASA Virtual Cluster Management

One of the benefits of using ASA virtual clustering is the ease of management. This section describes how to manage the cluster.

Management Network

We recommend connecting all nodes to a single management network. This network is separate from the cluster control link.

Management Interface

Use the Management 0/0 interface for management.

Note You cannot enable dynamic routing for the management interface. You must use a static route.

You can use either static addressing or DHCP for the management IP address.

If you use static addressing, you can use a Main cluster IP address that is a fixed address for the cluster that always belongs to the current control node. For each interface, you also configure a range of addresses so that each node, including the current control node, can use a Local address from the range. The Main cluster IP address provides consistent management access to an address; when a control node changes, the Main cluster IP address moves to the new control node, so management of the cluster continues seamlessly. The Local IP address is used for routing, and is also useful for troubleshooting. For example, you can manage the cluster by connecting to the Main cluster IP address, which is always attached to the current control node. To manage an individual member, you can connect to the Local IP address. For outbound management traffic such as TFTP or syslog, each node, including the control node, uses the Local IP address to connect to the server.

If you use DHCP, you do not use a pool of Local addresses or have a Main cluster IP address.

Note To-the-box traffic needs to be directed to the node's management IP address; to-the-box traffic is not forwarded over the cluster control link to any other node.

Control Node Management Vs. Data Node Management

All management and monitoring can take place on the control node. From the control node, you can check runtime statistics, resource usage, or other monitoring information of all nodes. You can also issue a command to all nodes in the cluster, and replicate the console messages from data nodes to the control node.

You can monitor data nodes directly if desired. Although also available from the control node, you can perform file management on data nodes (including backing up the configuration and updating images). The following functions are not available from the control node:

- Monitoring per-node cluster-specific statistics.
- Syslog monitoring per node (except for syslogs sent to the console when console replication is enabled).
- SNMP
- NetFlow

Crypto Key Replication

When you create a crypto key on the control node, the key is replicated to all data nodes. If you have an SSH session to the Main cluster IP address, you will be disconnected if the control node fails. The new control node uses the same key for SSH connections, so that you do not need to update the cached SSH host key when you reconnect to the new control node.

ASDM Connection Certificate IP Address Mismatch

By default, a self-signed certificate is used for the ASDM connection based on the Local IP address. If you connect to the Main cluster IP address using ASDM, then a warning message about a mismatched IP address might appear because the certificate uses the Local IP address, and not the Main cluster IP address. You can ignore the message and establish the ASDM connection. However, to avoid this type of warning, you can

enroll a certificate that contains the Main cluster IP address and all the Local IP addresses from the IP address pool. You can then use this certificate for each cluster member. See https://www.cisco.com/c/en/us/td/docs/security/asdm/identity-cert/cert-install.html for more information.

Inter-Site Clustering

For inter-site installations, you can take advantage of ASA virtual clustering as long as you follow the recommended guidelines.

You can configure each cluster chassis to belong to a separate site ID. Site IDs are used to enable flow mobility using LISP inspection, director localization to improve performance and reduce round-trip time latency for inter-site clustering for data centers, and site redundancy for connections where a backup owner of a traffic flow is always at a different site from the owner.

See the following sections for more information about inter-site clustering:

- Sizing the Data Center Interconnect—Requirements and Prerequisites for ASA Virtual Clustering, on page 513
- Inter-Site Guidelines—Guidelines for ASA Virtual Clustering, on page 514
- Configure Cluster Flow Mobility—Configure Cluster Flow Mobility, on page 532
- Enable Director Localization—Configure Basic ASA Cluster Parameters, on page 527
- Enable Site Redundancy—Configure Basic ASA Cluster Parameters, on page 527
- Inter-Site Examples—Individual Interface Routed Mode North-South Inter-Site Example, on page 542

Licenses for ASA Virtual Clustering

Each cluster node requires the same model license. We recommend using the same number of CPUs and memory for all nodes, or else peformance 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.

Note If you deregister the ASA virtual so that it is unlicensed, then it will revert to a severely rate-limited state if you reload the ASA virtual. An unlicensed, low performing cluster node will impact the performance of the entire cluster negatively. Be sure to keep all cluster nodes licensed, or remove any unlicensed nodes.

Requirements and Prerequisites for ASA Virtual Clustering

Model Requirements

- ASAv30, ASAv50, ASAv100
- The following private cloud services:
  - KVM

- VMware

- A maximum of 16 nodes in a cluster on *two* hosts in a 2x8 deployment configuration. We recommend you to deploy a maximum of *eight* ASAvs on each of the *two* hosts (2x8), which results in a cluster of 16 nodes.

ASA Virtual Platform and Software Requirements

All nodes in a cluster:

- Must be the same model. We recommend using the same number of CPUs and memory for all nodes, or else peformance will be limited on all nodes to match the least capable node.
- Must run the identical software except at the time of an image upgrade. Hitless upgrade is supported. Mismatched software versions can lead to poor performance, so be sure to upgrade all nodes in the same maintenance window.
- New cluster members must use the same SSL encryption setting (the **ssl encryption** command) as the control node for initial cluster control link communication before configuration replication.

Guidelines for ASA Virtual Clustering

Failover

Failover is not supported with clustering.

IPv6

The cluster control link is only supported using IPv4.

Additional Guidelines

- When significant topology changes occur (such as enabling or disabling an interface on the ASA 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 nodes, 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.
- We do not support VXLANs for data interfaces; only the cluster control link supports VXLAN.
- It takes time to replicate changes to all the nodes in a cluster. If you make a large change, for example, adding an access control rule that uses object groups (which, when deployed, are broken out into multiple rules), the time needed to complete the change can exceed the timeout for the cluster nodes to respond with a success message. If this happens, you might see a "failed to replicate command" message. You can ignore the message.

Defaults for ASA Virtual Clustering

- 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 rebalancing is disabled by default. If you enable connection rebalancing, the default time between load information exchanges is 5 seconds.
- Connection replication delay of 5 seconds is enabled by default for HTTP traffic.

Configure the ASA Virtual Clustering Using a Day0 Configuration

Control Node Day0 Configuration

The following Day0 configuration for the control node includes the bootstrap configuration followed by interface configuration that will be replicated to the data nodes. Bold text shows the values you need to change for the data node Day0 configuration.



Note This configuration only includes the cluster-centric configuration. Your Day0 configuration should also include other settings for licensing, SSH access, ASDM access and more. See the getting started guide for more information about Day0 configurations.

```
!BOOTSTRAP
! Cluster interface mode
cluster interface mode individual
!
! VXLAN peer group
object-group network cluster-peers
network-object host 10.6.6.51
network-object host 10.6.6.52
network-object host 10.6.6.53
network-object host 10.6.6.54
!
! Alternate object group representation
! object-network xyz
! range 10.6.6.51 10.6.6.54
! object-group network cluster-peers
! network-object object xyz
!
! Cluster control link physical interface (VXLAN tunnel endpoint (VTEP) src interface)
interface gigabitethernet 0/7
description CCL VTEP src ifc
nve-only cluster
nameif ccl
security-level 0
ip address 10.6.6.51 255.255.255.0
no shutdown
!
! VXLAN Network Identifier (VNI) interface
```

```
interface vni1
segment-id 1
vtep-nve 1
!
! Set the CCL MTU
mtu ccl 1654
!
! Network Virtualization Endpoint (NVE) association with VTEP src interface
nve 1
encapsulation vxlan
source-interface ccl
peer-group cluster-peers
!
! Management Interface Using DHCP
interface management 0/0
nameif management
ip address dhcp setroute
no shutdown
!
! Alternate Management Using Static IP
! ip local pool mgmt_pool 10.1.1.1 10.10.10.4
! interface management 0/0
! nameif management
! ip address 10.1.1.25 255.255.255.0 cluster-pool mgmt_pool
! no shutdown
!
! Cluster Config
cluster group cluster1
local-unit A
cluster-interface vni1 ip 10.2.2.1 255.255.255.0
priority 1
enable noconfirm
!
! INTERFACES
!
ip local pool inside_pool 10.10.10.11 10.10.10.14
ip local pool outside_pool 10.11.11.11 10.11.11.14
!
interface GigabitEthernet0/1
nameif inside
security-level 100
ip address 10.10.10.10 255.255.255.0 cluster-pool inside_pool
!
interface GigabitEthernet0/0
nameif outside
security-level 0
ip address 10.11.11.10 255.255.255.0 cluster-pool outside_pool
!
!JUMBO FRAME RESERVATION for CCL MTU
jumbo-frame reservation
```

Data Node Day0 Configuration

The following Day0 configuration for the data node includes only the bootstrap configuration. Bold text shows the values you need to change from the control node Day0 configuration.



Note This configuration only includes the cluster-centric configuration. Your Day0 configuration should also include other settings for licensing, SSH access, ASDM access and more. See the getting started guide for more information about Day0 configurations.

```
!BOOTSTRAP
! Cluster interface mode
cluster interface mode individual
!
! VXLAN peer group
object-group network cluster-peers
network-object host 10.6.6.51
network-object host 10.6.6.52
network-object host 10.6.6.53
network-object host 10.6.6.54
!
! Alternate object group representation
! object-network xyz
! range 10.6.6.51 10.6.6.54
! object-group network cluster-peers
! network-object object xyz
!
! Cluster control link physical interface (VXLAN tunnel endpoint (VTEP) src interface)
interface gigabitethernet 0/7
description CCL VTEP src ifc
nve-only cluster
nameif ccl
security-level 0
ip address 10.6.6.52 255.255.255.0
no shutdown
!
! VXLAN Network Identifier (VNI) interface
interface vni1
segment-id 1
vtep-nve 1
!
! Set the CCL MTU
mtu ccl 1654
!
! Network Virtualization Endpoint (NVE) association with VTEP src interface
nve 1
encapsulation vxlan
source-interface ccl
peer-group cluster-peers
!
! Management Interface Using DHCP
interface management 0/0
nameif management
ip address dhcp setroute
no shutdown
!
! Alternate Management Using Static IP
! ip local pool mgmt_pool 10.1.1.1 10.10.10.4
! interface management 0/0
! nameif management
! ip address 10.1.1.25 255.255.255.0 cluster-pool mgmt_pool
! no shutdown
!
! Cluster Config
cluster group cluster1
local-unit B
cluster-interface vni1 ip 10.2.2.2 255.255.255.0
priority 2
enable noconfirm
!
! INTERFACES
!
ip local pool inside_pool 10.10.10.11 10.10.10.14
```

```
ip local pool outside_pool 10.11.11.11 10.11.11.14
!
interface GigabitEthernet0/1
nameif inside
security-level 100
ip address 10.10.10.10 255.255.255.0 cluster-pool inside_pool
!
interface GigabitEthernet0/0
nameif outside
security-level 0
ip address 10.11.11.10 255.255.255.0 cluster-pool outside_pool
!
!JUMBO FRAME RESERVATION for CCL MTU
jumbo-frame reservation
```

Configure ASA Virtual Clustering after Deployment

To configure clustering after you deploy your ASA virtuals, perform the following tasks.

Back Up Your Configurations (Recommended)

When you enable clustering on a data unit, the current configuration is replaced with one synced from the active unit. If you ever want to leave the cluster entirely, it may be useful to have a backup configuration with a usable management interface configuration.

Before you begin

Perform a backup on each unit.

Procedure

Step 1 Choose **Tools** > **Backup Configurations**.

Step 2 Back up at least the running configuration. See Back Up and Restore Configurations or Other Files, on page 1086 for a detailed procedure.

Configure Interface Settings

Configure the cluster interface mode as well as interfaces on the control node. The interface configuration will be replicated to data nodes when they join the cluster. Note that configuration of the cluster control link is covered in the bootstrap configuration procedure.

Configure the Cluster Interface Mode on the Control Node

Before you enable clustering, you need to convert the firewall to use Individual interfaces. Because clustering limits the types of interfaces you can use, this process lets you check your existing configuration for incompatible interfaces and then prevents you from configuring any unsupported interfaces.

Note If you do not add data nodes from the control node, you must set the interface mode manually on all nodes according to this section, not just the control node; if you add data nodes from the control node, ASDM sets the interface mode automatically on the data node.

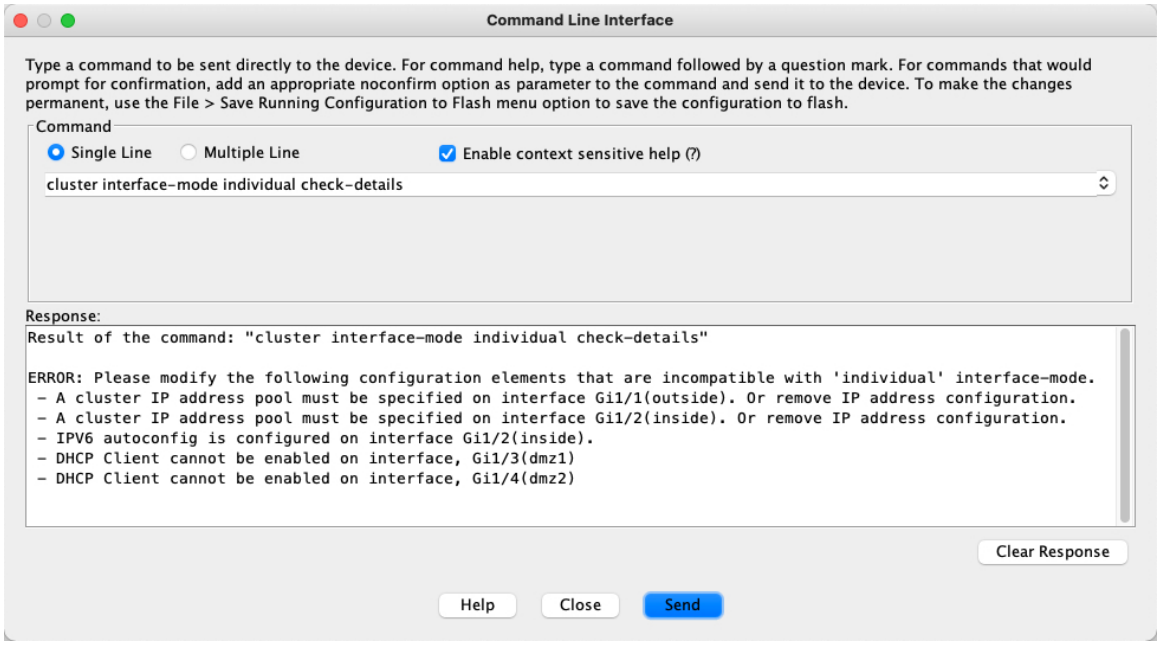
Procedure

Step 1 In ASDM on the control node, choose **Tools** > **Command Line Interface**.Show any incompatible configuration so that you can force the interface mode and fix your configuration later; the mode is not changed with this command:

cluster interface-mode individual check-details

Example:

Figure 70: Command Line Interface Output



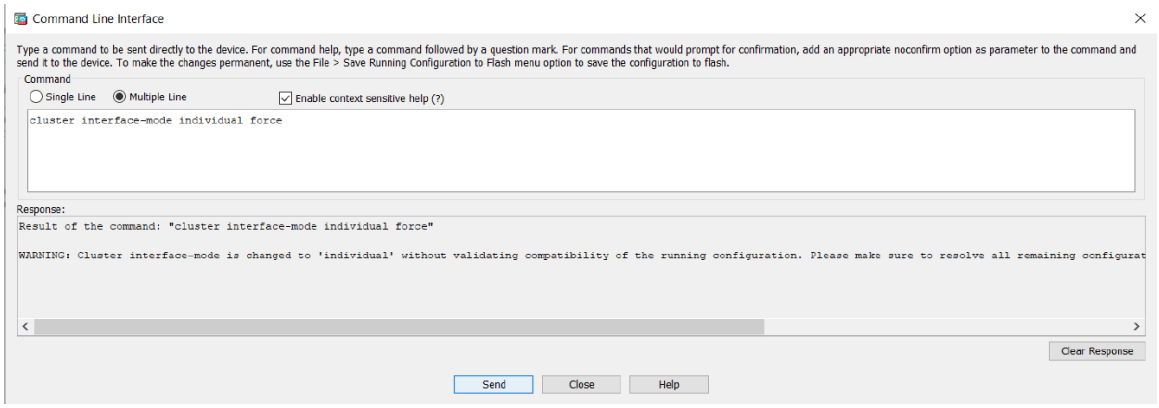
Caution
After you set the interface mode, you can continue to connect to the interface; however, if you reload the ASA before you configure your management interface to comply with clustering requirements (for example, adding a cluster IP pool or getting the IP address from DHCP), you will not be able to reconnect because cluster-incompatible interface configuration is removed. In that case, you will have to connect to the console port, if available, to fix the interface configuration.

Step 2 Set the interface mode for clustering:

cluster interface-mode individual force

Example:

Figure 71: Set the Interface Mode



There is no default setting; you must explicitly choose the mode. If you have not set the mode, you cannot enable clustering.

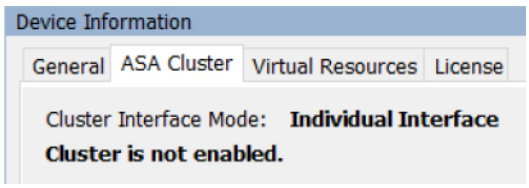
The **force** option changes the mode without checking your configuration for incompatible settings. You need to manually fix any configuration issues after you change the mode. Because any interface configuration can only be fixed after you set the mode, we recommend using the **force** option so that you can at least start from the existing configuration. You can re-run the **check-details** option after you set the mode for more guidance.

Without the **force** option, if there is any incompatible configuration, you are prompted to clear your configuration and reload, thus requiring you to connect to the console port (if available) to reconfigure your management access. If your configuration is compatible (rare), the mode is changed and the configuration is preserved. If you do not want to clear your configuration, you can exit the command by typing **n**.

To remove the interface mode, enter the **no cluster interface-mode** command.

Step 3 Quit ASDM and reload. ASDM needs to be restarted to correctly account for the cluster interface mode. After you reload, you see the ASA Cluster tab on the home page:

Figure 72: ASDM Needs Updating



Configure the Cluster Control Link on the Control Node

Configure a VXLAN interface for the cluster control link interface before you run the wizard. For more information about VXLAN and the cluster control link, see Cluster Control Link, on page 510.

Before you begin

Enable jumbo frame reservation for use with the cluster control link, so you can set the cluster control link MTU to the recommended value. Enabling jumbo frames causes the ASA to reload. See the **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces** screen.



Note You must enable jumbo frame reservation on each node separately.

Procedure

Step 1 Identify the VXLAN tunnel endpoint (VTEP) peer IP addresses in a network object group.

See the **Configuration** > **Firewall** > **Objects** > **Network Objects/Groups** page, and the "Objects for Access Control" chapter in the ASA firewall configuration guide for more information about network object groups.

The underlying IP network between VTEPs is independent of the cluster control link network that the VXLAN Network Identifier (VNI) interfaces use. 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.

Step 2 Configure the VTEP source interface.

The VTEP source interface is a regular ASA 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.

a) Choose **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces**, and edit the interface you want to use for the VTEP source interface.
b) Configure the **Interface Name**.
c) Check the **VTEP Source Interface (cluster)** check box.
d) Check **Enable Interface**.
e) Configure a static IPv4 address.

   The IP address should be included as one of the peers in the network object group.

f) Click the **Advanced** tab, and set the **MTU** to be at least 154 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 (100 bytes) and VXLAN overhead (54 bytes). Set the MTU between 1554 and 9198 bytes. The default MTU is 1554 bytes. We suggest setting the cluster control link MTU to 1654 when data interfaces are set to 1500; this value requires jumbo frame reservation, which required a reload.

   For example, when using jumbo frames, 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.

g) Click **OK**.

Step 3 Associate the VTEP source interface with the Network Virtualization Endpoint (NVE) instance.

a) Choose **Configuration** > **Device Setup** > **Interface Settings** > **VXLAN**.
b) (Optional) Enter a **VXLAN Destination Port** value if you want to change from the default 4789.
c) Check the **Enable Network Virtualization Endpoint encapsulation using VXLAN** check box.
d) Choose the **VTEP Tunnel Interface** from the drop-down list.
e) Check the **Configure Packet Recipient** check box, click the **Peer Group** radio button, and choose the peer group you created.
f) Click **Apply**.

Step 4 Create the 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.

a) Choose **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces**, and click **Add** > **VNI Interface**.

b) Enter the **VNI ID**, between 1 and 10000.

This ID is just an internal interface identifier.

c) Enter the **VNI Segment ID**, between 1 and 16777215.

The segment ID is used for VXLAN tagging.

d) Check the **NVE Mapped to VTEP Interface** check box.

This setting associates the VNI interface with the VTEP source interface.

e) Click **OK**, and then **Apply**.

Configure Individual Interfaces

You must modify any interface that is currently configured with an IP address to be cluster-ready before you enable clustering. At a minimum, you may need to modify the management interface to which ASDM is currently connected when you use a static IP address for management. For other interfaces, you can configure them before or after you enable clustering; we recommend pre-configuring all of your interfaces so that the complete configuration is synced to new cluster nodes.

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.

For the Management interface, you can configure an IP address pool or you can use DHCP; only the Management interface supports getting an address from DHCP. To use DHCP, do not use this procedure; instead configure it as usual (see Configure General Routed Mode Interface Parameters, on page 637).

Before you begin

- (Optional) Configure subinterfaces.
- For the management interface, you can use a static address or you can use DHCP. If you are using static IP addresses and connecting remotely to the management interface using ASDM, the current IP address of prospective data nodes are for temporary use.
  - Each member will be assigned an IP address from the cluster IP pool defined on the control node.
  - The cluster IP pool cannot include addresses already in use on the network, including prospective secondary IP addresses.

    For example:

    1. You configure the control node to use 10.1.1.1.
    2. Other nodes use 10.1.1.2, 10.1.1.3, and 10.1.1.4.

3. When you configure the cluster IP pool on the control node, you cannot include the .2, .3, or .4 addresses in the pool, because they are in use.
4. Instead, you need to use other IP addresses on the network, such as .5, .6, .7, and .8.

Note The pool needs as many addresses as there are members of the cluster, including the control node; the original .1 address is the main cluster IP address that belongs to the current control node.

5. After you join the cluster, the old, temporary addresses are relinquished and can be used elsewhere.

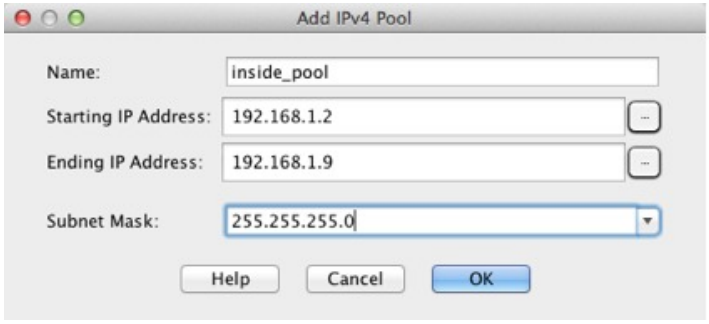
Procedure

Step 1 Choose the **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces** pane.

Step 2 Choose the interface row, and click **Edit**. Choose **Use Static IP**. DHCP and PPPoE are not supported.

Step 3 To add the IPv4 cluster IP pool, MAC address pool, and site-specific MAC addresses, click the **Advanced** tab and set **ASA Cluster** area parameters.

a) Create a cluster IP pool by clicking the **...** button next to the **IP Address Pool** field. The valid range shown is determined by the Main IP address you set on the General tab.
b) Click **Add**.
c) Configure a range of addresses that does not include the Main cluster IP address, and that does not include any addresses currently in-use on your network. You should make the range large enough for the size of the cluster, for example, 8 addresses.



d) Click **OK** to create the new pool.
e) Select the new pool you created, and click **Assign**, and then click **OK**.

The pool name appears in the **IP Address Pool** field.

f) (Optional) (Optional) Configure a **MAC Address Pool** if you want to manually configure MAC addresses.

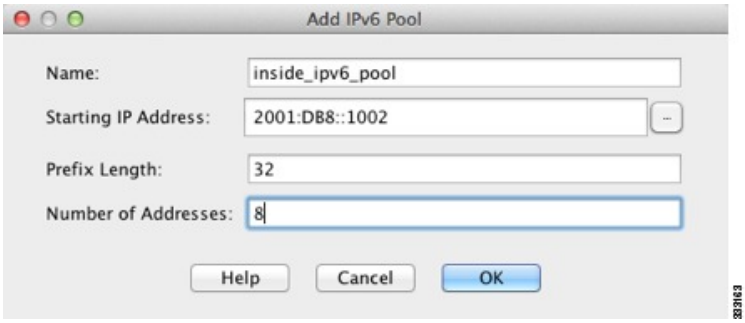
Step 4 To configure an IPv6 address, click the **IPv6** tab.

a) Check the **Enable IPv6** check box.
b) In the **Interface IPv6 Addresses** area, click **Add**.

The **Enable address autoconfiguration** option is not supported. Manually configuring the link-local address is also not supported.

The **Add IPv6 Address for Interface** dialog box appears.

c) In the **Address/Prefix Length** field, enter the global IPv6 address and the IPv6 prefix length. For example, 2001:0DB8::BA98:0:3210/48.
d) Click the **...** button to configure the cluster IP pool.
e) Click **Add**.



f) Configure the starting IP address (network prefix), prefix length, and number of addresses in the pool.
g) Click **OK** to create the new pool.
h) Select the new pool you created, and click **Assign**, and then click **OK**.

The pool appears in the **ASA Cluster IP Pool** field.

i) Click **OK**.

Step 5 Click **OK** to return to the Interfaces pane.

Step 6 Click **Apply**.

Create or Join a Cluster Using the High Availability Wizard

Each node in the cluster requires a bootstrap configuration to join the cluster. Run the High Availability and Scalability wizard on one node (that will become the control node) to create the cluster, and then add data nodes to it.

Before you begin

- The VXLAN VTEP source interface you intend to use for the cluster control link interface must be in an up state on the connected switch.
- When you add a node to a running cluster, you may see temporary, limited packet/connection drops; this is expected behavior.

Procedure

Step 1 Choose **Wizards** > **High Availability and Scalability Wizard**. See select wizard guidelines in the following steps.

Step 2 On the ASA Cluster Configuration screen, configure bootstrap settings including:

- **Member Priority**—Sets the priority of this node for control node elections, between 1 and 100, where 1 is the highest priority.
- **Site Index**—If you use inter-site clustering, set the site ID for this node so it uses a site-specific MAC address, between 1 and 8.
- (Optional) **Shared Key**—Sets an encryption 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 encryption key. This parameter does not affect datapath traffic, including connection state update and forwarded packets, which are always sent in the clear. You must configure this parameter if you also enable the password encryption service.
- (Optional) **Enable connection rebalancing for TCP traffic across all the ASAs in the cluster**—Enables connection rebalancing. This parameter is disabled by default. If enabled, ASAs in a cluster exchange load information periodically, and offload new connections from more loaded devices to less loaded devices. The frequency, between 1 and 360 seconds, specifies how often the load information is exchanged. This parameter is not part of the bootstrap configuration, and is replicated from the control node to the data nodes.

  Note
  Do not configure connection rebalancing for inter-site topologies; you do not want connections rebalanced to cluster members at a different site.

- (Optional) **Enable health monitoring of this device within the cluster**—Enables the cluster node health check feature. To determine node health, the ASA 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 holdtime period, the peer node is considered unresponsive or dead.

  Note
  When any topology changes occur (such as adding or removing a data interface, enabling or disabling an interface on the ASA or the switch) you must disable the health check 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 health check.

  - **Time to Wait Before Device Considered Failed**—This value determines the amount of time between node keepalive status messages, between .3 and 45 seconds; The default is 3 seconds.

- (Optional) **Replicate console output**—Enables console replication from data nodes to the control node. This feature is disabled by default. The ASA may print out some messages directly to the console for certain critical events. If you enable console replication, data nodes send the console messages to the control node so that you only need to monitor one console port for the cluster. This parameter is not part of the bootstrap configuration, and is replicated from the control node to the data nodes.
- **Cluster Control Link**—Specifies the cluster control link interface.
  - **MTU**—Specify the maximum transmission unit for the VTEP source interface to be at least 154 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 (100 bytes) and VXLAN overhead (54 bytes). Set the MTU between 1554 and 9198 bytes. The default MTU is 1554 bytes. We suggest setting the cluster control link MTU to 1654 when data interfaces are set to 1500; this value requires jumbo frame reservation. For example, when using jumbo frames, 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. This parameter

is not part of the bootstrap configuration, and is replicated from the control node to the data nodes. **Note:** If you have not pre-enabled jumbo frame reservation, you should quit the wizard, enable jumbo frames, and then restart this procedure.

Step 3 On the **Interfaces for Health Monitoring** screen, you can exempt some interfaces from monitoring for failure. You might want to disable health monitoring of non-essential interfaces, for example, the management interface.

Note
When any topology changes occur (such as adding or removing a data interface, enabling or disabling an interface on the ASA or the switch) you must disable the health check 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 health check.

Step 4 On the **Interface Auto Rejoin settings** screen, customize the auto-rejoin settings in case of an interface or cluster control link failure. For each type, you can set the following:

- **Maximum Rejoin Attempts**—Define the number of attempts at rejoining the cluster by setting **Unlimited** or a value between 0 and 65535. **0** disables auto-rejoining. The default value is **Unlimited** for the cluster-interface and **3** for the data-interface.
- **Rejoin Interval**—Define the interval duration in minutes between rejoin attempts by setting the interval 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**—Define if the interval duration increases by setting the interval variation 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.

Step 5 Click **Finish**.

Step 6 The ASA scans the running configuration for incompatible commands for features that are not supported with clustering, including commands that may be present in the default configuration. Click **OK** to delete the incompatible commands. If you click **Cancel**, then clustering is not enabled.

After a period of time while ASDM enables clustering and reconnects to the ASA, the Information screen appears confirming that the ASA was added to the cluster.

Note
In some cases, there might be an error when joining the cluster after you finish the wizard. If ASDM was disconnected, ASDM will not receive any subsequent errors from the ASA. If clustering remains disabled after you reconnect ASDM, you should connect to the ASA console port to determine the exact error condition that disabled clustering; for example, the cluster control link might be down.

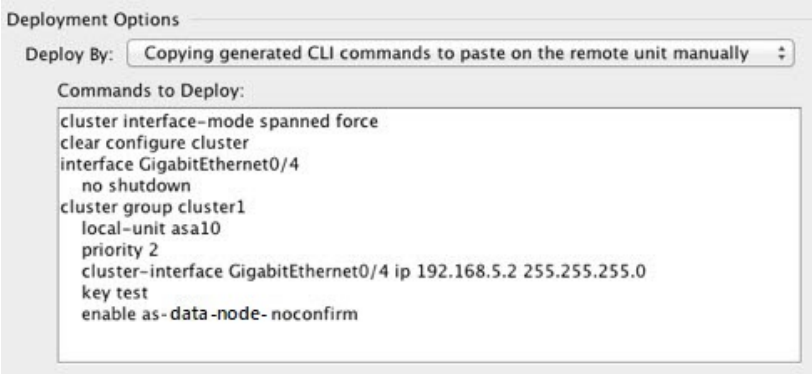
Step 7 To add a data node, click **Yes**.

If you are re-running the wizard from the control node, you can add data nodes by choosing the **Add another member to the cluster** option when you first start the wizard.

Step 8 In the **Deployment Options** area, choose one of the following **Deploy By** options:

- **Sending CLI commands to the remote unit now—**Send the bootstrap configuration to the data node (temporary) management IP address. Enter the data node management IP address, username, and password.

- **Copying generated CLI commands to paste on the remote unit manually**—Generates the commands so that you can cut and paste them at the data node CLI or using the CLI tool in ASDM. In the Commands to Deploy box, select and copy the generated commands for later use.



Customize the Clustering Operation

You can customize clustering health monitoring, TCP connection replication delay, flow mobility and other optimizations, either as part of the Day 0 configuration or after you deploy the cluster.

Perform these procedures on the control node.

Configure Basic ASA Cluster Parameters

You can customize cluster settings on the control node. If you do not use the wizard to add a node to the cluster, you can configure the cluster parameters manually. If you already enabled clustering, you can edit some cluster parameters; others that cannot be edited while clustering is enabled are grayed out. This procedure also includes advanced parameters that are not included in the wizard.

Before you begin

- If you did not use the wizard, and want to manually join the cluster, you need to pre-configure the cluster control link on each node before joining the cluster. See Configure the Cluster Control Link on the Control Node, on page 520.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster**.

If your device is already in the cluster, and is the control node, then this pane is on the **Cluster Configuration** tab.

Step 2 Check the **Configure ASA cluster settings** check box.

If you uncheck the check box, the settings are erased. Do not check **Participate in ASA cluster** until after you set all your parameters.

Note
After you enable clustering, do not uncheck the **Configure ASA cluster settings** check box without understanding the consequences. This action clears all cluster configuration, and also shuts down all interfaces including the management interface to which ASDM is connected. To restore connectivity in this case, you need to access the CLI at the console port.

Step 3 Configure the following bootstrap parameters:

- **Cluster Name**—Names the cluster. The name must be an ASCII string from 1 to 38 characters. You can only configure one cluster per node. All members of the cluster must use the same name.
- **Member Name**—Names this member of the cluster with a unique ASCII string from 1 to 38 characters.
- **Member Priority**—Sets the priority of this node for control node elections, between 1 and 100, where 1 is the highest priority.
- **Site Index**—If you use inter-site clustering, set the site ID for this node so it uses a site-specific MAC address, between 1 and 8.
- (Optional) **Shared Key**—Sets an encryption 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 encryption key. This parameter does not affect datapath traffic, including connection state update and forwarded packets, which are always sent in the clear. You must configure this parameter if you also enable the password encryption service.
- (Optional) **Enable connection rebalancing for TCP traffic across all the ASAs in the cluster**—Enables connection rebalancing. This parameter is disabled by default. This parameter is not part of the bootstrap configuration, and is replicated from the control node to the data nodes. If enabled, ASAs exchange information about the connections per second periodically, and offload new connections from devices with more connections per second to less loaded devices. Existing connections are never moved. Moreover, because this command only rebalances based on connections per second, the total number of established connections on each node is not considered, and the total number of connections may not be equal. The frequency, between 1 and 360 seconds, specifies how often the load information is exchanged. The default is 5 seconds.

  Once a connection is offloaded to a different node, it becomes an asymmetric connection.

  Do not configure connection rebalancing for inter-site topologies; you do not want new connections rebalanced to cluster members at a different site.
- **Enable cluster load monitor**—You can monitor the traffic load for cluster members, including total connection count, CPU and memory usage, and buffer drops. If the load is too high, you can choose to manually disable clustering on the node if the remaining nodes can handle the load, or adjust the load balancing on the external switch. This feature is enabled by default. You can periodically monitor the traffic load. If the load is too high, you can choose to manually disable clustering on the node.

  Set the following values:

  - **Time Interval**—Sets the time in seconds between monitoring messages, between 10 and 360 seconds. The default is 20 seconds.
  - **Number of Intervals**—Sets the number of intervals for which the ASA maintains data, between 1 and 60. The default is 30.

See **Monitoring** > **ASA Cluster** > **Cluster Load-Monitoring** to view the traffic load.

- (Optional) **Enable health monitoring of this device within the cluster**—Enables the cluster node health check feature, and determines the amount of time between node heartbeat status messages, between .3 and 45 seconds; The default is 3 seconds. **Note**: When you are adding new nodes to the cluster, and making topology changes on the ASA or the switch, you should disable this feature temporarily until the cluster is complete, and also disable interface monitoring for the disabled interfaces (**Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** > **Cluster Interface Health Monitoring**). You can re-enable this feature after cluster and topology changes are complete. To determine node health, the ASA 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 holdtime period, the peer node is considered unresponsive or dead.
- (Optional) **Debounce Time**—Configures the debounce time before the ASA considers an interface to be failed and the node is removed from the cluster. This feature allows 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 ASA waits the number of milliseconds specified before marking the interface as failed and the node is removed from the cluster. The default debounce time is 500 ms, with a range of 300 ms to 9 seconds.
- (Optional) **Replicate console output**—Enables console replication from data nodes to the control node. This feature is disabled by default. The ASA may print out some messages directly to the console for certain critical events. If you enable console replication, data nodes send the console messages to the control node so that you only need to monitor one console port for the cluster. This parameter is not part of the bootstrap configuration, and is replicated from the control node to the data nodes.
- (Optional) **Enable Clustering Flow Mobility**. See Configure LISP Inspection, on page 533.
- (Optional) **Enable Director Localization for inter-DC cluster**—To improve performance and reduce round-trip time latency for inter-site clustering for data centers, you can enable director localization. New connections are typically load-balanced and owned by cluster members within a given site. However, the ASA assigns the Director role to a member at *any* site. Director localization enables additional Director roles: a Local Director at the same site as the Owner, and a Global Director that can be at any site. Keeping the Owner and Director at the same site improves performance. Also, if the original Owner fails, the Local Director will choose a new connection Owner at the same site. The Global Director is used if a cluster member receives packets for a connection that is owned on a different site.
- (Optional) **Site Redundancy**—To protect flows from a site failure, you can enable site redundancy. If the connection backup owner is at the same site as the owner, then an additional backup owner will be chosen from another site to protect flows from a site failure. Director localization and site redundancy are separate features; you can configure one or the other, or configure both.
- (Optional) **Enable config sync acceleration**—When a data node has the same configuration as the control node, it will skip syncing the configuration and will join faster. This feature is enabled by default. This feature is configured on each node, and is not replicated from the control node to the data node.

  Note
  Some configuration commands are not compatible with accelerated cluster joining; if these commands are present on the node, even if accelerated cluster joining is enabled, configuration syncing will always occur. You must remove the incompatible configuration for accelerated cluster joining to work. Use the **show cluster info unit-join-acceleration incompatible-config** to view incompatible configuration.

- **Enable parallel configuration replicate**—Enable the control node to sync configuration changes with data nodes in parallel. Otherwise, synching occurs sequentially, and can take more time.

- **Flow State Refresh Keepalive Interval**—Set the keepalive interval for flow state refresh messages (clu_keepalive and clu_update messages) from the flow owner to the director and backup owner, between 15 and 20 seconds. The default is 15. You may want to set the interval to be longer than the default to reduce the amount of traffic on the cluster control link.
- **Cluster Control Link**—Specifies the cluster control link interface.
  - **Interface**—Specifies the VNI interface.
  - **IP Address**—Specifies an IPv4 address for the IP address; IPv6 is not supported for this interface.
  - **Subnet Mask**—Specifies the subnet mask.
  - **MTU**—Specifies the maximum transmission unit for the VTEP source interface to be at least 154 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 (100 bytes) and VXLAN overhead (54 bytes). Set the MTU between 1554 and 9198 bytes. The default MTU is 1554 bytes. We suggest setting the cluster control link MTU to 1654 when data interfaces are set to 1500; this value requires jumbo frame reservation. For example, when using jumbo frames, 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. This parameter is not part of the bootstrap configuration, and is replicated from the control node to the data nodes. **Note:** If you have not pre-enabled jumbo frame reservation, enable jumbo frames, and then restart this procedure.

Step 4 Check the **Participate in ASA cluster** check box to join the cluster.

Step 5 Click **Apply**.

Configure Interface Health Monitoring and Auto-Rejoin Settings

You might want to disable health monitoring of non-essential interfaces, for example, the management interface. Health monitoring is not performed on VLAN subinterfaces. You cannot configure monitoring for the cluster control link; it is always monitored.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** > **Cluster Interface Health Monitoring**.

Step 2 In the **Monitored Interfaces** box, select an interface, and click **Add** to move it to the **Unmonitored Interfaces** box.

Interface status messages detect link failure. If a node does not receive interface status messages within the holdtime, then the amount of time before the ASA removes a member from the cluster depends on whether the node is an established member or is joining the cluster. Health check is enabled by default for all interfaces.

You might want to disable health monitoring of non-essential interfaces, for example, the management interface. Health monitoring is not performed on VLAN subinterfaces. You cannot configure monitoring for the cluster control link; it is always monitored.

When any topology changes occur (such as adding or removing a data interface, enabling or disabling an interface on the ASA or the switch) you should disable the health check feature (**Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster**) 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 health check feature.

Step 3 Click the **Auto Rejoin** tab to customize the auto-rejoin settings in case of an interface, system, or cluster control link failure. For each type, click **Edit** to set the following:

- **Maximum Rejoin Attempts**—Define the number of attempts at rejoining the cluster by setting **Unlimited** or a value between 0 and 65535. **0** disables auto-rejoining. The default value is **Unlimited** for the cluster-interface and **3** for the data-interface and system.
- **Rejoin Interval**—Define the interval duration in minutes between rejoin attempts by setting the interval 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**—Define if the interval duration increases by setting the interval variation 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.

Click **Restore Defaults** to restore the default settings.

Step 4 Click **Apply**.

Configure the Cluster TCP Replication Delay

Enable the cluster replication delay for TCP connections to help eliminate the “unnecessary work” related to short-lived flows by delaying the director/backup flow creation. Note that if a unit fails before the director/backup flow is created, then those flows cannot be recovered. Similarly, if traffic is rebalanced to a different unit before the flow is created, then the flow cannot be recovered. You should not enable the TCP replication delay for traffic on which you disable TCP randomization.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster Replication**.

Step 2 Click **Add** and set the following values:

- **Replication delay**—Set the seconds between 1 and 15.
- **HTTP**—Set the delay for all HTTP traffic.
- **Source Criteria**
  - **Source**—Set the source IP address.
  - **Service**—(Optional) Set the source port. Typically you set either the source or the destination port, but not both.

- **Destination Criteria**
  - **Source**—Set the destination IP address.
  - **Service**—(Optional) Set the destination port. Typically you set either the source or the destination port, but not both.

Step 3 Click **OK**.

Step 4 Click **Apply**.

Configure Inter-Site Features

For inter-site clustering, you can customize your configuration to enhance redundancy and stability.

Configure Cluster Flow Mobility

You can inspect LISP traffic to enable flow mobility when a server moves between sites.

About LISP Inspection

You can inspect LISP traffic to enable flow mobility between sites.

About LISP

Data center virtual machine mobility such as VMware VMotion enables servers to migrate between data centers while maintaining connections to clients. To support such data center server mobility, routers need to be able to update the ingress route towards the server when it moves. Cisco Locator/ID Separation Protocol (LISP) architecture separates the device identity, or endpoint identifier (EID), from its location, or routing locator (RLOC), into two different numbering spaces, making server migration transparent to clients. For example, when a server moves to a new site and a client sends traffic to the server, the router redirects traffic to the new location.

LISP requires routers and servers in certain roles, such as the LISP egress tunnel router (ETR), ingress tunnel router (ITR), first hop routers, map resolver (MR), and map server (MS). When the first hop router for the server senses that the server is connected to a different router, it updates all of the other routers and databases so that the ITR connected to the client can intercept, encapsulate, and send traffic to the new server location.

ASA LISP Support

The ASA does not run LISP itself; it can, however, inspect LISP traffic for location changes and then use this information for seamless clustering operation. Without LISP integration, when a server moves to a new site, traffic comes to an ASA cluster member at the new site instead of to the original flow owner. The new ASA forwards traffic to the ASA at the old site, and then the old ASA has to send traffic back to the new site to reach the server. This traffic flow is sub-optimal and is known as “tromboning” or “hair-pinning.”

With LISP integration, the ASA cluster members can inspect LISP traffic passing between the first hop router and the ETR or ITR, and can then change the flow owner to be at the new site.

LISP Guidelines

- The ASA cluster members must reside between the first hop router and the ITR or ETR for the site. The ASA cluster itself cannot be the first hop router for an extended segment.

- Only fully-distributed flows are supported; centralized flows, semi-distributed flows, or flows belonging to individual nodes are not moved to new owners. Semi-distributed flows include applications, such as SIP, where all child flows are owned by the same ASA that owns the parent flow.
- The cluster only moves Layer 3 and 4 flow states; some application data might be lost.
- For short-lived flows or non-business-critical flows, moving the owner may not be worthwhile. You can control the types of traffic that are supported with this feature when you configure the inspection policy, and should limit flow mobility to essential traffic.

ASA LISP Implementation

This feature includes several inter-related configurations (all of which are described in this chapter):

1. (Optional) Limit inspected EIDs based on the host or server IP address—The first hop router might send EID-notify messages for hosts or networks the ASA cluster is not involved with, so you can limit the EIDs to only those servers or networks relevant to your cluster. For example, if the cluster is only involved with 2 sites, but LISP is running on 3 sites, you should only include EIDs for the 2 sites involved with the cluster.
2. LISP traffic inspection—The ASA inspects LISP traffic on UDP port 4342 for the EID-notify message sent between the first hop router and the ITR or ETR. The ASA maintains an EID table that correlates the EID and the site ID. For example, you should inspect LISP traffic with a source IP address of the first hop router and a destination address of the ITR or ETR. Note that LISP traffic is not assigned a director, and LISP traffic itself does not participate in cluster state sharing.
3. Service Policy to enable flow mobility on specified traffic—You should enable flow mobility on business-critical traffic. For example, you can limit flow mobility to only HTTPS traffic, and/or to traffic to specific servers.
4. Site IDs—The ASA uses the site ID for each cluster node to determine the new owner.
5. Cluster-level configuration to enable flow mobility—You must also enable flow mobility at the cluster level. This on/off toggle lets you easily enable or disable flow mobility for a particular class of traffic or applications.

Configure LISP Inspection

You can inspect LISP traffic to enable flow mobility when a server moves between sites.

Before you begin

- Assign each cluster unit to a site ID according to Configure Basic ASA Cluster Parameters, on page 527.
- LISP traffic is not included in the default-inspection-traffic class, so you must configure a separate class for LISP traffic as part of this procedure.

Procedure

Step 1 (Optional) Configure a LISP inspection map to limit inspected EIDs based on IP address, and to configure the LISP pre-shared key:

a) Choose **Configuration** > **Firewall** > **Objects** > **Inspect Maps** > **LISP**.

b) Click **Add** to add a new map.
c) Enter a name (up to 40 characters) and description.
d) For the **Allowed-EID access-list**, click **Manage**.

The **ACL Manager** opens.

The first hop router or ITR/ETR might send EID-notify messages for hosts or networks the ASA cluster is not involved with, so you can limit the EIDs to only those servers or networks relevant to your cluster. For example, if the cluster is only involved with 2 sites, but LISP is running on 3 sites, you should only include EIDs for the 2 sites involved with the cluster.

e) Add an ACL with at least one ACE according to the firewall configuration guide.
f) If necessary, enter the **Validation Key**.

If you copied an encrypted key, click the **Encrypted** radio button.

g) Click **OK**.

Step 2 Add a service policy rule to configure LISP inspection:

a) Choose **Configuration** > **Firewall** > **Service Policy Rules**.
b) Click **Add**.
c) On the **Service Policy** page, apply the rule to an interface or globally.

If you have an existing service policy you want to use, add a rule to that policy. By default, the ASA includes a global policy called **global_policy**. You can also create one service policy per interface if you do not want to apply the policy globally. LISP inspection is applied to traffic bidirectionally so you do not need to apply the service policy on both the source and destination interfaces; all traffic that enters or exits the interface to which you apply the rule is affected if the traffic matches the class for both directions.

d) On the **Traffic Classification Criteria** page, click **Create a new traffic class**, and under **Traffic Match Criteria**, check **Source and Destination IP Address (uses ACL)**.
e) Click **Next**.
f) Specify the traffic you want to inspect. You should specify traffic between the first hop router and the ITR or ETR on UDP port 4342. Both IPv4 and IPv6 ACLs are accepted.
g) Click **Next**.
h) On the **Rule Actions** wizard page or tab, select the **Protocol Inspection** tab.
i) Check the **LISP** check box,.
j) (Optional) Click **Configure** to choose the inspection map you created.
k) Click **Finish** to save the service policy rule.

Step 3 Add a service policy rule to enable Flow Mobility for critical traffic:

a) Choose **Configuration** > **Firewall** > **Service Policy Rules**.
b) Click **Add**.
c) On the **Service Policy** page, choose the same service policy you used for LISP inspection.
d) On the **Traffic Classification Criteria** page, click **Create a new traffic class**, and under **Traffic Match Criteria**, check **Source and Destination IP Address (uses ACL)**.
e) Click **Next**.
f) Specify the business critical traffic that you want to re-assign to the most optimal site when servers change sites. For example, you can limit flow mobility to only HTTPS traffic, and/or to traffic to specific servers. Both IPv4 and IPv6 ACLs are accepted.
g) Click **Next**.
h) On the **Rule Actions** wizard page or tab, select the **Cluster** tab.

i) Check the **Enable Cluster flow-mobility triggered by LISP EID messages** check box.
j) Click **Finish** to save the service policy rule.

Step 4 Choose **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** > **Cluster Configuration**, and check the **Enable Clustering flow mobility** check box.

Step 5 Click **Apply**.

Manage Cluster Nodes

After you deploy the cluster, you can change the configuration and manage cluster nodes.

Add a New Data Node from the Control Node

You can add additional data nodes to the cluster from the control node. You can also add data nodes using the High Availability and Scalability wizard. Adding a data node from the control node has the benefit of configuring the cluster control link and setting the cluster interface mode on each data node you add.

You can alternatively log into the data node and configure clustering directly on the node. However, after you enable clustering, your ASDM session will be disconnected, and you will have to reconnect.

Before you begin

- If you want to send the bootstrap configuration over the management network, be sure the data node has an accessible IP address.

Procedure

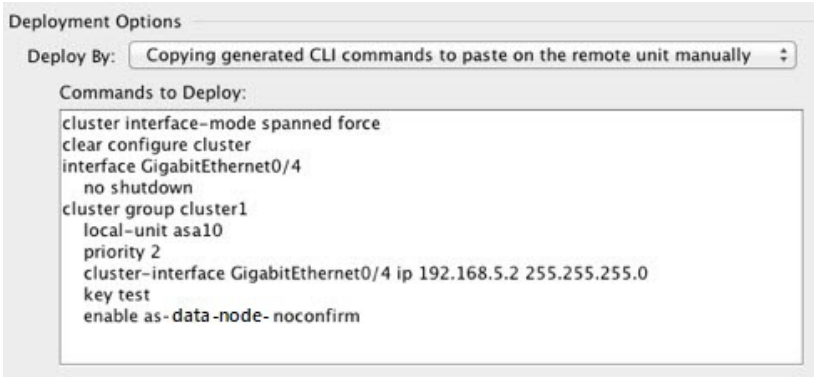
Step 1 Choose **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** > **Cluster Members**.

Step 2 Click **Add**.

Step 3 Configure the following parameters:

- **Member Name**—Names this member of the cluster with a unique ASCII string from 1 to 38 characters.
- **Member Priority**—Sets the priority of this node for control node elections, between 1 and 100, where 1 is the highest priority.
- **Cluster Control Link > IP Address**—Specifies a unique IP address for this member for the cluster control link, on the same network as the control node cluster control link.
- In the **Deployment Options** area, choose one of the following **Deploy By** options:
  - **Sending CLI commands to the remote unit now**—Send the bootstrap configuration to the data node (temporary) management IP address. Enter the data node management IP address, username, and password.

- **Copying generated CLI commands to paste on the remote unit manually**—Generates the commands so that you can cut and paste them at the data node CLI or using the CLI tool in ASDM. In the Commands to Deploy box, select and copy the generated commands for later use.



Step 4 Click **OK**, then **Apply**.

Become an Inactive Node

To become an inactive member of the cluster, disable clustering on the node while leaving the clustering configuration intact.

Note When an ASA becomes inactive (either manually or through a health check failure), all data interfaces are shut down; only the management-only interface can send and receive traffic. To resume traffic flow, re-enable clustering; or you can remove the node altogether from the cluster. The management interface remains up using the IP address the node received from the cluster IP pool. However if you reload, and the node is still inactive in the cluster (for example, you saved the configuration with clustering disabled), then the management interface is disabled. You must use the console port for any further configuration.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** > **Cluster Configuration**.

Step 2 Uncheck the **Participate in ASA cluster** check box.

Note
Do not uncheck the **Configure ASA cluster settings** check box; this action clears all cluster configuration, and also shuts down all interfaces including the management interface to which ASDM is connected. To restore connectivity in this case, you need to access the CLI at the console port.

Step 3 Click **Apply**.

Deactivate a Data Node from the Control Node

To deactivate a data node, perform the following steps.

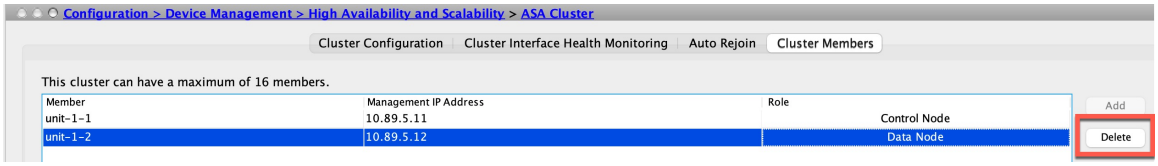
Note When an ASA becomes inactive, all data interfaces are shut down; only the management-only interface can send and receive traffic. To resume traffic flow, re-enable clustering. The management interface remains up using the IP address the node received from the cluster IP pool. However if you reload, and the node is still inactive in the cluster (for example, if you saved the configuration with clustering disabled), the management interface is disabled. You must use the console port for any further configuration.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** > **Cluster Members**.

Step 2 Select the data node that you want to remove, and click **Delete**.

Figure 73: Delete Node



The data node bootstrap configuration remains intact, so that you can later re-add the data node without losing your configuration.

Step 3 Click **Apply**.

Rejoin the Cluster

If a node was removed from the cluster, for example for a failed interface or if you manually deactivated a member, you must manually rejoin the cluster.

Procedure

Step 1 If you still have ASDM access, you can reenable clustering in ASDM by connecting ASDM to the node you want to reenable.

You cannot reenable clustering for a data node from the control node unless you add it as a new member.

a) Choose **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster**.
b) Check the **Participate in ASA cluster** check box.
c) Click **Apply**.

Step 2 If you cannot use ASDM: At the console, enter cluster configuration mode:

cluster group *name*

Example:

```
ciscoasa(config)# cluster group pod1
```

Step 3 Enable clustering.

enable

Leave the Cluster

If you want to leave the cluster altogether, you need to remove the entire cluster bootstrap configuration. Because the current configuration on each node is the same (synced from the active unit), leaving the cluster also means either restoring a pre-clustering configuration from backup, or clearing your configuration and starting over to avoid IP address conflicts.

Procedure

Step 1 For a data node, disable clustering:

cluster group *cluster_name*
no enable

Example:

```
ciscoasa(config)# cluster group cluster1
ciscoasa(cfg-cluster)# no enable
```

You cannot make configuration changes while clustering is enabled on a data node.

Step 2 Clear the cluster configuration:

clear configure cluster

The ASA shuts down all interfaces including the management interface and cluster control link.

Step 3 Disable cluster interface mode:

no cluster interface-mode

The mode is not stored in the configuration and must be reset manually.

Step 4 If you have a backup configuration, copy the backup configuration to the running configuration:

copy *backup_cfg* **running-config**

Example:

```
ciscoasa(config)# copy backup_cluster.cfg running-config

Source filename [backup_cluster.cfg]?

Destination filename [running-config]?
ciscoasa(config)#
```

Step 5 Save the configuration to startup:

write memory

Step 6 If you do not have a backup configuration, reconfigure management access. Be sure to change the interface IP addresses, and restore the correct hostname, for example.

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 node you want to become the control node, use the procedure in this section. Note, however, that for centralized features, if you force a control node change using this procedure, 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 Choose **Monitoring** > **ASA Cluster** > **Cluster Summary**.

Step 2 From the drop-down list, choose a data node to become control, and click the button to make it the control node.

Step 3 You are prompted to confirm the control node change. Click **Yes**.

Step 4 Quit ASDM, and reconnect using the Main cluster IP address.

Execute a Command Cluster-Wide

To send a command to all nodes in the cluster, or to a specific node, perform the following steps. Sending a **show** command to all nodes collects all output and displays it on the console of the current node. Other commands, such as **capture** and **copy**, can also take advantage of cluster-wide execution.

Before you begin

Perform this procedure at the Command Line Interface tool: choose **Tools > Command Line Interface**.

Procedure

Send a command to all nodes, or if you specify the node name, a specific node:

cluster exec [**unit** *node_name*] *command*

Example:

```
ciscoasa# cluster exec show xlate
```

To view node names, enter **cluster exec unit ?** (to see all names except the current node), or enter the **show cluster info** command.

Examples

To copy the same capture file from all nodes in the cluster at the same time to a TFTP server, enter the following command on the control node:

```
ciscoasa# cluster exec copy /pcap capture: tftp://10.1.1.56/capture1.pcap
```

Multiple PCAP files, one from each node, are copied to the TFTP server. The destination capture file name is automatically attached with the node name, such as capture1_asa1.pcap, capture1_asa2.pcap, and so on. In this example, asa1 and asa2 are cluster node names.

Monitoring the ASA Virtual Cluster

You can monitor and troubleshoot cluster status and connections.

Monitoring Cluster Status

See the following screens for monitoring cluster status:

- **Monitoring > ASA Cluster > Cluster Summary**

  This pane shows cluster information about the node to which you are connected, as well as other nodes in the cluster. You can also change the primary node from this pane.

- **Cluster Dashboard**

  On the home page on the primary node, you can monitor the cluster using the Cluster Dashboard and the Cluster Firewall Dashboard.

Capturing Packets Cluster-Wide

See the following screen for capturing packets in a cluster:

Wizards > Packet Capture Wizard

To support cluster-wide troubleshooting, you can enable capture of cluster-specific traffic on the control node, which is then automatically enabled on all of the data nodes in the cluster.

Monitoring Cluster Resources

See the following screens for monitoring cluster resources:

- **Monitoring > ASA Cluster > System Resources Graphs > CPU**

  This pane lets you create graphs or tables showing the CPU utilization across the cluster nodes.

- **Monitoring > ASA Cluster > System Resources Graphs > Memory**. This pane lets you create graphs or tables showing the Free Memory and Used Memory across the cluster nodes.

Monitoring Cluster Traffic

See the following screens for monitoring cluster traffic:

- **Monitoring > ASA Cluster > Traffic Graphs > Connections**.

  This pane lets you create graphs or tables showing the Connections across the cluster members.

- **Monitoring > ASA Cluster > Traffic Graphs > Throughput**.

  This pane lets you create graphs or tables showing the traffic throughput across the cluster members.

- **Monitoring** > **ASA Cluster** > **Cluster Load-Monitoring**

  This section includes the **Load Monitor-Information** and **Load-Monitor Details** panes. **Load Monitor-Information** shows the traffic load for cluster members for the last interval and also the average over total number of intervals configured (30 by default). Use the **Load-Monitor Details** pane to view the value for each measure at each interval.

Monitoring the Cluster Control Link

See the following screen for monitoring cluster status:

Monitoring > Properties > System Resources Graphs > Cluster Control Link.

This pane lets you create graphs or tables showing the cluster control link Receival and Transmittal capacity utilization.

Monitoring Cluster Routing

See the following screen for cluster routing:

- **Monitoring** > **Routing** > **LISP-EID Table**

  Shows the ASA EID table showing EIDs and site IDs.

Configuring Logging for Clustering

See the followingscreen for configuring logging for clustering:

Configuration > Device Management > Logging > Syslog Setup

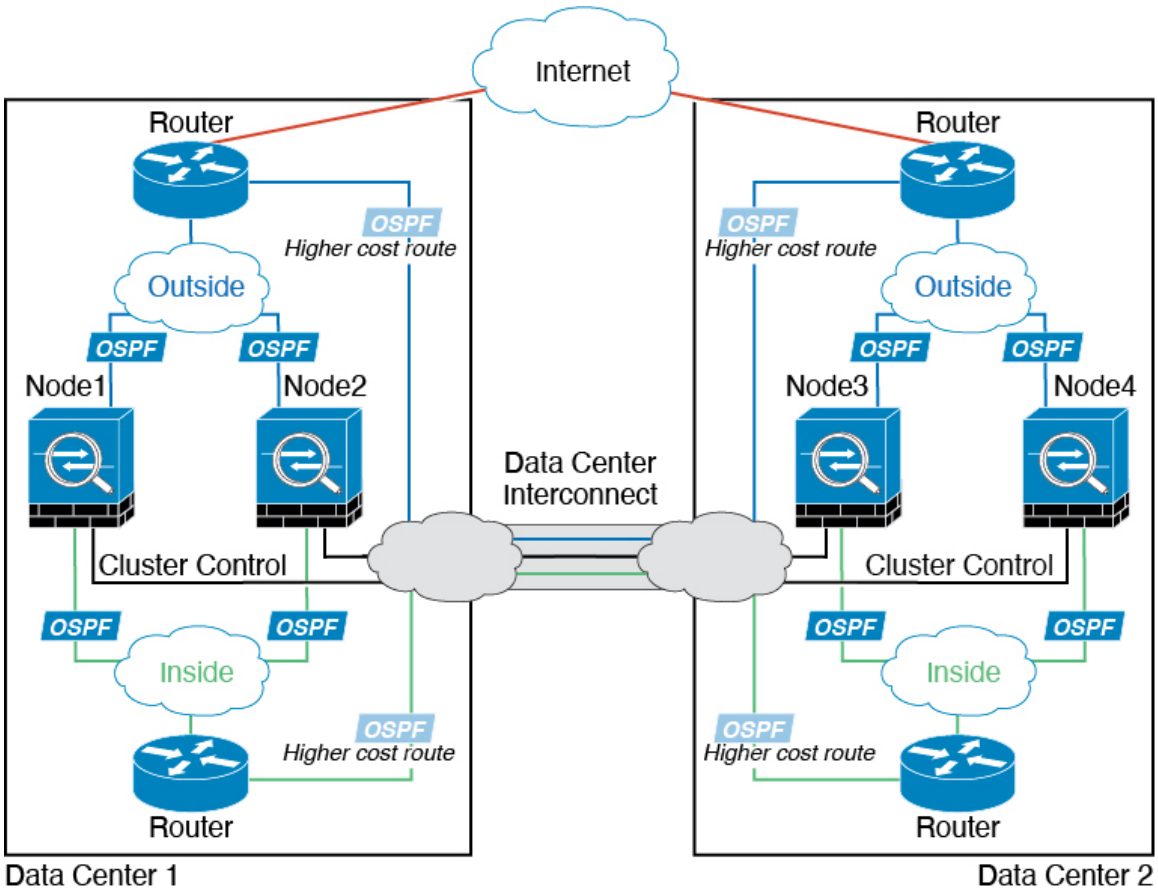
Each node in the cluster generates syslog messages independently. You can generate syslog messages with identical or different device IDs to make messages appear to come from the same or different nodes in the cluster.

Examples for ASA Virtual Clustering

These examples include all cluster-related ASA configuration for typical deployments.

Individual Interface Routed Mode North-South Inter-Site Example

The following example shows 2 ASA cluster nodes at each of 2 data centers placed between inside and outside routers (North-South insertion). The cluster nodes are connected by the cluster control link over the DCI. The inside and outside routers at each data center use OSPF and PBR or ECMP to load balance the traffic between cluster members. By assigning a higher cost route across the DCI, traffic stays within each data center unless all ASA cluster nodes at a given site go down. In the event of a failure of all cluster nodes at one site, traffic goes from each router over the DCI to the ASA cluster nodes at the other site.



Reference for Clustering

This section includes more information about how clustering operates.

ASA Features and Clustering

Some ASA features are not supported with ASA clustering, and some are only supported on the control node. 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.

- Unified Communication features that rely on TLS Proxy
- Remote access VPN (SSL VPN and IPsec VPN)
- Virtual Tunnel Interfaces (VTIs)
- The following application inspections:
  - CTIQBE

  - H323, H225, and RAS
  - IPsec passthrough
  - MGCP
  - MMP
  - RTSP
  - SCCP (Skinny)
  - WAAS
  - WCCP
- Botnet Traffic Filter
- Auto Update Server
- DHCP client, server, and proxy. DHCP relay is supported.
- VPN load balancing
- Failover on Azure
- Integrated Routing and Bridging
- FIPS 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.

- The following application inspections:
  - DCERPC
  - ESMTP
  - IM
  - NetBIOS
  - PPTP
  - RADIUS

  - RSH
  - SNMP
  - SQLNET
  - SUNRPC
  - TFTP
  - XDMCP
- Static route monitoring
- Authentication and Authorization for network access. Accounting is decentralized.
- Filtering Services
- Site-to-site VPN
- Multicast routing

Features Applied to Individual Nodes

These features are applied to each ASA node, instead of the cluster as a whole or to the control node.

- QoS—The QoS policy is synced across the cluster as part of configuration replication. However, the policy is enforced on each node independently. For example, if you configure policing on output, then the conform rate and conform burst values are enforced on traffic exiting a particular ASA. In a cluster with 3 nodes and with traffic evenly distributed, the conform rate actually becomes 3 times the rate for the cluster.
- Threat detection—Threat detection works on each node independently; for example, the top statistics is node-specific. Port scanning detection, for example, does not work because scanning traffic will be load-balanced between all nodes, and one node will not see all traffic.
- Resource management—Resource management in multiple context mode is enforced separately on each node based on local usage.
- LISP traffic—LISP traffic on UDP port 4342 is inspected by each receiving node, but is not assigned a director. Each node adds to the EID table that is shared across the cluster, but the LISP traffic itself does not participate in cluster state sharing.

AAA for Network Access and Clustering

AAA for network access consists of three components: authentication, authorization, and accounting. Authentication and authorization are implemented as centralized features on the clustering control node with replication of the data structures to the cluster data nodes. If a control node is elected, the new control node will have all the information it needs to continue uninterrupted operation of the established authenticated users and their associated authorizations. Idle and absolute timeouts for user authentications are preserved when a control node change occurs.

Accounting is implemented as a distributed feature in a cluster. Accounting is done on a per-flow basis, so the cluster node owning a flow will send accounting start and stop messages to the AAA server when accounting is configured for a flow.

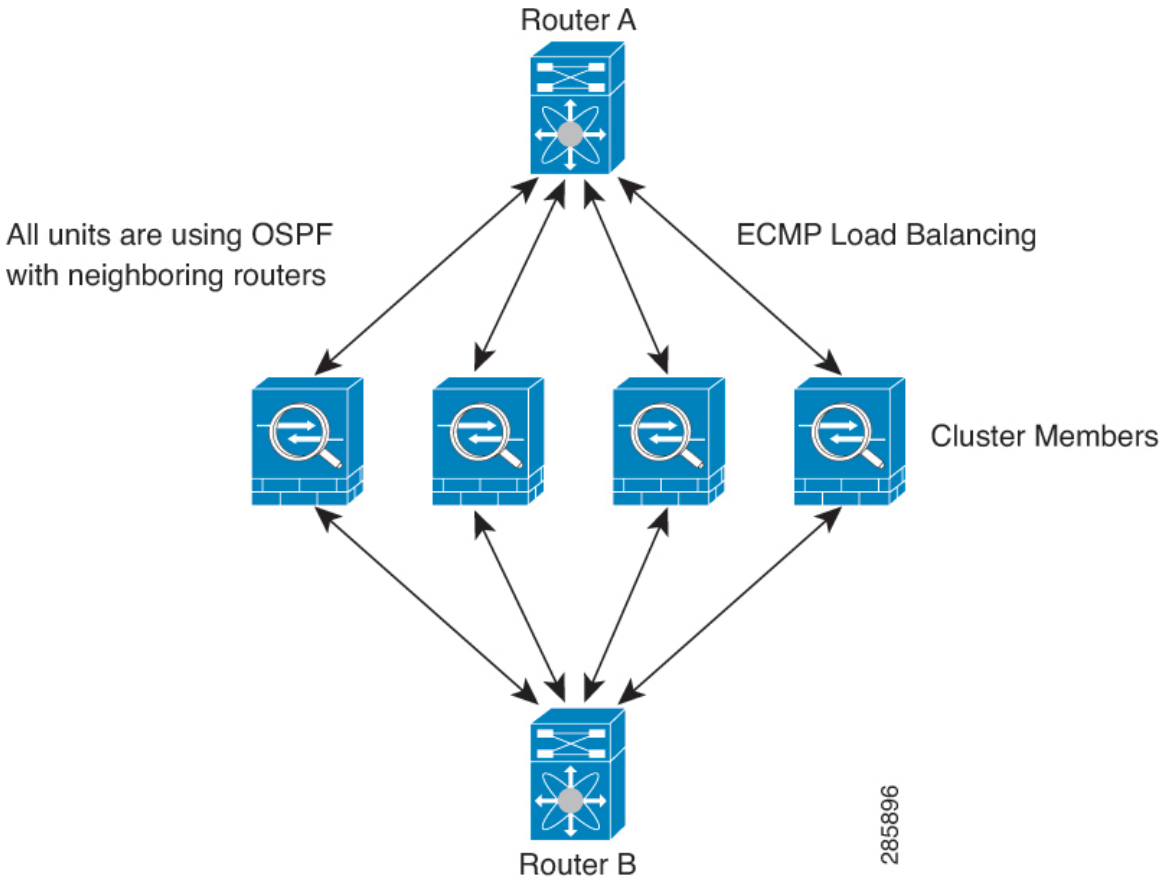
Connection Settings and Clustering

Connection limits are enforced cluster-wide (see **Configuration** > **Firewall** > **Service Policy** page). 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 74: 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 Configure a Traffic Zone, on page 684.

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.
- If you use AAA for FTP access, then the control channel flow is centralized on the control node.

ICMP Inspection and Clustering

The flow of ICMP and ICMP error packets through the cluster varies depending on whether ICMP/ICMP error inspection is enabled. Without ICMP inspection, ICMP is a one-direction flow, and there is no director flow support. With ICMP inspection, the ICMP flow becomes two-directional and is backed up by a director/backup flow. One difference for an inspected ICMP flow is in the director handling of a forwarded packet: the director will forward the ICMP echo reply packet to the flow owner instead of returning the packet to the forwarder.

Multicast Routing and Clustering

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 ASAs 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 ASA 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. 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.
- Per-session PAT feature—Although not exclusive to clustering, the per-session PAT feature improves the scalability of PAT and, for clustering, allows each data node to own PAT connections; by contrast, multi-session PAT connections have to be forwarded to and owned by the control node. By default, all TCP traffic and UDP DNS traffic use a per-session PAT xlate, whereas ICMP and all other UDP traffic uses multi-session. You can configure per-session NAT rules to change these defaults for TCP and UDP, but you cannot configure per-session PAT for ICMP. For example, with an increasing use of Quic protocol over UDP/443 as a best performance alternative compared to HTTPS TLS over TCP/443, you should enable per-session PAT for UDP/443. For traffic that benefits from multi-session PAT, such as H.323, SIP, or Skinny, you can disable per-session PAT for the associated TCP ports (the UDP ports for those H.323 and SIP are already multi-session by default). For more information about per-session PAT, see the firewall configuration guide.

- No static PAT for the following inspections—
  - FTP
  - PPTP
  - 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.

SCTP and Clustering

An SCTP association can be created on any node (due to load balancing); its multi-homing connections must reside on the same node.

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.

TLS Proxy configuration is not supported.

SNMP and Clustering

An SNMP agent polls each individual ASA by its Local IP address. You cannot poll consolidated data for the cluster.

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 re-add them on the control node to force the users to replicate to the new node, or directly on the data node.

STUN and Clustering

STUN inspection is supported in failover and cluster modes, as pinholes are replicated. However, the transaction ID is not replicated among nodes. In the case where a node fails after receiving a STUN Request and another node received the STUN Response, the STUN Response will be dropped.

Syslog and NetFlow and Clustering

- Syslog—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.

- NetFlow—Each node in the cluster generates its own NetFlow stream. The NetFlow collector can only treat each ASA as a separate NetFlow exporter.

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.

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 ASA Virtual Cluster

The ASA virtual 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 550 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.

When you enable health monitoring, all physical interfaces are monitored by default; you can optionally disable monitoring per interface. Only named interfaces can be monitored.

A node is removed from the cluster if its monitored interfaces fail. The amount of time before the ASA removes a member from the cluster depends on whether the node is an established member or is joining the cluster. The ASA 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 ASA to be removed from the cluster. The node is removed after 500 ms, regardless of the node 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 ASA automatically tries to rejoin the cluster, depending on the failure event.



Note When the ASA becomes inactive and fails to automatically rejoin the cluster, all data interfaces are shut down; only the management-only interface can send and receive traffic. The management interface remains up using the IP address the node received from the cluster IP pool. However if you reload, and the node is still inactive in the cluster, the management interface is disabled. You must use the console port for any further configuration.

Rejoining the Cluster

After a cluster node 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 ASA automatically tries to rejoin every 5 minutes, indefinitely. This behavior is configurable.
- Failed data interface—The ASA 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 ASA disables clustering. After you resolve the problem with the data interface, you have to manually enable clustering. This behavior is configurable.
- 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 and clustering is still enabled. The ASA attempts to rejoin the cluster every 5 seconds.
- Internal error—Internal failures include: application sync timeout; inconsistent application statuses; and so on. A node will attempt to rejoin the cluster automatically at the following intervals: 5 minutes, 10 minutes, and then 20 minutes. This behavior is configurable.

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 24: 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	Includes AAA rules (uauth).
IPv6 Neighbor database	Yes	—
Dynamic routing	Yes	—
SNMP Engine ID	**No**	—
Distributed VPN (Site-to-Site) for Firepower 4100/9300	Yes	Backup session becomes the active session, then a new backup session is created.

How the ASA Virtual 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.

  If you enable director localization for inter-site clustering, then there are two backup owner roles: the local backup and the global backup. The owner always chooses a local backup at the same site as itself (based on site ID). The global backup can be at any site, and might even be the same node as the local backup. The owner sends connection state information to both backups.

  If you enable site redundancy, and the backup owner is at the same site as the owner, then an additional backup owner will be chosen from another site to protect flows from a site failure. Chassis backup and site backup are independent, so in some cases a flow will have both a chassis backup and a site backup.

- 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.

  If you enable director localization for inter-site clustering, then there are two director roles: the local director and the global director. The owner always chooses a local director at the same site as itself (based on site ID). The global director can be at any site, and might even be the same node as the local director. If the original owner fails, then the local director chooses a new connection owner at the same site.

  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. If you enable director localization, then the forwarder always queries the local director. The forwarder only queries the global director if the local director does not know the owner, for example, if a cluster member receives packets for a connection that is owned on a different site. 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.

When a connection uses Port Address Translation (PAT), then the PAT type (per-session or multi-session) influences which member of the cluster becomes the owner of a new connection:

- Per-session PAT—The owner is the node that receives the initial packet in the connection.

  By default, TCP and DNS UDP traffic use per-session PAT.

- Multi-session PAT—The owner is always the control node. If a multi-session PAT connection is initially received by a data node, then the data node forwards the connection to the control node.

  By default, UDP (except for DNS UDP) and ICMP traffic use multi-session PAT, so these connections are always owned by the control node.

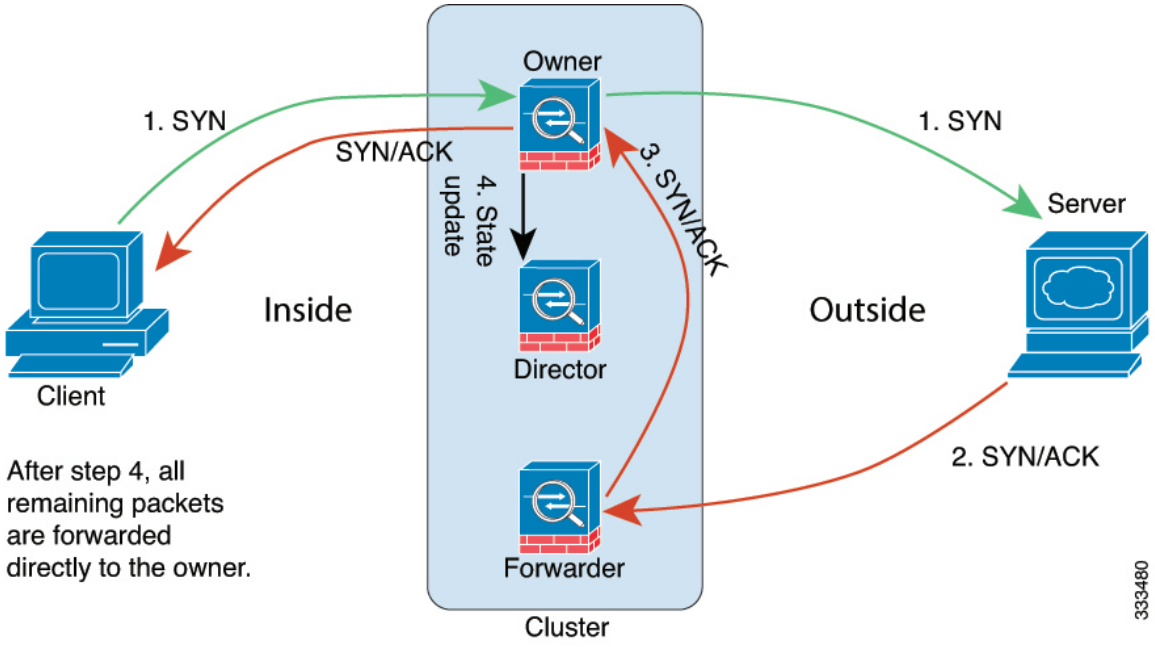
You can change the per-session PAT defaults for TCP and UDP so connections for these protocols are handled per-session or multi-session depending on the configuration. For ICMP, you cannot change from the default multi-session PAT. For more information about per-session PAT, see the firewall configuration guide.

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. For best performance, proper external load balancing is required for both directions of a flow to arrive at the same node, and for flows to be distributed evenly between nodes. 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.



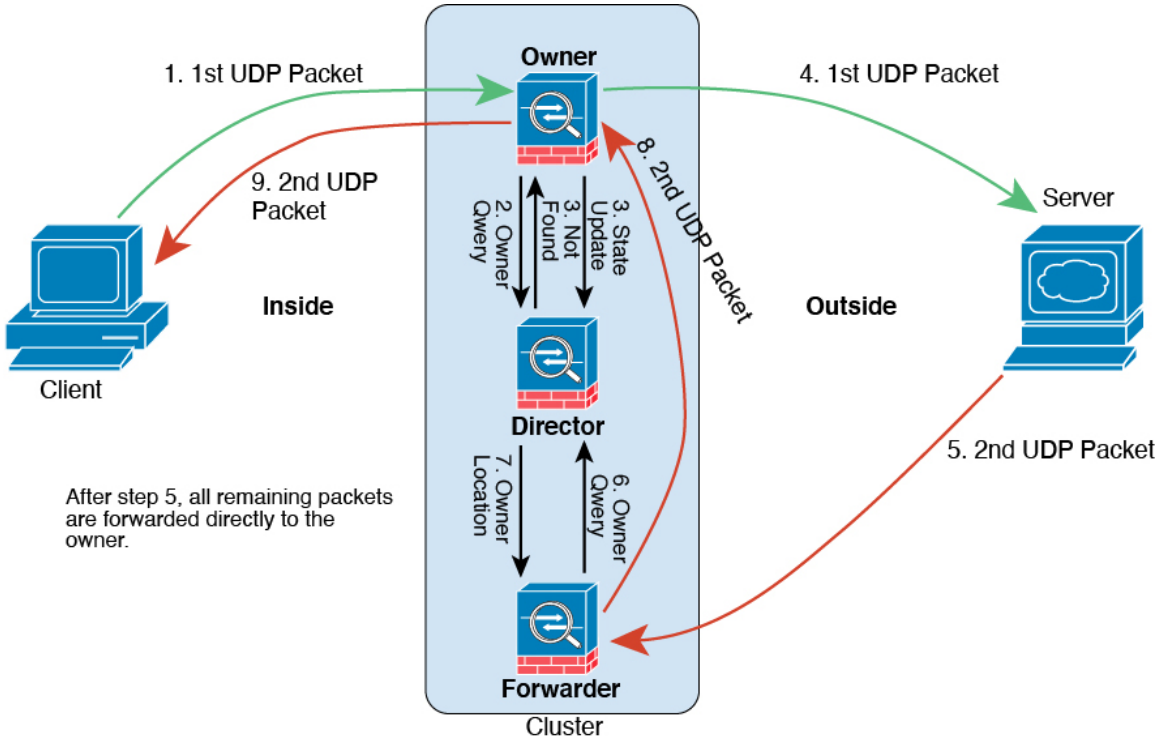
1. The SYN packet originates from the client and is delivered to one ASA (based on the load balancing method), which becomes the owner. The owner creates a flow, encodes owner information into a SYN cookie, and forwards the packet to the server.

2. The SYN-ACK packet originates from the server and is delivered to a different ASA (based on the load balancing method). This ASA 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 75: ICMP and UDP Data Flow***



The first UDP packet originates from the client and is delivered to one ASA (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.

Rebalancing New TCP Connections Across the Cluster

If the load balancing capabilities of the upstream or downstream routers result in unbalanced flow distribution, you can configure new connection rebalancing so nodes with higher new connections per second will redirect new TCP flows to other nodes. No existing flows will be moved to other nodes.

Because this command only rebalances based on connections per second, the total number of established connections on each node is not considered, and the total number of connections may not be equal.

Once a connection is offloaded to a different node, it becomes an asymmetric connection.

Do not configure connection rebalancing for inter-site topologies; you do not want new connections rebalanced to cluster members at a different site.

History for ASA Virtual Clustering

Feature Name	Version	Feature Information
Configurable cluster keepalive interval for flow status	9.20(1)	The flow owner sends keepalives (clu_keepalive messages) and updates (clu_update messages) to the director and backup owner to refresh the flow state. You can now set the keepalive interval. The default is 15 seconds, and you can set the interval between 15 and 55 seconds. You may want to set the interval to be longer to reduce the amount of traffic on the cluster control link. New/Modified screens: **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** > **Cluster Configuration**
Removal of biased language	9.19(1)	Commands, command output, and syslog messages that contained the terms "Master" and "Slave" have been changed to "Control" and "Data." New/Modified commands: **cluster control-node**, **enable as-data-node**, **prompt**, **show cluster history**, **show cluster info**

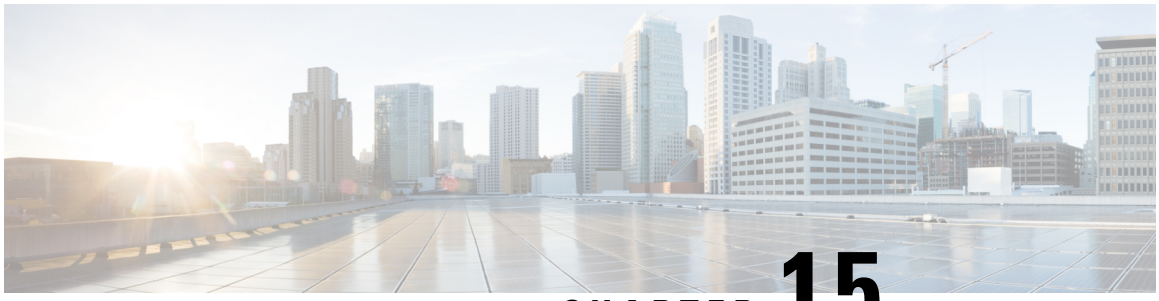
Feature Name	Version	Feature Information
ASAv30, ASAv50, and ASAv100 clustering for VMware and KVM	9.17(1)	The ASA virtual clustering lets you group up to 16 ASA virtuals together as a single logical device. A cluster provides all the convenience of a single device (management, integration into a network) while achieving the increased throughput and redundancy of multiple devices. The ASA virtual clustering supports Individual Interface mode in routed firewall mode; Spanned EtherChannels are not supported. The ASA virtual uses a VXLAN virtual interface (VNI) for the cluster control link. New/Modified screens: • **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces** • **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster**



PART III

Interfaces

- Basic Interface Configuration, on page 561
- Basic Interface Configuration for Firepower 1010 and Secure Firewall 1210/1220 Switch Ports, on page 575
- EtherChannel Interfaces, on page 585
- Loopback Interfaces, on page 597
- VLAN Subinterfaces, on page 605
- VXLAN Interfaces, on page 611
- Routed and Transparent Mode Interfaces, on page 633
- Advanced Interface Configuration, on page 665
- Traffic Zones, on page 675



CHAPTER 15

Basic Interface Configuration

This chapter includes basic interface configuration including Ethernet settings and Jumbo frame configuration.

Note For multiple context mode, complete all tasks in this section in the system execution space. If you are not already in the system execution space, in the Configuration > Device List pane, double-click **System** under the active device IP address.

Note For the Firepower 4100/9300 chassis, you configure basic interface settings in the FXOS operating system. See the configuration or getting started guide for your chassis for more information.

- About Basic Interface Configuration, on page 561
- Guidelines for Basic Interface Configuration, on page 563
- Default Settings for Basic Interface Configuration, on page 564
- Enable the Physical Interface and Configure Ethernet Parameters, on page 565
- Enable Jumbo Frame Support (ASA Virtual, ISA 3000), on page 567
- Manage the Network Module for the Secure Firewall 3100/4200, on page 567
- Examples for Basic Interfaces, on page 571
- History for Basic Interface Configuration, on page 572

About Basic Interface Configuration

This section describes interface features and special interfaces.

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.

Management Interface

The management interface, depending on your model, is a separate interface just for management traffic.

Management Interface Overview

You can manage the ASA by connecting to:

- Any through-traffic interface
- A dedicated Management *Slot*/*Port* interface (if available for your model)

You may need to configure management access to the interface according to Management Access, on page 1027.

Management *Slot*/*Port* Interface

The following table shows the Management interfaces per model.

Table 25: Management Interfaces Per Model

Model	Management 0/0	Management 1/1	Management 1/2	Configurable for Through Traffic	Subinterfaces Allowed
Firepower 1000	—	**Yes**	—	**Yes**	**Yes**
Secure Firewall 1200	—	**Yes**	—	**Yes**	**Yes**
Secure Firewall 3100	—	**Yes**	—	**Yes**	**Yes**
Secure Firewall 4200	—	**Yes**	**Yes**	**Yes**	**Yes**
Firepower 4100/9300	N/A The interface ID depends on the physical mgmt-type interface that you assigned to the ASA logical device	—	—	—	**Yes**
ISA 3000	—	**Yes**	—	—	—
ASAv	**Yes**	—	—	**Yes**	—

Use Any Interface for Management-Only Traffic

You can use any interface as a dedicated management-only interface by configuring it for management traffic, including an EtherChannel interface.

Management Interface for Transparent Mode

In transparent firewall mode, in addition to the maximum allowed through-traffic interfaces, you can also use the Management interface (either the physical interface, or a subinterface (if supported for your model)) as a separate management-only interface.

You cannot use any other interface types as Management interfaces.

For the Firepower 4100/9300 chassis, the management interface ID depends on the mgmt-type interface that you assigned to the ASA logical device.



Note The management interface is not part of a normal bridge group. Note that for operational purposes, it is part of a non-configurable bridge group.

Management-Only Traffic

The Management interface is for to- and from-the-box traffic only and cannot pass through traffic. In any case, since bridge groups cannot communicate with any other bridge groups, it wouldn't be possible for the Management interface to route to a bridge group or vice versa even if it did support through traffic.

Shared MAC Address Table and Switch Connections

In transparent firewall mode, the Management interface updates the ASA's MAC address table in the same manner as a data interface; therefore you should not connect both a Management and a data interface to the same switch unless you configure one of the switch ports as a routed port (by default Catalyst switches share a MAC address for all VLAN switch ports). Otherwise, if traffic arrives on the Management interface from the physically-connected switch, then the ASA updates the MAC address table to use the *Management* interface to access the switch, instead of the data interface. This action causes a temporary traffic interruption; the ASA will not re-update the MAC address table for packets from the switch to the data interface for at least 30 seconds for security reasons.

Multiple Context Mode

In multiple context mode, you cannot share any interfaces, including the Management interface, across contexts. To provide management per context on supported models, you can create subinterfaces of the Management interface and allocate a Management subinterface to each context. Some models do not allow subinterfaces on the Management interface, so per-context management for these models requires you to connect to a data interface. For the Firepower 4100/9300 chassis, the management interface and its subinterfaces are not recognized as specially-allowed management interfaces within the contexts; you must treat a management subinterface as a data interface in this case and add it to a BVI.

Guidelines for Basic Interface Configuration

Transparent Firewall Mode

For multiple context, transparent mode, each context must use different interfaces; you cannot share an interface across contexts.

Failover

You cannot share a failover or state interface with a data interface.

Additional Guidelines

Some management-related services are not available until a non-management interface is enabled, and the the ASA achieves a “System Ready” state. The ASA generates the following syslog message when it is in a “System Ready” state:

```
%ASA-6-199002: Startup completed.  Beginning operation.
```

Default Settings for Basic Interface Configuration

This section lists default settings for interfaces if you do not have a factory default configuration.

Default State of Interfaces

The default state of an interface depends on the type and the context mode.

In multiple context mode, all allocated interfaces are enabled by default, no matter what the state of the interface is in the system execution space. However, for traffic to pass through the interface, the interface also has to be enabled in the system execution space. If you shut down an interface in the system execution space, then that interface is down in all contexts that share it.

In single mode or in the system execution space, interfaces have the following default states:

- Physical interfaces—Disabled.
- VLAN subinterfaces—Enabled. However, for traffic to pass through the subinterface, the physical interface must also be enabled.
- VXLAN VNI interfaces—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 (Other models)—Disabled.



Note For the Firepower 4100/9300, you can administratively enable and disable interfaces in both the chassis and on the ASA. 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 the ASA.

Default Speed and Duplex

- By default, the speed and duplex for copper (RJ-45) interfaces are set to auto-negotiate.

Default Connector Type

Some models include two connector types: copper RJ-45 and fiber SFP. RJ-45 is the default. You can configure the ASA to use the fiber SFP connectors.

Default MAC Addresses

By default, the physical interface uses the burned-in MAC address, and all subinterfaces of a physical interface use the same burned-in MAC address.

Enable the Physical Interface and Configure Ethernet Parameters

This section describes how to:

- Enable the physical interface
- Set a specific speed and duplex (if available)
- (Secure Firewall 1200/3100/4200) Enable pause frames for flow control
- (Secure Firewall 3100/4200) Set Forward Error Correction

Before you begin

For multiple context mode, complete this procedure in the system execution space. If you are not already in the System configuration mode, in the **Configuration** > **Device List** pane, double-click **System** under the active device IP address.

Procedure

Step 1 Depending on your context mode:

- For single mode, choose the **Configuration > Device Setup > Interface Settings > Interfaces** pane.
- For multiple mode in the System execution space, choose the **Configuration > Context Management > Interfaces** pane.

By default, all physical interfaces are listed.

Step 2 Click a physical interface that you want to configure, and click **Edit**.

The **Edit Interface** dialog box appears.

Note
In single mode, this procedure only covers a subset of the parameters on the **Edit Interface** dialog box. Note that in multiple context mode, before you complete your interface configuration, you need to allocate interfaces to contexts.

Step 3 To enable the interface, check the **Enable Interface** check box.

Step 4 To add a description, enter text in the Description field.

The description can be up to 240 characters on a single line, without carriage returns. In the case of a failover or state link, the description is fixed as "LAN Failover Interface," "STATE Failover Interface," or "LAN/STATE Failover Interface," for example. You cannot edit this description. The fixed description overwrites any description you enter here if you make this interface a failover or state link.

Step 5 (Secure Firewall 1200/3100/4200) To enable pause (XOFF) frames for flow control, check the **Flow-Control** check box.

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 ASA 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 ASA 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 6 (Optional) To set the media type, duplex, speed, and enable pause frames for flow control, click **Configure Hardware Properties**.

a) To set the **Duplex** for RJ-45 interfaces, choose **Full**, **Half**, or **Auto**, depending on the interface type from the drop-down list.

Note
SFP interfaces only support full duplex.

b) To set the **Speed**, choose a value from the drop-down list (varies depending on the model).

For Firepower 1000 SFP interfaces, **Negotiate** sets the speed to 1000 Mbps and enables link negotiation for flow-control parameters and remote fault information. For 10 Gbps interfaces, this option sets the speed down to 1000 Mbps. The **Nonegotiate** option disables link negotiation. For Secure Firewall 1200/3100/4200 auto-negotiation options, see the **Auto-negotiate** check box on the **Advanced** tab, which lets you enable or disable auto-negotiation on any interface 1000 Mbps and higher. For Secure Firewall 1250, you can set a maximum speed of 2500 Mbps.

(Secure Firewall 1200/3100/4200) 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.

c) (Secure Firewall 1210CE/1210CP/3100/4200) To set the **FEC Mode** for 25 Gbps and higher interfaces, choose a value from the drop-down list.

For an EtherChannel member interface, you must configure Forward Error Correction before you add it to the EtherChannel.

d) Click **OK** to accept the **Hardware Properties** changes.

Step 7 Click **OK** to accept the **Interface** changes.

Enable Jumbo Frame Support (ASA Virtual, ISA 3000)

A jumbo frame is an Ethernet packet larger than the standard maximum of 1518 bytes (including Layer 2 header and VLAN header), up to 9216 bytes. You can enable support for jumbo frames for all interfaces by increasing the amount of memory to process Ethernet frames. Assigning more memory for jumbo frames might limit the maximum use of other features, such as ACLs. Note that the ASA MTU sets the payload size not including the Layer 2 (14 bytes) and VLAN header (4 bytes), so the maximum MTU is 9198, depending on your model.

This procedure only applies to the ISA 3000 and the ASA virtual. Other models support jumbo frames by default.

Jumbo frames are not supported on the ASAv5 and ASAv10 with less than 8GB RAM.

Before you begin

- In multiple context mode, set this option in the system execution space.
- Changes in this setting require you to reload the ASA.
- Be sure to set the MTU for each interface that needs to transmit jumbo frames to a higher value than the default 1500; for example, set the value to 9198. In multiple context mode, set the MTU within each context.
- Be sure to adjust the TCP MSS, either to disable it for non-IPsec traffic, or to increase it in accord with the MTU.

Procedure

Depending on your context mode:

- Multiple mode—To enable jumbo frame support, choose **Configuration > Context Management > Interfaces**, and click the **Enable jumbo frame support** check box.
- Single mode—Setting the MTU larger than 1500 bytes automatically enables jumbo frames. To manually enable or disable this setting, choose **Configuration > Device Setup > Interface Settings > Interfaces**, and click the **Enable jumbo frame support** check box.

Manage the Network Module for the Secure Firewall 3100/4200

If you install a network module before you first power on the firewall, no action is required; the network module is enabled and ready for use.

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.

If an interface is already in use in your configuration, you will have to manually remove any configuration related to interfaces that will no longer be present.

Before you begin

- You must use a supported breakout cable. See the hardware installation guide for more information.
- For clustering or failover, make sure the cluster/failover link is not using the parent interface (for breaking out) or the child interface (for rejoining); you cannot make changes to the interface if it is in use for the cluster/failover link.

Procedure

Step 1 Break out 10GB ports from one or more 40GB or higher interfaces by choosing **Configuration** > **Device Management** > **Advanced** > **EPM**, and entering one or more **Port Numbers** that you want to break out separated by commas (with no spaces).

The slot is always **2**.

For example, to break out the Ethernet2/1 and Ethernet 2/2 interfaces, you would specify **1,2** in the **Port Number** field. The resulting child interfaces will be identified as Ethernet2/1/1, Ethernet2/1/2, Ethernet2/1/3, Ethernet2/1/4, Ethernet2/2/1, Ethernet2/2/2, Ethernet2/2/3, and Ethernet2/2/4.

For clustering or failover, perform this step on the control node/active unit; the interface changes are replicated to the other nodes.

Step 2 Rejoin the breakout ports to restore the interface by choosing **Configuration** > **Device Management** > **Advanced** > **EPM**, and removing one or more **Port Numbers**.

For clustering or failover, perform this step on the control node/active unit; the module state is replicated to the other nodes.

You must rejoin all child ports for a given interface.

Step 3 Click **Apply**.

The configuration is applied to the firewall.

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 reload. For clustering or failover, zero downtime is not supported, so make sure to perform this procedure during a maintenance window.

Procedure

Step 1 Install the network module according to the hardware installation guide. You can install the network module while the firewall is powered on.

For clustering or failover, install the network module on all nodes.

Step 2 Reload the firewall; see **Tools > System Reload**.

For clustering or failover, reload all nodes. Because nodes with different network modules cannot join the cluster/failover pair, you need to reload all nodes with the new module before they can reform the cluster/failover pair.

Step 3 Enable the network module by choosing **Configuration** > **Device Management** > **Advanced** > **EPM**, and unchecking **Disable Netmod**.

For clustering or failover, perform this step on the control node/active unit; the module state is replicated to the other nodes.

Step 4 Click **Apply**.

The configuration is applied 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 reload. 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 failover, you cannot disable a network module if the cluster control link/failover link is on the module.

Procedure

Step 1 For clustering or failover, 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 392); then disable clustering on the node. See Become an Inactive Node, on page 388 or Deactivate a Data Node from the Control Node, on page 389.

  If the cluster control link is on the network module, you must leave the cluster. See Leave the Cluster, on page 391. Disabling the network module with an active cluster control link is not allowed.

- **Failover**—Ensure the unit you want to perform the hot swap on is the standby node. See Force Failover, on page 316.

  If the failover link is on the network module, you must disable failover. See Disable Failover, on page 316. Disabling the network module with an active failover link is not allowed.

Step 2 Disable the network module by choosing **Configuration** > **Device Management** > **Advanced** > **EPM**, and checking **Disable Netmod**.

Step 3 Click **Apply**.

The configuration is applied to the firewall.

Step 4 Replace the network module according to the hardware installation guide. You can replace the network module while the firewall is powered on.

Step 5 Enable the network module by choosing **Configuration** > **Device Management** > **Advanced** > **EPM**, and unchecking **Disable Netmod**.

Step 6 Click **Apply**.

The configuration is applied to the firewall.

Step 7 For clustering or failover, perform the following steps.

- **Clustering**—Add the node back to the cluster. See Rejoin the Cluster, on page 390 or Add a New Data Node from the Control Node, on page 387.
- **Failover**—If you disabled failover, then reform failover.

Replace the Network Module with a Different Type

If you replace a network module with a different type, then a reload 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 failover, zero downtime is not supported, so make sure to perform this procedure during a maintenance window.

Procedure

Step 1 Disable the network module by choosing **Configuration** > **Device Management** > **Advanced** > **EPM**, and checking **Disable Netmod**.

For clustering or failover, perform this step on the control node/active unit; the module state is replicated to the other nodes.

Step 2 Click **Apply**.

The configuration is applied to the firewall. Do not save the configuration; when you reload, the module will be enabled using the saved configuration.

Step 3 Replace the network module according to the hardware installation guide. You can replace the network module while the firewall is powered on.

For clustering or failover, install the network module on all nodes.

Step 4 Reload the firewall; see **Tools > System Reload**.

For clustering or failover, reload all nodes. Because nodes with different network modules cannot join the cluster/failover pair, you need to reload all nodes with the new module before they can reform the cluster/failover pair.

Step 5 If you saved the configuration before reloading, you will have to reenable the module.

Remove the Network Module

If you want to permanently remove the network module, follow these steps. Removing a network module requires a reload. For clustering or failover, zero downtime is not supported, so make sure to perform this procedure during a maintenance window.

Before you begin

For clustering or failover, make sure the cluster/failover link is not on the network module; you cannot remove the module in this case.

Procedure

Step 1 Disable the network module by choosing **Configuration** > **Device Management** > **Advanced** > **EPM**, and checking **Disable Netmod**.

For clustering or failover, perform this step on the control node/active unit; the module state is replicated to the other nodes.

Step 2 Click **Apply**, and the **Save**.

The configuration is saved to the firewall.

Step 3 Remove the network module according to the hardware installation guide. You can remove the network module while the firewall is powered on.

For clustering or failover, remove the network module on all nodes.

Step 4 Reload the firewall; see **Tools > System Reload**.

For clustering or failover, reload all nodes. Because nodes with different network modules cannot join the cluster/failover pair, you need to reload all nodes without the module before they can reform the cluster/failover pair.

Examples for Basic Interfaces

See the following configuration examples.

Physical Interface Parameters Example

The following example configures parameters for the physical interface in single mode:

```
interface gigabitethernet 0/1
speed 1000
duplex full
```

```
no shutdown
```

Multiple Context Mode Example

The following example configures interface parameters in multiple context mode for the system configuration, and allocates the gigabitethernet 0/1.1 subinterface to contextA:

```
interface gigabitethernet 0/1
speed 1000
duplex full
no shutdown
interface gigabitethernet 0/1.1
vlan 101
context contextA
allocate-interface gigabitethernet 0/1.1
```

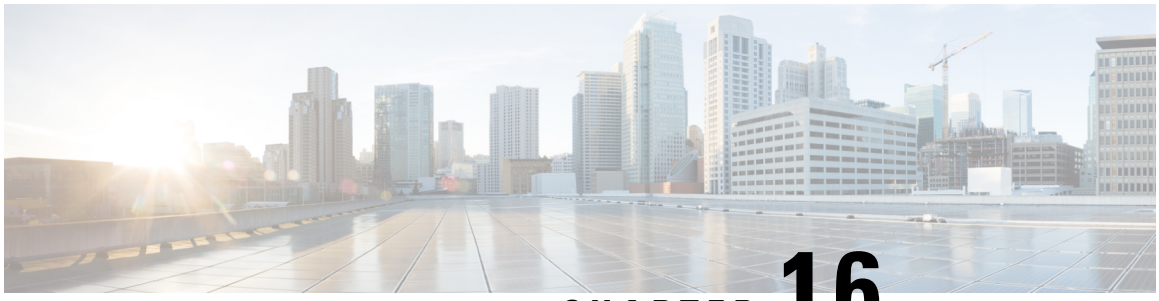
History for Basic Interface Configuration

Table 26: History for Interfaces

Feature Name	Releases	Feature Information
Support for Secure Firewall 1200 Series	9.22(1)	Features such as flow control, FEC, detect SFP, and auto-negotiation are supported.
Default Forward Error Correction (FEC) on Secure Firewall 3100 fixed ports changed to cl108-rs from cl74-fc for 25 GB+ SR, CSR, and LR transceivers	9.18(3) / 9.19(1)	When you set the FEC to Auto on the Secure Firewall 3100 fixed ports, the default type is now set to cl108-rs instead of cl74-fc for 25 GB SR, CSR, and LR transceivers. New/Modified screens: **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces** > **Edit Interface** > **Configure Hardware Properties** > **FEC Mode**
Pause Frames for Flow Control for the Secure Firewall 3100	9.18(1)	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: **Configuration** > **Device Settings** > **Interfaces** > **General**
Breakout ports for the Secure Firewall 3130 and 3140	9.18(1)	You can now configure four 10GB breakout ports for each 40GB interface on the Secure Firewall 3130 and 3140. New/Modified screens: **Configuration** > **Device Management** > **Advanced** > **EPM**

Feature Name	Releases	Feature Information
Support for hot swapping the network module for the Secure Firewall 3100	9.17(1)	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: **Configuration** > **Device Management** > **Advanced** > **EPM**
Support for Forward Error Correction for the Secure Firewall 3100	9.17(1)	Secure Firewall 3100 25 Gbps interfaces support Forward Error Correction (FEC). FEC is enabled by default and set to Auto. New/Modified screens: **Configuration** > **Device Settings** > **Interfaces** > **Edit Interface** > **Configure Hardware Properties**
Support for setting the speed based on the SFP for the Secure Firewall 3100	9.17(1)	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: **Configuration** > **Device Settings** > **Interfaces** > **Edit Interface** > **Configure Hardware Properties**
Secure Firewall 3100 auto-negotiation can be enabled or disabled for 1Gigabit and higher interfaces.	9.17(1)	Secure Firewall 3100 auto-negotiation can be enabled or disabled separately from speed for 1Gigabit and higher interfaces. New/Modified screens: **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces** > **Advanced**
Speed auto-negotiation can be disabled on fiber interfaces on the Firepower 1100 and 2100	9.14(1)	You can now configure a Firepower 1100 or 2100 fiber interface to disable auto-negotiation. For 10GB interfaces, you can configure the speed down to 1GB without auto-negotiation; you cannot disable auto-negotiation for an interface with the speed set to 10GB. New/Modified screens: **Configuration** > **Device Settings** > **Interfaces** > **Edit Interface** > **Configure Hardware Properties** > **Speed**
Through traffic support on the Management 0/0 interface for the ASA virtual	9.6(2)	You can now allow through traffic on the Management 0/0 interface on the ASA virtual. Previously, only the ASA virtual on Microsoft Azure supported through traffic; now all ASA virtuals support through traffic. You can optionally configure this interface to be management-only, but it is not configured by default.

Feature Name	Releases	Feature Information
Support for Pause Frames for Flow Control on Gigabit Ethernet Interfaces	8.2(5)/8.4(2)	You can now enable pause (XOFF) frames for flow control for Gigabit Ethernet interfaces on all ASA models. We modified the following screens: (Single Mode) Configuration > Device Setup > Interface Settings > Interfaces > Add/Edit Interface > General (Multiple Mode, System) Configuration > Interfaces > Add/Edit Interface.
Support for Pause Frames for Flow Control on the ASA 5580 Ten Gigabit Ethernet Interfaces	8.2(2)	You can now enable pause (XOFF) frames for flow control. This feature is also supported on the ASA 5585-X. We modified the following screens: (Single Mode) Configuration > Device Setup > Interface Settings > Interfaces > Add/Edit Interface > General (Multiple Mode, System) Configuration > Interfaces > Add/Edit Interface.
Jumbo packet support for the ASA 5580	8.1(1)	The ASA 5580 supports jumbo frames. A jumbo frame is an Ethernet packet larger than the standard maximum of 1518 bytes (including Layer 2 header and FCS), up to 9216 bytes. You can enable support for jumbo frames for all interfaces by increasing the amount of memory to process Ethernet frames. Assigning more memory for jumbo frames might limit the maximum use of other features, such as ACLs. This feature is also supported on the ASA 5585-X. We modified the following screen: Configuration > Device Setup > Interface Settings > Interfaces > Add/Edit Interface > Advanced.
Gigabit Ethernet Support for the ASA 5510 Security Plus License	7.2(3)	The ASA 5510 now supports GE (Gigabit Ethernet) for port 0 and 1 with the Security Plus license. If you upgrade the license from Base to Security Plus, the capacity of the external Ethernet0/0 and Ethernet0/1 ports increases from the original FE (Fast Ethernet) (100 Mbps) to GE (1000 Mbps). The interface names will remain Ethernet 0/0 and Ethernet 0/1.
Increased interfaces for the Base license on the ASA 5510	7.2(2)	For the Base license on the ASA 5510, the maximum number of interfaces was increased from 3 plus a management interface to unlimited interfaces.



CHAPTER 16

Basic Interface Configuration for Firepower 1010 and Secure Firewall 1210/1220 Switch Ports

You can configure each Firepower 1010 or Secure Firewall 1210/1220 interface to run as a regular firewall interface or as a Layer 2 hardware switch port. This chapter 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 chapter also describes how to customize Power over Ethernet (PoE) on supported interfaces.

- About Switch Ports, on page 575
- Guidelines and Limitations for Switch Ports, on page 576
- Configure Switch Ports and Power Over Ethernet, on page 578
- Monitoring Switch Ports, on page 582
- History for Switch Ports, on page 583

About Switch Ports

This section describes the switch ports of the 1010/1210/1220.

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.
- 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 ASA 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, or EtherChannel interfaces. When a switch port needs to communicate with another network, then the ASA 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 ASA 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.

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

Context Mode

- The Firepower 1010 does not support multiple context mode.
- The Secure Firewall 1210/1220 only supports multiple context mode when you do not use switch ports.

Failover and Clustering

- No cluster support.

- Active/Standby failover support only.
- You should not use the switch port functionality when using Failover. Because the switch ports operate in hardware, they continue to pass traffic on both the active *and* the standby units. Failover is designed to prevent traffic from passing through the standby unit, but this feature does not extend to switch ports. In a normal Failover 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 Failover, 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.
- 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 Manual MAC Address, MTU, and TCP MSS, on page 670.
  - 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 Manual MAC Address, MTU, and TCP MSS, on page 670.

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
- Policy based routing
- Equal-Cost Multi-Path routing (ECMP)
- VXLAN
- EtherChannels—Switch ports cannot be part of an EtherChannel. PoE is also not supported on a port in an EtherChannel.
- Failover and state link
- Traffic zones
- 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 the 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.

Configure a VLAN Interface

This section describes how to configure VLAN interfaces (SVIs) for use with associated switch ports. You can create up to 60 VLAN interfaces associated with switch ports.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces**, and choose **Add** > **VLAN Interface**.

Step 2 In the **VLAN ID** field, enter the VLAN ID for this interface, between 1 and 4070, excluding IDs in the range 3968 to 4047, which are reserved for internal use.

Step 3 (Optional) In the **Block Traffic From this Interface to** drop-down list, choose the VLAN to which this VLAN interface cannot initiate traffic.

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 use the **Block Traffic From this Interface to** option on the home VLAN; the business network can access the home network, but the home network cannot access the business network.

Step 4 Click **OK**.

Step 5 Click **Apply**.

Configure Switch Ports as Access Ports

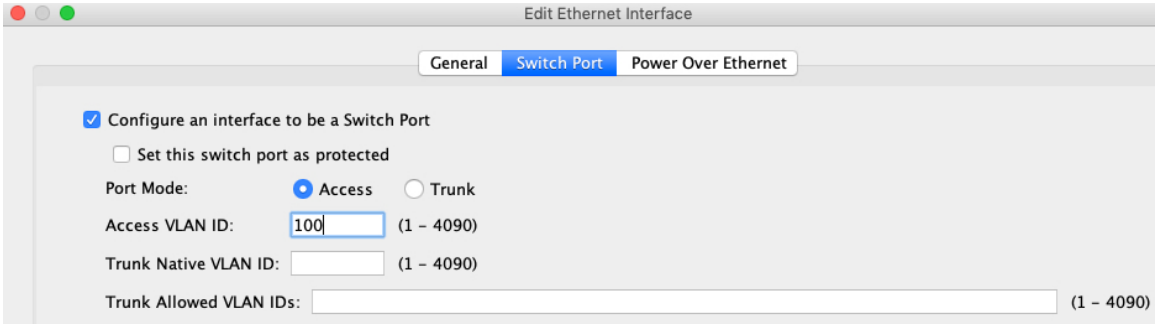
To assign a switch port to a single VLAN, configure it as an access port. Access ports accept only untagged traffic. By default, Ethernet 1/2 through Ethernet 1/8 switch ports are enabled and assigned to VLAN 1 on Firepower 1010 and Secure Firewall 1210. On the Secure Firewall 1220, by default, Ethernet 1/2 through Ethernet 1/10 switch ports are enabled and assigned to VLAN 1.

Note The device does not support Spanning Tree Protocol for loop detection in the network. Therefore you must ensure that any connection with the ASA does not end up in a network loop.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces**, select the interface you want to edit, and click **Edit**.

Step 2 Click **Switch Port**.



Step 3 Check the **Configure an interface to be a Switch Port** check box.

Step 4 (Optional) Check the **Set this switch port as protected** check box to 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 apply the **Set this switch port as protected** option to 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 5 For the **Port Mode**, click the **Access** radio button.

Step 6 Enter the **Access VLAN ID** associated with this switch port, between 1 and 4070.

The default is VLAN 1.

Step 7 Click **General**.

Step 8 Check **Enable Interface**.

Note
Other fields on the **General** page, such as the **Interface Name**, are not applicable to switch ports.

Step 9 (Optional) Set hardware properties.

a) Click **Configure Hardware Properties**.

b) Choose the **Duplex**.

The default is **Auto**.

c) Choose the **Speed**.

The default is **Auto**.

d) Click **OK**.

Step 10 Click **OK**.

Step 11 Click **Apply**.

Configure Switch Ports as Trunk Ports

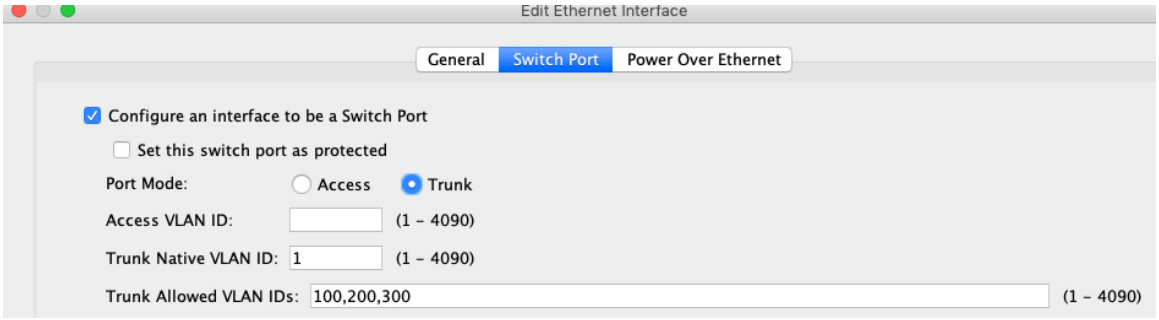
This procedure describes how to create a trunk port that can carry multiple VLANs using 802.1Q tagging. Trunk ports accept untagged and tagged traffic. Traffic on allowed VLANs pass through the trunk port unchanged.

When the trunk receives untagged traffic, it tags it to the native VLAN ID so that the ASA can forward the traffic to the correct switch ports, or can route it to another firewall interface. When the ASA 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 Choose **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces**, select the interface you want to edit, and click **Edit**.

Step 2 Click **Switch Port**.



Step 3 Check the **Configure an interface to be a Switch Port** check box.

Step 4 (Optional) Check the **Set this switch port as protected** check box to 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 apply the **Set this**

switch port as protected option to 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 5 For the **Port Mode**, click the **Trunk** radio button.

Step 6 Enter the **Trunk Native VLAN ID**, between 1 and 4070. The default is VLAN 1.

Each port can only have one native VLAN, but every port can have either the same or a different native VLAN.

Step 7 Enter the **Trunk Allowed VLAN IDs** associated with this switch port, separated by commas, between 1 and 4070.

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 Click **General**.

Step 9 Check **Enable Interface**.

Note
Other fields on the **General** page, such as the **Interface Name**, are not applicable to switch ports.

Step 10 (Optional) Set hardware properties.

a) Click **Configure Hardware Properties**.
b) Choose the **Duplex**.

The default is **Auto**.

c) Choose the **Speed**.

The default is **Auto**.

d) Click **OK**.

Step 11 Click **OK**.

Step 12 Click **Apply**.

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.

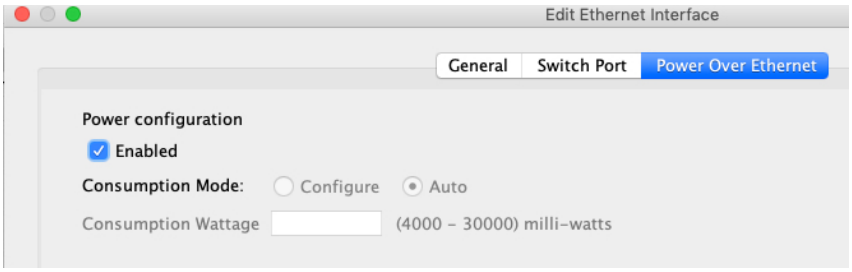
Before you begin

In multiple context mode, perform these steps in the system execution space.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces**, select the interface you want to edit, and click **Edit**.

Step 2 Click **Power Over Ethernet**.



Step 3 Check **Enabled**.

Step 4 Click the **Consumption Mode**: **Configure** or **Auto** radio button.

- **Auto**—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.
- **Configure**—Manually specifies the wattage in milliwatts in the **Consumption Wattage** field, from 4000 to 30000 (1010) or 90000 (1210CP). Use this option if you want to set the watts manually and disable LLDP negotiation.

Step 5 Click **OK**.

Step 6 Click **Apply**.

Step 7 Choose **Monitor** > **Interfaces** > **Power Over Ethernet** to view the current PoE status.

Monitoring Switch Ports

- **Monitoring** > **Interfaces** > **ARP Table**

  Displays the ARP table, including static and dynamic entries. The ARP table includes entries that map a MAC address to an IP address for a given interface.

- **Monitoring** > **Interfaces** > **MAC Address Table**

  Shows the static and dynamic MAC address entries.

- **Monitoring** > **Interfaces** > **Interface Graphs**

  Shows interface statistics in graph or table form.

- **Monitoring** > **Interfaces** > **L2 Switching**

  Shows the VLAN–to–switch port association and the static and dynamic MAC address entries.

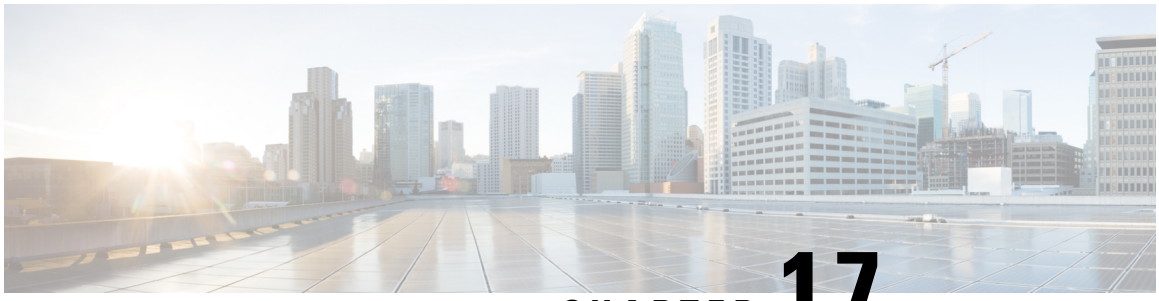
- **Monitoring** > **Interfaces** > **Power Over Ethernet**

  Shows the PoE+ status.

History for Switch Ports

Table 27: History for Switch Ports

Feature Name	Version	Feature Information
Secure Firewall 1210CP IEEE 802.3bt support (PoE++ and Hi-PoE)	9.23(1)	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: • **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces** > **Edit** > **Power Over Ethernet** • **Monitoring** > **Interfaces** > **Power Over Ethernet**
Secure Firewall 1210/1220 hardware switch support	9.22(1)	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	9.22(1)	The Secure Firewall 1210CP supports Power over Ethernet+ (PoE+) on Ethernet ports 1/5-1/8.
Firepower 1010 hardware switch support	9.13(1)	The Firepower 1010 supports setting each Ethernet interface to be a switch port or a firewall interface. New/Modified screens: • **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces** > **Edit** > **Switch Port** • **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces** > **Add VLAN Interface** • **Monitoring** > **Interfaces** > **L2 Switching**
Firepower 1010 PoE+ support on Ethernet 1/7 and Ethernet 1/8	9.13(1)	The Firepower 1010 supports Power over Ethernet+ (PoE+) on Ethernet 1/7 and Ethernet 1/8. New/Modified screens: • **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces** > **Edit** > **Power Over Ethernet** • **Monitoring** > **Interfaces** > **Power Over Ethernet**



CHAPTER 17

EtherChannel Interfaces

This chapter tells how to configure EtherChannels interfaces.

Note For multiple context mode, complete all tasks in this section in the system execution space. If you are not already in the system execution space, in the Configuration > Device List pane, double-click **System** under the active device IP address.

For ASA cluster interfaces, which have special requirements, see ASA Cluster for the Secure Firewall 3100/4200, on page 341.

Note For Firepower 4100/9300 chassis, EtherChannel interfaces are configured in the FXOS operating system. See the configuration or getting started guide for your chassis for more information.

- About EtherChannels, on page 585
- Guidelines for EtherChannels, on page 588
- Default Settings for EtherChannels Interfaces, on page 590
- Configure an EtherChannel, on page 590
- Examples for EtherChannels, on page 594
- History for EtherChannels, on page 595

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.

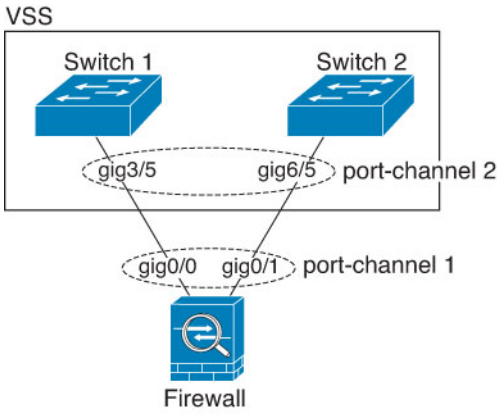
The EtherChannel aggregates the traffic across all the available active interfaces in the channel. The interface is selected using a proprietary hash algorithm, based on source or destination MAC addresses, IP addresses, TCP and UDP port numbers and VLAN numbers.

Connecting to an EtherChannel on Another Device

The device to which you connect the ASA 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 ASA 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.

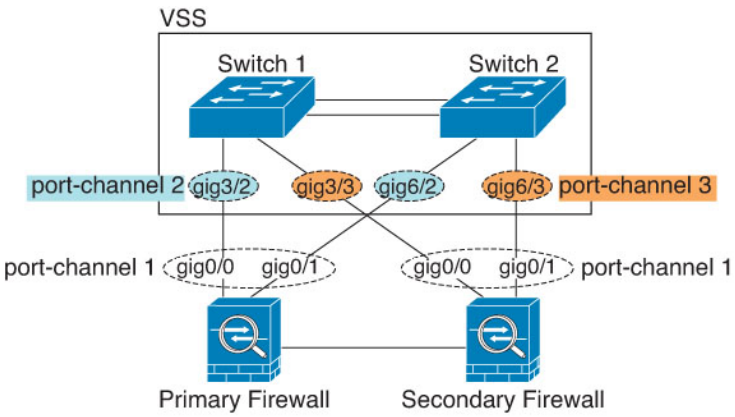
Figure 76: Connecting to a VSS/vPC



Note If the ASA device is in transparent firewall mode, and you place the ASA device between two sets of VSS/vPC switches, then be sure to disable Unidirectional Link Detection (UDLD) on any switch ports connected to the ASA 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 ASA device in an Active/Standby failover deployment, then you need to create separate EtherChannels on the switches in the VSS/vPC, one for each ASA device. On each ASA device, a single EtherChannel connects to both switches. Even if you could group all switch interfaces into a single EtherChannel connecting to both ASA devices (in this case, the EtherChannel will not be established because of the separate ASA system IDs), a single EtherChannel would not be desirable because you do not want traffic sent to the standby ASA device.

Figure 77: 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 ASA 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.

For a spanned EtherChannel in clustering, load balancing occurs on a per ASA basis. For example, if you have 32 active interfaces in the spanned EtherChannel across 8 ASAs, with 4 interfaces per ASA in the EtherChannel, then load balancing only occurs across the 4 interfaces on the ASA.

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.

Related Topics

Customize the EtherChannel, on page 592

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. In multiple context mode, you can automatically assign unique MAC addresses to *shared* interfaces, including an EtherChannel port 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, ASA-defined EtherChannels are not supported as bridge group members. EtherChannels on the Firepower 4100/9300 can be bridge group members.

Failover

- When you use an EtherChannel interface as a Failover link, it must be pre-configured on both units in the Failover pair; you cannot configure it on the primary unit and expect it to replicate to the secondary unit because *the Failover 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 Failover. 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 Failover. 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 Failover 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 Failover link. To alter the configuration, you need to temporarily disable Failover, which prevents Failover from occurring for the duration.

Model Support

- You cannot add EtherChannels in ASA for the Firepower 4100/9300, or the ASA 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.

Clustering

- To configure a spanned EtherChannel or an individual cluster interface, see the clustering chapter.

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, 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, 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 ASA EtherChannel must also support 802.3ad EtherChannels.
- The ASA 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 ASA device will drop the tagged LACPDUs. Be sure to disable native VLAN tagging on the neighboring switch. In multiple context mode, these messages are not included in a packet capture, so that you cannot diagnose the issue easily.
- 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, ASA did not support connecting an EtherChannel to a switch stack. With default switch settings, if the ASA 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 ASA configuration refers to the logical EtherChannel interface instead of the member physical interfaces.
- You must first remove the breakout ports from the port channel membership before you can delete an port channel with breakout ports. Otherwise, the breakout ports will show as unassociated when joining them back after deleting the port channel. This is not applicable if a port channel has only fixed ports and no breakout ports.

Default Settings for EtherChannels Interfaces

This section lists default settings for interfaces if you do not have a factory default configuration.

Default State of Interfaces

The default state of an interface depends on the type and the context mode.

In multiple context mode, all allocated interfaces are enabled by default, no matter what the state of the interface is in the system execution space. However, for traffic to pass through the interface, the interface also has to be enabled in the system execution space. If you shut down an interface in the system execution space, then that interface is down in all contexts that share it.

In single mode or in the system execution space, interfaces have the following default states:

- Physical interfaces—Disabled.
- EtherChannel port-channel interfaces—Enabled. However, for traffic to pass through the EtherChannel, the channel group physical interfaces must also be enabled.

Configure an EtherChannel

This section describes how to create an EtherChannel port-channel interface, assign interfaces to the EtherChannel, and customize the EtherChannel.

Add Interfaces to the EtherChannel

This section describes how to create an EtherChannel port-channel interface and assign interfaces to the EtherChannel. By default, port-channel interfaces are enabled.

Before you begin

- You can configure up to 48 EtherChannels, depending on how many interfaces your model has.
- 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.
- To configure a spanned EtherChannel for clustering, see the clustering chapter instead of this procedure.
- All interfaces in the channel group must be 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, except for the Secure Firewall 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..
- You cannot add a physical interface to the channel group if you configured a name for it. You must first remove the name in the **Configuration > Device Setup > Interface Settings > Interfaces** pane.
- For multiple context mode, complete this procedure in the system execution space. If you are not already in the System configuration mode, in the **Configuration** > **Device List** pane, double-click **System** under the active device IP address.



Caution 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 (Optional) (Secure Firewall 3100/4200 only) Set the LACP data unit receive rate for a physical interface in the channel group.

lacp rate {**normal** | **fast**}

The default is **normal** (slow, every 30 seconds). The **fast** option receives LACP data units every second. You should match the setting on the connected switch.

Step 2 Depending on your context mode:

- For single mode, choose the **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces** pane.
- For multiple mode in the System execution space, choose the **Configuration** > **Context Management** > **Interfaces** pane.

Step 3 Choose **Add** > **EtherChannel Interface**.

The **Add EtherChannel Interface** dialog box appears.

Note
In single mode, this procedure only covers a subset of the parameters on the Edit EtherChannel Interface dialog box. Note that in multiple context mode, before you complete your interface configuration, you need to allocate interfaces to contexts. See Configure Multiple Contexts, on page 260.

Step 4 In the **Port Channel ID** field, enter a number between 1 and 48 (1 and 8 for the Firepower 1010 and 1210, and between 1 and 10 for the 1220).

Step 5 In the **Available Physical Interface** area, click an interface and then click **Add** to move it to the **Members in Group** area.

In transparent mode, if you create a channel group with multiple Management interfaces, then you can use this EtherChannel as the management-only interface.

Note
If you want to set the EtherChannel mode to On, then you must include only one interface initially. After you complete this procedure, edit the member interface, and set the mode to **On**. Apply your changes, then edit the EtherChannel to add more member interfaces.

Step 6 Repeat for each interface you want to add to the channel group.

Make sure all interfaces are the same type and speed. The first interface you add determines the type and speed of the EtherChannel. Any non-matching interfaces you add will be put into a suspended state. ASDM does not prevent you from adding non-matching interfaces.

Step 7 Click **OK**.

You return to the **Interfaces** pane. The member interfaces now show a lock to the left of the interface ID showing that only basic parameters can be configured for it. The EtherChannel interface is added to the table.



Step 8 (Optional) (Secure Firewall 3100/4200 only) Set the LACP data unit receive rate for a physical interface in the channel group.

a) Click the physical interface in the **Interfaces** table, and click **Edit**.

The **Edit Interface** dialog box appears.

b) Click the **Advanced** tab.

c) In the **EtherChannel** area, from the **Rate** drop down list, choose **Normal** or **Fast**.

The default is **Normal** (slow, every 30 seconds). The **Fast** option received LACP updates every second. You should match the setting on the connected switch.

Step 9 Click **Apply**. All member interfaces are enabled automatically.

Related Topics

Link Aggregation Control Protocol, on page 587
Customize the EtherChannel, on page 592

Customize the EtherChannel

This section describes how to set the maximum number of interfaces in the EtherChannel, the minimum number of operating interfaces for the EtherChannel to be active, the load balancing algorithm, and other optional parameters.

Procedure

Step 1 Depending on your context mode:

- For single mode, choose the **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces** pane.
- For multiple mode in the System execution space, choose the **Configuration** > **Context Management** > **Interfaces** pane.

Step 2 Click the port-channel interface you want to customize, and click **Edit**.

The **Edit Interface** dialog box appears.

Step 3 To override the media type, duplex, speed, and pause frames for flow control for all member interfaces, click **Configure Hardware Properties**. This method provides a shortcut to set these parameters because these parameters must match for all interfaces in the channel group.

Step 4 (Optional; ISA 3000 only) To customize the EtherChannel, click the **Advanced** tab.

a) In the **EtherChannel** area, from the **Minimum** 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.

b) From the **Maximum** 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.

c) From the **Load Balance** drop-down list, select the criteria used to load balance the packets across the group channel interfaces. By default, the ASA 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 587.

d) For **Secure Group Tagging** settings, see the firewall configuration guide.

e) For **ASA Cluster** settings, see (Recommended; Required in Multiple Context Mode) Configure Interfaces on the Control Node, on page 364.

Step 5 Click **OK**.

You return to the **Interfaces** pane.

Step 6 To set the mode and priority for a physical interface in the channel group:

a) Click the physical interface in the **Interfaces** table, and click **Edit**.

The **Edit Interface** dialog box appears.

b) Click the **Advanced** tab.

c) In the **EtherChannel** area, from the **Mode** drop down list, choose **Active**, **Passive**, or **On**. We recommend using Active mode (the default).

d) (Optional; ISA 3000 only) In the **LACP Port Priority** field, set the port priority between 1 and 65535. The default is 32768. The higher the number, the lower the priority. The ASA uses this setting to decide which interfaces are active and which are standby if you assign more interfaces than can be used. If the port priority setting is the same for all interfaces, then the priority is determined by the interface ID (slot/port). The lowest interface ID is the highest priority. For example, GigabitEthernet 0/0 is a higher priority than GigabitEthernet 0/1.

If you want to prioritize an interface to be active even though it has a higher interface ID, then set this command to have a lower value. For example, to make GigabitEthernet 1/3 active before GigabitEthernet 0/7, then make the priority value be 12345 on the 1/3 interface vs. the default 32768 on the 0/7 interface.

If the device at the other end of the EtherChannel has conflicting port priorities, the system priority is used to determine which port priorities to use. See Step 9 to set the system priority.

Step 7 Click **OK**.

You return to the **Interfaces** pane.

Step 8 Click **Apply**.

Step 9 (Optional; ISA 3000 only) To set the LACP system priority, perform the following steps. If the device at the other end of the EtherChannel has conflicting port priorities, the system priority is used to determine which port priorities to use. See Step 6d for more information.

a) Depending on your context mode:

- For single mode, choose the **Configuration** > **Device Setup** > **EtherChannel** pane.
- For multiple mode in the System execution space, choose the **Configuration** > **Context Management** > **EtherChannel** pane.

b) In the **LACP System Priority** field, enter a priority between 1 and 65535.

The default is 32768.

Related Topics

Load Balancing, on page 587
Add Interfaces to the EtherChannel, on page 590

Examples for EtherChannels

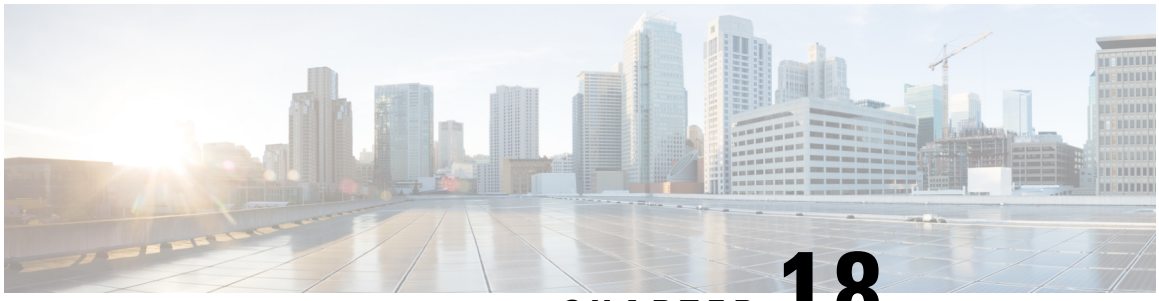
The following example configures three interfaces as part of an EtherChannel. It also sets the system priority to be a higher priority, and GigabitEthernet 0/2 to be a higher priority than the other interfaces in case more than eight interfaces are assigned to the EtherChannel.

```
lacp system-priority 1234
interface GigabitEthernet0/0
  channel-group 1 mode active
interface GigabitEthernet0/1
  channel-group 1 mode active
interface GigabitEthernet0/2
  lacp port-priority 1234
  channel-group 1 mode passive
interface Port-channel1
  lacp max-bundle 4
  port-channel min-bundle 2
  port-channel load-balance dst-ip
```

History for EtherChannels

Table 28: History for EtherChannels

Feature Name	Releases	Feature Information
EtherChannel support	8.4(1)	You can configure up to 48 802.3ad EtherChannels of eight active interfaces each. We modified or introduced the following screens: Configuration > Device Setup > Interface Settings > Interfaces Configuration > Device Setup > Interface Settings > Interfaces > Add/Edit EtherChannel Interface Configuration > Device Setup > Interface Settings > Interfaces > Add/Edit Interface Configuration > Device Setup > EtherChannel **Note** EtherChannel is not supported on the ASA 5505.
Support for 16 active links in an EtherChannel	9.2(1)	You can now configure up to 16 active links in an EtherChannel. Previously, you could have 8 active links and 8 standby links. Be sure that your switch can support 16 active links (for example the Cisco Nexus 7000 with F2-Series 10 Gigabit Ethernet Module). **Note** If you upgrade from an earlier ASA version, the maximum active interfaces is set to 8 for compatibility purposes. We modified the following screen: Configuration > Device Setup > Interface Settings > Interfaces > Add/Edit EtherChannel Interface > Advanced.



CHAPTER 18

Loopback Interfaces

This chapter tells how to configure loopback interfaces.

- About Loopback Interfaces, on page 597
- Guidelines for Loopback Interfaces, on page 598
- Configure a Loopback Interface, on page 598
- Rate-Limit Traffic to the Loopback Interface, on page 599
- History for Loopback Interfaces, on page 603

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
- SNMP
- SSH
- Static and dynamic VTI tunnels
- Syslog
- Telnet

The ASA 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 ASA's physical interfaces. You cannot configure a static route on the ASA that specifies the loopback interface.

Guidelines for Loopback Interfaces

Failover and Clustering

- No clustering support.

Context Mode

- VTI is supported in single context mode only. Other loopback uses are supported in multiple context mode.

Additional Guidelines and Limitations

- TCP sequence randomization is always disabled for traffic from the physical interface to the loopback interface.

Configure a Loopback Interface

Add a loopback interface.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces**.

Step 2 Choose **Add** > **Loopback Interface**.

The **Add Loopback Interface** dialog box appears.

Step 3 In the **Loopback ID** field, enter an integer between 0 and 10413.

Step 4 If the interface is not already enabled, check the **Enable Interface** check box.

The interface is enabled by default.

Step 5 (Optional) Enter a description in the **Description** field.

Step 6 Configure the name and IP address. See Routed and Transparent Mode Interfaces, on page 633.

Step 7 Click **OK**.

You return to the **Interfaces** pane.

Step 8 Configure rate-limiting for loopback traffic. See Rate-Limit Traffic to the Loopback Interface, on page 599.

Rate-Limit Traffic to the Loopback Interface

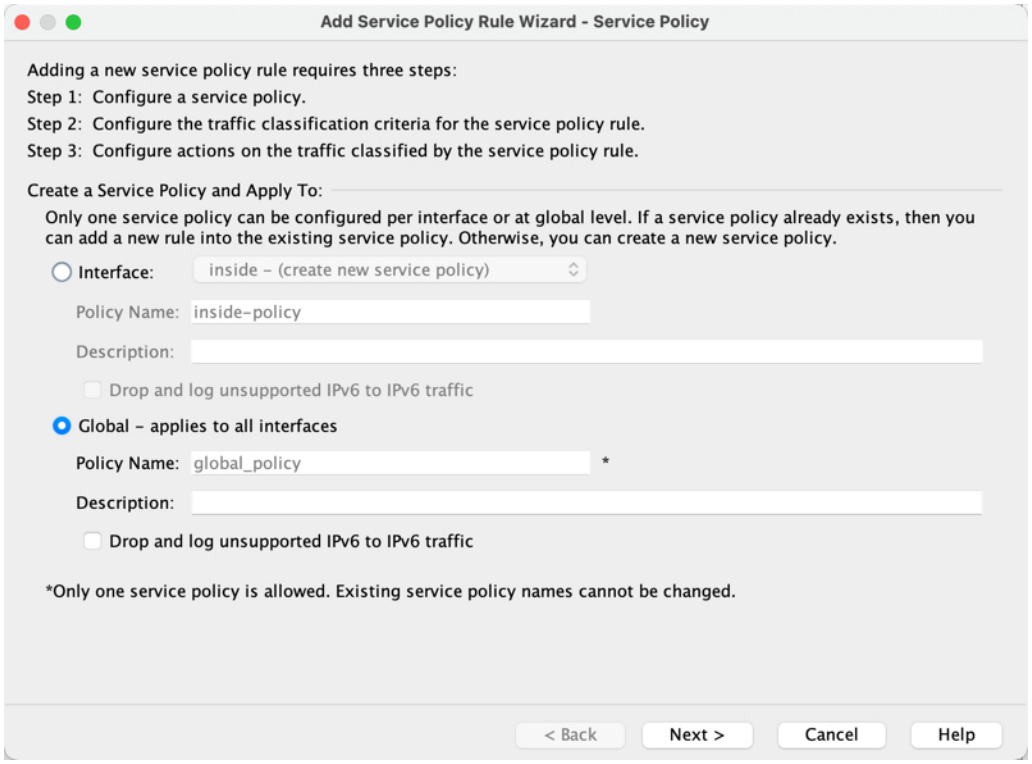
You should rate-limit traffic going to the loopback interface IP address to prevent excessive load on the system. You can add a connection limit rule to the global service policy. This procedure shows adding to the default global policy (global_policy).

Procedure

Step 1 Choose **Configuration** > **Firewall** > **Service Policy**, and click **Add** > **Add Service Policy Rule**.

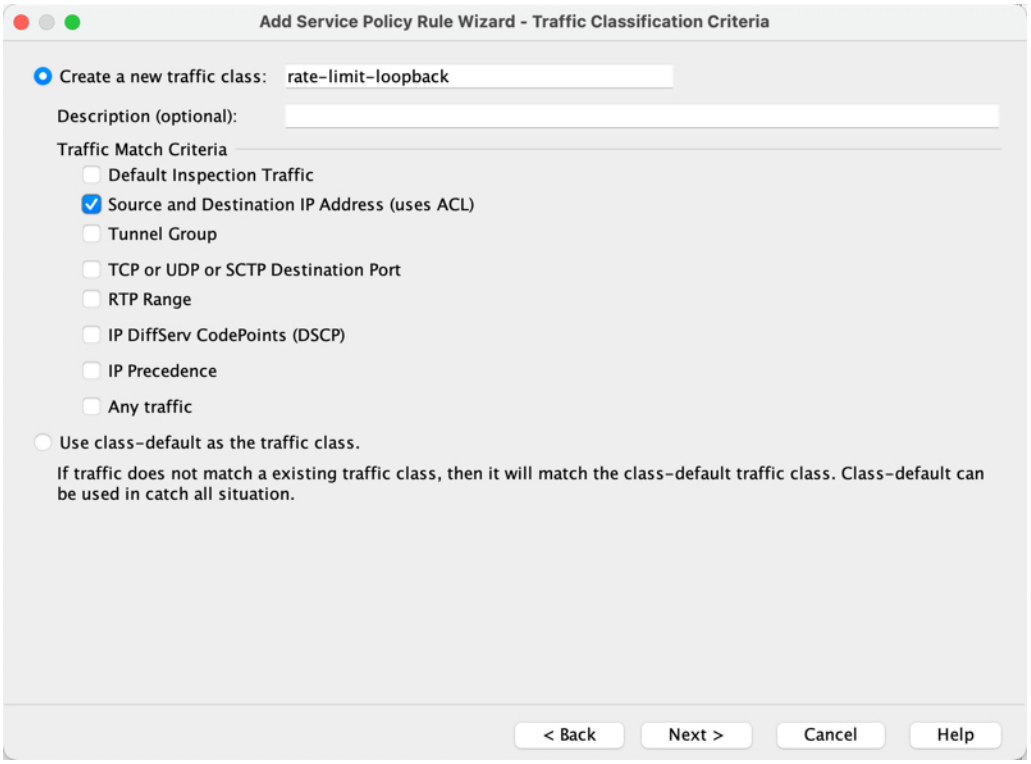
Step 2 Choose the **Global** policy and click **Next**.

Figure 78: Service Policy



Step 3 On the **Traffic Classification Criteria** page, set the following values and click **Next**.

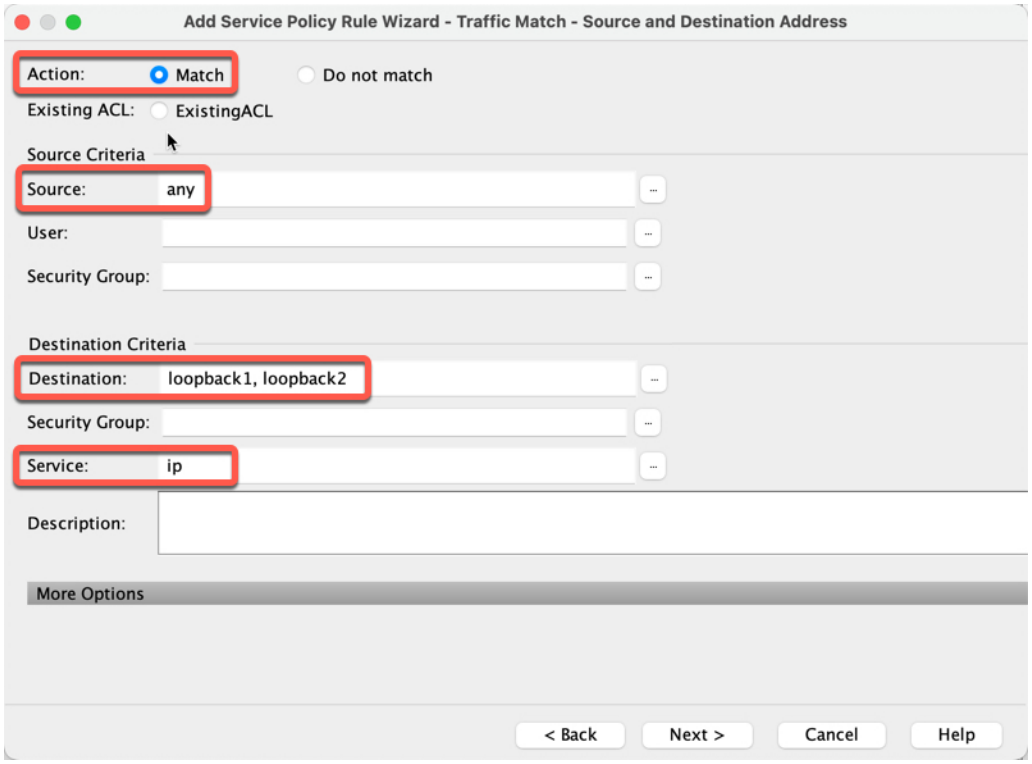
Figure 79: Traffic Classification Criteria



- **Create a new traffic class**—Name the loopback traffic class.
- **Source and Destination IP Address (uses ACL)**

Step 4 On the **Traffic Match - Source and Destination Address** page, define the access control list to specify all IP traffic going to the loopback IP address, and click **Next**.

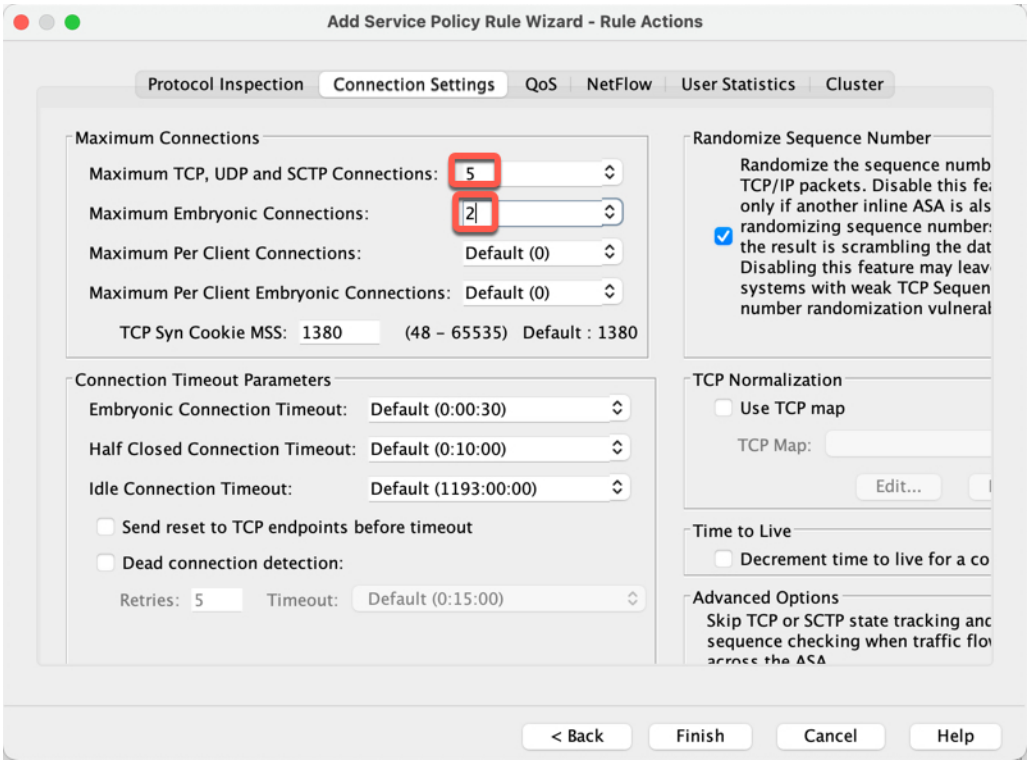
Figure 80: Traffic Match - Source and Destination Address



- **Action: Match**
- **Source**—any. You can also narrow this access list by specifiying the source IP addresses instead of **any**.
- **Destination**—The loopback interface IP addresses
- **Service**—ip

Step 5 On the **Rule Actions** page, click the **Connection Settings** tab, and in the **Maximum Connections** area, set the following values.

Figure 81: Rule Actions

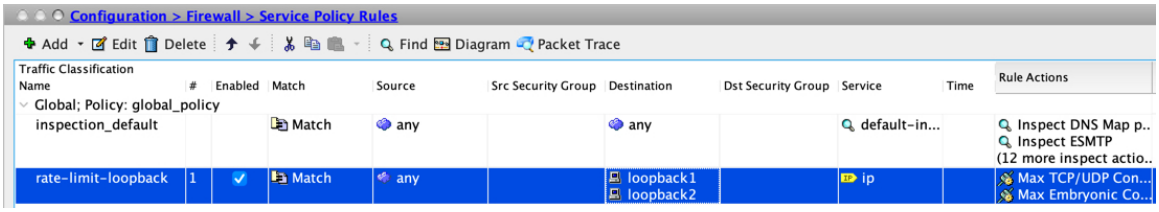


- **Maximum TCP, UDP and SCTP Connections**—Set the maximum connections to the expected number of connections for the loopback interface, and the embryonic connections to a lower number. For example, you can set it to **5/2**, or **10/5**, or **1024/512**, depending on the expected loopback interface sessions you need.
- **Embryonic Connections**—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 6 Click **Finish**.

The rule is added to the global policy.

Figure 82: Service Policy Rules Table

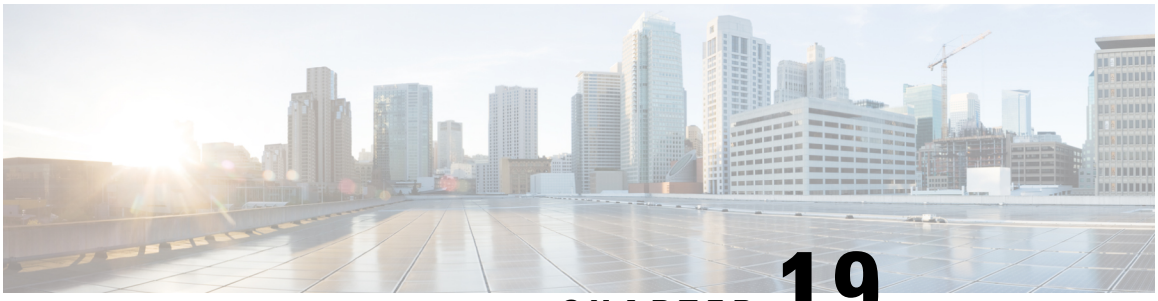


Step 7 Click **Apply**.

History for Loopback Interfaces

Table 29: History for Loopback Interfaces

Feature Name	Version	Feature Information
Loopback interface support for DNS, HTTP, ICMP, and IPsec Flow Offload	9.20(1)	You can now add a loopback interface and use it for: • DNS • HTTP • ICMP • IPsec Flow Offload
Loopback interface support for VTI	9.19(1)	A loopback interface provides redundancy of static and dynamic VTI VPN tunnels. You can now set a loopback interface as the source interface for a VTI. The VTI interface can also inherit the IP address of a loopback interface instead of a statically configured IP address. The loopback interface helps to overcome path failures. If an interface goes down, you can access all interfaces through the IP address of the loopback interface. New/Modified screens: **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces** > **Add VTI Interface** > **Advanced**
ASDM support for loopback interfaces	9.19(1)	ASDM now supports loopback interfaces. New/Modified screens: **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces** > **Add Loopback Interface**
Support for loopback interface	9.18(2)	You can now add a loopback interface and use it for: • BGP • AAA • SNMP • Syslog • SSH • Telnet New/Modified commands: **interface loopback**, **logging host**, **neighbor update-source**, **snmp-server host**, **ssh**, **telnet** No ASDM support.



CHAPTER 19

VLAN Subinterfaces

This chapter tells how to configure VLAN subinterfaces.

Note For multiple context mode, complete all tasks in this section in the system execution space. If you are not already in the system execution space, in the Configuration > Device List pane, double-click **System** under the active device IP address.

- About VLAN Subinterfaces, on page 605
- Licensing for VLAN Subinterfaces, on page 605
- Guidelines and Limitations for VLAN Subinterfaces, on page 606
- Default Settings for VLAN Subinterfaces, on page 607
- Configure VLAN Subinterfaces and 802.1Q Trunking, on page 607
- Examples for VLAN Subinterfaces, on page 609
- History for VLAN Subinterfaces, on page 610

About VLAN Subinterfaces

VLAN subinterfaces let you divide a physical 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 allow you to 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 ASAs. This feature is particularly useful in multiple context mode so that you can assign unique interfaces to each context.

You can configure a primary VLAN, as well as one or more secondary VLANs. When the ASA receives traffic on the secondary VLANs, it maps it to the primary VLAN.

Licensing for VLAN Subinterfaces

Model	License Requirement
Firepower 1010	Essentials License: 60
Firepower 1120	Essentials License: 512

Model	License Requirement
Firepower 1140, 1150	Essentials License: 1024
Secure Firewall 1210, 1220	Essentials License: 1024
Secure Firewall 1230, 1240, 1250	Essentials License: 1024
Secure Firewall 3100	Essentials License: 1024
Firepower 4100	Essentials License: 1024
Secure Firewall 4200	Essentials License: 1024
Firepower 9300	Essentials License: 1024
ASA Virtual	Throughput capability: 100 Mbps: 25 1 Gbps: 50 2 Gbps: 200 10 Gbps: 1024
ISA 3000	Essentials License: 5 Security Plus License: 100

Note For an interface to count against the VLAN limit, you must assign a VLAN to it.

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.
- For ASA models, you cannot configure subinterfaces on the Management interface. See Management Slot/Port Interface, on page 562 for subinterface support.

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 for EtherChannel links. Because the physical or EtherChannel interface must be enabled for the subinterface to pass traffic, ensure that the physical

or EtherChannel interface does not pass traffic by not configuring a name for the interface. If you want to let the physical or EtherChannel interface pass untagged packets, you can configure the name as usual.

- All subinterfaces on the same parent interface must be either bridge group members or routed interfaces; you cannot mix and match.
- The ASA 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 ASA, 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 ASA. You can automatically generate unique MAC addresses; see Automatically Assign MAC Addresses, on page 669.



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.

Default Settings for VLAN Subinterfaces

This section lists default settings for interfaces if you do not have a factory default configuration.

Default State of Interfaces

The default state of an interface depends on the type and the context mode.

In multiple context mode, all allocated interfaces are enabled by default, no matter what the state of the interface is in the system execution space. However, for traffic to pass through the interface, the interface also has to be enabled in the system execution space. If you shut down an interface in the system execution space, then that interface is down in all contexts that share it.

In single mode or in the system execution space, interfaces have the following default states:

- Physical interfaces—Disabled.
- VLAN subinterfaces—Enabled. However, for traffic to pass through the subinterface, the physical interface must also be enabled.

Configure VLAN Subinterfaces and 802.1Q Trunking

Add a VLAN subinterface to a physical or EtherChannel interface.

Before you begin

For multiple context mode, complete this procedure in the system execution space. If you are not already in the System configuration mode, in the **Configuration > Device List** pane, double-click **System** under the active device IP address.

Procedure

Step 1 Depending on your context mode:

- For single mode, choose the **Configuration > Device Setup > Interface Settings > Interfaces** pane.
- For multiple mode in the System execution space, choose the **Configuration > Context Management > Interfaces** pane.

Step 2 Choose **Add** > **Interface**.

The **Add Interface** dialog box appears.

Note
In single mode, this procedure only covers a subset of the parameters on the Edit Interface dialog box; to configure other parameters, see Routed and Transparent Mode Interfaces, on page 633. Note that in multiple context mode, before you complete your interface configuration, you need to allocate interfaces to contexts. See Configure Multiple Contexts, on page 260.

Step 3 From the **Hardware Port** drop-down list, choose the physical or port-channel interface to which you want to add the subinterface.

Step 4 If the interface is not already enabled, check the **Enable Interface** check box.

The interface is enabled by default.

Step 5 In the **VLAN ID** field, enter the VLAN ID between 1 and 4094.

Some VLAN IDs might be reserved on connected switches, so check the switch documentation for more information. For multiple context mode, you can only set the VLAN in the system configuration.

Step 6 In the **Secondary VLAN ID** field, enter one or more VLAN IDs separated by spaces, commas, or dashes (for a contiguous range).

When the ASA receives traffic on the secondary VLANs, it maps the traffic to the primary VLAN.

Step 7 In the **Subinterface ID** field, enter the subinterface ID as an integer between 1 and 4294967293.

The number of subinterfaces allowed depends on your platform. You cannot change the ID after you set it.

Step 8 (Optional) In the **Description** field, enter a description for this interface.

The description can be up to 240 characters on a single line, without carriage returns. For multiple context mode, the system description is independent of the context description. In the case of a failover or state link, the description is fixed as “LAN Failover Interface,” “STATE Failover Interface,” or “LAN/STATE Failover Interface,” for example. You cannot edit this description. The fixed description overwrites any description you enter here if you make this interface a failover or state link.

Step 9 Click **OK**.

You return to the **Interfaces** pane.

Related Topics

Licensing for VLAN Subinterfaces, on page 605

Examples for VLAN Subinterfaces

The following example configures parameters for a subinterface in single mode:

```
interface gigabitethernet 0/1
  no nameif
  no security-level
  no ip address
  no shutdown
interface gigabitethernet 0/1.1
  vlan 101
  nameif inside
  security-level 100
  ip address 192.168.6.6 255.255.255.0
  no shutdown
```

The following example shows how VLAN mapping works with the Catalyst 6500. Consult the Catalyst 6500 configuration guide on how to connect nodes to PVLANS.

ASA Configuration

```
interface GigabitEthernet1/1
  description Connected to Switch GigabitEthernet1/5
  no nameif
  no security-level
  no ip address
  no shutdown
!
interface GigabitEthernet1/1.70
  vlan 70 secondary 71 72
  nameif vlan_map1
  security-level 50
  ip address 10.11.1.2 255.255.255.0
  no shutdown
!
interface GigabitEthernet1/2
  nameif outside
  security-level 0
  ip address 172.16.171.31 255.255.255.0
  no shutdown
```

Catalyst 6500 Configuration

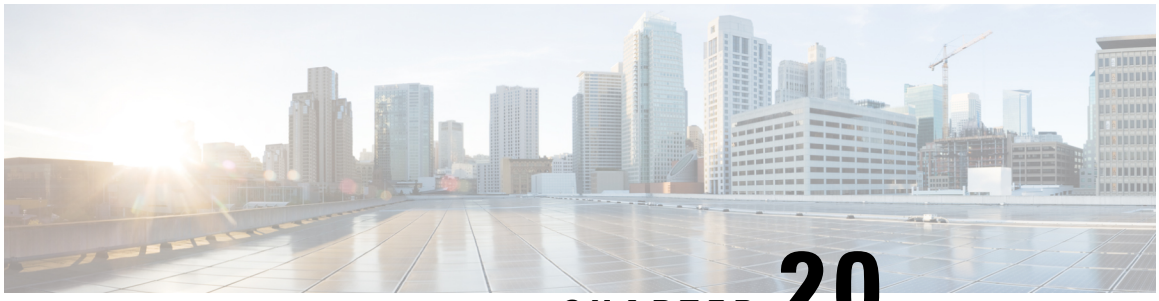
```
vlan 70
  private-vlan primary
  private-vlan association 71-72
!
vlan 71
  private-vlan community
!
vlan 72
  private-vlan isolated
!
interface GigabitEthernet1/5
  description Connected to ASA GigabitEthernet1/1
  switchport
  switchport trunk encapsulation dot1q
  switchport trunk allowed vlan 70-72
  switchport mode trunk
```

!

History for VLAN Subinterfaces

Table 30: History for VLAN Subinterfaces

Feature Name	Version	Feature Information
Increased VLANs	7.0(5)	Increased the following limits: • ASA5510 Base license VLANs from 0 to 10. • ASA5510 Security Plus license VLANs from 10 to 25. • ASA5520 VLANs from 25 to 100. • ASA5540 VLANs from 100 to 200.
Increased VLANs	7.2(2)	VLAN limits were increased for the ASA 5510 (from 10 to 50 for the Base license, and from 25 to 100 for the Security Plus license), the ASA 5520 (from 100 to 150), the ASA 5550 (from 200 to 250).
Increased VLANs for the ASA 5580	8.1(2)	The number of VLANs supported on the ASA 5580 are increased from 100 to 250.
Support to map a Secondary VLANs to a Primary VLAN	9.5(2)	You can now configure one or more secondary VLANs for a subinterface. When the ASA receives traffic on the secondary VLANs, it maps it to the primary VLAN. We modified the following screens: **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces** **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces** > **Add Interface** > **General**
Increased VLANs for the ISA 3000	9.13(1)	The maximum VLANs for the ISA 3000 with the Security Plus license increased from 25 to 100.



CHAPTER 20

VXLAN Interfaces

This chapter tells how to configure Virtual eXtensible LAN (VXLAN) interfaces. VXLANs act as Layer 2 virtual networks over Layer 3 physical networks to stretch Layer 2 networks.

- About VXLAN Interfaces, on page 611
- Requirements and Prerequisites for VXLAN Interfaces, on page 620
- Guidelines for VXLAN Interfaces, on page 620
- Default Settings for VXLAN Interfaces, on page 621
- Configure VXLAN Interfaces, on page 622
- Configure Geneve Interfaces, on page 624
- Allow Gateway Load Balancer Health Checks, on page 626
- Examples for VXLAN Interfaces, on page 626
- History for VXLAN Interfaces, on page 630

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

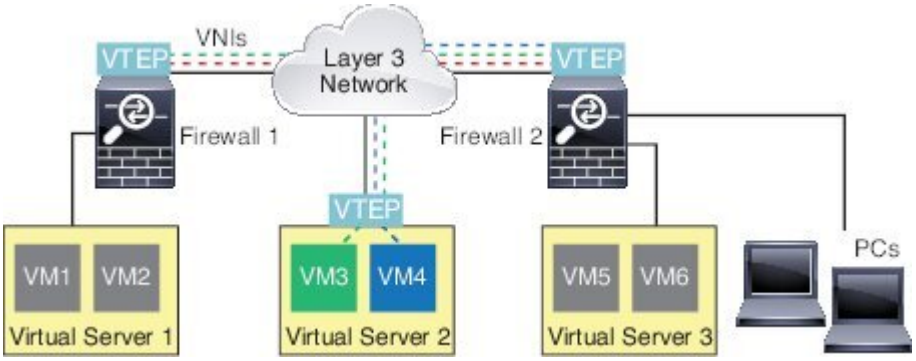
The ASA supports two types of VXLAN encapsulation:

- 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 (ASA 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 ASAs and Virtual Server 2 acting as VTEPs across a Layer 3 network, extending the VNI 1, 2, and 3 networks between sites. The ASAs 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 ASA 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 ASA interface (physical, EtherChannel, or even VLAN) with which you plan to associate all VNI interfaces. You can configure one VTEP source interface per ASA/security context. 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 ASA 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 ASA 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 ASA 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 ASA 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 VTEP

When the ASA sends a packet to a device behind a peer VTEP, the ASA needs two important pieces of information:

- The destination MAC address of the remote device
- The destination IP address of the peer VTEP

The ASA 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 ASA can find this information:

- A single peer VTEP IP address can be statically configured on the ASA.

  You cannot manually define multiple peers.

  For IPv4: The ASA then sends a VXLAN-encapsulated ARP broadcast to the VTEP to learn the end node MAC address.

  For IPv6: The ASA 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).



Note This option is not supported with Geneve.

  For IPv4: The ASA 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 ASA to learn both the remote VTEP IP address along with the destination MAC address of the remote end node.

  For IPv6: The ASA sends a Multicast Listener Discovery (MLD) Report message through the VTEP source interface to indicate that the ASA is listening on the VTEP interface for the multicast address traffic.

Geneve Peer

The ASA virtual only supports statically defined peers. You can define the ASA virtual peer IP address on the AWS Gateway Load Balancer. Because the ASA virtual never initiates traffic to the Gateway Load Balancer, you do not also have to specify the Gateway Load Balancer IP address on the ASA virtual; it learns the peer IP address when it receives Geneve traffic. Multicast groups are not supported with Geneve.

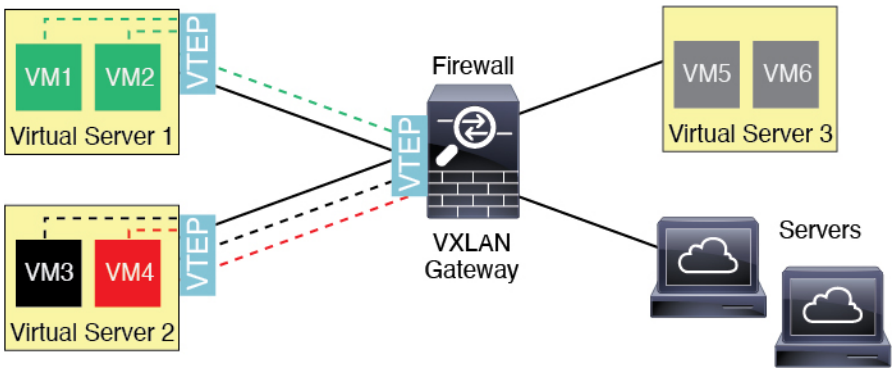
VXLAN Use Cases

This section describes the use cases for implementing VXLAN on the ASA.

VXLAN Bridge or Gateway Overview

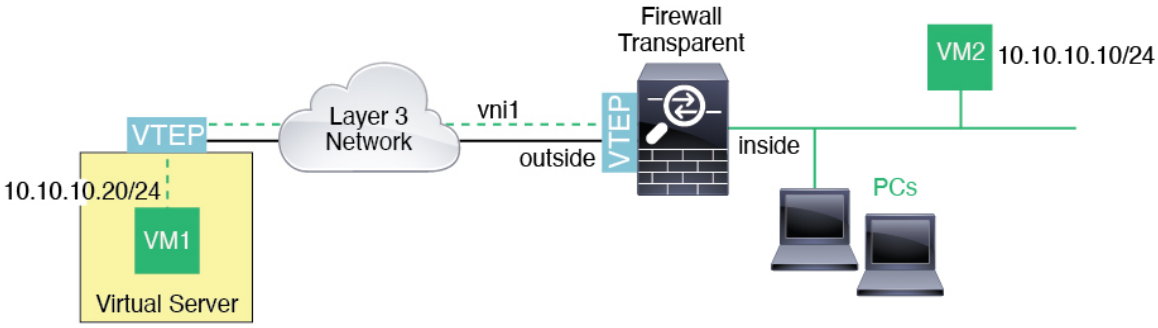
Each ASA VTEP acts as a bridge or gateway between end nodes such as VMs, servers, and PCs and the VXLAN overlay network. For incoming frames received with VXLAN encapsulation over the VTEP source interface, the ASA strips out the VXLAN header and forwards it to a physical interface connected to a non-VXLAN network based on the destination MAC address of the inner Ethernet frame.

The ASA always processes VXLAN packets; it does not just forward VXLAN packets untouched between two other VTEPs.



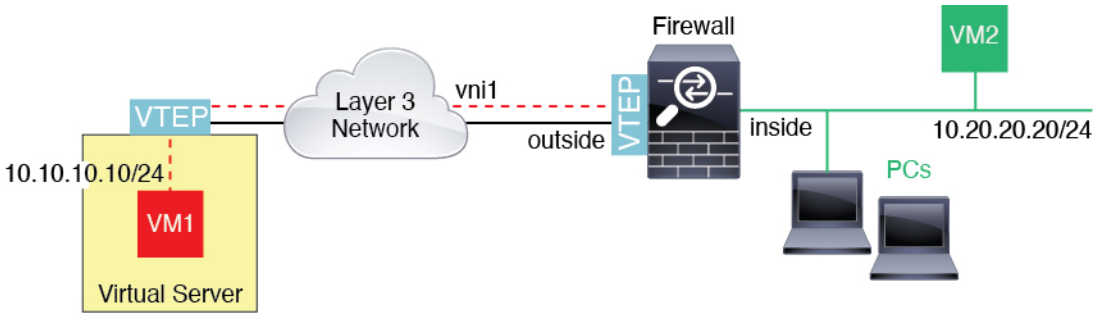
VXLAN Bridge

When you use a bridge group (transparent firewall mode, or optionally routed mode), the ASA 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.



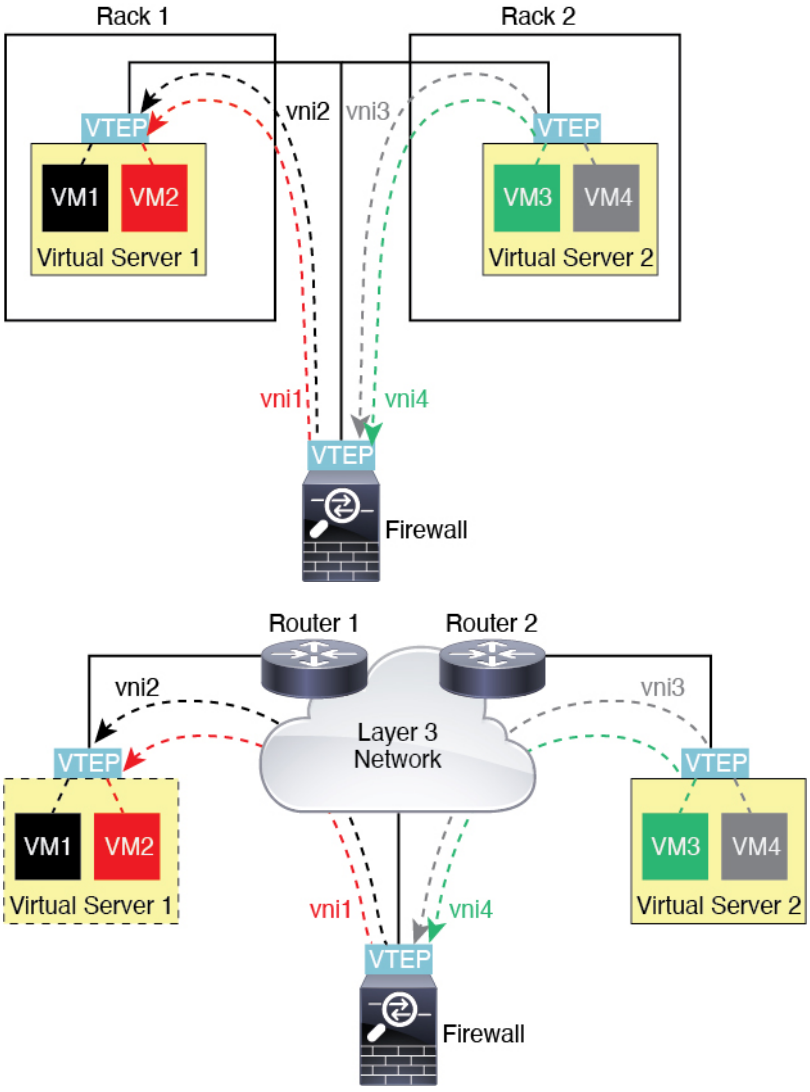
VXLAN Gateway (Routed Mode)

The ASA 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 ASA as its gateway while the ASA is not on the same rack, or even when the ASA 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 ASA MAC address, because the ASA 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 ASA.
3. When the ASA receives the packets, it decapsulates the packets to get the inner frames.
4. The ASA 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 ASA sends an encapsulated ARP broadcast on the multicast group IP on VNI 2.

Note The ASA must use dynamic VTEP peer discovery because it has multiple VTEP peers in this scenario.

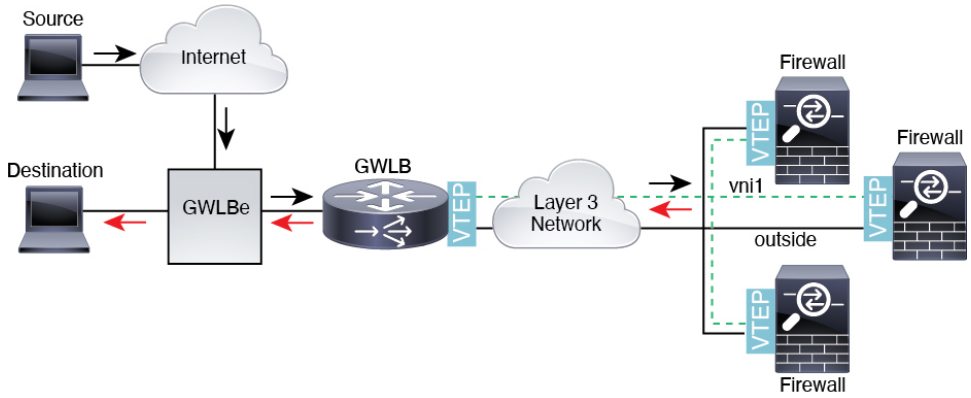
5. The ASA encapsulates the packets again with the VXLAN tag for VNI 2 and sends the packets to Virtual Server 1. Before encapsulation, the ASA changes the inner frame destination MAC address to be the MAC of VM1 (multicast-encapsulated ARP might be needed for the ASA to learn the VM1 MAC address).
6. When Virtual Server 1 receives the VXLAN packets, it decapsulates the packets and delivers the inner frames to VM1.

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 ASA 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 ASA 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 83: Geneve Single-Arm Proxy



AWS Gateway Load Balancer and Geneve Dual-Arm Proxy

Note This use case is the only currently supported use case for Geneve interfaces.

The ASA virtual supports the Gateway Load Balancer centralized control plane with a distributed data plane (Gateway Load Balancer endpoint) with single-arm and dual-arm mode. The following figure shows outbound traffic (traffic inspected by ASA virtual) directly forwarded to the destination (Internet) without the need for traffic hop to the GWLB and GWLB endpoint. The ASA virtual inspect the outbound and perform NAT of the traffic before either dropping it or sending it back to the internet via 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 84: Geneve Dual-Arm Proxy - Egress Traffic from single-VPC

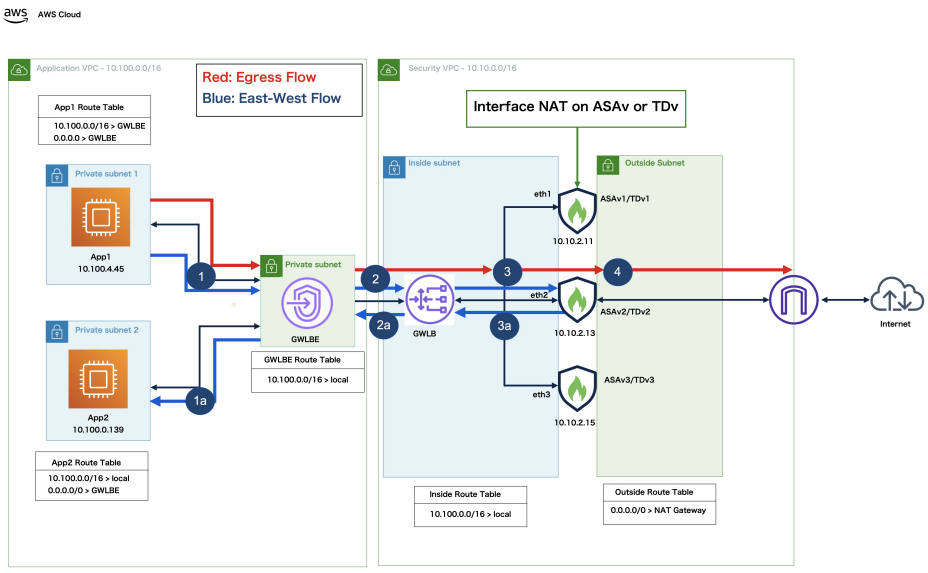
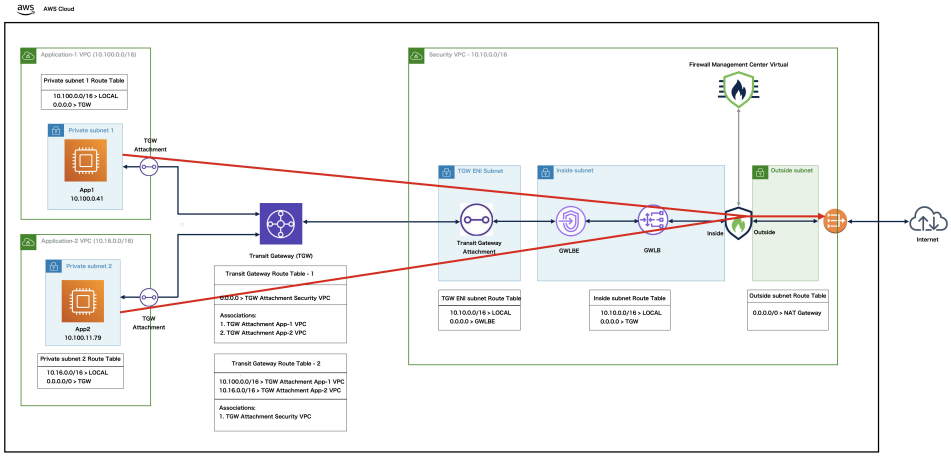


Figure 85: Geneve Dual-Arm Proxy - Egress Traffic from Multiple-VPC

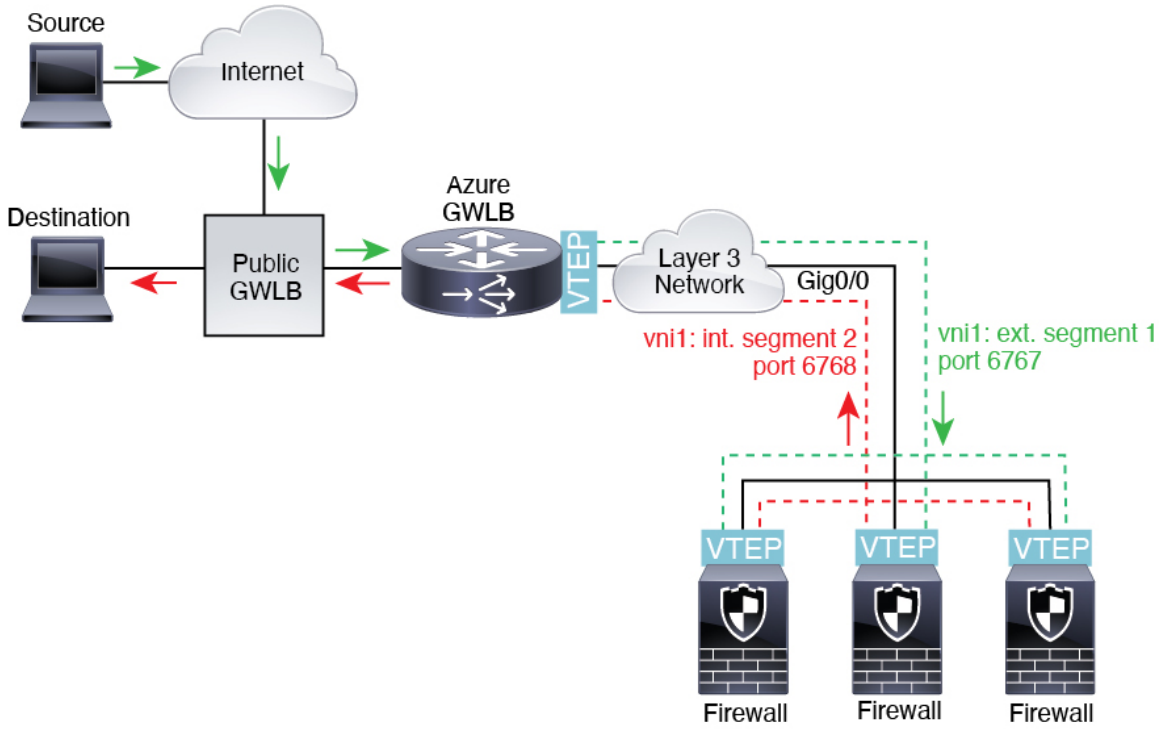


Azure Gateway Load Balancer and Paired Proxy

In an Azure service chain, ASA virtuals act as a transparent gateway that can intercept packets between the internet and the customer service. The ASA 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 ASA 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 86: Azure Gateway Load Balancer with Paired Proxy



Requirements and Prerequisites for VXLAN Interfaces

Model Requirements

- Firepower 1010 and Secure Firewall 1210/1220 switch ports and VLAN interfaces are not supported as VTEP interfaces.
- Geneve encapsulation is supported for the following models: ASAv30, ASAv50, ASAv100 on Amazon Web Services (AWS)
- VXLAN in paired proxy mode is supported for the following models:
  - ASA virtual in Azure

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 ASA virtual cluster control link VTEP source interface only supports IPv4.

  For Geneve, the VTEP source interfaces only supports IPv4.

Clustering and Multiple Context Mode

- ASA clustering does not support VXLAN in Individual Interface mode except for the cluster control link (ASA virtual only). Only Spanned EtherChannel mode supports VXLAN.

  An exception is made for the ASA virtualon 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.
- Geneve interfaces are only supported in single context mode. They are not supported with multiple context mode.

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 ASA automatically raises the MTU. 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. This MTU requires you to enable jumbo frame reservation on some models; see Enable Jumbo Frame Support (ASA Virtual, ISA 3000), on page 567.
- Geneve encapsulation—If the source interface MTU is less than 1806 bytes, then the ASA 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 you to enable jumbo frame reservation on some models; see Enable Jumbo Frame Support (ASA Virtual, ISA 3000), on page 567.

Default Settings for VXLAN Interfaces

VNI interfaces are enabled by default.

Configure VXLAN Interfaces

To configure VXLAN, perform the following steps.

Note You can configure either VXLAN or Geneve (ASA virtual only). For Geneve interfaces, see Configure Geneve Interfaces, on page 624.

Procedure

Step 1 Configure the VTEP Source Interface, on page 622.

Step 2 Configure the VNI Interface, on page 623

Step 3 (Azure GWLB) Allow Gateway Load Balancer Health Checks, on page 626.

Configure the VTEP Source Interface

You can configure one VTEP source interface per ASA or per security context. The VTEP is defined as a Network Virtualization Endpoint (NVE). An exception is made for clustering on the ASA 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.

Before you begin

For multiple context mode, complete the tasks in this section in the context execution space. In the Configuration > Device List pane, double-click the context name under the active device IP address.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces**, and edit the interface you want to use for the VTEP source interface.

Step 2 (Transparent Mode) Check the **VTEP Source Interface** check box.

This setting lets you configure an IP address for the interface. This command is optional for routed mode where this setting restricts traffic to VXLAN only on this interface.

Step 3 Configure the source interface name and IPv4 and/or IPv6 address, and click **OK**.

The ASA virtual cluster control link does not support IPv6.

Step 4 Choose **Configuration** > **Device Setup** > **Interface Settings** > **VXLAN**.

Step 5 (Optional) Enter a **VXLAN Destination Port** value if you want to change from the default 4789.

In multiple context mode, configure this setting in the System execution space.

Step 6 From the **Enable Network Virtualization Endpoint encapsulation using**drop-down menu, choose **VXLAN**.

Step 7 Choose the **VTEP Tunnel Interface** from the drop-down list.

Note
If the VTEP interface MTU is less than 1554 bytes for IPv4 or 1574 bytes for IPv6, then the ASA automatically raises the MTU to 1554 bytes or 1574 bytes.

Step 8 (Optional) Check the **Configure Packet Recipient** check box.

- (Multiple context mode; Optional for single mode) Enter the **Specify Peer VTEP IP Address** to manually specify the peer VTEP IP address

  If you specify the peer IP address, you cannot use multicast group discovery. Multicast is not supported in multiple context mode, so manual configuration is the only option. You can only specify one peer for the VTEP.

- (Single mode only) Enter the **Multicast traffic to default multicast address** to 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 9 Click **Apply**.

Configure the VNI Interface

Add a VNI interface, associate it with the VTEP source interface, and configure basic interface parameters.

For the ASA 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 **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces**, and click **Add** > **VNI Interface**.

Step 2 Enter the **VNI ID**, between 1 and 10000.

This ID is just an internal interface identifier.

Step 3 Enter the **VNI Segment ID**, between 1 and 16777215.

The segment ID is used for VXLAN tagging.

Step 4 (Transparent Mode) Choose the **Bridge Group** to which you want to assign this interface.

See Configure Bridge Group Interfaces, on page 641 to configure the BVI interface and associate regular interfaces to this bridge group.

Step 5 Enter the **Interface Name**.

The name is a text string up to 48 characters, and is not case-sensitive. You can change the name by reentering this command with a new value.

Step 6 Enter the **Security Level**, between 0 (lowest) and 100 (highest). See Security Levels, on page 633.

Step 7 (Single Mode) 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. Multicast is not supported in multiple context mode.

Step 8 Check the **Map to VTEP Tunnel Interface** check box.

This setting associates the VNI interface with the VTEP source interface.

Step 9 Check the **Enable Interface** check box. This setting is enabled by default.

Step 10 (Routed Mode) In the **IP Address** area, configure an IPv4 address. To configure IPv6, click the **IPv6** tab.

Step 11 Click **OK**, and then **Apply**.

Configure Geneve Interfaces

To configure Geneve interfaces for the ASA virtual, perform the following steps.

Note You can configure either VXLAN or Geneve. For VXLAN interfaces, see Configure VXLAN Interfaces, on page 622.

Procedure

Step 1 Configure the VTEP Source Interface for Geneve, on page 624.

Step 2 Configure the VNI Interface for Geneve, on page 625

Step 3 Allow Gateway Load Balancer Health Checks, on page 626.

Configure the VTEP Source Interface for Geneve

You can configure one VTEP source interface per ASA virtual. The VTEP is defined as a Network Virtualization Endpoint (NVE).

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces**, and edit the interface you want to use for the VTEP source interface.

Step 2 (Optional) Check the **VTEP Source Interface** check box.

This setting restricts traffic to VXLAN only on this interface.

Step 3 Configure the source interface name and IPv4 address, and click **OK**.

Step 4 Choose **Configuration** > **Device Setup** > **Interface Settings** > **VXLAN**.

Step 5 From the **Enable Network Virtualization Endpoint encapsulation using** drop-down menu, choose **Geneve**.

Step 6 Do not change the **Geneve Port**; AWS requires a port of 6081.

Step 7 Choose the **VTEP Tunnel Interface** from the drop-down list.

Note
If the VTEP interface MTU is less than 1806 bytes, then the ASA automatically raises the MTU to 1806 bytes.

Step 8 Click **Apply**.

Configure the VNI Interface for Geneve

Add a VNI interface, associate it with the VTEP source interface, and configure basic interface parameters.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces**, and click **Add** > **VNI Interface**.

Step 2 Enter the **VNI ID**, between 1 and 10000.

This ID is just an internal interface identifier.

Step 3 Enter the **Interface Name**.

The name is a text string up to 48 characters, and is not case-sensitive. You can change the name by reentering this command with a new value.

Step 4 Enter the **Security Level**, between 0 (lowest) and 100 (highest). See Security Levels, on page 633.

Step 5 Check the **Map to VTEP Tunnel Interface** check box.

This setting associates the VNI interface with the VTEP source interface.

Step 6 Check the **Enable Interface** check box. This setting is enabled by default.

Step 7 Check **Enable Single-Arm Proxy**.

Step 8 In the **IP Address** area, configure an IPv4 address. To configure IPv6, click the **IPv6** tab.

Step 9 Click **OK**.

Step 10 To allow traffic to enter and exit the same interface., check **Enable traffic between two or more hosts connected to the same interface**.

Step 11 Click **Apply**.

Allow Gateway Load Balancer Health Checks

The AWS or Azure Gateway Load Balancer requires appliances to answer a health check properly. The AWS Gateway Load Balancer will only send traffic to appliances that are considered healthy.

You must configure the ASA virtual to respond to an SSH, Telnet, HTTP, or HTTPS health check.

SSH Connection

For SSH, allow SSH from the Gateway Load Balancer. The Gateway Load Balancer will attempt to establish a connection to the ASA virtual, and the ASA virtual's prompt to log in is taken as proof of health.

Note An SSH login attempt will time out after 1 minute. You will need to configure a longer health check interval on the Gateway Load Balancer to accommodate this timeout.

Telnet Connection

For Telnet, allow Telnet from the Gateway Load Balancer. The Gateway Load Balancer will attempt to establish a connection to the ASA virtual, and the ASA virtual's prompt to log in is taken as proof of health.

Note You cannot Telnet to the lowest security level interface, so this method may not be practical.

HTTP(S) Cut-Through Proxy

You can configure the ASA to prompt the Gateway Load Balancer for an HTTP(S) login.

HTTP(S) Redirection Using Static Interface NAT with Port Translation

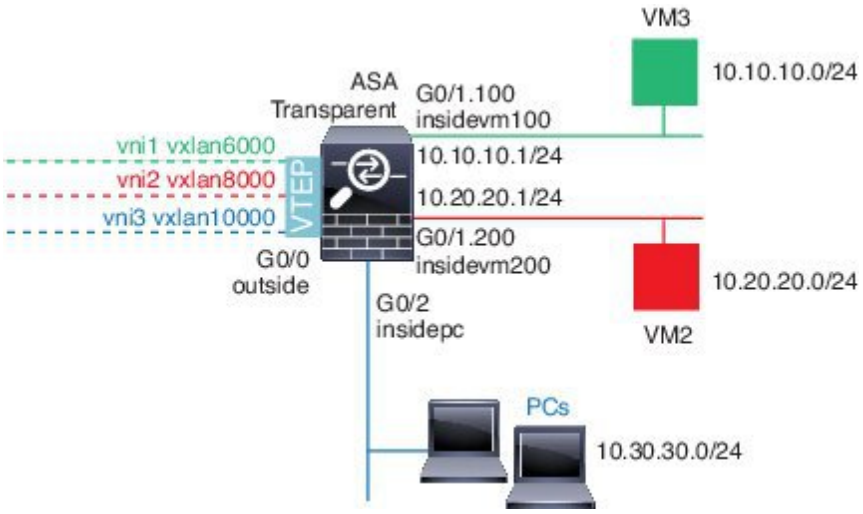
You can configure the ASA virtual to redirect health checks to a metadata HTTP(S) server. For HTTP(S) health checks, the HTTP(S) server must reply to the Gateway Load Balancer with a status code in the range 200 to 399. Because the ASA virtual has limits on the the number of simultaneous management connections, 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 Gateway Load Balancer with a destination of the ASA virtual outside interface so that it appears to be from the ASA virtual outside interface with a destination of the HTTP server. The ASA virtual then forwards the packet to the mapped destination address. The HTTP server responds to the ASA virtual outside interface, and then the ASA virtual forwards the response back to the Gateway Load Balancer. You need an access rule that allows traffic from the Gateway Load Balancer to the HTTP server.

Examples for VXLAN Interfaces

See the following configuration examples for VXLAN.

Transparent VXLAN Gateway Example



See the following description of this example:

- The outside interface on GigabitEthernet 0/0 is used as the VTEP source interface, and it is connected to the Layer 3 network.
- The insidevm100 VLAN subinterface on GigabitEthernet 0/1.100 is connected to the 10.10.10.0/24 network, on which VM3 resides. When VM3 communicates with VM1 (not shown; both have 10.10.10.0/24 IP addresses), the ASA uses VXLAN tag 6000.
- The insidevm200 VLAN subinterface on GigabitEthernet 0/1.200 is connected to the 10.20.20.0/24 network, on which VM2 resides. When VM2 communicates with VM4 (not shown; both have 10.20.20.0/24 IP addresses), the ASA uses VXLAN tag 8000.
- The insidepc interface on GigabitEthernet 0/2 is connected to the 10.30.30.0/24 network on which a few PCs reside. When those PCs communicate with VMs/PCs (not shown) behind a remote VTEP that belongs to same network (all have 10.30.30.0/24 IP addresses), the ASA uses VXLAN tag 10000.

ASA Configuration

```
firewall transparent
vxlan port 8427
!
interface gigabitethernet0/0
  nve-only
  nameif outside
  ip address 192.168.1.30 255.255.255.0
  no shutdown
!
nve 1
  encapsulation vxlan
  source-interface outside
!
interface vni1
  segment-id 6000
  nameif vxlan6000
  security-level 0
  bridge-group 1
```

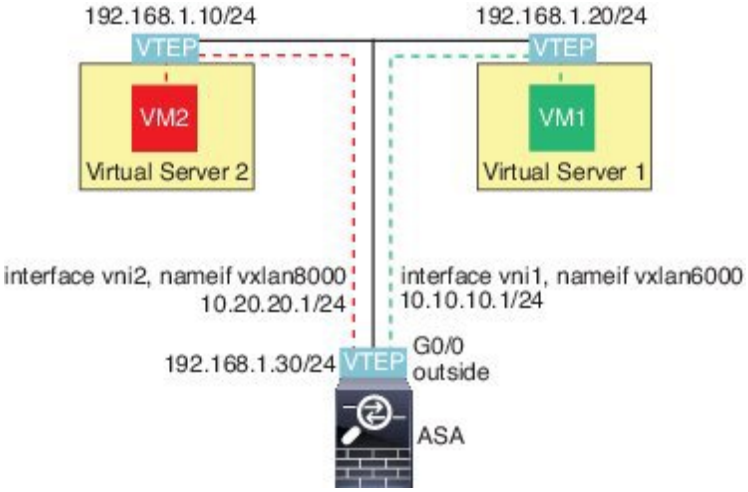
```
  vtep-nve 1
  mcast-group 235.0.0.100
!
interface vni2
  segment-id 8000
  nameif vxlan8000
  security-level 0
  bridge-group 2
  vtep-nve 1
  mcast-group 236.0.0.100
!
interface vni3
  segment-id 10000
  nameif vxlan10000
  security-level 0
  bridge-group 3
  vtep-nve 1
  mcast-group 236.0.0.100
!
interface gigabitethernet0/1.100
  nameif insidevm100
  security-level 100
  bridge-group 1
!
interface gigabitethernet0/1.200
  nameif insidevm200
  security-level 100
  bridge-group 2
!
interface gigabitethernet0/2
  nameif insidepc
  security-level 100
  bridge-group 3
!
interface bvi 1
  ip address 10.10.10.1 255.255.255.0
!
interface bvi 2
  ip address 10.20.20.1 255.255.255.0
!
interface bvi 3
  ip address 10.30.30.1 255.255.255.0
```

Notes

- For VNI interfaces vni1 and vni2, the inner VLAN tag is removed during encapsulation.
- VNI interfaces vni2 and vni3 share the same multicast IP address for encapsulated ARP over multicast. This sharing is allowed.
- The ASA bridges the VXLAN traffic to non-VXLAN-supported interfaces based on the above BVIs and bridge group configurations. For each of the stretched Layer 2 network segments (10.10.10.0/24, 10.20.20.0/24 and 10.30.30.0/24), the ASA serves as a bridge.
- It is allowed to have more than one VNI or more than one regular interface (VLAN or just physical interface) in a bridge group. The forwarding or association between VXLAN segment ID to the VLAN ID (or a physical interface) is decided by the destination MAC address and which interface connects to the destination.

- The VTEP source-interface is a Layer 3 interface in transparent firewall mode indicated by **nve-only** in the interface configuration. The VTEP source interface is not a BVI interface or a management interface, but it has an IP address and uses the routing table.

VXLAN Routing Example



See the following description of this example:

- VM1 (10.10.10.10) is hosted on Virtual Server 1, and VM2 (10.20.20.20) is hosted on Virtual Server 2.
- The default gateway for VM1 is the ASA, which is not in the same pod as Virtual Server 1, but VM1 is not aware of it. VM1 only knows that its default gateway IP address is 10.10.10.1. Similarly, VM2 only knows that its default gateway IP address is 10.20.20.1.
- The VTEP-supported hypervisors on Virtual Server 1 and 2 are able to communicate with the ASA over the same subnet or through a Layer 3 network (not shown; in which case, the ASA and uplinks of virtual servers have different network addresses).
- VM1's packet will be encapsulated by its hypervisor's VTEP and sent to its default gateway over VXLAN tunneling.
- When VM1 sends a packet to VM2, the packet will be sent through default gateway 10.10.10.1 from its perspective. Virtual Server1 knows 10.10.10.1 is not local, so the VTEP encapsulates the packet over VXLAN and sends it to ASA's VTEP.
- On the ASA, the packet is decapsulated. The VXLAN segment ID is learned during decapsulation. The ASA then re-injects the inner frame to the corresponding VNI interface (vni1) based on the VXLAN segment ID. The ASA then conducts a route lookup and sends the inner packet through another VNI interface, vni2. All egressing packets through vni2 are encapsulated with VXLAN segment 8000 and sent through the VTEP to outside.
- Eventually the encapsulated packet is received by the VTEP of Virtual Server 2, which decapsulates it and forwards it to VM2.

ASA Configuration

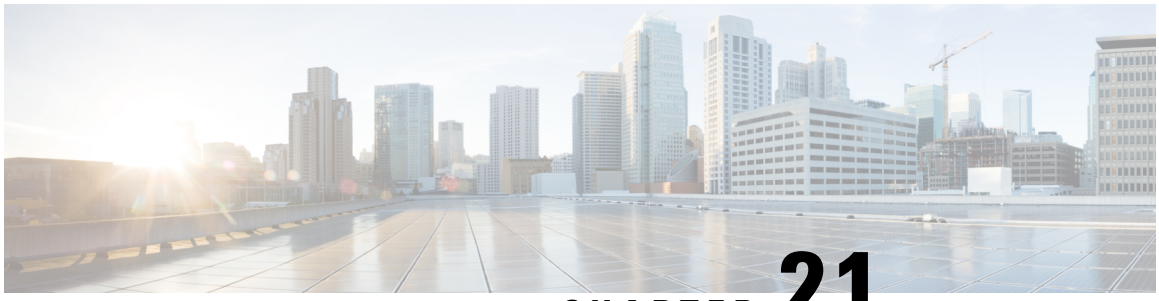
```
interface gigabitethernet0/0
  nameif outside
  ip address 192.168.1.30 255.255.255.0
  no shutdown
!
nve 1
  encapsulation vxlan
  source-interface outside
  default-mcast-group 235.0.0.100
!
interface vni1
  segment-id 6000
  nameif vxlan6000
  security-level 0
  vtep-nve 1
  ip address 10.20.20.1 255.255.255.0
!
interface vni2
  segment-id 8000
  nameif vxlan8000
  security-level 0
  vtep-nve 1
  ip address 10.10.10.1 255.255.255.0
!
```

History for VXLAN Interfaces

Table 31: History for VXLAN Interfaces

Feature Name	Releases	Feature Information
VXLAN VTEP IPv6 support	9.20(1)	You can now specify an IPv6 address for the VXLAN VTEP interface. IPv6 is not supported for the ASA virtual cluster control link or for Geneve encapsulation. New/Modified screens: • **Configuration** > **Device Setup** > **Interface Settings** > **VXLAN** • **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces** > **Add** > **VNI Interface**
Paired proxy VXLAN for the ASA virtual for the Azure Gateway Load Balancer	9.19(1)	You can configure a paired proxy mode VXLAN interface for the ASA virtual in Azure for use with the Azure Gateway Load Balancer (GWLB). The ASA virtual defines an external interface and an internal interface on a single NIC by utilizing VXLAN segments in a paired proxy. New/Modified commands: **external-port, external-segment-id, internal-port, internal-segment-id, proxy paired** No ASDM support.

Feature Name	Releases	Feature Information
Geneve support for the ASA virtual on AWS for the AWS Gateway Load Balancer	9.17(1)	Geneve encapsulation support was added for the ASAv30, ASAv50, and ASAv100 to support single-arm proxy for the AWS Gateway Load Balancer. New/Modified screens: **Configuration > Device Setup > Interface Settings > Interfaces > Add > VNI Interface** **Configuration > Device Setup > Interface Settings > VXLAN**
VXLAN support	9.4(1)	VXLAN support was added, including VXLAN tunnel endpoint (VTEP) support. You can define one VTEP source interface per ASA or security context. We introduced the following screens: **Configuration > Device Setup > Interface Settings > Interfaces > Add > VNI Interface** **Configuration > Device Setup > Interface Settings > VXLAN**



CHAPTER 21

Routed and Transparent Mode Interfaces

This chapter includes tasks to complete the interface configuration for all models in routed or transparent firewall mode.

Note For multiple context mode, complete the tasks in this section in the context execution space. In the Configuration > Device List pane, double-click the context name under the active device IP address.

- About Routed and Transparent Mode Interfaces, on page 633
- Guidelines and Limitations for Routed and Transparent Mode Interfaces, on page 635
- Configure Routed Mode Interfaces, on page 637
- Configure Bridge Group Interfaces, on page 641
- Configure IPv6 Addressing, on page 645
- Monitoring Routed and Transparent Mode Interfaces, on page 656
- Examples for Routed and Transparent Mode Interfaces, on page 658
- History for Routed and Transparent Mode Interfaces, on page 661

About Routed and Transparent Mode Interfaces

The ASA supports two types of interfaces: routed and bridged.

Each Layer 3 routed interface requires an IP address on a unique subnet.

Bridged interfaces belong to a bridge group, and all interfaces are on the same network. The bridge group is represented by a Bridge Virtual Interface (BVI) that has an IP address on the bridge network. Routed mode supports both routed and bridged interfaces, and you can route between routed interfaces and BVIs. Transparent firewall mode only supports bridge group and BVI interfaces.

Security Levels

Each interface must have a security level from 0 (lowest) to 100 (highest), including bridge group member interfaces. For example, you should assign your most secure network, such as the inside host network, to level 100. While the outside network connected to the Internet can be level 0. Other networks, such as DMZs can be in between. You can assign interfaces to the same security level.

Whether you assign a security level to a BVI depends on the firewall mode. In transparent mode, the BVI interface does not have a security level because it does not participate in routing between interfaces. In routed mode, BVI interfaces have a security level if you choose to route between the BVIs and other interfaces. For routed mode, the security level on a bridge group member interface only applies for communication within the bridge group. Similarly, the BVI security level only applies for inter-BVI/Layer 3 interface communication.

The level controls the following behavior:

- Network access—By default, there is an implicit permit from a higher security interface to a lower security interface (outbound). Hosts on the higher security interface can access any host on a lower security interface. You can limit access by applying an ACL to the interface.

  If you enable communication for same-security interfaces, there is an implicit permit for interfaces to access other interfaces on the same security level or lower.

- Inspection engines—Some application inspection engines are dependent on the security level. For same-security interfaces, inspection engines apply to traffic in either direction.
  - NetBIOS inspection engine—Applied only for outbound connections.
  - SQL*Net inspection engine—If a control connection for the SQL*Net (formerly OraServ) port exists between a pair of hosts, then only an inbound data connection is permitted through the ASA.

Dual IP Stack (IPv4 and IPv6)

The ASA 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 ASAs 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 in Spanned clustering mode, excluding the management interface and the Cluster Control Link.

You cannot use a 31-bit subnet mask in Individual clustering mode on any interface.

31-Bit Subnet and Failover

For failover, when you use a 31-bit subnet for the ASA 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 ASA 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 ASA, 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

Context Mode

- In multiple context mode, you can only configure context interfaces that you already assigned to the context in the system configuration according to Configure Multiple Contexts, on page 260.
- PPPoE is not supported in multiple context mode.
- For multiple context mode in transparent mode, each context must use different interfaces; you cannot share an interface across contexts.
- For multiple context mode in transparent mode, each context typically uses a different subnet. You can use overlapping subnets, but your network topology requires router and NAT configuration to make it possible from a routing standpoint.
- DHCPv6 and prefix delegation options are not supported with multiple context mode.
- In routed firewall mode, bridge group interfaces are not supported in multiple context mode.

Failover, Clustering

- Do not configure failover links with the procedures in this chapter. See the Failover chapter for more information.
- For cluster interfaces, see the clustering chapter for requirements.
- When you use Failover, you must set the IP address and standby address for data interfaces manually; DHCP and PPPoE are not supported.

IPv6

- IPv6 is supported on all interfaces.

- You can only configure IPv6 addresses manually in transparent mode.
- The ASA does not support IPv6 anycast addresses.
- DHCPv6 and prefix delegation options are not supported with multiple context mode, transparent mode, clustering, or Failover.

Model Guidelines

- For the ASAv50, bridge groups are not supported in either transparent or routed mode.

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 ASA 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 ASA. 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 the ASAv50 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.
- 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 ASA 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.
- In transparent mode, PPPoE is not supported for the Management interface.
- In routed mode, to route between bridge groups and other routed interfaces, you must name the BVI.
- In routed mode, ASA-defined EtherChannel and VNI 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 ASA when using bridge group members. If there are two neighbors on either side of the ASA running BFD, then the ASA will drop BFD echo packets because they have the same source and destination IP address and appear to be part of a LAND attack.

Default Security Level

The default security level is 0. If you name an interface "inside," and you do not set the security level explicitly, then the ASA sets the security level to 100.



Note If you change the security level of an interface, and you do not want to wait for existing connections to time out before the new security information is used, you can clear the connections using the **clear conn** command.

Additional Guidelines and Requirements

- The ASA supports only one 802.1Q header in a packet and does not support multiple headers (known as Q-in-Q support).
- 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

To configure routed mode interfaces, perform the following steps.

Configure General Routed Mode Interface Parameters

This procedure describes how to set the name, security level, IPv4 address, and other options.

Before you begin

In multiple context mode, complete this procedure in the context execution space. To change from the system to a context configuration, in the Configuration > Device List pane, double-click the context name under the active device IP address.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces**.

Step 2 Choose the interface row, and click **Edit**.

The **Edit Interface** dialog box appears with the **General** tab selected.

Note
For the Firepower 1010, you cannot configure switch ports as routed mode interfaces.

Step 3 In the **Interface Name** field, enter a name up to 48 characters in length.

Step 4 In the **Security level** field, enter a level between 0 (lowest) and 100 (highest).

Note
For loopback interfaces, you do not set the security level because the interface is only supported for to/from the device traffic.

Step 5 (Optional) To set this interface as a management-only interface, check the **Dedicate this interface to management-only** check box.

Through traffic is not accepted on a management-only interface.

Note
The Channel Group field is read-only and indicates if the interface is part of an EtherChannel.

Note
For loopback interfaces, you do not set the management mode because the interface is only supported for to/from the device traffic.

Step 6 If the interface is not already enabled, check the **Enable Interface** check box.

Step 7 To set the IP address, use one of the following options.

Note
For failover and clustering, and for loopback interfaces, you must set the IP address manually; DHCP and PPPoE are not supported.

- To set the IP address manually, click the **Use Static IP** radio button and enter the IP address and mask.

  For failover, set the standby IP addresses on the **Configuration > Device Management > High Availability > Failover > Interfaces** tab. 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.

  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. You cannot set the standby IP address in this case.
- To obtain an IP address from a DHCP server, click the **Obtain Address via DHCP** radio button.

  a. To force a MAC address to be stored inside a DHCP request packet for option 61, click the **Use MAC Address** radio button.

  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.

  b. To use a generated string for option 61, click **Use "Cisco-<MAC>-<interface_name>-<host>"**.

  c. (Optional) To obtain the default route from the DHCP server, check **Obtain Default Route Using DHCP**.

  d. (Optional) To assign an administrative distance to the learned route, enter a value between 1 and 255 in the **DHCP Learned Route Metric** field. If this field is left blank, the administrative distance for the learned routes is 1.

  e. (Optional) To enable tracking for DHCP-learned routes, check **Enable Tracking for DHCP Learned Routes**. Set the following values:

  Track ID—A unique identifier for the route tracking process. Valid values are from 1 to 500.

Track IP Address—Enter the IP address of the target being tracked. Typically, this would be the IP address of the next hop gateway for the route, but it could be any network object available off of that interface.

Note
Route tracking is only available in single, routed mode.

SLA ID—A unique identifier for the SLA monitoring process. Valid values are from 1 to 2147483647.

Monitor Options—Click this button to open the **Route Monitoring Options** dialog box. In the **Route Monitoring Options** dialog box you can configure the parameters of the tracked object monitoring process.

f. (Optional) To set the broadcast flag to 1 in the DHCP packet header when the DHCP client sends a discover requesting an IP address, check **Enable DHCP Broadcast flag for DHCP request and discover messages**.

The DHCP server listens to this broadcast flag and broadcasts the reply packet if the flag is set to 1.

g. (Optional) To renew the lease, click **Renew DHCP Lease**.

- (Single mode only) To obtain an IP address using PPPoE, check **Use PPPoE**.

a. In the **Group Name** field, specify a group name.

b. In the **PPPoE Username** field, specify the username provided by your ISP.

c. In the **PPPoE Password** field, specify the password provided by your ISP.

d. In the **Confirm Password** field, retype the password.

e. For PPP authentication, click either the **PAP**, **CHAP**, or **MSCHAP** radio button.

PAP passes 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.

f. (Optional) To store the username and password in flash memory, check the **Store Username and Password in Local Flash** check box.

The ASA stores the username and password in a special location of NVRAM. If an Auto Update Server sends a **clear configure** command to the ASA, and the connection is then interrupted, the ASA can read the username and password from NVRAM and re-authenticate to the Access Concentrator.

g. (Optional) To display the **PPPoE IP Address and Route Settings** dialog box where you can choose addressing and tracking options, click **IP Address and Route Settings**.

Step 8 (Optional) In the **Description** field, enter a description for this interface.

The description can be up to 240 characters on a single line, without carriage returns. In the case of a failover or state link, the description is fixed as "LAN Failover Interface," "STATE Failover Interface," or "LAN/STATE Failover Interface," for example. You cannot edit this description. The fixed description overwrites any description you enter here if you make this interface a failover or state link.

Step 9 Click **OK**.

Related Topics

Configure IPv6 Addressing, on page 645
Enable the Physical Interface and Configure Ethernet Parameters, on page 565
Configure PPPoE, on page 640

Configure 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.

Procedure

Step 1 Choose **Configuration** > **Interfaces** > **Add/Edit Interface** > **General**, and then click **PPPoE IP Address and Route Settings**.

Step 2 In the **IP Address** area, choose one of the following:

- **Obtain IP Address using PPP**—Dynamically configure the IP address.
- **Specify an IP Address**—Manually configure the IP address.

Step 3 In the **Route Settings Area**, configure the following:

- **Obtain default route using PPPoE**—Set the default routes when the PPPoE client has not yet established a connection. When using this option, you cannot have a statically defined route in the configuration.
- **PPPoE learned route metric**—Assign an administrative distance to the learned route. Valid values are from 1 to 255. If this field is left blank, the administrative distance for the learned routes is 1.
- **Enable tracking**—Enable route tracking for PPPoE-learned routes. Route tracking is only available in single, routed mode.
- **Primary Track**—Configure the primary PPPoE route tracking.
- **Track ID**—A unique identifier for the route tracking process. Valid values are from 1 to 500.
- **Track IP Address**—Enter the IP address of the target being tracked. Typically, this would be the IP address of the next hop gateway for the route, but it could be any network object available off of that interface.
- **SLA ID**—A unique identifier for the SLA monitoring process. Valid values are from 1 to 2147483647.
- **Monitor Options**—Click this button to open the **Route Monitoring Options** dialog box. In the **Route Monitoring Options** dialog box you can configure the parameters of the tracked object monitoring process.
- **Secondary Track**—Configure the secondary PPPoE route tracking.
- **Secondary Track ID**—A unique identifier for the route tracking process. Valid values are from 1 to 500.

Step 4 Click **OK**.

Configure Bridge Group Interfaces

A bridge group is a group of interfaces that the ASA 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 217.

To configure bridge groups and associated interfaces, perform these steps.

Configure the Bridge Virtual Interface (BVI)

Each bridge group requires a BVI for which you configure an IP address. The ASA 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.

Some models include a bridge group and BVI in the default configuration. You can create additional bridge groups and BVIs and reassign member interfaces between the groups.

Note For a separate management interface in transparent mode (for supported models), a non-configurable bridge group (ID 301) is automatically added to your configuration. This bridge group is not included in the bridge group limit.

Procedure

Step 1 Choose **Configuration** > **Interfaces**, and then choose **Add** > **Bridge Group Interface**.

Step 2 In the **Bridge Group ID** field, enter the bridge group ID between 1 and 250.

You will later assign physical interfaces to this bridge group number.

Step 3 (Routed Mode) In the **Interface 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.

Step 4 (Routed Mode) In the **Security level** field, enter a level between 0 (lowest) and 100 (highest).

Step 5 (Transparent Mode) Set the IP address.

a) In the **IP Address** field, enter the IPv4 address.

b) In the **Subnet Mask** field, enter the subnet mask or choose one from the menu.

Do not assign a host address (/32 or 255.255.255.255) to the transparent firewall. 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 ASA 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 ASA drops the ARP request from the downstream router to the upstream router.

Step 6 (Routed Mode) To set the IP address, use one of the following options.

For failover and clustering, you must set the IP address manually; DHCP is not supported.

- To set the IP address manually, click the **Use Static IP** radio button and enter the IP address and mask.
- To obtain an IP address from a DHCP server, click the **Obtain Address via DHCP** radio button.
  - **a.** To force a MAC address to be stored inside a DHCP request packet for option 61, click the **Use MAC Address** radio button.

    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.
  - **b.** To use a generated string for option 61, click **Use "Cisco-<MAC>-<interface_name>-<host>"**.
  - **c.** (Optional) To obtain the default route from the DHCP server, check **Obtain Default Route Using DHCP**.
  - **d.** (Optional) To set the broadcast flag to 1 in the DHCP packet header when the DHCP client sends a discover requesting an IP address, check **Enable DHCP Broadcast flag for DHCP request and discover messages**.

    The DHCP server listens to this broadcast flag and broadcasts the reply packet if the flag is set to 1.
  - **e.** (Optional) To renew the lease, click **Renew DHCP Lease**.

Step 7 (Optional) In the **Description** field, enter a description for this bridge group.

Step 8 Click **OK**.

A Bridge Virtual Interface (BVI) is added to the interface table, along with the physical and subinterfaces.

Configure General Bridge Group Member Interface Parameters

This procedure describes how to set the name, security level, and bridge group for each bridge group member interface.

Before you begin

- The same bridge group can include different types of interfaces: physical interfaces, VLAN subinterfaces, VNI interfaces, and EtherChannels. The Management interface is not supported. In routed mode, EtherChannels and VNIs are not supported.
- In multiple context mode, complete this procedure in the context execution space. To change from the system to a context configuration, in the Configuration > Device List pane, double-click the context name under the active device IP address.
- For transparent mode, do not use this procedure for Management interfaces; see Configure a Management Interface for Transparent Mode, on page 643 to configure the Management interface.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces**.

BVIs appear in the table alongside physical interfaces, subinterfaces, and EtherChannel port-channel interfaces. In multiple context mode, only interfaces that were assigned to the context in the System execution space appear in the table.

Step 2 Choose the row for a non-BVI interface, and click **Edit**.

The **Edit Interface** dialog box appears with the **General** tab selected.

Note
For the Firepower 1010, you cannot configure switch ports as bridge group members.

You cannot mix logical VLAN interfaces and physical router interfaces in the same bridge group.

Note
In routed mode, the **port-channel** and **vni** interfaces are not supported as bridge group members.

Step 3 In the **Bridge Group** drop-down menu, choose the bridge group to which you want to assign this interface.

Step 4 In the **Interface Name** field, enter a name up to 48 characters in length.

Step 5 In the **Security level** field, enter a level between 0 (lowest) and 100 (highest).

Step 6 If the interface is not already enabled, check the **Enable Interface** check box.

Note
The **Channel Group** field is read-only and indicates if the interface is part of an EtherChannel.

Step 7 (Optional) If you install a module, and you want to demonstrate the module functionality on a non-production ASA, check the **Forward traffic to the ASA module for inspection and reporting check box**. See the module chapter or quick start guide for more information.

Step 8 (Optional) In the **Description** field, enter a description for this interface.

The description can be up to 240 characters on a single line, without carriage returns. In the case of a failover or state link, the description is fixed as "LAN Failover Interface," "STATE Failover Interface," or "LAN/STATE Failover Interface," for example. You cannot edit this description. The fixed description overwrites any description you enter here if you make this interface a failover or state link.

Step 9 Click **OK**.

Related Topics

Configure the Manual MAC Address, MTU, and TCP MSS, on page 670

Configure a Management Interface for Transparent Mode

In transparent firewall mode, all interfaces must belong to a bridge group. The only exception is the Management interface (either the physical interface, a subinterface (if supported for your model), or an EtherChannel interface comprised of Management interfaces (if you have multiple Management interfaces)) which you can configure as a separate management interface; for the Firepower 4100/9300 chassis, the management interface ID depends on the mgmt-type interface that you assigned to the ASA logical device. You cannot use any other

interface types as management interfaces. You can configure one management interface in single mode or per context. For more information see Management Interface for Transparent Mode, on page 563.

The Management interface is for to- and from-the-box traffic only and cannot pass through traffic.

Before you begin

- Do not assign this interface to a bridge group; a non-configurable bridge group (ID 301) is automatically added to your configuration. This bridge group is not included in the bridge group limit.
- For the Firepower 4100/9300 chassis, the management interface ID depends on the mgmt-type interface that you assigned to the ASA logical device.
- In multiple context mode, you cannot share any interfaces, including the Management interface, across contexts. You must connect to a data interface.
- In multiple context mode, complete this procedure in the context execution space. To change from the system to a context configuration, in the Configuration > Device List pane, double-click the context name under the active device IP address.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces**.

Step 2 Choose the row for a Management interface, subinterface, or EtherChannel port-channel interface comprised of Management interfaces, and click **Edit**.

The **Edit Interface** dialog box appears with the **General** tab selected.

For the Firepower 4100/9300 chassis, the management interface ID depends on the mgmt-type interface (individual or EtherChannel) that you assigned to the ASA logical device.

Step 3 In the **Bridge Group** drop-down menu, leave the default **--None--**. You cannot assign a management interface to a bridge group.

Step 4 In the **Interface Name** field, enter a name up to 48 characters in length.

Step 5 In the **Security level** field, enter a level between 0 (lowest) and 100 (highest).

Note
The **Dedicate this interface to management only** check box is enabled by default and is non-configurable.

Step 6 If the interface is not already enabled, check the **Enable Interface** check box.

Step 7 To set the IP address, use one of the following options.

Note
For use with failover, you must set the IP address and standby address manually; DHCP is not supported. Set the standby IP addresses on the **Configuration > Device Management > High Availability > Failover > Interfaces** tab.

- To set the IP address manually, click the **Use Static IP** radio button and enter the IP address and mask.
- To obtain an IP address from a DHCP server, click the **Obtain Address via DHCP** radio button.
  - To force a MAC address to be stored inside a DHCP request packet for option 61, click the **Use MAC Address** radio button.

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.

- To use a generated string for option 61, click **Use "Cisco-<MAC>-<interface_name>-<host>"**.
- (Optional) To obtain the default route from the DHCP server, check **Obtain Default Route Using DHCP**.
- (Optional) To set the broadcast flag to 1 in the DHCP packet header when the DHCP client sends a discover requesting an IP address, check **Enable DHCP Broadcast flag for DHCP request and discover messages**.

  The DHCP server listens to this broadcast flag and broadcasts the reply packet if the flag is set to 1.
- (Optional) To renew the lease, click **Renew DHCP Lease**.

Step 8 (Optional) In the **Description** field, enter a description for this interface.

The description can be up to 240 characters on a single line, without carriage returns.

Step 9 Click **OK**.

Configure IPv6 Addressing

This section describes how to configure IPv6 addressing.

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 ASA 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 ASA 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 ASA 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 ASA can then subnet and assign to its inside interfaces.

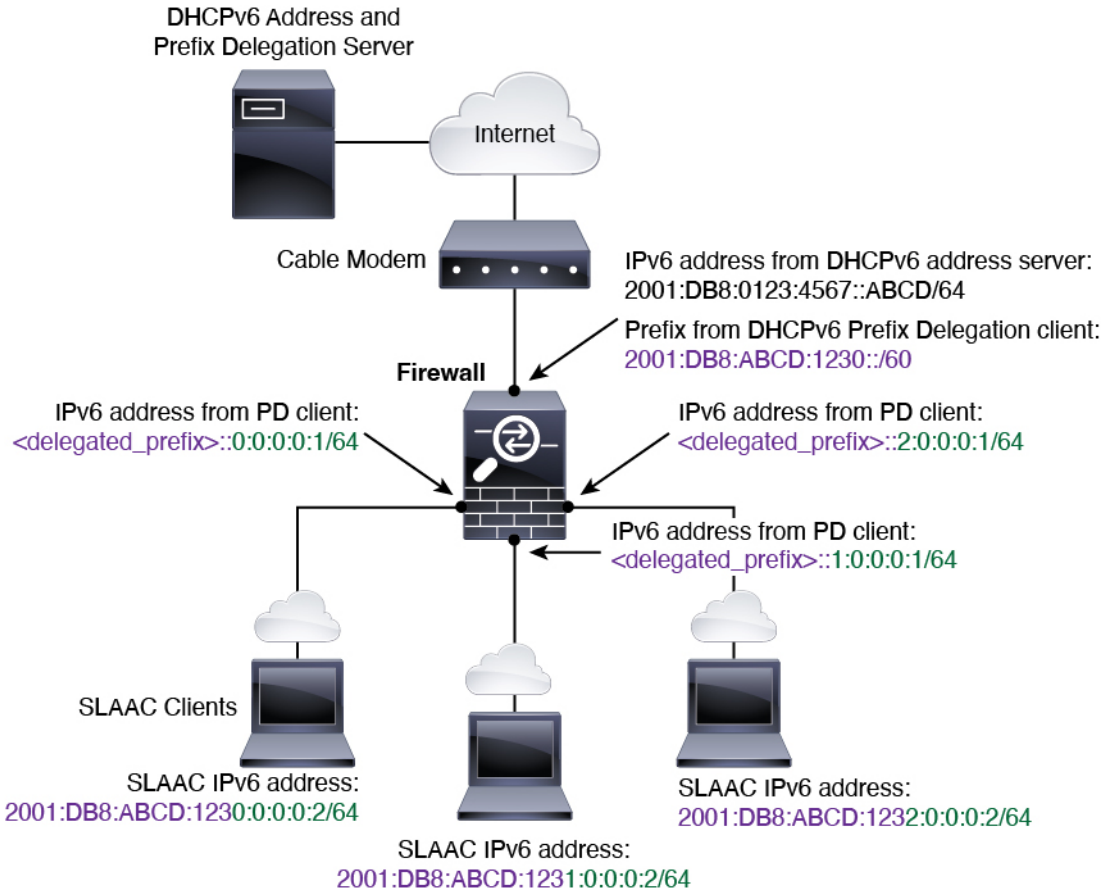
About IPv6 Prefix Delegation

The ASA 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 ASA 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 ASA interfaces do not in turn act as Prefix Delegation servers; the ASA can only provide global IP addresses to SLAAC clients. For example, if a router is connected to the ASA, 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 ASA includes a light DHCPv6 server so the ASA can provide information such as the DNS server and domain name to SLAAC clients when they send Information Request (IR) packets to the ASA. The ASA 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 ASA received using Prefix Delegation.

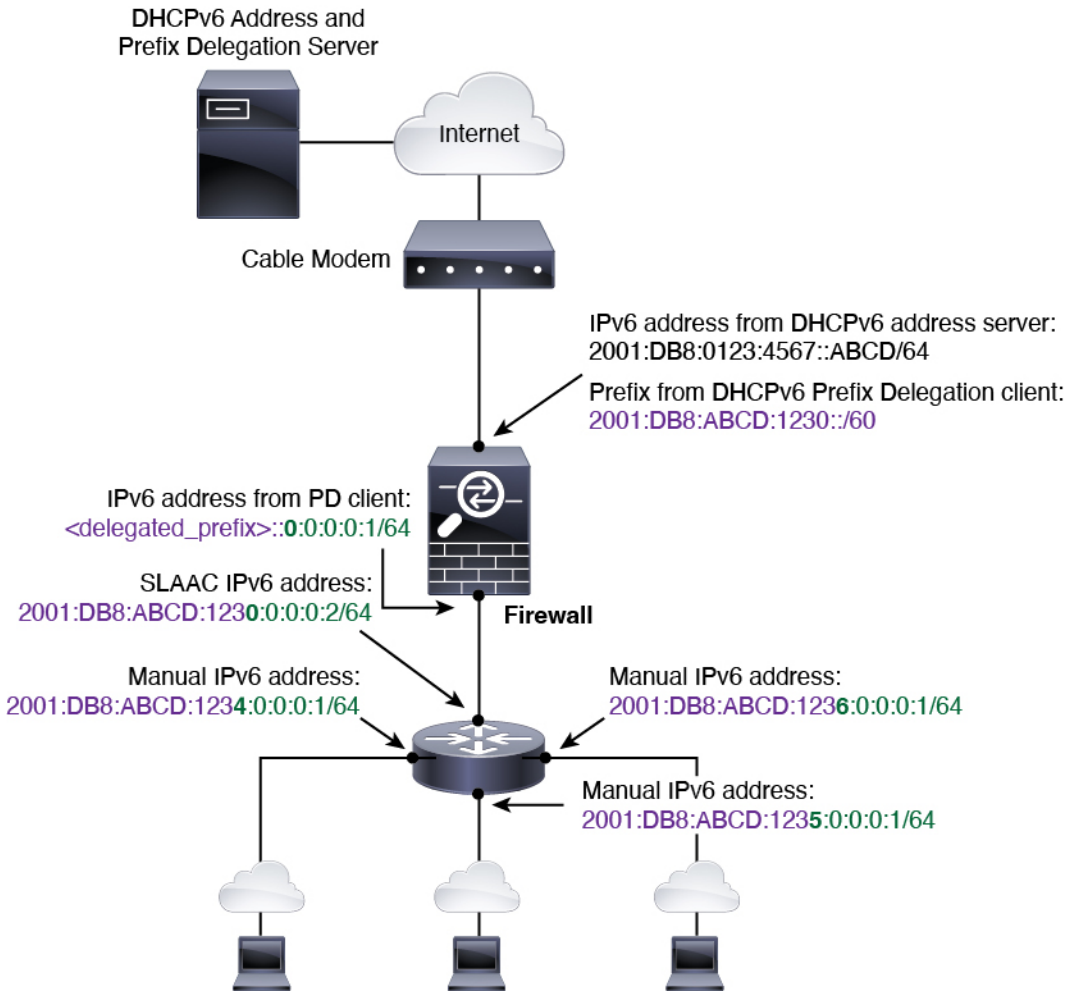
IPv6 Prefix Delegation /64 Subnet Example

The following example shows the ASA 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 ASA 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 ASA 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 ASA 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 ASA, 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 ASA 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 ASA 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 ASA interfaces use addresses derived from the delegated prefix.

Before you begin

- This feature is only supported in routed firewall mode.
- This feature is not supported in multiple context mode.
- This feature is not supported in clustering.
- You cannot configure this feature on a management-only interface.
- When you use Prefix Delegation, you must set the ASA 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 ASA RA interval to be 150 seconds. To set the preferred

lifetime, use the **show ipv6 general-prefix** command. To set the ASA RA interval, see Configure IPv6 Neighbor Discovery, on page 653; the default is 200 seconds.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces**.

Step 2 Choose an interface, and click **Edit**.

The **Edit Interface** dialog box appears with the **General** tab selected.

Step 3 Click the **IPv6** tab.

Step 4 In the **Interface IPv6 DHCP** area, click the **Client Prefix Delegation Name** radio button, and enter the prefix name.

Step 5 (Optional) In the **Prefix Hint** field, provide one or more hints about the delegated prefix you want to receive.

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 (2001:DB8:ABCD:1230::/60). 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**.

You return to the **Configuration > Device Setup > Interface Settings > Interfaces** pane.

Step 7 Click **Apply**.

Step 8 See Configure a Global IPv6 Address, on page 649 to assign a subnet of the prefix as the global IP address for an ASA interface.

Step 9 (Optional) See Configure the DHCPv6 Stateless Server, on page 720 to provide domain-name and server parameters to SLAAC clients.

Step 10 (Optional) See Configure IPv6 Network Settings, on page 842 to advertise the prefix(es) with BGP.

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.

DHCPv6 and prefix delegation options are not supported with multiple context mode.



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, 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 ASA. See Configure the Manual MAC Address, MTU, and TCP MSS, on page 670.

Before you begin

- In multiple context mode, complete this procedure in the context execution space. To change from the system to a context configuration, in the Configuration > Device List pane, double-click the context name under the active device IP address.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces**.

Step 2 Choose an interface, and click **Edit**.

The **Edit Interface** dialog box appears with the **General** tab selected.

In transparent mode or for a bridge group in routed mode, select a BVI. For transparent mode, you can also select a management-only interface.

Step 3 Click the **IPv6** tab.

Step 4 Check the **Enable IPv6** check box.

Step 5 (Optional) 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.

Step 6 (Routed interface) 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 not supported.

- Stateless autoconfiguration—In the **Interface IPv6 Addresses** area, check the **Enable address autoconfiguration** check box.

  Enabling stateless autconfiguration 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.

  Note
  Although RFC 4862 specifies that hosts configured for stateless autoconfiguration do not send Router Advertisement messages, the ASA does send Router Advertisement messages in this case. Check the **Suppress RA** check box to suppress messages.

  If you want to install a default route, choose **DHCP** or **Ignore** from the drop-down menu. **DHCP** specifies the ASA only uses a default route from Router Advertisements that come from a trusted source (in other words, from the same server that provided the IPv6 address). **Ignore** specifies that Router Advertisements can be sourced from another network, which can be a riskier method.

- Manual configuration—To manually configure a global IPv6 address:

  a. In the **Interface IPv6 Addresses** area, click **Add**.

  The **Add IPv6 Address for Interface** dialog box appears.

  b. In the **Address/Prefix Length** field, the value you enter depends on the method you want to use:

  - Full global address—If you want to manually enter the entire address, enter the full address plus the prefix length.

- Modified EUI 64 format—Enter the IPv6 prefix and length, and then check the **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).
- Delegated Prefix—To derive the IPv6 prefix from the delegated prefix, enter the IPv6 address and length. Then enter the prefix name that you configured for the DHCPv6 Prefix Delegation client (See Enable the IPv6 Prefix Delegation Client, on page 648) in the **Prefix Name** field, and click **Add**.

  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.

c. Click **OK**.

- Obtain an address using DHCPv6:

a. In the **Interface IPv6 DHCP** area, check the **Enable DHCP** check box.

b. (Optional) Check the **Enable Default** check box to obtain a default route from Router Advertisements.

Step 7 (BVI interface) Manually assign a global address to the BVI. For a management interface in Transparent mode, use this method as well.

a) In the **Interface IPv6 Addresses** area, click **Add**.

The **Add IPv6 Address for Interface** dialog box appears.

b) In the **Address/Prefix Length** field, enter the full global IPv6 address along with the IPv6 prefix length.

c) Click **OK**.

Step 8 Click **OK**.

You return to the **Configuration > Device Setup > Interface Settings > Interfaces** pane.

(Optional) Configure the Link-Local Addresses Automatically

If you do not want to configure a global address, and only need to configure a link-local address, you have the option of generating the link-local addresses based on the interface MAC addresses (Modified EUI-64 format. Because MAC addresses use 48 bits, additional bits must be inserted to fill the 64 bits required for the interface ID.)

To automatically configure the link-local addresses for an interface, perform the following steps.

Before you begin

Supported in routed mode only.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces**.

Step 2 Select an interface, and click **Edit**.

For bridge groups in routed mode, choose the BVI.

The **Edit Interface** dialog box appears with the **General** tab selected.

Step 3 Click the **IPv6** tab.

Step 4 In the **IPv6 configuration** area, check the **Enable IPv6** check box.

This option enables IPv6 and automatically generates the link-local address using the Modified EUI-64 interface ID based on the interface MAC address.

For bridge groups in routed mode, enabling IPv6 for the BVI generates link-local addresses for all member interfaces.

Step 5 Click **OK**.

(Optional) Configure the Link-Local Addresses Manually

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.

To assign a link-local address to an interface, perform the following steps.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces**.

Step 2 Select an interface, and click **Edit**.

For bridge groups, choose a bridge group member interface.

The **Edit Interface** dialog box appears with the **General** tab selected.

Step 3 Click the **IPv6** tab.

Step 4 (Optional) 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.

Step 5 To 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. See IPv6 Addresses, on page 1215 for more information about IPv6 addressing.

Step 6 Click **OK**.

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.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces**.

Step 2 Choose the IPv6 interface on which to configure IPv6 neighbor settings, and click **Edit**.

Step 3 Click the **IPv6** tab.

Step 4 Enter the number of allowed **DAD Attempts**.

Values range from 0 to 600. A 0 value disables DAD processing on the specified interface. The default is 1 message.

DAD ensures the uniqueness of new unicast IPv6 addresses before they are assigned, and ensures that duplicate IPv6 addresses are detected in the network on a link basis. The ASA uses neighbor solicitation messages to perform DAD.

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 5 Enter the **NS Interval** in milliseconds to set the interval between IPv6 neighbor solicitation retransmissions.

Valid values for the value argument range from 1000 to 3600000 milliseconds.

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 6 Enter the **Reachable Time** in seconds to set how long a remote IPv6 node is reachable.

Set the reachable time between 0 to 3600000 milliseconds. When you set the time to 0, then 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 7 Enter the **RA Lifetime** in seconds to set the length of time that nodes on the local link consider the ASA as the default router on the link.

Values range from 0 to 9000 seconds. Entering 0 indicates that the ASA should not be considered a default router on the selected interface.

Step 8 Check the **Suppress RA** check box to suppress router advertisements.

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 ASA to supply the IPv6 prefix (for example, the outside interface).

Enabling this option causes the ASA to appear as a regular IPv6 neighbor on the link and not as an IPv6 router.

Step 9 Enter the **RA Interval** to set the interval between IPv6 router advertisement transmissions.

Valid values range from 3 to 1800 seconds. The default is 200 seconds.

To add a router advertisement transmission interval value in milliseconds instead, check the **RA Interval in Milliseconds** check box, and enter a value from 500 to 1800000.

To prevent synchronization with other IPv6 nodes, the ASA randomly adjusts the value that you set (jitter).

Step 10 Check the **Hosts should use DHCP for address config** check box to inform IPv6 autoconfiguration clients that they should use DHCPv6 to obtain addresses, in addition to the derived stateless autoconfiguration address.

This option sets the Managed Address Config flag in the IPv6 router advertisement packet.

Step 11 Check the **Hosts should use DHCP for non-address config** check box to inform IPv6 autoconfiguration clients that they should use DHCPv6 to obtain additional information from DHCPv6, such as the DNS server address.

This option sets the Other Address Config flag in the IPv6 router advertisement packet.

Step 12 Configure which IPv6 prefixes are included in IPv6 router advertisements.

a) In the **Interface IPv6 Prefixes** area, click **Add**.
b) Enter the **Address/Prefix Length** or check the **Default** check box to use the default prefix.
c) Check the **No Auto-Configuration** check box to force hosts to configure the IPv6 address manually. Hosts on the local link with the specified prefix cannot use IPv6 autoconfiguration.
d) Check the **No Advertisements** check box to disable prefix advertisement. For the **Default** prefix, this setting only applies to on-link prefixes. Off-link prefixes will still be advertised unless you specify **No Advertisements** for a specific off-link prefix.
e) Check the **Off Link** check box to configure the specified prefix as off-link. The prefix will be advertised with the L-bit clear. The prefix will not be inserted into the routing table as a Connected prefix.
f) In the **Prefix Lifetime** area, specify a **Lifetime Duration** or **Lifetime Expiration Date**.

After the preferred lifetime expires, the address goes into a deprecated state; while an address is in a deprecated state, its use is discouraged, but not strictly forbidden. After the valid lifetime expires, the address becomes invalid and cannot be used. The valid lifetime must be greater than or equal to the preferred lifetime.

- **Lifetime Duration**—Values range from 0 to 4294967295. The default valid lifetime is 2592000 (30 days). The default preferred lifetime is 604800 (7 days). The maximum value represents infinity.
- **Lifetime Expiration Date**—Choose a valid and preferred month and day from the drop-down lists, and then enter a time in hh:mm format.

g) Click **OK** to save your settings.

Step 13 Click **OK**.

Step 14 Configure a static IPv6 neighbor.

The following guidelines and limitations apply for configuring a static IPv6 neighbor:

- This feature is similar to adding a static ARP entry. If an entry for the specified IPv6 address already exists in the neighbor discovery cache—learned through the IPv6 neighbor discovery process—the entry is automatically converted to a static entry. These entries are stored in the configuration when the copy command is used to store the configuration.
- Static entries in the IPv6 neighbor discovery cache are not modified by the neighbor discovery process.
- The ICMP syslogs generated are caused by a regular refresh of IPv6 neighbor entries. The ASA default timer for IPv6 neighbor entry is 30 seconds, so the ASA would generate ICMPv6 neighbor discovery and response packets about every 30 seconds. If the ASA has both failover LAN and state interfaces configured with IPv6 addresses, then every 30 seconds, ICMPv6 neighbor discovery and response packets will be generated by both ASAs for both configured and link-local IPv6 addresses. In addition, each packet will generate several syslogs (ICMP connection and local-host creation or teardown), so it may appear that constant ICMP syslogs are being generated. The refresh time for IPV6 neighbor entry is configurable on the regular data interface, but not configurable on the failover interface. However, the CPU impact for this ICMP neighbor discovery traffic is minimal.

See also View and Clear Dynamically Discovered Neighbors, on page 656.

a) Choose **Configuration** > **Device Management** > **Advanced** > **IPv6 Neighbor Discovery Cache**.
b) Click **Add**.

The **Add IPv6 Static Neighbor** dialog box appears.

c) From the **Interface Name** drop-down list, choose an interface on which to add the neighbor.
d) In the **IP Address** field, enter the IPv6 address that corresponds to the local data-link address, or click the ellipsis (...) to browse for an address.
e) In the **MAC address** field, enter the local data-line (hardware) MAC address.
f) Click **OK**.

Step 15 Click **Apply** to save the running configuration.

View and Clear Dynamically Discovered Neighbors

When a host or node communicates with a neighbor, the neighbor is added to the neighbor discovery cache. The neighbor is removed from the cache when there is no longer any communication with that neighbor.

To view dynamically discovered neighbors and clear these neighbors from the IPv6 neighbor discovery cache, perform the following steps:

Procedure

Step 1 Choose **Monitoring** > **Interfaces** > **IPv6 Neighbor Discovery Cache**.

You can view all static and dynamically discovered neighbors from the IPv6 Neighbor Discovery Cache pane.

Step 2 To clear all dynamically discovered neighbors from the cache, click **Clear Dynamic Neighbor Entries**.

The dynamically discovered neighbor is removed from the cache.

Note
This procedure clears only dynamically discovered neighbors from the cache; it does not clear static neighbors.

Monitoring Routed and Transparent Mode Interfaces

You can monitor interface statistics, status, PPPoE, and more.



Note For Firepower and Secure Firewall models, some statistics are not shown using the ASA commands. You must view more detailed interface statistics using FXOS commands.

- /eth-uplink/fabric# **show interface**
- /eth-uplink/fabric# **show port-channel**
- /eth-uplink/fabric/interface# **show stats**

See the FXOS troubleshooting guide for more information.

Interface Statistics and Information

- **Monitoring > Interfaces > Interface Graphs**

  Lets you view interface statistics in graph or table form. If an interface is shared among contexts, the ASA shows only statistics for the current context. The number of statistics shown for a subinterface is a subset of the number of statistics shown for a physical interface.

- **Monitoring > Interfaces > Interface Graphs > Graph/Table**

Shows a graph for the selected statistics. The Graph window can show up to four graphs and tables at a time. By default, the graph or table displays the real-time statistics. If you enable History Metrics, you can view statistics for past time periods.

DHCP Information

- **Monitoring** > **Interfaces** > **DHCP** > **DHCP Client Lease Information**.

  This screen displays configured DHCP client IP addresses.

- **Monitoring** > **Interfaces** > **DHCP** > **IPV6 DHCP Client PD Statistics**

  This screen shows DHCPv6 Prefix Delegation client statistics and shows the output of the number of messages sent and received.

- **Monitoring** > **Interfaces** > **DHCP** > **IPV6 DHCP Client Statistics**

  This screen shows DHCPv6 client statistics and shows the output of the number of messages sent and received.

- **Monitoring** > **Interfaces** > **DHCP** > **IPV6 DHCP Interface Statistics**

  This screen displays DHCPv6 information for all interfaces. If the interface is configured for DHCPv6 stateless server configuration (see Configure the DHCPv6 Stateless Server, on page 720), this screen lists the DHCPv6 pool that is being used by the server. If the interface has DHCPv6 address client or Prefix Delegation client configuration, this screen shows the state of each client and the values received from the server. This screen also shows message statistics for the DHCP server or client.

- **Monitoring** > **Interfaces** > **DHCP** > **IPV6 DHCP HA Statistics**

  This screen shows the transaction statistics between failover units, including how many times the DUID information was synced between the units.

Static Route Tracking

- **Monitoring > Interfaces > interface connection > Track Status**

  Displays information about the tracked object.

- **Monitoring > Interfaces > interface connection > Monitoring Statistics**

  Displays statistics for the SLA monitoring process.

PPPoE

- **Monitoring > Interfaces > PPPoE Client > PPPoE Client Lease Information**

  Displays information about current PPPoE connections.

Dynamic ACLs

Monitoring > Interfaces > Dynamic ACLs

Shows a table of the Dynamic ACLs, which are functionally identical to the user-configured ACLs except that they are created, activated and deleted automatically by the ASA. These ACLs do not show up in the configuration and are only visible in this table. They are identified by the "(dynamic)" keyword in the ACL header.

Examples for Routed and Transparent Mode Interfaces

Transparent Mode Example with 2 Bridge Groups

The following example for transparent mode includes two bridge groups of three interfaces each, plus a management-only interface:

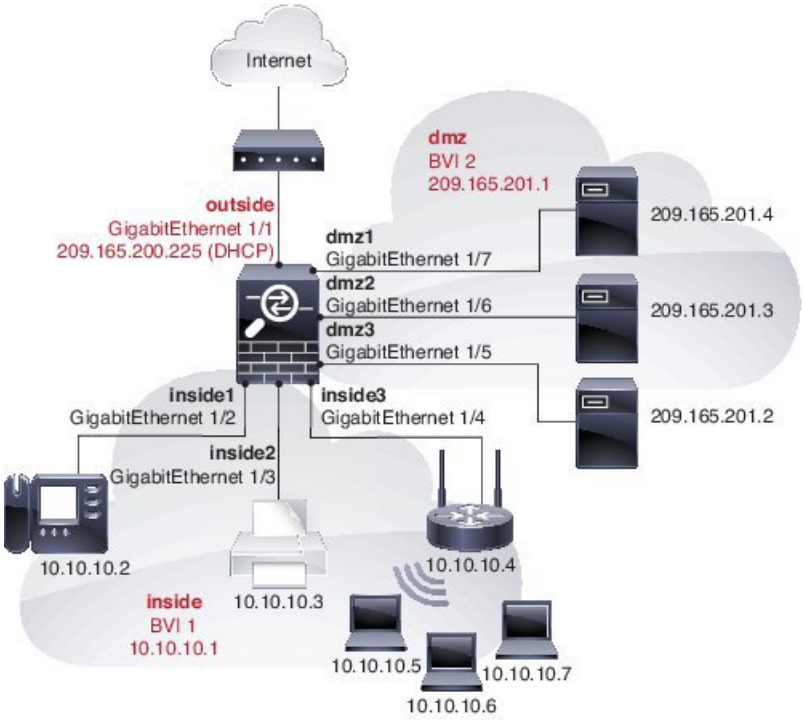
```
interface gigabitethernet 0/0
  nameif inside1
  security-level 100
  bridge-group 1
  no shutdown
interface gigabitethernet 0/1
  nameif outside1
  security-level 0
  bridge-group 1
  no shutdown
interface gigabitethernet 0/2
  nameif dmz1
  security-level 50
  bridge-group 1
  no shutdown
interface bvi 1
  ip address 10.1.3.1 255.255.255.0 standby 10.1.3.2

interface gigabitethernet 1/0
  nameif inside2
  security-level 100
  bridge-group 2
  no shutdown
interface gigabitethernet 1/1
  nameif outside2
  security-level 0
  bridge-group 2
  no shutdown
interface gigabitethernet 1/2
  nameif dmz2
  security-level 50
  bridge-group 2
  no shutdown
interface bvi 2
  ip address 10.3.5.8 255.255.255.0 standby 10.3.5.9

interface management 0/0
  nameif mgmt
  security-level 100
  ip address 10.2.1.1 255.255.255.0 standby 10.2.1.2
  no shutdown
```

Switched LAN Segment Example with 2 Bridge Groups

The following example configures 2 bridge groups with 3 interfaces each and one regular routed interface for outside. Bridge group 1 is inside and bridge group 2 is dmz with public web servers. The bridge group member interfaces can communicate freely within the bridge group because each member is at the same security level, and we enabled same security communication. Although the inside member security level is 100 and the dmz member security level is also 100, these security levels do not apply to inter-BVI communications; only the BVI security levels affect inter-BVI traffic. The security levels of the BVIs and outside (100, 50, and 0) implicitly permit traffic from inside to dmz and inside to outside; and from dmz to outside. An access rule is applied to outside to allow traffic to the servers on dmz.




```
interface gigabitethernet 1/1
  nameif outside
  security-level 0
  ip address dhcp setroute
  no shutdown
!
interface gigabitethernet 1/2
  nameif inside1
  security-level 100
  bridge-group 1
  no shutdown
interface gigabitethernet 1/3
  nameif inside2
  security-level 100
  bridge-group 1
  no shutdown
interface gigabitethernet 1/4
  nameif inside3
  security-level 100
  bridge-group 1
  no shutdown
```

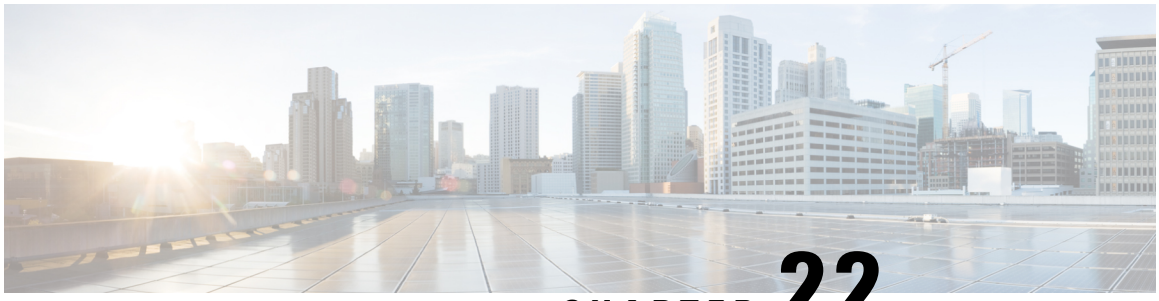
```
!
interface bvi 1
  nameif inside
  security-level 100
  ip address 10.10.10.1 255.255.255.0
!
interface gigabitethernet 1/5
  nameif dmz1
  security-level 100
  bridge-group 2
  no shutdown
interface gigabitethernet 1/6
  nameif dmz2
  security-level 100
  bridge-group 2
  no shutdown
interface gigabitethernet 1/7
  nameif dmz3
  security-level 100
  bridge-group 2
  no shutdown
!
interface bvi 2
  nameif dmz
  security-level 50
  ip address 209.165.201.1 255.255.255.224
!
same-security-traffic permit inter-interface
!
# Assigns IP addresses to inside hosts
dhcpd address 10.10.10.2-10.10.10.200 inside
dhcpd enable inside
!
# Applies interface PAT for inside traffic going outside
nat (inside1,outside) source dynamic any interface
nat (inside2,outside) source dynamic any interface
nat (inside3,outside) source dynamic any interface
!
# Allows outside traffic to each server for specific applications
object network server1
  host 209.165.201.2
object network server2
  host 209.165.201.3
object network server3
  host 209.165.201.4
!
# Defines mail services allowed on server3
object-group service MAIL
  service-object tcp destination eq pop3
  service-object tcp destination eq imap4
  service-object tcp destination eq smtp
!
# Allows access from outside to servers on the DMZ
access-list SERVERS extended permit tcp any object server1 eq www
access-list SERVERS extended permit tcp any object server2 eq ftp
access-list SERVERS extended permit tcp any object server3 object-group MAIL
access-group SERVERS in interface outside
```

History for Routed and Transparent Mode Interfaces

Feature Name	Platform Releases	Feature Information
IPv6 Neighbor Discovery	7.0(1)	We introduced this feature. We introduced the following screens: Monitoring > Interfaces > IPv6 Neighbor Discovery Cache. Configuration > Device Management > Advanced > IPv6 Neighbor Discovery Cache. Configuration > Device Setup > Interface Settings > Interfaces > IPv6.
IPv6 support for transparent mode	8.2(1)	IPv6 support was introduced for transparent firewall mode.
Bridge groups for transparent mode	8.4(1)	If you do not want the overhead of security contexts, or want to maximize your use of security contexts, you can group interfaces together in a bridge group, and then configure multiple bridge groups, one for each network. Bridge group traffic is isolated from other bridge groups. You can configure up to eight bridge groups of four interfaces each in single mode or per context. We modified or introduced the following screens: **Configuration > Device Setup > Interface Settings > Interfaces** **Configuration > Device Setup > Interface Settings > Interfaces > Add/Edit Bridge Group Interface** **Configuration > Device Setup > Interface Settings > Interfaces > Add/Edit Interface**
Address Config Flags for IPv6 DHCP Relay	9.0(1)	We modified the following screen: Configuration > Device Setup > Interfaces > IPv6.
Transparent mode bridge group maximum increased to 250	9.3(1)	The bridge group maximum was increased from 8 to 250 bridge groups. You can configure up to 250 bridge groups in single mode or per context in multiple mode, with 4 interfaces maximum per bridge group. We modified the following screens: **Configuration > Device Setup > Interface Settings > Interfaces** **Configuration > Device Setup > Interface Settings > Interfaces > Add/Edit Bridge Group Interface** **Configuration > Device Setup > Interface Settings > Interfaces > Add/Edit Interface**
Transparent mode maximum interfaces per bridge group increased to 64	9.6(2)	The maximum interfaces per bridge group was increased from 4 to 64. We did not modify any screens.

Feature Name	Platform Releases	Feature Information
IPv6 DHCP	9.6(2)	The ASA now supports the following features for IPv6 addressing: • DHCPv6 Address client—The ASA obtains an IPv6 global address and optional default route from the DHCPv6 server. • DHCPv6 Prefix Delegation client—The ASA obtains delegated prefix(es) from a DHCPv6 server. The ASA can then use these prefixes to configure other ASA interface addresess 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 ASA provides other information such as the domain name to SLAAC clients when they send Information Request (IR) packets to the ASA. The ASA only accepts IR packets, and does not assign addresses to the clients. We added or modified the following screens: **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces** > **Add Interface** > **IPv6** **Configuration** > **Device Management** > **DHCP** > **DHCP Pool** **Configuration** > **Device Setup** > **Routing** > **BGP** > **IPv6 Family** > **Networks** **Monitoring** > **interfaces** > **DHCP**

Feature Name	Platform Releases	Feature Information
Integrated Routing and Bridging	9.7(1)	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 ASA bridges instead of routes. The ASA is not a true bridge in that the ASA 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 ASA 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: multiple context mode, ASA clustering. The following features are also not supported on BVIs: dynamic routing and multicast routing. We modified the following screens: **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces** **Configuration** > **Device Setup** > **Routing** > **Static Routes** **Configuration** > **Device Management** > **DHCP** > **DHCP Server** **Configuration** > **Firewall** > **Access Rules** **Configuration** > **Firewall** > **EtherType Rules**
31-bit Subnet Mask	9.7(1)	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 ASAs 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. We modified the following screens: **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces** > **Add Interface** > **General**



CHAPTER 22

Advanced Interface Configuration

This chapter describes how to configure MAC addresses for interfaces, how to set the maximum transmission unit (MTU), and set the TCP maximum segment size (TCP MSS), and how to allow same security level communication. Setting the correct MTU and maximum TCP segment size is essential for the best network performance.

- About Advanced Interface Configuration, on page 665
- Automatically Assign MAC Addresses, on page 669
- Configure the Manual MAC Address, MTU, and TCP MSS, on page 670
- Allow Same Security Level Communication, on page 672
- Monitoring the ARP and MAC Address Table, on page 672
- History for Advanced Interface Configuration, on page 672

About Advanced Interface Configuration

This section describes advanced interface settings.

About MAC Addresses

You can manually assign MAC addresses to override the default. For multiple context mode, you can automatically generate unique MAC addresses (for all interfaces assigned to a context) and single context mode (for subinterfaces)..

Note You might want to assign unique MAC addresses to subinterfaces defined on the ASA, 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 ASA device.

Default MAC Addresses

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 Manual MAC Address, MTU, and TCP MSS, on page 670.

  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 Manual MAC Address, MTU, and TCP MSS, on page 670.

- 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—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 ASA.

Automatic MAC Addresses

In multiple context mode, auto-generation assigns unique MAC addresses to all interfaces assigned to a context.

If you manually assign a MAC address and also enable auto-generation, then the manually assigned MAC address is used. If you later remove the manual MAC address, the auto-generated address is used, if enabled.

In the rare circumstance that the generated MAC address conflicts with another private MAC address in your network, you can manually set the MAC address for the interface.

Because auto-generated addresses (when using a prefix) start with A2, you cannot start manual MAC addresses with A2 if you also want to use auto-generation.

The ASA generates the MAC address using the following format:

A2*xx.yyzz.zzzz*

Where *xx.yy* is a user-defined prefix or an autogenerated prefix based on the last two bytes of the interface MAC address, and *zz.zzzz* is an internal counter generated by the ASA. For the standby MAC address, the address is identical except that the internal counter is increased by 1.

For an example of how the prefix is used, if you set a prefix of 77, then the ASA converts 77 into the hexadecimal value 004D (*yyxx*). When used in the MAC address, the prefix is reversed (*xxyy*) to match the ASA native form:

A2**4D.00***zz.zzzz*

For a prefix of 1009 (03F1), the MAC address is:

A2**F1.03**zz.zzzz

Note The MAC address format without a prefix is a legacy version. See the **mac-address auto** command in the command reference for more information about the legacy format.

About the MTU

The MTU specifies the maximum frame *payload* size that the ASA 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 VXLAN or Geneve, the entire Ethernet datagram is being encapsulated, so the new IP packet is larger and requires a larger MTU: you should set the ASA VTEP source interface MTU to be the network MTU + 54 bytes (for VXLAN) or + 306 bytes (Geneve).

Path MTU Discovery

The ASA 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 ASA is 1500 bytes. This value does not include the 18-22 bytes for the Ethernet header, VLAN tagging, or other overhead.

When you enable VXLAN on the VTEP source interface, if the MTU is less than 1554 bytes, then the ASA automatically raises the MTU to 1554 bytes. In this case, the entire Ethernet datagram is being encapsulated, so the new packet is larger and requires a larger MTU. In general, you should set the ASA source interface MTU to be the network MTU + 54 bytes.

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 668.

For UDP or ICMP, the application should take the MTU into account to avoid fragmentation.

Note The ASA 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 ASA 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 ASA for through traffic; by default, the maximum TCP MSS is set to 1380 bytes. This setting is useful when the ASA 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 ASA.

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 ASA, then the ASA overwrites the TCP MSS in the request packet with the ASA maximum. If the host or server does not request a TCP MSS, then the ASA 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 ASA maximum TCP MSS is 1380 (the default), then the ASA changes the MSS value in the TCP request packet to 1380. The server then sends packets with 1380-byte payloads. The ASA 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 ASA 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 ASA 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 ASA 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 ASA acts as an IPv4 IPsec VPN endpoint and has an MTU of 1500. When the ASA 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 ASA as an IPsec VPN endpoint, then you should change the TCP 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.

Inter-Interface Communication

Allowing interfaces on the same security level to communicate with each other provides the following benefits:

- You can configure more than 101 communicating interfaces.

  If you use different levels for each interface and do not assign any interfaces to the same security level, you can configure only one interface per level (0 to 100).

- You want traffic to flow freely between all same security interfaces without ACLs.

If you enable same security interface communication, you can still configure interfaces at different security levels as usual.

Intra-Interface Communication (Routed Firewall Mode)

Intra-interface communication might be useful for VPN traffic that enters an interface, but is then routed out the same interface. The VPN traffic might be unencrypted in this case, or it might be reencrypted for another VPN connection. For example, if you have a hub and spoke VPN network, where the ASA is the hub, and remote VPN networks are spokes, for one spoke to communicate with another spoke, traffic must go into the ASA and then out again to the other spoke.



Note All traffic allowed by this feature is still subject to firewall rules. Be careful not to create an asymmetric routing situation that can cause return traffic not to traverse the ASA.

Automatically Assign MAC Addresses

This section describes how to configure auto-generation of MAC addresses. For multiple context mode, this feature assigns unique MAC addresses to all interface types that are assigned to a context. For single mode, this feature assigns unique MAC addresses to VLAN subinterfaces.

Before you begin

- When you configure a name for the interface, the new MAC address is generated immediately. If you enable this feature after you configure interfaces, then MAC addresses are generated for all interfaces immediately after you enable it. If you disable this feature, the MAC address for each interface reverts to the default MAC address. For example, subinterfaces of GigabitEthernet 0/1 revert to using the MAC address of GigabitEthernet 0/1.
- In the rare circumstance that the generated MAC address conflicts with another private MAC address in your network, you can manually set the MAC address for the interface.

- For multiple context mode, complete this procedure in the system execution space. If you are not already in the System configuration mode, in the **Configuration** > **Device List** pane, double-click **System** under the active device IP address.

Procedure

Step 1 For multiple context mode: Complete the following steps in the System.

a) Choose **Configuration** > **Context Management** > **Security Contexts**.
b) Check **Mac-Address auto**.
c) (Optional) Check the **Prefix** check box, and in the field, enter a decimal value between 0 and 65535.

This prefix is converted to a four-digit hexadecimal number, and used as part of the MAC address. If you do not enter a prefix, then the ASA autogenerates the prefix based on the last two bytes of the interface MAC address.

Step 2 For single context mode: Complete the following steps.

a) Choose **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces**.
b) At the bottom of the page, check the **Enable auto-generation of MAC addresses for subinterfaces** check box.
c) (Optional) In the **Prefix** field, enter a decimal value between 0 and 65535.

This prefix is converted to a four-digit hexadecimal number, and used as part of the MAC address. If you do not enter a prefix, then the ASA autogenerates the prefix based on the last two bytes of the interface MAC address.

Step 3 Click **Apply**.

Configure the Manual MAC Address, MTU, and TCP MSS

Before you begin

- In multiple context mode, complete this procedure in the context execution space. To change from the system to a context configuration, in the Configuration > Device List pane, double-click the context name under the active device IP address.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces**.

Step 2 Choose the interface row, and click **Edit**.

The **Edit Interface** dialog box appears with the **General** tab selected.

Step 3 Click the **Advanced** tab.

Step 4 To set the MTU or to enable jumbo frame support (supported models only), enter the value in the **MTU** field. The minimum and maximum depend on your platform.

The default is 1500 bytes.

Note
When you set the MTU for a port-channel interface, the ASA applies the setting to all member interfaces.

- For models that support jumbo frames in single mode—If you enter a value for any interface that is greater than 1500, then you enable jumbo frame support automatically for all interfaces. If you set the MTU for all interfaces back to a value under 1500, then jumbo frame support is disabled.
- For models that support jumbo frames in multiple mode—If you enter a value for any interface that is greater than 1500, then be sure to enable jumbo frame support in the system configuration, if required for your model. See Enable Jumbo Frame Support (ASA Virtual, ISA 3000), on page 567.

Note
For some models, enabling or disabling jumbo frame support requires you to reload the ASA.

Step 5 To manually assign a MAC address to this interface, enter a MAC address in the **Active Mac Address** field 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 first two bytes of a manual MAC address cannot be A2 if you also want to use auto-generated MAC addresses.

Step 6 If you use failover, enter the standby MAC address in the **Standby Mac Address** field. 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 7 To set the TCP MSS, choose **Configuration** > **Firewall** > **Advanced** > **TCP Options**. Set the following options:

- Send reset reply for denied outside TCP packets—Enables ASA to send reset replies for all outbound TCP sessions that attempt to transit the ASA and are denied by the ASA based on access lists or AAA settings.
- Force Maximum Segment Size for TCP—Sets the maximum TCP segment size in bytes, between 48 and any maximum number. The default value is 1380 bytes. You can disable this feature by setting the bytes to 0.
- Force Minimum Segment Size for TCP—Overrides the maximum segment size to be no less than the number of bytes you set, between 48 and any maximum number. This feature is disabled by default (set to 0).
- TCP Maximum unprocessed segments—Check this check box and specify the maximum number of unprocessed TCP segments. The default value is 6. The range is from 6 to 24.

Step 8 For **Secure Group Tagging** settings, see the firewall configuration guide.

Step 9 (Secure Firewall 3100) Click **Auto-negotiate** to negotiate the link status and flow control for 1 Gigabit and higher interfaces.

Step 10 For **ASA Cluster** settings, see (Recommended; Required in Multiple Context Mode) Configure Interfaces on the Control Node, on page 364.

Allow Same Security Level Communication

By default, interfaces on the same security level cannot communicate with each other, and packets cannot enter and exit the same interface. This section describes how to enable inter-interface communication when interfaces are on the same security level, and how to enable intra-interface communication.

Procedure

Step 1 To enable interfaces on the same security level to communicate with each other, from the **Configuration** > **Interfaces** pane, check **Enable traffic between two or more interfaces which are configured with same security level**.

Step 2 To enable communication between hosts connected to the same interface, check **Enable traffic between two or more hosts connected to the same interface**.

Monitoring the ARP and MAC Address Table

- **Monitoring > Interfaces > ARP Table**

  Displays the ARP table, including static and dynamic entries. The ARP table includes entries that map a MAC address to an IP address for a given interface.

- **Monitoring > Interfaces > MAC Address Table**

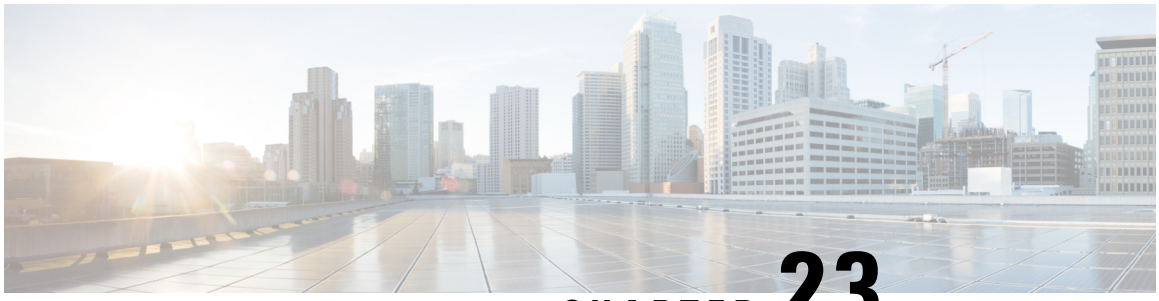
  Shows the static and dynamic MAC address entries.

History for Advanced Interface Configuration

Table 32: History for Advanced Interface Configuration

Feature Name	Releases	Feature Information
Maximum MTU is now 9198 bytes	9.1(6), 9.2(1)	The maximum MTU that the ASA can use is 9198 bytes (check for your model's exact limit at the CLI help). This value does not include the Layer 2 header. Formerly, the ASA let you specify the maximum MTU as 65535 bytes, which was inaccurate and could cause problems. If your MTU was set to a value higher than 9198, then the MTU is automatically lowered when you upgrade. In some cases, this MTU change can cause an MTU mismatch; be sure to set any connecting equipment to use the new MTU value. We modified the following screen: **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces** > **Edit Interface** > **Advanced**

Feature Name	Releases	Feature Information
Increased MTU size for the ASA on the Firepower 4100/9300 chassis	9.6(2)	You can set the maximum MTU to 9184 bytes on the Firepower 4100 and 9300; formerly, the maximum was 9000 bytes. This MTU is supported with FXOS 2.0.1.68 and later. We modified the following screen: **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces** > **Advanced**
Unique MAC address generation for single context mode	9.8(3), 9.8(4), 9.9(2)	You can now enable unique MAC address generation for VLAN subinterfaces in single context mode. Normally, subinterfaces share the same MAC address with the main interface. Because IPv6 link-local addresses are generated based on the MAC address, this feature allows for unique IPv6 link-local addresses. New or modified command: **mac-address auto** No ASDM support.
ASDM support for unique MAC address generation for single context mode	ASDM 7.15(1)	You can now enable unique MAC address generation for VLAN subinterfaces in single context mode in ASDM. Normally, subinterfaces share the same MAC address with the main interface. Because IPv6 link-local addresses are generated based on the MAC address, this feature allows for unique IPv6 link-local addresses. New or modified screen: **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces**
Secure Firewall 3100 auto-negotiation can be enabled or disabled for 1Gigabit and higher interfaces.	9.17(1)	Secure Firewall 3100 auto-negotiation can be enabled or disabled for 1Gigabit and higher interfaces. For other model SFP ports, the **no speed nonegotiate** option sets the speed to 1000 Mbps; the new command means you can set auto-negotiation and speed independently. New/Modified screens: **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces** > **Advanced**



CHAPTER 23

Traffic Zones

You can assign multiple interfaces to a *traffic zone*, which lets traffic from an existing flow exit or enter the ASA on any interface within the zone. This capability allows Equal-Cost Multi-Path (ECMP) routing on the ASA as well as external load balancing of traffic to the ASA across multiple interfaces.

- About Traffic Zones, on page 675
- Prerequisites for Traffic Zones, on page 681
- Guidelines for Traffic Zones, on page 682
- Configure a Traffic Zone, on page 684
- Monitoring Traffic Zones, on page 684
- Example for Traffic Zones, on page 686
- History for Traffic Zones, on page 689

About Traffic Zones

This section describes how you should use traffic zones in your network.

Non-Zoned Behavior

The Adaptive Security Algorithm takes into consideration the state of a packet when deciding to permit or deny the traffic. One of the enforced parameters for the flow is that traffic enters and exits the same interface. Any traffic for an existing flow that enters a different interface is dropped by the ASA.

Traffic zones let you group multiple interfaces together so that traffic entering or exiting *any* interface in the zone fulfills the Adaptive Security Algorithm security checks.

Related Topics

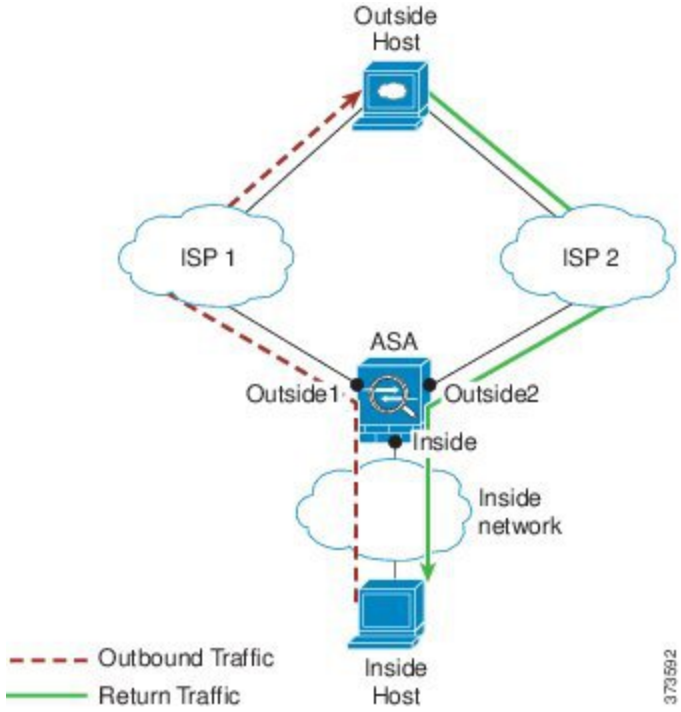
Stateful Inspection Overview, on page 13

Why Use Zones?

You can use zones to accommodate several routing scenarios.

Asymmetric Routing

In the following scenario, a connection was established between an inside host and an outside host through ISP 1 on the Outside1 interface. Due to asymmetric routing on the destination network, return traffic arrived from ISP 2 on the Outside2 interface.

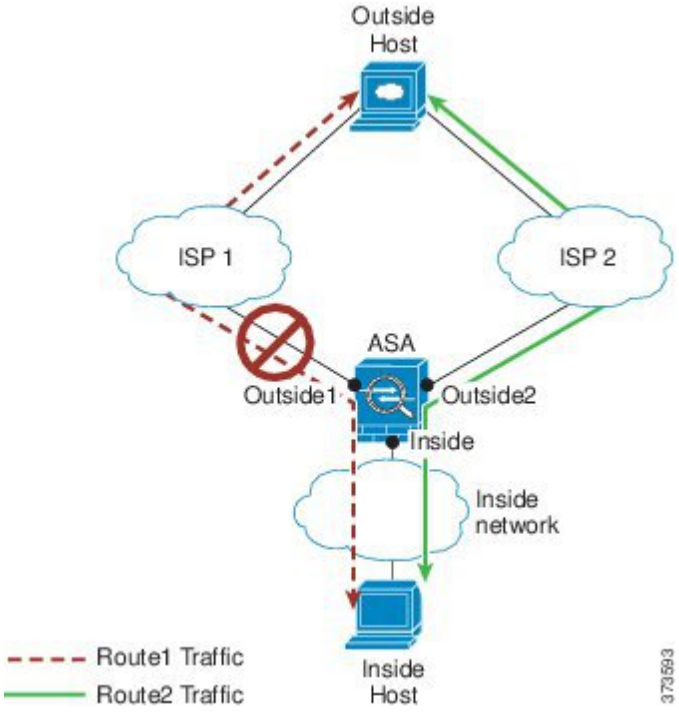


Non-Zoned Problem: The ASA maintains the connection tables on a per-interface basis. When the returning traffic arrives at Outside2, it will not match the connection table and will be dropped. For an ASA cluster, asymmetric routing when the cluster has multiple adjacencies to the same router can lead to unacceptible traffic loss.

Zoned Solution: The ASA maintains connection tables on a per-zone basis. If you group Outside1 and Outside2 into a zone, then when the returning traffic arrives at Outside2, it will match the per-zone connection table, and the connection will be allowed.

Lost Route

In the following scenario, a connection was established between an inside host and an outside host through ISP 1 on the Outside1 interface. Due to a lost or moved route between Outside1 and ISP 1, traffic needs to take a different route through ISP 2.

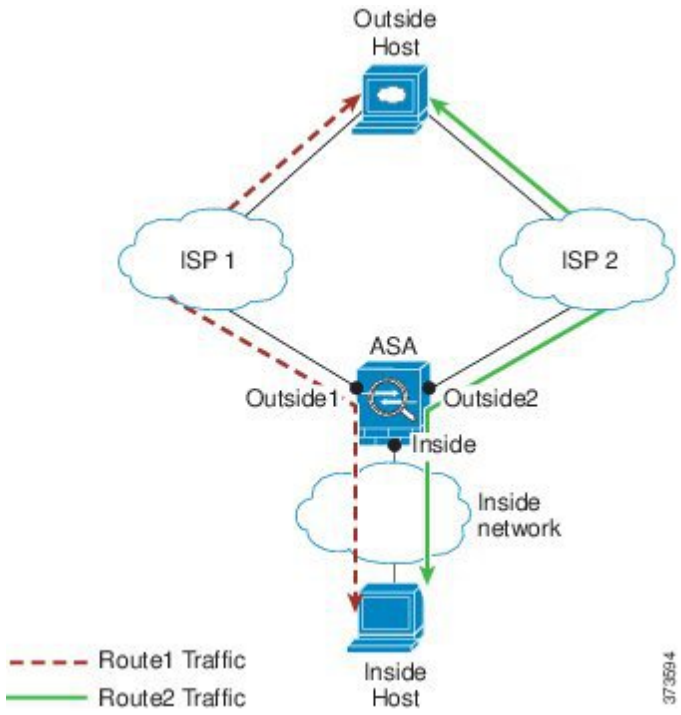


Non-Zoned Problem: The connection between the inside and outside host will be deleted; a new connection must be established using a new next-best route. For UDP, the new route will be used after a single packet drop, but for TCP, a new connection has to be reestablished.

Zoned Solution: The ASA detects the lost route and switches the flow to the new path through ISP 2. Traffic will be seamlessly forwarded without any packet drops.

Load Balancing

In the following scenario, a connection was established between an inside host and an outside host through ISP 1 on the Outside1 interface. A second connection was established through an equal cost route through ISP 2 on Outside2.



Non-Zoned Problem: Load-balancing across interfaces is not possible; you can only load-balance with equal cost routes on one interface.

Zoned Solution: The ASA load-balances connections across up to eight equal cost routes on all the interfaces in the zone.

Per-Zone Connection and Routing Tables

The ASA maintains a per-zone connection table so that traffic can arrive on any of the zone interfaces. The ASA also maintains a per-zone routing table for ECMP support.

ECMP Routing

The ASA supports Equal-Cost Multi-Path (ECMP) routing.

Non-Zoned ECMP Support

Without zones, you can have up to 8 equal cost static or dynamic routes per interface. For example, you can configure three default routes on the outside interface that specify different gateways:

```
route outside 0 0 10.1.1.2
route outside 0 0 10.1.1.3
route outside 0 0 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 is not supported across multiple interfaces, so you cannot define a route to the same destination on a different interface. The following route is disallowed when configured with any of the routes above:

```
route outside2 0 0 10.2.1.1
```

Zoned ECMP Support

With zones, you can have up to 8 equal cost static or dynamic routes across up to 8 interfaces within a zone. For example, you can configure three default routes across three interfaces in the zone:

```
route outside1 0 0 10.1.1.2
route outside2 0 0 10.2.1.2
route outside3 0 0 10.3.1.2
```

Similarly, your dynamic routing protocol can automatically configure equal cost routes. The ASA load-balances traffic across the interfaces with a more robust load balancing mechanism.

When a route is lost, the ASA seamlessly moves the flow to a different route.

How Connections Are Load-Balanced

The ASA load balances connections across equal cost routes using a hash made from the packet 6-tuple (source and destination IP address, source and destination port, protocol, and ingress interface). Unless the route is lost, a connection will stay on the chosen interface for its duration.

Packets within a connection are not load-balanced across routes; a connection uses a single route unless that route is lost.

The ASA does not consider the interface bandwidth or other parameters when load balancing. You should make sure all interfaces within the same zone have the same characteristics such as MTU, bandwidth, and so on.

The load-balancing algorithm is not user configurable.

Falling Back to a Route in Another Zone

When a route is lost on an interface, if there are no other routes available within the zone, then the ASA will use a route from a different interface/zone. If this backup route is used, then you may experience packet drops as with non-zoned routing support.

Interface-Based Security Policy

Zones allow traffic to and from any interface in the zone, but the security policy itself (access rules, NAT, and so on) is still applied per interface, not per zone. If you configure the same security policy for all interfaces within the zone, then you can successfully implement ECMP and load balancing for that traffic. For more information about required parallel interface configuration, see Prerequisites for Traffic Zones, on page 681.

Supported Services for Traffic Zones

The following services are supported with zones:

- Access Rules

- NAT
- Service Rules, except for QoS traffic policing.
- Routing

You can also configure to- and from-the-box services listed in To- and From-the-Box Traffic, on page 681, although full zoned support is not available.

Do not configure other services (such as VPN or Botnet Traffic Filter) for interfaces in a traffic zone; they may not function or scale as expected.



Note For detailed information about how to configure the security policy, see Prerequisites for Traffic Zones, on page 681.

Security Levels

The first interface that you add to a zone determines the security level of the zone. All additional interfaces must have the same security level. To change the security level for interfaces in a zone, you must remove all but one interface, and then change the security levels, and re-add the interfaces.

Primary and Current Interface for the Flow

Each connection flow is built based on the initial ingress and egress interfaces. These interfaces are the *primary* interfaces.

If a new egress interface is used because of route changes or asymmetric routing, then the new interfaces are the *current* interfaces.

Joining or Leaving a Zone

When you assign an interface to a zone, any connections on that interface are deleted. The connections must be reestablished.

If you remove an interface from a zone, any connections that have the interface as the primary interface are deleted. The connections must be reestablished. If the interface is the current interface, the ASA moves the connections back to the primary interface. The zone route table is also refreshed.

Intra-Zone Traffic

To allow traffic to *enter* one interface and *exit* another in the same zone, enable **Configuration > Device Setup > Interface Settings > Interfaces > Enable traffic between two or more hosts connected to the same interface**, which allows traffic to enter and exit the same interface, as well as **Configuration > Device Setup > Interface Settings > Interfaces > Enable traffic between two or more interfaces which are configured with same security level**, which allows traffic between same-security interfaces. Otherwise, a flow cannot be routed between two interfaces in the same zone.

To- and From-the-Box Traffic

- You cannot add management-only or management-access interfaces to a zone.
- For management traffic on regular interfaces in a zone, only asymmetric routing on existing flows is supported; there is no ECMP support.
- You can configure a management service on only one zone interface, but to take advantage of asymmetric routing support, you need to configure it on all interfaces. Even when the configurations are parallel on all interfaces, ECMP is not supported.
- The ASA supports the following to- and from-the-box services in a zone:
  - Telnet
  - SSH
  - HTTPS
  - SNMP
  - Syslog

Overlapping IP Addresses Within a Zone

For non-zoned interfaces, the ASA supports overlapping IP address networks on interfaces so long as you configure NAT properly. However, overlapping networks are not supported on interfaces in the same zone.

Prerequisites for Traffic Zones

- Configure all interface parameters including the name, IP address, and security level. Note that the security level must match for all interfaces in the zone. You should plan to group together like interfaces in terms of bandwidth and other Layer 2 properties.
- Configure the following services to match on all zone interfaces:
  - Access Rules—Apply the same access rule to all zone member interfaces, or use a global access rule.

    For example:

    ```
    access-list ZONE1 extended permit tcp any host WEBSERVER1 eq 80
    access-group ZONE1 in interface outside1
    access-group ZONE1 in interface outside2
    access-group ZONE1 in interface outside3
    ```

  - NAT—Configure the same NAT policy on all member interfaces of the zone or use a global NAT rule (in other words, use “any” to represent the zone interfaces in the NAT rule).

    Interface PAT is not supported.

    For example:

    ```
    object network WEBSERVER1
    ```

```
host 10.9.9.9 255.255.255.255
nat (inside,any) static 209.165.201.9
```

Note When you use interface-specific NAT and PAT pools, the ASA cannot switch connections over in case of the original interface failure.

If you use interface-specific PAT pools, multiple connections from the same host might load-balance to different interfaces and use different mapped IP addresses. Internet services that use multiple concurrent connections may not work correctly in this case.

- Service Rules—Use the global service policy, or assign the same policy to each interface in a zone.

  QoS traffic policing is not supported.

  For example:

```
service-policy outside_policy interface outside1
service-policy outside_policy interface outside2
service-policy outside_policy interface outside3
```

Note For VoIP inspections, zone load balancing can cause increased out-of-order packets. This situation can occur because later packets might reach the ASA before earlier packets that take a different path. Symptoms of out-of-order packets include:

- Higher memory utilization at intermediate nodes (firewall and IDS) and the receiving end nodes if queuing is used.
- Poor video or voice quality.

To mitigate these effects, we recommend that you use IP addresses only for load distribution for VoIP traffic.

- Configure routing with ECMP zone capabilities in mind.

Guidelines for Traffic Zones

Firewall Mode

Supported in routed firewall mode only. Does not support transparent firewall mode or bridge group interfaces in routed mode.

Failover

- You cannot add the failover or state link to a zone.

- In Active/Active failover mode, you can assign an interface in each context to an asymmetrical routing (ASR) group. This service allows traffic returning on a similar interface on the peer unit to be restored to the original unit. You cannot configure both ASR groups and traffic zones within a context. If you configure a zone in a context, none of the context interfaces can be part of an ASR group. See Configure Support for Asymmetrically Routed Packets (Active/Active Mode), on page 311 for more information about ASR groups.
- Only the primary interfaces for each connection are replicated to the standby unit; current interfaces are not replicated. If the standby unit becomes active, it will assign a new current interface if necessary.

Clustering

- You cannot add the cluster control link to a zone.

Model Guidelines

You cannot add Firepower 1010 and Secure Firewall 1210/1220 switch ports and VLAN interfaces to a zone.

Additional Guidelines

- You can create a maximum of 256 zones.
- You can add the following types of interfaces to a zone:
  - Physical
  - VLAN
  - EtherChannel
- You cannot add the following types of interfaces:
  - Management-only
  - Management-access
  - Failover or state link
  - Cluster control link
  - Member interfaces in an EtherChannel
  - VNI; also, if a regular data interface is marked as nve-only, it cannot be a member of a zone.
  - BVI, or bridge group member interfaces.
- An interface can be a member of only one zone.
- You can include up to 8 interfaces per zone.
- For ECMP, you can add up to 8 equal cost routes per zone, across all zone interfaces. You can also configure multiple routes on a single interface as part of the 8 route limit.
- When you add an interface to a zone, all static routes for those interfaces are removed.
- You cannot enable DHCP Relay on an interface in a traffic zone.

- The ASA does not support fragmented packet reassembly for fragments that are load-balanced to separate interfaces; those fragments will be dropped.
- PIM/IGMP Multicast routing is not supported on interfaces in a zone.

Configure a Traffic Zone

Configure a named zone, and assign interfaces to the zone.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Interface Settings** > **Zones**, and click **Add**.

You can alternately assign an interface to a zone from the **Configuration > Device Setup > Interface Settings > Interfaces > Add Interface** dialog box.

Step 2 Name the zone with a name up to 48 characters in length.

Step 3 Add one or more interfaces to the **Member** area. Ensure all interfaces have the same security level.

Step 4 Click **Apply**.

Monitoring Traffic Zones

This section describes how to monitor traffic zones.

Zone Information

- **show zone** [*name*]

  Shows zone ID, context, security level, and members.

  See the following output for the **show zone** command:

```
ciscoasa# show zone outside-zone

Zone: zone-outside id: 2
Security-level: 0
Context: test-ctx
 Zone Member(s) : 2
  outside1       GigabitEthernet0/0
  outside2       GigabitEthernet0/1
```

- **show nameif zone**

  Shows the interface names and zone names.

  See the following output for the **show nameif zone** command:

```
ciscoasa# show nameif zone
Interface                Name                   zone-name      Security
GigabitEthernet0/0       inside-1               inside-zone     100
GigabitEthernet0/1.21    inside                 inside-zone     100
GigabitEthernet0/1.31    4                                        0
GigabitEthernet0/2       outside                outside-zone   0
Management0/0            lan                                      0
```

Zone Connections

- **show conn** [**long** | **detail**] [**zone** *zone_name* [**zone** *zone_name*] [...]]

  The **show conn zone** command displays connections for a zone. The **long** and **detail** keywords show the primary interface on which the connection was built and the one in the brackets is the current interface used to forward the traffic or the interface the last packet came from. Thus, the current interface in case of a connection coming from multiple interfaces can show different interfaces at different times depending on when the show conn command was issued.

  See the following output for the **show conn long zone** command:

```
ciscoasa# show conn long zone zone-inside zone zone-outside

TCP outside-zone:outside1(outside2): 10.122.122.1:1080
inside-zone:inside1(inside2): 10.121.121.1:34254, idle 0:00:02, bytes 10, flags UO
```

- **show asp table zone**

  Shows the accelerated security path tables for debugging purposes.

- **show local-host** [**zone** *zone_name* [**zone** *zone_name*] [...]]

  Shows the network states of local hosts within a zone.

  See the following output for the **show local-host zone** command. The primary interface is listed first, and the current interface is in parentheses.

```
ciscoasa# show local-host zone outside-zone

Zone:outside-zone: 4 active, 5 maximum active, 0 denied
local host: <10.122.122.1>,
    TCP flow count/limit = 3/unlimited
    TCP embryonic count to host = 0
    TCP intercept watermark = unlimited
    UDP flow count/limit = 0/unlimited

  Conn:
 TCP outside-zone:outside1(outside2): 10.122.122.1:1080
inside-zone:inside1(inside2): 10.121.121.1:34254, idle 0:00:02, bytes 10, flags UO
```

Zone Routing

- **show route zone**

  Shows the routes for zone interfaces.

See the following output for the **show route zone** command:

```
ciscoasa# show route zone

Codes: C - connected, S - static, I - IGRP, 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, E - EGP
       i - IS-IS, 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

Gateway of last resort is not set

S    192.168.105.1 255.255.255.255 [1/0] via 172.16.1.1, outside-zone:outside1
C    192.168.212.0 255.255.255.0 is directly connected, lan-zone:inside,
C    172.16.1.0 255.255.255.0 is directly connected, wan-zone:outside2
S    10.5.5.0 255.255.255.0 [1/0] via 172.16.1.1, wan-zone:outside2
O    10.2.2.1 255.255.255.255 [110/11] via 192.168.212.3, 2:09:24, lan-zone:inside
O    10.1.1.1 255.255.255.255 [110/11] via 192.168.212.2, 2:09:24, lan-zone:inside
```

- **show asp table routing**

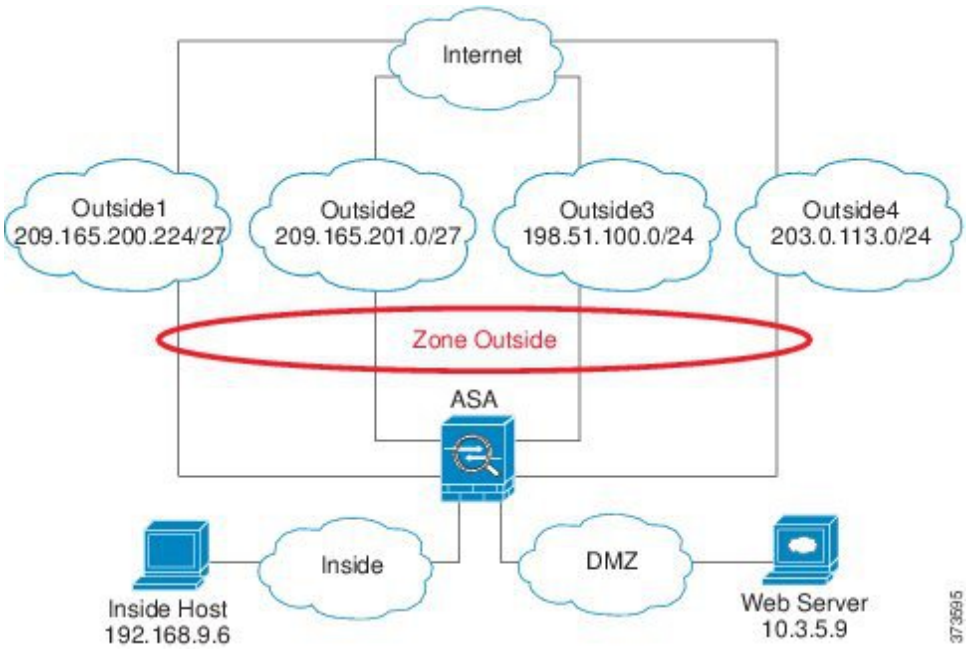
  Shows the accelerated security path tables for debugging purposes, and shows the zone associated with each route.

  See the following output for the **show asp table routing** command:

```
ciscoasa# show asp table routing
route table timestamp: 60
in   255.255.255.255 255.255.255.255 identity
in   10.1.0.1       255.255.255.255 identity
in   10.2.0.1       255.255.255.255 identity
in   10.6.6.4         255.255.255.255 identity
in   10.4.4.4         255.255.255.255 via 10.4.0.10 (unresolved, timestamp: 49)
in   172.0.0.67       255.255.255.255 identity
in   172.0.0.0       255.255.255.0   wan-zone:outside2
in   10.85.43.0      255.255.255.0   via 10.4.0.3 (unresolved, timestamp: 50)
in   10.85.45.0      255.255.255.0   via 10.4.0.20 (unresolved, timestamp: 51)
in   192.168.0.0        255.255.255.0   mgmt
in   192.168.1.0       255.255.0.0     lan-zone:inside
out  255.255.255.255 255.255.255.255 mgmt
out  172.0.0.67       255.255.255.255 mgmt
out  172.0.0.0        255.255.255.0   mgmt
out  10.4.0.0       240.0.0.0       mgmt
out  255.255.255.255 255.255.255.255 lan-zone:inside
out  10.1.0.1       255.255.255.255 lan-zone:inside
out  10.2.0.0       255.255.0.0     lan-zone:inside
out  10.4.0.0       240.0.0.0       lan-zone:inside
```

Example for Traffic Zones

The following example assigns 4 VLAN interfaces to the outside zone, and configures 4 equal cost default routes. PAT is configured for the inside interface, and a web server is available on a DMZ interface using static NAT.




```
interface gigabitethernet0/0
  no shutdown
  description outside switch 1
interface gigabitethernet0/1
  no shutdown
  description outside switch 2

interface gigabitethernet0/2
  no shutdown
  description inside switch

zone outside

interface gigabitethernet0/0.101
  vlan 101
  nameif outside1
  security-level 0
  ip address 209.165.200.225 255.255.255.224
  zone-member outside
  no shutdown

interface gigabitethernet0/0.102
  vlan 102
  nameif outside2
  security-level 0
  ip address 209.165.201.1 255.255.255.224
  zone-member outside
  no shutdown

interface gigabitethernet0/1.201
  vlan 201
  nameif outside3
  security-level 0
  ip address 198.51.100.1 255.255.255.0
  zone-member outside
  no shutdown
```

```
interface gigabitethernet0/1.202
  vlan 202
  nameif outside4
  security-level 0
  ip address 203.0.113.1 255.255.255.0
  zone-member outside
  no shutdown

interface gigabitethernet0/2.301
  vlan 301
  nameif inside
  security-level 100
  ip address 192.168.9.1 255.255.255.0
  no shutdown

interface gigabitethernet0/2.302
  vlan 302
  nameif dmz
  security-level 50
  ip address 10.3.5.1 255.255.255.0
  no shutdown

# Static NAT for DMZ web server on any destination interface
object network WEBSERVER
  host 10.3.5.9 255.255.255.255
  nat (dmz,any) static 209.165.202.129 dns

# Dynamic PAT for inside network on any destination interface
object network INSIDE
  subnet 192.168.9.0 255.255.255.0
  nat (inside,any) dynamic 209.165.202.130

# Global access rule for DMZ web server
access-list WEB-SERVER extended permit tcp any host WEBSERVER eq 80
access-group WEB-SERVER global

# 4 equal cost default routes for outside interfaces
route outside1 0 0 209.165.200.230
route outside2 0 0 209.165.201.10
route outside3 0 0 198.51.100.99
route outside4 0 0 203.0.113.87
# Static routes for NAT addresses - see redistribute static command
route dmz 209.165.202.129 255.255.255.255 10.3.5.99
route inside 209.165.202.130 255.255.255.255 192.168.9.99

# The global service policy
class-map inspection_default
  match default-inspection-traffic
policy-map type inspect dns preset_dns_map
  parameters
    message-length maximum client auto
    message-length maximum 512
    dns-guard
    protocol-enforcement
    nat-rewrite
policy-map global_policy
  class inspection_default
    inspect dns preset_dns_map
    inspect ftp
    inspect h323 h225 _default_h323_map
    inspect h323 ras _default_h323_map
    inspect ip-options _default_ip_options_map
    inspect netbios
    inspect rsh
```

```
    inspect rtsp
    inspect skinny
    inspect esmtp _default_esmtp_map
    inspect sqlnet
    inspect sunrpc
    inspect tftp
    inspect sip
    inspect xdmcp
service-policy global_policy global
```

History for Traffic Zones

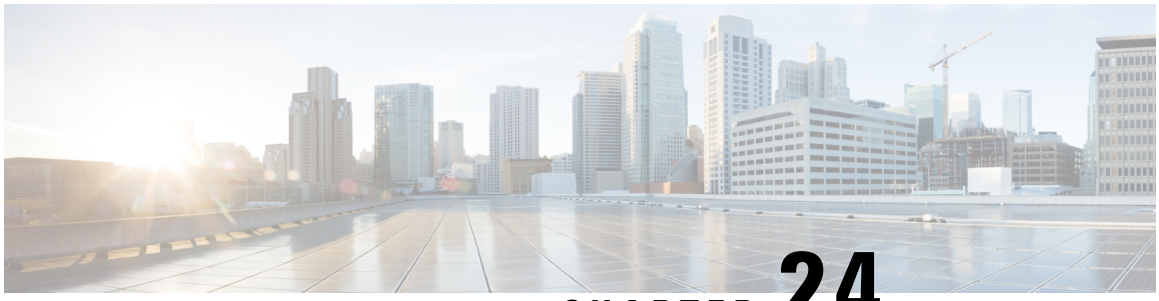
Feature Name	Platform Releases	Description
Traffic Zones	9.3(2)	You can group interfaces together into a traffic zone to accomplish traffic load balancing (using Equal Cost Multi-Path (ECMP) routing), route redundancy, and asymmetric routing across multiple interfaces. **Note** You cannot apply a security policy to a named zone; the security policy is interface-based. When interfaces in a zone are configured with the same access rule, NAT, and service policy, then load-balancing and asymmetric routing operate correctly. We introduced or modified the following screens: **Configuration > Device Setup > Interface Parameters > Zones** **Configuration > Device Setup > Interface Parameters > Interfaces**.
clear local-host command	9.14(1)	The **clear local-host** command and all of its attributes and keywords were deprecated. They will be removed in a future release.



PART IV

Basic Settings

- Basic Settings, on page 693
- DHCP and DDNS Services, on page 713
- Digital Certificates, on page 731
- ARP Inspection and the MAC Address Table, on page 763



CHAPTER 24

Basic Settings

This chapter describes how to configure basic settings on the ASA that are typically required for a functioning configuration.

- Set the Hostname, Domain Name, and the Enable and Telnet Passwords, on page 693
- Set the Date and Time, on page 695
- Configure the Master Passphrase, on page 698
- Configure the DNS Servers, on page 701
- Configure the Hardware Bypass and Dual Power Supply (Cisco ISA 3000), on page 704
- Adjust ASP (Accelerated Security Path) Performance and Behavior, on page 705
- Monitoring the DNS Cache, on page 707
- History for Basic Settings, on page 708

Set the Hostname, Domain Name, and the Enable and Telnet Passwords

To set the hostname, domain name, and the enable and Telnet passwords, perform the following steps.

Before you begin

Before you set the hostname, domain name, and the enable and Telnet passwords, check the following requirements:

- In multiple context mode, you can configure the hostname and domain name in both the system and context execution spaces.
- For the enable and Telnet passwords, set them in each context; they are not available in the system.
- To change from the system to a context configuration, in the **Configuration > Device List** pane, double-click the context name under the active device IP address.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Device Name/Password**.

Step 2 Enter the hostname. The default hostname is “ciscoasa.”

The hostname appears in the command line prompt, and if you establish sessions to multiple devices, the hostname helps you keep track of where you enter commands. The hostname is also used in syslog messages.

For multiple context mode, the hostname that you set in the system execution space appears in the command line prompt for all contexts. The hostname that you optionally set within a context does not appear in the command line; it can be used for a banner.

Step 3 Enter the domain name. The default domain name is default.domain.invalid.

The ASA appends the domain name as a suffix to unqualified names. For example, if you set the domain name to "example.com" and specify a syslog server by the unqualified name of "jupiter," then the ASA qualifies the name to "jupiter.example.com."

Step 4 Change the privileged mode (enable) password. The default password is blank, but you are prompted to change it the first time you enter the **enable** command at the CLI.

The enable password lets you enter privileged EXEC mode if you do not configure enable authentication. The enable password also lets you log into ASDM with a blank username if you do not configure HTTP authentication. ASDM does not enforce the enable password change like CLI access does.

a) Check the **Change the privileged mode password** check box.
b) Enter the new password, and then confirm the new password. Set a case-sensitive password of 8 to 127 characters long. It can be any combination of ASCII printable characters (character codes 32-126), with the following exceptions:

- No spaces
- No question marks
- You cannot use three or more consecutive sequential or repetitive ASCII characters. For example, the following passwords will be rejected:
  - **abc**user1
  - user**543**
  - user**aaaa**
  - user2**666**

You cannot reset the password to a blank value.

Step 5 Set the login password for Telnet access. There is no default password.

The login password is used for Telnet access when you do not configure Telnet authentication.

a) Check the **Change the password to access the console of the security appliance** check box.
b) Enter the old password (for a new ASA, leave this field blank), new password, then confirm the new password. The password can be up to 16 characters long. It can be any combination of ASCII printable characters (character codes 32-126), with the exception of spaces and the question mark.

Step 6 Click **Apply** to save your changes.

Set the Date and Time

Note Do not set the date and time for the Firepower 4100/9300; the ASA receives these settings from the chassis.

Set the Date and Time Using an NTP Server

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 multiple NTP servers. The ASA chooses the server with the lowest stratum—a measure of how reliable the data is.

Time derived from an NTP server overrides any time set manually.

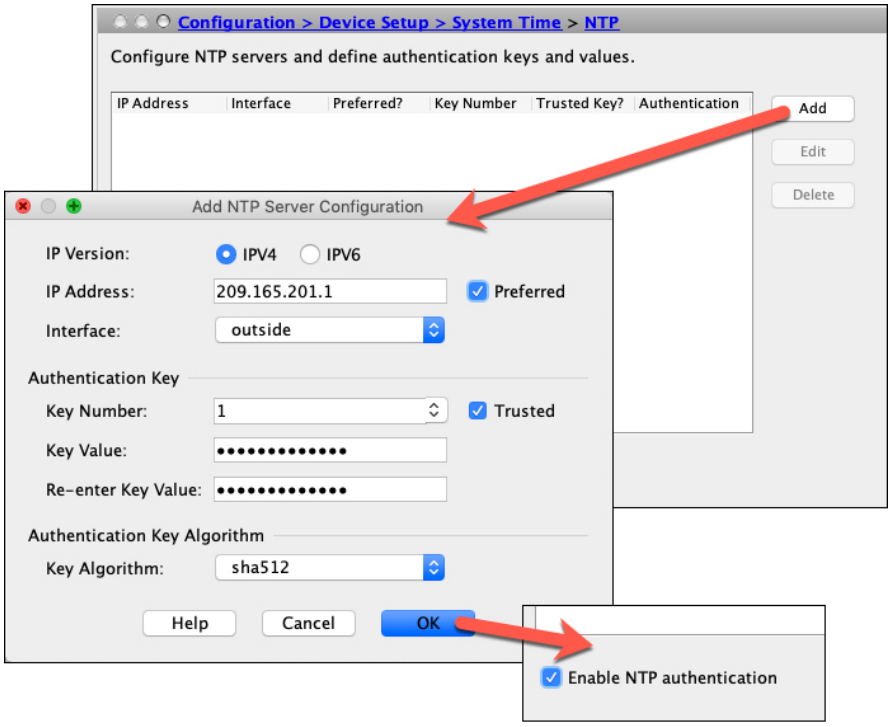
The ASA supports NTPv4.

Before you begin

In multiple context mode, you can set the time in the system configuration only.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **System Time** > **NTP**.



Step 2 Click **Add** to display the **Add NTP Server Configuration** dialog box.

Step 3 Enter the NTP server **IPv4** or **IPv6 IP Address**.

You cannot enter a hostname for the server; the ASA does not support DNS lookup for the NTP server.

Step 4 (Optional) Check the **Preferred** check box to set this server as a preferred server.

NTP uses an algorithm to determine which server is the most accurate and synchronizes to it. If servers are of similar accuracy, then the preferred server is used. However, if a server is significantly more accurate than the preferred one, the ASA uses the more accurate one.

Step 5 (Optional) Choose the **Interface** from the drop-down list.

This setting specifies the outgoing interface for NTP packets. If the interface is blank, then the ASA uses the default admin context interface according to the management routing table.

Step 6 (Optional) Configure NTP authentication.

a) Enter a **Key Number** between 1 and 4294967295, or choose an existing key number from the drop-down list if you previously created a key for another NTP server that you want to reuse.

This setting specifies the key ID for this authentication key, which enables you to use authentication to communicate with the NTP server. The NTP server packets must also use this key ID.

b) Check the **Trusted** check box.
c) Enter the **Key Value**, which is a string up to 32 characters long, and then re-enter the key value.
d) Choose a **Key Algorithm** from the drop-down list.
e) Click **OK**.

Step 7 Check the **Enable NTP authentication** check box to turn on NTP authentication.

Step 8 Click **Apply** to save your changes.

Set the Date and Time Manually

To set the date and time manually, perform the following steps:

Before you begin

In multiple context mode, you can set the time in the system configuration only.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **System Time** > **Clock**.

Step 2 Choose the time zone from the drop-down list. This setting specifies the time zone as GMT plus or minus the appropriate number of hours. If you select the Eastern Time, Central Time, Mountain Time, or Pacific Time zone, then the time adjusts automatically for daylight savings time, from 2:00 a.m. on the second Sunday in March to 2:00 a.m. on the first Sunday in November.

Note
Changing the time zone on the ASA may drop the connection to intelligent SSMs.

Step 3 Click the **Date** drop-down list to display a calendar. Then find the correct date using the following methods:

- Click the name of the month to display a list of months, then click the desired month. The calendar updates to that month.
- Click the year to change the year. Use the up and down arrows to scroll through the years, or enter a year in the entry field.
- Click the arrows to the right and left of the month and year to scroll the calendar forward and backward one month at a time.
- Click a day on the calendar to set the date.

Step 4 Enter the time manually in hours, minutes, and seconds.

Step 5 Click **Update Display Time** to update the time shown in the bottom right corner of the ASDM pane. The current time updates automatically every ten seconds.

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 ASA device to be a transparent clock. The ASA device does not synchronize its clock with the PTP clocks. The ASA 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

- This feature is only available on the ISA 3000.
- Use of PTP is supported in single context mode only.
- Cisco PTP supports multicast PTP messages only.
- PTP is enabled on all ISA 3000 interfaces in transparent mode by default. In routed mode, you must add the necessary configuration to ensure that the PTP packets are allowed to flow through the device.
- PTP is available only for IPv4 networks, not for IPv6 networks.
- PTP configuration is supported on physical Ethernet interfaces, whether stand-alone or bridge group members. It is not supported on:
  - Management interface.
  - Subinterfaces, EtherChannels, BVIs. 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. In transparent firewall mode, the access list configuration to allow PTP traffic is configured by default. PTP traffic is identified by UDP ports 319 and 320, and destination IP address 224.0.1.129, so in routed firewall mode any ACL that allows this traffic should be acceptable.
- In routed firewall mode, you must also enable multicast routing for PTP multicast groups:
  - Enter the global configuration mode command **multicast-routing**.
  - And for each interface that is not a bridge group member, and on which PTP is enabled, enter the interface configuration command **igmp join-group 224.0.1.129** to statically enable PTP multicast group membership. This command is not supported or needed for bridge group members.

Procedure

Step 1 Select **Configuration** > **Device Management** > **PTP**.

Step 2 Enter the **Domain value**.

This is the domain number for all ports on the device. Packets received on a different domain are treated like regular multicast packets and will not undergo any PTP processing. This value can be from zero to 255; the default value is zero. Enter the domain number that is configured on the PTP devices in your network.

Step 3 (Optional) Select **Enable End-to-End Transparent Clock Mode** to enable End-to-End Transparent mode on all PTP-enabled interfaces.

A transparent clock is a clock which compensates for its delays by measuring the residence times and updating the `correctionField` in the PTP packet.

Step 4 Enable PTP on one or more device interfaces by selecting an interface and clicking **Enable** or **Disable**.

Enable PTP on each interface through which the system can contact a PTP clock in the configured domain.

Step 5 Click **Apply**.

What to do next

You can choose **Monitoring** > **Properties** > **PTP** to view PTP clock and interface/port information.

Configure the Master Passphrase

The master passphrase allows you to securely store plain text passwords in encrypted format and provides a key that is used to universally encrypt or mask all passwords, without changing any functionality. Features that use the master passphrase include the following:

- OSPF
- EIGRP

- VPN load balancing
- VPN (remote access and site-to-site)
- Failover
- AAA servers
- Logging
- Shared licenses

Add or Change the Master Passphrase

To add or change the master passphrase, perform the following steps.

Before you begin

- This procedure will only be accepted in a secure session, for example by console, SSH, or ASDM via HTTPS.
- If failover is enabled but no failover shared key is set, an error message appears if you change the master passphrase, informing you that you must enter a failover shared key to protect the master passphrase changes from being sent as plain text.

  Choose **Configuration > Device Management > High Availability > Failover**, enter any character in the **Shared Key** field or 32 hexadecimal numbers (0-9A-Fa-f) if a failover hexadecimal key is selected, except a backspace. Then click **Apply**.
- Enabling or changing password encryption in Active/Standby failover causes a **write standby**, which replicates the active configuration to the standby unit. Without this replication, the encrypted passwords on the standby unit will differ even though they use the same passphrase; configuration replication ensures that the configurations are the same. For Active/Active failover, you must manually enter **write standby**. A **write standby** can cause traffic interruption in Active/Active mode, because the configuration is cleared on the secondary unit before the new configuration is synced. You should make all contexts active on the primary ASA using the **failover active group 1** and **failover active group 2** commands, enter **write standby**, and then restore the group 2 contexts to the secondary unit using the **no failover active group 2** command.

Procedure

Step 1 Choose one of the following options:

- In single context mode, choose **Configuration > Device Management > Advanced > Master Passphrase**.
- In multiple context mode, choose **Configuration > Device Management > Device Administration > Master Passphrase**.

Step 2 Check the **Advanced Encryption Standard (AES) password encryption** check box.

If no master passphrase is in effect, a warning message appears when you click Apply. You can click OK or Cancel to continue.

If you later disable password encryption, all existing encrypted passwords are left unchanged, and as long as the master passphrase exists, the encrypted passwords will be decrypted, as required by the application.

Step 3 Check the **Change the encryption master passphrase** check box to enable you to enter and confirm your new master passphrases. By default, they are disabled.

Your new master passphrase must be between 8 and 128 characters long.

If you are changing an existing passphrase, you must enter the old passphrase before you can enter a new one.

Leave the **New** and **Confirm master passphrase** fields blank to delete the master passphrase.

Step 4 Click **Apply**.

Disable the Master Passphrase

Disabling the master passphrase reverts encrypted passwords into plain text passwords. Removing the passphrase might be useful if you downgrade to a previous software version that does not support encrypted passwords.

Before you begin

- You must know the current master passphrase to disable it.
- This procedure works only in a secure session; that is, by Telnet, SSH, or ASDM via HTTPS.

To disable the master passphrase, perform the following steps:

Procedure

Step 1 Choose one of the following options:

- In single context mode, choose **Configuration > Device Management > Advanced > Master Passphrase**.
- In multiple context mode, choose **Configuration > Device Management > Device Administration > Master Passphrase**.

Step 2 Check the **Advanced Encryption Standard (AES) password encryption** check box.

If no master passphrase is in effect, a warning statement appears when you click Apply. Click OK or Cancel to continue.

Step 3 Check the Change the encryption master passphrase check box.

Step 4 Enter the old master passphrase in the **Old master passphrase** field. You must provide the old master passphrase to disable it.

Step 5 Leave the **New master passphrase** and the **Confirm master passphrase** fields empty.

Step 6 Click **Apply**.

Configure the DNS Servers

You need to configure a DNS server so that the ASA can resolve host names to IP addresses. You also must configure a DNS server to use fully qualified domain names (FQDN) network objects in access rules.

Some ASA features require use of a DNS server to access external servers by domain name. Other features, such as the **ping** or **traceroute** command, let you enter a name that you want to ping or traceroute, and the ASA can resolve the name by communicating with a DNS server. Many SSL VPN and certificate commands also support names.

By default, there is a default DNS server group called DefaultDNS. You can create multiple DNS server groups: one group is the default, while other groups can be associated with specific domains. A DNS request that matches a domain associated with a DNS server group will use that group. For example, if you want traffic destined to inside eng.cisco.com servers to use an inside DNS server, you can map eng.cisco.com to an inside DNS group. All DNS requests that do not match a domain mapping will use the default DNS server group, which has no associated domains. For example, the DefaultDNS group can include a public DNS server available on the outside interface. Other DNS server groups can be configured for VPN tunnel groups. See the **tunnel-group** command in the command reference for more information.



Note The ASA has limited support for using the DNS server, depending on the feature.

Before you begin

Make sure that you configure the appropriate routing and access rules for any interface on which you enable DNS domain lookup so you can reach the DNS server.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **DNS** > **DNS Client**.

Step 2 Choose one of the following options in the **DNS Setup** area:

- **Configure one DNS server group**—This option defines the servers in the DefaultDNS group.
- **Configure multiple DNS server groups**—With this option, you can configure the DefaultDNS group as well as other groups that you can associate with specific domains, and groups for use with remote access SSL VPN group policies. Even if you configure the DefaultDNS group only, you must select this option if you want to alter the timeout and other characteristics used with the group.

Step 3 If you select **Configure one DNS server group**, configure the servers in the DefaultDNS group.

a) In **Primary DNS Server**, enter the IP address of the DNS server that should be used whenever it is available. For this server and each secondary server, optionally specify the *interface_name* through which the ASA communicates with the server. If you do not specify the interface, the ASA checks the data routing table; if there are no matches, it then checks the management-only routing table.

b) Click **Add** to add secondary DNS servers.

You can add up to six DNS servers. The ASA tries each DNS server in order until it receives a response. Use the **Move Up**/**Move Down** buttons to put the servers in priority order.

c) Enter a DNS domain name appended to the hostname if it is not fully qualified.

Step 4 If you select **Configure multiple DNS server groups**, define the server group properties.

a) Click **Add** to create a new group, or select a group and click **Edit**.

The DefaultDNS group is always listed.

b) Configure the group properties.

- **Server IP Address to Add**, **Source Interface**—Enter the IP address of a DNS server and click **Add>>**. For each server, optionally specify the *interface_name* through which the ASA communicates with the server. If you do not specify the interface, the ASA checks the management-only routing table; if there are no matches, it then checks the data routing table.

  You can add up to six DNS servers. The ASA tries each DNS server in order until it receives a response. Use the **Move Up**/**Move Down** buttons to put the servers in priority order.

- **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 ASA retries the list of servers, this timeout doubles.

- **Retries**—The number of times, from 0 to 10, to retry the list of DNS servers when the ASA does not receive a response.

- **Expire Entry Timer** (DefaultDNS or active group only)—The minimum 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. This option is used when resolving FQDN network objects only.

- **Poll Timer** (DefaultDNS or active group only)—The time, in minutes, of the polling cycle used to resolve FQDN network/host objects to IP addresses. FQDN objects are resolved only if they are used in a firewall policy. The timer determines the maximum time between resolutions; the DNS entry's time-to-live (TTL) value is also used to determine when to update to IP address resolution, so individual FQDNs might be resolved more frequently than the polling cycle. The default is 240 (four hours). The range is 1 to 65535 minutes.

- **Domain Name** (DefaultDNS or active group only)—The domain name appended to the hostname if it is not fully qualified.

c) Click **OK**.

d) If you have multiple groups, you can change the default group by selecting it and clicking **Set Active**.

You can only use a group as the default if it does not have any domains mapped to it (see Step 8, on page 703).

Step 5 Ensure that DNS lookup is enabled on at least one interface. In the **DNS Lookup** interface list, below the DNS server group table, click in the **DNS Enabled** column and select **True** to enable lookup on the interface.

Make sure to enable DNS lookup on all interfaces that will be used to access DNS servers.

If you do not enable DNS lookup on an interface, then the DNS server **Source Interface** or the interface found using the routing table cannot be used.

Step 6 (Optional) Under **Trusted DNS Server**, configure the options for determining which servers to trust when resolving domain names in network-service objects.

a) (Optional) Add or remove explicitly-configured trusted DNS servers.

- Click **Add** to add a new server, then select the IP type (IPv4 or IPv6), enter the IP address of the server, and click **OK**.
- Select a server and click **Edit** to change the address.
- Select a server and click **Delete** to remove it from the trusted server list.

b) Select or deselect the following options:

- **Any**—Trust every DNS server, snoop them all. This option is disabled by default.
- **Configured-Servers**—Whether servers configured in DNS server groups should be trusted. This option is enabled by default.
- **DHCP-Client**—Whether the servers that are learned by snooping messages between a DHCP client and DHCP server are considered trusted DNS servers. This option is enabled by default.
- **DHCP-Pools**—Whether the DNS servers that are configured in the DHCP pools for clients that obtain addresses through DHCP servers running on the device interfaces should be trusted. This option is enabled by default.
- **DHCP-Relay**—Whether the servers that are learned by snooping DHCP relay messages between a DHCP client and DHCP server are considered trusted DNS servers. This option is enabled by default.

Step 7 (Optional) Check the **Enable DNS Guard on all interfaces** check box to enforce one DNS response per query.

You can also set DNS Guard when configuring DNS inspection. For a given interface, the DNS Guard setting configured in DNS inspection takes precedence over this global setting. By default, DNS inspection is enabled on all interfaces with DNS Guard enabled.

Step 8 (Optional) Map domains to specific DNS server groups.

You can map up to 30 domains. You cannot map the same domain to multiple DNS server groups, but you can map multiple domains to the same server group. Do not map any domains to the group you want to use for the default (for example, DefaultDNS).

a) In the **DNS Group Map** area, check **Enable DNS Group Map**.
b) Click **Add**.

The **Add Domains to DNS Server Group** dialog box appears.

c) In the **DNS server group to domain name mapping** drop-down list, choose the DNS server group name.
d) In the **Domain Name** field, enter the domain name that you want to map to the DNS group.
e) Click **OK**.
f) Repeat these steps to add more mappings.

Step 9 Click **Apply** to save your changes.

Configure the Hardware Bypass and Dual Power Supply (Cisco ISA 3000)

You can enable the hardware bypass so that traffic continues to flow between an interface pair during a power outage. Supported interface pairs are copper GigabitEthernet 1/1 & 1/2; and GigabitEthernet 1/3 & 1/4. When the hardware bypass is active, no firewall functions are in place, so make sure you understand the risks of allowing traffic through. See the following hardware bypass guidelines:

- This feature is only available on the Cisco ISA 3000 appliance.
- If you have a fiber Ethernet model, only the copper Ethernet pair (GigabitEthernet 1/1 & 1/2) supports hardware bypass.
- When the ISA 3000 loses power and goes into hardware bypass mode, only the supported interface pairs can communicate; when using the default configuration, inside1 <---> inside2, and outside1 <---> outside2 can no longer communicate. Any existing connections between these interfaces will be lost.
- We suggest that you disable TCP sequence randomization (as described in this procedure). If randomization is enabled (the default), then when the hardware bypass is activated, TCP sessions will need to be re-established. By default, the ISA 3000 rewrites the initial sequence number (ISN) of TCP connections passing through it to a random number. When the 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. Even with TCP sequence randomization disabled, some TCP connections will have to be re-established because of the link that is temporarily down during the switchover.
- Cisco TrustSec connections on hardware bypass interfaces are dropped when hardware bypass is activated. When the ISA 3000 powers on and hardware bypass is deactivated, the connections are renegotiated.
- When the hardware bypass is deactivated, and traffic resumes going through the ISA 3000 data path, some existing TCP sessions need to be re-established because of the link that is temporarily down during the switchover.
- When hardware bypass is active, the Ethernet PHYs are disconnected, so the ASA is unable to determine the interface status. Interfaces may appear to be in a down state.

For dual power supplies in the ISA 3000, you can establish dual power supplies as the expected configuration in the ASA OS. If one power supply fails, the ASA issues an alarm. By default, the ASA expects a single power supply and won't issue an alarm as long as it includes one working power supply.

Before you begin

- You must attach the hardware bypass interfaces to access ports on the switch. Do not attach them to trunk ports.

Procedure

Step 1 To configure hardware bypass, choose **Configuration** > **Device Management** > **Hardware Bypass**.

Step 2 Configure the hardware bypass to activate for each interface pair by checking the **Enable Bypass during Power Down** check box.

Step 3 (Optional) Configure each interface pair to remain in hardware bypass mode after the power comes back and the appliance boots up by checking the **Stay in Bypass after Power Up** check box.

When the hardware bypass is deactivated, there is a brief connection interruption as the ASA takes over the flows. In this case, you need to manually turn off the hardware bypass when you are ready; this option lets you control when the brief interruption occurs.

Step 4 For an interface pair, manually activate or deactivate the hardware bypass by checking the **Bypass Immediately** check box.

Step 5 (Optional) Configure the hardware bypass to remain active until after the ASA FirePOWER module boots up by checking the **Stay in Bypass Mode until after the ASA Firepower Module Boots Up** check box.

You must enable hardware bypass without the **Stay in Bypass after Power Up** option for the boot delay to operate. Without this option, the hardware bypass is likely to become inactive before the ASA FirePOWER module finishes booting up. This scenario can cause traffic to be dropped if you configured the module to fail-close, for example.

Step 6 Click **Apply**.

Step 7 Disable TCP randomization. This example shows how to disable randomization for all traffic by adding the setting to the default configuration.

a) Choose **Configuration** > **Firewall** > **Service Policy**.
b) Select the **sfrclass** rule, and click **Edit**.
c) Click **Rule Actions**, and then click **Connection Settings**.
d) Uncheck the **Randomize Sequence Number** check box.
e) Click **OK**, and then **Apply**.

Step 8 To establish dual power supplies as the expected configuration, choose **Configuration** > **Device Management** > **Power Supply**, check the **Enable Redundant Power Supply** check box, and click **Apply**.

This screen also shows the available power supplies.

Step 9 Click **Save**.

The behavior of hardware bypass after the system comes online is determined by the configuration setting in the startup configuration, so you must save your running configuration.

Adjust ASP (Accelerated Security Path) Performance and Behavior

The ASP is an implementation layer that puts your policies and configurations into action. It is not of direct interest except during troubleshooting with the Cisco Technical Assistance Center. However, there are a few behaviors related to performance and reliability that you can adjust.

Choose a Rule Engine Transactional Commit Model

By default, when you change a rule-based policy (such as access rules), the changes become effective immediately. However, this immediacy comes with a slight cost in performance. The performance cost is more noticeable for very large rule lists in a high connections-per-second environment, for example, when you change a policy with 25,000 rules while the ASA is handling 18,000 connections per second.

The performance is affected because the rule engine compiles rules to enable faster rule lookup. By default, the system also searches uncompiled rules when evaluating a connection attempt so that new rules can be applied; because the rules are not compiled, the search takes longer.

You can change this behavior so that the rule engine uses a transactional model when implementing rule changes, continuing to use the old rules until the new rules are compiled and ready for use. With the transactional model, performance should not drop during the rule compilation. The following table clarifies the behavioral difference.

Model	Before Compilation	During Compilation	After Compilation
Default	Matches old rules.	Match new rules. (The rate for connections per second decreases.)	Matches new rules.
Transactional	Matches old rules.	Match old rules. (The rate for connections per second is unaffected.)	Matches new rules.

An additional benefit of the transactional model is that, when replacing an ACL on an interface, there is no gap between deleting the old ACL and applying the new one. This feature reduces the chances that acceptable connections may be dropped during the operation.

Before you begin

- Transactional commit is not recommended for access control rules if you use FQDN objects for hostnames whose resolution can frequently change, as the access group compilation might never completely resolve due to DNS churn. If you still want to use transactional commit, consider lengthening the time for DNS expiration.
- If you enable the transactional model for a rule type, syslogs to mark the beginning and the end of the compilation are generated. These syslogs are numbered 780001 through 780004.

Procedure

Choose **Configuration > Device Management > Advanced > Rule Engine** and select the desired options:

- **Access group**—Access rules applied globally or to interfaces.
- **NAT**—Network address translation rules.

Enable ASP Load Balancing

The ASP load balancing mechanism helps avoid the following issues:

- Overruns caused by sporadic traffic spikes on flows
- Overruns caused by bulk flows oversubscribing specific interface receive rings
- Overruns caused by relatively heavily overloaded interface receive rings, in which a single core cannot sustain the load.

ASP load balancing allows multiple cores to work simultaneously on packets that were received from a single interface receive ring. If the system drops packets, and the **show cpu** command output is far less than 100%, then this feature may help your throughput if the packets belong to many unrelated connections.



Note ASP load balancing is disabled on the ASA virtual. With the integration of DPDK (Dataplane Development Kit) into the ASA virtual's accelerated security path (ASP), the ASA virtual shows better performance with this feature disabled.

Procedure

Step 1 To enable the automatic switching on and off of ASP load balancing, choose **Configuration** > **Device Management** > **Advanced** > **ASP Load Balancing**, and check the **Dynamically enable or disable ASP load balancing based on traffic monitoring** check box .

Step 2 To manually enable or disable ASP load balancing, check or uncheck the **Enable ASP load balancing** check box.

When you manually enable ASP load balancing, it is enabled until you manually disable it, even if you also have the Dynamic option enabled. Manually disabling ASP load balancing applies only if you manually enabled ASP load blancing. If you also enabled the Dynamic option, then the system reverts to automatically enabling or disabling ASP load balancing.

Monitoring the DNS Cache

The ASA provides a local cache of DNS information from external DNS queries that are sent for certain clientless SSL VPN and certificate commands. Each DNS translation request is first looked for in the local cache. If the local cache has the information, the resulting IP address is returned. If the local cache can not resolve the request, a DNS query is sent to the various DNS servers that have been configured. If an external DNS server resolves the request, the resulting IP address is stored in the local cache with its corresponding hostname.

See the following command for monitoring the DNS cache:

- **show dns-hosts**

This command shows the DNS cache, which includes dynamically learned entries from a DNS server as well as manually entered name and IP addresses using the name command.

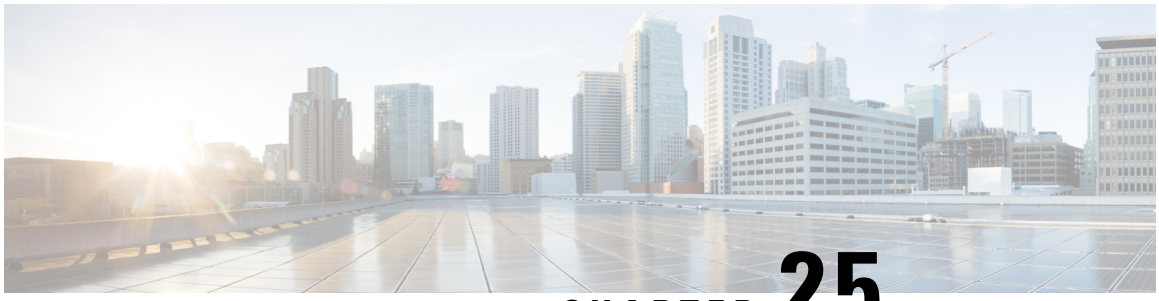
History for Basic Settings

Feature Name	Platform Releases	Description
Multiple DNS server groups	9.18(1)	You can now use multiple DNS server groups: one group is the default, while other groups can be associated with specific domains. A DNS request that matches a domain associated with a DNS server group will use that group. For example, if you want traffic destined to inside eng.cisco.com servers to use an inside DNS server, you can map eng.cisco.com to an inside DNS group. All DNS requests that do not match a domain mapping will use the default DNS server group, which has no associated domains. For example, the DefaultDNS group can include a public DNS server available on the outside interface. New/Modified screens: **Configuration** > **Device Management** > **DNS** > **DNS Client**
Trusted DNS servers for network-service object domain resolution.	9.17(1)	You can specify which DNS servers the system should trust when resolving domain names in network-service objects. This feature ensures that any DNS domain name resolutions acquire IP addresses from trusted sources. New/Modified screens: **Configuration** > **Device Management** > **DNS** > **DNS Client**
Change in DNS entry TTL behavior	9.17(1)	Formerly, the configured value was added to the existing TTL of each entry (the default was 1 minute). Now, 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. New/Modified screens: **Configuration** > **Device Management** > **DNS** > **DNS Client** > **Configure multiple DNS server groups**
Stronger local user and enable password requirements	9.17(1)	For local users and the enable password, the following password requirements were added: • Password length—Minimum 8 characters. Formerly, the minimum was 3 characters. • Repetitive and sequential characters—Three or more consecutive sequential or repetitive ASCII characters are disallowed. For example, the following passwords will be rejected: • **abc**user1 • user**543** • user**aaaa** • user2**666** New/Modified screens: • **Configuration** > **Device Management** > **Users/AAA** > **User Accounts** • **Configuration** > **Device Setup** > **Device Name/Password**

Feature Name	Platform Releases	Description
NTPv4 support	9.14(1)	The ASA now supports NTPv4. No modified screens.
Additional NTP authentication algorithms	9.13(1)	Formerly, only MD5 was supported for NTP authentication. The ASA now supports the following algorithms: • MD5 • SHA-1 • SHA-256 • SHA-512 • AES-CMAC New/Modified screens: **Configuration** > **Device Setup** > **System Time** > **NTP** > **Add** button **> Add NTP Server Configuration** dialog box **> Key Algorithm** drop-down list
NTP support on IPv6	9.12(1)	You can now specify an IPv6 address for the NTP server. New/Modified screens: **Configuration** > **Device Setup** > **System Time** > **NTP** > **Add** button **> Add NTP Server Configuration** dialog box
enable password change now required on login	9.12(1)	The default **enable** password is blank. When you try to access privileged EXEC mode on the ASA, you are now required to change the password to a value of 3 to 127 characters. You cannot keep it blank. The **no enable password** command is no longer supported. At the CLI, you can access privileged EXEC mode using the **enable** command, the **login** command (with a user at privilege level 2+), or an SSH or Telnet session when you enable **aaa authorization exec auto-enable**. All of these methods require you to set the enable password. This password change requirement is not enforced for ASDM logins. In ASDM, by default you can log in without a username and with the **enable** password. No modified screens.
ASP load balancing is disabled on the ASA virtual	9.10(1)	With the recent integration of DPDK (Dataplane Development Kit) into the ASA virtual's accelerated security path (ASP), the ASA virtual shows better performance with this feature disabled.
Automatic ASP load balancing now supported for the ASA virtual	9.8(1)	Formerly, you could only manually enable and disable ASP load balancing. We modified the following screen: **Configuration > Device Management > Advanced > ASP Load Balancing**

Feature Name	Platform Releases	Description
PBKDF2 hashing for all local **username** and **enable** passwords	9.7(1)	Local **username** and **enable** passwords of all lengths are stored in the configuration using a PBKDF2 (Password-Based Key Derivation Function 2) hash using SHA-512. Previously, passwords 32 characters and shorter used the MD5-based hashing method. Already existing passwords continue to use the MD5-based hash unless you enter a new password. See the "Software and Configurations" chapter in the General Operations Configuration Guide for downgrading guidelines. We modified the following screens: **Configuration** > **Device Setup** > **Device Name/Password** > **Enable Password** **Configuration** > **Device Management** > **Users/AAA** > **User Accounts** > **Add/Edit User Account** > **Identity**
Dual power supply support for the ISA 3000	9.6(1)	For dual power supplies in the ISA 3000, you can establish dual power supplies as the expected configuration in the ASA OS. If one power supply fails, the ASA issues an alarm. By default, the ASA expects a single power supply and won't issue an alarm as long as it includes one working power supply. We introduced the following screen: **Configuration** > **Device Management** > **Power Supply**
Longer password support for local **username** and **enable** passwords (up to 127 characters)	9.6(1)	You can now create local **username** and **enable** passwords up to 127 characters (the former limit was 32). When you create a password longer than 32 characters, it is stored in the configuration using a PBKDF2 (Password-Based Key Derivation Function 2) hash. Shorter passwords continue to use the MD5-based hashing method. We modified the following screens: **Configuration** > **Device Setup** > **Device Name/Password** > **Enable Password** **Configuration** > **Device Management** > **Users/AAA** > **User Accounts** > **Add/Edit User Account** > **Identity**
ISA 3000 hardware bypass	9.4(1.225)	The ISA 3000 supports a hardware bypass function to allow traffic to continue flowing through the appliance when there is a loss of power. We introduced the following screen: **Configuration** > **Device Management** > **Hardware Bypass** *This feature is not available in Version 9.5(1).*
Automatic ASP Load Balancing	9.3(2)	You can now enable automatic switching on and off of the ASP load balancing feature. **Note** The automatic feature is not supported on the ASA virtual; only manual enabling and disabling is supported. We modified the following screen: **Configuration > Device Management > Advanced > ASP Load Balancing**

Feature Name	Platform Releases	Description
Removal of the default Telnet password	9.0(2), 9.1(2)	To improve security for management access to the ASA, the default login password for Telnet was removed; you must manually set the password before you can log in using Telnet. **Note** The login password is only used for Telnet if you do not configure Telnet user authentication. Previously, when you cleared the password, the ASA restored the default of "cisco." Now when you clear the password, the password is removed. The login password is also used for Telnet sessions from the switch to the ASASM (see the **session** command). For initial ASASM access, you must use the **service-module session** command, until you set a login password. We did not modify any ASDM screens.
Password Encryption Visibility	8.4(1)	We modified the **show password encryption** command.
Master Passphrase	8.3(1)	We introduced this feature. The master passphrase allows you to securely store plain text passwords in encrypted format and provides a key that is used to universally encrypt or mask all passwords, without changing any functionality. We introduced the following screens: **Configuration > Device Management > Advanced > Master Passphrase** **Configuration > Device Management > Device Administration > Master Passphrase**



CHAPTER 25

DHCP and DDNS Services

This chapter describes how to configure the DHCP server or DHCP relay as well as dynamic DNS (DDNS) update methods.

- About DHCP and DDNS Services, on page 713
- Guidelines for DHCP and DDNS Services, on page 715
- Configure the DHCP Server, on page 717
- Configure the DHCP Relay Agent, on page 721
- Configure Dynamic DNS, on page 722
- Monitoring DHCP and DDNS Services, on page 726
- History for DHCP and DDNS Services, on page 728

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 ASA can provide a DHCP server to DHCP clients attached to ASA 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 ASA.

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 648), you can configure the ASA to provide information such as the DNS server or domain name when they send Information Request (IR) packets to the ASA. The ASA 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 ASA 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 ASA because it does not forward broadcast traffic. The DHCP relay agent lets you configure the interface of the ASA that is receiving the broadcasts to forward DHCP requests to a DHCP server on another device.

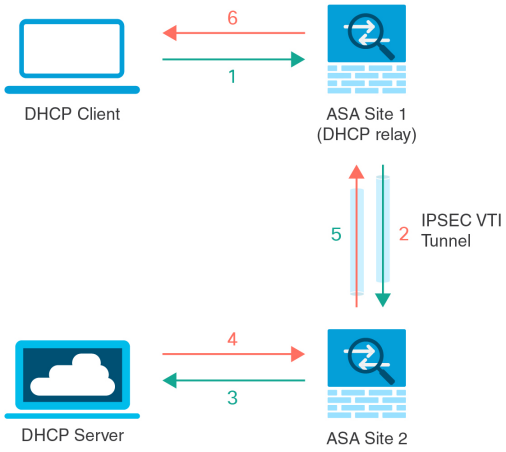
DHCP Relay Server Support on VTI

You can configure DHCP relay agent on an ASA interface to receive and forward DHCP messages between a DHCP client and a DHCP server. However, a DHCP relay server to forward messages through a logical interface was not supported.

Following figure shows the DISCOVER process of the DHCP Client and DHCP Server using DHCP relay over VTI VPN. The DHCP relay agent, configured on VTI interface of ASA Site 1, receives DHCPDISCOVER packet from the DHCP Client and sends the packet through the VTI tunnel. ASA Site 2 forwards the

DHCPDISCOVER packet to the DHCP Server. The DHCP Server replies with a DHCPOFFER to ASA Site 2. ASA Site 2 forwards it to DHCP relay (ASA Site1), which forwards it to the DHCP Client.

Figure 87: DHCP Relay Server over VTI



The same procedure is followed for a DHCPREQUEST and DHCPACK/NACK requirements.

Guidelines for DHCP and DDNS Services

This section includes guidelines and limitations that you should check before configuring DHCP and DDNS services.

Context Mode

- DHCPv6 stateless server is not supported in multiple context mode.

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.
- You can reserve a DHCP address for an interface. The ASA assigns a specific address from the address pool to a DHCP client based on the client's MAC address.
- The ASA does not support QIP DHCP servers for use with the DHCP proxy service.
- 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, with a maximum of 4 servers per interface.
- You can configure a maximum of 10 DHCPv6 relay servers. 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 ASA, 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 ASA and cannot send requests through another relay agent or a router. For IPv6, the ASA supports packets from another relay server.
- The DHCP clients must be on different interfaces from the DHCP servers to which the ASA 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 DHCP Server

This section describes how to configure a DHCP server provided by the ASA.

Procedure

Step 1 Enable the DHCPv4 Server, on page 717.
Step 2 Configure Advanced DHCPv4 Options, on page 719.
Step 3 Configure the DHCPv6 Stateless Server, on page 720.

Enable the DHCPv4 Server

To enable the DHCP server on an ASA interface, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Device Management** > **DHCP** > **DHCP Server**.

Step 2 Choose an interface, then click **Edit**.

In transparent mode, choose a bridge group member interface. In routed mode, choose a routed interface or a BVI; do not choose the bridge group member interface.

a) Check the **Enable DHCP Server** check box to enable the DHCP server on the selected interface.
b) Enter the range of IP addresses from lowest to highest that is used by the DHCP server in the **DHCP Address Pool** field. 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.
c) Set the following in the **Optional Parameters** area:

- The DNS servers (1 and 2) configured for the interface.
- The WINS servers (primary and secondary) configured for the interface.
- The domain name of the interface.
- The time in milliseconds that the ASA will wait for an ICMP ping response on the interface.
- The duration of time that the DHCP server configured on the interface allows DHCP clients to use an assigned IP address.

- The interface on a DHCP client that provides DNS, WINS, and domain name information for automatic configuration if the ASA is acting as a DHCP client on a specified interface (usually outside).
- Click **Advanced** to display the **Advanced DHCP Options** dialog box to configure more DHCP options. See Configure Advanced DHCPv4 Options, on page 719 for more information.

d) Check the **Update DNS Clients** check box in the **Dynamic Settings for DHCP Server** area to specify that in addition to the default action of updating the client PTR resource records, the selected DHCP server should also perform the following update actions:

- Check the **Update Both Records** check box to specify that the DHCP server should update both the A and PTR RRs.
- Check the **Override Client Settings** check box to specify that DHCP server actions should override any update actions requested by the DHCP client.

e) Click **OK** to close the **Edit DHCP Server** dialog box.

Step 3 (Optional) (Routed mode) Check the **Enable Auto-configuration from interface** check box in the **Global DHCP Options** area below the DHCP Server table to enable DHCP auto configuration only if the ASA is acting as a DHCP client on a specified interface (usually outside).

DHCP auto configuration enables the DHCP Server to provide DHCP clients with DNS server, domain name, and WINS server information obtained from a DHCP client that is running on the specified interface. If information obtained through auto configuration is also specified manually in the **Global DHCP Options** area, the manually specified information takes precedence over the discovered information.

Step 4 Choose the auto-configuration interface from the drop-down list.

Step 5 Check the **Allow VPN override** check box to override the interface DHCP or PPPoE client WINS parameter with the VPN client parameter.

Step 6 Enter the IP address of the primary DNS server for a DHCP client in the **DNS Server 1** field.

Step 7 Enter the IP address of the alternate DNS server for a DHCP client in the **DNS Server 2** field.

Step 8 Enter the DNS domain name for DHCP clients (for example, example.com) in the **Domain Name** field.

Step 9 Enter the amount of time, in seconds, in the **Lease Length** field 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).

Step 10 Enter the IP address of the primary WINS server for a DHCP client in the **Primary WINS Server** field.

Step 11 Enter the IP address of the alternate WINS server for a DHCP client in the **Secondary WINS Server** field.

Step 12 To avoid address conflicts, the ASA sends two ICMP ping packets to an address before assigning that address to a DHCP client. enter the amount of time, in milliseconds, in the **Ping Timeout** field that the ASA waits to time out a DHCP ping attempt. Valid values range from 10 to 10000 milliseconds. The default value is 50 milliseconds.

Step 13 Click **Advanced** to display the **Configuring Advanced DHCP Options** dialog box to specify additional DHCP options and their parameters. For more information, see Configure Advanced DHCPv4 Options, on page 719.

Step 14 You configure the DDNS update settings for the DHCP server in the **Dynamic DNS Settings for DHCP Server** area. Check the **Update DNS Clients** check box to specify that, in addition to the default action of updating the client PTR resource records, the selected DHCP server should also perform the following update actions:

- Check the **Update Both Records** check box to specify that the DHCP server should update both the A and PTR RRs.
- Check the **Override Client Settings** check box to specify that the DHCP server actions should override any update actions requested by the DHCP client.

Step 15 Click **Apply** to save your changes.

Configure Advanced DHCPv4 Options

The ASA 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.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **DHCP** > **DHCP Server**, then click **Advanced**.

Step 2 Choose the option code from the drop-down list.

Step 3 Choose the options that you want to configure. Some options are standard. For standard options, the option name is shown in parentheses after the option number and the option parameters are limited to those supported by the option. For all other options, only the option number is shown and you must choose the appropriate parameters to supply with the option. For example, if you choose DHCP Option 2 (Time Offset), you can only enter a hexadecimal value for the option. For all other DHCP options, all of the option value types are available and you must choose the appropriate one.

Step 4 Specify the type of information that the option returns to the DHCP client in the **Option Data** area. For standard DHCP options, only the supported option value type is available. For all other DHCP options, all of the option value types are available. Click **Add** to add the option to the DHCP option list. Click **Delete** to remove the option from the DHCP option list.

- Click **IP Address** to indicate that an IP address is returned to the DHCP client. You can specify up to two IP addresses. IP Address 1 and IP Address 2 indicate an IP address in dotted-decimal notation.

  Note
  The name of the associated IP address fields can change based on the DHCP option that you chose. For example, if you choose DHCP Option 3 (Router), the fields names change to Router 1 and Router 2.

- Click **ASCII** to specify that an ASCII value is returned to the DHCP client. Enter an ASCII character string in the **Data** field. The string cannot include spaces.

  Note
  The name of the associated Data field can change based on the DHCP option that you chose. For example, if you choose DHCP Option 14 (Merit Dump File), the associated Data field names change to File Name.

- Click **Hex** to specify that a hexadecimal value is returned to the DHCP client. Enter a hexadecimal string with an even number of digits and no spaces in the **Data** field. You do not need to use a 0x prefix.

  Note
  The name of the associated Data field can change based on the DHCP option you chose. For example, if you choose DHCP Option 2 (Time Offset), the associated Data field becomes the Offset field.

Step 5 Click **OK** to close the **Advanced DHCP Options** dialog box.

Step 6 Click **Apply** to save your changes.

Configure 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 648), you can configure the ASA to provide information such as the DNS server or domain name when they send Information Request (IR) packets to the ASA. The ASA 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 ASA received using Prefix Delegation.

Before you begin

This feature is only supported in single, routed mode. This feature is not supported in clustering.

Procedure

Step 1 Configure the IPv6 DHCP pool that contains the information you want the DHCPv6 server to provide:

a) Choose **Configuration** > **Device Management** > **DHCP** > **DHCP Pool**, and click **Add**.

b) In the **DHCP Pool Name** field, enter a name.

c) For each parameter on each tab, either check the **Import** check box or manually enter a value in the field and click **Add**.

The **Import** option uses one or more parameters that the ASA 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 specify **Import**.

d) Click **OK**, and then **Apply**.

Step 2 Choose **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces**.

Step 3 Choose an interface, and click **Edit**.

The **Edit Interface** dialog box appears with the **General** tab selected.

Step 4 Click the **IPv6** tab.

Step 5 In the **Interface IPv6 DHCP** area, click the **Server DHCP Pool Name** radio button, and enter the IPv6 DHCP pool name.

Step 6 Check the **Hosts should use DHCP for non-address config** check box to set the Other Address Config flag in the IPv6 router advertisement packet.

This flag informs IPv6 autoconfiguration clients that they should use DHCPv6 to obtain additional information from DHCPv6, such as the DNS server address.

Step 7 Click **OK**.

You return to the **Configuration > Device Setup > Interface Settings > Interfaces** pane.

Step 8 Click **Apply**.

Configure the DHCP Relay Agent

When a DHCP request enters an interface, the DHCP servers to which the ASA relays the request depends on your configuration. You may configure the following types of servers:

- Interface-specific DHCP servers—When a DHCP request enters a particular interface, then the ASA relays the request only to the interface-specific servers.
- Global DHCP servers—When a DHCP request enters an interface that does not have interface-specific servers configured, the ASA relays the request to all global servers. If the interface has interface-specific servers, then the global servers are not used.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **DHCP** > **DHCP Relay**.

Step 2 Check the check boxes for the services you want for each interface in the **DHCP Relay Agent** area.

- **IPv4 > DHCP Relay Enabled**.
- **IPv4 > Set Route**— Changes the default gateway address in the DHCP message from the server to that of the ASA 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 ASA even if the DHCP server specifies a different router. If there is no default router option in the packet, the ASA adds one containing the interface address.
- **IPv6 > DHCP Relay Enabled**.
- **Trusted Interface**—Specifies a DHCP client interface that you want to trust. 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 ASA 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 ASA will drop that packet by default. You can now preserve Option 82 and forward the packet by identifying an interface as a trusted interface. You can alternatively trust all interfaces by checking the **Set dhcp relay information as trusted on all interfaces** check box.

Step 3 Add one or more DHCP servers to which DHCP requests are relayed in the **Global DHCP Relay Servers** area,

a) Click **Add**. The **Add Global DHCP Relay Server** dialog box appears.
b) Enter the IPv4 or IPv6 address of the DHCP server in the **DHCP Server** field.
c) Choose the interface to which the specified DHCP server is attached from the **Interface** drop-down list.
d) Click **OK**.

The newly added global DHCP relay server appears in the **Global DHCP Relay Servers** list.

Step 4 (Optional) In the **IPv4 Timeout** field, enter the amount of time, in seconds, allowed for DHCPv4 address handling. Valid values range from 1 to 3600 seconds. The default value is 60 seconds.

Step 5 (Optional) In the **IPv6 Timeout** field, enter the amount of time, in seconds, allowed for DHCPv6 address handling. Valid values range from 1 to 3600 seconds. The default value is 60 seconds.

Step 6 In the **DHCP Relay Interface Servers** area, add one or more interface-specific DHCP servers to which DHCP requests on a given interface are relayed:

a) Click **Add**. The **Add DHCP Relay Server** dialog box appears.

b) From the **Interface** drop-down list, choose the interface connected to the DHCP clients. Note that you do not specify the egress interface for the requests, as for a Global DHCP Server; instead, the ASA uses the routing table to determine the egress interface.

c) In the **Server to** field, enter the IPv4 address of the DHCP server, and click **Add**. The server is added to the right-hand list. Add up to 4 servers, if available out of the overall maximum. IPv6 is not supported for interface-specific servers.

d) Click **OK**.

The newly added interface DHCP relay servers appear in the **DHCP Relay Interface Servers** list.

Step 7 To configure all interfaces as trusted interfaces, check the **Set dhcp relay information as trusted on all interfaces** check box. You can alternatively trust individual interfaces.

Step 8 Click **Apply** to save your settings.

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 ASA supports the following DDNS update methods:

- Standard DDNS—The standard DDNS update method is defined by RFC 2136.

  With this method, the ASA and the DHCP server use DNS requests to update the DNS RRs. The ASA 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 ASA or DHCP server then sends an update request directly to the main DNS server. See the following typical scenarios.

  - The ASA updates the A RR, and the DHCP server updates the PTR RR.

    Typically, the ASA "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 ASA 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 ASA does not have the authority to update the A RR. When the IP address or hostname changes, the ASA 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 ASA 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 ASA sends an HTTP request directly to a DNS provider with which you have an account.

Note DDNS is not supported on the BVI or bridge group member interfaces.

Before you begin

- Configure a DNS server on **Configuration** > **Device Management** > **DNS** > **DNS Client**. See Configure the DNS Servers, on page 701.
- Configure the device hostname and domain name on **Configuration** > **Device Setup** > **Device Name/Password**. See Set the Hostname, Domain Name, and the Enable and Telnet Passwords, on page 693. If you do not specify the hostname per interface, then the device hostname is used. If you do not specify an FQDN, then for static or PPPoE IP addressing, the system domain name or the DNS server domain name is appended to the hostname.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **DNS** > **Dynamic DNS**.

Step 2 Standard DDNS method: Configure a DDNS update method to enable DNS requests from the ASA.

You do not need to configure a DDNS update method if the DHCP server will perform all requests.

a) In the **Update Methods** area, click **Add**.
b) Specify a **Name** for this method.
c) (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**.
d) Choose **DDNS Record Type** > **Standard DDNS**.
e) Under **Records to Update**, specify the standard DDNS records that you want the ASA to update.

   This setting only affects the records you want to update directly from the ASA; to determine the records you want the DHCP server to update, configure the DHCP client settings per interface or globally. See Step Step 4, on page 724.

   - **Both (PTR and A records)**—Sets the ASA to update both A and PTR RRs. Use this option for static or PPPoE IP addressing.
   - **A records only**—Sets the ASA 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 Step 4, on page 724.

Step 3 Web method: Configure a DDNS update method to enable HTTP update requests from the ASA.

a) In the **Update Methods** area, click **Add**.
b) Specify a **Name** for this method.
c) (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**.
d) Choose **DDNS Record Type** > **Web**.
e) In the **Web** field, 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) For the **Web Update Type**, specify the address types (IPv4 or IPv6) that you want to update.

- **Both All**—(Default) Updates all IPv4 and IPv6 addresses.
- **Both**—Updates the IPv4 address and the latest IPv6 address.
- **IPv4**—Updates only the IPv4 address.
- **IPv6**—Updates only the latest IPv6 address.
- **IPv6 All**—Updates all IPv6 addresses.

g) In **Reference Identity Name**, enter the reference identity name that is configured to validate the server certificate identity.
h) Click **OK**.
i) Assign this method to the interface in Step Step 4, on page 724.
j) 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 Step 6, on page 725.

Step 4 Configure interface settings for DDNS, including setting the update method, DHCP client settings, and the hostname for this interface.

a) In the **Dynamic DNS Interface Settings** area, click **Add**.
b) Choose the **Interface** from the drop-down list.
c) Choose the **Method Name** that you created in the **Update Methods** area.

(Standard DDNS method) You do not need to assign a method if you want the DHCP server to perform all updates.

d) Set the **Hostname** for this interface.

If you do not set the hostname, the device hostname is used. If you do not specify an FQDN, then the system domain name or 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 Server Record Updates** to determine which records you want the DHCP server to update.

The ASA 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). Even if the client does not request DDNS updates, the DHCP server can be configured to send updates anyway.

For static or PPPoE IP addressing, these settings are ignored.

Note
You can also set these values globally for all interfaces on the main **Dynamic DNS** page. The per-interface settings take precedence over the global settings.

- **Default (PTR Records)**—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 (PTR Records and A 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.
- **None**—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.

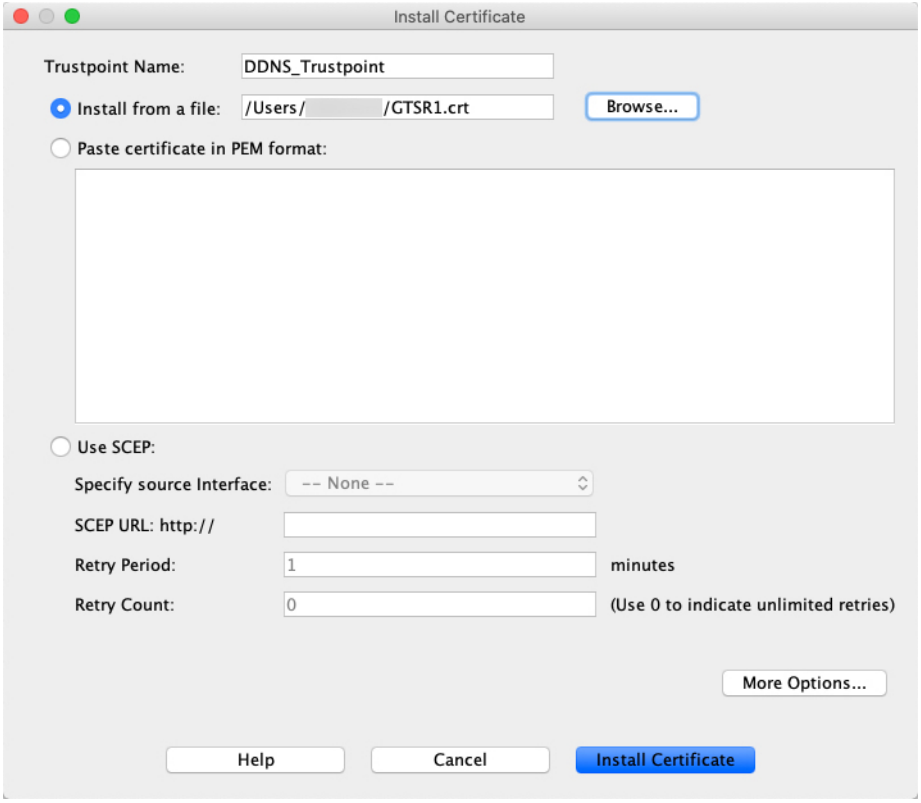
f) Click **OK**.

Step 5 Click **Apply** to save your changes, or click **Reset** to discard them and enter new ones.

Step 6 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 file import, but you can also paste it in PEM format.
b) Choose **Configuration** > **Device Management** > **Certficate Management** > **CA Certificates**, and click **Add**.



c) Enter a **Trustpoint Name**.
d) Click **Install from a file**, and browse to the certificate file.
e) Click **Install Certificate**.

Monitoring DHCP and DDNS Services

This section includes the procedures to monitor both DHCP and DDNS services.

Monitoring DHCP Services

- **Monitoring > Interfaces > DHCP > DHCP Client Lease Information**.

  This pane displays configured DHCP client IP addresses.

- **Monitoring > Interfaces > DHCP > DHCP Server Table**

  This pane displays configured dynamic DHCP client IP addresses.

- **Monitoring > Interfaces > DHCP > DHCP Statistics**

  This pane displays DHCPv4 message types, counters, values, directions, messages received, and messages sent.

- **Monitoring > Interfaces > DHCP > IPV6 DHCP Relay Statistics**

This pane displays DHCPv6 Relay message types, counters, values, directions, messages received, and messages sent.

- **Monitoring > Interfaces > DHCP > IPV6 DHCP Relay Binding**

  This pane displays DHCPv6 Relay bindings.

- **Monitoring** > **Interfaces** > **DHCP** > **IPV6 DHCP Interface Statistics**

  This screen displays DHCPv6 information for all interfaces. If the interface is configured for DHCPv6 stateless server configuration (see Configure the DHCPv6 Stateless Server, on page 720), this screen lists the DHCPv6 pool that is being used by the server. If the interface has DHCPv6 address client or Prefix Delegation client configuration, this screen shows the state of each client and the values received from the server. This screen also shows message statistics for the DHCP server or client.

- **Monitoring** > **Interfaces** > **DHCP** > **IPV6 DHCP HA Statistics**

  This screen shows the transaction statistics between failover units, including how many times the DUID information was synced between the units.

- **Monitoring** > **Interfaces** > **DHCP** > **IPV6 DHCP Server Statistics**

  This screen shows the DHCPv6 stateless server statistics.

Monitoring DDNS Status

See the following command for monitoring DDNS status. Enter the commands on **Tools > Command Line Interface**.

- **show ddns update** {**interface** *if_name* | **method** [*name*]}

  This command shows the DDNS update status.

  The following example show details about the DDNS update method:

```
ciscoasa# show ddns update method ddns1

Dynamic DNS Update Method: ddns1
  IETF standardized Dynamic DNS 'A' record update
```

  The following example shows details about the web update method:

```
ciscoasa# show ddns update method web1

Dynamic DNS Update Method: web1
 Dynamic DNS updated via HTTP(s) protocols
  URL used to update record:
https://cdarwin:*****@ddns.cisco.com/update?hostname=<h>&myip=<a>
```

  The following example shows information about the DDNS interface:

```
ciscoasa# show ddns update interface outside

Dynamic DNS Update on outside:
  Update Method Name            Update Destination
  test                          not available
```

  The following example shows a successful web type update:

```
ciscoasa# show ddns update interface outside

Dynamic DNS Update on outside:
  Update Method Name            Update Destination
  test                          not available

Last Update attempted on 09:01:52.729 UTC Mon Mar 23 2020
Status : Success
FQDN : asa1.example.com
IP addresses(s): 10.10.32.45,2001:DB8::1
```

The following example shows a web type failure:

```
ciscoasa# show ddns update interface outside

Dynamic DNS Update on outside:
  Update Method Name            Update Destination
  test                          not available

Last Update attempted on 09:01:52.729 UTC Mon Mar 23 2020
Status : Failed
Reason : Could not establish a connection to the server
```

The following example shows that the DNS server returned an error for the web type update:

```
ciscoasa# show ddns update interface outside

Dynamic DNS Update on outside:
  Update Method Name            Update Destination
  test                          not available

Last Update attempted on 09:01:52.729 UTC Mon Mar 23 2020
Status : Failed
Reason : Server error (Error response from server)
```

The following example shows that a web update was not yet attempted due to the IP address unconfigured or the DHCP request failed, for example:

```
ciscoasa# show ddns update interface outside

Dynamic DNS Update on outside:
  Update Method Name            Update Destination
  test                          not available

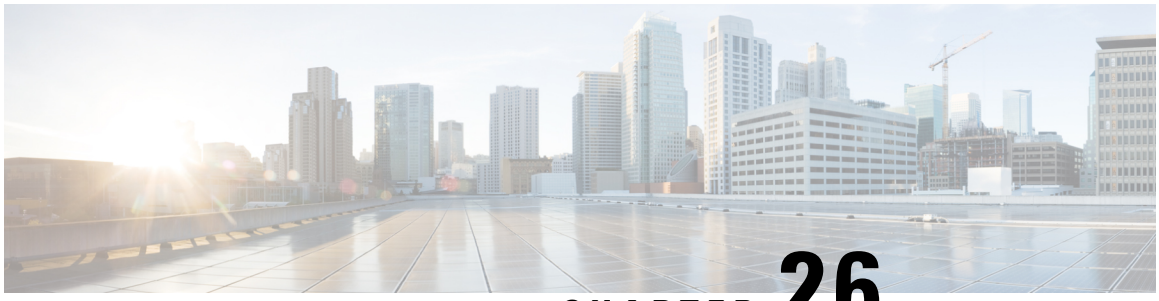
Last Update Not attempted
```

History for DHCP and DDNS Services

Feature Name	Platform Releases	Description
DDNS support for the web update method	9.15(1)	You can now configure an interface to use DDNS with the web update method. New/Modified screens: **Configuration** > **Device Management** > **DNS** > **Dynamic DNS**

Feature Name	Platform Releases	Description
DHCP relay server support on VTIs	9.14(1)	You can now enable DHCP relay on VTIs. New/Modified screens: **Configuration** > **Device Management** > **DHCP** > **DHCP Relay** > **DHCP Relay Interface Servers**
DHCP reservation	9.13(1)	ASA supports DHCP reservation. The DHCP server assigns a static IP address from the defined address pool to a DHCP client based on the client's MAC address. New/Modified commands: **dhcpd reserve-address** . No ASDM support.
IPv6 DHCP	9.6(2)	The ASA now supports the following features for IPv6 addressing: • DHCPv6 Address client—The ASA obtains an IPv6 global address and optional default route from the DHCPv6 server. • DHCPv6 Prefix Delegation client—The ASA obtains delegated prefix(es) from a DHCPv6 server. The ASA can then use these prefixes to configure other ASA interface addresess 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 ASA provides other information such as the domain name to SLAAC clients when they send Information Request (IR) packets to the ASA. The ASA only accepts IR packets, and does not assign addresses to the clients. We added or modified the following screens: **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces** > **Add Interface** > **IPv6** **Configuration** > **Device Management** > **DHCP** > **DHCP Pool** **Configuration** > **Device Setup** > **Routing** > **BGP** > **IPv6 Family** > **Networks** **Monitoring** > **interfaces** > **DHCP**
DHCPv6 monitoring	9.4(1)	You can now monitor DHCP statistics for IPv6 and DHCP bindings for IPv6. We introduced the following screens: DHCPv6 monitoring Monitoring > Interfaces > DHCP > IPV6 DHCP Statistics, Monitoring > Interfaces > DHCP > IPV6 DHCP Binding.
DHCP Relay server validates the DHCP Server identifier for replies	9.2(4)/ 9.3(3)	If the ASA DHCP relay server receives a reply from an incorrect DHCP server, it now verifies that the reply is from the correct server before acting on the reply. We did not introduce or modify any commands. We did not modify any ASDM screens. We did not modify any ASDM screens.
DHCP rebind function	9.1(4)	During the DHCP rebind phase, the client now tries to rebind to other DHCP servers in the tunnel group list. Before this release, the client did not rebind to an alternate server when the DHCP lease fails to renew. We did not modify any ASDM screens.

Feature Name	Platform Releases	Description
DHCP trusted interfaces	9.1(2)	You can now 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 ASA 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 ASA will drop that packet by default. You can now preserve Option 82 and forward the packet by identifying an interface as a trusted interface. We modified the following screen: Configuration > Device Management > DHCP > DHCP Relay.
DHCP relay servers per interface (IPv4 only)	9.1(2)	You can now configure DHCP relay servers per-interface, so requests that enter a given interface are relayed only to servers specified for that interface. IPv6 is not supported for per-interface DHCP relay. We modified the following screen: Configuration > Device Management > DHCP > DHCP Relay.
DHCP relay for IPv6 (DHCPv6)	9.0(1)	DHCP relay support for IPv6 was added. We modified the following screen: Configuration > Device Management > DHCP > DHCP Relay.
DDNS	7.0(1)	We introduced this feature. We introduced the following screens: Configuration > Device Management> DNS > DNS Client. Configuration > Device Management > DNS > Dynamic DNS.
DHCP	7.0(1)	The ASA can provide a DHCP server or DHCP relay services to DHCP clients attached to ASA interfaces. We introduced the following screens: Configuration > Device Management > DHCP > DHCP Relay. Configuration > Device Management > DHCP > DHCP Server.



CHAPTER 26

Digital Certificates

This chapter describes how to configure digital certificates.

- About Digital Certificates, on page 731
- Guidelines for Digital Certificates, on page 738
- Configure Digital Certificates, on page 741
- How to Set Up Specific Certificate Types, on page 742
- Set a Certificate Expiration Alert (for Identity or CA Certificates), on page 758
- Monitoring Digital Certificates, on page 758
- History for Certificate Management, on page 759

About Digital Certificates

Digital certificates provide digital identification for authentication. A digital certificate includes information that identifies a device or user, such as the name, serial number, company, department, or IP address. CAs are responsible for managing certificate requests and issuing digital certificates. CAs are trusted authorities that "sign" certificates to verify their authenticity, thereby guaranteeing the identity of the device or user.

A digital certificate also includes a copy of the public key for the user or device. A CA can be a trusted third party, such as VeriSign, or a private (in-house) CA that you establish within your organization. CAs issue digital certificates in the context of a PKI, which uses public-key or private-key encryption to ensure security.

For authentication using digital certificates, at least one identity certificate and its issuing CA certificate must exist on an ASA. This configuration allows multiple identities, roots, and certificate hierarchies. The ASA evaluates third-party certificates against CRLs, also called authority revocation lists, all the way from the identity certificate up the chain of subordinate certificate authorities.

Descriptions of several different types of available digital certificates follow:

- A CA certificate is used to sign other certificates. It is self-signed and called a root certificate. A certificate that is issued by another CA certificate is called a subordinate certificate.
- CAs also issue identity certificates, which are certificates for specific systems or hosts.
- Code-signer certificates are special certificates that are used to create digital signatures to sign code, with the signed code itself revealing the certificate origin.

In a hierarchy with a root and two intermediate CA certificates, a CRL validation for remote access fail on the headend running 9.13 or later when the ID certificate is signed by one intermediate CA, but the CRL is

signed by another intermediate CA. The failure happens even if the headend trusts both the intermediates where both are signed by the same root.

Hence, each signer must maintain their own CRL. Each signer would then specify the location of the CRL in the url list of each certificate it signs. Alternatively, you can configure a url override in the trustpoint of each signer pointing to the correct CRL location.

The local CA integrates an independent certificate authority feature on the ASA, deploys certificates, and provides secure revocation checking of issued certificates. The local CA provides a secure, configurable, in-house authority for certificate authentication with user enrollment through a website login page.



Note CA certificates and identity certificates apply to both site-to-site VPN connections and remote access VPN connections. Procedures in this document refer to remote access VPN use in the ASDM GUI.



Tip For an example of a scenario that includes certificate configuration and load balancing, see the following URL: https://supportforums.cisco.com/docs/DOC-5964.

Public Key Cryptography

Digital signatures, enabled by public key cryptography, provide a way to authenticate devices and users. In public key cryptography, such as the RSA encryption system, each user has a key pair containing both a public and a private key. The keys act as complements, and anything encrypted with one of the keys can be decrypted with the other.

In simple terms, a signature is formed when data is encrypted with a private key. The signature is attached to the data and sent to the receiver. The receiver applies the public key of the sender to the data. If the signature sent with the data matches the result of applying the public key to the data, the validity of the message is established.

This process relies on the receiver having a copy of the public key of the sender and a high degree of certainty that this key belongs to the sender, not to someone pretending to be the sender.

Obtaining the public key of a sender is normally handled externally or through an operation performed at installation. For example, most web browsers are configured with the root certificates of several CAs by default. For VPN, the IKE protocol, a component of IPsec, can use digital signatures to authenticate peer devices before setting up security associations.

Certificate Scalability

Without digital certificates, you must manually configure each IPsec peer for each peer with which it communicates; as a result, each new peer that you add to a network would require a configuration change on each peer with which it needs to communicate securely.

When you use digital certificates, each peer is enrolled with a CA. When two peers try to communicate, they exchange certificates and digitally sign data to authenticate each other. When a new peer is added to the network, you enroll that peer with a CA and none of the other peers need modification. When the new peer attempts an IPsec connection, certificates are automatically exchanged and the peer can be authenticated.

With a CA, a peer authenticates itself to the remote peer by sending a certificate to the remote peer and performing some public key cryptography. Each peer sends its unique certificate, which was issued by the CA. This process works because each certificate encapsulates the public key for the associated peer, each certificate is authenticated by the CA, and all participating peers recognize the CA as an authenticating authority. The process is called IKE with an RSA signature.

The peer can continue sending its certificate for multiple IPsec sessions, and to multiple IPsec peers, until the certificate expires. When its certificate expires, the peer administrator must obtain a new one from the CA.

CAs can also revoke certificates for peers that no longer participate in IPsec. Revoked certificates are not recognized as valid by other peers. Revoked certificates are listed in a CRL, which each peer may check before accepting a certificate from another peer.

Some CAs have an RA as part of their implementation. An RA is a server that acts as a proxy for the CA, so that CA functions can continue when the CA is unavailable.

Key Pairs

Key pairs are RSA or Elliptic Curve Signature Algorithm (ECDSA) keys, which have the following characteristics:

- RSA keys can be used for SSH or SSL.
- SCEP enrollment supports the certification of RSA keys.
- The maximum RSA key size is 4096, and the default is 2048.
- The maximum ECDSA key length is 521, and the default is 384.
- You can generate a general purpose RSA key pair, used for both signing and encryption, or you can generate separate RSA key pairs for each purpose. Separate signing and encryption keys help to reduce exposure of the keys, because 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.

Trustpoints

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.

After you have defined a trustpoint, you can reference it by name in commands requiring that you specify a CA. You can configure many trustpoints.



Note If the ASA has multiple trustpoints that share the same CA, only one of these trustpoints sharing the CA can be used to validate user certificates. To control which trustpoint sharing a CA is used for validation of user certificates issued by that CA, use the **support-user-cert-validation** command.

For automatic enrollment, a trustpoint must be configured with an enrollment URL, and the CA that the trustpoint represents must be available on the network and must support SCEP.

You can export and import the keypair and issued certificates associated with a trustpoint in PKCS12 format. This format is useful to manually duplicate a trustpoint configuration on a different ASA.

Certificate Enrollment

The ASA needs a CA certificate for each trustpoint and one or two certificates for itself, depending upon the configuration of the keys used by the trustpoint. If the trustpoint uses separate RSA keys for signing and encryption, the ASA needs two certificates, one for each purpose. In other key configurations, only one certificate is needed.

The ASA supports automatic enrollment with SCEP and with manual enrollment, which lets you paste a base-64-encoded certificate directly into the terminal. For site-to-site VPNs, you must enroll each ASA. For remote access VPNs, you must enroll each ASA and each remote access VPN client.

You can configure Automated Certificate Management Environment (ACME) protocol to ASA trustpoint to manage the TLS device certificates. The ACME-enabled trustpoints also supports manual certificate enrollment. ACME server validates domain ownership using ASA port 80. ACME is supported only in non-clustered single-context ASA deployments.

Proxy for SCEP Requests

The ASA can proxy SCEP requests between Secure Client and a third-party CA. The CA only needs to be accessible to the ASA if it is acting as the proxy. For the ASA to provide this service, the user must authenticate using any of the methods supported by AAA before the ASA sends an enrollment request. You can also use host scan and dynamic access policies to enforce rules of eligibility to enroll.

The ASA supports this feature only with an Secure Client SSL or IKEv2 VPN session. It supports all SCEP-compliant CAs, including Cisco IOS CS, Windows Server 2003 CA, and Windows Server 2008 CA.

Clientless (browser-based) access does not support SCEP proxy, although WebLaunch—clientless-initiated Secure Client—does support it.

The ASA does not support polling for certificates.

The ASA supports load balancing for this feature.

Revocation Checking

When a certificate is issued, it is valid for a fixed period of time. Sometimes a CA revokes a certificate before this time period expires; for example, because of security concerns or a change of name or association. CAs periodically issue a signed list of revoked certificates. Enabling revocation checking forces the ASA to check that the CA has not revoked a certificate each time that it uses the certificate for authentication.

When you enable revocation checking, the ASA checks certificate revocation status during the PKI certificate validation process, which can use either CRL checking, OCSP, or both. OCSP is only used when the first method returns an error (for example, indicating that the server is unavailable).

With CRL checking, the ASA retrieves, parses, and caches CRLs, which provide a complete list of revoked (and unrevoked) certificates with their certificate serial numbers. The ASA evaluates certificates according to CRLs, also called authority revocation lists, from the identity certificate up the chain of subordinate certificate authorities.

OCSP offers a more scalable method of checking revocation status in that it localizes certificate status through a validation authority, which it queries for status of a specific certificate.

Supported CA Servers

The ASA supports the following CA servers:

Cisco IOS CS, ASA Local CA, and third-party X.509 compliant CA vendors including, but not limited to:

- Baltimore Technologies
- Entrust
- Digicert
- Geotrust
- GoDaddy
- iPlanet/Netscape
- Microsoft Certificate Services
- RSA Keon
- Thawte
- VeriSign

CRLs

CRLs provide the ASA with one way of determining whether a certificate that is within its valid time range has been revoked by the issuing CA. CRL configuration is part of configuration of a trustpoint.

You can configure the ASA to make CRL checks mandatory when authenticating a certificate by using the **revocation-check crl** command. You can also make the CRL check optional by using the **revocation-check crl none** command, which allows the certificate authentication to succeed when the CA is unavailable to provide updated CRL data.



Note The **revocation-check crl none**, which was removed in 9.13(1), was restored.

The ASA can retrieve CRLs from CAs using HTTP, SCEP, or LDAP. CRLs retrieved for each trustpoint are cached for a configurable amount of time for each trustpoint.



Note Though the CRL server responds with HTTP flag "Connection: Keep-alive" to indicate a persistent connection, ASA does not request support for persistent connection. Change the settings on the CRL server to respond with "Connection: Close" when the list is sent.

When the ASA has cached a CRL for longer than the amount of time it is configured to cache CRLs, the ASA considers the CRL too old to be reliable, or “stale.” The ASA tries to retrieve a newer version of the CRL the next time that a certificate authentication requires a check of the stale CRL.

You could receive a *revocation check* failure for a user connection/certificate if you exceed the CRL size limit of 16 MB.

The ASA caches CRLs for an amount of time determined by the following two factors:

- The number of minutes specified with the **cache-time** command. The default value is 60 minutes.

- The NextUpdate field in the CRLs retrieved, which may be absent from CRLs. You control whether the ASA requires and uses the NextUpdate field with the **enforcenextupdate** command.

The ASA uses these two factors in the following ways:

- If the NextUpdate field is not required, the ASA marks CRLs as stale after the length of time defined by the **cache-time** command.
- If the NextUpdate field is required, the ASA marks CRLs as stale at the sooner of the two times specified by the **cache-time** command and the NextUpdate field. For example, if the **cache-time** command is set to 100 minutes and the NextUpdate field specifies that the next update is 70 minutes away, the ASA marks CRLs as stale in 70 minutes.

If the ASA has insufficient memory to store all CRLs cached for a given trustpoint, it deletes the least recently used CRL to make room for a newly retrieved CRL. Large CRLs require significant computational overhead to parse them. Hence, for better performance, use many CRLs of smaller size rather than few large CRLs, or preferably, use OCSP.

See the following the cache sizes:

- Single context mode—128 MB
- Multiple context mode—16 MB per context

OCSP

OCSP provides the ASA with a way of determining whether a certificate that is within its valid time range has been revoked by the issuing CA. OCSP configuration is part of trustpoint configuration.

OCSP localizes certificate status on a validation authority (an OCSP server, also called the *responder*) which the ASA queries for the status of a specific certificate. This method provides better scalability and more up-to-date revocation status than does CRL checking, and helps organizations with large PKI installations deploy and expand secure networks.



Note The ASA allows a five-second time skew for OCSP responses.

You can configure the ASA to make OCSP checks mandatory when authenticating a certificate by using the **revocation-check ocsp** command. You can also make the OCSP check optional by using the **revocation-check ocsp none** command, which allows the certificate authentication to succeed when the validation authority is unavailable to provide updated OCSP data.



Note The **revocation-check ocsp none**, which was removed in 9.13(1), was restored.

OCSP provides three ways to define the OCSP server URL. The ASA uses these servers in the following order:

1. The OCSP URL defined in a match certificate override rule by using the **match certificate** command).
2. The OCSP URL configured by using the **ocsp url** command.
3. The AIA field of the client certificate.

Note To configure a trustpoint to validate a self-signed OCSP responder certificate, you import the self-signed responder certificate into its own trustpoint as a trusted CA certificate. Then you configure the **match certificate** command in the client certificate validating trustpoint to use the trustpoint that includes the self-signed OCSP responder certificate to validate the responder certificate. Use the same procedure for configuring validating responder certificates external to the validation path of the client certificate.

The OCSP server (responder) certificate usually signs the OCSP response. After receiving the response, the ASA tries to verify the responder certificate. The CA normally sets the lifetime of the OCSP responder certificate to a relatively short period to minimize the chance of being compromised. The CA usually also includes an ocsp-no-check extension in the responder certificate, which indicates that this certificate does not need revocation status checking. However, if this extension is not present, the ASA tries to check revocation status using the same method specified in the trustpoint. If the responder certificate is not verifiable, revocation checks fail. To avoid this possibility, use the **revocation-check none** command to configure the responder certificate validating trustpoint, and use the **revocation-check ocsp** command to configure the client certificate.

Certificates and User Login Credentials

The following section describes the different methods of using certificates and user login credentials (username and password) for authentication and authorization. These methods apply to IPsec, Secure Client, and Clientless SSL VPN.

In all cases, LDAP authorization does not use the password as a credential. RADIUS authorization uses either a common password for all users or the username as a password.

User Login Credentials

The default method for authentication and authorization uses the user login credentials.

- Authentication
  - Enabled by the authentication server group setting in the tunnel group (also called ASDM Connection Profile)
  - Uses the username and password as credentials
- Authorization
  - Enabled by the authorization server group setting in the tunnel group (also called ASDM Connection Profile)
  - Uses the username as a credential

Certificates

If user digital certificates are configured, the ASA first validates the certificate. It does not, however, use any of the DNs from certificates as a username for the authentication.

If both authentication and authorization are enabled, the ASA uses the user login credentials for both user authentication and authorization.

- Authentication

- Enabled by the authentication server group setting
- Uses the username and password as credentials

- Authorization
  - Enabled by the authorization server group setting
  - Uses the username as a credential

If authentication is disabled and authorization is enabled, the ASA uses the primary DN field for authorization.

- Authentication
  - DISABLED (set to None) by the authentication server group setting
  - No credentials used

- Authorization
  - Enabled by the authorization server group setting
  - Uses the username value of the certificate primary DN field as a credential



Note If the primary DN field is not present in the certificate, the ASA uses the secondary DN field value as the username for the authorization request.

For example, consider a user certificate that includes the following Subject DN fields and values:

```
Cn=anyuser,OU=sales;O=XYZCorporation;L=boston;S=mass;C=us;ea=anyuser@example.com
```

If the Primary DN = EA (E-mail Address) and the Secondary DN = CN (Common Name), then the username used in the authorization request would be anyuser@example.com.

Guidelines for Digital Certificates

This section includes guidelines and limitations that you should check before configuring digital certificates.

Context Mode Guidelines

- Supported in single context mode only for third-party CAs.

Failover Guidelines

- Does not support replicating sessions in Stateful Failover.
- Does not support failover for local CAs.
- Certificates are automatically copied to the standby unit if you configure stateful failover. If you find a certificate is missing, use the **write standby** command on the active unit.

IPv6 Guidelines

Supports IPv6 OCSP and CRL URLs. You must enclose IPv6 addresses in square brackets, for example: *http://[0:0:0:0:0.18:0a01:7c16]*.

Local CA Certificates

- Make sure that the ASA is configured correctly to support certificates. An incorrectly configured ASA can cause enrollment to fail or request a certificate that includes inaccurate information.
- Make sure that the hostname and domain name of the ASA are configured correctly. To view the currently configured hostname and domain name, enter the **show running-config** command.
- Make sure that the ASA clock is set accurately before configuring the CA. Certificates have a date and time that they become valid and expire. When the ASA enrolls with a CA and obtains a certificate, the ASA checks that the current time is within the valid range for the certificate. If it is outside that range, enrollment fails.
- Thirty days before the local CA certificate expires, a rollover replacement certificate is generated, and a syslog message informs the administrator that it is time for local CA rollover. The new local CA certificate must be imported onto all necessary devices before the current certificate expires. If the administrator does not respond by installing the rollover certificate as the new local CA certificate, validations may fail.
- The local CA certificate rolls over automatically after expiration using the same keypair. The rollover certificate is available for export in base 64 format.

  The following example shows a base 64 encoded local CA certificate:

```
MIIXlwIBAzCCF1EGCSqGSIb3DQEHAaCCF0IEghc+MIIXOjCCFzYGCSqGSIb3DQEHBqCCFycwghcjAgEAMIIXHA
YJKoZIhvcNAQcBMBsGCiqGSIb3DQEMAQMwDQQIjph4SxJoyTgCAQGAghbw3v4bFy+GGG2dJnB4OLphsUM+IG3S
DOiDwZG9n1SvtMieoxd7Hxknxbum06JDrujWKtHBIqkrm+td34qlNE1iGeP2YC94/NQ2z+4kS+uZzwcRhl1KEZ
TS1E4L0fSaC3uMTxJq2NUHYWmoc8pi4CIeLj3h7VVMy6qbx2AC8I+q57+QG5vG5l5Hi5imwtYfaWwPEdPQxaWZ
PrzoG1J8BFqdPa1jBGhAzzuSmElm3j/2dQ3Atro1G9nIsRHgV39fcBgwz4fEabHG7/Vanb+fj81d5nlOiJjDYY
bP86tvbZ2yOVZR6aKFVI0b2AfCr6PbwfC9U8Z/aF3BCyM2sN2xPJrXva94CaYrqyotZdAkSYA5KWScyEcgdqmu
BeGDKOncTknfgy0XM+fG5rb3qAXy1GkjyFI5Bm9Do6RUROoG1DSrQrKeq/hj….

END OF CERTIFICATE
```

SCEP Proxy Support

- Ensure that the ASA and the Cisco ISE Policy Service Nodes are synchronized using the same NTP server.
- Secure Client 3.0 or later must be running at the endpoint.
- The authentication method, configured in the connection profile for your group policy, must be set to use both AAA and certificate authentication.
- An SSL port must be open for IKEv2 VPN connections.
- The CA must be in auto-grant mode.

Additional Guidelines

- The type of certificate you can use is constrained by the certificate types supported by the applications that will use the certificate. RSA certificates are generally supported by all applications that use certificates. But EDDSA certificates might not be supported by workstation operating systems, browsers, ASDM, or Secure Client. For example, you need to use an RSA certificate for remote access VPN identity and authentication. For site-to-site VPN, where the ASA is the application that uses the certificate, EDDSA is supported.
- For ASAs that are configured as CA servers or clients, limit the validity period of the certificate to less than the recommended end date of 03:14:08 UTC, January 19, 2038. This guideline also applies to imported certificates from third-party vendors.
- The ASA establishes LDAP/SSL connection only if one of the following certification criteria is satisfied:
  - The LDAP server certificate is trusted (exists in a trustpoint or the ASA trustpool) and is valid.
  - A CA certificate from servers issuing chain is trusted (exists in a trustpoint or the ASA trustpool) and all subordinate CA certificates in the chain are complete and valid.
- When a certificate enrollment is completed, the ASA stores a PKCS12 file containing the user's keypair and certificate chain, which requires about 2 KB of flash memory or disk space per enrollment. The actual amount of disk space depends on the configured RSA key size and certificate fields. Keep this guideline in mind when adding a large number of pending certificate enrollments on an ASA with a limited amount of available flash memory, because these PKCS12 files are stored in flash memory for the duration of the configured enrollment retrieval timeout. We recommend using a key size of at least 2048.
- You should configure the ASA to use an identity certificate to protect ASDM traffic and HTTPS traffic to the management interface. Identity certificates that are automatically generated with SCEP are regenerated after each reboot, so make sure that you manually install your own identity certificates. For an example of this procedure that applies only to SSL, see the following URL: http://www.cisco.com/en/US/products/ps6120/products_configuration_example09186a00809fcf91.shtml.
- The ASA and the Secure Client can only validate certificates in which the X520Serialnumber field (the serial number in the Subject Name) is in PrintableString format. If the serial number format uses encoding such as UTF8, the certificate authorization will fail.
- Use only valid characters and values for certificate parameters when you import them on the ASA. In ASA, these certificates are decoded to build them into internal data structures. Certificates with blank fields are construed as non-compliant with the decoding standards, and hence the installation validation fails. However, from version 9.16, blank values of optional fields does not impact decoding and installation validation criteria.
- To use a wildcard (*) symbol, make sure that you use encoding on the CA server that allows this character in the string value. Although RFC 5280 recommends using either a UTF8String or PrintableString, you should use UTF8String because PrintableString does not recognize the wildcard as a valid character. The ASA rejects the imported certificate if an invalid character or value is found during the import. For example:

```
ERROR: Failed to parse or verify imported certificate ciscoasa(config)# Read 162*H÷ytes
 as CA certificate:0U0= \Ivr"phÕV°3é¼þ0 CRYPTO_PKI(make trustedCerts list)
CERT-C: E ../cert-c/source/certlist.c(302): Error #711h
CRYPTO_PKI: Failed to verify the ID certificate using the CA certificate in trustpoint
 mm.
CERT-C: E ../cert-c/source/p7contnt.c(169): Error #703h
```

```
crypto_certc_pkcs7_extract_certs_and_crls failed (1795):
crypto_certc_pkcs7_extract_certs_and_crls failed
CRYPTO_PKI: status = 1795: failed to verify or insert the cert into storage
```

Configure Digital Certificates

The following topics explain how to configure digital certificates.

Configure Reference Identities

When the ASA is acting as a TLS client, it supports rules for verification of an application server's identity as defined in RFC 6125. This RFC specifies procedures for representing the reference identities (configured on the ASA) and verifying them against the presented identities (sent from the application server). If the presented identity cannot be matched against the configured reference identity, the connection is not established and an error is logged.

The server presents its identity by including one or more identifiers in the server certificate presented to the ASA while establishing the connection. Reference identities are configured on the ASA, to be compared to the identity presented in a server certificate during connection establishment. These identifiers are specific instances of the four identifier types specified in RFC 6125. The four identifier types are:

- CN_ID: A Relative Distinguished Name (RDN) in a certificate subject field that contains only one attribute-type-and-value pair of type Common Name (CN), where the value matches the overall form of a domain name. The CN value cannot be free text. A CN-ID reference identifier does not identify an application service.
- DNS-ID: A subjectAltName entry of type dNSName. This is a DNS domain name. A DNS-ID reference identifier does not identify an application service.
- SRV-ID: A subjectAltName entry of type otherName whose name form is SRVName as defined in RFC 4985. A SRV-ID identifier may contain both a domain name and an application service type. For example, a SRV-ID of "_imaps.example.net" would be split into a DNS domain name portion of "example.net" and an application service type portion of "imaps."
- URI-ID: A subjectAltName entry of type uniformResourceIdentifier whose value includes both (i) a "scheme" and (ii) a "host" component (or its equivalent) that matches the "reg-name" rule specified in RFC 3986. A URI-ID identifier must contain the DNS domain name, not the IP address, and not just the hostname. For example, a URI-ID of "sip:voice.example.edu" would be split into a DNS domain name portion of "voice.example.edu" and an application service type of "sip."

A reference identity is created when configuring one with a previously unused name. Once a reference identity has been created, the four identifier types and their associated values can be added or deleted from the reference identity. The reference identifiers MAY contain information identifying the application service and MUST contain information identifying the DNS domain name.

Before you begin

- Reference identities are used when connecting to the Syslog Server and the Smart Licensing server only. No other ASA SSL client mode connections currently support the configuration or use of reference identities.

- ASA implements all the rules for matching the identifiers described in RFC 6125 except for pinned certificates and fallback for interactive clients.
- Ability to pin certificates is not implemented. Therefore, `No Match Found, Pinned Certificate` will not occur. Also, a user will not be given the opportunity to pin a certificate if a match is not found since our implementation is not an interactive client.

Procedure

Step 1 Go to **Configuration > Remote Access VPN > Advanced > Reference Identity**.

Configured Reference Identities are listed. You may **Add** a new one, choose and **Edit** an existing one, or choose and **Delete** an existing one. A reference identity that is in use, cannot be deleted.

Step 2 Create or modify the reference-ids by choosing **Add** or **Edit**.

Use this Add or Edit Reference Identity dialog box to choose and specify your reference ids.

- Multiple reference-ids of any type may be added to the reference identity.
- You cannot modify the name once it is set, delete and re-create a reference identity to change the name.

What to do next

Use the reference identity when configuring the Syslog and the Smart Call Home server connections.

How to Set Up Specific Certificate Types

After you have established trusted certificates, you can begin other fundamental tasks such as establishing identity certificates or more advanced configurations such as establishing local CA or code signing certificates.

Before you begin

Read about digital certificate information and establish trusted certificates. CA certificates with no private key are used by all VPN protocols and webvpn, and are configured in trustpoints to validate incoming client certificates. Similarly, a trustpool is a list of trusted certificates used by webvpn features to validate proxied connections to https servers and to validate the smart-call-home certificate.

Procedure

Step 1 An identity certificate is a certificate that is configured on the ASA along with a corresponding private key. It is used for outbound encryption or for signature generation when enabling SSL and IPsec services on the ASA and is obtained through trustpoint enrollment. To configure identity certificates, refer to Identity Certificates, on page 743.

Step 2 A local CA allows VPN clients to enroll for certificates directly from the ASA. This advanced configuration converts the ASA into a CA. To configure CAs, refer to CA Certificates, on page 751.

Step 3 If you are planning to use identity certificates as part of the webvpn java code signing feature, refer to Code Signer Certificate, on page 757.

What to do next

Set up a certificate expiration alert or monitor digital certificates and certificate management history.

Identity Certificates

An identity certificate can be used to authenticate VPN access through the ASA.

In the Identity Certificates Authentication pane, you can perform the following tasks:

- Add or Import an Identity Certificate, on page 743.
- Enable CMPv2 Enrollments as a Request from a CA
- Display details of an identity certificate.
- Delete an existing identity certificate.
- Export an Identity Certificate, on page 748.
- Set certificate expiration alerts.
- Enroll for an identity certificate with Entrust Generate a Certificate Signing Request, on page 749.

Add or Import an Identity Certificate

To add or import a new identity certificate configuration, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Remote Access VPN** > **Certificate Management** > **Identity Certificates**.

Step 2 Click **Add**.

The **Add Identity Certificate** dialog box appears, with the selected trustpoint name displayed at the top.

Step 3 Click the **Import the identity certificate from a file (PKCS12 format with Certificate(s) + Private Key)** radio button to import an identity certificate from an existing file.

Step 4 Enter the passphrase used to decrypt the PKCS12 file.

Step 5 Enter the path name of the file, or click **Browse** to display the **Import ID Certificate File** dialog box. Find the certificate file, then click **Import ID Certificate File**.

Step 6 Click the **Add a new identity certificate** radio button to add a new identity certificate.

Step 7 Click **New** to display the **Add Key Pair** dialog box.

Note
Named key pairs are not supported for ACME certificate enrollment.

Step 8 Choose the **RSA**, **ECDSA**, or **EdDSA** key type.

Step 9 If you choose **EdDSA**, the **Edwards Curve** option appears. Click the EdDSA1 radio button.

Step 10 Click the **Use default keypair name** radio button to use the default key pair name.

Step 11 Click the **Enter a new key pair name** radio button, then enter the new name.

Step 12 Choose the modulus size from the drop-down list. If **Edwards Curve** is selected, choose Ed25519. If you are not sure of the modulus size, consult Entrust.

For ASA 9.16(1) and higher versions, ensure that you select an RSA modulus size of 2048 or bigger. CA Certificate validation fails when the RSA key size is smaller than 2048 bits. To override this restriction, enable the permit weak crypto option. (See Permit Weak Crypto for CA Certificates, on page 756).

Step 13 Choose the key pair usage by clicking the **General purpose** radio button (default) or **Special** radio button. When you choose the **Special** radio button, the ASA generates two key pairs, one for signature use and one for encryption use. This selection indicates that two certificates are required for the corresponding identity.

Step 14 Click **Generate Now** to create new key pairs, then click **Show** to display the **Key Pair Details** dialog box, which includes the following display-only information:

- The name of the key pair whose public key is to be certified.
- The time of day and the date when the key pair is generated.
- The usage of an RSA key pair.
- The modulus size (bits) of the key pairs: 2048, 3072, and 4096. The default is 2048.
- The key data, which includes the specific key data in text format.

Step 15 Click **OK** when you are done.

Step 16 Choose a certificate subject DN to form the DN in the identity certificate, then click **Select** to display the **Certificate Subject DN** dialog box.

Step 17 Choose one or more DN attributes that you want to add from the drop-down list, enter a value, then click **Add**. Available X.500 attributes for the Certificate Subject DN are the following:

- **Common Name (CN)**
- **Department (OU)**
- **Company Name (O)**
- **Country (C)**
- **State/Province (ST)**
- **Location (L)**
- **E-mail Address (EA)**

For ACME certification, **Common Name (CN)** value must be a valid FQDN.

Step 18 Click **OK** when you are done.

Step 19 Check the **Generate self-signed certificate** check box to create self-signed certificates.

Step 20 Check the **Act as local certificate authority and issue dynamic certificates to TLS proxy** check box to have the identity certificate act as the local CA.

Step 21 Click **Advanced** to establish additional identity certificate settings.

The **Advanced Options** dialog box appears, with the following three tabs: **Certificate Parameters**, **Enrollment Mode**, and **SCEP Challenge Password**.

Note
Enrollment mode settings and the SCEP challenge password are not available for self-signed certificates.

Step 22 Click the **Certificate Parameters** tab, and then enter the following information:

- The FQDN, an unambiguous domain name, to indicate the position of the node in the DNS tree hierarchy. For ACME enrollment, you can provide additional FQDN values separated by space in the **Additional FQDNs** text box.
- The e-mail address associated with the identity certificate. This field is not supported for ACME.
- The ASA IP address on the network in four-part, dotted-decimal notation. This field is not supported for ACME.
- Check the **Include serial number of the device** check box to add the ASA serial number to the certificate parameters.

Step 23 Click the **Enrollment Mode** tab, then enter the following information:

- Choose the enrollment method by clicking the **Request by manual enrollment** radio button or the **Request from a CA** radio button. If you choose **Request from a CA**, for information about enabling CMPV2 enrollments, refer to Enable CMPv2 Enrollments as a Request from a CA , on page 746. For information about enabling ACME enrollments, refer to Enable ACME Enrollments as a Request from a CA, on page 747.
- Choose the enrollment protocol—scep, cmp, acme, or est.

  Note
  Only RSA and ECDSA keys are supported for CMP and ACME certificate enrollment.

- The enrollment URL of the certificate to be automatically installed through SCEP.
- The maximum number of minutes allowed to retry installing an identity certificate. The default is one minute.
- The maximum number of retries allowed for installing an identity certificate. The default is zero, which indicates an unlimited number of retries within the retry period.

Step 24 Click the **SCEP Challenge Password** tab, then enter the following information:

- The SCEP password
- The SCEP password confirmation

Note
SCEP password is not supported in ACME certificate enrollment.

Step 25 Click **OK** when you are done.

Step 26 Check the **Enable CA flag in basic constraints extension** if this certificate should be able to sign other certificates.

The basic constraints extension identifies whether the subject of the certificate is a Certificate Authority (CA), in which case the certificate can be used to sign other certificates. The CA flag is part of this extension. The

presence of these items in a certificate indicates that the certificate's public key can be used to validate certificate signatures. There is no harm in leaving this option selected.

Step 27 Click **Add Certificate** in the **Add Identity Certificate** dialog box.

The new identity certificate appears in the Identity Certificates list.

Step 28 Click **Apply** to save the new identity certificate configuration.

Step 29 Click **Show Details** to display the **Certificate Details** dialog box, which includes the following three display-only tabs:

- The **General** tab displays the values for type, serial number, status, usage, public key type, CRL distribution point, the times within which the certificate is valid, and associated trust points. The values apply to both available and pending status.
- The **Issued to** tab displays the X.500 fields of the subject DN or certificate owner and their values. The values apply only to available status.
- The **Issued by** tab displays the X.500 fields of the entity granting the certificate. The values apply only to available status.

Step 30 To remove an identity certificate configuration, select it, then click **Delete**.

Note
After you delete a certificate configuration, it cannot be restored. To recreate the deleted certificate, click **Add** to reenter all of the certificate configuration information.

Enable CMPv2 Enrollments as a Request from a CA

To be positioned as a Security Gateway device in wireless LTE networks, ASA supports some certificate management functions using the Certificate Management Protocol (CMPv2). Using CMPv2 for enrollment of ASA device certificates, you can perform manual enrollment, for the first and secondary certificate from the CMPv2-enabled CA, or manual certificate updates, for replacement of a previously issued certificate using the same key pair. The received certificates are stored outside of the conventional configuration and are used in certificate-enabled IPsec configurations.

Note You will not have the full CMPv2 functionality on the ASA.

An initial request establishes trust with the CA and obtains the first certificate. A CA certificate must be preconfigured in a trustpoint. Authentication occurs when you acknowledge the fingerprint of the certificate that is being installed.

After clicking **Request from a CA** on the Enrollment Mode tab of the Advanced Options window, complete the following steps specific for CMPv2 enrollments:

Before you begin

- Follow the steps in Add or Import an Identity Certificate, on page 743.
- Carrier license must be included to enable CMP enrollment.

Procedure

Step 1 Choose CMP as the enrollment protocol and enter the CMP URL in the http:// area.

Step 2 To automatically generate new keypairs for all CMP manual and automatic enrollments, choose either **RSA** or **EDCSA**.

If you choose RSA, choose a value from the Modulus drop-down menu. If you choose EDCSA, choose a value from the elliptic-curve drop-down menu.

Step 3 (Optional) Click **Regenerate the key pair** to generate a key pair while renewing the certificate or prior to building the enrollment request.

Step 4 Click **Shared Key** and enter a value provided out of band by the CA. This value is used by the CA and ASA to confirm the authenticity and integrity of the messages that they exchange.

Step 5 Click **Signing Trustpoint** and enter the name of the trustpoint which contains a previously-issued device certificate used to sign the CMP enrollment request.

These options are only available when the trustpoint enrollment protocol is set to CMP. When a CMP trustpoint, the shared secret or the signing certificate can be specified, but not both.

Step 6 Click **Browse Certificate** to specify the CA certificate.

Step 7 (Optional) Click the **Auto Enroll** checkbox to trigger auto-enroll of CMPv2.

Step 8 At the Auto Enroll Lifetime field, enter the percentage of the absolute lifetime of the certificate after which auto-enroll will be necessary.

Step 9 Click **Auto Enroll Regenerate Key** to generate a new key while renewing the certificate.

Enable ACME Enrollments as a Request from a CA

To automate certificate management of TLS device certificates, ASA supports certificate management functions using the Automatic Certificate Management Environment (ACME). ACME enables ASA to simplify certificate management through auto renewal, domain validation, and easy enrolling and revoking of certificates. As with CMP, the received certificates are stored outside the conventional configuration and are used in certificate-enabled IPsec configurations.

An initial request establishes trust with the CA and obtains the first certificate. A CA certificate must be preconfigured in a trustpoint. Authentication occurs only through the HTTP-01 method for the certificate that is being installed.

In the **Add Identity Certificate** page, choose **Add a new Identity Certificate** > **Advanced**. Under the **Enrollment Mode** tab, click **Request from a CA** and complete the following steps, which are specific for ACME enrollments.

Before you begin

- Follow the steps provided in Add or Import an Identity Certificate, on page 743.
- Configure DNS server on the device for a sucessful ACME enrollment. For information on DNS server configuration procedure, see Configure the DNS Servers, on page 701.

Procedure

Step 1 From the **Source Interface** drop-down list, choose the interface for the ACME certification.

Step 2 From the **Enrollment Protocol** drop-down list, choose **ACME**.

Step 3 To use the Let's Encrypt CA, check the **Let's Encrypt** check box.

The CA server URL: *https://acme-v02.api.letsencrypt.org:443/directory* is displayed in the **https://** field.

To use any other ACME server, enter the corresponding url in the **https://** field.

Step 4 From the **Authentication Method** drop-down list, choose **http01**.

Step 5 From the **Authentication Interface** drop-down list, choose the interface that ACME uses for the domain authentication.

Step 6 To automatically generate new key pairs for all ACME manual and automatic enrollments, choose either **RSA** or **EDCSA**.

If you choose **RSA**, choose a value from the **Modulus** drop-down list. If you choose **EDCSA**, choose a value from the **elliptic-curve** drop-down list.

Step 7 (Optional) Click **Regenerate the key pair** to generate a key pair while renewing the certificate or before building the enrollment request.

Step 8 Check **Install CA Certificate** check box, and then click **Browse Certificate**.

Step 9 Select the CA certificate to be used and click **Choose CA Certificate**.

For the first time of ACME enrollment, you must install CA certificate.

Step 10 For subsequent trustpoints in the already enrolled device, do not check the **Install CA Certificate** check box if you want to use the already enrolled ACME certificate. However, ensure that the CA in the signature path of the issued ACME certificate is present in the trustpoint or trustpool.

Step 11 (Optional) Click **Auto Enroll** to trigger auto enroll for the specified auto enroll lifetime.

Step 12 In the **Auto Enroll Lifetime** field, enter the percentage of the absolute lifetime of the certificate after which auto-enroll will be necessary.

Step 13 Click **Auto Enroll Regenerate Key** to generate a new key while renewing the certificate.

Step 14 Click **Ok**.

Step 15 In the **Add Identity Certificate** dialog box, click **Add certificate**.

For other trustpoints in the device, the certificate is automatically enrolled without manually installing the CA certificate.

Step 16 To send the certificate to ASA, click **Send**.

Export an Identity Certificate

To export an identity certificate, perform the following steps:

Procedure

Step 1 Click **Export** to display the **Export Certificate** dialog box.

Step 2 Enter the name of the PKCS12 format file to use in exporting the certificate configuration. Alternatively, click **Browse** to display the **Export ID Certificate File** dialog box to find the file to which you want to export the certificate configuration.

Step 3 Choose the certificate format by clicking the **PKCS12 Format** radio button or the **PEM Format** radio button.

Step 4 Enter the passphrase used to encrypt the PKCS12 file for export.

Step 5 Confirm the encryption passphrase.

Step 6 Click **Export Certificate** to export the certificate configuration.

An information dialog box appears, informing you that the certificate configuration file has been successfully exported to the location that you specified.

Generate a Certificate Signing Request

To generate a certificate signing request to send to Entrust, perform the following steps:

Procedure

Step 1 Click **Enroll ASA SSL VPN with Entrust** to display the **Generate Certificate Signing Request** dialog box.

Step 2 Perform the following steps in the **Key Pair** area:

a) Choose one of the configured key pairs from the drop-down list.
b) Click **Show** to display the **Key Details** dialog box, which provides information about the selected key pair, including date and time generated, usage (general or special purpose), modulus size, and key data.
c) Click **OK** when you are done.
d) Click **New** to display the **Add Key Pair** dialog box. When you generate the key pair, you can send it to the ASA or save it to a file.

Step 3 Enter the following information in the **Certificate Subject DN** area:

a) The FQDN or IP address of the ASA.
b) The name of the company.
c) The two-letter country code.

Step 4 Perform the following steps in the **Optional Parameters** area:

a) Click **Select** to display the **Additional DN Attributes** dialog box.
b) Choose the attribute to add from the drop-down list, then enter a value.
c) Click **Add** to add each attribute to the attribute table.
d) Click **Delete** to remove an attribute from the attribute table.
e) Click **OK** when you are done.

The added attributes appear in the **Additional DN Attributes** field.

Step 5 Enter additional fully qualified domain name information if the CA requires it.

Step 6 Click **Generate Request** to generate the certificate signing request, which you can then send to Entrust, or save to a file and send later.

The **Enroll with Entrust** dialog box appears, with the CSR displayed.

Step 7 Complete the enrollment process by clicking the **request a certificate from Entrust** link. Then copy and paste the CSR provided and submit it through the Entrust web form, provided at http://www.entrust.net/cisco/. Alternatively, to enroll at a later time, save the generated CSR to a file, then click the **enroll with Entrust** link on the **Identity Certificates** pane.

Step 8 Entrust issues a certificate after verifying the authenticity of your request. which may take several days. You then need to install the certificate by selecting the pending request in the **Identity Certificate** pane and clicking **Install**.

Step 9 Click **Close** to close the **Enroll with Entrust** dialog box.

Install Identity Certificates

To install a new identity certificate, perform the following steps:

Procedure

Step 1 Click **Add**in the **Identity Certificates** pane to display the **Add Identity Certificate** dialog box.

Step 2 Click the **Add a new identity certificate** radio button.

Step 3 Change the key pair or create a new key pair. A key pair is required.

Step 4 Enter the certificate subject DN information, then click **Select** to display the **Certificate Subject DN** dialog box.

Step 5 Specify all of the subject DN attributes required by the CA involved, then click **OK** to close the **Certificate Subject DN** dialog box.

Step 6 In the **Add Identity Certificate** dialog box, click **Advanced** to display the **Advanced Options** dialog box.

Step 7 To continue, see Steps 17 through 23 of the Add or Import an Identity Certificate, on page 743.

Step 8 In the **Add Identity Certificate** dialog box, click **Add Certificate**.

The **Identity Certificate Request** dialog box appears.

Step 9 Enter the CSR file name of type, text, such as c:\verisign-csr.txt, then click **OK**.

Step 10 Send the CSR text file to the CA. Alternatively, you can paste the text file into the CSR enrollment page on the CA website.

Step 11 When the CA returns the identity certificate to you, go to the **Identity Certificates** pane, select the pending certificate entry, then click **Install**.

The **Install Identity Certificate** dialog box appears.

Step 12 Choose one of the following options by clicking the applicable radio button:

- **Install from a file**.

  Alternatively, click **Browse** to search for the file.

- **Paste the certificate data in base-64 format**.

Paste the copied certificate data into the area provided.

Step 13 Click **Install Certificate**.

Step 14 Click **Apply** to save the newly installed certificate with the ASA configuration.

Step 15 To show detailed information about the selected identity certificate, click **Show Details** to display the **Certificate Details** dialog box, which includes the following three display-only tabs:

The **General** tab displays the values for type, serial number, status, usage, public key type, CRL distribution point, the times within which the certificate is valid, and associated trustpoints. The values apply to both available and pending status.

The **Issued to** tab displays the X.500 fields of the subject DN or certificate owner and their values. The values apply only to available status.

The **Issued by** tab displays the X.500 fields of the entity granting the certificate. The values apply only to available status.

Step 16 To remove a code signer certificate configuration, select it, and then click **Delete**.

Note
After you delete a certificate configuration, it cannot be restored. To recreate the deleted certificate, click **Import** to reenter all of the certificate configuration information.

CA Certificates

This page is where you manage CA certificates. The following topics explain what you can do.

Add or Install a CA Certificate

To add or install a CA certificate, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Remote Access VPN** > **Certificate Management** > **CA Certificates**.

Step 2 Click **Add**.

The **Install Certificate** dialog box appears.

Step 3 Click the **Install from a file** radio button to add a certificate configuration from an existing file (this is the default setting).

Step 4 Enter the path and file name, or click **Browse** to search for the file. Then click **Install Certificate**.

Step 5 The **Certificate Installation** dialog box appears with a confirmation message indicating that the certificate was successfully installed. Click **OK** to close this dialog box.

Step 6 Click the **Paste certificate in PEM format** radio button to enroll manually.

Step 7 Copy and paste the PEM format (base64 or hexadecimal) certificate into the area provided, then click **Install Certificate**.

Step 8 The **Certificate Installation** dialog box appears with a confirmation message indicating that the certificate was successfully installed. Click **OK** to close this dialog box.

Step 9 Click the **Use SCEP** radio button to enroll automatically. The ASA contacts the CA using SCEP, obtains the certificates, and installs them on the device. To use SCEP, you must enroll with a CA that supports SCEP, and you must enroll via the Internet. Automatic enrollment using SCEP requires that you provide the following information:

- The path and file name of the certificate to be automatically installed.
- The maximum number of minutes to retry certificate installation. The default is one minute.
- The number of retries for installing a certificate. The default is zero, which indicates unlimited retries within the retry period.

Step 10 Click **More Options** to display additional configuration options for new and existing certificates.

The **Configuration Options for CA Certificates** pane appears.

Step 11 To change an existing CA certificate configuration, select it, then click **Edit**.

Step 12 To remove a CA certificate configuration, select it, then click **Delete**.

Note
After you delete a certificate configuration, it cannot be restored. To recreate the deleted certificate, click **Add** to reenter all of the certificate configuration information.

Step 13 Click **Show Details** to display the **Certificate Details** dialog box, which includes the following three display-only tabs:

- The **General** tab displays the values for type, serial number, status, usage, public key type, CRL distribution point, the times within which the certificate is valid, and associated trust points. The values apply to both available and pending status.
- The **Issued to** tab displays the X.500 fields of the subject DN or certificate owner and their values. The values apply only to available status.
- The **Issued by** tab displays the X.500 fields of the entity granting the certificate. The values apply only to available status.

Configure CA Certificates for Revocation

To configure CA certificates for revocation, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Site-to-Site VPN** > **Certificate Management** > **CA Certificates** > **Add** to display the **Install Certificates** dialog box. Then click **More Options**.

Step 2 Click the **Revocation Check** tab.

Step 3 Click the **Do not check certificates for revocation** radio button to disable revocation checking of certificates.

Step 4 Click the **Check certificates for revocation** radio button to select one or more revocation checking methods (CRL or OCSP).

Step 5 Click **Add** to move a revocation method to the right and make it available. Click **Move Up** or **Move Down** to change the method order.

The methods you choose are implemented in the order in which you add them. If a method returns an error, the next revocation checking method activates.

Step 6 Check the **Consider certificate valid if revocation information cannot be retrieved** check box to ignore revocation checking errors during certificate validation.

Step 7 Click **OK** to close the **Revocation Check** tab.

Configure CRL Retrieval Policy

To configure the CRL retrieval policy, perform the following steps:

Before you begin

- To assign static URLs,

Procedure

Step 1 Choose **Configuration** > **Site-to-Site VPN** > **Certificate Management** > **CA Certificates** > **Add** to display the **Install Certificates** dialog box. Then click **More Options**.

Step 2 Click the **CRL Retrieval Policy** tab.

Step 3 Check the **Use CRL Distribution Point from the certificate** check box to direct revocation checking to the CRL distribution point from the certificate being checked.

Step 4 Check the **Use Static URLs configured below** check box to list specific URLs to be used for CRL retrieval. The URLs you select are implemented in the order in which you add them. If an error occurs with the specified URL, the next URL in order is taken.

Step 5 Click **Add** in the **CRL Distribution Point** tab.

Step 6 In the **Add CDP Rule** dialog box, select the certificate map, enter the index, and the URL.

Note
Ensure that you have configured certificate map on the device—Go to **Remote Access VPN** > **Network (Client) Access** > **Advanced** > **IPSec** > **Certificate to Connection Map** > **Rules** > **Add** .

ASA supports IPv4 or IPv6 based CDP and Static URLs. Enclose IPv6 addresses in square brackets, for example: *http://[0:0:0:0:0.18:0a01:7c16]*.

Step 7 Click **OK** to close this dialog box.

Configure CRL Retrieval Methods

To configure CRL retrieval methods, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Site-to-Site VPN** > **Certificate Management** > **CA Certificates** > **Add** to display the **Install Certificates** dialog box. Then click **More Options**.

Step 2 Click the **CRL Retrieval Methods** tab in the **Configuration Options for CA Certificates** pane.

Step 3 Choose one of the following three retrieval methods:

- To enable LDAP for CRL retrieval, check the **Enable Lightweight Directory Access Protocol (LDAP)** check box. With LDAP, CRL retrieval starts an LDAP session by connecting to a named LDAP server, accessed by a password. The connection is on TCP port 389 by default. Enter the following required parameters:
  - **Name**
  - **Password**
  - **Confirm Password**
  - **Default Server** (server name)
  - **Default Port** (389)
- To enable HTTP for CRL retrieval, check the **Enable HTTP** check box.

Step 4 Click **OK** to close this tab.

Configure OCSP Rules

To configure OCSP rules for obtaining revocation status of an X.509 digital certificate, perform the following steps.

Before you begin

Make sure that you have configured a certificate map before you try to add OCSP rules. If a certificate map has not been configured, an error message appears.

Procedure

Step 1 Choose **Configuration** > **Site-to-Site VPN** > **Certificate Management** > **CA Certificates** > **Add** to display the **Install Certificates** dialog box. Then click **More Options**.

Step 2 Click the **OCSP Rules** tab in the **Configuration Options for CA Certificates** pane.

Step 3 Choose the certificate map to match to this OCSP rule. Certificate maps match user permissions to specific fields in a certificate. The name of the CA that the ASA uses to validate responder certificates appears in the **Certificate** field. The priority number for the rule appears in the **Index** field. The URL of the OCSP server for this certificate appears in the **URL** field.

Step 4 Click **Add**.

The **Add OCSP Rule** dialog box appears.

Step 5 Choose the certificate map to use from the drop-down list.

Step 6 Choose the certificate to use from the drop-down list.

Step 7 Enter the priority number for the rule.

Step 8 Enter the URL of the OCSP server for this certificate.

Step 9 When you are done, click **OK** to close this dialog box.

The newly added OCSP rule appears in the list.

Step 10 Click **OK** to close this tab.

Configure Advanced CRL and OCSP Settings

To configure additional CRL and OCSP settings, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Site-to-Site VPN** > **Certificate Management** > **CA Certificates** > **Add** to display the **Install Certificates** dialog box. Then click **More Options**.

Step 2 Click the **Advanced** tab in the **Configuration Options for CA Certificates** pane.

Step 3 Enter the number of minutes between cache refreshes in the **CRL Options** area. The default is 60 minutes. The range is 1-1440 minutes. To avoid having to retrieve the same CRL from a CA repeatedly, the ASA can store retrieved CRLs locally, which is called CRL caching. The CRL cache capacity varies by platform and is cumulative across all contexts. If an attempt to cache a newly retrieved CRL would exceed its storage limits, the ASA removes the least recently used CRL until more space becomes available.

Step 4 Check the **Enforce next CRL update** check box to require valid CRLs to have a Next Update value that has not expired. Uncheck the **Enforce next CRL update** check box to let valid CRLs with no Next Update value or a Next Update value that has expired.

Step 5 Enter the URL for the OCSP server in the **OCSP Options** area. The ASA uses OCSP servers according to the following order:

a) OCSP URL in a match certificate override rule
b) OCSP URL configured in the selected OCSP Options attribute
c) AIA field of a user certificate

Step 6 By default, the **Disable nonce extension** check box is checked, which cryptographically binds requests with responses to avoid replay attacks. This process works by matching the extension in the request to that in the response, ensuring that they are the same. Uncheck the **Disable nonce extension** check box if the OCSP server you are using sends pregenerated responses that do not include this matching nonce extension.

Step 7 Choose one of the following options in the **Other Options** area:

- Check the **Accept certificates issued by this CA** check box to indicate that the ASA should accept certificates from the specified CA.
- Check the **Accept certificates issued by the subordinate CAs of this CA** check box to indicate that the ASA should accept certificates from the subordinate CA.

Step 8 Click **OK** to close this tab, then click **Apply** to save your configuration changes.

CA Server Management

Manage User Certificates

To change the certificate status, perform the following steps:

Procedure

Step 1 Select specific certificates by username or by certificate serial number in the **Manage User Certificates** pane.

Step 2 Choose one of the following options:

- Click **Revoke**to remove user access if the user certificate lifetime period runs out. The local CA also marks the certificate as revoked in the certificate database, automatically updates the information, and reissues the CRL.
- Select a revoked certificate and click **Unrevoke** to restore access. The local CA also marks the certificate as unrevoked in the certificate database, automatically updates the certificate information, and reissues an updated CRL.

Step 3 Click **Apply** when you are done to save your changes.

Permit Weak Crypto for CA Certificates

When the following attributes are present, the CA certificate validation operation fails:

- Certificates signed with the SHA-1 with RSA encryption algorithm.
- Certificates with RSA key sizes smaller than 2048 bits.

However, you can override these restrictions by configuring the permit weak crypto option. When enabled, the ASA allows using the above attributes when validating certificates. We do not recommend permitting weak-crypto keys, because such keys are not as secure as the ones with higher key sizes.

Procedure

Step 1 Browse to **Configuration** > **Device Management** > **Certificate Management** > **Identity Certificate**, or **Configuration** > **Remote Access VPN** > **Certificate Management** > **Identity Certificate**, or **Configuration** > **Remote Access VPN** > **Certificate Management** > **Code Signer**.

Step 2 To permit key sizes that are smaller than 2048 bits and SHA-1 signature algorithms, under **Weak Crypto Configurations**, click the **Permit Weak Key Sizes and Hash Algorithms** check box.

Code Signer Certificate

Import a Code Signer Certificate

To import a code signer certificate, perform the following steps:

Procedure

Step 1 In the **Code Signer** pane, click **Import** to display the **Import Certificate** dialog box.

Step 2 Enter the passphrase used to decrypt the PKCS12-format file.

Step 3 Enter the name of the file to import, or click **Browse** to display the **Import ID Certificate File** dialog box and search for the file.

Step 4 Select the file to import and click **Import ID Certificate File**.

The selected certificate file appears in the **Import Certificate** dialog box.

Step 5 Click **Import Certificate**.

The imported certificate appears in the **Code Signer** pane.

Step 6 Click **Apply** to save the newly imported code signer certificate configuration.

Export a Code Signer Certificate

To export a code signer certificate, perform the following steps:

Procedure

Step 1 In the **Code Signer** pane, click **Export** to display the **Export Certificate** dialog box.

Step 2 Enter the name of the PKCS12 format file to use in exporting the certificate configuration.

Step 3 In the **Certificate Format** area, to use the public key cryptography standard, which can be base64 encoded or in hexadecimal format, click the **PKCS12 format** radio button. Otherwise, click the **PEM format** radio button.

Step 4 Click **Browse** to display the **Export ID Certificate File** dialog box to find the file to which you want to export the certificate configuration.

Step 5 Select the file and click **Export ID Certificate File**.

The selected certificate file appears in the **Export Certificate** dialog box.

Step 6 Enter the passphrase used to decrypt the PKCS12 format file for export.

Step 7 Confirm the decryption passphrase.

Step 8 Click **Export Certificate** to export the certificate configuration.

Set a Certificate Expiration Alert (for Identity or CA Certificates)

ASA checks all the CA and ID certificates in the trust points for expiration once every 24 hours. If a certificate is nearing expiration, a syslog will be issued as an alert.

In addition to the renewal reminder, if an already expired certificate is found in the configuration, a syslog is generated once every day to rectify the configuration by either renewing the certificate or removing the expired certificate.

For example, assume that the expiration alerts are configured to begin at 60 days and repeat every 6 days after that. If the ASA is rebooted at 40 days, an alert is sent on that day, and the next alert is sent on the 36th day.



Note Expiration checking is not done on trust pool certificates. The Local CA trust point is treated as a regular trustpoint for expiration checking too.

Procedure

Step 1 Browse to **Configuration** > **Device Management** > **Certificate Management** > **Identity Certificate/CA Certificate**.

Step 2 Check the **Enable Certificate Expiration Alert** check box.

Step 3 Fill in the desired number of days:

- Send the first alert before—Configure the number of days (1 to 90) before expiration at which the first alert will go out.
- Repeat the alert for—Configure the alert frequency (1 to 14 days) if the certificate is not renewed. By default, the first alert is sent 60 days prior to expiration and once every week after until the certificate is renewed and removed. Additionally, an alert is sent on the day of the expiration and once every day after that, and irrespective of the alert configuration, an alert is sent every day during the last week of expiration.

Monitoring Digital Certificates

See the following commands for monitoring digital certificate status:

- **Monitoring > Properties > CRL**

  This pane shows CRL details.

- **Tools > Command Line Interface**

  This pane allows you to issue various non-interactive commands and view results.

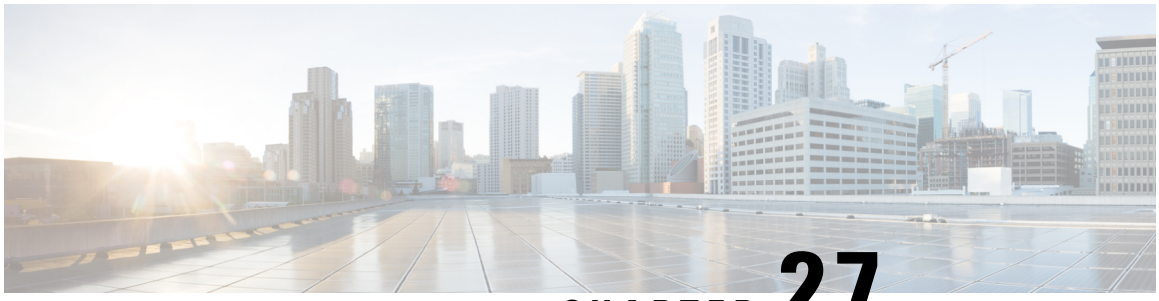
History for Certificate Management

Table 33: History for Certificate Management

Feature Name	Platform Releases	Description
Certificate management	7.0(1)	Digital certificates (including CA certificates, identity certificates, and code signer certificates) provide digital identification for authentication. A digital certificate includes information that identifies a device or user, such as the name, serial number, company, department, or IP address. CAs are trusted authorities that "sign" certificates to verify their authenticity, thereby guaranteeing the identity of the device or user. CAs issue digital certificates in the context of a PKI, which uses public-key or private-key encryption to ensure security. We introduced the following screens: Configuration > Remote Access VPN > Certificate Management Configuration > Site-to-Site VPN > Certificate Management. We introduced or modified the following screens: Configuration > Firewall > Advanced > Certificate Management > CA Certificates Configuration > Device Management > Certificate Management > CA Certificates.
Certificate management	7.2(1)	
Certificate management	8.0(2)	
SCEP proxy	8.4(1)	We introduced this feature, which provides secure deployment of device certificates from third-party CAs.
Reference Identities	9.6(2)	TLS client processing now supports rules for verification of a server identity defined in RFC 6125, Section 6. Identity verification will be done during PKI validation for TLS connections to the Syslog Server and the Smart Licensing server only. If the presented identity cannot be matched against the configured reference identity, the connection is not established. We modifed the following screens: Configuration > Remote Access VPN > Advanced Configuration > Device Management > Logging > Syslog Servers > Add/Edit Configuration > Device Management > Smart Call Home

Feature Name	Platform Releases	Description
Local CA Server	9.12(1)	To make the FQDN of the enrollment URL configurable instead of using the ASA's configured FQDN, a new CLI option is introduced. This new option is added to the smpt mode of crypto ca server. We deprecated Local CA Server and will be removing in a later release—When ASA is configured as local CA server, it is enabled to issue digital certificates, publish Certificate Revocation Lists (CRLs), and securely revoke issued certificates. This feature has become obsolete and hence the crypto ca server command is deprecated.
Local CA Server	9.13(1)	We removed the local ca server support. Thus, the **crypto ca server** command and its subcommands are removed. We removed the following commands: **crypto ca server** and all of its subcommands.
Modifications to the CRL Distribution Point commands	9.13(1)	The static CDP URL configuration commands are removed and moved to the match certificate command. New/Modified screens: **Configuration** > **Device Management** > **Certificate Management** > **CA Certificates**
CRL cache size increased	9.13(1)	To prevent failure of large CRL downloads, the cache size was increased, and the limit on the number of entries in an individual CRL was removed. • Increased the total CRL cache size to 16 MB per context for multi-context mode. • Increased the total CRL cache size to 128 MB for single-context mode.
Restoration of bypass certificate validity checks option	9.15(1)	The option to bypass revocation checking due to connectivity problems with the CRL or OCSP server that was removed in 9.13(1) was restored.
Modifications to Match Certificate commands to support static CRL Distribution Point URL	9.15(1)	The static CDP URL configuration command allowed static CDPs to be mapped uniquely to each certificate in a chain that is being validated. However, only one such mapping was supported for each certificate. This modification allows statically configured CDPs to be mapped to a chain of certificates for authentication.

Feature Name	Platform Releases	Description
Modifications to the trustpoint keypair and crypto key generate commands	9.16(1)	Support for certificates with key sizes smaller than 2048 was removed. Any configuration using 512, 768 or 1024-bit options are transitioned to 2048 with due notification. Support to use SHA1 hashing algorithm for certification was removed. **Note** **crypto ca permit-weak-crypto** command was introduced to override these restrictions. The new key option - EDDSA was added to the existing RSA and ECDSA options.
Support for OCSP and CRL IPv6 URL	9.20(1)	Support to use IPv6 OCSP and CRL URLs were added. The IPv6 addresses must be enclosed in square brackets.



CHAPTER 27

ARP Inspection and the MAC Address Table

This chapter describes how to customize the MAC address table and configure ARP Inspection for bridge groups.

- About ARP Inspection and the MAC Address Table, on page 763
- Default Settings, on page 764
- Guidelines for ARP Inspection and the MAC Address Table, on page 764
- Configure ARP Inspection and Other ARP Parameters, on page 765
- Customize the MAC Address Table for Bridge Groups, on page 767
- History for ARP Inspection and the MAC Address Table, on page 768

About ARP Inspection and the MAC Address Table

For interfaces in a bridge group, ARP inspection prevents a “man-in-the-middle” attack. You can also customize other ARP settings. You can customize the MAC address table for bridge groups, including adding a static ARP entry to guard against MAC spoofing.

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 ASA 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 ASA drops the packet.

- If the ARP packet does not match any entries in the static ARP table, then you can set the ASA 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 ASA 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 ASA adds the MAC address to its table. The table associates the MAC address with the source interface so that the ASA knows to send any packets addressed to the device out the correct interface. Because traffic between bridge group members is subject to the ASA security policy, if the destination MAC address of a packet is not in the table, the ASA 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 ASA 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 ASA generates a ping to the destination IP address so that it can learn which interface receives the ping reply.

The original packet is dropped.

For routed mode, you can optionally enable flooding of non-IP packets on all interfaces.

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 ASA adds corresponding entries to the MAC address table.

Note Secure Firewall ASA 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 ARP Inspection and Other ARP Parameters

For bridge groups, you can enable ARP inspection. You can also configure other ARP parameters for both bridge groups and for routed mode interfaces.

Procedure

Step 1 Add static ARP entries according to Add a Static ARP Entry and Customize Other ARP Parameters, on page 765. ARP inspection compares ARP packets with static ARP entries in the ARP table, so static ARP entries are required for this feature. You can also configure other ARP parameters.

Step 2 Enable ARP inspection according to Enable ARP Inspection, on page 766.

Add a Static ARP Entry and Customize Other ARP Parameters

By default for bridge groups, all ARP packets are allowed between bridge group member interfaces. You can control the flow of ARP packets by enabling ARP inspection. 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 ASA only uses dynamic ARP entries in the ARP table for traffic to and from the ASA, such as management traffic.

You can also set the ARP timeout and other ARP behavior.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Advanced** > **ARP** > **ARP Static Table**.

Step 2 Click **Add** to add a static ARP entry.

The **Add ARP Static Configuration** dialog box appears.

a) Choose the interface attached to the host network, from the **Interface** drop-down list.
b) Enter the IP address of the host, in the **IP Address** field.
c) Enter the MAC address of the host, in the **MAC Address** field; for example, 00e0.1e4e.3d8b.
d) Check the **Proxy ARP** check box, to perform proxy ARP for this address.

If the ASA receives an ARP request for the specified IP address, then it responds with the specified MAC address.

e) Click **OK**.

Step 3 Enter a value in the **ARP Timeout** field, to set the ARP timeout for *dynamic* ARP entries.

This field sets the amount of time before the ASA rebuilds the ARP table, between 60 to 4294967 seconds. The default is 14400 seconds. Rebuilding the ARP table automatically updates new host information and removes old host information. You might want to reduce the timeout because the host information changes frequently.

Step 4 To allow non-connected subnets, check the **Allow non-connected subnets** check box. The ASA ARP cache only contains entries from directly-connected subnets by default. You can enable the ARP cache to also include non-directly-connected subnets. We do not recommend enabling this feature unless you know the security risks. This feature could facilitate denial of service (DoS) attack against the ASA; a user on any interface could send out many ARP replies and overload the ASA ARP table with false entries.

You may want to use this feature if you use:

- Secondary subnets.
- Proxy ARP on adjacent routes for traffic forwarding.

Step 5 Enter a value in the **ARP Rate-Limit** field to control the number of ARP packets per second on all interfaces.

Enter a value between 10 and 32768. The default value depends on your ASA model. You can customize this value to prevent an ARP storm attack.

Step 6 Click **Apply**.

Enable ARP Inspection

This section describes how to enable ARP inspection for bridge groups.

Procedure

Step 1 Choose the **Configuration** > **Device Management** > **Advanced** > **ARP** > **ARP Inspection** pane.

Step 2 Choose the interface row on which you want to enable ARP inspection, and click **Edit**.

The Edit ARP Inspection dialog box appears.

Step 3 Check the **Enable ARP Inspection** check box, to enable ARP inspection.

Step 4 (Optional) Check the **Flood ARP Packets** check box, to flood non-matching ARP packets.

By default, packets that do not match any element of a static ARP entry are flooded out all interfaces except the originating interface. If there is a mismatch between the MAC address, the IP address, or the interface, then the ASA drops the packet.

If you uncheck this check box, all non-matching packets are dropped, which restricts ARP through the ASA to only static entries.

Note

The Management 0/0 or 0/1 interface or subinterface, if present, never floods packets even if this parameter is set to flood.

Step 5 Click **OK**, and then **Apply**.

Customize the MAC Address Table for Bridge Groups

This section describes how you can customize the MAC address table for bridge groups.

Add a Static MAC Address for Bridge Groups

Normally, MAC addresses are added to the MAC address table dynamically as traffic from a particular MAC address enters an interface. You can 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 ASA drops the traffic and generates a system message. When you add a static ARP entry (see Add a Static ARP Entry and Customize Other ARP Parameters, on page 765), a static MAC address entry is automatically added to the MAC address table.

To add a static MAC address to the MAC address table, perform the following steps.

Procedure

Step 1 Choose the **Configuration** > **Device Setup** > **Bridging** > **MAC Address Table** pane.

Step 2 (Optional) Enter a value in the Dynamic Entry Timeout field, to set the time a MAC address entry stays in the MAC address table before timing out.

This value is between 5 and 720 minutes (12 hours). 5 minutes is the default.

Step 3 Click **Add**.

The Add MAC Address Entry dialog box appears.

Step 4 Choose the source interface associated with the MAC address, from the Interface Name drop-down list.

Step 5 Enter the MAC address, in the MAC Address field.

Step 6 Click **OK**, and then **Apply**.

Configure MAC Address Learning

By default, each interface automatically learns the MAC addresses of entering traffic, and the ASA adds corresponding entries to the MAC address table. You can disable MAC address learning if desired, however, unless you statically add MAC addresses to the table, no traffic can pass through the ASA. In routed mode, you can enable flooding of non-IP packets on all interfaces.

To configure MAC address learning, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Advanced** > **Bridging** > **MAC Learning** .

Step 2 To disable MAC learning, choose an interface row, and click **Disable**.

Step 3 To reenable MAC learning, click **Enable**.

Step 4 To enable flooding of non-IP packets, check **Enable flooding for unknown MAC address for non IPv4-IPv6 packets**.

Step 5 Click **Apply**.

History for ARP Inspection and the MAC Address Table

Feature Name	Platform Releases	Feature Information
ARP inspection	7.0(1)	ARP inspection compares the MAC address, IP address, and source interface in all ARP packets to static entries in the ARP table. This feature is available for Transparent Firewall Mode, and for interfaces in a bridge group in both Transparent and Routed modes starting in 9.7(1). We introduced the following commands: **arp**, **arp-inspection**, and **show arp-inspection**.
MAC address table	7.0(1)	You might want to customize the MAC address table for transparent mode, and for interfaces in a bridge group in both Transparent and Routed modes starting in 9.7(1). We introduced the following commands: **mac-address-table static**, **mac-address-table aging-time**, **mac-learn disable**, and **show mac-address-table**.
ARP cache additions for non-connected subnets	8.4(5)/9.1(2)	The ASA ARP cache only contains entries from directly-connected subnets by default. You can now enable the ARP cache to also include non-directly-connected subnets. We do not recommend enabling this feature unless you know the security risks. This feature could facilitate denial of service (DoS) attack against the ASA; a user on any interface could send out many ARP replies and overload the ASA ARP table with false entries. You may want to use this feature if you use: • Secondary subnets. • Proxy ARP on adjacent routes for traffic forwarding. We modified the following screen: Configuration > Device Management > Advanced > ARP > ARP Static Table.

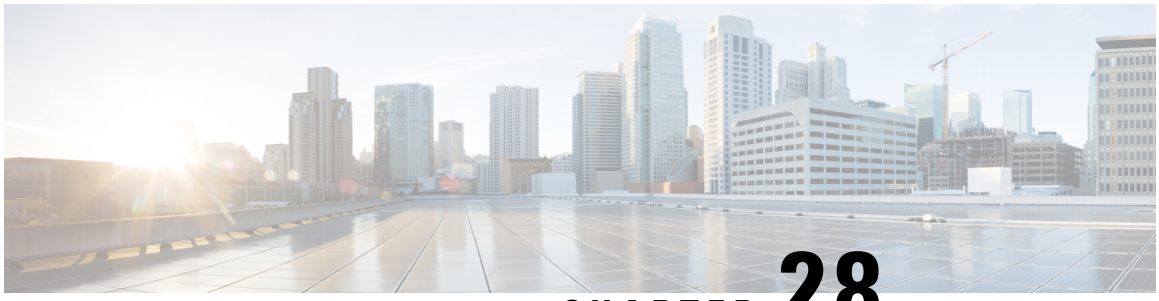
Feature Name	Platform Releases	Feature Information
Customizable ARP rate limiting	9.6(2)	You can set the maximum number of ARP packets allowed per second. The default value depends on your ASA model. You can customize this value to prevent an ARP storm attack. We modified the following screen: **Configuration** > **Device Management** > **Advanced** > **ARP** > **ARP Static Table**
Integrated Routing and Bridging	9.7(1)	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 ASA bridges instead of routes. The ASA is not a true bridge in that the ASA 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 ASA 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: multiple context mode, ASA clustering. The following features are also not supported on BVIs: dynamic routing and multicast routing. We modified the following screens: **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces** **Configuration** > **Device Setup** > **Routing** > **Static Routes** **Configuration** > **Device Management** > **DHCP** > **DHCP Server** **Configuration** > **Firewall** > **Access Rules** **Configuration** > **Firewall** > **EtherType Rules**



PART V

IP Routing

- Routing Overview, on page 773
- Static and Default Routes, on page 785
- Policy Based Routing, on page 793
- Route Maps, on page 803
- Bidirectional Forwarding Detection Routing, on page 813
- BGP, on page 821
- OSPF, on page 847
- IS-IS, on page 895
- EIGRP, on page 919
- Multicast Routing, on page 945



CHAPTER 28

Routing Overview

This chapter describes how routing behaves within the ASA.

- Path Determination, on page 773
- Supported Route Types, on page 774
- Supported Internet Protocols for Routing, on page 775
- Routing Table, on page 776
- Routing Table for Management Traffic, on page 781
- Equal-Cost Multi-Path (ECMP) Routing, on page 783
- Disable Proxy ARP Requests, on page 783
- Display the Routing Table, on page 784
- History for Route Overview, on page 784

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.

Note Asymmetric routing is only supported for Active/Active failover in multiple context mode.

Supported Route Types

There are several route types that a router can use. The ASA 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 ASA 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).

- Intermediate System to Intermediate System (IS-IS)

IS-IS is a link state Interior Gateway Protocol (IGP). Link-state protocols are characterized by the propagation of the information required to build a complete network connectivity map on each participating router. That map is then used to calculate the shortest path to destinations.

Routing Table

The ASA 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 781.

How the Routing Table Is Populated

The ASA routing table can be populated by statically defined routes, directly connected routes, and routes discovered by the dynamic routing protocols. Because the ASA 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 ASA 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 ASA 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 ASA 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 ASA.

Table 34: 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
Unknown	255

The smaller the administrative distance value, the more preference is given to the protocol. For example, if the ASA 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 ASA 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 ASA 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 ASA 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 ASA. 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 Failover

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 Failover 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 Failover bulk synchronization and continuous replication processes.

Dynamic Routing and Clustering

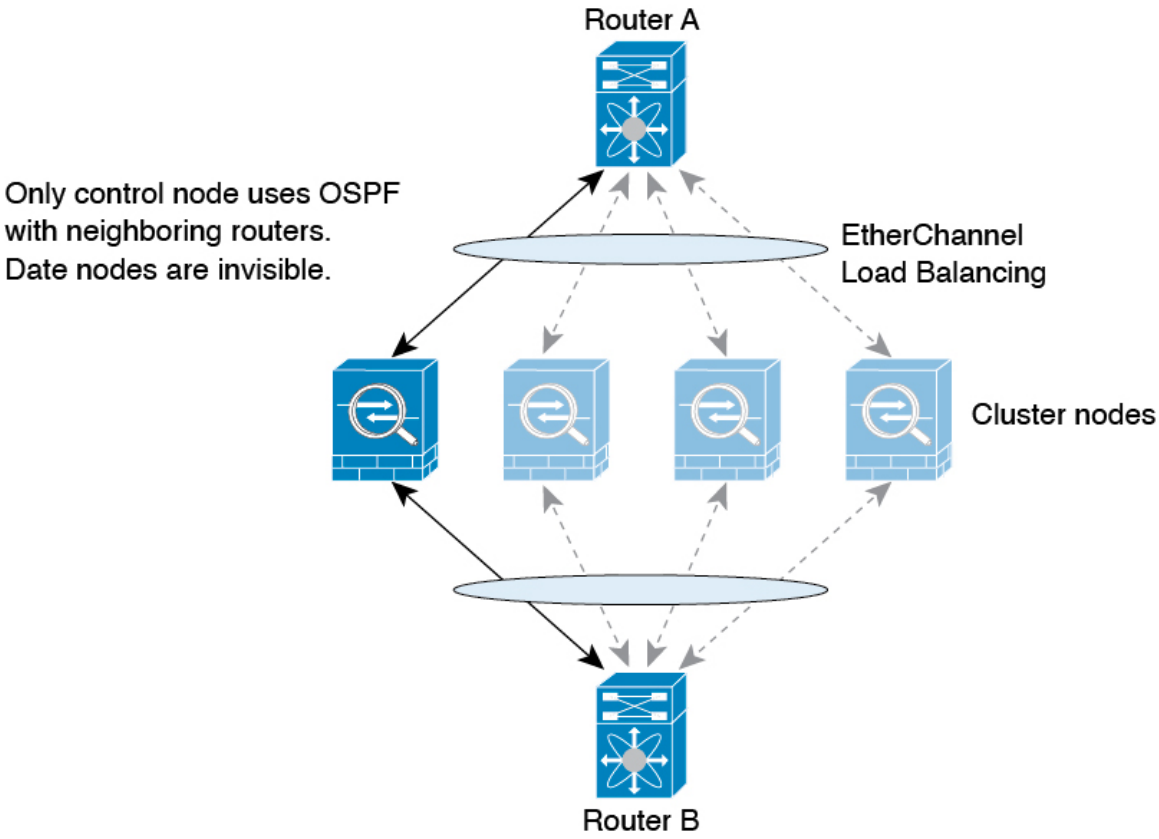
This section describes how to use dynamic routing with clustering.

Dynamic Routing in Spanned EtherChannel Mode

Note IS-IS is not supported in Spanned EtherChannel mode.

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 88: Dynamic Routing in Spanned EtherChannel Mode



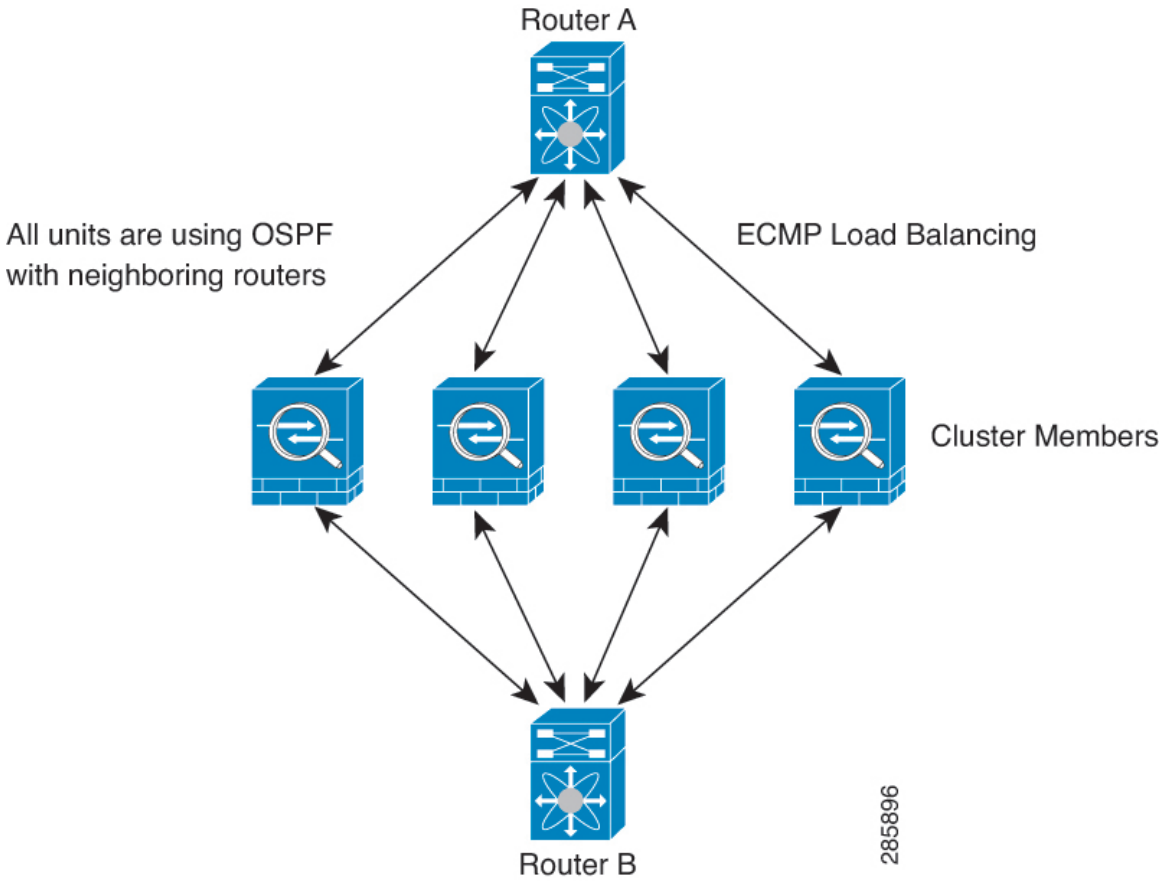
After the data node learn the routes from the control node, each node makes forwarding decisions independently.

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 89: 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 Configure a Traffic Zone, on page 684.

Dynamic Routing in Multiple Context Mode

In multiple context mode, each context maintains a separate routing table and routing protocol databases. This enables you to configure OSPFv2 and EIGRP independently in each context. You can configure EIGRP in some contexts and OSPFv2 in the same or different contexts. In mixed context mode, you can enable any of the dynamic routing protocols in contexts that are in routed mode. RIP and OSPFv3 are not supported in multiple context mode.

The following table lists the attributes for EIGRP, OSPFv2, route maps used for distributing routes into OSPFv2 and EIGRP processes, and prefix lists used in OSPFv2 to filter the routing updates entering or leaving an area when they are used in multiple context mode:

EIGRP	OSPFv2	Route Maps and Prefix Lists
One instance is supported per context.	Two instances are supported per context.	N/A
It is disabled in the system context.		N/A
Two contexts may use the same or different autonomous system numbers.	Two contexts may use the same or different area IDs.	N/A
Shared interfaces in two contexts may have multiple EIGRP instances running on them.	Shared interfaces in two contexts may have multiple OSPF instances running on them.	N/A
The interaction of EIGRP instances across shared interfaces is supported.	The interaction of OSPFv2 instances across shared interfaces is supported.	N/A
All CLIs that are available in single mode are also available in multiple context mode.		
Each CLI has an effect only in the context in which it is used.		

Route Resource Management

A resource class called *routes* specifies the maximum number of routing table entries that can exist in a context. This resolves the problem of one context affecting the available routing table entries in another context and also allows you greater control over the maximum route entries per context.

Because there is no definitive system limit, you can only specify an absolute value for this resource limit; you may not use a percentage limit. Also, there are no minimum and maximum limits per context, so the default class does not change. If you add a new route for any of the static or dynamic routing protocols (connected, static, OSPF, EIGRP, and RIP) in a context and the resource limit for that context is exhausted, then the route addition fails and a syslog message is generated.

Routing Table for Management Traffic

As a standard security practice, it is often necessary to segregate and isolate management (from-the-device) traffic from data traffic. To achieve this isolation, the ASA device uses a separate routing table for

management-only traffic vs. data traffic. Separate routing tables means that you can create separate default routes for data and management as well.

Types of Traffic for Each Routing Table

Through-the-device traffic always uses the data routing table.

From-the-device traffic, depending on the type, uses either the management-only routing table or the data routing table by default. If a match is not found in the default routing table, it checks the other routing table.

- Management-only table from-the-device traffic includes features that open a remote file using HTTP, SCP, TFTP, the **copy** command, Smart Licensing, Smart Call Home, **trustpoint**, **trustpool**, and so on.
- Data table from-the-device traffic includes all other features like ping, DNS, DHCP, and so on.

Interfaces Included in the Management-Only Routing Table

Management-only interfaces include any the Management x/x interfaces as well as any interfaces that you have configured to be management-only.

Fallback to the Other Routing Table

If a match is not found in the default routing table, it checks the other routing table.

Using the Non-Default Routing Table

If you need from-the-box traffic to go out an interface that isn't in its default routing table, then you might need to specify that interface when you configure it, rather than relying on the fall back to the other table. The ASA will only check routes for the specified interface. For example, if you need a ping to go out a management-only interface, then specify the interface in the ping function. Otherwise, if there is a default route in the data routing table, then it will match the default route and never fall back to the management routing table.

Dynamic Routing

The management-only routing table supports dynamic routing separate from the data interface routing table. A given dynamic routing process must run on either the management-only interface or the data interface; you cannot mix both types. When upgrading from an earlier release without the separate management routing table, if you have a mix of data and management interfaces using the same dynamic routing process, management interfaces will be dropped.

Management-Access Feature for VPN Requirements

If you configure the management-access feature that allows management access to an interface other than the one from which you entered the ASA when using VPN, then due to routing considerations with the separate management and data routing tables, the VPN termination interface and the management access interface need to be the same type: both need to be management-only interfaces or regular data interfaces.

Management Interface Identification

An interface configured with management-only is considered a management interface.

In the following configuration, both the interfaces GigabitEthernet0/0 and Management0/0 are considered as management interfaces.

```
a/admin(config-if)# show running-config int g0/0
!
interface GigabitEthernet0/0
 management-only
 nameif inside
```

```
 security-level 100
 ip address 10.10.10.123 255.255.255.0
 ipv6 address 123::123/64
a/admin(config-if)# show running-config int m0/0
!
interface Management0/0
 management-only
 nameif mgmt
 security-level 0
 ip address 10.106.167.118 255.255.255.0
a/admin(config-if)#
```

Equal-Cost Multi-Path (ECMP) Routing

The ASA 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 ASA 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.

Disable Proxy ARP Requests

When a host sends IP traffic to another device on the same Ethernet network, the host needs to know the MAC address of the device. ARP is a Layer 2 protocol that resolves an IP address to a MAC address. A host sends an ARP request asking “Who is this IP address?” The device owning the IP address replies, “I own that IP address; here is my MAC address.”

Proxy ARP is used when a device responds to an ARP request with its own MAC address, even though the device does not own the IP address. The ASA uses proxy ARP when you configure NAT and specify a mapped address that is on the same network as the ASA interface. The only way traffic can reach the hosts is if the ASA uses proxy ARP to claim that the MAC address is assigned to destination mapped addresses.

Under rare circumstances, you might want to disable proxy ARP for NAT addresses.

If you have a VPN client address pool that overlaps with an existing network, the ASA by default sends proxy ARP requests on all interfaces. If you have another interface that is on the same Layer 2 domain, it will see the ARP requests and will answer with the MAC address of its interface. The result of this is that the return traffic of the VPN clients towards the internal hosts will go to the wrong interface and will get dropped. In this case, you need to disable proxy ARP requests for the interface on which you do not want them.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Routing** > **Proxy ARP/Neighbor Discovery**.

The Interface field lists the interface names. The Enabled field shows whether or not proxy ARP/Neighbor Discovery is enabled (Yes) or disabled (No) for NAT global addresses.

Step 2 To enable proxy ARP/Neighbor Discovery for the selected interface, click **Enable**. By default, proxy ARP/Neighbor discovery is enabled for all interfaces.

Step 3 To disable proxy ARP/Neighbor Discovery for the selected interface, click **Disable**.

Step 4 Click **Apply** to save your settings to the running configuration.

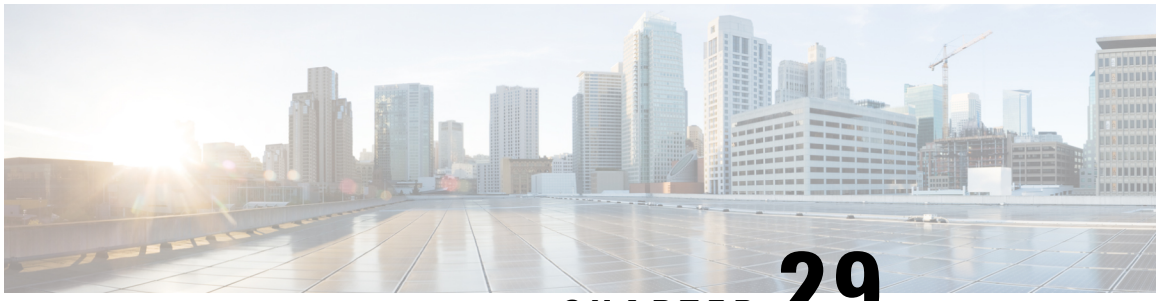
Display the Routing Table

To show all routes in ASDM that are in the routing table, choose **Monitoring** > **Routing** > **Routes**. Each row represents one route.

History for Route Overview

Table 35: History for Route Overview

Feature Name	Platform Releases	Feature Information
Routing Table for Management Interface	9.5(1)	To segregate and isolate, management traffic from data traffic, a separate table is added for management traffic. Separate routing tables, for manage data respectively, are created for both IPv4 and IPv6, for each context, of Further, for each context of the ASA, two extra routing tables are added i RIB and FIB. We updated the following screens:



CHAPTER 29

Static and Default Routes

This chapter describes how to configure static and default routes on the ASA.

- About Static and Default Routes, on page 785
- Guidelines for Static and Default Routes, on page 787
- Configure Default and Static Routes, on page 788
- Monitoring a Static or Default Route, on page 791
- Examples for Static or Default Routes, on page 791
- History for Static and Default Routes, on page 792

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 ASA 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.

Because the ASA device uses separate routing tables for data traffic and for management traffic, you can optionally configure a default route for data traffic and another default route for management traffic. Note that from-the-device traffic uses either the management-only or data routing table by default depending on the type, but will fall back to the other routing table if a route is not found. Default routes will always match traffic, and will prevent a fall back to the other routing table. In this case, you must specify the interface you want to use for egress traffic if that interface is not in the default routing table.

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 ASA.
- You are using a feature that does not support dynamic routing protocols.

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 783.
- 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 ASA 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 ASA knows out of which bridge group member interface to send traffic. Traffic that originates on the ASA 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 220 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 ASA 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 ASA implements static route tracking by associating a static route with a monitoring target host on the destination network that the ASA 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 ASA 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.

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 ASA does not support Class E routing, so a Class E network is not routable in static routes.

Clustering and Multiple Context Mode

- In clustering, static route tracking is only supported on the control node.
- Static route tracking is not supported in multiple context mode.

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.

Configure Default and Static Routes

At a minimum, you should configure a default route. You may need to configure static routes as well. In this section we will configure a default route, configure a static route and track a static route.

Configure a Default Route

A default route is simply a static route with 0.0.0.0/0 as the destination IP address. You should always have a default route, either configured manually with this procedure, or derived from a DHCP server or other routing protocol.

Before you begin

See the following guidelines for the Tunneled option:

- Do not enable unicast RPF on the egress interface of a tunneled route, because this setting causes the session to fail.
- Do not enable TCP intercept on the egress interface of the tunneled route, because this setting causes the session to fail.
- Do not use the VoIP inspection engines (CTIQBE, H.323, GTP, MGCP, RTSP, SIP, SKINNY), the DNS inspection engine, or the DCE RPC inspection engine with tunneled routes, because these inspection engines ignore the tunneled route.
- You cannot define more than one default route with the tunneled option.
- ECMP for tunneled traffic is not supported.
- Tunneled routes are not supported for bridge groups, which do not support VPN termination for through traffic.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Routing** > **Static Routes**, and click **Add**.

Step 2 Choose the **IP Address Type**, **IPv4** or **IPv6.**

Step 3 Choose the **Interface** through which you want to send the traffic.

For transparent mode, specify a bridge group member interface name. For routed mode with bridge groups, specify the BVI name.

Step 4 For the **Network**, type **any4** or **any6**, depending on the type.

Step 5 Enter the **Gateway IP** where you want to send the traffic.

Step 6 Set the **Metric** to set the administrative distance for the route.

The default is **1**. Administrative distance is a parameter 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 connect 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 routes take precedence. Connected routes always take precedence over static or dynamically discovered routes.

Step 7 (Optional) In the **Options** area, set the following:

- **Tunneled**—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 ASA that cannot be routed using learned or static routes, is sent to this route. This option is not supported for bridge groups.
- **Tracked**—(IPv4 only) For information about tracking routes, see Configure Static Route Tracking, on page 790.

Step 8 Click **OK**.

Configure a Static Route

A static route defines where to send traffic for specific destination networks.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Routing** > **Static Routes**, and click **Add**.

Step 2 Choose the **IP Address Type**, **IPv4** or **IPv6.**

Step 3 Choose the **Interface** through which you want to send the traffic.

To drop unwanted traffic, choose the **Null0** interface. For transparent mode, specify a bridge group member interface name. For routed mode with bridge groups, specify the BVI name.

Step 4 For the **Network**, enter the destination network for which you want to route traffic.

Step 5 Enter the **Gateway IP** where you want to send the traffic.

Step 6 Set the **Metric** to set the administrative distance for the route.

The default is **1**. Administrative distance is a parameter 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 connect 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 routes take precedence. Connected routes always take precedence over static or dynamically discovered routes.

Step 7 (Optional) In the **Options** area, set the following:

- **Tunneled**—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 ASA that cannot be routed using learned or static routes, is sent to this route.

- **Tracked**—(IPv4 only) For information about tracking routes, see Configure Static Route Tracking, on page 790.

Step 8 Click **OK**.

Configure Static Route Tracking

To configure static route tracking, complete the following steps.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Routing** > **Static Routes** and add or edit a static route according to Configure a Static Route, on page 789.

Step 2 Click the **Tracked** radio button in the **Options** area.

Step 3 In the **Track ID** field, enter a unique identifier for the route tracking process.

Step 4 In the **Track IP Address/DNS Name** field, enter the IP address or hostname of the target being tracked. Typically, this would be the IP address of the next hop gateway for the route, but it could be any network object available from that interface.

Step 5 In the **SLA ID** field, enter a unique identifier for the SLA monitoring process.

Step 6 (Optional) Click **Monitoring Options**.

The **Route Monitoring Options** dialog box appears. From here, change the following tracking object monitoring properties:

- **Frequency**—Sets how often, in seconds, the ASA should test for the presence of the tracking target. Valid values range from 1 to 604800 seconds. The default value is 60 seconds.
- **Threshold**—Sets the amount of time, in milliseconds, that indicates an over-threshold event. This value cannot be more than the timeout value.
- **Timeout**—Sets the amount of time, in milliseconds, that the route monitoring operation should wait for a response from the request packets. Valid values range from 0 to 604800000 milliseconds. The default value is 5000 milliseconds.
- **Data Size**—Sets the size of data payload to use in the echo request packets. The default value is 28. Valid values range from 0 to 16384.

  Note
  This setting specifies the size of the payload only; it does not specify the size of the entire packet.

- **ToS**—Sets a value for the type of service byte in the IP header of the echo request. Valid values are from 0 to 255. The default value is 0.
- **Number of Packets**—Sets the number of echo requests to send for each test. Valid values range from 1 to 100. The default value is 1.

Click **OK**.

Step 7 Click **OK** to save your route, and then click **Apply**.

The monitoring process begins as soon as you apply the tracked route.

Step 8 Create an untracked backup route.

The backup route is a static route to the same destination as the tracked route, but through a different interface or gateway. You must assign this route a higher administrative distance (metric) than your tracked route.

Monitoring a Static or Default Route

- **Monitoring > Routing > Routes**.

  In the **Routes** pane, each row represents one route. You can filter by IPv4 connections, IPv6 connections, or both. The routing information includes the protocol, the route type, the destination IP address, the netmask or prefix length, the gateway IP address, the interface through which the route is connected, and the administrative distance.

Examples for Static or Default Routes

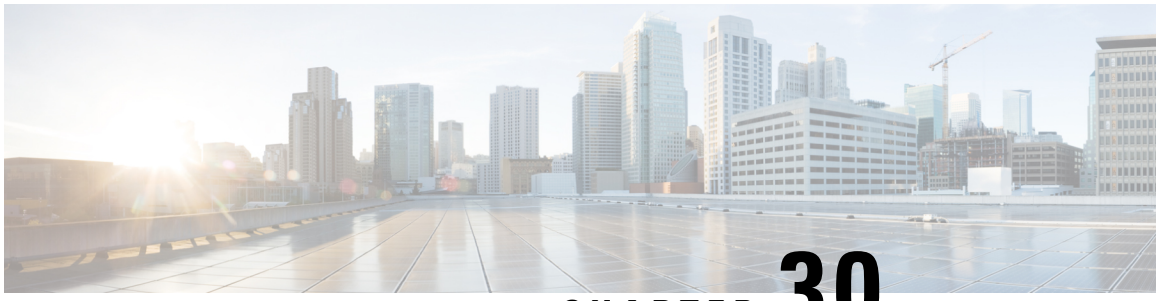
The following example shows how to create a static route that sends all traffic destined for 10.1.1.0/24 to the router 10.1.2.45, which is connected to the inside interface, defines three equal cost static routes that direct traffic to three different gateways on the dmz interface, and adds a default route for tunneled traffic and one for regular traffic.

```
route inside 10.1.1.0 255.255.255.0 10.1.2.45
route dmz 10.10.10.0 255.255.255.0 192.168.2.1
route dmz 10.10.10.0 255.255.255.0 192.168.2.2
route dmz 10.10.10.0 255.255.255.0 192.168.2.3
route outside 0 0 209.165.201.1
route inside 0 0 10.1.2.45 tunneled
```

History for Static and Default Routes

Table 36: Feature History for Static and Default Routes

Feature Name	Platform Releases	Feature Information
Static Route Tracking	7.2(1)	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. We introduced or modified the following screens: **Configuration > Device Setup > Routing > Static Routes > Add Static Route Configuration > Device Setup > Routing > Static Routes > Add Static Route > Route Monitoring Options**
Static null0 route to drop traffic	9.2(1)	Sending traffic to a null0 interface results in dropping the packets destined to the specified network. This feature is useful in configuring Remotely Triggered Black Hole (RTBH) for BGP. We modified the following screen: **Configuration > Device Setup > Routing > Static Routes > Add Static Route**



CHAPTER 30

Policy Based Routing

This chapter describes how to configure the ASA to support policy based routing (PBR). The following sections describe policy based routing, guidelines for PBR, and configuration for PBR.

- About Policy Based Routing, on page 793
- Guidelines for Policy Based Routing, on page 795
- Path Monitoring, on page 797
- Configure Policy Based Routing, on page 798
- History for Policy Based Routing, on page 801

About Policy Based Routing

Traditional routing is destination-based, meaning packets are routed based on destination IP address. However, it is difficult to change the routing of specific traffic in a destination-based routing system. With Policy Based Routing (PBR), you can define routing based on criteria other than destination network—PBR lets you route traffic based on source address, source port, destination address, destination port, protocol, or a combination of these.

Policy Based Routing:

- Lets you provide Quality of Service (QoS) to differentiated traffic.
- Lets you distribute interactive and batch traffic across low-bandwidth, low-cost permanent paths and high-bandwidth, high-cost switched paths.
- Allows Internet service providers and other organizations to route traffic originating from various sets of users through well-defined Internet connections.

Policy Based Routing can implement QoS by classifying and marking traffic at the network edge, and then using PBR throughout the network to route marked traffic along a specific path. This permits routing of packets originating from different sources to different networks, even when the destinations are the same, and it can be useful when interconnecting several private networks.

Why Use Policy Based Routing?

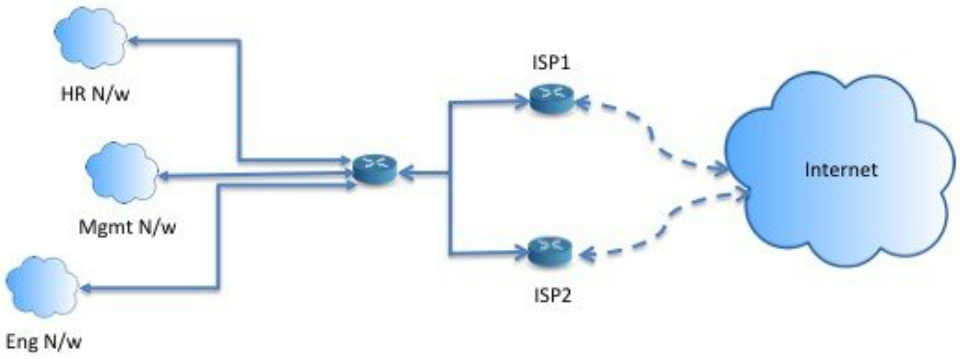
Consider a company that has two links between locations: one a high-bandwidth, low-delay expensive link, and the other a low-bandwidth, higher-delay, less-expensive link. While using traditional routing protocols, the higher-bandwidth link would get most, if not all, of the traffic sent across it based on the metric savings

obtained by the bandwidth and/or delay (using EIGRP or OSPF) characteristics of the link. PBR allows you to route higher priority traffic over the high-bandwidth/low-delay link, while sending all other traffic over the low-bandwidth/high-delay link.

Some applications of policy based routing are:

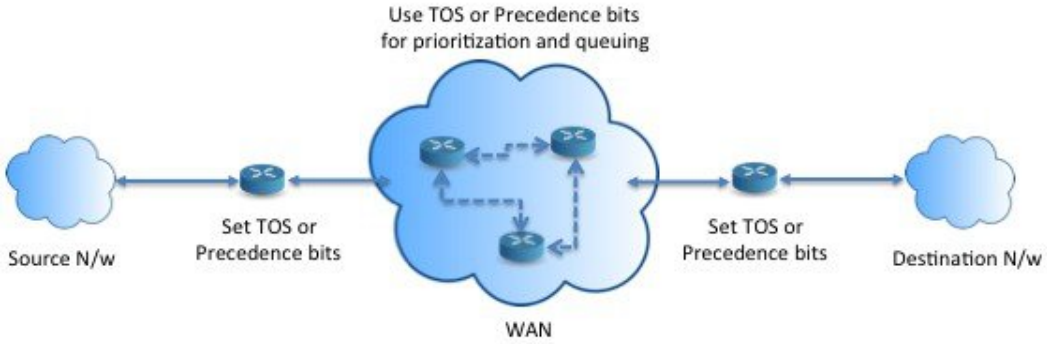
Equal-Access and Source-Sensitive Routing

In this topology, traffic from HR network & Mgmt network 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.



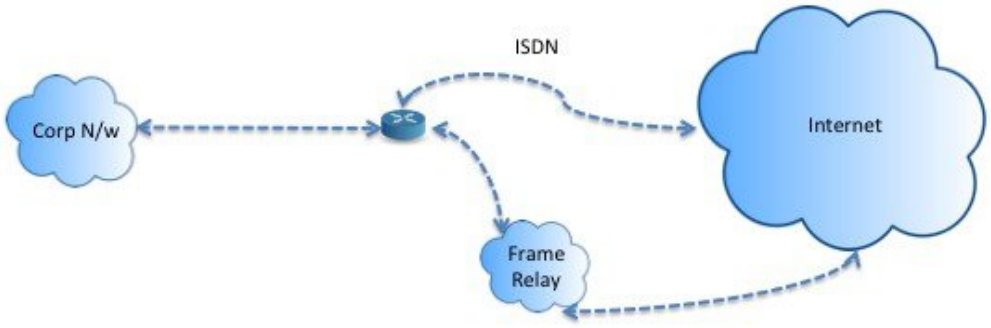
Quality of Service

By tagging packets with policy based routing, network administrators can classify the network traffic at the perimeter of the network for various classes of service and then implementing those classes of service in the core of the network using priority, custom or weighted fair queuing (as shown in the figure below). This setup improves network performance by eliminating the need to classify the traffic explicitly at each WAN interface in the core of backbone network.



Cost Saving

An organization can direct the bulk traffic associated with a specific activity to use a higher-bandwidth high-cost link for a short time and continues basic connectivity over a lower-bandwidth low-cost link for interactive traffic by defining the topology, as show 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 load share the traffic from HR network through ISP1 and traffic from Eng network through ISP2.

Implementation of PBR

The ASA uses ACLs to match traffic and then perform routing actions on the traffic. Specifically, you configure a route map that specifies an ACL for matching, and then you specify one or more actions for that traffic. Finally, you associate the route map with an interface where you want to apply PBR on all incoming traffic.

Note 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 asymetirc routing, specifically when NAT and VPN are in use.

Guidelines for Policy Based Routing

Firewall Mode

Supported only in routed firewall mode. Transparent firewall mode is not supported.

Per-flow Routing

Since the ASA performs routing on a per-flow basis, policy routing is applied on the first packet and the resulting routing decision is stored in the flow created for the packet. All subsequent packets belonging to the same connection simply match this flow and are routed appropriately.

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.

Clustering

- Clustering is supported.
- In a cluster scenario, without static or dynamic routes, with ip-verify-reverse path enabled, asymmetric traffic may get dropped. So disabling ip-verify-reverse path is recommended.

IPv6 Support

IPv6 is supported

Path Monitoring Guidelines

Following are the guidelines for configuring the path monitoring on the interfaces:

- Interfaces must have an interface name.
- Management-only interfaces cannot be configured with the path monitoring. To configure the path monitoring, you must uncheck the **Dedicate this interface to management only** check box.
- Path monitoring is not supported on devices in Transparent or multicontext system mode.
- Auto monitoring types (auto, auto4, and auto6) are not supported for Tunnel interfaces.
- Path monitoring cannot be configured for the following interfaces:
  - BVI
  - Loopback
  - DVTI

Additional Guidelines

- All existing route map related configuration restrictions and limitations will be carried forward.
- Do not use route maps containing match policy lists for policy based routing. The match policy-list is only used for BGP.

- 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

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.

The metrics on the interfaces are collected dynamically using ICMP probe messages to the interface's default gateway or a specified remote peer.

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.

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, on page 798.

Configure Path Monitoring

You can configure path monitoring to perform Policy Based Routing based on the network service groups. To use path monitoring without NSG, you can navigate to the **Interface** > **Edit** page and specify the path monitoring type. See Configure Policy Based Routing.

Procedure

Step 1 In ASDM, choose **Configuration** > **Device Setup** > **Interface Settings** > **Path Monitoring**.

Step 2 Select interface from **Interface** drop-down.

Step 3 Select the network service group (NSG) in the **Available Network Service Groups** box. To select multiple NSGs, use the control key and click on the required NSGs.

Step 4 Click **Add** to add the Network Service Groups.

Step 5 Click **Apply**.

Step 6 To remove the configuration, select the NSGs from the **Added Network Service Groups** box and click **Remove**, and then click **Apply**.

Configure Policy Based Routing

A route map is comprised of one or more route-map statements. Each statement has a sequence number, as well as a permit or deny clause. Each route-map statement contains match and set commands. The match command denotes the match criteria to be applied on the packet. The set command denotes the action to be taken on the packet.

- When a route map is configured with both IPv4 and IPv6 match/set clauses or when a unified ACL matching IPv4 and IPv6 traffic is used, the set actions will be applied based on destination IP version.
- When multiple next-hops or interfaces are configured as a set action, all options are evaluated one after the other until a valid usable option is found. No load balancing will be done among the configured multiple options.
- The verify-availability option is not supported in multiple context mode.

Procedure

Step 1 In ASDM, configure one or more standard or extended ACLs to identify traffic on which you want to perform Policy Based Routing. See **Configuration** > **Firewall** > **Advanced** > **ACL Manager** .

Step 2 Choose **Configuration** > **Device Setup** > **Routing** > **Route Maps**, and click **Add**.

The **Add Route Map** dialog box appears.

Step 3 Enter the route map name and sequence number. You will use this same name for optional additional route map statements. The sequence number is the order in which the ASA assesses the route maps.

Step 4 Click **Deny** or **Permit**.

The ACL also includes its own permit and deny statements. For Permit/Permit matches between the route map and the ACL, the Policy Based Routing processing continues. For Permit/Deny matches, processing ends for this route map, and other route maps are checked. If the result is still Permit/Deny, then the regular routing table is used. For Deny/Deny matches, the Policy Based Routing processing continues.

Step 5 Click the **Match Clause** tab to identify the ACLs you created.

In the **IPv4** section, choose **Access List** from the drop-down menu, and then select one or more standard or extended ACLs from the dialog box.

Note
Ensure that the access list does not contain any inactive rules. You cannot set match ACL with inactive rules to a PBR.

If you use a standard ACL, matching is done on the destination address only. If you use an extended ACL, you can match on source, destination, or both.

Use the IPv4 section for both IPv4 and IPv6 ACLs. For the extended ACL, you can specify IPv4, IPv6, Identity Firewall, or Cisco TrustSec parameters. You can also include network-service objects. For complete syntax, see the ASA command reference.

Step 6 Click the **Policy Based Routing** tab to define policy for traffic flows.

Check one or more of the following set actions to perform for the matching traffic flows:

- **Set PBR next hop address**—For IPv4 and IPv6, you can configure multiple next-hop IP addresses in which case they are evaluated in the specified order until a valid routable next-hop IP address is found. The configured next-hops should be directly connected; otherwise the set action will not be applied.
- **Set default next-hop IP address**—For IPv4 and IPv6, if the normal route lookup fails for matching traffic, then the ASA forwards the traffic using this specified next-hop IP address.
- **Recursively find and set next-hop IP address—**Both the next-hop address and the default next-hop address require that the next-hop be found on a directly connected subnet. With this option, the next-hop address does not need to be directly connected. Instead 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 routing path in use on the router.
- **Configure Next Hop Verifiability**—Verify if the next IPv4 hops of a route map are available. You can configure an SLA monitor tracking object to verify the reachability of the next-hop. Click **Add** to add next-hop IP address entries, and specify the following information.
  - **Sequence Number**—Entries are assessed in order using the sequence number.
  - **IP Address**—Enter the next hop IP address.
  - **Tracking Object ID**—Enter a valid ID.
- **Set interfaces**—This option configures the interface through which the matching traffic is forwarded. You can configure multiple interfaces, in which case they are evaluated in the specified order until a valid interface is found. When you specify **null0**, all traffic matching the route map will be dropped. There must be a route for the destination that can be routed through the specified interface (either static or dynamic).
- **Set Clause** > **Adaptive Interface Cost**—This option is on the Set Clause tab rather than the **Policy Based Routing** tab. This option sets the output interface based on the interface's cost. Click the **Available Interfaces** field and select the interfaces that should be considered. The egress interface is selected from

the list of interfaces. If the costs of the interfaces are the same, it is an active-active configuration and packets are load-balanced (round-robin) on the egress interfaces. If the costs are different, the interface with the lowest cost is selected. Interfaces are considered only if they are up.

- **Set null0 interface as the default interface**—If a normal route lookup fails, the ASA forwards the traffic null0, and the traffic will be dropped.
- **Set do-not-fragment bit to either 1or 0**—Select the appropriate radio button.
- **Set differential service code point (DSCP) value in QoS bits**—Select a value from the IPv4 or IPv6 drop-down list.

Step 7 Click **OK**, and then click **Apply**.

Step 8 Choose **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces**, and configure the ingress interfaces that should apply this route map to determine the egress interfaces.

a) Select an ingress interface and click **Edit**.
b) In **Route Map**, select the policy-based route map that should be applied.
c) If you used **Adaptive Interface Cost** to select the output interface in the route map, set the **Cost** value on the interface.

The value can be 1-65535. The default is 0, which you can reset by deleting the value from this field. The lower the number, the higher the priority. For example, 1 has priority over 2.

d) For PBR to use flexible metrics in identifying the best path for routing packets, from the **Path Monitoring** drop-down list, select the relevant monitoring type:

- **auto**—Sends ICMP probes to the IPv4 default gateway of the interface, if it exists (same as Auto IPv4). Else, sends to the IPv6 default gateway of the interface (same as Auto IPv6).
- **ipv4**—Sends ICMP probes to the specified peer IPv4 address (next-hop IP) for monitoring. If you select this option, the adjacent field is enabled. Enter the IPv4 address in the field.
- **ipv6**—Sends ICMP probes to the specified peer IPv4 address (next-hop IP) for monitoring. If you select this option, the adjacent field is enabled. Enter the IPv4 address in the field.
- **auto4**—Sends ICMP probes to the IPv4 default gateway of the interface.
- **auto6**—Send ICMP probes to the IPv6 default gateway of the interface.
- **None**—To disable path monitoring for the interface.

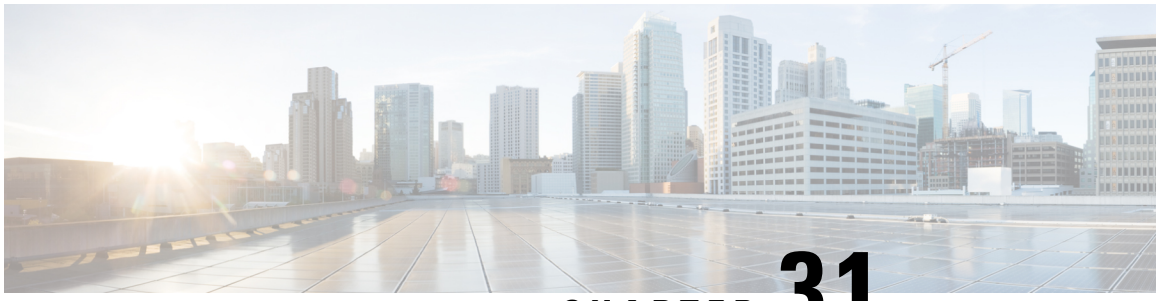
e) Click **OK**, then **Apply**.

History for Policy Based Routing

Table 37: History for Route Maps

Feature Name	Platform Releases	Feature Information
Path monitoring through HTTP client	9.20(1)	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 path-monitoring can be configured on the interface using Network Service Group objects. New/Modified screens: **Configuration** > **Device Setup** > **Interface Settings** > **Path Monitoring**
Path monitoring metrics in PBR.	9.18(1)	PBR uses the metrics to determine the best path (egress interface) for forwarding the traffic. Path monitoring periodically notifies PBR with 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. New/Modified screens: **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces**
Policy based routing	9.4(1)	Policy Based Routing (PBR) is a mechanism by which traffic is routed through specific paths with a specified QoS using ACLs. ACLs let traffic be classified based on the content of the packet's Layer 3 and Layer 4 headers. This solution lets administrators provide QoS to differentiated traffic, distribute interactive and batch traffic among low-bandwidth, low-cost permanent paths and high-bandwidth, high-cost switched paths, and allows Internet service providers and other organizations to route traffic originating from various sets of users through well-defined Internet connections. We updated the following screens: Configuration > Device Setup > Routing > Route Maps > Policy Based Routing, Configuration > Device Setup > Routing > Interface Settings > Interfaces
IPv6 support for Policy Based Routing	9.5(1)	IPv6 addresses are now supported for Policy Based Routing. We modified the following screens: **Configuration > Device Setup > Routing > Route Maps > Add Route Map > Policy Based Routing** **Configuration > Device Setup > Routing > Route Maps > Add Route Maps > Match Clause**

Feature Name	Platform Releases	Feature Information
VXLAN support for Policy Based Routing	9.5(1)	You can now enable Policy Based Routing on a VNI interface. We modified the following screen: **Configuration > Device Setup > Interface Settings > Interfaces > Add/Edit Interface > General**
Policy Based Routing support for Identity Firewall and Cisco Trustsec	9.5(1)	You can configure Identity Firewall and Cisco TrustSec and then use Identity Firewall and Cisco TrustSec ACLs in Policy Based Routing route maps. We modified the following screen: **Configuration > Device Setup > Routing > Route Maps > Add Route Maps > Match Clause**



CHAPTER 31

Route Maps

This chapter describes how to configure and customize route-maps, for ASA.

- About Route Maps, on page 803
- Guidelines for Route Maps, on page 805
- Define a Route Map, on page 805
- Customize a Route Map, on page 808
- Example for Route Maps, on page 810
- History for Route Maps, on page 811

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).

Guidelines for Route Maps

Firewall Mode

Supported only in routed firewall mode. Transparent firewall mode is not supported.

Additional Guidelines

Route maps do not support ACLs that include a user, user group, or fully qualified domain name objects.

Define a Route Map

You must define a route map when specifying which of the routes from the specified routing protocol are allowed to be redistributed into the target routing process. In ASDM, you can define a route map by adding, editing, or deleting a route map name, sequence number, or redistribution.

Procedure

Step 1 In ASDM, choose **Configuration** > **Device Setup** > **Routing** > **Route Maps**.

Step 2 Click **Add**.

The **Add Route Map** or **Edit Route Map** dialog box appears.

Step 3 Enter the route map name and sequence number. The route map name is the name that you assign to a particular route. The sequence number is the order in which you add or delete the route map entries into the ASA.

Note
If you are editing an existing route map, the fields for Route Map name and sequence number are already filled in.

Step 4 To reject route matches from redistribution, click **Deny**. If you use an ACL in a route map Deny clause, routes that are permitted by the ACL are not redistributed. To allow route matches for redistribution. click **Permit**. If you use an ACL in a route map Permit clause, routes that are permitted by the ACL are redistributed.

In addition, if you use an ACL in a route map Permit or Deny clause, and the ACL denies a route, then the route map clause match is not found and the next route map clause is evaluated.

Step 5 Click the **Match Clause** tab to choose routes to which this clause should be applied, and set the following parameters:

- Check the **Match first hop interface of route** check box to enable or disable matching the first hop interface of a route or to match any routes with the specified next hop interface. If you specify more than one interface, then the route can match either interface.
  - Enter the interface name in the Interface field, or click the ellipses to display the Browse Interface dialog box.
  - Choose one or more interfaces, click **Interface**, then click **OK**.

- In the IPv4 and IPv6 sections, do one or more of the following:
  - Check the **Match Address** check box to enable or disable the Match address of a route or match packet.
  - Check the **Match Next Hop** check box to enable or disable the Match next hop address of a route.
  - Check the **Match Route Source** check box to enable or disable the Match advertising source address of the route.
  - Choose **Access List to Prefix List** from the drop-down list to match the IP address.
  - According to the previous selection, click the ellipses to display the **Browse Access List** or **Browse Prefix List** dialog box.

    Note
    Prefix lists are not supported in OSPF.

  - Choose the ACL or prefix list that you want.
- Check the **Match metric of route** check box to enable or disable matching the metric of a route.
  - In the **Metric Value** field, type the metric values. You can enter multiple values, separated by commas. This setting allows you to match any routes that have a specified metric. The metric value can range from 0 to 4294967295.
- Check the **Match Route Type** check box to enable or disable matching of the route type. Valid route types are External1, External2, Internal, Local, NSSA-External1, and NSSA-External2. When enabled, you can choose more than one route type from the list.

Step 6 Click the **Set Clause** tab to modify the following information, which will be redistributed to the target protocol:

- Check the **Set Metric Clause** check box to enable or disable the metric value for the destination routing protocol, and type the value in the Value field.
- Check the **Set Metric Type** check box to enable or disable the type of metric for the destination routing protocol, and choose the metric type from the drop-down list.
- **Adaptive Interface Metric Type**—This option relates to Policy Based Routing. This option sets the output interface based on metric values collected on the interfaces, namely, cost, Round Trip Time (RTT), jitter, Mean Opinion Score (MOS), and lost (packet loss).
- Click the **Available Interfaces** field and select the interfaces that should be used for routing. The egress interface is selected from the list of interfaces. If the costs of the interfaces are the same, it is an active-active configuration and packets are load-balanced (round-robin) on the egress interfaces. If the costs are different, the interface with the lowest cost is selected. Similar to the cost metric, the other values are applied based on the metric type, minimal jitter, minimal RTT, minimal packet loss, and maximum MOS. Interfaces are considered only if they are up.

Step 7 Click the BGP Match Clause tab to choose routes to which this clause should be applied, and set the following parameters:

- Check the **Match AS path access lists** check box 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.

- Check the **Match Community** check box to enable matching the BGP community with the specified community. If you specify more than one community, then the route can match either community. Any route that does not match at least one Match community will not be advertised for outbound route maps.
  - Check the Match the specified community exactly check box to enable matching the BGP community exactly with the specified community.
- Check the **Match Policy list** check box to configure a route map to evaluate and process a BGP policy. If you specify more than one policy list, then the route can process either policy list.

Step 8 Click the BGP **Set Clause** tab to modify the following information, which will be redistributed to the BGP protocol:

- Check the **Set AS Path** check box to modify an autonomous system path for BGP routes.
  - Check the Prepend AS path check box 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 numbers.
  - Check the Prepend Last AS to the AS Path check box to prepend the AS path with the last AS number. Enter a value for the AS number from 1 to 10.
  - Check the Convert route tag into AS Path check box to convert the tag of a route into an autonomous system path.
- Check the **Set Community check box to set the BGP communities attributes.**
  - Click Specify Community to enter a community number, if applicable. Valid values are from 1 to 4294967200, internet, no-advertise and no-export.
  - Check Add to the existing communities to add the community to the already existing communities.
  - Click None to remove the community attribute from the prefixes that pass the route map.
- Check the **Set local preference** check box to specify a preference value for the autonomous system path.
- Check the Set weight check box to specify the BGP weight for the routing table. Enter a value between 0 and 65535.
- Check the Set origin check box to specify the BGP origin code. Valid values are Local IGP and Incomplete.
- Check the Set next hop check box to specify the output address of packets that fulfill the match clause of a route map.
  - Click Specify IP address to enter the IP address of the next hop to which packets are output. It need not be an adjacent router. If you specify more than one IP address then the packets can output at either IP address.
  - Click Use peer address to set the next hop to be the BGP peer address.

Step 9 Click OK.

Customize a Route Map

This section describes how to customize the route map.

Define a Route to Match a Specific Destination Address

Procedure

Step 1 In ASDM, choose **Configuration** > **Device Setup** > **Routing** > **Route Maps**.

Step 2 Click **Add**.

The **Add Route Map** dialog box appears. From this dialog box you can assign or choose the route map name, the sequence number and its redistribution access (that is, permit or deny). Route map entries are read in order. You can identify the order using the sequence number, or the ASA uses the order in which you add the entries.

Step 3 Click the **Match Clause** tab to choose routes to which this clause should be applied, and set the following parameters:

- Check the **Match first hop interface of route** check box to enable or disable matching the first hop interface of a route or to match any routes with the specified next hop interface. If you specify more than one interface, then the route can match either interface.
  - Enter the interface name in the **Interface** field, or click the ellipses to display the **Browse Interface** dialog box.
  - Choose the interface type (**inside** or **outside**), click **Selected Interface**, then click **OK**.
  - Check the **Match IP Address** check box to enable or disable the Match address of a route or match packet.
  - Check the **Match Next Hop** check box to enable or disable the Match next hop address of a route.
  - Check the **Match Route Source** check box to enable or disable the Match advertising source address of the route.
  - Choose **Access List to Prefix List** from the drop-down list to match the IP address.
  - According to the previous selection, click the ellipses to display the **Browse Access List** or **Browse Prefix List** dialog box.

    Note
    Prefix lists are not supported in OSPF.

  - Choose the ACL or prefix list that you want.
- Check the **Match metric of route** check box to enable or disable matching the metric of a route.
  - In the **Metric Value** field, type the metric values. You can enter multiple values, separated by commas. This setting allows you to match any routes that have a specified metric. The metric value can range from 0 to 4294967295.

- Check the **Match Route Type** check box to enable or disable matching of the route type. Valid route types are External1, External2, Internal, Local, NSSA-External1, and NSSA-External2. When enabled, you can choose more than one route type from the list.

Configure Prefix Rules

Note You must configure a prefix list before you may configure a prefix rule.

To configure prefix rules, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Routing** > **IPv4 Prefix Rules** or **IPv6 Prefix Rules**.

Step 2 Click **Add** and choose Add Prefix Rule.

The **Add Prefix Rule** dialog box appears. From this dialog box, you can add a sequence number, select an IP version- IPv4 or IPv6, specify a prefix for the network, its redistribution access (that is, permit or deny) and the minimum and maximum prefix length.

Step 3 Enter an optional **Sequence Number** or accept the default value.

Step 4 Specify the **Prefix** number in the format of IP address/mask length.

Step 5 Click the **Permit** or **Deny** radio button to indicate the redistribution access.

Step 6 Enter the optional **Minimum length** and **Maximum length**.

Step 7 Click **OK** when you are done.

The new or revised prefix rule appears in the list.

Step 8 Click **Apply** to save your changes.

Configure Prefix Lists

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.

Note Prefix lists are not supported with OSPF.

To add prefix lists, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Routing** > **IPv4 Prefix Rules** or **IPv6 Prefix Rules**.

Step 2 Click **Add** > **Add Prefix List**.

The **Add Prefix List** dialog box appears.

Step 3 Enter the prefix name and description, then click **OK**.

Configure the Metric Values for a Route Action

To configure the metric value for a route action, perform the following steps:

Procedure

Step 1 In ASDM, choose **Configuration** > **Device Setup** > **Routing** > **Route Maps**.

Step 2 Click **Add**.

The **Add Route Map** or **Edit Route Map** dialog box appears. From this dialog box, you can assign or select the route map name, the sequence number and its redistribution access (that is, permit or deny). Route map entries are read in order. You can identify the order using the sequence number, or the ASA uses the order in which you add route map entries.

Step 3 Click the **Set Clause** tab to modify the following information, which will be redistributed to the target protocol:

- Check the **Set Metric Clause** check box to enable or disable the metric value for the destination routing protocol, and enter the value in the Value field.
- Check the **Set Metric Type** check box to enable or disable the type of metric for the destination routing protocol, and choose the metric type from the drop-down list.

Example for Route Maps

The following example shows how to redistribute routes with a hop count equal to 1 into OSPF.

1. In ASDM, choose **Configuration** > **Device Setup** > **Routing** > **Route Maps**.
2. Click **Add**.
3. Enter **1-to-2** in the **Route Map Name** field.
4. Enter the routing sequence number in the **Sequence Number** field.
5. Click the **Permit** radio button.

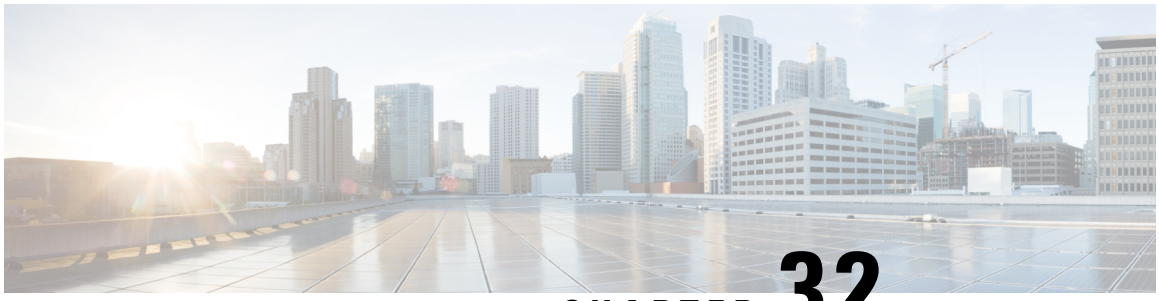
By default, this tab is on top.

6. Click the **Match Clause** tab.
7. Check the **Match Metric of Route** check box and type **1** for the metric value.
8. Click the **Set Clause** tab.
9. Check the **Set Metric Value** check box, and type **5** for the metric value.
10. Check the **Set Metric-Type** check box, and choose **Type-1**.

History for Route Maps

Table 38: Feature History for Route Maps

Feature Name	Platform Releases	Feature Information
Route maps	7.0(1)	We introduced this feature. We introduced the following screen: Configuration > Device Setup > Routing > Route Maps.
Enhanced support for static and dynamic route maps	8.0(2)	Enhanced support for dynamic and static route maps was added.
Dynamic Routing in Multiple Context Mode	9.0(1)	Route maps are supported in multiple context mode.
Support for BGP	9.2(1)	We introduced this feature. We updated the following screen: Configuration > Device Setup > Routing > Route Maps with 2 additional tabs BGP match clause and BGP set clause.
IPv6 support for Prefix Rule	9.3.2	We introduced this feature. We updated the following screens: **Configuration** > **Device Setup** > **Routing** > **IPv4 Prefix Rules** and **IPv6 Prefix Rules**



CHAPTER 32

Bidirectional Forwarding Detection Routing

This chapter describes how to configure the ASA to use the Bidirectional Forwarding Detection (BFD) routing protocol.

- About BFD Routing, on page 813
- Guidelines for BFD Routing, on page 816
- Configure BFD, on page 817
- History for BFD Routing, on page 820

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.

BFD Asynchronous Mode and Echo Function

BFD can operate in asynchronous mode with or without the echo function enabled.

Asynchronous Mode

In asynchronous mode, the systems periodically send BFD control packets to one another, and if a number of those packets in a row are not received by the other system, the session is declared to be down. Pure asynchronous mode (without the Echo function) is useful because it requires half as many packets to achieve a particular detection time as the Echo function requires.

BFD Echo Function

The BFD echo function sends echo packets from the forwarding engine to the directly-connected single-hop BFD neighbor. The echo packets are sent by the forwarding engine and forwarded back along the same path to perform detection. The BFD session at the other end does not participate in the actual forwarding of the echo packets. Because the echo function and the forwarding engine are responsible for the detection process, the number of BFD control packets that are sent out between BFD neighbors is reduced. And

also because the forwarding engine is testing the forwarding path on the remote neighbor system without involving the remote system, the inter-packet delay variance is improved. This results in quicker failure detection times.

When the echo function is enabled, BFD can use the slow timer to slow down the asynchronous session and reduce the number of BFD control packets that are sent between BFD neighbors, which reduces processing overhead while at the same time delivering faster failure detection.

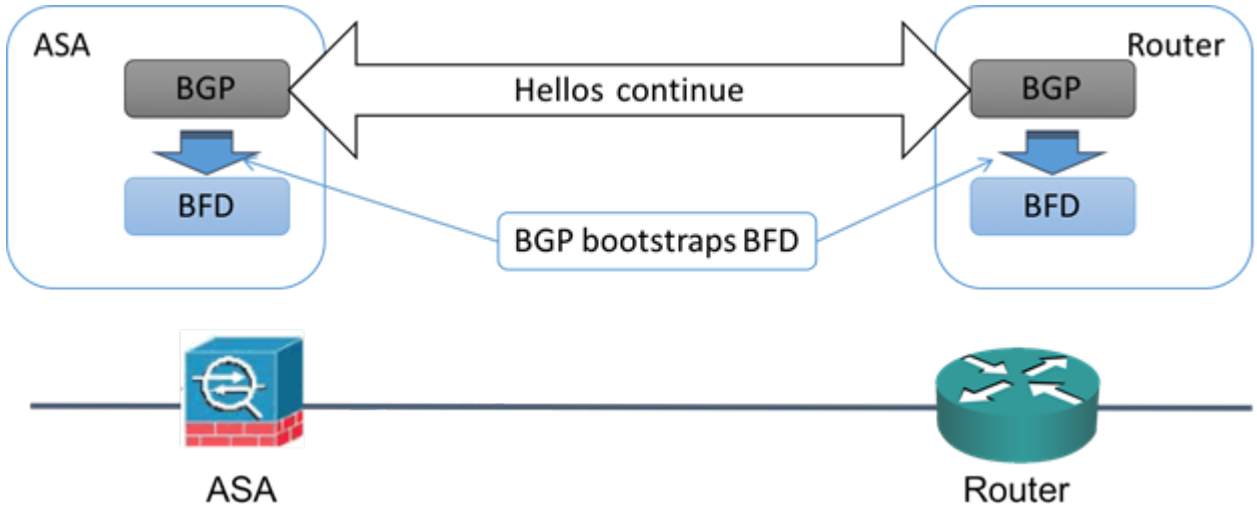
Note The echo function is not supported for IPv4 multi-hop or IPv6 single-hop BFD neighbors.

You can enable BFD at the interface and routing protocol levels. You must configure BFD on both systems (BFD peers). After you enable BFD on the interfaces and at the router level for the appropriate routing protocols, a BFD session is created, BFD timers are negotiated, and the BFD peers begin to send BFD control packets to each other at the negotiated level.

BFD Session Establishment

The following example shows the ASA and a neighboring router running Border Gateway Protocol (BGP). At the time when both devices come up, there is no BFD session established between them.

Figure 90: Established BFD Session



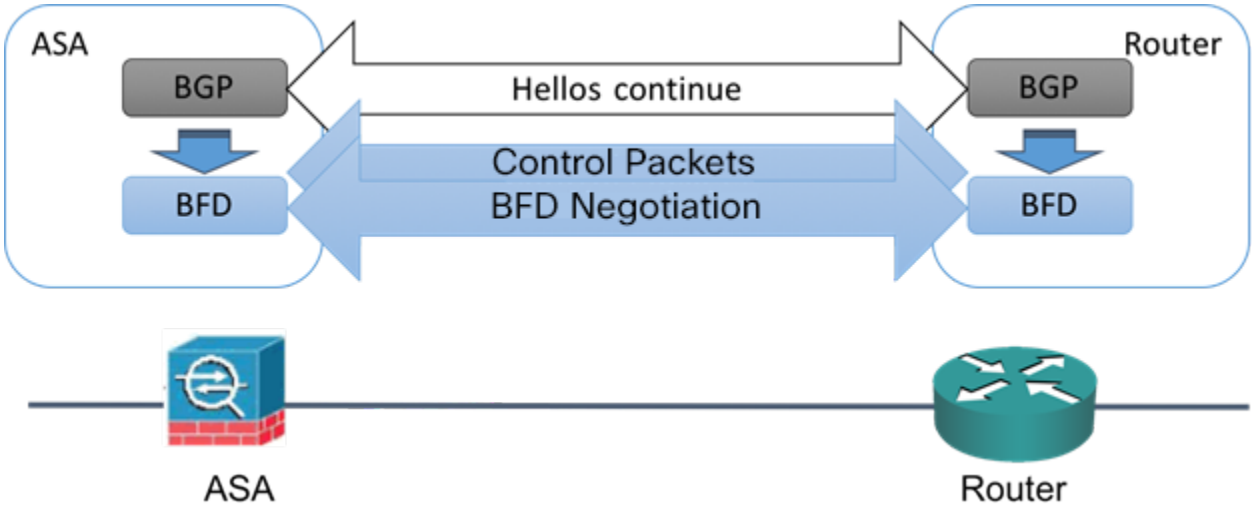
After BGP identifies its BGP neighbor, it bootstraps the BFD process with the IP address of the neighbor. BFD does not discover its peers dynamically. It relies on the configured routing protocols to tell it which IP addresses to use and which peer relationships to form.

The BFD on the router and the BFD on the ASA form a BFD control packet and start sending the packets to each other at a one-second interval until the BFD session is established. The initial control packets from either system are very similar, for example, the Vers, Diag, H, D, P, and F bits are all set to zero, and the State is set to Down. The My Discriminator field is set to a value that is unique on the transmitting device. The Your Discriminator field is set to zero because the BFD session has not yet been established. The TX and RX timers are set to the values found in the configuration of the device.

After the remote BFD device receives a BFD control packet during the session initiation phase, it copies the value of the My Discriminator field into its own Your Discriminator field and the transition from Down state to Init state and then eventually to Up state occurs. Once both systems see their own Discriminators in each other's control packets, the session is officially established.

The following illustration shows the established BFD connection.

Figure 91: BGP With No BFD Session Established



BFD Timer Negotiation

BFD devices must negotiate the BFD timers to control and synchronize the send rate of BFD control packets. A device needs to ensure the following before it can negotiate a BFD timer:

- That its peer device saw the packet containing the proposed timers of the local device
- That it never sends BFD control packets faster than the peer is configured to receive them
- That the peer never sends BFD control packets faster than the local system is configured to receive them

The setting of the Your Discriminator field and the H bit are sufficient to let the local device that the remote device has seen its packets during the initial timer exchange. After receiving a BFD control packet, each system takes the Required Min RX Interval and compares it to its own Desired Min TX Interval, and then takes the greater (slower) of the two values and uses it as the transmission rate for its BFD packets. The slower of the two systems determines the transmission rate.

When these timers have been negotiated, they can be renegotiated at any time during the session without causing a session reset. The device that changes its timers sets the P bit on all subsequent BFD control packets until it receives a BFD control packet with the F bit set from the remote system. This exchange of bits guards against packets that might otherwise be lost in transit.

Note The setting of the F bit by the remote system does not mean that it accepts the newly proposed timers. It indicates that the remote system has seen the packets in which the timers were changed.

BFD Failure Detection

When the BFD session and timers have been negotiated, the BFD peers send BFD control packets to each other at the negotiated interval. These control packets act as a heartbeat that is very similar to IGP Hello protocol except that the rate is more accelerated.

As long as each BFD peer receives a BFD control packet within the configured detection interval (Required Minimum RX Interval), the BFD session stays up and any routing protocol associated with BFD maintains its adjacencies. If a BFD peer does not receive a control packet within this interval, it informs any clients participating in that BFD session about the failure. The routing protocol determines the appropriate response to that information. The typical response is to terminate the routing protocol peering session and reconverge and thus bypass a failed peer.

Each time a BFD peer successfully receives a BFD control packet in a BFD session, the detection timer for that session is reset to zero. Thus the failure detection is dependent on received packets and NOT when the receiver last transmitted a packet.

BFD Deployment Scenarios

The following describes how BFD operates in these specific scenarios.

Failover

In a failover scenario, BFD sessions are established and maintained between the active unit and the neighbor unit. Standby units do not maintain any BFD sessions with the neighbors. When a failover happens, the new active unit must initiate session establishment with the neighbor because session information is not synched between active and standby units.

For a graceful restart/NSF scenario, the client (BGP IPv4/IPv6) is responsible for notifying its neighbor about the event. When the neighbor receives the information, it keeps the RIB table until failover is complete. During failover, the BFD and the BGP sessions go down on the device. When the failover is complete, a new BFD session between the neighbors is established when the BGP session comes up.

Spanned EtherChannel and L2 Cluster

In a Spanned EtherChannel cluster scenario, the BFD session is established and maintained between the primary unit and its neighbor. Subordinate units do not maintain any BFD sessions with the neighbors. If a BFD packet is routed to the subordinate unit because of load balancing on the switch, the subordinate unit must forward this packet to the primary unit through the cluster link. When a cluster switchover happens, the new primary unit must initiate session establishment with the neighbor because session information is not synched between primary and subordinate units.

Individual Interface Mode and L3 Cluster

In an individual interface mode cluster scenario, individual units maintain their BFD sessions with their neighbors.

Guidelines for BFD Routing

Context Mode Guidelines

Supported in single and multiple context modes.

Firewall Mode Guidelines

Supported in routed firewall mode; support for standalone, failover, and cluster modes. BFD is not supported on failover and cluster interfaces. In clustering this feature is only supported on the primary unit. BFD is not supported in transparent mode.

IPv6 Guidelines

Echo mode is not supported for IPv6.

Additional Guidelines

OSPFv2, OSPFv3, BGP IPv4, and BGP IPv6 protocol are supported.

IS-IS and EiGRP protocols are not supported.

BFD for Static Routes is not supported.

BFD on Transfer and Tunnel is not supported.

Note For optimal routing, do not configure BFD when BGP graceful restart for NSF is configured on the device.

Configure BFD

This section describes how to enable and configure the BFD routing process on your system.

Procedure

Step 1 Create the BFD Template, on page 817.
Step 2 Configure BFD Interfaces, on page 819.
Step 3 Configure BFD Maps, on page 820.

Create the BFD Template

This section describes the steps required to create a BFD template and enter BFD configuration mode.

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. You can enable Echo on single-hop only.

Procedure

Step 1 In ASDM, choose **Configuration** > **Device Setup** > **Routing** > **BFD** > **Template**.
Step 2 Click **Add** or **Edit**.

Use the **Add BFD Template** dialog box to create a new BFD template. Use the **Edit BFD Template** dialog box to change existing parameters.

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.
- **Configuration Mode**—Select **single-hop** or **multi-hop** from the drop-down list.
- **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 **None**, **Both**, **Microseconds**, or **Milliseconds**.

b) If you selected **Both**, configure the following options:

- **Multiplier Values**—The value used to compute the hold down time. Specifies 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.
- **Both Transmit and Receive Values**—The minimum transmit and receive interval capability. The range is 50 to 999 milliseconds.

c) If you selected **Microseconds**, you can click the **Both** radio button and configure the following:

- **Multiplier Values**—The value used to compute the hold down time. Specifies 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.
- **Minimum Transmit Values**—The minimum transmit interval capability. The range is 50,000 to 999,000 microseconds.
- **Minimum Receive Values**—The minimum receive interval capability. The range is 50,000 to 999,000 microseconds.

d) If you selected **Milliseconds**, configure the following:

- **Multiplier Values**—Specifies 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.
- **Minimum Transmit Values**—The minimum transmit interval capability. The range is 50 to 999 milliseconds.
- **Minimum Receive Values**—The minimum receive interval capability. The range is 50 to 999 milliseconds.

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,
- **Key Value**—The authentication string that must be sent and received in the packets using the routing protocol being authenticated. The valid value is a string containing 1 to 17 uppercase and lowercase alphanumeric characters, except that the first character CANNOT be a number.
- **Key ID**—The shared key ID that matches the key value.

Step 6 Click **OK**.

Step 7 Click **Apply** to save the BFD template configuration.

Configure BFD Interfaces

You can bind a BFD template to an interface, configure the baseline BFD session parameters per interface, and enable echo mode per interface.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **BFD** > **Interface**.

Step 2 Click **Add** or **Edit**.

Use the **Add Interface** dialog box to configure a new BFD interface. Use the **Edit Interface** dialog box to change existing parameters.

Step 3 From the **Interface** drop-down list, select the interface you want to configure with BFD.

Step 4 Check the **Template Name** check box, and choose a BFD template from the drop-down list.

Step 5 Configure the following BFD intervals:

- **Minimum Transmit Values**— The minimum transmit interval capability. The range is 50 to 999 milliseconds.
- **Minimum Receive Values**— The minimum receive interval capability. The range is 50 to 999 milliseconds.
- **Multiplier**— Specifies 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.

Step 6 (Optional) Check the **Echo** check box if you want to have Echo mode on this interface. You can only enable Echo on single-hop templates.

Step 7 Click **OK**.

Configure BFD Maps

You can create a BFD map containing destinations that you can associate with a multi-hop template. You must have a multi-hop BFD template already configured.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **BFD** > **Map**.

Step 2 Click **Add** or **Edit**.

Use the **Add Map** dialog box to configure a new BFD map. Use the **Edit Map** dialog box to change existing parameters.

Step 3 From the **Template Name** drop-down list, select a BFD template.

Step 4 Configure the following BFD intervals:

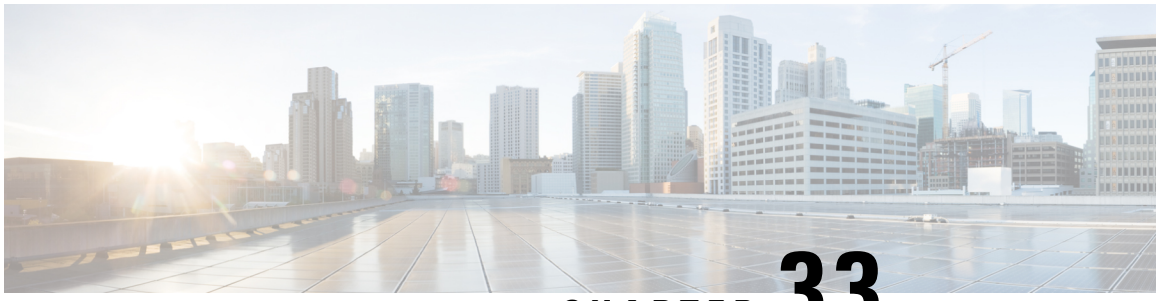
- **Minimum Transmit Values**—The minimum transmit interval capability. The range is 50 to 999 milliseconds.
- **Minimum Receive Values**— The minimum receive interval capability. The range is 50 to 999 milliseconds.
- **Multiplier**—Specifies the number of consecutive BFD control packets that must be 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.

Step 5 Click **OK**.

History for BFD Routing

Table 39: Feature History for BFD Routing

Feature Name	Platform Releases	Feature Information
BFD routing support	9.6(2)	The ASA now supports the BFD routing protocol. Support was added for configuring BFD templates, interfaces, and maps. Support for BGP routing protocol to use BFD was also added. We added or modified the following screens: **Configuration > Device Setup > Routing > BFD > Template** **Configuration > Device Setup > Routing > BFD > Interface** **Configuration > Device Setup > Routing > BFD > Map** **Configuration > Device Setup > Routing > BGP > IPv6 Family > Neighbor**



CHAPTER 33

BGP

This chapter describes how to configure the ASA to route data, perform authentication, and redistribute routing information using the Border Gateway Protocol (BGP).

- About BGP, on page 821
- Guidelines for BGP, on page 824
- Configure BGP, on page 825
- Monitoring BGP, on page 843
- History for BGP, on page 844

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).

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.

Note When a BGPv6 device joins the cluster, it generates a soft traceback when logging level 7 is enabled.

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.

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 823.
- 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.

Guidelines for BGP

Context Mode Guidelines

- Supported in single and multiple context mode.
- Only one Autonomous System (AS) number is supported for all contexts.

Firewall Mode Guidelines

Does not support transparent firewall mode. BGP is supported only in routed mode.

IPv6 Guidelines

Supports IPv6.

Additional Guidelines

- 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.

Configure BGP

This section describes how to enable and configure the BGP process on your system.

Procedure

Step 1 Enable BGP, on page 825.
Step 2 Define the Best Path for a BGP Routing Process, on page 827.
Step 3 Configure Policy Lists, on page 827.
Step 4 Configure AS Path Filters, on page 829.
Step 5 Configure Community Rules, on page 829.
Step 6 Configure IPv4 Address Family Settings, on page 830.
Step 7 Configure IPv6 Address Family Settings, on page 837.

Enable BGP

This section describes the steps required to enable BGP routing, establish a BGP routing process and configure general BGP parameters.

Procedure

Step 1 For single-mode, in ASDM, choose **Configuration** > **Device Setup** > **Routing** > **BGP** > **General**.

Note

For multi-mode, in ASDM choose Configuration > Context Management > BGP. After enabling BGP, switch to a security context and enable BGP by choosing **Configuration > Device Setup > Routing > BGP > General**.

Step 2 Check the **Enable BGP Routing** check box.

Step 3 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 XX.YY.

Step 4 (Optional) Check the **Limit the number of AS numbers in the AS_PATH attribute of received routes** check box to restrict the number of AS numbers in AS_PATH attribute to a specific number. Valid values are from 1 to 254.

Step 5 (Optional) 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.

Step 6 (Optional) Check the Use TCP path MTU discovery check box to use the Path MTU Discovery technique to determine the maximum transmission unit (MTU) size on the network path between two IP hosts. This avoids IP fragmentation.

Step 7 (Optional) Check the Enable fast external failover check box to reset the external BGP session immediately upon link failure.

Step 8 (Optional) Check the Enforce that 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.

Step 9 (Optional) Check the Use dot notation for AS numbers 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.

Step 10 Specify the timer information in the Neighbor timers area:

a) In the Keepalive interval field, enter the time interval for which the BGP neighbor remains active after not sending a keepalive message. 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) In the Hold Time field, enter the time interval for which the BGP neighbor remains active while a BGP connection is being initiated and configured.The default values is 180 seconds.

c) (Optional) In the Min. Hold Time field, enter the minimum time interval for which the he BGP neighbor remains active while a BGP connection is being initiated and configured. Specify a value from 0 to 65535.

Note
A hold time of less than 20 seconds increases the possibility of peer flapping.

Step 11 (Optional) In the Non Stop Forwarding section do the following:

a) Check the Enable Graceful Restart check box to enable ASA peers to avoid a routing flap following a switchover.

b) In the Restart Time field, enter the time duration that ASA peers will wait to delete stale routes before a BGP open message is received. The default value is 120 seconds. Valid values are between 1 and 3600 seconds.

c) In the Stale Path Time field, enter the time duration that the ASA will wait before deleting stale routes after an end of record (EOR) message is received from the restarting ASA. The default value is 360 seconds. Valid values are between 1 and 3600 seconds.

Step 12 Click **OK**.

Step 13 Click **Apply**.

Define the Best Path for a BGP Routing Process

This section describes the steps required to configure the BGP best path. For more information on the best path, see BGP Path Selection, on page 823.

Procedure

Step 1 In ASDM, choose **Configuration** > **Device Setup** > **Routing** > **BGP** > **Best Path**.

The Best Path configuration pane appears.

Step 2 In the Default Local Preference field, specify a value 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.

Step 3 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.

Step 4 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.

Step 5 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.add a new network entry. The comparison between MEDs is made only if no external autonomous systems are there in the path.

Step 6 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 this path the least desirable; therefore, a path with a missing MED is least preferred.

Step 7 Click **OK**.

Step 8 Click Apply.

Configure Policy Lists

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. This section describes the steps required to configure policy lists.

Procedure

Step 1 In ASDM, choose **Configuration** > **Device Setup** > **Routing** > **BGP** > **Policy Lists**.

Step 2 Click **Add**.

The Add Policy List dialog box appears. From this dialog box, you can add a policy list name, its redistribution access (that is, permit or deny), match interfaces, specify IP addresses, match the AS path, match community names list, match metrices, and match tag numbers.

Step 3 In the **Policy List Name** field, enter a name for the policy list.

Step 4 Click the **Permit** or **Deny** radio button to indicate the redistribution access.

Step 5 Check the **Match Interfaces** check box to distribute routes that have their next hop out of one of the interfaces specified and do one of the following:

- In the **Interface** field, enter the interface name.
- In the Interface field, click the ellipses to manually browse and locate the interface. Choose one or more interfaces, click **Interface**, then click **OK**.

Step 6 In the Specify IP area, configure the following:

a) Check the Match Address check box to redistribute any routes that have a destination network number address that is permitted by a standard access list or prefix list, and performs policy routing on packets.

Specify an access list / prefix list or click the ellipses to manually browse and locate an access list. Choose one or more access lists, click Access List, then click OK.

b) Check the Match Next Hop check box to redistribute any routes that have a next hop router address passed by one of the access lists or prefix lists specified.

Specify an access list/ prefix list or click the ellipses to manually browse and locate an access list. Choose one or more access lists, click Access List, then click OK.

c) Check the Match Route Source check box to redistribute routes that have been advertised by routers and access servers at the address specified by the access lists or prefix list.

Specify an access list/ prefix list or click the ellipses to manually browse and locate an access list. Choose one or more access lists, click Access List, then click OK.

Step 7 Check the Match AS Path check box to match a BGP autonomous system path.

Specify an AS path filter or click the ellipses to manually browse and locate the AS Path Filter. Choose one or more AS Path Filters, click AS Path Filter, then click OK.

Step 8 Check the Match Community Names List check box to match a BGP community.

a) Specify a community rule or click the ellipses to manually browse and locate the community rules.Choose one or more community rules, click Community Rules, then click OK.
b) Check the Match the specified community exactly check box to match a specific BGP community.

Step 9 Check the **Match Metrices** check box to redistribute routes with the metric specified. If you specify more than one metric, the routes can be matched with either metric.

Step 10 Check the **Match Tag Numbers** check box to redistribute routes in the routing table that match the specified tags. If you specify more than one tag number, routes can be matched with either metric.

Step 11 Click OK.

Step 12 Click Apply.

Configure AS Path Filters

An AS path filter allows you to filter the routing update message by using access lists and look at the individual prefixes within an update message. If a prefix within the update message matches the filter criteria then that individual prefix is filtered out or accepted depending on what action the filter entry has been configured to carry out. This section describes the steps required to configure AS path filters.



Note The as-path access-lists are not the same as the regular firewall ACLs.

Procedure

Step 1 In ASDM, choose **Configuration** > **Device Setup** > **Routing** > **BGP** > **AS Path Filters**.

Step 2 Click Add.

The Add Filter dialog box appears. From this dialog box, you can add a filter name, its redistribution access (that is, permit or deny), and regular expression.

Step 3 In the Name field, enter a name for the AS Path Filter.

Step 4 Click the **Permit** or **Deny** radio button to indicate the redistribution access.

Step 5 Specify the regular expression. Click**Build** to build regular expression.

Step 6 Click **Test** to test if a regular expression matches a string of your choice.

Step 7 Click**OK**.

Step 8 Click **Apply**.

Configure Community Rules

A community is a group of destinations that share some common attribute. You can use community lists to create groups of communities to use in a match clause of a route map. Just like an access list, a series of community lists can be created. Statements are checked until a match is found. As soon as one statement is satisfied, the test is concluded. This section describes the steps required to configure community rules.

Procedure

Step 1 In ASDM, choose **Configuration** > **Device Setup** > **Routing** > **BGP** > **Community Rules** > **.**

Step 2 Click Add.

The Add Community Rule dialog box appears. From this dialog box, you can add a rule name, rule type, its redistribution access (that is, permit or deny) and specific communities.

Step 3 In the Rule Name field, enter a name for the community rule.

Step 4 Click Standard or Expanded radio button to indicate the community rule type.

Step 5 Click **Permit** or **Deny** radio button to indicate the redistribution access.

Step 6 To add Standard Community Rules:

a) In the Communities field, specify a community number. Valid values are from 1 to 4294967200.
b) (Optional) Check the Internet (well-known community) check box to specify the Internet community. Routes with this community are advertised to all peers (internal and external).
c) (Optional) Check the Do not advertise to any peers (well-known community) check box to specify the no-advertise community. Routes with this community are not advertised to any peer (internal or external).
d) (Optional) Check the Do not export to next AS (well-known community) check box to specify the no-export 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 7 To add expanded community rules:

a) In the**Regular Expression** field, enter a regular expression. Alternately, Click **Build** to build regular expression.
b) Click **Test** to examine if the regular expression built, matches a string of your choice.

Step 8 Click OK.

Step 9 Click Apply.

Configure IPv4 Address Family Settings

The IPv4 settings for BGP can be set up from the IPv4 family option within the BGP configuration setup. The IPv4 family section includes subsections for General settings, Aggregate address settings, Filtering settings and Neighbor settings. Each of these subsections enable you to customize parameters specific to the IPv4 family.

Configure IPv4 Family General Settings

This section describes the steps required to configure the general IPv4 settings.

Procedure

Step 1 In ASDM, choose **Configuration** > **Device Setup** > **Routing** > **BGP** > **IPv4 Family**.

Step 2 Click General.

The General IPv4 family BGP parameters configuration pane is displayed.

Step 3 Specify External, Internal and Local distances in the Administrative Distances area.

Step 4 Choose a route map name from the Learned Routes Map drop-down list. Click **Manage** to add and configure route maps.

Step 5 (Optional) Check the**Generate Default Route** check box to configure a BGP routing process to distribute a default route (network 0.0.0.0).

Step 6 (Optional) Check the **Summarize subnet routes into network-level routes** check box to configure automatic summarization of subnet routes into network-level routes.

Step 7 (Optional) Check the **Advertise inactive routes** check box to advertise routes that are not installed in the routing information base (RIB).

Step 8 (Optional) Check the **Redistribute iBGP into an IGP** check box to configure iBGP redistribution into an interior gateway protocol (IGP), such as IS-IS or OSPF.

Step 9 (Optional) Enter a scanning interval (in seconds) for BGP routers for next-hop validation in the **Scanning Interval** field. Valid values are from 5 to 60 seconds.

Step 10 (Optional) Check the **Enable address tracking** check box to enable BGP next hop address tracking. Specify the delay interval between checks on updated next-hop routes installed in the routing table in the **Delay Interval** field.

Step 11 (Optional) Specify the maximum number of parallel internal Border Gateway Protocol (iBGP) routes that can be installed in a routing table in the **Number of paths** field and check the **iBGP multipaths** check box.

Step 12 Click**Apply** .

Configure IPv4 Family Aggregate Address Settings

This section describes the steps required to define the aggregation of specific routes into one route.

Procedure

Step 1 In ASDM, choose **Configuration** > **Device Setup** > **Routing** > **BGP** > **IPv4 Family**.

Step 2 Click **Aggregate Address**.

The Aggregate Address parameters configuration pane is displayed.

Step 3 Click**Add** .

The Add Aggregate Address pane is displayed.

Step 4 Specify a network object in the **Network** field.

Step 5 Check the **Generate autonomous system set path information** check box to generate autonomous system set path information.

Step 6 Check the **Filters all more- specific routes from the updates** check box to filter all more-specific routes from updates.

Step 7 Choose a route-map from the Attribute Map drop-down list. Click **Manage** to add or configure a route map.

Step 8 Choose a route-map from the Advertise Map drop-down list. Click **Manage** to add or configure a route.

Step 9 Choose a route-map from the Suppress Map drop-down list. Click **Manage** to add or configure a route.

Step 10 Click **OK**.

Step 11 Specify a value for the aggregate timer (in seconds) in the **Aggregate Timer** field. Valid values are 0 or any value between 6 and 60.

Step 12 Click **Apply**.

Configure IPv4 Family Filtering Settings

This section describes the steps required to filter routes or networks received in incoming BGP updates.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Routing** > **BGP** > **IPv4 Family**.

Step 2 Click **Filtering**.

The Define filters for BGP updates pane is displayed.

Step 3 Click **Add**.

The Add Filter pane is displayed.

Step 4 Choose a direction from the Direction drop-down list. The direction will specify if the filter should be applied to inbound updates or outbound updates.

Step 5 Choose a standard access list from the Access List drop-down list. Click **Manage** to add a new ACL.

Step 6 For outbound filters, you can optionally specify what types of route are distributed.

a) Choose an option from the Protocol drop-down list.

You can choose a routing protocol, such as **BGP**, **EIGRP**, **OSPF**, or **RIP**.

Choose **Connected** to filter on peers and networks learned through connected routes.

Choose **Static** to filter on peers and networks learned through static routes.

b) If you chose BGP, EIGRP, or OSPF, also choose the **Process ID** for that protocol.

Step 7 Click **OK**.

Step 8 Click **Apply**.

Configure IPv4 Family BGP Neighbor Settings

This section describes the steps required to define BGP neighbors and neighbor settings.

Procedure

Step 1 In ASDM, choose **Configuration** > **Device Setup** >**Routing** **BGP** > **IPv4 Family**.

Step 2 Click **Neighbor**.

Step 3 Click **Add**.

Step 4 Click **General** in the left pane.

Step 5 Enter a BGP neighbor IP address in the **IP Address** field. This IP address is added to the BGP neighbor table.

Step 6 Enter the autonomous system to which the BGP neighbor belongs in the **Remote AS** field.

Step 7 (Optional) Enter a description for the BGP neighbor in the **Description** field.

Step 8 (Optional) Check the **Shutdown neighbor administratively** check box to disable a neighbor or peer group.

Step 9 (Optional) Check the **Enable address family** check box to enable communication with the BGP neighbor.

Step 10 (Optional) Check the **Global Restart Functionality for this peer** check box to enable or disable the Border Gateway Protocol (BGP) graceful restart capability for a ASA neighbor or peer group.

Note

This option is enabled when the device is in HA mode or when L2 cluster (all nodes from the same network) is configured.

Step 11 (Optional) To update an interface as the source for the BGP neighborship, select the interface from the **Update-Source** drop-down box.

Note
If you update the loopback interface as source for BGP neighborship, IP address of the loopback interface is advertised across the network. The loopback interface acts as eBGP peer and participate in routing. As the loopback interface is stable when enabled and remains available until administratively shut-down, the ASA is always reachable on the loopback interface IP address.

Step 12 Click **Filtering** in the left pane.

Step 13 (Optional) Choose the appropriate incoming or outgoing access control list in the Filter routes using an access list area, to distribute BGP neighbor information. Click **Manage** to add an ACL and ACEs as required.

Step 14 (Optional) Choose the appropriate incoming or outgoing route maps in the Filter routes using a route map area, to apply a route map to incoming or outgoing routes. Click **Manage** to configure a route map.

Step 15 (Optional) Choose the appropriate incoming or outgoing prefix list in the Filter routes using a prefix list area, to distribute BGP neighbor information. Click **Manage** to configure prefix lists.

Step 16 (Optional) Choose the appropriate incoming or outgoing AS path filter in the Filter routes using AS path filter area, to distribute BGP neighbor information. Click **Manage** to configure AS path filters.

Step 17 (Optional) Check the **Limit the number of prefixes allowed from the neighbor** check box 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 to 100. The default value is 75.
- (Optional) Check the **Control prefixes received from a peer** check box to specify additional controls for the prefixes received from a peer. Do one of the following:
  - Click **Terminate peering when prefix limit is exceeded** 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.
  - Click **Give only warning message when prefix limit is exceeded** to generate a log message when the maximum prefix limit is exceeded. Here, the BGP neighbor will not be terminated.

Step 18 Click **Routes** in the left pane.

Step 19 Enter the minimum interval (in seconds) between the sending of BGP routing updates in the **Advertisement Interval** field.

Step 20 (Optional) Check the **Generate Default** route 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.

- Choose the route map that allows the route 0.0.0.0 to be injected conditionally from the Route map drop-down list. Click **Manage** to add and configure a route map.

Step 21 (Optional) To add conditionally advertised routes do the following:

a) Click **Add** in the Conditionally Advertised Routes section.

b) Choose the route map that will be advertised from the Advertise Map drop-down list, if the conditions of the exist map or the non-exist map are met.

c) Do one of the following:

- Click **Exist Map** and choose a route map. This route map will be compared with the routes in the BGP table, to determine whether or not the advertise map route is advertised.
- Click **Non-exist Map** and choose a route map. This route map will be compared with the routes in the BGP table, to determine whether or not the advertise map route is advertised.

d) Click **Ok**.

Step 22 (Optional) Check the **Remove private autonomous system (AS) numbers from outbound routing updates** check box to exclude the private AS numbers from being advertised on outbound routes.

Step 23 Click **Timers** in the left pane.

Step 24 (Optional) Check the Set timers for the BGP peer check box to set the keepalive frequency, hold time and minimum hold time.

- Enter the frequency (in seconds) with which the ASA sends keepalive messages to the neighbor, in the **Keepalive frequency** field. Valid values are between 0 and 65535. The default value is 60 seconds.
- Enter the interval (in seconds) after not receiving a keepalive message that the ASA declares a peer dead, in the **Hold time** field. The default value is 180 seconds.
- (Optional) Enter the minimum interval (in seconds) after not receiving a keepalive message that the ASA declares a peer dead, in the Min Hold time field.

Note
A hold time of less than 20 seconds increases the possibility of peer flapping.

Step 25 Click Advanced in the left pane.

Step 26 (Optional) Check the Enable Authentication check box to enable MD5 authentication on a TCP connection between two BGP peers.

- Choose an encryption type from the Encryption Type drop-down list.
- 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. You cannot specify a password in the format number-space-anything. The space after the number can cause authentication to fail.

Step 27 (Optional) Check the **Send Community Attribute to this neighbor** check box.

Step 28 (Optional) Check the **Use ASA as next hop for neighbor** check box to configure the router as the next-hop for a BGP speaking neighbor or peer group.

Step 29 Do one of the following:

- Click **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.

- (Optional) Check the **Disable connection verification** check box to disable connection verification to establish an eBGP peering session with a single-hop peer that uses a loopback interface.

- Click **Limit number of TTL hops to neighbor** to enable you 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.

Step 30 (Optional) Enter a weight for the BGP neighbor connection in the **Weight** field.

Step 31 Choose the BGP version that the ASA will accept from the **BGP version** drop-down list.

Note
The version can be set to 2 to force the software to use only Version 2 with the specified neighbor. The default is to use Version 4 and dynamically negotiate down to Version 2 if requested.

Step 32 (Optional) Check the **TCP Path MTU Discovery** check box to enable a TCP transport session for a BGP session.

Step 33 Choose the TCP connection mode from the **TCP transport mode** drop-down list.

Step 34 Click **Migration** in the left pane

Step 35 (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.

- Enter the local autonomous system number in the **Local AS Number** field. Valid values are between 1 and 65535.
- (Optional) Check the **Do not prepend local AS number** for routes received from neighbor check box. The local AS number will not be prepended to any routes received from eBGP peer.
- (Optional) Check the **Replace real AS number with local AS numbern routes received from neighbor** check box. The AS number from the local routing process is not prepended.
- (Optional) Check the **Accept either real AS number or local AS number in routes received from neighbor** check box.

Step 36 Click **OK**.

Step 37 Click **Apply**.

Configure IPv4 Network Settings

This section describes the steps required to define the networks to be advertised by the BGP routing process.

Procedure

Step 1 In ASDM, choose **Configuration** > **Device Setup** > **Routing** > **BGP** > **IPv4 Family**.

Step 2 Click **Networks**.

The Define networks to be advertised by the BGP routing process configuration pane appears.

Step 3 Click **Add**.

The Add Network pane is displayed.

Step 4 Specify the network that BGP will advertise in the**Address** field.

Note
For a network prefix to be advertised, a route to the device must exist on the routing table.

Step 5 (Optional) Choose a network or subnetwork mask from the **Netmask** drop-down list.

Step 6 Choose a route map that should be examined to filter the networks to be advertised from the **Route Map** drop-down list. Click Manage to configure or add a route map.

Step 7 Click **OK**.

Step 8 Click **Apply**.

Configure IPv4 Redistribution Settings

This section describes the steps required to define the conditions for redistributing routes from another routing domain into BGP.

Procedure

Step 1 In ASDM, choose **Configuration** > **Device Setup** > **Routing** > **BGP** > **IPv4 Family** > **.**

Step 2 Click **Redistribution**.

The Redistribution pane is displayed.

Step 3 Click**Add** .

The Add Redistribution pane is displayed.

Step 4 Choose the protocol from which you want to redistribute routes into the BGP domain from the **Source Protocol** drop-down list.

Step 5 Choose a process ID for the source protocol from the **Process ID** drop-down list.

Step 6 (Optional) Enter a metric for the redistributed route in the**Metric** field.

Step 7 Choose a route map that should be examined to filter the networks to be redistributed from the **Route Map** drop-down list. Click **Manage** to configure or add a route map.

Step 8 Check one or more of the**Internal** , **External** and **NSSA External Match** check boxes to redistribute routes from an OSPF network.

This step is only applicable for redistribution from OSPF networks.

Step 9 Click **OK**.

Step 10 Click**Apply** .

Configure IPv4 Route Injection Settings

This section describes the steps required to define the routes to be conditionally injected into the BGP routing table.

Procedure

Step 1 In ASDM, choose **Configuration** > **Device Setup** > **Routing** > **BGP** > **IPv4 Family** > **.**

Step 2 Click **Route Injection**.

The Route Injection pane is displayed.

Step 3 Click**Add** .

The Add Conditionally injected route pane is displayed.

Step 4 Choose the route map that specifies the prefixes to inject into the local BGP routing table from the **Inject Map** drop-down list.

Step 5 Choose the route map containing the prefixes that the BGP speaker will track from the **Exist Map** drop-down list.

Step 6 Check the **Injected routes will inherit the attributes of the aggregate route** check box to configure the injected route to inherit attributes of the aggregate route.

Step 7 Click**OK** .

Step 8 Click **Apply**.

Configure IPv6 Address Family Settings

The IPv6 settings for BGP can be set up from the IPv6 family option within the BGP configuration setup. The IPv6 family section includes subsections for General settings, Aggregate address settings and Neighbor settings. Each of these subsections enable you to customize parameters specific to the IPv6 family.

This section describes how to customize the BGP IPv6 family settings.

Configure IPv6 Family General Settings

This section describes the steps required to configure the general IPv6 settings.

Procedure

Step 1 In ASDM, choose **Configuration** > **Device Setup** > **Routing** > **BGP** > **IPv6 Family**.

Step 2 Click **General**.

The general IPv6 family BGP parameters configuration pane is displayed.

Step 3 Specify external, internal and local distances in the **Administrative Route Distances** area.

Step 4 (Optional) Check the**Generate Default Route** check box to configure a BGP routing process to distribute a default route (network 0.0.0.0).

Step 5 (Optional) Check the **Advertise inactive routes** check box to advertise routes that are not installed in the routing information base (RIB).

Step 6 (Optional) Check the **Redistribute iBGP into an IGP** check box to configure iBGP redistribution into an interior gateway protocol (IGP), such as IS-IS or OSPF.

Step 7 (Optional) Enter a scanning interval (in seconds) for BGP routers for next-hop validation in the**Scanning Interval** field. Valid values are from 5 to 60 seconds.

Step 8 (Optional) Specify the maximum number of Border Gateway Protocol routes that can be installed in a routing table in the **Number of paths** field.

Step 9 (Optional) Check the**IBGP multipaths** check box and specify the maximum number of parallel internal Border Gateway Protocol (iBGP) routes that can be installed in a routing table in the**Number of paths** field.

Step 10 Click **Apply**.

Configure IPv6 Family Aggregate Address Settings

This section describes the steps required to define the aggregation of specific routes into one route.

Procedure

Step 1 In ASDM, choose **Configuration** > **Device Setup** > **Routing** > **BGP** > **IPv6 Family**.

Step 2 Click **Aggregate Address**.

The Aggregate Address parameters configuration pane is displayed.

Step 3 Click**Add** .

The Add Aggregate Address pane is displayed.

Step 4 Specify an IPv6 address in the**IPv6/Address Mask** field. Alternately, browse to add a network object.

Step 5 Check the **Generate autonomous system set path information** check box to generate autonomous system set path information. The path advertised for this route will be an AS_SET consisting of all elements contained in all paths that are being summarized.

Note
Do not use this form of the aggregate-address command when aggregating many paths, because this route must be continually withdrawn and updated as autonomous system path reachability information for the summarized routes changes.

Step 6 Check the**Filters all more- specific routes from the updates** check box to filter all more-specific routes from updates. This not only creates the aggregate route but also suppresses advertisements of more-specific routes to all neighbors.

Step 7 Choose a route-map from the **Attribute Map** drop-down list. Click **Manage** to add or configure a route map. This allows attributes of the aggregate route to be changed.

Step 8 Choose a route-map from the**Advertise Map** drop-down list. Click **Manage** to add or configure a route. This selects specific routes that will be used to build different components of the aggregate route.

Step 9 Choose a route-map from the **Suppress Map** drop-down list. Click **Manage** to add or configure a route. This creates the aggregate route but suppresses advertisement of specified routes.

Step 10 Click**OK** .

Step 11 Specify a value for the aggregate timer (in seconds) in the **Aggregate Timer** field. Valid values are 0 or any value between 6 and 60. This specifies the interval at which the routes will be aggregated. The default value is 30 seconds.

Step 12 Click**Apply** .

Configure IPv6 Family BGP Neighbor Settings

This section describes the steps required to define BGP neighbors and neighbor settings.

Procedure

Step 1 In ASDM, choose**Configuration** > **Device Setup** > **Routing** > **BGP** > **IPv6 Family**.

Step 2 Click**Neighbor** .

Step 3 Click**Add** .

Step 4 Click**General** in the left pane.

Step 5 Enter a BGP neighbor IPv6 address in the**IPv6 Address** field. This IPv6 address is added to the BGP neighbor table.

Step 6 Enter the autonomous system to which the BGP neighbor belongs in the **Remote AS** field.

Step 7 (Optional) Enter a description for the BGP neighbor in the **Description** field.

Step 8 (Optional) Check the **Shutdown neighbor administratively** check box to disable a neighbor or peer group.

Step 9 (Optional) Check the **Enable address family** check box to enable communication with the BGP neighbor.

Step 10 (Optional) Check the **Global Restart Functionality for this peer** check box to enable or disable the Border Gateway Protocol (BGP) graceful restart capability for a ASA neighbor or peer group.

Note
This option is enabled when the device is in HA mode or when L2 cluster (all nodes from the same network) is configured.

Step 11 (Optional) To update an interface as the source for BGP neighborship, select the interface from the **Update-Source** drop-down box.

Note
If you update the loopback interface as source for BGP neighborship, IP address of the loopback interface is advertised across the network. The loopback interface acts as eBGP peer and participate in routing. As the loopback interface is stable when enabled and remains available until administratively shut-down, the ASA is always reachable on the loopback interface IP address.

Step 12 Click **Filtering** in the left pane.

Step 13 (Optional) Choose the appropriate incoming or outgoing route maps in the **Filter routes using a route map** area, to apply a route map to incoming or outgoing routes. Click **Manage** to configure a route map.

Step 14 (Optional) Choose the appropriate incoming or outgoing prefix list in the **Filter routes using a prefix list** area, to distribute BGP neighbor information. Click **Manage** to configure prefix lists.

Step 15 (Optional) Choose the appropriate incoming or outgoing AS path filter in the **Filter routes using AS path filter** area, to distribute BGP neighbor information. Click **Manage**to configure AS path filters.

Step 16 (Optional) Check the **Limit the number of prefixes allowed from the neighbor** check box to control the number of prefixes that can be received from a neighbor.

Step 17 Enter the maximum number of prefixes allowed from a specific neighbor in the **Maximum prefixes** field.

Step 18 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 to 100. The default value is 75.

Step 19 (Optional) Check the **Control prefixes received from a peer** check box to specify additional controls for the prefixes received from a peer. Do one of the following:

- Click **Terminate peering when prefix limit is exceeded** 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.
- Click **Give only warning message when prefix limit is exceeded** to generate a log message when the maximum prefix limit is exceeded. Here, the BGP neighbor will not be terminated.

Step 20 Click **Routes**in the left pane.

Step 21 Enter the minimum interval (in seconds) between the sending of BGP routing updates in the **Advertisement Interval**field.

Step 22 (Optional) Check the **Generate Default route**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.

Step 23 Choose the route map that allows the route 0.0.0.0 to be injected conditionally from the **Route map** drop-down list. Click **Manage** to add and configure a route map.

Step 24 (Optional) To add conditionally advertised routes do the following:

a) Click **Add** in the Conditionally Advertised Routes section.

b) Choose the route map that will be advertised from the **Advertise Map** drop-down list, if the conditions of the exist map or the non-exist map are met.

c) Do one of the following:

- Click **Exist Map** and choose a route map. This route map will be compared with the routes in the BGP table, to determine whether or not the advertise map route is advertised.
- Click **Non-exist Map** and choose a route map. This route map will be compared with the routes in the BGP table, to determine whether or not the advertise map route is advertised.

d) Click **Ok**.

Step 25 (Optional) Check the **Remove private autonomous system (AS) numbers from outbound routing updates** check box to exclude the private AS numbers from being advertised on outbound routes.

Step 26 Click **Timers** in the left pane.

Step 27 (Optional) Check the **Set timers for the BGP peer** check box to set the keepalive frequency, hold time and minimum hold time.

Step 28 Enter the frequency (in seconds) with which the ASA sends keepalive messages to the neighbor in the **Keepalive frequency** field. Valid values are between 0 and 65535. The default value is 60 seconds.

Step 29 Enter the interval (in seconds) after not receiving a keepalive message that the ASA declares a peer dead, in the**Hold time** field. The default value is 180 seconds.

Step 30 (Optional) Enter the minimum interval (in seconds) after not receiving a keepalive message that the ASA declares a peer dead, in the **Min Hold time** field.

Note
A hold time of less than 20 seconds increases the possibility of peer flapping.

Step 31 Click **Advanced** in the left pane.

Step 32 (Optional) Check the**Enable Authentication** check box to enable MD5 authentication on a TCP connection between two BGP peers.

Step 33 Choose an encryption type from the **Encryption Type** drop-down list.

Step 34 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. You cannot specify a password in the format number-space-anything. The space after the number can cause authentication to fail.

Step 35 (Optional) Check the **Send Community Attribute to this neighbor** check box.

Step 36 (Optional) Check the **Use ASA as next hop for neighbor** check box to configure the router as the next-hop for a BGP speaking neighbor or peer group.

Step 37 Do one of the following:

- Click **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.
  - (Optional) Check the**Disable connection verification** check box to disable connection verification to establish an eBGP peering session with a single-hop peer that uses a loopback interface.
- Click **Limit number of TTL hops to neighbor** to enable you 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

Step 38 (Optional) Enter a weight for the BGP neighbor connection in the **Weight** field.

Step 39 Choose the BGP version that the ASA will accept from the BGP version drop-down list.

Note
The version can be set to 2 to force the software to use only Version 2 with the specified neighbor. The default is to use Version 4 and dynamically negotiate down to Version 2 if requested.

Step 40 (Optional) Check the **TCP Path MTU Discovery** check box to enable a TCP transport session for a BGP session.

Step 41 Choose the **TCP connection mode** from the TCP transport mode drop-down list.

Step 42 Click**Migration** in the left pane

Step 43 (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.

- Enter the local autonomous system number in the Local AS Number field. Valid values are between 1 and 65535.
- (Optional) Check the Do not prepend local AS number for routes received from neighbor check box. The local AS number will not be prepended to any routes received from eBGP peer.
- (Optional) Check the Replace real AS number with local AS number in routes received from neighbor check box. The AS number from the local routing process is not prepended.
- (Optional) Check the Accept either real AS number or local AS number in routes received from neighbor check box.

Step 44 Click **OK**.

Step 45 Click**Apply** .

Configure IPv6 Network Settings

This section describes the steps required to define the networks to be advertised by the BGP routing process.

Procedure

Step 1 In ASDM, choose **Configuration** > **Device Setup** > **Routing** > **BGP** > **IPv6 Family**.

Step 2 Click **Networks**.

The Define the networks to be advertised by the BGP routing process configuration pane appears.

Step 3 Click **Add**.

The Add Network pane is displayed.

Step 4 (Optional) In the Prefix Name field, specify the prefix name for the DHCPv6 Prefix Delegation client (see Enable the IPv6 Prefix Delegation Client, on page 648).

Step 5 In the**IPv6 Address/mask** field, specify the network that BGP will advertise.

If you specify the **Prefix Name**, enter the subnet prefix and mask; the advertised network consists of the delegated prefix + the subnet prefix.

Step 6 Choose a route map that should be examined to filter the networks to be advertised from the **Route Map** drop-down list. Optionally, click **Manage** to configure or add a route map.

Step 7 Click **OK**.

Step 8 Click**Apply** .

Configure IPv6 Redistribution Settings

This section describes the steps required to define the conditions for redistributing routes from another routing domain into BGP.

Procedure

Step 1 In ASDM, choose **Configuration** > **Device Setup** > **Routing** > **BGP** > **IPv6 Family**.

Step 2 Click Redistribution.

Step 3 Click Add.

The Add Redistribution pane is displayed.

Step 4 In the Source Protocol drop-down list, choose the protocol from which you want to redistribute routes into the BGP domain.

Step 5 In the Process ID drop-down list, choose a process ID for the source protocol. This is available only for OSPF source protocol.

Step 6 (Optional) In the Metric field, enter a metric for the redistributed route.

Step 7 In the Route Map drop-down list, choose a route map that should be examined to filter the networks to be redistributed. Click Manage to configure or add a route map.

Step 8 Check one or more of the Match check-boxes -Internal, External 1, External 2, NSSA External 1, and NSSA External 2 check boxes to redistribute routes from an OSPF network.

This step is only applicable for redistribution from OSPF networks.

Step 9 Click OK.

Step 10 Click Apply.

Configure IPv6 Route Injection Settings

This section describes the steps required to define the routes to be conditionally injected into the BGP routing table.

Procedure

Step 1 In ASDM, choose **Configuration** > **Device Setup** > **Routing** > **BGP** > **IPv4 Family**.

Step 2 Click **Route Injection**.

Step 3 Click**Add** .

The Add Conditionally injected route pane is displayed.

Step 4 In the **Inject Map** drop-down list, choose the route map that specifies the prefixes to inject into the local BGP routing table.

Step 5 In the **Exist Map** drop-down list, choose the route map containing the prefixes that the BGP speaker will track.

Step 6 Check the **Injected routes will inherit the attributes of the aggregate route** check box to configure the injected route to inherit attributes of the aggregate route.

Step 7 Click **OK**.

Step 8 Click**Apply**.

Monitoring BGP

You can use the following commands to monitor the BGP routing process. For examples and descriptions of the command output, see the command reference. Additionally, you can disable the logging of neighbor change messages and neighbor warning messages.

To monitor various BGP routing statistics, enter one of the following commands:

Note To disable BGP Log messages, enter the **no bgp log-neighbor-changes** command in the router configuration mode. This disables the logging of neighbor change messages. Enter this command in router configuration mode for the BGP routing process. By default, neighbor changes are logged.

- **Monitoring > Routing > BGP Neighbors**

  Each row represents one BGP neighbor. For each neighbor, the list includes the IP address, the AS number, the router ID, the state (active, idle and so on), the uptime, graceful restart capability, the restart time and the stalepath time.

- **Monitoring > Routing > BGP Routes**

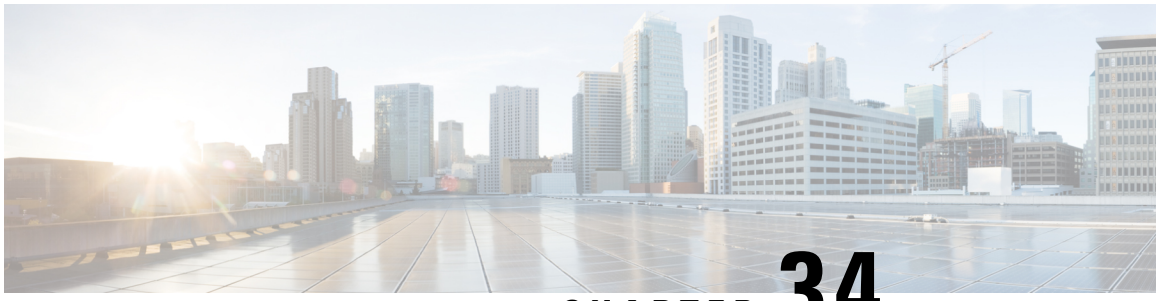
  Each row represents one BGP route. For each route, the list includes the status code, IP address, the next hop address, the route metric, the local preference values, the weight and the path.

History for BGP

Table 40: Feature History for BGP

Feature Name	Platform Releases	Feature Information
BGP Support	9.2(1)	Support was added for routing data, performing authentication, and redistributing and monitoring routing information using the Border Gateway Protocol. We introduced the following screens: Configuration > Device Setup > Routing > BGP Monitoring > Routing > BGP Neighbors, Monitoring > Routing > BGP Routes We modified the following screens: Configuration > Device Setup > Routing > Static Routes> Add > Add Static Route Configuration > Device Setup > Routing > Route Maps> Add > Add Route Map
BGP support for ASA clustering	9.3(1)	We added support for L2 and L3 clustering. We modified the following screen: Configuration > Device Setup > Routing > BGP > IPv4 Family > General
BGP support for nonstop forwarding	9.3(1)	We added support for Nonstop Forwarding. We modified the following screens: Configuration > Device Setup > Routing > BGP > General, Configuration > Device Setup > Routing > BGP > IPv4 Family > Neighbor, Monitoring > Routing > BGP Neighbors
BGP support for advertised maps	9.3(1)	We added support for BGPv4 advertised map. We modified the following screen: Configuration > Device Setup > Routing > BGP > IPv4 Family > Neighbor > Add BGP Neighbor > Routes
BGP support for IPv6	9.3(2)	We added support for IPv6. We introduced the following screen: Configuration > Device Setup > Routing > BGP > IPv6 Family

Feature Name	Platform Releases	Feature Information
IPv6 network advertisement for delegated prefixes	9.6(2)	The ASA now supports the DHCPv6 Prefix Delegation client. The ASA obtains delegated prefix(es) from a DHCPv6 server. The ASA can then use these prefixes to configure other ASA interface addresess so that StateLess Address Auto Configuration (SLAAC) clients can autoconfigure IPv6 addresses on the same network. You can configure the BGP router to advertise these prefixes. We modified the following screen: **Configuration** > **Device Setup** > **Routing** > **BGP** > **IPv6 Family** > **Networks**
Loopback interface support for BGP traffic	9.18(2)	You can now add a loopback interface and use it for BGP traffic. New/Modified commands: **interface loopback**, **neighbor update-source** New/Modified screens: • **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces** > **Add Loopback Interface** • **Configuration** > **Device Setup** > **Routing** > **BGP** > **IPv4 Family / IPv6 Family** > **Neighbor** > **Add** > **General** ASDM support was added in 7.19.
Graceful restart supported for IPv6	9.19(1)	We added graceful restart support for IPv6 address family.



CHAPTER 34

OSPF

This chapter describes how to configure the ASA to route data, perform authentication, and redistribute routing information using the Open Shortest Path First (OSPF) routing protocol.

- About OSPF, on page 847
- Guidelines for OSPF, on page 850
- Configure OSPFv2, on page 853
- Configure OSPFv2 Router ID, on page 855
- Customize OSPFv2, on page 857
- Configure OSPFv3, on page 874
- Configure Graceful Restart, on page 884
- Example for OSPFv2, on page 888
- Examples for OSPFv3, on page 890
- Monitoring OSPF, on page 891
- History for OSPF, on page 892

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 ASA 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 ASA 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 ASA 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 ASA as a designated router or a designated backup router. The ASA 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 ASA 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 ASA. Also, you should not mix public and private networks on the same ASA interface.

You can have two OSPF routing processes, one RIP routing process, and one EIGRP routing process running on the ASA 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.

About OSPF Support for Fast Hello Packets

The key concepts related to OSPF support for fast hello packets and the benefits of OSPF Fast Hello Packets are described below:

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 849.

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.

Guidelines for OSPF

Context Mode Guidelines

OSPFv2 supports single and multiple context mode.

- OSPFv2 instances cannot form adjacencies with each other across shared interfaces because, by default, inter-context exchange of multicast traffic is not supported across shared interfaces. However, you can use the static neighbor configuration under OSPFv2 process configuration under OSPFv2 process to bring up OSPFv2 neighbourship on a shared interface.
- Inter-context OSPFv2 on separate interfaces is supported.

(For point-to-point topology) In a multiple context mode, you can configure static OSPFv2 neighbors to advertise OSPFv2 routes over a point-to-point, non-broadcast network. For a successful OSPF adjacency formation, when you remove and reconfigure OSPF on a shared interface, execute the commands in the following order:

- Remove the OSPF configuration using `no router <ospf name>`, and then remove the OSPF point-to-point non-broadcast configuration on the interface using `no ospf network point-to-point non-broadcast`.
- When you reconfigure OSPF on the shared interface, configure the OSPF router using `router <ospf name>`, and then configure the interface with `ospf network point-to-point non-broadcast`.

OSPFv3 supports single mode only.

Key Chain Authentication Guidelines

OSPFv2 supports key chain authentication on both single and multiple mode, both in physical and virtual modes. However, in multiple mode, you can configure the key chain only in context mode.

- The rotating keys are applicable only for OSPFv2 protocol. OSPF area authentication with key chain is not supported.
- The existing MD5 authentication without time range in OSPFv2 is still supported along with new rotating keys.
- Though the platform supports SHA1 and MD5 cryptographic algorithms, only MD5 cryptographic algorithm is used for authentication.

Firewall Mode Guidelines

OSPF supports routed firewall mode only. OSPF does not support transparent firewall mode.

Failover Guidelines

OSPFv2 and OSPFv3 support Stateful Failover.

IPv6 Guidelines

- OSPFv2 does not support IPv6.
- OSPFv3 supports IPv6.
- OSPFv3 uses IPv6 for authentication.
- The ASA installs OSPFv3 routes into the IPv6 RIB, provided it is the best route.
- OSPFv3 packets can be filtered out using IPv6 ACLs in the **capture** command.

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.

Disable the opaque capability on the ASA to ensure that the neighbor relationship remains stable:

```
router ospf process_ID_number
no nsf ietf helper
no capability opaque
```



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.
- 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.
- The ASA virtual cannot use dynamic interior routing protocols like EIGRP and OSPF due to the nature of Azure cloud routing. The Effective Routing Table determines the next hop, regardless of whether a virtual client has any static/dynamic route configured.

  Currently, you cannot view either the Effective Routing Table or the System Routing Table.
- 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 how to enable an OSPFv2 process on the ASA.

After you enable OSPFv2, you need to define a route map. For more information, see Define a Route Map, on page 805. Then you generate a default route. For more information, see Configure a Static Route, on page 789.

After you have defined a route map for the OSPFv2 process, you can customize it for your particular needs, To learn how to customize the OSPFv2 process on the ASA, see Customize OSPFv2, on page 857.

To enable OSPFv2, you need to create an OSPFv2 routing process, specify the range of IP addresses associated with the routing process, then assign area IDs associated with that range of IP addresses.

You can enable up to two OSPFv2 process instances. Each OSPFv2 process has its own associated areas and networks.

To enable OSPFv2, perform the following steps:

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPF** > **Setup**.

In the OSPF Setup pane, you can enable OSPF processes, configure OSPF areas and networks, and define OSPF route summarization.

Step 2 The three tabs in ASDM used to enable OSPF are as follows:

- The Process Instances tab allows you to enable up to two OSPF process instances for each context. Single context mode and multiple context mode are both supported. After you check the **Enable Each OSPF Process** check box, you can enter a unique identifier numeric identifier for that OSPF process. This process ID is used internally and does not need to match the OSPF process ID on any other OSPF devices; valid values range from 1 to 65535. Each OSPF process has its own associated areas and networks.

  If you click **Advanced**, the Edit OSPF Process Advanced Properties dialog box appears. From here, you can configure the Router ID, cluster IP address pools in Spanned EtherChannel or Individual Interface clustering, Adjacency Changes, Administrative Route Distances, Timers, and Default Information Originate settings for each OSPF process. You can enable BFD on all interfaces that support OSPFv2 here, or enable BFD for specific OSPFv2 interfaces (see Configure OSPFv2 Interface Parameters, on page 861).

- The Area/Networks tab allows you to display the areas and the networks that they include for each OSPF process on the ASA. From this tab you can display the area ID, the area type, and the type of authentication set for the area. To add or edit the OSPF area or network, see Configure OSPFv2 Area Parameters, on page 864 for more information.

- The Route Summarization tab allows you to configure an ABR. In OSPF, an ABR will advertise networks in one area into another area. If the network numbers in an area are assigned in a way so that they are contiguous, you can configure the ABR to advertise a summary route that includes all the individual networks within the area that fall into the specified range. See Configure Route Summarization Between OSPFv2 Areas, on page 860 for more information.

Configure a Key Chain for Authentication

To enhance data security and protection of devices, you can enable rotating keys for authenticating IGP peers. 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. If the key lifetime expires and no active keys are found, OSPF uses the last valid key to maintain the adjacency with the peers.

This section describes how to create a key chain for OSPF peer authentication. This section also covers steps to add or edit the key chain attributes. After configuring a key chain object, you can use it in defining the OSPFv2 authentication for an interface and for a virtual link. Use the same authentication type (MD5 or Key Chain) and key ID for the peers to establish a successful adjacency. To learn how to define authentication for an interface, see Configure OSPFv2 Interface Parameters, on page 861; for a virtual link, see Configure a Virtual Link in OSPF, on page 873.

To configure a key chain, perform the following steps:

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Key Chain** .

Step 2 In the **Configure Key Chain** section, click **Add**.

Step 3 Enter the key chain name in the **Add Key Chain** dialog box, and click **Ok**.

The created key chain name is listed in the **Configure Key Chain** grid.

Step 4 Select the key chain name from the **Configure Key Chain** section, and in the **Configure Key** section, click **Add**. To edit an existing key, seclect the key name and click **Edit**.

The **Add Key** or **Edit Key**dialog box appears, depending on the action that you have selected.

Step 5 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.

Note
You cannot edit a saved key id.

Step 6 From the **Cryptographic Algorithm** drop-down, choose **MD5**. MD5 is the only algorithm supported for authenticating the key chain.

Step 7 Select the encryption type by clicking the **Plain Text** or **Encrypted** radio button, and then enter the password in the **Authentication Key** 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 8 Provide the lifetime values in the **Accept Lifetime** and **Send Lifetime** fields:

You can specify the time interval for the device to accept/send the key during key exchange with another device. The end time can be the duration, the absolute time when the accept/send lifetime ends, or never expires.

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 end lifetime.

Step 9 To save the key chain attibutes, click **Ok**. In the **Key Chain** page, click **Appy**.

What to do next

You can now apply the configured key chain to define the OSPFv2 authentication for an interface and for virtual link.

- Configure OSPFv2 Interface Parameters, on page 861
- Configure a Virtual Link in OSPF, on page 873

Configure OSPFv2 Router ID

The OSPF Router-ID is used to identify a specific device within an OSPF database. No two routers in an OSPF system can have the same router-id.

If a router-id is not configured manually in the OSPF routing process the router will automatically configure a router-id determined from the highest IP address of an active interface. When configuring a router-id, the neighbors will not be updated automatically until that router has failed or the OSPF process has been cleared and the neighbor relationship has been re-established.

Manually Configure OSPF Router-ID

This section describes how to manually configure router-id in OSPFv2 process on the ASA.

Procedure

Step 1 To use a fixed router ID, use the **router-id** command.

router-id *ip-address*

Example:

```
ciscoasa(config-router)# router-id 193.168.3.3
```

Step 2 To revert to the previous OSPF router ID behavior, use the **no router-id** command.

no router-id *ip-address*

Example:

```
ciscoasa(config-router)# no router-id 193.168.3.3
```

Router ID Behaviour while Migrating

While migrating OSPF configuration from one ASA , say ASA 1 to another ASA, say ASA 2, the following router id selection behaviour is observed:

1. ASA 2 does not use any IP address for OSPF router-id when all interfaces are in shutdown mode. The possibilities for configuring router-id when all interfaces are in "admin down" state or shutdown mode are:
   - If ASA 2 does not have any router-id configured before, you would see this message:

     ```
     %OSPF: Router process 1 is not running, please configure a router-id
     ```

     After the first interface is brought up, ASA 2 will take IP address of this interface as router id.
   - If ASA 2 had router-id configured before and all interfaces were in "admin down" state when "no router-id" command was issued, ASA 2 will use old router id. ASA 2 uses the old router id, even if IP addresses on the interface that is brought up is changed, until "clear ospf process" command is issued.

2. ASA 2 uses new router id, when ASA 2 had router-id configured before and at least one of interfaces were not in "admin down" state or shutdown mode when "no router-id" command was issued. ASA 2 will use new router id from the IP address of the interfaces even when interfaces are in "down/down" state.

Customize OSPFv2

This section explains how to customize the OSPFv2 processes.

Redistribute Routes Into OSPFv2

The ASA can control the redistribution of routes between OSPFv2 routing processes.



Note If you want to redistribute a route by defining which of the routes from the specified routing protocol are allowed to be redistributed into the target routing process, you must first generate a default route. See Configure a Static Route, on page 789, and then define a route map according to Define a Route Map, on page 805.

To redistribute static, connected, RIP, or OSPFv2 routes into an OSPFv2 process, perform the following steps:

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPF** > **Redistribution**.

The Redistribution pane displays the rules for redistributing routes from one routing process into an OSPF routing process. You can redistribute routes discovered by RIP and OSPF into the EIGRP routing process. You can also redistribute static and connected routes into the EIGRP routing process. You do not need to redistribute static or connected routes if they fall within the range of a network that has been configured through the Setup > Networks tab.

Step 2 Click **Add** or **Edit**.

Alternatively, double-clicking a table entry in the Redistribution pane (if any) opens the Add/Edit OSPF Redistribution Entry dialog box for the selected entry.

Note
All steps that follow are optional.

The Add/Edit OSPF Redistribution Entry dialog box lets you add a new redistribution rule or edit an existing redistribution rule in the Redistribution table. Some of the redistribution rule information cannot be changed when you are editing an existing redistribution rule.

Step 3 Choose the OSPF process associated with the route redistribution entry. If you are editing an existing redistribution rule, you cannot change this setting.

Step 4 Choose the source protocol from which the routes are being redistributed. You can choose one of the following options:

- Static—Redistributes static routes to the OSPF routing process.

- Connected—Redistributes connected routes (routes established automatically by virtue of having IP address enabled on the interface) to the OSPF routing process. Connected routes are redistributed as external to the AS.
- OSPF—Redistributes routes from another OSPF routing process. Choose the OSPF process ID from the list. If you choose this protocol, the Match options on this dialog box become visible. These options are not available when redistributing static, connected, RIP, or EIGRP routes. Skip to Step 5.
- RIP—Redistributes routes from the RIP routing process.
- BGP— Redistribute routes from the BGP routing process.
- EIGRP—Redistributes routes from the EIGRP routing process. Choose the autonomous system number of the EIGRP routing process from the list.

Step 5 If you have chosen OSPF for the source protocol, choose the conditions used for redistributing routes from another OSPF routing process into the selected OSPF routing process. These options are not available when redistributing static, connected, RIP, or EIGRP routes. The routes must match the selected condition to be redistributed. You can choose one or more of the following match conditions:

- Internal—The route is internal to a specific AS.
- External 1—Routes that are external to the autonomous system, but are imported into OSPF as Type 1 external routes.
- External 2—Routes that are external to the autonomous system, but are imported into OSPF as Type 2 external routes.
- NSSA External 1—Routes that are external to the autonomous system, but are imported into OSPF as Type 2 NSSA routes.
- NSSA External 2—Routes that are external to the autonomous system, but are imported into OSPF as Type 2 NSSA routes.

Step 6 In the **Metric Value** field, enter the metric value for the routes being redistributed. Valid values range from 1 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.

Note
Redistribution of static routes with a route-map matching a metric is not supported.

Step 7 Choose one of the following options for the Metric Type.

- If the metric is a Type 1 external route, choose **1**.
- If the metric is a Type 2 external route, choose **2**.

Step 8 Enter the tag value in the **Tag Value** field.

The tag value is a 32-bit decimal value attached to each external route that is not used by OSPF itself, but may be used to communicate information between ASBRs. Valid values range from 0 to 4294967295.

Step 9 Check the **Use Subnets** check box to enable the redistribution of subnetted routes. Uncheck this check box to cause only routes that are not subnetted to be redistributed.

Step 10 Choose the name of the route map to apply to the redistribution entry from the Route Map drop-down list.

Step 11 If you need to add or configure a route map, click **Manage**.

The Configure Route Map dialog box appears.

Step 12 Click **Add** or **Edit** to define which of the routes from the specified routing protocol are allowed to be redistributed into the target routing process. For more information, see Define a Route Map, on page 805.

Step 13 Click **OK**.

Configure Route Summarization When Redistributing Routes Into OSPFv2

When routes from other protocols are redistributed into OSPF, each route is advertised individually in an external LSA. However, you can configure the ASA 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.

Add a Route Summary Address

The Summary Address pane displays information about the summary addresses configured for each OSPF routing process.

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.

Note OSPF does not support summary-address 0.0.0.0 0.0.0.0.

To configure the software advertisement on one summary route for all redistributed routes included for a network address and mask, perform the following steps:

Procedure

Step 1 In the main ASDM home page, choose **Configuration** > **Device Setup** > **Routing** > **OSPF** > **Summary Address**.

Step 2 Click **Add**.

The Add OSPF Summary Address Entry dialog box appears. You can add new entries to existing entries in the Summary Address table. Some of the summary address information cannot be changed when editing an existing entry.

Step 3 Choose the specified OSPF Process ID associated with the summary address from the OSPF Process drop-down list. You cannot change this information when editing an existing entry.

Step 4 Enter the IP address of the summary address in the **IP Address** field. You cannot change this information when editing an existing entry.

Step 5 Choose the network mask for the summary address from the**Netmask** drop-down list. You cannot change this information when editing an existing entry.

Step 6 Check the **Advertise** check box to advertise the summary route. Uncheck this check box to suppress routes that fall under the summary address. By default, this check box is checked.

The Tag value displays 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.

Step 7 Click **OK**.

Add or Edit an OSPF Summary Address

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPF** > **Setup**.

Step 2 Click the **Route Summarization** tab.

The Add/Edit a Route Summarization Entry dialog box appears.

The Add/Edit a Route Summarization Entry dialog box allows you to add new entries to or modify existing entries in the Summary Address table. Some of the summary address information cannot be changed when editing an existing entry.

Step 3 Choose the specified OSPF Process ID associated with the summary address from the **OSPF Process** drop-down list. You cannot change this information when editing an existing entry.

Step 4 Enter the IP address of the summary address in the **IP Address** field. You cannot change this information when editing an existing entry.

Step 5 Enter the network mask for the summary address from the **Netmask** drop-down list. You cannot change this information when editing an existing entry.

Step 6 Check the **Advertise** check box to advertise the summary route. Uncheck this check box to suppress routes that fall under the summary address. By default, this check box is checked.

Configure Route Summarization Between OSPFv2 Areas

Route summarization is the consolidation of advertised addresses. This feature causes a single summary route to be advertised to other areas by an area boundary router. In OSPF, an area boundary router advertises networks in one area into another area. If the network numbers in an area are assigned in a way so that they are contiguous, you can configure the area boundary router to advertise a summary route that includes all the individual networks within the area that fall into the specified range.

To define an address range for route summarization, perform the following steps:

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPF** > **Setup**.

Step 2 Click the **Route Summarization** tab.

The Add/Edit a Route Summarization Entry dialog box appears.

The Add/Edit a Route Summarization Entry dialog box allows you to add new entries to or modify existing entries in the Summary Address table. Some of the summary address information cannot be changed when editing an existing entry.

Step 3 Enter the OSPF Area ID in the **Area ID** field. You cannot change this information when editing an existing entry.

Step 4 Enter the IP address of the summary address in the **IP Address** field. You cannot change this information when editing an existing entry.

Configure OSPFv2 Interface Parameters

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.

In ASDM, the Interface pane lets you configure interface-specific OSPF routing properties, such as OSPF message authentication and properties. There are two tabs that help you configure interfaces in OSPF:

- The Authentication tab displays the OSPF authentication information for the ASA interfaces.
- The Properties tab displays the OSPF properties defined for each interface in a table format.

To configure OSPFv2 interface parameters, perform the following steps:

Procedure

Step 1 Click the **Authentication** tab to display the authentication information for the ASA interfaces. Double-clicking a row in the table opens the Edit OSPF Authentication Interface dialog box for the selected interface.

Step 2 Click **Edit**.

The Edit OSPF Authentication Interface dialog box appears. The Edit OSPF Interface Authentication dialog box lets you configure the OSPF authentication type and parameters for the selected interface.

Step 3 Choose the Authentication type by clicking the relevant radio button:

- **No authentication** to disable OSPF authentication.
- **Area authentication, if defined** (Default) to use the authentication type specified for the area. See Configure OSPFv2 Area Parameters, on page 864 for information about configuring area authentication. Area authentication is disabled by default. Therefore, unless you have previously specified an area

authentication type, interfaces set to area authentication have authentication disabled until you configure this setting.

- **Password authentication** to use clear text password authentication (not recommended where security is a concern).
- **MD5 authentication** to use MD5 authentication.
- **Key chain authentication** to use key chain authentication (recommended). See Configure a Key Chain for Authentication, on page 854 for information about configuring key chain for authentication.

Step 4 If you have chosen password authentication, in the Authentication Password area, enter the password:

a) In the Enter Password field, type a text string of up to eight characters.
b) In the Re-enter Password field, retype the password.

Step 5 If you have chosen Key chain authentication, enter the key chain name in the Enter Key chain name field.

Step 6 Choose the settings for MD5 IDs and keys in the ID area, which includes the settings for entering the MD5 keys and parameters when MD5 authentication is enabled. All devices on the interface using OSPF authentication must use the same MD5 key and ID.

a) In the Key ID field, enter a numerical key identifier. Valid values range from 1 to 255. The Key ID displays for the selected interface.
b) In the Key field, enter an alphanumeric character string of up to 16 bytes. The key displays for the selected interface.
c) Click **Add** or **Delete** to add or delete the specified MD5 key to the MD5 ID and Key table.

Step 7 Click **OK.**

Step 8 Click the **Properties** tab.

Step 9 Choose the interface that you want to edit. Double-clicking a row in the table opens the Properties tab dialog box for the selected interface.

Step 10 Click **Edit**.

The Edit OSPF Interface Properties dialog box appears. The Interface field displays the name of the interface for which you are configuring OSPF properties. You cannot edit this field.

Step 11 Check or uncheck the **Broadcast** check box to specify 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.

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. See Define Static OSPFv2 Neighbors, on page 869 for more information.
- You need to define a static route pointing to the crypto endpoint. See Configure a Static Route, on page 789 for more information.
- 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.

Step 12 Configure the following options:

- Enter a value in the Cost field, which determines the cost of sending a packet through the interface. The default value is 10.
- In the Priority field, enter the OSPF router priority value.

  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.

  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.

  In multiple context mode, for shared interfaces, specify 0 to ensure the device does not become the designated router. OSPFv2 instances cannot form adjacencies with each other across shared interfaces.
- Check or uncheck the **MTU Ignore** check box.

  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 will not be established.
- Check or uncheck the **Database filter** check box.

  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.

Step 13 To enable BFD on this interface, from the **BFD** drop-down list, choose **Enable**. To enable BFD on all interfaces that are supporting OSPFv2, see Configure OSPFv2, on page 853.

Step 14 (Optional) Click **Advanced** to display the Edit OSPF Advanced Interface Properties dialog box, which lets you change the values for the OSPF hello interval, retransmit interval, transmit delay, and dead interval.

Typically, you only need to change these values from the defaults if you are experiencing OSPF problems on your network.

Step 15 In the Intervals section, enter values for the following:

- The Hello Interval, which specifies the interval, in seconds, between hello packets sent on an interface. 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. Valid values range from 1 to 8192 seconds. The default value is 10 seconds.
- The Retransmit Interval, which specifies the time, in seconds, between LSA retransmissions for adjacencies belonging to the interface. When a router sends an LSA to its neighbor, it keeps the LSA until it receives the acknowledgement message. If the router receives no acknowledgement, it will resend the LSA. Be conservative when setting this value, or needless retransmission can result. The value should be larger for serial lines and virtual links. Valid values range from 1 to 8192 seconds. The default value is 5 seconds.
- The Transmit Delay, which specifies the estimated time, in seconds, required to send an LSA packet on the interface. 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. Valid values range from 1 to 8192 seconds. The default value is 1 second.

Step 16 In the Detecting Lost Neighbors section, do one of the following:

- Click Configure interval within which hello packets are not received before the router declares the neighbor to be down. In the Dead Interval field, specify the interval, in seconds, in which no hello packets are received, causing neighbors to declare a router down. Valid values range from 1 to 8192 seconds. The default value of this setting is four times the interval that was set in the Hello Interval field.
- Click Send fast hello packets within 1 seconds dead interval. In the Hello multiplier field, specify the number of hello packets to be sent per second. Valid values are between 3 and 20.

Configure OSPFv2 Area Parameters

You can configure several OSPF area parameters. These area parameters (shown in the following task list) include setting authentication, defining stub areas, and assigning specific costs to the default summary route. 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 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPF** > **Setup**.

Step 2 Click the **Area/Networks** tab.

The Add OSPF Area dialog box appears.

Step 3 Choose one of the following Area Type options:

- **Normal** to make the area a standard OSPF area. This option is selected by default when you first create an area.
- **Stub** to make the area a stub area. Stub areas do not have any routers or areas beyond it. Stub areas prevent AS External LSAs (Type 5 LSAs) from being flooded into the stub area. When you create a stub area, you have the option of preventing summary LSAs (Types 3 and 4) from being flooded into the area by unchecking the Summary check box.
- **Summary** to prevent LSAs from being sent into the stub area when the area being defined is a stub area, uncheck this check box. By default, this check box is checked for stub areas.
- **NSSA** to make the area a not-so-stubby area. NSSAs accept Type 7 LSAs. When you create the NSSA, you have the option of preventing summary LSAs from being flooded into the area by unchecking the Summary check box. You can also disable route redistribution by unchecking the Redistribute check box and checking the Default Information Originate check box.

Step 4 Enter the IP address in the IP Address field of the network or host to be added to the area. Use **0.0.0.0** with a netmask of **0.0.0.0** to create the default area. You can only enter **0.0.0.0** in one area.

Step 5 Enter the network mask in the Network Mask field for the IP address or host to be added to the area. If adding a host, choose the **255.255.255.255** mask.

Step 6 Choose the OSPF Authentication type from the following options:

- **None** to disable OSPF area authentication. This is the default setting.
- **Password** to provide a clears text password for area authentication, which is not recommended where security is a concern.
- **MD5** to allow MD5 authentication.

Step 7 Enter a value in the Default Cost field to specify a default cost for the OSPF area.

Valid values range from 0 to 65535. The default value is 1.

Step 8 Click **OK**.

Configure OSPFv2 Filter Rules

Use the following procedure to filter routes or networks received or transmitted in OSPF updates.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Routing** > **OSPF** > **Filter Rules**.

Step 2 Click **Add**.

Step 3 Select the OSPF process ID in **OSPF AS**.

Step 4 Choose a standard access list from the Access List drop-down list. Click **Manage** to add a new ACL.

Step 5 Choose a direction from the Direction drop-down list. The direction will specify if the filter should be applied to inbound updates or outbound updates.

Step 6 For inbound filters, you can optionally specify an interface to limit the filter to updates received on that interface.

Step 7 For outbound filters, you can optionally specify what types of route are distributed.

a) Choose an option from the Protocol drop-down list.

You can choose a routing protocol, such as **BGP**, **EIGRP**, **OSPF**, or **RIP**.

Choose **Connected** to filter on peers and networks learned through connected routes.

Choose **Static** to filter on peers and networks learned through static routes.

b) If you chose BGP, EIGRP, or OSPF, also choose the **Process ID** for that protocol.

Step 8 Click **OK**.

Step 9 Click **Apply**.

Configure an OSPFv2 NSSA

The OSPFv2 implementation of an NSSA is similar to an OSPFv2 stub area. NSSA does not flood Type 5 external LSAs from the core into the area, but it can import autonomous system external routes in a limited way within the area.

NSSA imports Type 7 autonomous system external routes within an NSSA area by redistribution. These Type 7 LSAs are translated into Type 5 LSAs by NSSA ABRs, which are flooded throughout the whole routing domain. Summarization and filtering are supported during the translation.

You can simplify administration if you are an ISP or a network administrator that must connect a central site using OSPFv2 to a remote site that is using a different routing protocol with NSSA.

Before the implementation of NSSA, the connection between the corporate site border router and the remote router could not be run as an OSPFv2 stub area because routes for the remote site could not be redistributed into the stub area, and two routing protocols needed to be maintained. A simple protocol such as RIP was usually run and handled the redistribution. With NSSA, you can extend OSPFv2 to cover the remote connection by defining the area between the corporate router and the remote router as an NSSA.

Before you use this feature, consider these guidelines:

- You can set a Type 7 default route that can be used to reach external destinations. When configured, the router generates a Type 7 default into the NSSA or the NSSA area boundary router.
- Every router within the same area must agree that the area is NSSA; otherwise, the routers cannot communicate with each other.

Procedure

Step 1 From the main ASDM home page, choose **Configuration** > **Device Setup** > **Routing** > **OSPF** > **Setup**.

Step 2 Click the **Area/Networks** tab.

Step 3 Click **Add**.

The Add OSPF Area dialog box appears.

Step 4 Click the **NSSA** radio button in the Area Type area.

Choose this option to make the area a not-so-stubby area. NSSAs accept Type 7 LSAs. When you create the NSSA, you have the option of preventing summary LSAs from being flooded into the area by unchecking the Summary check box. You can also disable route redistribution by unchecking the Redistribute check box and checking the Default Information Originate check box.

Step 5 Enter the IP address in the IP Address field of the network or host to be added to the area. Use **0.0.0.0** with a netmask of **0.0.0.0** to create the default area. You can only enter **0.0.0.0** in one area.

Step 6 Enter the network mask in the Network Mask field for the IP address or host to be added to the area. If adding a host, choose the **255.255.255.255** mask.

Step 7 In the Authentication area, click the **None** radio button to disable OSPF area authentication.

Step 8 Enter a value in the Default Cost field to specify a default cost for the OSPF area.

Valid values range from 0 to 65535. The default value is 1.

Step 9 Click **OK**.

Configure an IP Address Pool for Clustering (OSPFv2 and OSPFv3)

You can assign a range of IPv4 addresses for the router ID cluster pool if you are using Individual Interface clustering.

To assign a range of IPv4 addresses for the router ID cluster pool in Individual Interface for OSPFv2, perform the following steps:

Procedure

Step 1 From the main ASDM home page, choose **Configuration** > **Device Setup** > **Routing** > **OSPF** > **Setup**.

Step 2 Click the **Process Instances** tab.

Step 3 Choose the OSPF process that you want to edit, then click **Advanced**.

The Edit OSPF Process Advanced Properties dialog box appears.

Step 4 Click the **Cluster Pool** radio button. If you are using clustering, then you do not need to specify an IP address pool for the router ID (that is, leave the field blank). If you do not enter an IP address pool, then the ASA uses the automatically generated router ID.

Step 5 Enter the name of the IP address pool, or click the ellipses to display the Select IP Address Pool dialog box.

Step 6 Double-click an existing IP address pool name to add it to the Assign field. Alternatively, click **Add** to create a new IP address pool.

The Add IPv4 Pool dialog box appears.

Step 7 Enter the new IP address pool name in the **Name** field.

Step 8 Enter the starting IP address or click the ellipses to display the Browse Starting IP Address dialog box.

Step 9 Double-click an entry to add it to the Starting IP Address field, then click **OK**.

Step 10 Enter the ending IP address or click the ellipses to display the Browse Ending IP Address dialog box

Step 11 Double-click an entry to add it to the Ending IP Address field, then click **OK**

Step 12 Choose the subnet mask from the drop-down list, then click **OK**.

The new IP address pool appears in the Select IP Address Pool list.

Step 13 Double-click the new IP address pool name to add it to the Assign field, then click **OK**.

The new IP address pool name appears in the Cluster Pool field of the Edit OSPF Process Advanced Properties dialog box.

Step 14 Click **OK**.

Step 15 If you want to change the newly added IP address pool settings, click **Edit**.

The Edit IPv4 Pool dialog box appears.

Step 16 Repeat Steps 4 through 14.

Note

You cannot edit or delete an existing IP address pool that has been assigned and is already being used by one or more connection profiles.

Step 17 Click **OK**.

Step 18 To assign a range of IPv4 addresses for the router ID cluster pool in Individual Interface clustering for OSPFv3, perform the following steps:

a) From the main ASDM home page, choose **Configuration** > **Device Setup** > **Routing** > **OSPFv3** > **Setup**.
b) Click the **Process Instances** tab.
c) Choose the OSPF process that you want to edit, then click **Advanced**.

The Edit OSPFv3 Process Advanced Properties dialog box appears.

d) Choose the Cluster Pool option from the Router ID drop-down list. If you do not need to specify an IP address pool for the router ID, choose the Automatic option. If you do not configure an IP address pool, then the ASA uses the automatically generated router ID.
e) Enter the IP address pool name. Alternatively, click the ellipses to display the Select IP Address Pool dialog box.
f) Double-click an existing IP address pool name to add it to the Assign field. Alternatively, click **Add** to create a new IP address pool.

The Add IPv4 Pool dialog box appears.

g) Enter the new IP address pool name in the **Name** field.
h) Enter the starting IP address or click the ellipses to display the Browse Starting IP Address dialog box.
i) Double-click an entry to add it to the Starting IP Address field, then click **OK**.
j) Enter the ending IP address or click the ellipses to display the Browse Ending IP Address dialog box.
k) Double-click an entry to add it to the Ending IP Address field, then click **OK**.
l) Choose the subnet mask from the drop-down list, then click **OK**.

The new IP address pool appears in the Select IP Address Pool list.

m) Double-click the new IP address pool name to add it to the Assign field, then click **OK**.

The new IP address pool name appears in the Cluster Pool field of the Edit OSPF Process Advanced Properties dialog box.

n) Click **OK**.
o) If you want to change the newly added cluster pool settings, click **Edit**.

The Edit IPv4 Pool dialog box appears.

p) Repeat Steps 4 through 14.

Note
You cannot edit or delete an existing IP address pool that has been assigned and is already being used by another OSPFv3 process.

q) Click **OK**.

Define Static OSPFv2 Neighbors

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.

Before you begin, you must create a static route to the OSPFv2 neighbor. See Configure a Static Route, on page 789 for more information about creating static routes.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPF** > **Static Neighbor**.

Step 2 Click **Add** or **Edit**.

The Add/Edit OSPF Neighbor Entry dialog box appears. This dialog box lets you define a new static neighbor or change information for an existing static neighbor. You must define a static neighbor for each point-to-point, nonbroadcast interface. Note the following restrictions:

- You cannot define the same static neighbor for two different OSPF processes.
- You need to define a static route for each static neighbor.

Step 3 From the OSPF Process drop-down list, choose the OSPF process associated with the static neighbor. If you are editing an existing static neighbor, you cannot change this value.

Step 4 In the **Neighbor** field, enter the IP address of the static neighbor.

Step 5 In the**Interface** field, choose the interface associated with the static neighbor. If you are editing an existing static neighbor, you cannot change this value.

Step 6 Click **OK**.

Configure Route Calculation Timers

You can configure the delay time between when OSPFv2 receives a topology change and when it starts an SPF calculation. You also can configure the hold time between two consecutive SPF calculations.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPF** > **Setup**.

Step 2 Click the **Process Instances** tab.

Step 3 Choose the OSPF process that you want to edit, then click **Advanced**.

The Edit OSPF Process Advanced Properties dialog box appears.

Step 4 The Timers area allows you to modify the settings that are used to configure LSA pacing and SPF calculation timers. In the Timers area, enter the following values:

- The Initial SPF Delay, specifies the time (in milliseconds) between when OSPF receives a topology change and when the SPF calculation starts. Valid values range from 0 to 600000 milliseconds.
- The Minimum SPF Hold Time, specifies the hold time (in milliseconds)between consecutive SPF calculations.Valid values range from 0 to 600000 milliseconds.
- The Maximum SPF Wait Time, specifies the maximum wait time between two consecutive SPF calculations. Valid values range from 0 to 600000 milliseconds.

Step 5 Click **OK**.

Log Neighbors Going Up or Down

By default, a syslog message is generated when an OSPFv2 neighbor goes up or down.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPF** > **Setup**.

Step 2 Click the **Process Instances** tab.

Step 3 Click **Advanced**.

The Edit OSPF Process Advanced Properties dialog box appears.

Step 4 The Adjacency Changes area includes settings that define the adjacency changes that cause syslog messages to be sent. In the Adjacency Changes area, enter the following values:

- Check the **Log Adjacency Changes** check box to cause the ASA to send a syslog message whenever an OSPFv2 neighbor goes up or down. This setting is checked by default.
- Check the **Log Adjacency Changes Detail** check box to cause the ASA 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.

Step 5 Click **OK**.

Note
Logging must be enabled for the neighbor up or down messages to be sent.

Configure a Key Chain for Authentication

To enhance data security and protection of devices, you can enable rotating keys for authenticating IGP peers. 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. If the key lifetime expires and no active keys are found, OSPF uses the last valid key to maintain the adjacency with the peers.

This section describes how to create a key chain for OSPF peer authentication. This section also covers steps to add or edit the key chain attributes. After configuring a key chain object, you can use it in defining the OSPFv2 authentication for an interface and for a virtual link. Use the same authentication type (MD5 or Key Chain) and key ID for the peers to establish a successful adjacency. To learn how to define authentication for an interface, see Configure OSPFv2 Interface Parameters, on page 861; for a virtual link, see Configure a Virtual Link in OSPF, on page 873.

To configure a key chain, perform the following steps:

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Key Chain** .

Step 2 In the **Configure Key Chain** section, click **Add**.

Step 3 Enter the key chain name in the **Add Key Chain** dialog box, and click **Ok**.

The created key chain name is listed in the **Configure Key Chain** grid.

Step 4 Select the key chain name from the **Configure Key Chain** section, and in the **Configure Key** section, click **Add**. To edit an existing key, seclect the key name and click **Edit**.

The **Add Key** or **Edit Key**dialog box appears, depending on the action that you have selected.

Step 5 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.

Note
You cannot edit a saved key id.

Step 6 From the **Cryptographic Algorithm** drop-down, choose **MD5**. MD5 is the only algorithm supported for authenticating the key chain.

Step 7 Select the encryption type by clicking the **Plain Text** or **Encrypted** radio button, and then enter the password in the **Authentication Key** 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 8 Provide the lifetime values in the **Accept Lifetime** and **Send Lifetime** fields:

You can specify the time interval for the device to accept/send the key during key exchange with another device. The end time can be the duration, the absolute time when the accept/send lifetime ends, or never expires.

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 end lifetime.

Step 9 To save the key chain attibutes, click **Ok**. In the **Key Chain** page, click **Appy**.

What to do next

You can now apply the configured key chain to define the OSPFv2 authentication for an interface and for virtual link.

- Configure OSPFv2 Interface Parameters, on page 861
- Configure a Virtual Link in OSPF, on page 873

Configure Filtering in OSPF

The Filtering pane displays the ABR Type 3 LSA filters that have been configured 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. This type of area filtering can be applied out of a specific OSPF area, into a specific OSPF area, or into and out of the same OSPF areas at the same time.

OSPF ABR Type 3 LSA filtering improves your control of route distribution between OSPF areas.



Note Only Type 3 LSAs that originate from an ABR are filtered.

To configure filtering in OSPF, perform the following steps:

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPF** > **Filtering**.

Step 2 Click **Add** or **Edit**.

The Add or Edit OSPF Filtering Entry dialog box lets you add new filters to the Filter table or modify an existing filter. Some of the filtering information cannot be changed when you edit an existing filter.

Step 3 Choose the OSPF process that is associated with the filter entry from the OSPF Process drop-down list.

Step 4 Choose the Area ID that is associated with the filter entry from the Area ID drop-down list. If you are editing an existing filter entry, you cannot modify this setting.

Step 5 Choose a prefix list from the Prefix List drop-down list.

Step 6 Choose the traffic direction being filtered from the Traffic Direction drop-down list.

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 **Manage** to display the Configure Prefix Lists dialog box, from which you can add, edit, or delete prefix lists and prefix rules. For more information, see Configure Prefix Lists, on page 809 and the Configure the Metric Values for a Route Action, on page 810.

Step 8 Click **OK**.

Configure a Virtual Link in OSPF

If you add an area to an OSPF network, and it is not possible to connect the area directly to the backbone area, you need to create a virtual link. A virtual link connects two OSPF devices that have a common area, called the transit area. One of the OSPF devices must be connected to the backbone area.

To define new virtual links or change the properties of existing virtual links, perform the following steps:

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPF** > **Virtual Link**.

Step 2 Click **Add** or **Edit**.

The Add or Edit OSPF Virtual Link dialog box appears, which allows you to define new virtual links or change the properties of existing virtual links.

Step 3 Choose the OSPF process ID that is associated with the virtual link from the OSPF Process drop-down list. If you are editing an existing virtual link entry, you cannot modify this setting.

Step 4 Choose the Area ID that is associated with the virtual link from the Area ID drop-down list.

Choose the area shared by the neighbor OSPF devices. The selected area cannot be an NSSA or a Stub area. If you are editing an existing virtual link entry, you cannot modify this setting.

Step 5 In the **Peer Router ID** field, enter the router ID of the virtual link neighbor.

If you are editing an existing virtual link entry, you cannot modify this setting.

Step 6 Click **Advanced** to edit advanced virtual link properties,

The Advanced OSPF Virtual Link Properties dialog box appears. You can configure the OSPF properties for the virtual link in this area. These properties include authentication and packet interval settings.

Step 7 In the Authentication area, choose the Authentication type by clicking the radio button next to one of the following options:

- **No authentication** to disable OSPF authentication.
- **Password authentication** to use clear text password authentication (not recommended where security is a concern).
- **MD5 authentication** to use MD5 authentication.
- **Key chain authentication** to use key chain authentication (recommended). See Configure a Key Chain for Authentication, on page 854 for information about configuring key chain for authentication.

Step 8 In the Authentication Password area, enter and re-enter a password when password authentication is enabled. Passwords must be a text string of up to 8 characters.

Step 9 In the MD5 IDs and Key area, enter the MD5 keys and parameters when MD5 authentication is enabled. All devices on the interface using OSPF authentication must use the same MD5 key and ID. Specify the following settings:

a) In the **Key ID**field, enter a numerical key identifier. Valid values range from 1 to 255. The Key ID displays for the selected interface.

b) In the **Key** field, enter an alphanumeric character string of up to 16 bytes. The Key ID displays for the selected interface.
c) Click **Add** or **Delete** to add or delete the specified MD5 key to the MD5 ID and Key table.

Step 10 In the Interval area, specify the interval timing for the packet by choosing from the following options:

- **Hello Interval** to specify the interval, in seconds, between hello packets sent on an interface. The smaller the hello interval, the faster topological changes are detected, but the more traffic is sent on the interface. This value must be the same for all routers and access servers on a specific interface. Valid values range from 1 to 65535 seconds. The default value is 10 seconds.
- **Retransmit Interval** to specify the time, in seconds, between LSA retransmissions for adjacencies belonging to the interface. When a router sends an LSA to its neighbor, it keeps the LSA until it receives the acknowledgement message. If the router receives no acknowledgement, it will resend the LSA. Be conservative when setting this value, or needless retransmission can result. The value should be larger for serial lines and virtual links. Valid values range from 1 to 65535 seconds. The default value is 5 seconds.
- **Transmit Delay** to specify the estimated time, in seconds, required to send an LSA packet on the interface. 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. Valid values range from 1 to 65535 seconds. The default value is 1 second.
- **Dead Interval** to specify the interval, in seconds, in which no hello packets are received, causing neighbors to declare a router down. Valid values range from 1 to 65535. The default value of this field is four times the interval set by the Hello Interval field.

Step 11 Click **OK**.

Configure OSPFv3

This section describes the tasks involved in configuring an OSPFv3 routing process.

Enable OSPFv3

To enable OSPFv3, you need to create an OSPFv3 routing process, create an area for OSPFv3, enable an interface for OSPFv3, then redistribute the route into the targeted OSPFv3 routing processes.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPFv3** > **Setup**.
Step 2 On the Process Instances tab, check the **Enable OSPFv3 Process** check box. You can enable up to two OSPF process instances. Only single context mode is supported.
Step 3 Enter a process ID in the Process ID field. The ID can be any positive integer.

Step 4 To enable BFD on all interfaces that support OSPFv3, click **Advanced**. In the **Edit OPPFv3 Process Advanced Properties** window, under **Enable BFD on all interfaces**, click the **Enable BFD** check box. To enable BFD on a specific OSPFv3 intreface, see Configure OSPFv3 Interface Parameters, on page 875.

Step 5 Click **Apply** to save your changes.

Step 6 To continue, see Configure OSPFv3 Area Parameters, on page 876.

Configure OSPFv3 Interface Parameters

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.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPFv3** > **Interfaces**.

Step 2 Click the **Authentication** tab.

Step 3 To specify the authentication parameters for an interface, select the interface and click **Edit**.

The **Edit OSPFv3 Interface Authentication** dialog box appears.

Step 4 Choose the authentication type from the **Authentication Type** drop-down list. The available options are Area, Interface, and None. The None option indicates that no authentication is used.

Step 5 Choose the authentication algorithm from the **Authentication Algorithm** drop-down list. Supported values are SHA-1 and MD5.

Step 6 Enter the authentication key in the **Authentication Key** field. 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.

Step 7 Choose the encryption algorithm from the **Encryption Algorithm** drop-down list. Supported values are AES-CDC, 3DES, and DES. The NULL entry indicates no encryption.

Step 8 Enter the encryption key in the **Encryption Key** field.

Step 9 Click **OK**.

Step 10 Click the **Properties** tab.

Step 11 Select the interface whose properties you want to modify, and click **Edit**.

The Edit OSPFv3 Interface Properties dialog box appears.

Step 12 Check the **Enable OSPFv3 on this interface** check box.

Step 13 Choose the process ID from the drop-down list.

Step 14 Choose the area ID from the drop-down list.

Step 15 (Optional) Specify 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 16 Choose the network type from the drop-down list. Supported options are Default, Broadcast, and Point-to-Point.

Step 17 Enter the cost of sending a packet on an interface in the Cost field.

Step 18 Enter the router priority, which helps determine the designated router for a network. in the Priority field. Valid values range from 0 to 255.

Step 19 To enable BFD on this interface, from the **BFD Configuration** drop-down list, choose **Enable**. To enable BFD on all interfaces that support OSPFv3, see Enable OSPFv3, on page 874.

Step 20 Check the **Disable MTU mismatch detection** check box to disable the OSPF MTU mismatch detection when DBD packets are received. OSPF MTU mismatch detection is enabled by default.

Step 21 Check the **Filter outgoing link state advertisements** check box to filter outgoing LSAs to an OSPFv3 interface. All outgoing LSAs are flooded to the interface by default.

Step 22 Check the **OSPF Flood Reduction** check box to reduce unnecessary flooding and refreshing of LSAs to the interface.

Step 23 In the **Timers** area, in the **Dead Interval** field, enter the 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.

Step 24 In the **Hello Interval** field, enter the interval in seconds between hello packets sent on the interface. The value must be the same for all nodes on a specific network and can range from 1 to 65535. The default interval is 10 seconds for Ethernet interfaces and 30 seconds for non-broadcast interfaces.

Step 25 In the **Retransmit Interval** field, enter the 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.

Step 26 In the **Transmit Delay** field, enter the 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.

Step 27 Click **OK**.

Step 28 Click **Apply** to save your changes.

Configure OSPFv3 Area Parameters

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPFv3** > **Setup**.

Step 2 Click the **Areas** tab.

Step 3 To add a new area, click **Add**. To modify an existing area, click **Edit**. To remove a selected area, click **Delete**.

The Add OSPFv3 Area dialog box or Edit OSPFv3 Area dialog box appears.

Step 4 From the OSPFv3 Process ID drop-down list, choose the process ID.

Step 5 Enter the area ID, which specifies the area for which routes are to be summarized, in the Area ID field.

Step 6 Choose the area type from the Area Type drop-down list. Available options are Normal, NSSA, and Stub.

Step 7 To allow the sending of summary LSAs into the area, check the **Allow sending of summary LSAs into the area** check box.

Step 8 To allow redistribution to import routes to normal and not so stubby areas, check the **Redistribution imports routes to normal and NSSA areas** check box.

Step 9 To generate a default external route into an OSPFv3 routing domain, check the **Default information originate** check box.

Step 10 Enter the metric used for generating the default route in the Metric field. The default value is 10. Valid metric values range from 0 to 16777214.

Step 11 Choose the metric type from the Metric Type drop-down list. 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 12 Enter the cost in the Default Cost field.

Step 13 Click **OK**.

Step 14 Click the **Route Summarization** tab.

Step 15 To specify a new range for consolidating and summarizing routes, click **Add**. To modify an existing range for consolidating and summarizing routes, click **Edit**.

The Add Route Summarization dialog box or Edit Route Summarization dialog box appears.

Step 16 Choose the process ID from the Process ID drop-down list.

Step 17 Choose the area ID from the Area ID drop-down list.

Step 18 Enter the IPv6 prefix and prefix length in the IPv6 Prefix/Prefix Length field.

Step 19 (Optional) Enter 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.

Step 20 Check the **Advertised** check box to set the address range status to advertised and generate a Type 3 summary LSA.

Step 21 Click **OK**.

Step 22 To continue, see Configure a Virtual Link Neighbor, on page 877.

Configure a Virtual Link Neighbor

To configure a virtual link neighbor, perform the following steps:

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPFv3** > **Virtual Link**.

Step 2 To add a new virtual link neighbor, click **Add**. To modify an existing virtual link neighbor, click **Edit**. To remove a selected virtual link neighbor, click **Delete**.

The Add Virtual Link dialog box or Edit Virtual Link dialog box appears.

Step 3 Choose the process ID from the Process ID drop-down list.

Step 4 Choose the area ID from the Area ID drop-down list.

Step 5 Enter the peer router ID (that is, the IP address) in the Peer Router ID field.

Step 6 (Optional) Enter the time-to-live (TTL) security hop count on a virtual link in the TTL Security field. The hop count value can range from 1 to 254.

Step 7 In the Timers area, enter the time in seconds that hello packets are not seen before a neighbor indicates that the router is down in the Dead Interval field. 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 to 8192.

Step 8 Enter the time in seconds between the hello packets that are sent on an interface in the Hello Interval field. 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 that are attached to a common network. Valid values range from 1 to 8192. The default is 10.

Step 9 Enter the time in seconds between LSA retransmissions for adjacencies that belong to the interface in the Retransmit Interval field. 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 8192. The default is 5.

Step 10 Enter the estimated time in seconds that is required to send a link-state update packet on the interface in the Transmit Delay field. The integer value must be greater than zero. LSAs in the update packet have their own ages incremented by this amount before transmission. The range of values can be from 1 to 8192. The default is 1.

Step 11 In the Authentication area, check the **Enable Authentication** check box to enable authentication.

Step 12 Enter the security policy index, which must be a number from 256 to 4294967295, in the Security Policy Index field.

Step 13 Choose the authentication algorithm from the Authentication Algorithm drop-down list. Supported values are SHA-1 and MD5. 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.

Step 14 Enter the authentication key in the Authentication Key field. The key must include 32 hexadecimal characters.

Step 15 Choose the encryption algorithm from the Encryption Algorithm drop-down list. Supported values are AES-CDC, 3DES, and DES. The NULL entry indicates no encryption.

Step 16 Enter the encryption key in the Encryption Key field.

Step 17 Click **OK**.

Step 18 Click **Apply** to save your changes.

Configure OSPFv3 Passive Interfaces

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPFv3** > **Setup**.

Step 2 Click the **Process Instances** tab.

Step 3 Choose the OSPFv3 process that you want to edit, then click **Advanced**.

The Edit OSPFv3 Process Advanced Properties dialog box appears.

Step 4 The Passive Interfaces area allows you to enable passive OSPFv3 routing on an interface. Passive routing assists in controlling the advertisement of OSPFv3 routing information and disables the sending and receiving of OSPFv3 routing updates on an interface. In the Passive Interfaces area, choose the following settings:

- Check the **Global passive** check box to make all of the interfaces listed in the table passive. Uncheck individual interfaces to make them non-passive.
- Uncheck the **Global passive** check box to make all of the interfaces non-passive. Check individual interfaces to make them passive.

Step 5 Click **OK**.

Step 6 Click **Apply** to save your changes.

Configure OSPFv3 Administrative Distance

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPFv3** > **Setup**.

Step 2 Click the **Process Instances** tab.

Step 3 Choose the OSPF process that you want to edit, then click **Advanced**.

The Edit OSPFv3 Process Advanced Properties dialog box appears.

The Administrative Route Distances area allows you to modify the settings that were used to configure administrative route distances. The administrative route distance is an integer from 10 to 254. In the Administrative Route Distances area, enter the following values:

- The Inter Area, which specifies the inter-area routes for OSPV for IPv6 routes.
- The Intra Area, which specifies the intra-area routes for OSPF for IPv6 routes.
- The External, which specifies the external type 5 and type 7 routes for OSPF for IP6 routes.

Step 4 Click **OK**.

Step 5 Click **Apply** to save your changes.

Configure OSPFv3 Timers

You can set LSA arrival, LSA pacing, and throttling timers for OSPFv3.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPFv3** > **Setup**.

Step 2 Click the **Process Instances** tab.

Step 3 Choose the OSPFv3 process that you want to edit, then click **Advanced**.

The Edit OSPFv3 Process Advanced Properties dialog box appears.

Step 4 The Timers area allows you to modify the settings that are used to configure LSA arrival, LSA pacing, LSA retransmission, LSA throttle, and SPF throttle times. In the Timers area, enter the following values:

- The LSA Arrival, which 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.

- The LSA Flood Pacing, which 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.
- The LSA Group Pacing, which 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.
- The LSA Retransmission Pacing, which 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.
- The LSA Throttle Initial, which specifies the delay in milliseconds to generate the first occurrence of the LSA. The default value is 0 milliseconds.
- The LSA Throttle Min Hold, which specifies the minimum delay in milliseconds to originate the same LSA. The default value is 5000 milliseconds.
- The LSA Throttle Max Wait, which 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.

- The SPF Throttle Initial, specifies the delay in milliseconds to receive a change to the SPF calculation. The default value is 5000 milliseconds.
- The SPF Throttle Min Hold, which specifies the delay in milliseconds between the first and second SPF calculations. The default value is 10000 milliseconds.
- The SPF Throttle Max Wait, which 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 5 Click **OK**.

Step 6 Click **Apply** to save your changes.

Define Static OSPFv3 Neighbors

You need to define static OSPFv3 neighbors to advertise OSPF routes over a point-to-point, non-broadcast network. This feature lets you broadcast OSPFv3 advertisements across an existing VPN connection without having to encapsulate the advertisements in a GRE tunnel.

Before you begin, you must create a static route to the OSPFv3 neighbor. See Configure a Static Route, on page 789 for more information about creating static routes.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPFv3** > **Static Neighbor**.

Step 2 Click **Add** or **Edit**.

The Add or Edit Static Neighbor dialog box appears. This dialog box lets you define a new static neighbor or change information for an existing static neighbor. You must define a static neighbor for each point-to-point, nonbroadcast interface. Note the following restrictions:

- You cannot define the same static neighbor for two different OSPFv3 processes.
- You need to define a static route for each static neighbor.

Step 3 From the Interface drop-down list, choose the interface associated with the static neighbor. If you are editing an existing static neighbor, you cannot change this value.

Step 4 In the Link-local Address field, enter the IPv6 address of the static neighbor.

Step 5 (Optional) In the Priority field, enter the priority level.

Step 6 (Optional) In the Poll Interval field, enter the poll interval in seconds.

Step 7 Click **OK**.

Send Syslog Messages

Configure the router to send a syslog message when an OSPFv3 neighbor goes up or down.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPFv3** > **Setup**.

Step 2 Click the **Process Instances** tab.

Step 3 Choose the OSPF process that you want to edit, then click **Advanced**.

The Edit OSPFv3 Process Advanced Properties dialog box appears.

The Adjacency Changes area allows you to modify the settings for sending syslog messages when an OSPFv3 neighbor goes up or down. In the Adjacency Changes area, do the following:

- To send a syslog message when an OSPFv3 neighbor goes up or down, check the **Log Adjacency Changes** check box.
- To sends a syslog message for each state, not only when an OSPFv3 neighbor goes up or down, check the **Include Details** check box.

Step 4 Click **OK**.

Step 5 Click **Apply** to save your changes.

Suppress Syslog Messages

To suppress the sending of syslog messages when the route receives unsupported LSA Type 6 multicast OSPF (MOSPF) packets, perform the following steps:

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPFv3** > **Setup**.

Step 2 Click the **Process Instances** tab.

Step 3 Choose the OSPFv3 process that you want to edit, then click **Advanced**.

The Edit OSPFv3 Process Advanced Properties dialog box appears.

Step 4 Check the **Ignore LSA MOSPF** check box, then click **OK**.

Calculate Summary Route Costs

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPFv3** > **Setup**.

Step 2 Click the **Process Instances** tab.

Step 3 Choose the OSPF process that you want to edit, then click **Advanced**.

The Edit OSPFv3 Process Advanced Properties dialog box appears.

Step 4 Check the **RFC1583 Compatible** check box, then click **OK**.

Generate a Default External Route into an OSPFv3 Routing Domain

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPFv3** > **Setup**.

Step 2 Click the **Process Instances** tab.

Step 3 Choose the OSPFv3 process that you want to edit, then click **Advanced**.

The Edit OSPFv3 Process Advanced Properties dialog box appears.

Step 4 In the Default Information Originate Area, do the following:

a) Check the **Enable** check box to enable the OSPFv3 routing process.

b) Check the **Always advertise** check box to always advertise the default route, whether or not one exists.

c) Enter the metric used for generating the default route in the Metric field. Valid metric values range from 0 to 16777214. The default value is 10.

d) From the Metric Type drop-down list, choose the external link type that is associated with the default route that is advertised into the OSPFv3 routing domain. Valid values are the following:

- 1—Type 1 external route
- 2—Type 2 external route

The default is the Type 2 external route.

e) From the Route Map drop-down list, choose the routing process that generates the default route if the route map is satisfied.

Step 5 Click **OK**.

Step 6 Click **Apply** to save your changes.

Configure an IPv6 Summary Prefix

Procedure

Step 1 In the ASDM main window, choose **Configuration** > **Device Setup** > **Routing** > **OSPFv3** > **Summary Prefix**.

Step 2 To add a new summary prefix, click **Add**. To modify an existing summary prefix, click **Edit**. To remove a summary prefix, click **Delete**.

The Add Summary Prefix dialog box or Edit Summary Prefix dialog box appears.

Step 3 Choose the process ID from the Process ID drop-down list.

Step 4 Enter the IPv6 prefix and prefix length in the IPv6 Prefix/Prefix Length field.

Step 5 Check the **Advertise** check box to advertise routes that match the specified prefix and mask pair. Uncheck this check box to suppress routes that match the specified prefix and mask pair.

Step 6 Enter the tag value that you can use as a match value for controlling redistribution through route maps in the Tag field.

Step 7 Click **OK**.

Step 8 Click **Apply** to save your changes.

Redistribute IPv6 Routes

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPFv3** > **Redistribution**.

Step 2 To add new parameters for redistributing connected routes into an OSPFv3 process, click **Add**. To modify existing parameters for redistributing connected routes into an OSPFv3 process, click **Edit**. To remove a selected set of parameters, click **Delete**.

The Add Redistribution dialog box or Edit Redistribution dialog box appears.

Step 3 Choose the process ID from the Process ID drop-down list.

Step 4 Choose the source protocol from which routes are being redistributed from the Source Protocol drop-down list. The supported protocols are connected, static, and OSPF.

Step 5 Enter the metric value in the Metric field. When redistributing routes from one OSPF process into another OSPF process on the same router, the metric is carried through from one process to the other if no metric value is specified. When redistributing other processes into an OSPF process, the default metric is 20 when no metric value is specified.

Step 6 Choose the metric type from the Metric Type drop-down list. The available options are None, 1, and 2.

Step 7 (Optional) Enter the tag value in the Tag field. This parameter specifies the 32-bit decimal value attached to each external route, 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.

Step 8 Choose the route map from the Route Map drop-down list to check 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.

Step 9 To include connected routes in the redistribution, check the **Include connected** check box.

Step 10 Check the **Match** check box to redistribute routes into other routing domains, then check one of the following check boxes:

- **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 11 Click **OK**.

Step 12 Click **Apply** to save your changes.

Configure Graceful Restart

The ASA 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.

In a high availability mode, the OSPF process restarts when the active unit becomes inactive and the standby unit becomes the new active. Similarly, in a cluster mode, the OSPF process restarts when the control unit becomes inactive and the data unit is elected as the new control unit. Such OSPF transitioning processes involve a considerable amount of delay. You can configure NSF to avoid traffic loss during the OSPF process state change. The NSF capability is also useful when there is a scheduled hitless software upgrade.

Graceful restart is supported on both OSPFv2 and OSPFv3. You can configure graceful restart on OSPFv2 by using either using NSF Cisco (RFC 4811 and RFC 4812) or NSF IETF (RFC 3623). You can configure graceful restart on OSPFv3 using graceful-restart (RFC 5187).

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.



Note When fast hellos are configured for OSPFv2, graceful restart does not occur when the active unit reloads and the standby unit becomes active. This is because the time taken for the role change is more than the configured dead interval.

Configuring Graceful Restart for OSPFv2

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.

Configure Cisco NSF Graceful Restart for OSPFv2

Configure Cisco NSF Graceful Restart for OSPFv2, for a NSF-capable or NSF-aware device.

Procedure

Step 1 In the main ASDM window, choose Configuration > Device Setup > Routing > OSPF > Setup > Advanced > Add NSF Properties.

Step 2 Under Configuring Cisco NSF, check the Enable Cisco nonstop forwarding (NSF) check box.

Step 3 (Optional) Check the Cancels NSF restart when non-NSF-aware neighboring networking devices are detected check box if required.

Step 4 (Optional) Under Configuring Cisco NSF helper, uncheck the Enable Cisco nonstop forwarding (NSF) for helper mode check box.

Note
This is checked by default. Uncheck this to disable the Cisco NSF helper mode on NSF-aware device.

Step 5 Click **OK**.

Step 6 Click **Apply** to save your changes.

Configure IETF NSF Graceful Restart for OSPFv2

Configure IETF NSF Graceful Restart for OSPFv2, for a NSF-capable or NSF-aware device.

Procedure

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPF > Setup > Advanced > Add NSF Properties**.

Step 2 Under Configuring IETF NSF, check the Enable IETF nonstop forwarding (NSF) check box.

Step 3 (Optional) Enter the restart interval in seconds in the Length of graceful restart interval field.

Note
The default value is 120 seconds. For a restart interval below 30 seconds, graceful restart will be terminated.

Step 4 (Optional) Under Configuring IETF NSF helper, uncheck the Enable IETF nonstop forwarding (NSF) for helper mode check box.

This is checked by default. Uncheck this to disable the IETF NSF helper mode on NSF-aware device.

Step 5 Click **OK**.

Step 6 Click **Apply** to save your changes.

Configuring Graceful Restart for OSPFv3

Configuring the NSF graceful-restart feature for OSPFv3 involves two steps; configuring a device to be NSF-capable and then configuring a device to be NSF-aware.

Procedure

Step 1 In the main ASDM window, choose **Configuration > Device Setup > Routing > OSPFv3 > Setup > Advanced > Add NSF Properties**.

Step 2 Under Configuring Graceful Restart, check the Enable Graceful Restart check box.

Step 3 (Optional) Enter a value for the restart interval in the Restart Interval field.

Note
The default value is 120 seconds. For a restart interval below 30 seconds, graceful restart will be terminated.

Step 4 Under Configuring Graceful Restart Helper, check the Enable Graceful Restart Helper check box.

This is checked by default. Uncheck this to disable the Graceful-restart helper mode on a NSF-aware device.

Step 5 (Optional) Check the Enable LSA checking 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 6 Click **OK**.

Step 7 Click **Apply** to save your changes.

Configuring Graceful Restart Wait Timer for OSPF

OSPF routers are expected to set the RS-bit in the EO-TLV attached to a Hello packet when it is not known that all neighbors are listed in the packet, but the restarting routers require to preserve their adjacencies. However, the RS-bit value must not be longer than the RouterDeadInterval seconds. Hence the **timers nsf wait** command is introduced to set the RS-bit in Hello packets lesser than RouterDeadInterval seconds. The default value of NSF wait timer is 20 seconds.

Before you begin

- To configure Cisco NSF wait time for OSPF, the device must be NSF-aware or NSF-capable.

Procedure

Step 1 Enter into OSPF router configuration mode.

Example:

```
ciscoasa(config)# router ospf
```

Step 2 Enter timers and specify nsf.

Example:

```
ciscoasa(config-router)# timers?
router mode commands/options:
  lsa       OSPF LSA timers
  nsf       OSPF NSF timer
  pacing    OSPF pacing timers
  throttle  OSPF throttle timers
ciscoasa(config-router)# timers nsf ?
```

Step 3 Enter the graceful restart wait interval. This value can range between 1 and 65535.

Example:

```
ciscoasa(config-router)# timers nsf wait 200
```

By using the graceful restart wait interval, you can ensure that the wait interval is not longer than the router dead interval.

Remove the OSPFv2 Configuration

Remove the OSPFv2 configuration.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPF** > **Setup**.

Step 2 Uncheck the **Enable this OSPF Process** check box.

Step 3 Click **Apply**.

Remove the OSPFv3 Configuration

Remove the OSPFv3 configuration.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPFv3** > **Setup**.

Step 2 Uncheck the **Enable OSPFv3 Process** check box.

Step 3 Click **Apply**.

Example for OSPFv2

The following example shows how to enable and configure OSPFv2 with various optional processes:

1. In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPF** > **Setup**.
2. Click the **Process Instances** tab and in the OSPF Process 1 field, type **2**.
3. Click the **Area/Networks** tab, and click **Add**.
4. Enter **0** in the Area ID field.
5. In the Area Networks area, enter **10.0.0.0** in the IP Address field.
6. Choose 255.0.0.0 from the Netmask drop-down list.
7. Click **OK**.
8. In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPF** > **Redistribution**.
9. Click **Add**.

   The Add/Edit OSPF Redistribution Entry dialog box appears.

10. In the Protocol area, click the **OSPF** radio button to choose the source protocol from which the routes are being redistributed. Choosing OSPF redistributes routes from another OSPF routing process.
11. Choose the OSPF process ID from the OSPF Process drop-down list.
12. In the Match area, check the **Internal** check box.
13. In the Metric Value field, enter **5** for the metric value for the routes being redistributed.
14. From the Metric Type drop-down list, choose 1 for the Metric Type value.
15. From the Route Map drop-down list, choose 1.
16. Click **OK**.
17. In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPF** > **Interface**.
18. From the Properties tab, choose the **inside** interface and click **Edit.**

    The Edit OSPF Properties dialog box appears.
19. In the Cost field, enter **20**.
20. Click **Advanced**.
21. In the Retransmit Interval field, enter **15**.
22. In the Transmit Delay field, enter **20**.
23. In the Hello Interval field, enter **10**.
24. In the Dead Interval field, enter **40**.
25. Click **OK**.
26. In the Edit OSPF Properties dialog box, enter **20** in the Priorities field, and click **OK**.
27. Click the **Authentication** tab.

    The Edit OSPF Authentication dialog box appears.
28. In the Authentication area, click the **MD5** radio button.
29. In the MD5 and Key ID area, enter **cisco** in the MD5 Key field, and **1** in the MD5 Key ID field.
30. Click **OK**.
31. Choose **Configuration** > **Device Setup** > **Routing** > **OSPF** > **Setup,** and click the **Area/Networks** tab.
32. Choose the **OSPF 2** process and click **Edit**.

    The Edit OSPF Area dialog box appears.
33. In the Area Type area, choose **Stub**.
34. In the Authentication area, choose **None**, and enter **20** in the Default Cost field.
35. Click **OK**.
36. In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPF** > **Setup**.
37. Click the **Process Instances** tab and check the **OSPF process 2** check box.

38. Click **Advanced**.

    The Edit OSPF Area dialog box appears.

39. In the Timers area, enter **10** in the SPF Delay Time field and **20** in the SPF Hold Time field.
40. In the Adjacency Changes area, check the **Log Adjacency Change Details** check box.
41. Click **OK**.
42. Click **Reset**.

Examples for OSPFv3

The following example shows how to configure OSPFv3 routing in ASDM:

1. In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPFv3** > **Setup**.
2. On the Process Instances tab, do the following:
   a. Check the **Enable OSPFv3 Process** check box.
   b. Enter **1** in the Process ID field.
3. Click the **Areas** tab. then click **Add** to display the Add OSPFv3 Area dialog box.
4. From the OSPFv3 Process ID drop-down list, choose **1**.
5. Enter **22** in the Area ID field.
6. Choose **Normal** from the Area Type drop-down list.
7. Enter **10** in the Default Cost field.
8. Check the **Redistribution imports routes to normal and NSSA areas** check box.
9. Enter **20** in the Metric field.
10. Choose **1** from the Metric Type drop-down list.
11. Check the **inside** check box as the specified interface being used.
12. Check the **Enable Authentication** check box.
13. Enter **300** in the Security Policy Index field.
14. Choose **SHA-1** from the Authentication Algorithm drop-down list.
15. Enter **12345ABCDE** in the Authentication Key field.
16. Choose **DES** from the Encryption Algorithm drop-down list.
17. Enter **1122334455aabbccddee** in the Encryption Key field.
18. Click **OK**.
19. Click the **Route Summarization** tab, then click **Add** to display the Add Route Summarization dialog box.

20. Choose **1** from the Process ID drop-down list.
21. Choose **22** from the Area ID drop-down list.
22. Enter **2000:122::/64** in the IPv6 Prefix/Prefix Length field.
23. (Optional) Enter **100** in the Cost field.
24. Check the **Advertised** check box.
25. Click **OK**.
26. In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **OSPFv3** > **Interface**.
27. Click the **Properties** tab.
28. Check the **inside** check box and click **Edit** to display the Edit OSPF Properties dialog box.
29. In the Cost field, enter **20**.
30. Enter **1** in the Priority field.
31. Check the **point-to-point** check box.
32. In the Dead Interval field, enter **40**.
33. In the Hello Interval field, enter **10**.
34. In the Retransmit Interval field, enter **15**.
35. In the Transmit Delay field, enter **20**.
36. Click **OK**.
37. In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **Redistribution**.
38. Choose **1** from the Process ID drop-down list.
39. Choose **OSPF** from the Source Protocol drop-down list.
40. Enter **50** in the Metric field.
41. Choose **1** from the Metric Type drop-down list.
42. Click **OK**.
43. Click **Apply** to save your changes.

Monitoring OSPF

You can display specific statistics such as the contents of IP routing tables, caches, and databases. You can also use the information provided to determine resource utilization and solve network problems. You can also display information about node reachability and discover the routing path that your device packets are taking through the network.

To monitor or display various OSPFv2 routing statistics in ASDM, perform the following steps:

1. In the main ASDM window, choose **Monitoring** > **Routing** > **OSPF LSAs**.

2. You can select and monitor OSPF LSAs, Types 1 through 5 and 7. Each pane shows one LSA type, as follows:
   - Type 1 LSAs represent the routes in an area under a process.
   - Type 2 LSAs show the IP address of the designated router that advertises the routers.
   - Type 3 LSAs show the IP address of the destination network.
   - Type 4 LSAs show the IP address of the AS boundary router.
   - Type 5 LSAs and Type 7 LSAs show the IP address of the AS external network.
3. Click **Refresh** to update each LSA type pane.
4. In the main ASDM window, choose **Monitoring** > **Routing** > **OSPF Neighbors**.

   In the OSPF Neighbors pane, each row represents one OSPF neighbor. In addition, the OSPF Neighbors pane shows the network on which the neighbor is running, the priority, the state, the amount of dead time in seconds, the IP address of the neighbor, and the interface on which it is running. For a list of possible states for an OSPF neighbor, see RFC 2328.
5. Click **Refresh** to update the OSPF Neighbors pane.

To monitor or display various OSPFv3 routing statistics in ASDM, perform the following steps:

1. In the main ASDM window, choose **Monitoring** > **Routing** > **OSPFv3 LSAs**.
2. You can select and monitor OSPFv3 LSAs. Choose a link-state type to display its status according to specified parameters from the Link State type drop-down list. The supported link-state types are router, network, inter-area prefix, inter-area router, AS external, NSSA, link, and intra-area prefix.
3. Click **Refresh** to update each link-state type.
4. In the main ASDM window, choose **Monitoring** > **Routing** > **OSPFv3 Neighbors**.

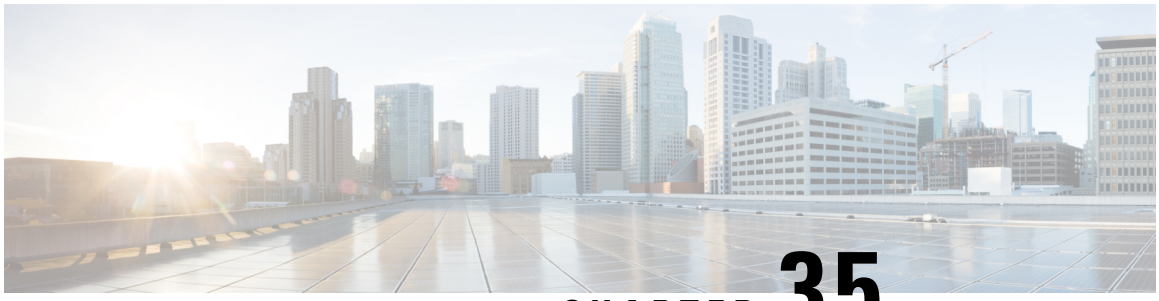
   In the OSPFv3 Neighbors pane, each row represents one OSPFv3 neighbor. In addition, the OSPFv3 Neighbors pane shows the IP address of the neighbor, the priority, the state, the amount of dead time in seconds, and the interface on which it is running. For a list of possible states for an OSPFv3 neighbor, see RFC 5340.
5. Click **Refresh** to update the OSPFv3 Neighbors pane.

History for OSPF

Table 41: Feature History for OSPF

Feature Name	Platform Releases	Feature Information
OSPF Support	7.0(1)	Support was added for route data, authentication, and redistribution and monitoring of routing information using the Open Shortest Path First (OSPF) routing protocol. We introduced the following screen: Configuration > Device Setup > Routing > OSPF.

Feature Name	Platform Releases	Feature Information
Dynamic Routing in Multiple Context Mode	9.0(1)	OSPFv2 routing is supported in multiple context mode. We modified the following screen: Configuration > Device Setup > Routing > OSPF > Setup
Clustering	9.0(1)	For OSPFv2 and OSPFv3, bulk synchronization, route synchronization, and Spanned EtherChannel load balancing are supported in the clustering environment.
OSPFv3 Support for IPv6	9.0(1)	OSPFv3 routing is supported for IPv6. We introduced the following screens: Configuration > Device Setup > Routing > OSPFv3 > Setup, Configuration > Device Setup > Routing > OSPFv3 > Interface, Configuration > Device Setup > Routing > OSPFv3 > Redistribution, Configuration > Device Setup > Routing > OSPFv3 > Summary Prefix, Configuration > Device Setup > Routing > OSPFv3 > Virtual Link, Monitoring > Routing > OSPFv3 LSAs, Monitoring > Routing > OSPFv3 Neighbors.
OSPF support for Fast Hellos	9.2(1)	OSPF supports the Fast Hello Packets feature, resulting in a configuration that results in faster convergence in an OSPF network. We modified the following screen: Configuration > Device Setup > Routing > OSPF > Interface > Edit OSPF Interface Advanced Properties
Timers	9.2(1)	New OSPF timers were added; old ones were deprecated. We modified the following screen: Configuration > Device Setup > Routing > OSPF > Setup > Edit OSPF Process Advanced Properties
Route filtering using access-list	9.2(1)	Route filtering using ACL is now supported. We introduced the following screen: Configuration > Device Setup > Routing > OSPF > Filtering Rules > Add Filter Rules
OSPF Monitoring enhancements	9.2(1)	Additional OSPF monitoring information was added.
OSPF redistribute BGP	9.2(1)	OSPF redistribution feature was added. We added the following screen: Configuration > Device Setup > Routing > OSPF > Redistribution
OSPF Support for Non-Stop Forwarding (NSF)	9.3(1)	OSPFv2 and OSPFv3 support for NSF was added. We added the following screens: Configuration > Device Setup > Routing > OSPF > Setup > NSF Properties, Configuration > Device Setup > Routing > OSPFv3 > Setup > NSF Properties
OSPF Support for Non-Stop Forwarding (NSF)	9.13(1)	NSF wait timer was added. We added a new command for setting the timer for the NSF restart interval. This command was introduced to ensure the wait interval is not longer that the router dead interval. We introduced the following command: **timers nsf wait <seconds>**



CHAPTER 35

IS-IS

This chapter describes the Intermediate System to Intermediate System (IS-IS) routing protocol.

- About IS-IS, on page 895
- Prerequisites for IS-IS, on page 901
- Guidelines for IS-IS, on page 901
- Configure IS-IS, on page 902
- Monitoring IS-IS, on page 917
- History for IS-IS, on page 917

About IS-IS

IS-IS routing protocol is a link state Interior Gateway Protocol (IGP). Link-state protocols are characterized by the propagation of the information required to build a complete network connectivity map on each participating device. That map is then used to calculate the shortest path to destinations. The IS-IS implementation supports IPv4 and IPv6.

You can divide a routing domain into one or more subdomains. Each subdomain is called an area and is assigned an area address. Routing within an area is known as Level-1 routing. Routing between Level-1 areas is known as Level-2 routing. A router is referred to as an Intermediate System (IS). An IS can operate at Level 1, Level 2, or both. ISes that operate at Level 1 exchange routing information with other Level-1 ISes in the same area. ISes that operate at Level 2 exchange routing information with other Level-2 routers regardless of whether they are in the same Level-1 area. The set of Level-2 routers and the links that interconnect them form the Level-2 subdomain, which must not be partitioned in order for routing to work properly.

About NET

An IS is identified by an address known as a Network Entity Title (NET). The NET is the address of a Network Service Access Point (NSAP), which identifies an instance of the IS-IS routing protocol running on an IS. The NET is 8 to 20 octets in length and has the following three parts:

- Area address—This field is 1 to 13 octets in length and is composed of high-order octets of the address.

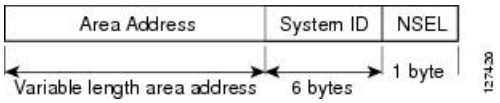
Note You can assign multiple area addresses to an IS-IS instance; in this case, all area addresses are considered synonymous. Multiple synonymous area addresses are useful when merging or splitting areas in the domain. Once the merge or split has been completed, you do not need to assign more than one area address to an IS-IS instance.

- System ID—This field is 6 octets long and immediately follows the area address. When the IS operates at Level 1, the system ID must be unique among all the Level-1 devices in the same area. When the IS operates at Level 2, the system ID must be unique among all devices in the domain.

Note You assign one system ID to an IS instance.

- NSEL—The N-selector field is 1 octet in length and immediately follows the system ID. It must be set to 00.

Figure 92: NET Format



IS-IS Dynamic Hostname

In the IS-IS routing domain, the system ID is used to represent each ASA. The system ID is part of the NET that is configured for each IS-IS ASA. For example, an ASA with a configured NET of 49.0001.0023.0003.000a.00 has a system ID of 0023.0003.000a. ASA-name-to-system-ID mapping is difficult for network administrators to remember during maintenance and troubleshooting on the ASAs.

The dynamic hostname mechanism uses link-state protocol (LSP) flooding to distribute the ASA-name-to-system-ID mapping information across the entire network. Every ASA on the network will try to install the system ID-to-ASA name mapping information in its routing table.

If an ASA that has been advertising the dynamic name type, length, value (TLV) on the network suddenly stops the advertisement, the mapping information last received will remain in the dynamic host mapping table for up to one hour, allowing the network administrator to display the entries in the mapping table during a time when the network experiences problems.

IS-IS PDU Types

ISes exchange routing information with their peers using protocol data units (PDUs). Intermediate System-to-Intermediate System Hello PDUs (IIHs), Link-State PDUs (LSPs), and Sequence Number PDUs (SNPs) types of PDUs are used.

IIHs

IIHs are exchanged between IS neighbors on circuits that have the IS-IS protocol enabled. IIHs include the system ID of the sender, the assigned area address(es), and the identity of neighbors on that circuit that are known to the sending IS. Additional optional information can also be included.

There are two types of IIHs:

- Level-1 LAN IIHs—These are sent on multiaccess circuits when the sending IS operates as a Level-1 device on that circuit.
- Level-2 LAN IIHs—These are sent on multiaccess circuits when the sending IS operates as a Level-2 device on that circuit.

LSPs

An IS generates LSPs to advertise its neighbors and the destinations that are directly connected to the IS. An LSP is uniquely identified by the following:

- System ID of the IS that generated the LSP
- Pseudonode ID—This value is always 0 except when the LSP is a pseudonode LSP
- LSP number (0 to 255)
- 32-bit sequence number

Whenever a new version of an LSP is generated, the sequence number is incremented.

Level-1 LSPs are generated by ISs that support Level 1. The Level-1 LSPs are flooded throughout the Level-1 area. The set of Level-1 LSPs generated by all Level-1 ISs in an area is the Level-1 LSP Database (LSPDB). All Level-1 ISs in an area have an identical Level-1 LSPDB and therefore have an identical network connectivity map for the area.

Level-2 LSPs are generated by ISs that support Level 2. Level-2 LSPs are flooded throughout the Level-2 subdomain. The set of Level-2 LSPs generated by all Level-2 ISs in the domain is the Level-2 LSP Database (LSPDB). All Level-2 ISs have an identical Level-2 LSPDB and therefore have an identical connectivity map for the Level-2 subdomain.

SNPs

SNPs contain a summary description of one or more LSPs. There are two types of SNPs for both Level 1 and Level 2:

- Complete Sequence Number PDUs (CSNPs) are used to send a summary of the LSPDB that an IS has for a given level.
- Partial Sequence Number PDUs (PSNPs) are used to send a summary of a subset of the LSPs for a given level that an IS either has in its database or needs to obtain.

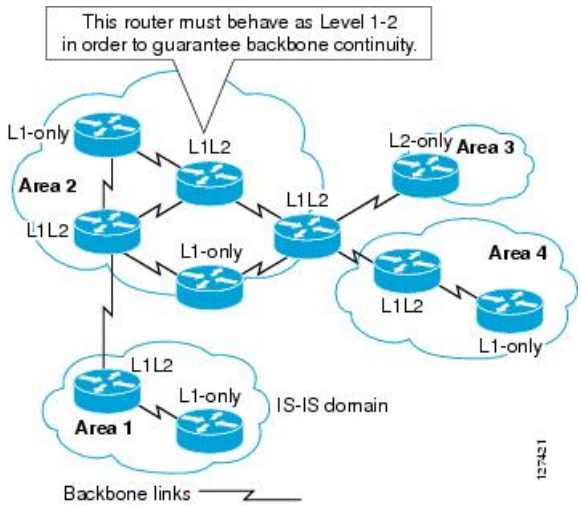
Operation of IS-IS on Multiaccess Circuits

Multiaccess circuits support multiple ISes, that is, two or more operating on the circuit. For multiaccess circuits a necessary prerequisite is the ability to address multiple systems using a multicast or broadcast address. An IS that supports Level 1 on a multiaccess circuit sends Level-1 LAN IIHs on the circuit. An IS that supports Level 2 on a multiaccess circuit sends Level-2 LAN IIHs on the circuit. ISes form separate adjacencies for each level with neighbor ISes on the circuit.

An IS forms a Level-1 adjacency with other ISes that support Level 1 on the circuit and has a matching area address. Two ISes with disjointed sets of area addresses supporting Level 1 on the same multiaccess circuit is NOT supported. An IS forms a Level-2 adjacency with other ISes that support Level 2 on the circuit.

The devices in the IS-IS network topology in the following figure perform Level 1, Level 2, or Level 1 and 2 routing along the backbone of the network.

Figure 93: Level-1, Level-2, Level 1-2 Devices in an IS-IS Network Topology



IS-IS Election of the Designated IS

If each IS advertised all of its adjacencies on a multiaccess circuit in its LSPs, the total number of advertisements required would be N 2 (where N is the number of ISes that operate at a given level on the circuit). To address this scalability issue, IS-IS defines a pseudonode to represent the multiaccess circuit. All ISes that operate on the circuit at a given level elect one of the ISes to act as the Designated Intermediate System (DIS) on that circuit. A DIS is elected for each level that is active on the circuit.

The DIS is responsible for issuing pseudonode LSPs. The pseudonode LSPs include neighbor advertisements for all of the ISes that operate on that circuit. All ISes that operate on the circuit (including the DIS) provide a neighbor advertisement to the pseudonode in their non-pseudonode LSPs and do not advertise any of their neighbors on the multiaccess circuit. In this way the total number of advertisements required varies as a function of N-the number of ISes that operate on the circuit.

A pseudonode LSP is uniquely classified by the following identifiers:

- System ID of the DIS that generated the LSP
- Pseudonode ID (ALWAYS NON-ZERO)
- LSP number (0 to 255)
- 32-bit sequence number

The nonzero pseudonode ID is what differentiates a pseudonode LSP from a non-pseudonode LSP and is chosen by the DIS to be unique among any other LAN circuits for which it is also the DIS at this level.

The DIS is also responsible for sending periodic CSNPs on the circuit. This provides a complete summary description of the current contents of the LSPDB on the DIS. Other ISes on the circuit can then perform the following activities, which efficiently and reliably synchronizes the LSPDBs of all ISes on a multiaccess circuit:

- Flood LSPs that are absent from or are newer than those that are described in the CSNPs sent by the DIS.

- Request an LSP by sending a PSNP for LSPs that are described in the CSNPs sent by the DIS that are absent from the local database or older than what is described in the CSNP set.

IS-IS LSPDB Synchronization

Proper operation of IS-IS requires a reliable and efficient process to synchronize the LSPDBs on each IS. In IS-IS this process is called the update process. The update process operates independently at each supported level. Locally generated LSPs are always new LSPs. LSPs received from a neighbor on a circuit may be generated by some other IS or may be a copy of an LSP generated by the local IS. Received LSPs can be older, the same age, or newer than the current contents of the local LSPDB.

Handling Newer LSPs

When a newer LSP is added to the local LSPDB, it replaces an older copy of the same LSP in the LSPDB. The newer LSP is marked to be sent on all circuits on which the IS currently has an adjacency in the UP state at the level associated with the newer LSP—excluding the circuit on which the newer LSP was received.

For multiaccess circuits, the IS floods the newer LSP once. The IS examines the set of CNSPs that are sent periodically by the DIS for the multiaccess circuit. If the local LSPDB contains one or more LSPs that are newer than what is described in the CSNP set (this includes LSPs that are absent from the CSNP set), those LSPs are reflooded over the multiaccess circuit. If the local LSPDB contains one or more LSPs that are older than what is described in the CSNP set (this includes LSPs described in the CSNP set that are absent from the local LSPDB), a PSNP is sent on the multiaccess circuit with descriptions of the LSPs that require updating. The DIS for the multiaccess circuit responds by sending the requested LSPs.

Handling Older LSPs

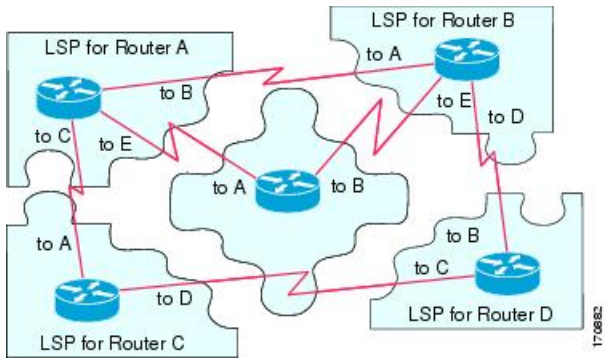
An IS may receive an LSP that is older than the copy in the local LSPDB. An IS may receive an SNP (complete or partial) that describes an LSP that is older than the copy in the local LSPDB. In both cases the IS marks the LSP in the local database to be flooded on the circuit on which the older LSP or SNP that contained the older LSP was received. Actions taken are the same as described above after a new LSP is added to the local database.

Handling Same-Age LSPs

Because of the distributed nature of the update process, it is possible than an IS may receive copies of an LSP that is the same as the current contents of the local LSPDB. In multiaccess circuits receipt of a same-age LSP is ignored. Periodic transmission of a CSNP set by the DIS for that circuit serves as an implicit acknowledgment to the sender that the LSP has been received.

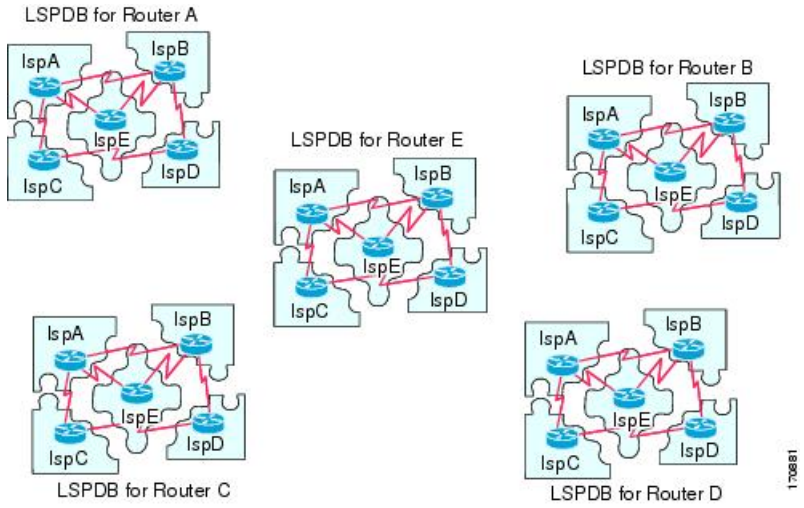
The following figure shows how LSPs are used to create a network map. Think of the network topology as a jigsaw puzzle. Each LSP (representing an IS) is one of the pieces. It is applicable to all Level-1 devices in an area or to all Level-2 devices in a Level-2 subdomain.

Figure 94: IS-IS Network Map



The following figure shows each device in the IS-IS network with its fully updated link-state database after the adjacencies have been formed among the neighbor devices. It is applicable to all Level-1 devices in an area or to all Level-2 devices in a Level-2 subdomain.

Figure 95: IS-IS Devices with Synchronized LSPDBs



IS-IS Shortest Path Calculation

When the contents of the LSPDB change, each IS independently reruns a shortest path calculation. The algorithm is based on the well-known Dijkstra algorithm for finding the shortest paths along a directed graph where the ISes are the vertices of the graph and the links between the ISes are edges with a nonnegative weight. A two-way connectivity check is performed before considering a link between two ISes as part of the graph. This prevents the use of stale information in the LSPDB, for example, when one IS is no longer operating in the network but did not purge the set of LSPs that it generated before stopping operation.

The output of the SPF is a set of tuples (destination, next hop). The destinations are protocol-specific. Multiple equal-cost paths are supported, in which case multiple next hops would be associated with the same destination.

Independent SPFs are performed for each level supported by the IS. When the same destination is reachable by both Level-1 and Level-2 paths, the Level-1 path is preferred.

A Level-2 IS that indicates that it has one or more Level-2 neighbors in other areas may be used by Level-1 devices in the same area as the path of last resort, also called the default route. The Level-2 IS indicates its attachment to other areas by setting an attached bit (ATT) in its Level-1 LSP 0.



Note An IS can generate up to 256 LSPs at each level. The LSPs are identified by the numbers 0 through 255. LSP 0 has special properties, including the significance of the setting of the ATT bit to indicate attachment to other areas. When LSPs that are numbered 1 through 255 have the ATT bit set, it is not significant.

IS-IS Shutdown Protocol

You can shut down IS-IS (placing it in an administrative down state) to make changes to the IS-IS protocol configuration without losing your configuration parameters. You can shut down IS-IS at the global IS-IS process level or at the interface level. If the device was rebooted when the protocol was turned off, the protocol would be expected to come back up in the disabled state. When the protocol is set to the administrative down state, network administrators are allowed to administratively turn off the operation of the IS-IS protocol without losing the protocol configuration, to make a series of changes to the protocol configuration without having the operation of the protocol transition through intermediate-and perhaps undesirable-states, and to then reenable the protocol at a suitable time.

Prerequisites for IS-IS

The following prerequisites are necessary before configuring IS-IS:

- Knowledge of IPv4 and IPv6.
- Knowledge of your network design and how you want traffic to flow through it before configuring IS-IS.
- Define areas, prepare an addressing plan for the devices (including defining the NETs), and determine the interfaces that will run IS-IS.
- Before you configure your devices, prepare a matrix of adjacencies that shows what neighbors should be expected in the adjacencies table. This will facilitate verification.

Guidelines for IS-IS

Firewall Mode Guidelines

Supported only in routed firewall mode. Transparent firewall mode is not supported.

Cluster Guidelines

Supported only in Individual Interface mode; Spanned EtherChannel mode is not supported.

Additional Guidelines

IS-IS is not supported with bidirectional forwarding.

Configure IS-IS

This section describes how to enable and configure the IS-IS process on your system.

Procedure

Step 1 Enable IS-IS Routing Globally, on page 902.
Step 2 Enable IS-IS Authentication, on page 903.
Step 3 Configure IS-IS LSP, on page 904
Step 4 Configure IS-IS Summary Addresses, on page 906.
Step 5 Configure IS-IS NET, on page 907.
Step 6 Configure IS-IS Passive Interfaces, on page 907.
Step 7 Configure IS-IS Interfaces, on page 908.
Step 8 Configure IS-IS IPv4 Address Family, on page 912.
Step 9 Configure IS-IS IPv6 Address Family, on page 915.

Enable IS-IS Routing Globally

Before you begin

In multiple context mode, complete this procedure in the context execution space. To change from the system to a context configuration, in the Configuration > Device List pane, double-click the context name under the active device IP address.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Routing** > **ISIS** > **General**.
Step 2 Check the **Configure ISIS** check box to enable IS-IS.
Step 3 Check the **Shutdown protocol** check box to enable shutdown protocol,

See IS-IS Shutdown Protocol, on page 901for more information on Shutdown protocol.

Step 4 To have IS-IS use a dynamic hostname, check the **Use dynamic hostname** check box.

Dynamic hostname is enabled by default. See IS-IS Dynamic Hostname, on page 896 for detailed information on the dynamic hostname in IS-IS.

Step 5 To prevent IS-IS from padding LAN hello PDUs, check the **Do not pad LAN hello PDUs** check box.

IS-IS hellos are padded to the full maximum transmission unit (MTU) size. This allows for early detection of errors that result from transmission problems with large frames or errors that result from mismatched MTUs on adjacent interfaces. You can disable hello padding to avoid wasting network bandwidth in case the MTU of both interfaces is the same or in the case of translational bridging.

Step 6 To advertise passive interfaces only, check the **Advertise passive only** check box.

It excludes IP prefixes of connected networks from LSP advertisements, which reduces IS-IS convergence time.

Step 7 Choose whether to have your ASA act as station router (Level 1), an area router (Level 2), or both (Level 1-2) by clicking the appropriate radio button.

See About IS-IS, on page 895for more information on IS-IS levels.

Step 8 In the **Topology priority** field, enter a number that indicates where the ASA's priority is in the topology. The range is from 0 to 127.

Step 9 In the **Route priority tag** field, enter a tag that indicates the ASA's route priority. The range is from 1 to 4294967295. The default value is 100. Higher values indicate higher preference. This preference is sent to all routers in the IS-IS system.

Step 10 To have the IS conditionally advertise as L2, choose a device from the drop-down menu, and click **Manage**.

See Define a Route Map for the procedure for adding a route map.

Step 11 Check the **Log changes in adjacency** check box to have the ASA send a log message whenever an IS-IS neighbor goes up or down.

This command is disabled by default. Logging adjacency changes is useful when monitoring large networks.

Step 12 To have changes included from non-IIH events, check the **Include changes generated by non-IIH events** check box.

Step 13 To set up the skeptical time interval, enter the amount of minutes in the **Skeptical interval** field. The range is 0 to 1440 minutes. The default is five minutes.

Step 14 Click **Apply**.

Enable IS-IS Authentication

IS-IS route authentication prevents the introduction of unauthorized or false routing messages from unapproved sources. You can set a password for each IS-IS area or domain to prevent unauthorized routers from injecting false routing information into the link-state database, or you can configure a type of IS-IS authentication, either IS-IS MD5 or enhanced clear text authentication. You can also set authentication per interface. All IS-IS neighbors on interfaces configured for IS-IS message authentication must be configured with the same authentication mode and key for adjacencies to be established.

See About IS-IS, on page 895 for more information on areas and domains.

Before you begin

Before you can enable IS-IS route authentication, you must enable IS-IS and set up an area. See Enable IS-IS Routing Globally, on page 902 for the procedure.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Routing** > **ISIS** > **Authentication**.

Step 2 Configure authentication parameters for Level 1 and Level 2:

- In the **Key** field, enter the key to authenticate IS-IS updates. The key can include up to 16 characters.
- Click the **Enable** or **Disable** radio button depending or whether you want to have Send Only enabled.

  Note
  ASAs will have more time for the keys to be configured on each ASA if authentication is inserted only on the packets being sent, not checked on packets being received.

- Choose the authentication mode by clicking either the **Disabled**, **MD5**, or **Plaintext** radio button.

Step 3 If you choose **Disabled**, enter an area password for the Level 1 area (subdomain) and/or a domain password for the Level 2 domain.

Step 4 Click **Apply**.

Configure IS-IS LSP

An IS generates LSPs to advertise its neighbors and the destinations that are directly connected to IS-IS. See IS-IS PDU Types, on page 896 for more detailed information on LSPs.

Use the following commands to configure LSPs so that you have a faster convergence configuration.

Before you begin

In multiple context mode, complete this procedure in the context execution space. To change from the system to a context configuration, in the Configuration > Device List pane, double-click the context name under the active device IP address.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Routing** > **ISIS** > **Link State Packet**.

Note
IS-IS must be enabled before you can configure LSP parameters. See Enable IS-IS Routing Globally, on page 902 for the procedure.

Step 2 To allow the ASA to ignore LSP packets that are received with internal checksum errors rather than purging the LSPs, check the **Ignore LSP errors** check box.

Step 3 To fast-flood and fill LSPs before running SPF, check the **Flood LSPs before running SPF**, and then in the **Number of LSPs to be flooded** field, enter a number. The range is 1 to 15. The default is 5.

This parameter sends a specified number of LSPs from the ASA. If no LSP number is specified, the default of 5 is used. The LSPs invoke SPF before running SPF. We recommend that you enable fast flooding, because then you speed up the LSP flooding process, which improves overall network convergence time.

Step 4 To suppress IP prefixes, check the **Suppress IP prefixes** check box, and then check one of the following:

- **Don't advertise IP prefixes learned form another ISIS level when ran out of LSP fragments**—Suppresses any routes coming from another level. For example, if the Level-2 LSP becomes full, routes from Level 1 are suppressed.

- **Don't advertise IP prefixes learned form other protocols when ran out of LSP fragments**—Suppresses any redistributed routes on the ASA.

In networks where there is no limit placed on the number of redistributed routes into IS-IS, it is possible that the LSP can become full and routes will be dropped. Use these options to control which routes are suppressed when the PDU becomes full.

Step 5 Configure the LSP generation intervals for Level 1 and Level 2:

- **LSP calculation interval**—Enter the interval of time in seconds between transmission of each LSP. The range is 1 to 120 seconds. The default is 5.

  The number should be greater than the expected round-trip delay between any two ASAs on the attached network. The number should be conservative or needless transmission results. Retransmissions occur only when LSPs are dropped. So setting the number to a higher value has little effect on reconvergence. The more neighbors the ASAs have, and the more paths over which LSPs can be flooded, the higher you can make this value.

- **Initial wait for LSP calculation**—Enter the time in milliseconds specifying the initial wait time before the first LSP is generated. The range is 1 to 120,000. The default is 50.

- **Minimum wait between first and second LSP calculation**—Enter the time in milliseconds between the first and second LSP generation. The range is 1 to 120,000. The default is 5000.

Step 6 If you want the values you configured for Level 1 to also apply to Level 2, check the **Use level 1 parameters also for level 2** check box.

Step 7 In the **Maximum LSP size** field, enter the maximum number of seconds between two consecutive occurrences of an LSP being generated. The range is 128 to 4352. The default is 1492.

Step 8 In the **LSP refresh interval** field, enter the number of seconds at which LSPs are refreshed. The range is 1 to 65,5535. The default is 900.

The refresh interval determines the rate at which the software periodically transmits in LSPs the route topology information that it originates. This is done to keep the database information from becoming too old.

Reducing the refresh interval reduces the amount of time that undetected link state database corruption can persist at the cost of increased link utilization. (This is an extremely unlikely event, however, because there are other safeguards against corruption.) Increasing the interval reduces the link utilization caused by the flooding of refreshed packets (although this utilization is very small).

Step 9 In the **Maximum LSP lifetime** field, enter the maximum number of seconds that LSPs can remain in a router's database without being refreshed. The range is 1 to 65,535. The default is 1200 (20 minutes).

You might need to adjust this parameter if you change the LSP refresh interval. LSPs must be periodically refreshed before their lifetimes expire. The value set for LSP refresh interval should be less that the value set for the maximum LSP lifetime; otherwise LSPs will time out before they are refreshed. If you make the LSP lifetime too low compared to the LSP refresh interval, the LSP refresh interval is automatically reduced to prevent the LSPs from timing out.

Step 10 Click **Apply**.

Configure IS-IS Summary Addresses

Multiple groups of addresses can be summarized for a given level. Routes learned from other routing protocols can also be summarized. The metric used to advertise the summary is the smallest metric of all the more specific routes. This helps to reduce the size of the routing table.

You need to manually define summary addresses if you want to create summary addresses that do not occur at a network number boundary or if you want to use summary addresses on an ASA with automatic route summarization disabled.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Routing** > **ISIS** > **Summary Address**.

The **Configure ISIS Summary Address** pane displays a table of the statically-defined IS-IS summary addresses. By default, IS-IS summarizes subnet routes to the network level. You can create statically defined IS-IS summary addresses to the subnet level from the **Configure ISIS Summary Address** pane.

Step 2 Click **Add** to add a new IS-IS summary address, or to click **Edit** to edit an existing IS-IS summary address in the table.

The **Add Summary Address** or **Edit Summary Address** dialog box is displayed. You can also double-click an entry in the table to edit that entry.

Step 3 In the **IP Address** field, enter the IP address of the summary route.

Step 4 In the **Netmask** field, choose or enter the network mask to apply to the IP address.

Step 5 Select the **Level 1**, **Level 2**, or **Level 1 and 2** radio button depending on which levels you want to receive summary addresses.

- **Level 1**—Summary routes are applied when redistributing routes into Level 1 and Level 2 and when Level 2 IS-IS advertises Level 1 routes as reachable in its area.
- **Level 2**—Routes learned by Level 1 routing are summarized into the Level 2 backbone with the configured address and mask value. Redistributed routes into Level 2 IS-IS are summarized also.
- **Level 1 and 2**—Summary routes are applied when redistributing routes into Level 1 and Level 2 and when Level 2 IS-IS advertises Level 1 routes as reachable in its area.

Step 6 In the **Tag** field, enter a number for the tag. The range is 1 to 4294967295.

The Tag field lets you specify a number to tag routes that are being summarized. If the routes have already been tagged on the **Configuration** > **Device Setup** > **Routing** > **ISIS** > **General** pane in the **Route priority tag** field, those routes are summarized, otherwise the tag is lost.

Step 7 In the **Metric** field, enter the metric that will be applied to the summary route. The range is 1 to 4294967295. The default value is 10.

The Metric value is assigned to the link and used to calculate the path cost via the links to destinations. You can configure this metric for Level 1 or Level 2 routing only.

Step 8 Click **OK**.

Step 9 Click **Apply**.

Configure IS-IS NET

Before you begin

In multiple context mode, complete this procedure in the context execution space. To change from the system to a context configuration, in the Configuration > Device List pane, double-click the context name under the active device IP address.

IS-IS uses addresses called Network Entity Titles (NET). They can be 8 to 20 bytes long, but are usually 10 bites long. You can add a NET entry on the NET page when clustering is not configured on the ASA. If your ASA has clustering configured, you must create a net pool entry on the **Configuration > Device Management > Advanced > Address Pools > NET Address Pools** pane. You can then reference the NET address pool on the NET pane.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Routing** > **ISIS** > **Network Entity Title (NET)**.

The **Configure Network Entity (NET)** pane displays a table of the NET addresses. You can add a NET entry here when clustering is NOT configured on the ASA. For an ASA with clustering configured, you must create a net pool entry at **Configuration** > **Device Management** > **Advanced** > **Address Pools** > **Net Address Pools**.

YOu can then reference the NET address pool on the Network Entity Title (NET) pane.

Step 2 Click **Add** to add a new IS-IS NET address, or to click **Edit** to edit an existing IS-IS NET address in the table.

The **Add Network Entity Title (NET)** or **Edit Network Entity Title (NET)** dialog box appears. You can also double-click an entry in the table to edit that entry.

Step 3 From the Network Entity Title (NET) drop-down list, choose a NET.

Step 4 In the **Maximum allowed Net** field, enter the maximum allowed NETs you want. The range is from 3 to 254. The default is 3.

In most cases only one NET is necessary, but in the case of merging multiple areas or splitting one area into multiple areas, you may need to use multiple area-addresses.

Step 5 Click **Apply**.

Configure IS-IS Passive Interfaces

You can disable IS-IS hello packets and routing updates on interfaces while still including the interface addresses in the topology database. These interfaces will not form IS-IS neighbor adjacencies

If you have an interface that you do not want to participate in IS-IS routing, but that is attached to a network that you want advertised, configure the passive interfaces to prevent that interface from using IS-IS. Additionally, you can specify the version of IS-IS that is used by the ASA for updates. Passive routing assists

in controlling the advertisement of IS-IS routing information and disables the sending and receiving of IS-IS routing updates on an interface.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Routing** > **IS-IS** > **Passive Interfaces**.

Step 2 To suppress routing updates on all interfaces, check the **Suppress routing updates on all Interfaces** check box.

This causes all interfaces to operate in passive mode.

Step 3 To configure individual interfaces to suppress routing updates, select the named routing interface in the left column and click **Add** to add it to the Suppress routing updates column.

Specifying an interface name sets only that interface to passive mode. In passive mode, IS-IS routing updates are accepted by, but not sent out of, the specified interface.

Note
Only interfaces that you have given a dynamic hostname can be suppressed from sending routing updates. See IS-IS Dynamic Hostname, on page 896 for more information.

Step 4 Click **Apply**.

Configure IS-IS Interfaces

This procedure describes how to modify individual ASA interfaces for IS-IS routing.

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Routing** > **ISIS** > **Interface**.

The **ISIS Interface Configuration** pane appears and displays the IS-IS interface configurations. You can configure hello padding per interface by checking/unchecking the **Hello Padding** check box.

IS-IS hellos are padded to the full maximum transmission unit (MTU) size. Padding IS-IS hellos to the full MTU allows for early detection of errors that result from transmission problems with large frames or errors that result from mismatched MTUs on adjacent interfaces.

Step 2 Choose an interface entry by double-clicking an interface entry, or choose the entry and click **Edit**.

The **Edit ISIS Interface** dialog box appears.

Step 3 On the **General** tab, configure the following:

- **Shutdown ISIS on this interface**—Lets you disable the IS-IS protocol for this interface without removing the configuration parameters. The IS-IS protocol does not form any adjacencies on this interface, and the IP address of this interface is put into the LSP that is generated by the ASA.
- **Enable ISIS on this interface**—Enables IS-IS protocol on this interface.

- **Enable IPv6 ISIS routing on this interface**—Enables IPv6 IS-IS routing on this interface.
- **Priority for level-1—**Lets you set a priority for Level 1. The priority is used to determine which router on a LAN will be the designated router or Designated Intermediate System (DIS). The priorities are advertised in the hello packets. The router with the highest priority becomes the DIS. The range is 0 to 127. The default is 64.

  Note
  In IS-IS, there is no backup designated router. Setting the priority to 0 lowers the chance of this system becoming the DIS, but does not prevent it. If a router with a higher priority comes on line, it takes over the role from the current DIS. In the case of equal priorities, the highest MAC address breaks the tie.

- **Priority for level-2—**Lets you set a priority for Level 2. The priority is used to determine which router on a LAN will be the designated router or Designated Intermediate System (DIS). The priorities are advertised in the hello packets. The router with the highest priority becomes the DIS. The range is 0 to 127. The default is 64.

  Note
  In IS-IS, there is no backup designated router. Setting the priority to 0 lowers the chance of this system becoming the DIS, but does not prevent it. If a router with a higher priority comes on line, it will take over the role from the current DIS. In the case of equal priorities, the highest MAC address breaks the tie.

- **Tag**—Sets a tag on the IP address configured for an interface when this IP prefix is put into an IS-IS LSP.
- **CSNP Interval for level-1—**Sets the Complete Sequence Number PDUs (CSNPs) interval in seconds between transmission of CSNPs on multiaccess networks for Level 1. This interval only applies for the designated ASA. The range is from 0 to 65535. The default is 10 seconds. It is unlikely that you will have to change the default.

  This option applies only for the designated router (DR) for a specified interface. Only DRs send CSNP packets to maintain database synchronization.

- **CSNP Interval for level-2—**Sets the Complete Sequence Number PDUs (CSNPs) interval in seconds between transmission of CSNPs on multiaccess networks for Level 2. This interval only applies for the designated ASA. The range is from 0 to 65535. The default is 10 seconds. It is unlikely that you will have to change the default.

  This option applies only for the designated router (DR) for a specified interface. Only DRs send CSNP packets to maintain database synchronization.

- **Adjacency filter—Filters the establishment of IS-IS adjacencies.**

  Filtering is performed by building NSAP addresses out of incoming IS-IS hello packets by combining each area address in the hello with the system ID. Each of these NSAP addresses is then passed through the filter. If any one NSAP matches, the filter is considered passed unless you specified**Match all area addresses**, in which case all addresses must pass. The functionality of **Match all area addresses** is useful in performing negative tests, such as accepting an adjacency only if a particular address is NOT present.

- **Match all area addresses—**(Optional) All NSAP addresses must match the filter to accept the adjacency. If not specified (the default), only one address must match the filter for the adjacency to be accepted.

Step 4 Click **OK**.

Step 5 On the **Authentication** tab, configure the following for Level 1 and/or Level 2:

- In the **Key** field, enter the key to authenticate IS-IS updates. The range is 0 to 8 characters.

  If no password is configured with the **Key** option, no key authentication is performed.

- For **Send only** click the **Enable** or **Disable** radio button.

  Choosing **Send only** causes the system only to insert the password into the SNPs, but not check the password in SNPs that it receives. Use this keyword during a software upgrade to ease the transition. The default is disabled.

- Choose the authentication mode by checking the **Mode** check box and then choosing **MD5** or **Text** from the drop-down list, and in the **Password** field, enter a password.

Step 6 Click **OK**.

Step 7 On the **Hello Padding** tab, configure the following:

- **Hello Padding**—Enables Hello Padding.

  IS-IS hellos are padded to the full maximum transmission unit (MTU) size. Padding IS-IS hellos to the full MTU allows for early detection of errors that result from transmission problems with large frames or errors that result from mismatched MTUs on adjacent interfaces.

- **Minimal holdtime 1 second for Level-1**—Enables the holdtime (in seconds) that the LSP remains valid for Level 1.

- **Hello Interval for level-1**—Specifies the length of time in seconds between hello packets for Level 1.

  By default, a value three times the hello interval seconds is advertised as the hold time in the hello packets sent. (Change the multiplier of 3 by checking the **Hello Multiplier** check box.) With smaller hello intervals, topological changes are detected faster, but there is more routing traffic. The range is 1 to 65535. The default is 10.

- **Minimal holdtime 1 second for Level-2**—Enables the holdtime (in seconds) that the LSP remains valid for Level 2.

- **Hello Interval for level-2**—Specifies the length of time in seconds between hello packets for Level 2.

  By default, a value three times the hello interval seconds is advertised as the hold time in the hello packets sent. (Change the multiplier of 3 by checking the **Hello Multiplier** check box.) With smaller hello intervals, topological changes are detected faster, but there is more routing traffic. The range is 1 to 65535. The default is 10.

- **Hello Multiplier for level-1**—Specifies the number of IS-IS hello packets a neighbor must miss before the ASA declares the adjacency is down for Level 1.

  The advertised hold time in IS-IS hello packets is set to the hello multiplier times the hello interval. Neighbors declare an adjacency to this ASA down after not having received any IS-IS hello packets during the advertised hold time. The hold time (and thus the hello multiplier and the hello interval) can be set on a per-interface basis, and can be different between different ASAs in one area. The range is 3 to 1000. The default is 3.

- **Hello Multiplier for level-2**—Specifies the number of IS-IS hello packets a neighbor must miss before the ASA declares the adjacency is down for Level 2.

  The advertised hold time in IS-IS hello packets is set to the hello multiplier times the hello interval. Neighbors declare an adjacency to this ASA down after not having received any IS-IS hello packets during the advertised hold time. The hold time (and thus the hello multiplier and the hello interval) can

be set on a per-interface basis, and can be different between different ASAs in one area. The range is 3 to 1000. The default is 3.

- **Configure Circuit Type**—Specifies whether the interface is configured for local routing (level 1), area routing (Level 2), or both local and area routing (Level 1-2).

Step 8 Click **OK**.

Step 9 On the **LSP Settings** tab, configure the following:

- **Advertise ISIS Prefix**—Allows the advertising of IP prefixes of connected networks in the LSP advertisements per IS-IS interface.

  Disabling this option is an IS-IS mechanism to exclude IP prefixed of connected network from LSP advertisements thereby reducing IS-IS convergence time.

- **Retransmit Interval**—Specifies the amount of time in seconds between retransmission of each IS-IS LSP.

  The number should be greater than the expected round-trip delay between any two ASAs on the attached network. The range is 0 to 65535. The default is 5.

- **Retransmit Throttle Interval**—Specifies the amount of time in milliseconds between retransmissions on each IS-IS LSP.

  This option may be useful in very large networks with many LSPs and many interfaces as a way of controlling LSP retransmission traffic. This option controls the rate at which LSPs can be resent on the interface. The range is 0 to 65535. The default is 33.

- **LSP Interval**—Specifies the time delay in millisecond between successive IS-IS LSP transmissions.

  In topologies with a large number of IS-IS neighbors and interfaces, an ASA may have difficulty with the CPU load imposed by LSP transmission and reception. This option allows the LSP transmission rate (and by implication the reception rate of other systems) to be reduced. The range is 1 to 4294967295. The default is 33.

Step 10 Click **OK**.

Step 11 On the **Metrics** tab, configure the following options for Level 1 and Level 2:

You can check the **Use the level 1 values also for level 2** check box if you want the same metrics for both Levels.

- **Use maximum metric value**— Specifies the metric assigned to the link and used to calculate the cost from each other router via the links in the network to other destinations.

- **Default metric**—Enter the number for the metric.

  The range is 1 to 16777214. The default value is 10.

Step 12 Click **OK**.

Step 13 Click **Apply**.

Configure IS-IS IPv4 Address Family

Routers are allowed to redistribute external prefixes or routes that are learned form any other routing protocol, static configuration, or connected interface. The redistributed routes are allowed in either a Level 1 router or a Level 2 router.

You can set up adjacency, Shortest Path First (SPF), and you can define conditions for redistributing routes from another routing domain into ISIS (redistribution) for IPv4 addresses.

Before you begin

Before you can enable IS-IS route authentication, you must enable IS-IS and set up an area. See Enable IS-IS Routing Globally, on page 902 for the procedure.

Make sure that IPv4 is enabled on at least one interface before trying to add a neighbor, or ASDM returns an error message indicating that the configuration failed.

Procedure

Step 1 Choose **Configuration > Device Setup > Routing > ISIS > IPv4 Address Family > General**.

a) Check the **Perform adjacency check** check box for the router to check on nearby IS routers.

b) In the **Administrative Distance** field, enter a distance assigned to routes discovered by IS-IS protocol.

Administrative distance is a parameter used to compare routes among different routing protocols. In general, the higher the value, the lower the trust rating. And administrative distance of 255 means that the routing information source cannot be trusted at all and should be ignored. The range is 1 to 255. The default is 1.

You can use the distance option to configure the administrative distances applied to IS-IS routes when they are inserted into the Routing Information Base (RIB), and influence the likelihood of these routes being preferred over routes to the same destination addresses discovered by other protocols.

c) In the **Maximum number of forward paths** field, enter the maximum number of IS routes that can be installed in a routing table. The range is 1 to 8.

d) Check the **Distribute default route** check box to configure an IS routing process to distribute a default route, and then choose the default route from the drop-down list or click **Manage** to create a new route. See Define a Route Map, on page 805 for the procedure for creating a new route.

Step 2 Configure IS-IS metrics:

a) In the **Global ISIS metric for level 1**, enter a number specifying the metric.

The range is 1 to 63. The default is 10.

When you need to change the default metric value for all IS-IS interfaces, we recommend that you use the **Global ISIS metric for level 1** option to configure all interfaces globally. Globally configuring the metric values prevents user errors, such as unintentionally removing a set metric from an interface without configuring a new value and unintentionally allowing the interface to revert to the default metric of 10, thereby becoming a highly preferred interface in the network.

b) In the **Global ISIS metric for level 2**, enter a number specifying the metric.

The range is 1 to 63. The default is 10.

When you need to change the default metric value for all IS-IS interfaces, we recommend that you use the **Global ISIS metric for level 1** option to configure all interfaces globally. Globally configuring the metric values prevents user errors, such as unintentionally removing a set metric from an interface without configuring a new value and unintentionally allowing the interface to revert to the default metric of 10, thereby becoming a highly preferred interface in the network.

c) Choose one of the following to configure Type, Length, and Values (TLVs):

- Check the **Send and accept both styles of TLVs during transition** check box.
- Click the **Use old style of TLVs with narrow metric** radio button.
- Click the **Use new style TLVs to carry wider metric** radio button.

  If you choose one of the radio buttons, you can also check the **Accept both styles of TLVs during transition** check box.

  We strongly recommend that you use the new-style TLV because TLVs that are used to advertise IPv4 information in LSPs are defined to use only extended metrics. The software provides support of a 24-bit metric field, the wide metric. Using the new metric style, link metrics now have a maximum value of 16777214 with a total path metric of 4261412864.

d) Check the **Apply metric style to** check box, and then check the **Level-1** and/or **Level-2** check box.

Step 3 Click **Apply**.

Step 4 Choose **Configuration > Device Setup > Routing > ISIS > IPv4 Address Family > SPF**.

a) Check the **Honour external metrics during SPF calculations** check box, to have the SPF calculations include external metrics.

b) Check the **Signal other routers not to use this router as an intermediate hop in their SPF calculations** check box if you want to exclude this device, and configure the following:

- Check the **Specify on-startup behavior** check box, and choose one of the following:
  - **Advertise ourself as overloaded until BGP has converged**
  - **Specify time to advertise ourself as overloaded after reboot**

    In the **Time to advertise ourself as overloaded** field, enter the seconds to wait until the router advertises that it is overloaded. The range is 5 to 86400 seconds.
- Check the **Don't advertise IP prefixes learned from other protocols when overload bit it set** check box to exclude IP prefixes.
- Check the **Don't advertise IP prefixes learned from another ISIS level when overload bit is set** check box to exclude IP prefixes.

c) Configure the partial route calculation (PRC) intervals:

- In the **PRC Interval** field, enter an amount of time for the router to wait between partial route calculations (PRCs). The range is 1 to 120 seconds. The default is 5 seconds.
- In the **Initial wait for PRC** field, enter the initial PRC calculation delay (in milliseconds) after a topology change. The range is 1 to 120.000 milliseconds. The default is 2000 milliseconds.
- In the **Minimum wait between first and second PRC** field, enter the amount of time in milliseconds that you want the router to wait between PRCs. The range is 1 to 120,000 milliseconds. The default is 5000 milliseconds.

d) Configure the intervals for SPF calculations for Level 1 and Level 2:

Note
Check the **Use level 1 values also for level 2** check box if you want both levels to have the same values.

- In the **SPF Calculation Interval** field, enter an amount of time for the router to wait between SPF calculations. The range is 1 to 120 seconds. The default is 10 seconds.
- In the **Initial wait for SPF calculation** field, enter the amount of time for the router to wait for an SPF calculation. The range is 1 to 120.000 milliseconds. The default is 5500 milliseconds.
- In the **Minimum wait between first and second SPF calculation** field, enter the amount of time in milliseconds that you want the router to wait between SPF calculations. The range is 1 to 120,000 milliseconds. The default is 5500 milliseconds.

Step 5 Click **Apply**.

Step 6 Choose **Configuration > Device Setup > Routing > ISIS > IPv6 Address Family > Redistribution**.

The **Redistribution** pane displays a table of the redistribution routes.

Step 7 Click **Add** to add a new redistribution route, or to click **Edit** to edit an redistribution route in the table.

The **Add Redistribution** or **Edit Redistribution** dialog box appears. You can also double-click an entry in the table to edit that entry.

a) From the **Source Protocol** drop-down list, choose the protocol (BGP, Connected, EIGRP, OSPF, RIP, or Static) from which you want to redistribute routes into the ISIS domain.
b) From the **Process ID** drop-down list, choose a process ID for the source protocol.
c) From the **Route Level** drop-down list, choose Level-1, Level- 2, or Level 1-2.
d) (Optional) In the **Metric** field, enter a metric for the redistributed route . The range is 1 to 4294967295.
e) For the **Metric Type**, click the internal or external radio button.
f) From the **Route Map** drop-down list, choose a route map that should be examined to filter the networks to be redistributed, or click **Manage** to add a new route map or edit an existing route map. See Define a Route Map for the procedure for configuring route maps.
g) Check one or more of the **Match** check boxes -Internal, External 1, External 2, NSSA External 1, and NSSA External 2 check boxes to redistribute routes from an OSPF network.

This step is only applicable for redistribution from OSPF networks.

Step 8 Click **OK**.

Step 9 Click **Apply**.

Attached Bit Configuration

In the following example, the attached-bit will stay set when the router matches 49.00aa in the L2 CLNS routing table:

```
ciscoasa(config)# router isis
 ciscoasa(config-router)# clns filter-set L2_backbone_connectivity permit 49.00aa
 ciscoasa(config-router)# route-map check-for-L2_backbone_connectivity
 ciscoasa(config-router)# match clns address L2_backbone_connectivity
ciscoasa(config)# router isis
 ciscoasa(config-router)#set-attached-bit route-map check-for-L2_backbone_connectivity
```

```
 ciscoasa(config-router)# end
 ciscoasa# show clns route 49.00aa

Known via "isis", distance 110, metric 30, Dynamic Entry
Routing Descriptor Blocks:
 via tr2, Serial0
   isis, route metric is 30, route version is 58
```

Configure IS-IS IPv6 Address Family

You can set up adjacency, SPF, and you can define conditions for redistributing routes from another routing domain into IS-IS (redistribution) for IPv6 addresses.

Before you begin

Before you can enable IS-IS route authentication, you must enable IS-IS and set up an area. See Enable IS-IS Routing Globally, on page 902 for the procedure.

Make sure that IPv6 is enabled on at least one interface before trying to add a neighbor, or ASDM returns an error message indicating that the configuration failed.

Procedure

Step 1 Choose **Configuration > Device Setup > Routing > ISIS > IPv6 Address Family > General**.

a) Check the **Perform adjacency check** check box for the router to check on nearby IS routers.

b) In the **Administrative Distance** field, enter a distance for the route. The range is 1 to 255. The default is 1.

Administrative distance is a parameter used to compare routes among different routing protocols. In general, the higher the value, the lower the trust rating. And administrative distance of 255 means that the routing information source cannot be trusted at all and should be ignored. The range is 1 to 255. The default is 1.

You can use the distance option to configure the administrative distances applied to IS-IS routes when they are inserted into the Routing Information Base (RIB), and influence the likelihood of these routes being preferred over routes to the same destination addresses discovered by other protocols.

c) In the **Maximum number of forward paths** field, enter the maximum number of IS routes that can be installed in a routing table. The range is 1 to 8.

d) Check the **Distribute default route** check box to configure an IS routing process to distribute a default route, and then choose the default route from the drop-down list or click **Manage** to create a new route. See Define a Route Map, on page 805for the procedure for creating a new route.

Step 2 Click **Apply**.

Step 3 Choose **Configuration > Device Setup > Routing > ISIS > IPv6 Address Family > SPF**.

a) Check the **Signal other routers not to use this router as an intermediate hop in their SPF calculations** check box if you want to exclude this device, and configure the following:

- Check the **Specify on-startup behavior** check box, and choose one of the following:
  - **Advertise ourself as overloaded until BGP has converged**
  - **Specify time to advertise ourself as overloaded after reboot**

In the **Time to advertise ourself as overloaded** field, enter the seconds to wait until the router advertises that it is overloaded. The range is 5 to 86,400 seconds.

- Check the **Don't advertise IP prefixes learned from other protocols when overload bit it set** check box to exclude IP prefixes.
- Check the **Don't advertise IP prefixes learned from another ISIS level when overload bit is set** check box to exclude IP prefixes.

b) Configure the partial route calculation (PRC) intervals: .

- In the **PRC Interval** field, enter an amount of time for the router to wait between partial route calculations (PRCs). The range is 1 to 120 seconds. The default is 5 seconds.
- In the **Initial wait for PRC** field, enter the amount of time for the router to wait for a PRC. The range is 1 to 120.000 milliseconds. The default is 2000 milliseconds.
- In the **Minimum wait between first and second PRC** field, enter the amount of time in milliseconds that you want the router to wait between PRCs. The range is 1 to 120.000 milliseconds. The default is 5000 milliseconds.

c) Configure the intervals for SPF calculations for Level 1 and Level 2:

Note
Check the **Use level 1 values also for level 2** check box if you want both levels to have the same values.

- In the **SPF Calculation Interval** field, enter an amount of time for the router to wait between SPF calculations. The range is 1 to 120 seconds. The default is 10 seconds.
- In the **Initial wait for SPF calculation** field, enter the amount of time for the router to wait for an SPF calculation. The range is 1 to 120.000 milliseconds. The default is 5500 milliseconds.
- In the **Minimum wait between first and second SPF calculation** field, enter the amount of time in milliseconds that you want the router to wait between SPF calculations. The range is 1 to 120,000 milliseconds. The default is 5500 milliseconds.

Step 4 Click **Apply**.

Step 5 Choose **Configuration > Device Setup > Routing > ISIS > IPv6 Address Family > Redistribution**.

The **Redistribution** pane displays a table of the redistribution routes.

Step 6 Click **Add** to add a new redistribution route, or to click **Edit** to edit an redistribution route in the table.

The **Add Redistribution** or **Edit Redistribution** dialog box appears. You can also double-click an entry in the table to edit that entry.

a) From the **Source Protocol** drop-down list, choose the protocol (BGP, Connected, EIGRP, OSPF, RIP, or Static) from which you want to redistribute routes into the ISIS domain.
b) From the **Process ID** drop-down list, choose a process ID for the source protocol.
c) From the **Route Level** drop-down list, choose Level-1, Level- 2, or Level 1-2.
d) (Optional) In the **Metric** field, enter a metric for the redistributed route . The range is 1 to 4294967295.
e) For the **Metric Type**, click the **internal** or **external** radio button to specify the type of metric for the destination routing protocol.

f) From the **Route Map** drop-down list, choose a route map that should be examined to filter the networks to be redistributed, or click **Manage** to add a new route map or edit an existing route map. See Define a Route Map for the procedure for configuring route maps.

g) Check one or more of the **Match** check boxes -Internal, External 1, External 2, NSSA External 1, and NSSA External 2 check boxes to redistribute routes from an OSPF network.

This step is only applicable for redistribution from OSPF networks.

Step 7 Click **OK**.

Step 8 Click **Apply**.

Monitoring IS-IS

You can use the following screens to monitor the IS-IS routing process.

- **Monitoring > Routing > ISIS Neighbors** This pane shows information about each IS-IS neighbor.

  Each row represents one IS-IS neighbor. For each neighbor, the list includes the system ID, type, interface, IP address, the state (active, idle and so on), the hold time, and the circuit ID.

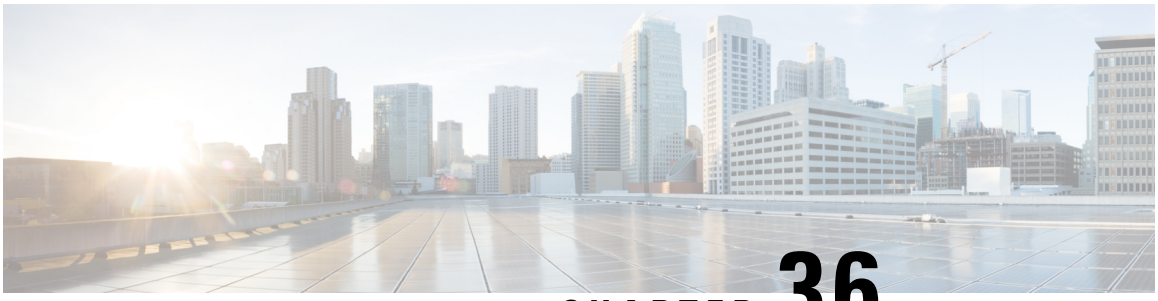
- **Monitoring > Routing > ISIS Rib** This pane displays the local IS-IS Routing Information Base (RIB) table.

- **Monitoring > Routing > ISIS IPv6 Rib** This pane displays the local IPv6 IS-IS RIB table.

History for IS-IS

Table 42: Feature History for IS-IS

Feature Name	Platform Releases	Feature Information
IS-IS routing	9.6(1)	The ASA now supports the Intermediate System to Intermediate System (IS-IS) routing protocol. Support was added for routing data, performing authentication, and redistributing and monitoring routing information using the IS-IS routing protocol. We introduced the following screens: **Configuration > Device Setup > Routing > ISIS** **Monitoring > Routing > ISIS**



CHAPTER 36

EIGRP

This chapter describes how to configure the ASA to route data, perform authentication, and redistribute routing information using the Enhanced Interior Gateway Routing Protocol (EIGRP).

- About EIGRP, on page 919
- Guidelines for EIGRP, on page 921
- Configure an EIGRP Process, on page 922
- Configure EIGRP, on page 922
- Customize EIGRP, on page 925
- Configure an EIGRPv6 Process, on page 937
- Monitoring for EIGRP, on page 942
- History for EIGRP, on page 943

About EIGRP

EIGRP is an enhanced version of IGRP developed by Cisco. 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 propagate until an alternate route is found. Its support for variable-length subnet masks permits routes to be automatically summarized on a network number boundary. In addition, EIGRP can be configured to summarize on any bit boundary at any interface. EIGRP does not make periodic updates. Instead, it sends partial updates only when the metric for a route changes. Propagation of partial updates is automatically bounded so that only those routers that need the information are updated. As a result of these two capabilities, EIGRP consumes significantly less bandwidth than IGRP.

Neighbor discovery is the process that the ASA uses to dynamically learn of other routers on directly attached networks. EIGRP routers send out multicast hello packets to announce their presence on the network. When the ASA 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 ASA.

The hello packets are sent out as multicast messages. No response is expected to a hello message. The exception to this is for statically defined neighbors. If you use the **neighbor** command, or configure the Hello Interval

in ASDM, to configure a neighbor, the hello messages sent to that neighbor are sent as unicast messages. Routing updates and acknowledgements are sent out 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. This is the time in which the ASA can expect to receive a hello packet from that neighbor. If the ASA does not receive a hello packet from that neighbor within the hold time advertised by that neighbor, the ASA considers that neighbor to be unavailable.

The EIGRP protocol uses four key algorithm technologies, four key technologies, including neighbor discovery/recovery, Reliable Transport Protocol (RTP), and DUAL, which is important for route computations. DUAL saves all routes to a destination in the topology table, 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 used for packet forwarding that has a least-cost path to a destination. The feasibility calculation guarantees that the path is not part of a routing loop.

If a feasible successor is not found in the topology table, a route recomputation must occur. During route recomputation, DUAL queries the EIGRP neighbors for a route, who in turn query their neighbors. Routers that do no have a feasible successor for the route return an unreachable message.

During route recomputation, DUAL marks the route as active. By default, the ASA waits for three minutes to receive a response from its neighbors. If the ASA 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.



Note EIGRP neighbor relationships are not supported through the IPsec tunnel without a GRE tunnel.

EIGRPv6

EIGRP for IPv6 can be configured just like EIGRP IPv4. EIGRPv6 communicates only with IPv6 peers and advertises only IPv6 routes. EIGRPv6 is similar to EIGRPv4 in many ways than one:

- DUAL is used for route calculation and selection with the same metrics.
- It is scalable to large network implementations.
- Neighbor, routing, and topology tables are maintained.
- Both equal-cost load balancing and unequal-cost load balancing are offered.

However, EIGRPv6 differ from EIGRPv4 in many ways, such as:

- The network command is not used in IPv6; EIGRP is configured using links.
- You must explicitly enable EIGRPv6 on each interface during configuration.

Null0 and EIGRP

By default, EIGRP advertises the Null0 route to the peer as summary route to prevent the router that is advertising the summary, from forwarding any packets that it does not have a route.

For example, consider the two routers, R1 and R2. The three interfaces on R1 have these networks- 192.168.0.0/24, 192.168.1.0/24, and 192.168.3.0/24. Configure R1 with summary route 192.168.0.0/22 and

advertise it to R2. When R2 has an IP packet for 192.168.2.x, it would forward it to R1. R1, would drop the packet as it does not have 192.168.2.x in its routing table. However, if R1 is also connected to an ISP and it has a default route pointing to the ISP, the 192.168.2.x packet is forwarded to the ISP. To prevent this forwarding action, EIGRP generates an entry that matches the summary route, pointing to Null0. Thus, when packets for 192.168.2.x are received, R1 will drop the packet instead of using the default route.

Guidelines for EIGRP

Firewall Mode Guidelines

Supported only in routed firewall mode. Transparent firewall mode is not supported.

Cluster Guidelines

For a cluster in individual interface mode, EIGRP can form neighbor relationships with cluster peers using cluster pools as router IDs.

IPv6 Guidelines

Supports IPv6 routing.

Context Guidelines

- EIGRP instances cannot form adjacencies with each other across shared interfaces because, by default, inter-context exchange of multicast traffic is not supported across shared interfaces. However, you can use the static neighbor configuration under EIGRP process configuration under EIGRP process to bring up EIGRP neighbourship on a shared interface.
- Inter-context EIGRP on separate interfaces is supported.

Redistribution Guidelines

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.

Additional Guidelines

- A maximum of one EIGRP process is supported.
- EIGRP adjacency flap occurs whenever a configuration change is applied which results in modifying the routing information (sent or received) from neighbors especially in distribute lists, offset lists, and changes to summarization. After the routers are synchronized, EIGRP reestablishes the adjacency between neighbors. When an adjacency is torn down and reestablished, all learned routes between the neighbors are erased and the entire synchronization between the neighbors is performed newly with the new distribute list.

- There is no restriction on the maximum number of EIGRP neignbours. However, to prevent unnecessary EIGRP flap, we recommend you to limit the number to 500 per unit.

Configure an EIGRP Process

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **EIGRP**.

Step 2 Enable the EIGRP routing process by checking the **Enable this EIGRP process** check box on the Process Instances tab. See Enable EIGRP, on page 923 or Enable EIGRP Stub Routing, on page 924.

Step 3 Define the networks and interfaces that will participate in EIGRP routing on the Setup > Networks tab. See Define a Network for an EIGRP Routing Process, on page 925 for more information.

Step 4 (Optional) Define route filters on the Filter Rules pane. Route filtering provides more control over the routes that are allowed to be sent or received in EIGRP updates. See Filter Networks in EIGRP, on page 932 for more information.

Step 5 (Optional) Define route redistribution in the Redistribution pane.

You can redistribute routes discovered by RIP and OSPF to the EIGRP routing process. You can also redistribute static and connected routes to the EIGRP routing process. See Redistribute Routes Into EIGRP, on page 930 for more information.

Step 6 (Optional) Define static EIGRP neighbors on the Static Neighbor pane.

See Define an EIGRP Neighbor, on page 930 for more information.

Step 7 (Optional) Define summary addresses on the Summary Address pane.

See Configure the Summary Aggregate Addresses on Interfaces, on page 927 for more information about defining summary addresses.

Step 8 (Optional) Define interface-specific EIGRP parameters on the Interfaces pane. These parameters include EIGRP message authentication, hold time, hello interval, delay metric, and the use of split-horizon. See Configure Interfaces for EIGRP, on page 926 for more information.

Step 9 (Optional) Control the sending and receiving of default route information in EIGRP updates on the Default Information pane. By default, default routes are sent and accepted. See Configure Default Information in EIGRP, on page 935 for more information.

Configure EIGRP

This section describes how to enable the EIGRP process on your system. After you have enabled EIGRP, see the following sections to learn how to customize the EIGRP process on your system.

Enable EIGRP

You can only enable one EIGRP routing process on the ASA.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **EIGRP** > **Setup**.

The EIGRP Setup pane appears.

The three tabs on the main EIGRP Setup pane used to enable EIGRP are as follows:

- The Process Instances tablets you enable an EIGRP routing process for each context. Single context mode and multiple context mode are both supported. See Enable EIGRP, on page 923 and the Enable EIGRP Stub Routing, on page 924 for more information.
- The Networks tablets you specify the networks used by the EIGRP routing process. For an interface to participate in EIGRP routing, it must fall within the range of addresses defined by the network entries. For directly connected and static networks to be advertised, they must also fall within the range of the network entries. See Define a Network for an EIGRP Routing Process, on page 925 for more information.
- The Passive Interfaces tablets you configure one or more interfaces as passive interfaces. In EIGRP, a passive interface does not send or receive routing updates.The Passive Interface table lists each interface that is configured as a passive interface.

Step 2 Check the **Enable this EIGRP process** check box.

You can only enable one EIGRP routing process on the device. You must enter an autonomous system number (AS) for the routing process in the EIGRP Process field before you can save your changes.

Step 3 In the EIGRP Process field, enter the autonomous system (AS) number for the EIGRP process. The AS number can be from 1 to 65535.

Step 4 (Optional) Click **Advanced** to configure the EIGRP process settings, such as the router ID, default metrics, stub routing, neighbor changes, and the administrative distances for the EIGRP routes.

Step 5 Click the **Networks** tab.

Step 6 To add a new network entry, click **Add**.

The **Add EIGRP Network** dialog box appears. To remove a network entry, choose an entry in the table and click **Delete**.

Step 7 Choose the AS number of the EIGRP routing process from the drop-down list.

Step 8 Enter the IP address of the networks to participate in the EIGRP routing process in the IP Address field.

Note
To change a network entry, you must first remove the entry and then add a new one. You cannot edit existing entries.

Step 9 Enter a network mask to apply to the IP address in the Network Mask field.

Step 10 Click **OK**.

Enable EIGRP Stub Routing

You can enable, and configure the ASA as an EIGRP stub router. Stub routing decreases memory and processing requirements on the ASA. As a stub router, the ASA does not need to maintain a complete EIGRP routing table because it forwards all nonlocal traffic to a distribution router. Generally, the distribution router need not send anything more than a default route to the stub router.

Only specified routes are propagated from the stub router to the distribution router. As a stub router, the ASA responds to all queries for summaries, connected routes, redistributed static routes, external routes, and internal routes with the message "inaccessible." When the ASA is configured as a stub, it sends a special peer information packet to all neighboring routers to report its status as a stub router. Any neighbor that receives a packet informing it of the stub status will not query the stub router for any routes, and a router that has a stub peer will not query that peer. The stub router depends on the distribution router to send the correct updates to all peers.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **EIGRP** > **Setup**.

The **EIGRP Setup** pane appears.

Step 2 Check the **Enable EIGRP routing** check box.

Step 3 In the **EIGRP Process** field, enter the autonomous system (AS) number for the EIGRP process. The AS number can be from 1 to 65535.

Step 4 Click **Advanced** to configure the EIGRP stub routing process.

The **Edit EIGRP Process Advanced Properties** dialog box appears.

Step 5 In the **Stub** area on the **Edit EIGRP Process Advanced Properties** dialog box, choose one or more of the following EIGRP stub routing processes:

- Stub Receive only—Configures the EIGRP stub routing process to receive route information from the neighbor routers but does not send route information to the neighbors. If this option is selected, you cannot select any of the other stub routing options.
- Stub Connected—Advertises connected routes.
- Stub Static—Advertises static routes.
- Stub Redistributed—Advertises redistributed routes.
- Stub Summary—Advertises summary routes.

Step 6 Click **OK**.

Step 7 Click the **Networks** tab.

Step 8 Click **Add** to add a new network entry.

The **Add EIGRP Network** dialog box appears. To remove a network entry, choose the entry in the table and click **Delete**.

Step 9 Choose the AS number of the EIGRP routing process from the drop-down list.

Step 10 Enter the IP address of the networks to participate in the EIGRP routing process in the **IP Address** field.

Note
To change a network entry, you must first remove the entry and then add a new one. You cannot edit existing entries.

Step 11 Enter a network mask to apply to the IP address in the **Network Mask** field.

Step 12 Click **OK**.

Customize EIGRP

This section describes how to customize the EIGRP routing.

Define a Network for an EIGRP Routing Process

The Network table lets you specify the networks used by the EIGRP routing process. For an interface to participate in EIGRP routing, it must fall within the range of addresses defined by the network entries. For directly connected and static networks to be advertised, they must also fall within the range of the network entries.

The Network table displays the networks configured for the EIGRP routing process. Each row of the table displays the network address and associated mask configured for the specified EIGRP routing process.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **EIGRP** > **Setup**.

The **EIGRP Setup** pane appears.

Step 2 Check the **Enable EIGRP routing** check box.

Step 3 In the **EIGRP Process** field, enter the autonomous system (AS) number for the EIGRP process. The AS number can be from 1 to 65535.

Step 4 Click the **Networks** tab.

Step 5 Click **Add** to add a new network entry.

The **Add EIGRP Network** dialog box appears. To remove a network entry, choose the entry in the table and click **Delete**.

Step 6 Choose the AS number of the EIGRP routing process from the drop-down list.

Step 7 Enter the IP address of the networks to participate in the EIGRP routing process in the **IP Address** field.

Note
To change a network entry, you must first remove the entry and then add a new one. You cannot edit existing entries.

Step 8 Enter a network mask to apply to the IP address in the **Network Mask** field.

Step 9 Click **OK**.

Configure Interfaces for EIGRP

If you have an interface that you do not want to have participate in EIGRP routing, but that is attached to a network that you want advertised, you can configure the ASA that includes the network to which the interface is attached, and prevent that interface from sending or receiving EIGRP updates.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **EIGRP** > **Setup**.

The **EIGRP Setup** pane appears.

Step 2 Check the **Enable EIGRP routing** check box.

Step 3 Click **OK**.

Step 4 Choose **Configuration** > **Device Setup** > **Routing** > **EIGRP** > **Interfaces**.

The **Interface** pane appears and displays the EIGRP interface configurations. The **Interface Parameters** table displays all of the interfaces on the ASA and lets you modify the following settings on a per-interface basis:

- Authentication key and mode.
- The EIGRP hello interval and hold time.
- The interface delay metric used in EIGRP metric calculations.
- The use of split-horizon on the interface.

Step 5 Choose an interface entry by double-clicking an interface entry, or choose the entry and click **Edit**.

The **Edit EIGRP Interface Entry** dialog box appears.

Step 6 In the **EIGRP Process** field, enter the AS number for the EIGRP process. The AS number can range from 1 to 65535.

Step 7 In the **Hello Interval** field, enter the interval between EIGRP hello packets sent on an interface.

Valid values range from 1 to 65535 seconds. The default value is 5 seconds.

Step 8 In the **Hold Time** field, enter the hold time, in seconds. Valid values range from 1 to 65535 seconds. The default value is 15 seconds.

Step 9 Check the **Enable** check box for Split Horizon.

Step 10 In the **Delay** field, enter the delay value. The delay time is in tens of microseconds. Valid values range from 1 to 16777215.

Step 11 Check the **Enable MD5 Authentication** check box to enable MD5 authentication of EIGRP process messages.

Step 12 Enter the Key or Key ID values.

- In the **Key** field, enter the key to authenticate EIGRP updates. The key can contain up to 16 characters.
- In the **Key ID** field, enter the key identification value. Valid values range from 1 to 255.

Step 13 Click **OK**.

Configure Passive Interfaces

You can configure one or more interfaces as passive interfaces. In EIGRP, a passive interface does not send or receive routing updates. In ASDM, the Passive Interface table lists each interface that is configured as a passive interface.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **EIGRP** > **Setup**.

The **EIGRP Setup** pane appears.

Step 2 Check the **Enable EIGRP routing** check box.

Step 3 Click **OK**.

Step 4 Click the **Passive Interfaces** tab.

Step 5 Choose the interface that you want to configure from the drop-down list.

Step 6 Check the **Suppress routing updates on all interfaces** check box to specify all interfaces as passive. Even if an interface is not shown in the Passive Interface table, it will be configured as passive when the check box is checked.

Step 7 Click **Add** to add a passive interface entry.

The **Add EIGRP Passive Interface** dialog box appears. Choose the interface that you want to make passive and click **Add**. To remove a passive interface, choose the interface in the table and click **Delete**.

Step 8 Click **OK**.

Configure the Summary Aggregate Addresses on Interfaces

You can configure a summary addresses on a per-interface basis. You need to manually define summary addresses if you want to create summary addresses that do not occur at a network number boundary or if you want to use summary addresses on an ASA with automatic route summarization disabled. If any more specific routes are in the routing table, EIGRP will advertise the summary address out the interface with a metric equal to the minimum of all more specific routes.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **EIGRP** > **Interfaces**.

The **Interface** pane shows the EIGRP interface configurations. The Interface Parameters table shows all of the interfaces on the ASA and lets you modify the settings on a per-interface basis. For more information about these settings, see Configure Interfaces for EIGRP, on page 926.

Step 2 To configure the EIGRP parameters for an interface, double-click an interface entry or select the entry and click **Edit**.

Step 3 Click **OK**.

Step 4 Choose **Configuration** > **Device Setup** > **Routing** > **EIGRP** > **Summary Address**.

The **Summary Address** pane displays a table of the statically-defined EIGRP summary addresses. By default, EIGRP summarizes subnet routes to the network level. You can create statically defined EIGRP summary addresses to the subnet level from the **Summary Address** pane.

Step 5 Click **Add** to add a new EIGRP summary address, or to click **Edit** to edit an existing EIGRP summary address in the table.

The **Add Summary Address** or **Edit Summary Address** dialog box appears. You can also double-click an entry in the table to edit that entry.

Step 6 In the **EIGRP Process** field, enter the autonomous system (AS) number for the EIGRP process. The AS number can be from 1 to 65535.

Step 7 In the **Interface** drop-down list, choose the interface from which the summary address is advertised.

Step 8 In the **IP Address** field, enter the IP address of the summary route.

Step 9 In the **Netmask** field, choose or enter the network mask to apply to the IP address.

Step 10 Enter the administrative distance for the route in the **Administrative Distance** field. If left blank, the route has the default administrative distance of 5.

Step 11 Click **OK**.

Change the Interface Delay Value

The interface delay value is used in EIGRP distance calculations. You can modify this value on a per-interface basis.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **EIGRP** > **Interface**.

The **Interface** pane shows the EIGRP interface configurations. The **Interface Parameters** table shows all of the interfaces on the ASA and lets you modify the settings on a per-interface basis. For more information about these settings, see Configure Interfaces for EIGRP, on page 926.

Step 2 Double-click an interface entry or choose the Interface entry and click **Edit** to configure the delay value in the EIGRP parameters for an interface.

The **Edit EIGRP Interface Entry** dialog box appears.

Step 3 In the **Delay** field, enter the delay time, which is in tens of microseconds. Valid values are from 1 to 16777215.

Step 4 Click **OK**.

Enable EIGRP Authentication on an Interface

EIGRP route authentication provides MD5 authentication of routing updates from the EIGRP routing protocol. The MD5 keyed digest in each EIGRP packet prevents the introduction of unauthorized or false routing messages from unapproved sources.

EIGRP route authentication is configured on a per-interface basis. All EIGRP neighbors on interfaces configured for EIGRP message authentication must be configured with the same authentication mode and key for adjacencies to be established.

Note Before you can enable EIGRP route authentication, you must enable EIGRP.

Procedure

Step 1 In the man ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **EIGRP** > **Setup**.

The **EIGRP Setup** pane appears.

Step 2 Check the **Enable EIGRP routing** check box.

Step 3 In the **EIGRP Process** field, enter the autonomous system (AS) number for the EIGRP process. The AS number can range from 1 to 65535.

Step 4 Click the **Networks** tab.

Step 5 Click **Add** to add a new network entry.

The **Add EIGRP Network** dialog box appears. To remove a network entry, choose the entry in the table and click **Delete**.

Step 6 Choose the AS number of the EIGRP routing process from the drop-down list.

Step 7 In the **IP Address** field, enter the IP address of the networks to participate in the EIGRP routing process.

Note
To change a network entry, you must first remove the entry and then add a new one. You cannot edit existing entries.

Step 8 In the **Network Mask** field, choose or enter a network mask to apply to the IP address.

Step 9 Click **OK**.

Step 10 Choose **Configuration** > **Device Setup** > **Routing** > **EIGRP** > **Interfaces**.

The **Interface** pane displays the EIGRP interface configurations. The **Interface Parameters** table displays all of the interfaces on the ASA and lets you modify the settings on a per-interface basis. For more information about these settings, see Configure Interfaces for EIGRP, on page 926.

Step 11 Check the **Enable MD5 Authentication** check box to enable MD5 authentication of EIGRP process messages. After you check this check box, provide one of the following:

- In the **Key** field, enter the key to authenticate EIGRP updates. The key can include up to 16 characters.
- In the **Key ID** field, enter the key identification value. Valid values range from 1 to 255.

Step 12 Click **OK**.

Define an EIGRP Neighbor

EIGRP hello packets are sent as multicast packets. If an EIGRP neighbor is located across a non broadcast network, such as a tunnel, you must manually define that neighbor. When you manually define an EIGRP neighbor, hello packets are sent to that neighbor as unicast messages.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **EIGRP** > **Setup**.

The **EIGRP Setup** pane appears.

Step 2 Check the **Enable EIGRP routing** check box.

Step 3 In the **EIGRP Process** field, enter the AS number for the EIGRP process. The AS number can range from 1 to 65535.

Step 4 Choose **Configuration** > **Device Setup** > **Routing** > **EIGRP** > **Static Neighbor**.

The **Static Neighbor** pane appears and displays the statically-defined EIGRP neighbors. An EIGRP neighbor sends EIGRP routing information to and receives EIGRP routing information from the ASA. Normally, neighbors are dynamically discovered through the neighbor discovery process. However, on point-to-point, nonbroadcast networks, you must statically define the neighbors.

Each row of the **Static Neighbor** table displays the EIGRP autonomous system number for the neighbor, the neighbor IP address, and the interface through which the neighbor is available.

From the **Static Neighbor** pane, you can add or edit a static neighbor.

Step 5 Click **Add** or **Edit** to add or edit a EIGRP static neighbor.

The **Add** or **Edit EIGRP Neighbor Entry** dialog box appears.

Step 6 Choose the **EIGRP AS** number from the drop-down list for the EIGRP process for which the neighbor is being configured.

Step 7 Choose the **Interface Name** from the **Interface Name** drop-down list, which is the interface through which the neighbor is available.

Step 8 Enter the IP address of the neighbor in the **Neighbor IP Address** field.

Step 9 Click **OK**.

Redistribute Routes Into EIGRP

You can redistribute routes discovered by RIP and OSPF into the EIGRP routing process. You can also redistribute static and connected routes into the EIGRP routing process. You do not need to redistribute connected routes if they fall within the range of a **network** statement in the EIGRP configuration.



Note For RIP only: Before you begin this procedure, you must create a route map to further define which routes from the specified routing protocol are redistributed in to the RIP routing process.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **EIGRP** > **Setup**.

The **EIGRP Setup** pane appears.

Step 2 Check the **Enable EIGRP routing** check box.

Step 3 In the **EIGRP Process** field, enter the AS number for the EIGRP process. The AS number can range from 1 to 65535.

Step 4 Choose **Configuration** > **Device Setup** > **Routing** > **EIGRP** > **Redistribution**.

The **Redistribution** pane displays the rules for redistributing routes from other routing protocols to the EIGRP routing process. When redistributing static and connected routes to the EIGRP routing process, metrics are not required to be configured, although this is recommended. Each row of the **Redistribution** pane table includes a route redistribution entry.

Step 5 Click **Add** to add a new redistribution rule. If you are editing an existing redistribution rule, go to Step 6.

The **Add EIGRP Redistribution Entry** dialog box appears.

Step 6 Choose the address in the table and click **Edit** to edit an existing EIGRP static neighbor, You can also double-click an entry in the table to edit that entry.

The **Edit EIGRP Redistribution Entry** dialog box appears.

Step 7 Choose the AS number of the EIGRP routing process to which the entry applies from the drop-down list.

Step 8 In the **Protocol** area, click the radio button next to one of the following protocols for the routing process:

- **Static** to redistribute static routes to the EIGRP routing process. Static routes that fall within the scope of a network statement are automatically redistributed into EIGRP; you do not need to define a redistribution rule for them.
- **Connected** to redistribute connected routes into the EIGRP routing process. Connected routes that fall within the scope of a network statement are automatically redistributed into EIGRP; you do not need to define a redistribution rule for them.
- **RIP** to redistributes routes discovered by the RIP routing process to EIGRP.
- **OSPF** to redistribute routes discovered by the OSPF routing process to EIGRP.

Step 9 In the **Optional Metrics** area, choose one of the following metrics used for the redistributed route:

- **Bandwidth**, which is the EIGRP bandwidth metric in kilobits per second. Valid values range from 1 to 4294967295.
- **Delay**, which is the EIGRP delay metric, in 10-microsecond units. Valid values range from 0 to 4294967295.
- **Reliability**, which is the EIGRP reliability metric. Valid values range from 0 to 255; 255 indicates 100 percent reliability.
- **Loading**, which is the EIGRP effective bandwidth (loading) metric. Valid values range from 1 to 255; 255 indicates 100 percent loaded.

- **MTU**, which is the MTU of the path. Valid values range from 1 to 65535.

Step 10 Choose the route map from the **Route Map** drop-down list to define which routes are redistributed into the EIGRP routing process. For more details about how to configure a route map, see Route Maps, on page 803.

Step 11 In the **Optional OSPF Redistribution** area, click one of the following OSPF radio buttons to further specify which OSPF routes are redistributed into the EIGRP routing process:

- **Match Internal** to match routes internal to the specified OSPF process.
- **Match External 1** to match type 1 routes external to the specified OSPF process.
- **Match External 2** to match type 2 routes external to the specified OSPF process.
- **Match NSSA-External 1** to match type 1 routes external to the specified OSPF NSSA.
- **Match NSSA-External 2** to match type 2 routes external to the specified OSPF NSSA.

Step 12 Click **OK**.

Filter Networks in EIGRP

Note Before you begin this process, you must create a standard ACL that defines the routes that you want to advertise. That is, create a standard ACL that defines the routes that you want to filter from sending or receiving updates.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **EIGRP** > **Setup**.

The **EIGRP Setup** pane appears.

Step 2 Check the **Enable EIGRP routing** check box.

Step 3 In the **EIGRP Process** field, enter the AS number for the EIGRP process. The AS number can range from 1 to 65535.

Step 4 Choose **Configuration** > **Device Setup** > **Routing** > **EIGRP** > **Filter Rules**.

The **Filter Rules** pane appears and displays the route filtering rules configured for the EIGRP routing process. Filter rules let you control which routes are accepted or advertised by the EIGRP routing process.

Each row of the **Filter Rule** table describes a filter rule for a specific interface or routing protocol. For example, a filter rule with a direction of in on the outside interface would apply filtering to any EIGRP updates received on the outside interface. A filter rule with a direction of out with OSPF 10 specified as the routing protocol would apply the filter rules to routes redistributed into the EIGRP routing process in outbound EIGRP updates.

Step 5 Click **Add** to add a filter rule. If you are editing an already existing filter rule, skip to Step 6.

The **Add Filter Rules** dialog box appears.

Step 6 To edit a filter rule, choose the filter rule in the table and click **Edit**.

The **Edit Filter Rules** dialog box appears. You can also double-click a filter rule to edit the rule. To remove a filter rule, choose the filter rule in the table and click **Delete**.

Step 7 Choose the AS number from the drop-down list of the EIGRP routing process to which the entry applies.

Step 8 Choose the direction of the filter routes from the drop-down list.

Choose **in** for rules that filter routes from incoming EIGRP routing updates. Choose **out** to filter routes from EIGRP routing updates that are sent by the ASA.

If you choose **out**, the **Routing** process field becomes active. Choose the type of route to be filtered. You can filter routes redistributed from static, connected, RIP, and OSPF routing processes. Filters that specify a routing process filter those routes from updates sent on all interfaces.

Step 9 Enter the OSPF process ID in the **ID** field.

Step 10 Click the **Interface** radio button and choose the interface to which the filter applies.

Step 11 Click **Add** or **Edit** to define an ACL for the filter rule. Clicking **Edit** opens the **Network Rule** dialog box for the selected network rule.

Step 12 In the **Action** drop-down list, choose **Permit** to allow the specified network to be advertised; choose **Deny** to prevent the specified network from being advertised.

Step 13 In the **IP Address** field, type IP address of the network being permitted or denied. To permit or deny all addresses, use the IP address **0.0.0.0** with a network mask of **0.0.0.0**.

Step 14 From the **Netmask** drop-down list, choose the network mask applied to the network IP address. You can type a network mask into this field or select one of the common masks from the list.

Step 15 Click **OK**.

Customize the EIGRP Hello Interval and Hold Time

The ASA periodically sends hello packets to discover neighbors and to learn when neighbors become unreachable or inoperative. By default, hello packets are sent every 5 seconds.

The hello packet advertises the ASA hold time. The hold time indicates to EIGRP neighbors the length of time the neighbor should consider the ASA reachable. If the neighbor does not receive a hello packet within the advertised hold time, then the ASA is considered unreachable. By default, the advertised hold time is 15 seconds (three times the hello interval).

Both the hello interval and the advertised hold time are configured on a per-interface basis. We recommend setting the hold time to be at minimum three times the hello interval.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **EIGRP** > **Setup**.

The **EIGRP Setup** pane appears.

Step 2 Check the **Enable EIGRP routing** check box.

Step 3 Click **OK**.

Step 4 Choose **Configuration** > **Device Setup** > **Routing** > **EIGRP** > **Interfaces**.

The **Interface** pane appears and displays all of the EIGRP interface configurations.

Step 5 Double-click an interface entry or choose the entry and click **Edit**.

The **Edit EIGRP Interface Entry** dialog box appears.

Step 6 Choose the EIGRP AS number from the drop-down list, which is populated from system numbers that were set up when you enabled the EIGRP routing process.

Step 7 In the **Hello Interval** field, enter the interval between EIGRP hello packets sent on an interface.

Valid values range from 1 to 65535 seconds. The default value is 5 seconds.

Step 8 In the **Hold Time** field, specify the hold time, in seconds.

Valid values range from 1 to 65535 seconds. The default value is 15 seconds.

Step 9 Click **OK**.

Disable Automatic Route Summarization

Automatic route summarization is enabled by default. The EIGRP routing process summarizes on network number boundaries. This can cause routing problems if you have noncontiguous networks.

For example, if you have a router with the networks 192.168.1.0, 192.168.2.0, and 192.168.3.0 connected to it, and those networks all participate in EIGRP, the EIGRP routing process creates the summary address 192.168.0.0 for those routes. If an additional router is added to the network with the networks 192.168.10.0 and 192.168.11.0, and those networks participate in EIGRP, they will also be summarized as 192.168.0.0. To prevent the possibility of traffic being routed to the wrong location, you should disable automatic route summarization on the routers creating the conflicting summary addresses.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **EIGRP** > **Setup**.

The **EIGRP Setup** pane appears.

Step 2 Check the **Enable EIGRP routing** check box.

Step 3 Click the **Process Instance** tab.

Step 4 Click **Advanced**.

Step 5 In the **Summary** area, uncheck the **Auto-Summary** check box.

Note
This setting is enabled by default.

Step 6 Click **OK**.

Configure Default Information in EIGRP

You can control the sending and receiving of default route information in EIGRP updates. By default, default routes are sent and accepted. Configuring the ASA to disallow default information to be received causes the candidate default route bit to be blocked on received routes. Configuring the ASA to disallow default information to be sent disables the setting of the default route bit in advertised routes.

In ASDM, the Default Information pane displays a table of rules for controlling the sending and receiving of default route information in EIGRP updates. You can have one in and one out rule for each EIGRP routing process (only one process is currently supported).

By default, default routes are sent and accepted. To restrict or disable the sending and receiving of default route information, perform the following steps:

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **EIGRP** > **Setup**.

The main **EIGRP Setup** pane appears.

Step 2 Check the **Enable EIGRP routing** check box.

Step 3 Click **OK**.

Step 4 Do one of the following:

- Click **Add** to create a new entry.
- To edit an entry, double-click the entry in the table or select an entry in the table and click **Edit**.

  The **Add Default Information** or **Edit Default Information** dialog box appears for that entry. The EIGRP AS number is automatically selected in the EIGRP field.

Step 5 In the **Direction** field, choose the direction for the rule from the following options:

- **in**—The rule filters default route information from incoming EIGRP updates.
- **out**—The rule filters default route information from outgoing EIGRP updates.

You can have one in rule and one out rule for each EIGRP process.

Step 6 Add network rules to the network rule table. The network rules define which networks are allowed and which are not when receiving or sending default route information. Repeat the following steps for each network rule you are adding to the default information filter rule.

a) Click **Add** to add a network rule. Double-click an existing network rule to edit the rule.
b) In the **Action** field, click **Permit** to allow the network or **Deny** to block the network.
c) Enter the IP address and network mask of the network being permitted or denied by the rule in the IP Address and Network Mask fields.

   To deny all default route information from being accepted or sent, enter **0.0.0.0** as the network address and choose **0.0.0.0** as the network mask.

d) Click **OK** to add the specified network rule to the default information filter rule.

Step 7 Click **OK** to accept the default information filter rule.

Disable EIGRP Split Horizon

Split horizon controls the sending of EIGRP update and query packets. When split horizon is enabled on an interface, update and query packets are not sent for destinations for which this interface is the next hop. Controlling update and query packets in this manner reduces the possibility of routing loops.

By default, split horizon is enabled on all interfaces.

Split horizon blocks route information from being advertised by a router out of any interface from which that information originated. This behavior usually optimizes communications among multiple routing devices, particularly when links are broken. However, with nonbroadcast networks, there may be situations where this behavior is not desired. For these situations, including networks in which you have EIGRP configured, you may want to disable split horizon.

If you disable split horizon on an interface, you must disable it for all routers and access servers on that interface.

To disable EIGRP split horizon, perform the following steps:

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **EIGRP** > **Interface**. For EIGRPv6, choose **Configuration** > **Device Setup** > **Routing** > **EIGRPv6** > **Interface**.

The **Interface** pane appears and displays the EIGRP interface configurations.

Step 2 Double-click an interface entry or choose the entry and click **Edit**.

The **Edit EIGRP Interface Entry** or **Edit EIGRPv6 Interface Entry**(EIGRPv6) dialog box appears.

Step 3 Choose the EIGRP Autonomous system (AS) number from the drop-down list, which is populated from system numbers that were set up when you enabled the EIGRP routing process.

Step 4 Uncheck the **Split Horizon** check box.

Step 5 Click **OK**.

Restart the EIGRP Process

You can restart an EIGRP process or clear redistribution or clear counters.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **EIGRP** > **Setup**.

The **EIGRP Setup** pane appears.

Step 2 Click **Reset**.

Configure an EIGRPv6 Process

This section describes how to enable and configure the EIGRP IPv6 process on your system.

Enable EIGRPv6

You can only enable one EIGRPv6 routing process on the ASA.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **EIGRPv6** > **Setup**.

The **EIGRPv6 Setup** pane appears.

Step 2 In the **Process Instances** tab, check the **Enable this EIGRPv6 process** check box.

You can only enable one EIGRP routing process on the device. You must enter an autonomous system number (AS) for the routing process in the EIGRP Process field before you can save your changes.

Step 3 In the **EIGRPv6 Process** field, enter the autonomous system (AS) number for the EIGRP process. The AS number can be from 1 to 65535.

Step 4 (Optional) Click **Advanced** to configure the EIGRP process settings, such as the router ID, default metrics, stub routing, neighbor changes, and the administrative distances for the EIGRP routes.

Step 5 Click the **Passive Interfaces** tab.

Step 6 Choose the interface that you want to configure from the drop-down list.

Step 7 Check the **Suppress routing updates on all interfaces** check box to specify all interfaces as passive. Even if an interface is not shown in the Passive Interface table, it will be configured as passive when the check box is checked.

Step 8 Click **Add** to add a passive interface entry.

The **Add EIGRPv6 Passive Interface** dialog box appears. Choose the interface that you want to make passive and click **Add**. To remove a passive interface, choose the interface in the table and click **Delete**.

Step 9 Click **OK**.

Filter Rules in EIGRPv6



Note Before you begin this process, you must create a standard ACL that defines the routes that you want to advertise. That is, create a standard ACL that defines the routes that you want to filter from sending or receiving updates.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **EIGRPv6** > **Setup**.

The **EIGRP Setup** pane appears.

Step 2 Check the **Enable this EIGRPv6 Process** check box.

Step 3 In the **EIGRPv6 Process** field, enter the AS number for the EIGRP process. The AS number can range from 1 to 65535.

Step 4 Choose **Configuration** > **Device Setup** > **Routing** > **EIGRPv6** > **Filter Rules**.

The **Filter Rules** pane appears and displays the route filtering rules configured for the EIGRP routing process. Filter rules let you control which routes are accepted or advertised by the EIGRPv6 routing process.

Each row of the **Filter Rule** table describes a filter rule for a specific interface or routing protocol. For example, a filter rule with a direction of in on the outside interface would apply filtering to any EIGRP updates received on the outside interface. A filter rule with a direction of out would apply the filter rules to routes advertised in outbound EIGRP updates.

Step 5 Click **Add** to add a filter rule. If you are editing an already existing filter rule, skip to next step.

The **Add EIGRPv6 Filter Rules** dialog box appears.

Step 6 To edit a filter rule, choose the filter rule in the table and click **Edit**.

The **Edit EIGRPv6 Filter Rules** dialog box appears. You can also double-click a filter rule to edit the rule. To remove a filter rule, choose the filter rule in the table and click **Delete**.

Step 7 Choose the AS number from the drop-down list of the EIGRPv6 routing process to which the entry applies.

Step 8 Choose the direction of the filter routes from the drop-down list.

Choose **in** for rules that filter routes from incoming EIGRP routing updates. Choose **out** to filter routes from EIGRP routing updates that are sent by the ASA.

Step 9 From the **Interface Name** drop-down, choose the interface to which the filter applies.

Step 10 Click **OK**.

Configure Interfaces for EIGRPv6

If you have an interface that you do not want to have participate in EIGRP routing, but that is attached to a network that you want advertised, you can configure the ASA that includes the network to which the interface is attached, and use the tab to prevent that interface from sending or receiving EIGRP updates.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **EIGRP** > **Setup**.

The **EIGRPv6 Setup** pane appears.

Step 2 Check the **Enable this EIGRPv6 Process** check box and enter the AS number.

Step 3 Click **Apply**.

Step 4 Choose **Configuration** > **Device Setup** > **Routing** > **EIGRPv6** > **Interface**.

The **Interface** pane appears and displays the EIGRPv6 interface configurations. The **Interface Parameters** table displays all of the interfaces on the ASA and lets you modify the following settings on a per-interface basis:

- The EIGRP hello interval and hold time.
- The use of split-horizon on the interface.
- Specify the statically-defined EIGRP summary addresses.

Step 5 Choose an interface entry by double-clicking an interface entry, or choose the entry and click **Edit**.

The **Edit EIGRPv6 Interface Entry** dialog box appears.

Step 6 In the **EIGRP Process** field, enter the AS number for the EIGRP process. The AS number can range from 1 to 65535.

Step 7 In the **Hello Interval** field, enter the interval between EIGRP hello packets sent on an interface.

Valid values range from 1 to 65535 seconds. The default value is 5 seconds.

Step 8 In the **Hold Time** field, enter the hold time, in seconds. Valid values range from 1 to 65535 seconds. The default value is 15 seconds.

Step 9 Check the **Enable** check box for **Split Horizon**.

Step 10 In the **Summary Address** field, enter the statically-defined EIGRP summary addresses. You can enter the address to the subnet level.

Step 11 Click **OK**.

Configure Passive Interfaces for EIGRPv6

You can configure one or more interfaces as passive interfaces. In EIGRPv6, a passive interface does not send or receive routing updates. In ASDM, the Passive Interface table lists each interface that is configured as a passive interface.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **EIGRPv6** > **Setup**.

The **EIGRPv6 Setup** pane appears.

Step 2 Check the **Enable the EIGRPv6 Process** check box and enter the AS number.

Step 3 Click the **Passive Interfaces** tab.

Step 4 Choose the process that you want to configure from the drop-down list.

Step 5 Check the **Suppress routing updates on all interfaces** check box to specify all interfaces as passive. Even if an interface is not shown in the Passive Interface table, it will be configured as passive when the check box is checked.

Step 6 Click **Add** to add a passive interface entry.

The **Add EIGRPv6 EIGRP Passive Interface** dialog box appears. Choose the AS process number and the interface that you want to make passive and click **Add**. To remove a passive interface, choose the interface in the table and click **Delete**.

Step 7 Click **OK**.

Redistribute Routes Into EIGRPv6

You can redistribute routes discovered OSPF, BGP, ISIS into the EIGRP IPv6 routing process. You can also redistribute static and connected routes into the EIGRP routing process.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **EIGRPv6** > **Setup**.

The **EIGRPv6 Setup** pane appears.

Step 2 Check the **Enable this EIGRPv6 Process** check box.

Step 3 In the **EIGRPIPv6 Process** field, enter the AS number for the EIGRP process. The AS number can range from 1 to 65535.

Step 4 Choose **Configuration** > **Device Setup** > **Routing** > **EIGRPv6** > **Redistribution**.

The **Redistribution** pane displays the rules for redistributing routes from other routing protocols to the EIGRP routing process. When redistributing static and connected routes to the EIGRP routing process, metrics are not required to be configured, although this is recommended. Each row of the **Redistribution** pane table includes a route redistribution entry.

Step 5 Click **Add** to add a new redistribution rule. If you are editing an existing redistribution rule, go to next step.

The **Add EIGRPv6 Redistribution Entry** dialog box appears.

Step 6 Choose the address in the table and click **Edit** to edit an existing EIGRP static neighbor, You can also double-click an entry in the table to edit that entry.

The **Edit EIGRPv6 Redistribution Entry** dialog box appears.

Step 7 Choose the AS number of the EIGRP routing process to which the entry applies from the drop-down list.

Step 8 In the **Protocol** area, click the radio button next to one of the following protocols for the routing process:

- **Static** to redistribute static routes to the EIGRP routing process. Static routes that fall within the scope of a network statement are automatically redistributed into EIGRP; you do not need to define a redistribution rule for them.
- **Connected** to redistribute connected routes into the EIGRP routing process. Connected routes that fall within the scope of a network statement are automatically redistributed into EIGRP; you do not need to define a redistribution rule for them.
- **BGP** to redistribute routes discovered by the BGP routing process to EIGRP.

- **ISIS** to redistribute routes discovered by the ISIS routing process to EIGRP. You can choose **Route Level** under **Optional Metrics**.
- **OSPF** to redistribute routes discovered by the OSPF routing process to EIGRP.

Step 9 In the **Optional Metrics** area, choose one of the following metrics used for the redistributed route:

- **Bandwidth**, which is the EIGRP bandwidth metric in kilobits per second. Valid values range from 1 to 4294967295.
- **Delay**, which is the EIGRP delay metric, in 10-microsecond units. Valid values range from 0 to 4294967295.
- **Reliability**, which is the EIGRP reliability metric. Valid values range from 0 to 255; 255 indicates 100 percent reliability.
- **Loading**, which is the EIGRP effective bandwidth (loading) metric. Valid values range from 1 to 255; 255 indicates 100 percent loaded.
- **MTU**, which is the MTU of the path. Valid values range from 1 to 65535.

Step 10 Choose the route map from the **Route Map** drop-down list to define which routes are redistributed into the EIGRP routing process. For more details about how to configure a route map, see Route Maps, on page 803.

Step 11 In the **Optional OSPF Redistribution** area, click one of the following OSPF radio buttons to further specify which OSPF routes are redistributed into the EIGRP routing process:

- **Match Internal** to match routes internal to the specified OSPF process.
- **Match External 1** to match type 1 routes external to the specified OSPF process.
- **Match External 2** to match type 2 routes external to the specified OSPF process.
- **Match NSSA-External 1** to match type 1 routes external to the specified OSPF NSSA.
- **Match NSSA-External 2** to match type 2 routes external to the specified OSPF NSSA.

Step 12 Click **OK**.

Define an EIGRPv6 Neighbor

EIGRP hello packets are sent as multicast packets. If an EIGRP neighbor is located across a non broadcast network, such as a tunnel, you must manually define that neighbor. When you manually define an EIGRP neighbor, hello packets are sent to that neighbor as unicast messages.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **EIGRPv6** > **Setup**.

The **EIGRPv6 Setup** pane appears.

Step 2 Check the **Enable this EIGRPv6 Process** check box.

Step 3 In the **EIGRPv6 Process** field, enter the AS number for the EIGRP process. The AS number can range from 1 to 65535.

Step 4 Choose **Configuration** > **Device Setup** > **Routing** > **EIGRPv6** > **Static Neighbor**.

The **Static Neighbor** pane appears and displays the statically-defined EIGRPv6 neighbors. An EIGRPv6 neighbor sends EIGRPv6 routing information to and receives EIGRPv6 routing information from the ASA. Normally, neighbors are dynamically discovered through the neighbor discovery process. However, on point-to-point, nonbroadcast networks, you must statically define the neighbors.

Each row of the **Static Neighbor** table displays the EIGRPv6 autonomous system number for the neighbor, the neighbor IP address, and the interface through which the neighbor is available.

From the **Static Neighbor** pane, you can add or edit a static neighbor.

Step 5 Click **Add** or **Edit** to add or edit a EIGRP static neighbor.

The **Add** or **Edit EIGRPv6 Neighbor Entry** dialog box appears.

Step 6 Choose the **EIGRP AS** number from the drop-down list for the EIGRP process for which the neighbor is being configured.

Step 7 Choose the **Interface Name** from the **Interface Name** drop-down list, which is the interface through which the neighbor is available.

Step 8 Enter the IP address of the neighbor in the **Neighbor IP Address** field.

Step 9 Click **OK**.

Monitoring for EIGRP

You can use the following commands to monitor the EIGRP routing process. For examples and descriptions of the command output, see the command reference. Additionally, you can disable the logging of neighbor change messages and neighbor warning messages.

To monitor or disable various EIGRP routing statistics, perform the following steps:

Procedure

Step 1 In the main ASDM window, choose **Monitoring** > **Routing** > **EIGRP Neighbor**.

Each row represents one EIGRP neighbor. For each neighbor, the list includes its IP address, the interface to which the neighbor is connected, the holdtime, the uptime, the queue length, the sequence number, the smoothed round trip time, and the retransmission timeout. The list of possible state changes are the following:

- NEW ADJACENCY—A new neighbor has been established.
- PEER RESTARTED—The other neighbor initiates the reset of the neighbor relationship. The router getting the message is not the one resetting the neighbor.
- HOLD TIME EXPIRED—The router has not heard any EIGRP packets from the neighbor within the hold-time limit.

- RETRY LIMIT EXCEEDED—EIGRP did not receive the acknowledgment from the neighbor for EIGRP reliable packets, and EIGRP has already tried to retransmit the reliable packet 16 times without any success.
- ROUTE FILTER CHANGED—The EIGRP neighbor is resetting because there is a change in the route filter.
- INTERFACE DELAY CHANGED—The EIGRP neighbor is resetting because there is a manual configuration change in the delay parameter on the interface.
- INTERFACE BANDWIDTH CHANGED—The EIGRP neighbor is resetting because there is a manual configuration change in the interface bandwidth on the interface.
- STUCK IN ACTIVE—The EIGRP neighbor is resetting because EIGRP is stuck in active state. The neighbor getting reset is the result of the stuck-in-active state.

Step 2 Click the EIGRP neighbor that you want to monitor.

Step 3 To remove the current list of neighbors, click **Clear Neighbors**.

Step 4 To refresh the current list of neighbors, click **Refresh**.

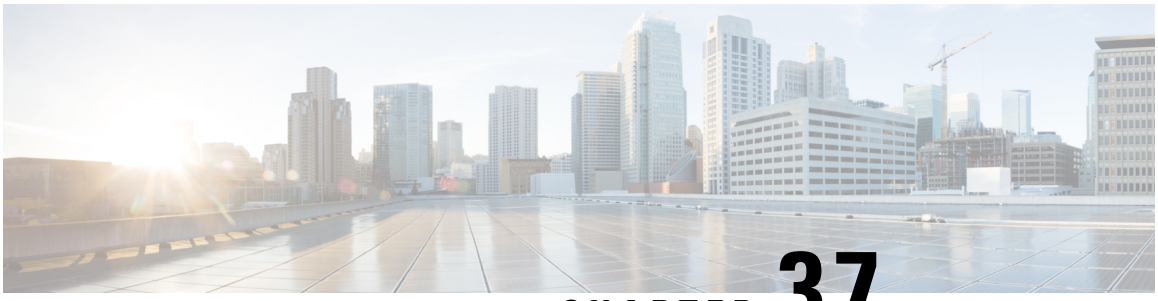
Note
By default, neighbor change and neighbor warning messages are logged.

History for EIGRP

Table 43: Feature History for EIGRP

Feature Name	Platform Releases	Feature Information
EIGRP Support	7.0(1)	Support was added for routing data, performing authentication, and redistributing and monitoring routing information using the Enhanced Interior Gateway Routing Protocol (EIGRP). We introduced the following screen:**Configuration** > **Device Setup** > **Routing** > **EIGRP** .
Dynamic Routing in Multiple Context Mode	9.0(1)	EIGRP routing is supported in multiple context mode. We modified the following screen: Configuration > Device Setup > Routing > EIGRP > Setup.
Clustering	9.0(1)	For EIGRP, bulk synchronization, route synchronization, and layer 2 load balancing are supported in the clustering environment.
EIGRP Auto-Summary	9.2(1)	For EIGRP, the Auto-Summary field is now disabled by default. We modified the following screen: Configuration > Device Setup > Routing > EIGRP > Setup > Edit EIGRP Process Advanced Properties.

Feature Name	Platform Releases	Feature Information
EIGRPv6 Support	9.20(1)	IPv6 support was added for routing data, performing authentication, and redistributing and monitoring routing information using the Enhanced Interior Gateway Routing Protocol (EIGRP). We introduced **Setup**, **Filter Rules**, **Interface**, **Redistribution**, and **Static Neighbor** screens under the following menu:**Configuration** > **Device Setup** > **Routing** > **EIGRPv6**.



CHAPTER 37

Multicast Routing

This chapter describes how to configure the ASA to use the multicast routing protocol.

- About Multicast Routing, on page 945
- Guidelines for Multicast Routing, on page 948
- Enable Multicast Routing, on page 949
- Customize Multicast Routing, on page 949
- Monitoring for PIM, on page 963
- Example for Multicast Routing, on page 964
- History for Multicast Routing, on page 965

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 ASA 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 ASA supports both stub multicast routing and PIM multicast routing. However, you cannot configure both concurrently on a single ASA.

Note The UDP and non-UDP transports are both supported for multicast routing. However, the non-UDP transport has no FastPath optimization.

Stub Multicast Routing

Stub multicast routing provides dynamic host registration and facilitates multicast routing. When configured for stub multicast routing, the ASA acts as an IGMP proxy agent. Instead of fully participating in multicast routing, the ASA forwards IGMP messages to an upstream multicast router, which sets up delivery of the

multicast data. When configured for stub multicast routing, the ASA cannot be configured for PIM sparse or bidirectional mode. You must enable PIM on the interfaces participating in IGMP stub multicast routing.

The ASA 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 ASA is the PIM RP, use the untranslated outside address of the ASA as the RP address.

PIM Source Specific Multicast Support

The ASA does not support PIM Source Specific Multicast (SSM) functionality and related configuration. However, the ASA 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.

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 ASA 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. For information about how to configure multicast groups, see Configure a Multicast Group, on page 959.

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. In Individual Interface clustering, units do not act independently. All data and routing packets are processed and forwarded by the control unit. Data units drop all packets that have been sent.

Guidelines for Multicast Routing

Context Mode

Supported in single context mode.

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 rule on the inbound interface to allow traffic to the multicast host, such as 224.1.2.3. However, you cannot specify a destination interface for the rule, or it cannot be applied to multicast connections during initial connection validation.
- PIM/IGMP Multicast routing is not supported on interfaces in a traffic zone.
- Do not configure ASA 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 ASA. 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.

Enable Multicast Routing

Enabling multicast routing on the ASA, enables IGMP and PIM on all data interfaces by default, but not on the management interface for most models (see Management Slot/Port Interface, on page 562 for interfaces that do not allow through traffic). 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.

To enable multicast routing on the management interface, you must explicitly set a multicast boundary on the management interface.

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 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **Multicast**.

Step 2 In the Multicast pane, check the **Enable Multicast** routing check box.

Checking this check box enables IP multicast routing on the ASA. 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.

Customize Multicast Routing

This section describes how to customize multicast routing.

Configure Stub Multicast Routing and Forward IGMP Messages

Note Stub multicast routing is not supported concurrently with PIM sparse and bidirectional modes.

An ASA acting as the gateway to the stub area does not need to participate in PIM sparse mode or bidirectional mode. Instead, you can configure it 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. To configure the ASA as an IGMP proxy agent, forward the host join and leave messages from the stub area interface to an upstream interface. You must also enable PIM on the interfaces participating in stub mode multicast routing.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **Multicast**.

Step 2 In the Multicast pane, check the **Enable Multicast routing** check box.

Step 3 Click **Apply** to save your changes.

Step 4 Choose **Configuration** > **Device Setup** > **Routing** > **Multicast** > **IGMP** > **Protocol**.

Step 5 To modify the specific interface from which you want to forward IGMP messages, select the interface and click **Edit**.

The Configure IGMP Parameters dialog box appears.

Step 6 From the **Forward Interface** drop-down list, choose the specific interface from which you want to forward IGMP messages.

Step 7 Click **OK** to close this dialog box, then click **Apply** to save your changes.

Configure a Static Multicast Route

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 ASA 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 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **Multicast** > **MRoute**.

Step 2 Choose **Add** or **Edit**.

The Add or Edit Multicast Route dialog box appears.

Use the Add Multicast Route dialog box to add a new static multicast route to the ASA. Use the Edit Multicast Route dialog box to change an existing static multicast route.

Step 3 In the Source Address field, enter the IP address of the multicast source. You cannot change this value when editing an existing static multicast route.

Step 4 Choose the network mask for the IP address of the multicast source from the Source Mask drop-down list.

Step 5 In the Incoming Interface area, click either the **RPF Interface** radio button to choose RPF to forward the route or the **Interface Name** radio button, then enter the following:

- In the Source Interface field, choose the incoming interface for the multicast route from the drop-down list.
- In the Destination Interface field, choose the destination interface that the route is forwarded through from the drop-down list.

Note
You can specify the interface or the RPF neighbor, but not both at the same time.

Step 6 In the Administrative Distance field, choose the administrative distance of the static multicast route. If the static multicast route has the same administrative distance as the unicast route, then the static multicast route takes precedence.

Step 7 Click **OK**.

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 setting on a per-interface basis.

Disable IGMP on an Interface

You can disable IGMP on specific interfaces. This information is useful if you know that there are no multicast hosts on a specific interface and you want to prevent the ASA from sending host query messages on that interface.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **Multicast** > **IGMP** > **Protocol**.

The Protocol pane displays the IGMP parameters for each interface on the ASA.

Step 2 Choose the interface that you want to disable and click **Edit**.

Step 3 To disable the specified interface, uncheck the **Enable IGMP** check box.

Step 4 Click **OK**.

The Protocol pane displays Yes if IGMP is enabled on the interface, or No if IGMP is disabled on the interface.

Configure IGMP Group Membership

You can configure the ASA to be a member of a multicast group. Configuring the ASA 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 If you want to forward multicast packets for a specific group to an interface without the ASA accepting those packets as part of the group, see Configure a Statically Joined IGMP Group, on page 952.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **Multicast** > **IGMP** > **Join Group**.

Step 2 Click **Add** or **Edit** > **in the** > **Join Group** > **pane**.

The Add IGMP Join Group dialog box allows you to configure an interface to be a member of a multicast group. The Edit IGMP Join Group dialog allows you to change existing membership information.

Step 3 In the Interface Name field, choose the interface name from the drop-down list. If you are editing an existing entry, you cannot change this value.

Step 4 In the Multicast Group Address field, enter the address of a multicast group to which the interface belongs. Valid group addresses range from 224.0.0.0 to 239.255.255.255.

Step 5 Click **OK**.

Configure a Statically Joined IGMP Group

Sometimes a group member cannot report its membership in the group because of some configuration, or there may be no members of a group on the network segment. However, 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.

In the main ASDM window, choose **Configuration > Routing > Multicast > IGMP > Static Group** to configure the ASA to be a statically connected member of a group. With this method, the ASA 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 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **Multicast** > **IGMP** > **Static Group**.

Step 2 Click **Add** or **Edit** in the **Static Group** pane.

Use the Add IGMP Static Group dialog box to statically assign a multicast group to an interface. Use the Edit IGMP Static Group dialog box to change existing static group assignments.

Step 3 In the Interface Name field, choose the interface name from the drop-down list. If you are editing an existing entry, you cannot change this value.

Step 4 In the Multicast Group Address field, enter the address of a multicast group to which the interface belongs. Valid group addresses range from 224.0.0.0 to 239.255.255.255.

Step 5 Click **OK**.

Control Access to Multicast Groups

You can control access to multicast groups by using access control lists.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **Multicast** > **IGMP** > **Access Group**.

The Access Group pane appears. The table entries in the Access Group pane are processed from the top down. Place more specific entries near the top of the table and more generic entries further down. For example, place an access group entry that permits a specific multicast group near the top of the table and an access group entry below that denies a range of multicast groups, including the group in the permit rule. The group is permitted because the permit rule is enforced before the deny rule.

Double-clicking an entry in the table opens the Add or Edit Access Group dialog box for the selected entry.

Step 2 Click **Add** or **Edit**.

The Add Access Group or Edit Access Group dialog box appears. The Add Access Group dialog box lets you add a new access group to the Access Group Table. The Edit Access Group dialog box lets you change information for an existing access group entry. Some fields may be dimmed when editing existing entries.

Step 3 Choose the interface name with which the access group is associated from the Interface drop-down list. You cannot change the associated interface when you are editing an existing access group.

Step 4 Choose permit from the Action drop-down list to allow the multicast group on the selected interface. Choose deny from the Action drop-down list to filter the multicast group from the selected interface.

Step 5 In the Multicast Group Address field, enter the address of the multicast group to which the access group applies.

Step 6 Enter the network mask for the multicast group address, or choose one of the common network masks from the Netmask drop-down list.

Step 7 Click **OK**.

Limit the Number of IGMP States on an Interface

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.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **Multicast** > **IGMP** > **Protocol**.

Step 2 Choose the interface you want to limit from the table on the Protocol pane, and click **Edit**.

The Configure IGMP Parameters dialog box appears.

Step 3 Enter the maximum number of host that can join on an interface, in the Group Limit field.

The default value is 500.Valid values range from 0 to 5000.

Note
Setting this value to 0 prevents learned groups from being added, but manually defined memberships are still permitted.

Step 4 Click **OK**.



Note When you change the IGMP limit on the interface with active joins on it, the new limit is not applicable to the existing groups. ASA validates the limit only when a new group is added to the interface or when the IGMP join timers expire. To apply the new limit with immediate effect, you must disenable and re-enable IGMP on the interface.

Modify the Query Messages to Multicast Groups

The ASA 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 ASA. If the ASA discovers that there are no local members of a multicast group still attached to an interface, it stops forwarding multicast packet 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 ASA does not receive a response to a host query within this amount of time, it deletes the group.

To change the query interval, query response time, and query timeout value, perform the following steps:

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **Multicast** > **IGMP** > **Protocol**.

Step 2 Choose the interface you want to limit from the table on the Protocol pane, and click **Edit**.

The Configure IGMP Parameters dialog box appears.

Step 3 Enter the interval in seconds, at which the designated router sends IGMP host-query messages, in the **Query Interval** field.

Valid values range from 1 to 3600 seconds. The default value is 125 seconds.

Note
If the ASA does not hear a query message on an interface for the specified timeout value, then the ASA becomes the designated router and starts sending the query messages.

Step 4 Enter the period of time, in seconds, in the **Query Timeout** field before which the ASA takes over as the requester for the interface after the previous requester has stopped doing so.

Valid values range from 60 to 300 seconds. The default value is 255 seconds.

Step 5 In the **Response Time** field, enter the maximum query response time advertised in IGMP queries, in seconds.

Values range from 1 to 25 seconds. The default value is 10 seconds.

Step 6 Click **OK**.

Change the IGMP Version

By default, the ASA runs IGMP Version 2, which enables several additional features.

All multicast routers on a subnet must support the same version of IGMP. The ASA does not automatically detect Version 1 routers and switch to Version 1. However, a mix of IGMP Version 1 and 2 hosts on the subnet works; the ASA running IGMP Version 2 works correctly when IGMP Version 1 hosts are present.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **Multicast** > **IGMP** > **Protocol**.

Step 2 Choose the interface whose version of IGMP you want to change from the table on the Protocol pane, and click **Edit**.

The Configure IGMP Interface dialog box appears.

Step 3 Choose the version number from the Version drop-down list.

Step 4 Click **OK**.

Configure PIM Features

Routers use PIM to maintain forwarding tables to use for forwarding multicast diagrams. When you enable multicast routing on the ASA, 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.

Enable and Disable PIM on an Interface

You can enable or disable PIM on specific interfaces.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **Multicast** > **PIM** > **Protocol**.

Step 2 Choose the interface on which you want to enable PIM from the table on the Protocol pane, and click **Edit**.

The Edit PIM Protocol dialog box appears.

Step 3 Check the **Enable PIM** check box. To disable PIM, uncheck this check box.

Step 4 Click **OK**.

Configure a Static Rendezvous Point Address

All routers within a common PIM sparse mode or bidir domain require knowledge of the PIM RP address. The address is statically configured using the **pim rp-address** command.

Note The ASA does not support Auto-RP.

You can configure the ASA to serve as 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).

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **Multicast** > **PIM** > **Rendezvous Points**.

Step 2 Click **Add** or **Edit**.

The Add or Edit Rendezvous Point dialog box appears. The Add Rendezvous Point dialog box lets you add a new entry to the Rendezvous Point table. The Edit Rendezvous Point dialog box lets you change an existing RP entry. Additionally, you can click **Delete** to remove the selected multicast group entry from the table.

These restrictions apply to RPs:

- You cannot use the same RP address twice.
- You cannot specify All Groups for more than one RP.

Step 3 In the Rendezvous Point Address field, enter the IP address for the RP.

When editing an existing RP entry, you cannot change this value.

Step 4 Check the **Use bi-directional forwarding** check box if the specified multicast groups are to operate in bidirectional mode. The Rendezvous Point pane displays Yes if the specified multicast groups are to operate in bidirectional mode and displays No if the specified groups are to operate in sparse mode. In bidirectional mode, if the ASA receives a multicast packet and has no directly connected members or PIM neighbors present, it sends a prune message back to the source.

Step 5 Click the **Use this RP for All Multicast Groups** radio button to use the specified RP for all multicast groups on the interface, or the **Use this RP for the Multicast Groups as specified below** radio button to designate the multicast groups to use with the specified RP.

For more information about multicast groups, see Configure a Multicast Group, on page 959.

Step 6 Click **OK**.

Configure the Designated Router 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, selecting 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 ASA has a DR priority of 1. You can change this value.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **Multicast** > **PIM** > **Protocol**.

Step 2 Choose the interface that you want to enable for PIM from the table on the Protocol pane, and click **Edit**.

The Edit PIM Protocol dialog box appears.

Step 3 In the DR Priority field, type the value for the designated router priority 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 ASA interface ineligible to become the default router.

Step 4 Click **OK**.

Configure and Filter PIM Register Messages

When the ASA is acting as an RP, you can restrict specific multicast sources from registering with it to prevent unauthorized sources from registering with the RP. The Request Filter pane lets you define the multicast sources from which the ASA will accept PIM register messages.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **Multicast** > **PIM** > **Request Filter**.

Step 2 Click **Add**.

The Request Filter Entry dialog box lets you define the multicast sources that are allowed to register with the ASA when the ASA acts as an RP. You create the filter rules based on the source IP address and the destination multicast address.

Step 3 From the Action drop-down list, choose Permit to create a rule that allows the specified source of the specified multicast traffic to register with the ASA, or choose Deny to create a rule that prevents the specified source of the specified multicast traffic from registering with the ASA.

Step 4 Type the IP address for the source of the register message, in the Source IP Address field.

Step 5 Type or choose the network mask from the drop-down list for the source of the register message, in the Source Netmask field.

Step 6 Type the multicast destination address, in the Destination IP Address field.

Step 7 Type or choose the network mask from the drop-down list for the multicast destination address, in the Destination Netmask field.

Step 8 Click **OK**.

Configure PIM Message Intervals

Router query messages are used to select 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 ASA sends PIM join or prune messages.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **Multicast** > **PIM** > **Protocol**.

Step 2 Choose the interface that you want to enable for PIM from the table on the Protocol pane, and click **Edit**.

The Edit PIM Protocol dialog box appears.

Step 3 Type the frequency, in seconds, at which the interface sends PIM hello messages, in the Hello Interval field.

Step 4 Type the frequency, in seconds, at which the interface sends PIM join and prune advertisements, in the Prune Interval field.

Step 5 Click **OK**.

Configure a Route Tree

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 ASA should join the shortest-path tree or use the shared tree, either for all multicast groups or only for specific multicast addresses.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **Multicast** > **PIM** > **Route Tree**.

Step 2 Click one of the following radio buttons:

- **Use Shortest Path Tree for All Groups**—Choose this option to use the shortest-path tree for all multicast groups.
- **Use Shared Tree for All Groups**—Choose this option to use the shared tree for all multicast groups.
- **Use Shared Tree for the Groups specified below**—Choose this option to use the shared tree for the groups specified in the Multicast Groups table. 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.

  To edit a multicast group, see Configure a Multicast Group, on page 959.

Configure a Multicast Group

Multicast groups are lists of access rules that define which multicast addresses are part of a group. A multicast group can include a single multicast address or a range of multicast addresses. Use the Add Multicast Group dialog box to create a new multicast group rule. Use the Edit Multicast Group dialog box to modify an existing multicast group rule.

To configure a multicast group, perform the following steps:

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **Multicast** > **PIM** > **Rendezvous Points**.

Step 2 The Rendezvous Point pane appears. Click the group that you want to configure.

The Edit Rendezvous Point dialog box appears.

Step 3 Click the **Use this RP for the Multicast Groups as specified below** radio button to designate the multicast groups to use with the specified RP.

Step 4 Click **Add** or **Edit**.

The Add or Edit Multicast Group dialog box appears.

Step 5 From the Action drop-down list, choose Permit to create a group rule that allows the specified multicast addresses, or choose Deny to create a group rule that filters the specified multicast addresses.

Step 6 In the Multicast Group Address field, type the multicast address associated with the group.

Step 7 From the Netmask drop-down list, choose the network mask for the multicast group address.

Step 8 Click **OK**.

Filter PIM Neighbors

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 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **Multicast** > **PIM** > **Neighbor Filter**.

Step 2 Choose the PIM neighbor that you want to configure from the table by clicking **Add**/**Edit**/**Insert**.

The Add/Edit/Insert Neighbor Filter Entry dialog box appears. It lets you create the ACL entries for the multicast boundary ACL. You can also delete a selected PIM neighbor entry.

Step 3 Choose the interface name from the Interface Name drop-down list.

Step 4 From the Action drop-down list, choose Permit or Deny for the neighbor filter ACL entry.

Choosing Permit allows the multicast group advertisements to pass through the interface. Choosing Deny 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 Enter the IP address of the multicast PIM group being permitted or denied, in the IP Address field. Valid group addresses range from 224.0.0.0 to 239.255.255.255.

Step 6 From the Netmask drop-down list, choose the netmask for the multicast group address.

Step 7 Click **OK**.

Configure a Bidirectional Neighbor Filter

The Bidirectional Neighbor Filter pane shows the PIM bidirectional neighbor filters, if any, that are configured on the ASA. A PIM bidirectional neighbor filter is an ACL that defines the neighbor devices that can participate in the DF election. If a PIM bidirectional neighbor filter is not configured for an interface, then 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.

When a PIM bidirectional neighbor filter configuration is applied to the ASA, an ACL appears in the running configuration with the name *interface-name*_multicast, in which the *interface-name* is the name of the interface to which the multicast boundary filter is applied. If an ACL with that name already exists, a number is appended to the name (for example, inside_multicast_1). This ACL defines which devices can become PIM neighbors of the ASA.

Bidirectional PIM allows multicast routers to keep reduced state information. All of the multicast routers in a segment must be bidirectionally enabled for bidir to elect a DF.

The PIM bidirectional neighbor filters enable the transition from a sparse-mode-only network to a bidir network by letting you specify the routers that should participate in the DF election, while still allowing all routers to participate in the sparse-mode domain. The bidir-enabled routers can elect a DF from among themselves, even when there are non-bidir routers on the segment. Multicast boundaries on the non-bidir routers prevent PIM messages and data from the bidir groups from leaking in or out of the bidir subset cloud.

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 bidir, then the DF election does not occur.
- If a denied neighbor supports bidir, then the DF election does not occur.
- If a denied neighbor does not support bidir, the DF election can occur.

Procedure

Step 1 In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **Multicast** > **PIM** > **Bidirectional Neighbor Filter**.

Step 2 Double-click an entry in the PIM Bidirectional Neighbor Filter table to access the Edit Bidirectional Neighbor Filter Entry dialog box for that entry.

Step 3 Choose the PIM neighbor that you want to configure from the table by clicking **Add**/**Edit**/**Insert**.

The Add/Edit/Insert Bidirectional Neighbor Filter Entry dialog box appears, which lets you create ACL entries for the PIM bidirectional neighbor filter ACL

Step 4 Choose the interface name from the Interface Name drop-down list. Select the interface for which you are configuring the PIM bidirectional neighbor filter ACL entry.

Step 5 From the Action drop-down list, choose Permit or Deny for the neighbor filter ACL entry.

Choose Permit to allow the specified devices to participate in the DF election process. Choose Deny to prevent the specified devices from participating in the DF election process.

Step 6 Enter the IP address of the multicast PIM group being permitted or denied. Valid group addresses range from 224.0.0.0 to 239.255.255.255, in the IP Address field.

Step 7 From the Netmask drop-down list, choose the netmask for the multicast group address.

Step 8 Click **OK**.

Configure the ASA as a Candidate BSR

You can configure the ASA as a candidate BSR.

Procedure

Step 1 In ASDM, choose **Configuration** > **Device Setup** > **Routing** > **Multicast** > **PIM** > **Bootstrap Router**.

Step 2 Check the **Configure this ASA as a candidate bootstrap router (CBSR)** check box to perform the CBSR set up.

a) Select the interface on the ASA from which the BSR address is derived to make it a candidate from the**Select Interface** drop-down list.

Note
This interface must be enabled with PIM.

b) Enter the length of a mask (32 bits maximum) that is to be ANDed with the group address before the hash function is called in the **Hash mask length** field. All groups with the same seed hash (correspond) to the same Rendezvous Point (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.

c) Enter the priority of the candidate BSR in the **Priority** field. 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 default value is 0.

Step 3 (Optional) Select an interface on which no PIM BSR messages will be sent or received, in the **Configure this ASA as a Border Bootstrap Router**section.

Step 4 Click **Apply**.

Configure a Multicast Boundary

Address scoping defines domain boundaries 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 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 is set up, no multicast data packets are allowed to flow across the boundary from either direction. The boundary 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 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 In the main ASDM window, choose **Configuration** > **Routing** > **Multicast** > **MBoundary**.

The MBoundary pane lets you 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 2 Click **Edit**.

The Edit Boundary Filter dialog box appears and displays the multicast boundary filter ACL. You can add and remove boundary filter ACL entries using this dialog box.

When the boundary filter configuration is applied to the ASA, the ACL appears in the running configuration with the name *interface-name*_multicast, where the *interface-name* is the name of the interface to which the multicast boundary filter is applied. If an ACL with that name already exists, a number is appended to the name (for example, inside_multicast_1).

Step 3 Choose the interface for which you are configuring the multicast boundary filter ACL from the Interface drop-down list.

Step 4 Check the **Remove any Auto-RP group range** check box to filter Auto-RP messages from sources denied by the boundary ACL. If the **Remove any Auto-RP group range** check box is unchecked, all Auto-RP messages are passed.

Step 5 Click **OK**.

Monitoring for PIM

To monitor or disable various PIM routing statistics, perform the following steps:

Procedure

Step 1 In the main ASDM window, choose **Monitoring** > **Routing** > **PIM** > **BSR Router**

The BSR Router configuration information is displayed.

Step 2 In the main ASDM window, choose **Monitoring** > **Routing** > **PIM** > **Multicast Routing Table**

The contents of the multicast routing table are displayed.

Step 3 In the main ASDM window, choose **Monitoring** > **Routing** > **PIM** > **MFIB**

The summary information about the number of IPv4 PIM multicast forwarding information base entries and interfaces are displayed.

Step 4 In the main ASDM window, choose **Monitoring** > **Routing** > **PIM** > **MFIB Active**

The summary information from the Multicast Forwarding Information Base (MFIB) about the rate at which active multicast sources are sending to multicast groups is displayed.

Step 5 In the main ASDM window, choose **Monitoring** > **Routing** > **PIM** > **Group Map**

The summary information from the Multicast Forwarding Information Base (MFIB) about the rate at which active multicast sources are sending to multicast groups is displayed.

a) Select **RP Timers** from the **Select PIM Group** drop-down list, to view the timer information for each group-to-PIM mode mapping.

Step 6 In the main ASDM window, choose **Monitoring** > **Routing** > **PIM** > **Neighbors**

The Protocol Independent Multicast (PIM) neighbor information is displayed.

Example for Multicast Routing

The following example shows how to enable and configure multicast routing with various optional processes:

1. In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **Multicast**.
2. In the Multicast pane, check the **Enable Multicast** routing check box, and click **Apply**.
3. In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **Multicast** > **MRoute**.
4. Click **Add** or **Edit**.

   The Add or Edit Multicast Route dialog box appears.

   Use the Add Multicast Route dialog box to add a new static multicast route to the ASA. Use the Edit Multicast Route dialog box to change an existing static multicast route.
5. In the Source Address field, enter the IP address of the multicast source. You cannot change this value when editing an existing static multicast route.
6. Choose the network mask for the IP address of the multicast source from the Source Mask drop-down list.
7. In the Incoming Interface area, click either the **RPF Interface** radio button to choose RPF to forward the route or the **Interface Name** radio button, then enter the following:
   - In the Source Interface field, choose the incoming interface for the multicast route from the drop-down list.
   - In the Destination Interface field, choose the destination interface to which the route is forwarded through the selected interface from the drop-down list.



Note You can specify the interface or the RPF neighbor, but not both at the same time.

8. In the Administrative Distance field, choose the administrative distance of the static multicast route. If the static multicast route has the same administrative distance as the unicast route, then the static multicast route takes precedence.
9. Click **OK**.
10. In the main ASDM window, choose **Configuration** > **Device Setup** > **Routing** > **Multicast** > **IGMP** > **Join Group**.

    The Join Group pane appears.
11. Click **Add** or **Edit**.

    The Add IGMP Join Group dialog box allows you to configure an interface to be a member of a multicast group. The Edit IGMP Join Group dialog box allows you to change existing membership information.
12. In the Interface Name field, choose the interface name from the drop-down list. If you are editing an existing entry, you cannot change this value.
13. In the Multicast Group Address field, enter the address of a multicast group to which the interface belongs. Valid group addresses range from 224.0.0.0 to 239.255.255.255.
14. Click **OK**.

History for Multicast Routing

Table 44: Feature History for Multicast Routing

Feature Name	Platform Releases	Feature Information
Multicast routing support	7.0(1)	Support was added for multicast routing data, authentication, and redistribution and monitoring of routing information using the multicast routing protocol. We introduced the following screen: Configuration > Device Setup > Routing > Multicast.
Clustering support	9.0(1)	Support was added for clustering.
Protocol Independent Multicast Source-Specific Multicast (PIM-SSM) pass-through support	9.5(1)	Support was added to allow PIM-SSM packets to pass through when multicast routing is enabled, unless the ASA is the Last-Hop Router. This allows greater flexibility in choosing a multicast group while also protecting against different attacks; hosts only receive traffic from explicitly-requested sources. We did not change any screens.

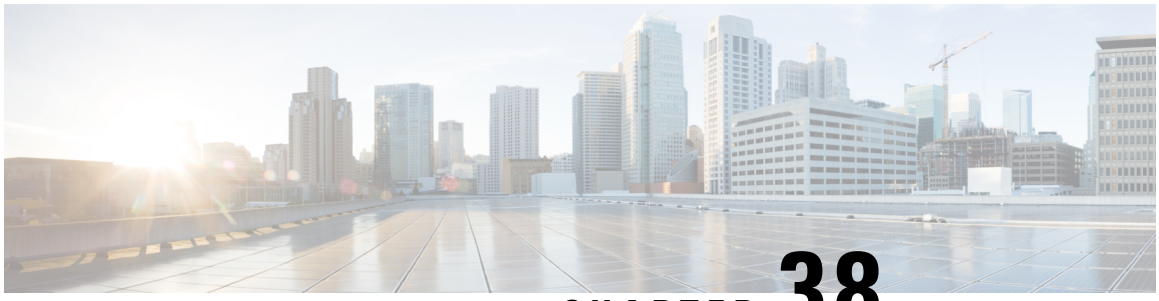
Feature Name	Platform Releases	Feature Information
Protocol Independent Multicast Bootstrap Router(BSR)	9.5(2)	Support was added for a new dynamic Rendezvous Point (RP) selection model that uses candidate routers for Rendezvous Point function and for relaying the Rendezvous Point information for a group. This feature provides a means of dynamically learning Rendezvous Points (RPs), which is very essential in large complex networks where an RP can periodically go down and come up. We introduced the following screens: Configuration > Device Setup > Routing > Multicast > PIM > Bootstrap Router
igmp limit increased	9.15(1) *Also in 9.12(4)*	igmp limit increased from 500 to 5000. We did not change any screens.



PART VI

AAA Servers and the Local Database

- AAA and the Local Database, on page 969
- RADIUS Servers for AAA, on page 981
- TACACS+ Servers for AAA, on page 1001
- LDAP Servers for AAA, on page 1009
- RSA SecurID Servers for AAA, on page 1019



CHAPTER 38

AAA and the Local Database

This chapter describes authentication, authorization, and accounting (AAA, pronounced “triple A”). AAA is a a set of services for controlling access to computer resources, enforcing policies, assessing usage, and providing the information necessary to bill for services. These processes are considered important for effective network management and security.

This chapter also describes how to configure the local database for AAA functionality. For external AAA servers, see the chapter for your server type.

- About AAA and the Local Database, on page 969
- Guidelines for the Local Database, on page 973
- Add a User Account to the Local Database, on page 974
- Test Local Database Authentication and Authorization, on page 975
- Monitoring the Local Database, on page 976
- History for the Local Database, on page 976

About AAA and the Local Database

This section describes AAA and the local database.

Authentication

Authentication provides a way to identify a user, typically by having the user enter a valid username and valid password before access is granted. The AAA server compares a user's authentication credentials with other user credentials stored in a database. If the credentials match, the user is permitted access to the network. If the credentials do not match, authentication fails and network access is denied.

You can configure the ASA to authenticate the following items:

- All administrative connections to the ASA, including the following sessions:
  - Telnet
  - SSH
  - Serial console
  - ASDM using HTTPS
  - VPN management access

- The **enable** command
- Network access
- VPN access

Authorization

Authorization is the process of enforcing policies: determining what types of activities, resources, or services a user is permitted to access. After a user is authenticated, that user may be authorized for different types of access or activity.

You can configure the ASA to authorize the following items:

- Management commands
- Network access
- VPN access

Accounting

Accounting measures the resources a user consumes during access, which may include the amount of system time or the amount of data that a user has sent or received during a session. Accounting is carried out through the logging of session statistics and usage information, which is used for authorization control, billing, trend analysis, resource utilization, and capacity planning activities.

Interaction Between Authentication, Authorization, and Accounting

You can use authentication alone or with authorization and accounting. Authorization always requires a user to be authenticated first. You can use accounting alone, or with authentication and authorization.

AAA Servers and Server Groups

The AAA server is a network server that is used for access control. Authentication identifies the user. Authorization implements policies that determine which resources and services an authenticated user may access. Accounting keeps track of time and data resources that are used for billing and analysis.

If you want to use an external AAA server, you must first create a AAA server group for the protocol that the external server uses, and add the server to the group. You can create more than one group per protocol, and separate groups for all protocols that you want to use. Each server group is specific to one type of server or service.

See the following topics for details on how to create the groups:

- Configure RADIUS Server Groups, on page 993
- Configure TACACS+ Server Groups, on page 1003
- Configure LDAP Server Groups, on page 1014
- Configure RSA SecurID AAA Server Groups, on page 1020

See the VPN configuration guide for more information on using HTTP Form.

Note From ASA 9.22.1, the Kerberos protocol is not supported. You can no longer use Kerberos for AAA services. We recommend that you to use other supported servers listed in the below table.

The following table summarizes the supported types of server and their uses, including the local database.

Table 45: Supported Services for AAA Servers

Server Type and Service	Authentication	Authorization	Accounting
Local Database			
Administrators	Yes	Yes	No
VPN Users	Yes	No	No
Firewall Sessions (AAA rules)	Yes	Yes	No
RADIUS			
Administrators	Yes	Yes	Yes
VPN Users	Yes	Yes	Yes
Firewall Sessions (AAA rules)	Yes	Yes	Yes
TACACS+			
Administrators	Yes	Yes	Yes
VPN Users	Yes	No	Yes
Firewall Sessions (AAA rules)	Yes	Yes	Yes
LDAP			
Administrators	Yes	No	No
VPN Users	Yes	Yes	No
Firewall Sessions (AAA rules)	Yes	No	No
SDI (RSA SecurID)			
Administrators	Yes	No	No
VPN Users	Yes	No	No
Firewall Sessions (AAA rules)	Yes	No	No
HTTP Form			
Administrators	No	No	No

Server Type and Service	Authentication	Authorization	Accounting
VPN Users	Yes	No	No
Firewall Sessions (AAA rules)	No	No	No

Notes

- RADIUS—Accounting for administrators does not include command accounting.
- RADIUS—Authorization for firewall sessions is supported with user-specific access lists only, which are received or specified in a RADIUS authentication response.
- TACACS+—Accounting for administrators includes command accounting.
- HTTP Form—Authentication and SSO operations for clientless SSL VPN user sessions only.

About the Local Database

The ASA maintains a local database that you can populate with user profiles. You can use a local database instead of AAA servers to provide user authentication, authorization, and accounting.

You can use the local database for the following functions:

- ASDM per-user access
- Console authentication
- Telnet and SSH authentication
- **enable** command authentication

  This setting is for CLI-access only and does not affect the Cisco ASDM login.
- Command authorization

  If you turn on command authorization using the local database, then the ASA refers to the user privilege level to determine which commands are available. Otherwise, the privilege level is not generally used. By default, all commands are either privilege level 0 or level 15. ASDM allows you to enable three predefined privilege levels, with commands assigned to level 15 (Admin), level 5 (Read Only), and level 3 (Monitor Only). If you use the predefined levels, then assign users to one of these three privilege levels.
- Network access authentication
- VPN client authentication

For multiple context mode, you can configure usernames in the system execution space to provide individual logins at the CLI using the **login** command; however, you cannot configure any AAA rules that use the local database in the system execution space.

Note You cannot use the local database for network access authorization.

Fallback Support

The local database can act as a fallback method for several functions. This behavior is designed to help you prevent accidental lockout from the ASA.

When a user logs in, the servers in the group are accessed one at a time, starting with the first server that you specify in the configuration, until a server responds. If all servers in the group are unavailable, the ASA tries the local database if you have configured it as a fallback method (for management authentication and authorization only). If you do not have a fallback method, the ASA continues to try the AAA servers.

For users who need fallback support, we recommend that their usernames and passwords in the local database match their usernames and passwords on the AAA servers. This practice provides transparent fallback support. Because the user cannot determine whether a AAA server or the local database is providing the service, using usernames and passwords on AAA servers that are different than the usernames and passwords in the local database means that the user cannot be certain which username and password should be given.

The local database supports the following fallback functions:

- Console and enable password authentication—If the servers in the group are all unavailable, the ASA uses the local database to authenticate administrative access, which can also include enable password authentication.
- Command authorization—If the TACACS+ servers in the group are all unavailable, the local database is used to authorize commands based on privilege levels.
- VPN authentication and authorization—VPN authentication and authorization are supported to enable remote access to the ASA if AAA servers that normally support these VPN services are unavailable. When a VPN client of an administrator specifies a tunnel group configured to fallback to the local database, the VPN tunnel can be established even if the AAA server group is unavailable, provided that the local database is configured with the necessary attributes.

How Fallback Works with Multiple Servers in a Group

If you configure multiple servers in a server group and you enable fallback to the local database for the server group, fallback occurs when no server in the group responds to the authentication request from the ASA. To illustrate, consider this scenario:

You configure an LDAP server group with two Active Directory servers, server 1 and server 2, in that order. When the remote user logs in, the ASA attempts to authenticate to server 1.

If server 1 responds with an authentication failure (such as user not found), the ASA does not attempt to authenticate to server 2.

If server 1 does not respond within the timeout period (or the number of authentication attempts exceeds the configured maximum), the ASA tries server 2.

If both servers in the group do not respond, and the ASA is configured to fall back to the local database, the ASA tries to authenticate to the local database.

Guidelines for the Local Database

Make sure that you prevent a lockout from the ASA when using the local database for authentication or authorization.

Add a User Account to the Local Database

To add a user to the local database, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Users/AAA** > **User Accounts**, then click **Add**.

The **Add User Account-Identity** dialog box appears.

Step 2 Enter a username from 4 to 64 characters long.

Step 3 (Optional) Enter a password between 8 and 127 characters.

Passwords are case-sensitive. It can be any combination of ASCII printable characters (character codes 32-126), with the following exceptions:

- No spaces
- No question marks
- You cannot use three or more consecutive sequential or repetitive ASCII characters. For example, the following passwords will be rejected:
  - **abc**user1
  - user**543**
  - user**aaaa**
  - user2**666**

The field displays only asterisks. You might want to create a username without a password if you are using SSH public key authentication, for example.

Note
To configure the enable password from the **User Accounts** pane, change the password for the enable_15 user. The enable_15 user is always present in the **User Accounts** pane, and represents the default username. This method of configuring the enable password is the only method available in ASDM for the system configuration. If you configured other enable level passwords at the CLI (enable password 10, for example), then those users are listed as enable_10, and so on.

Step 4 Reenter the password.

For security purposes, only asterisks appear in the password fields.

Step 5 Check the **User authenticated using MSCHAP** check box if you are using MSCHAP for authentication.

Step 6 Set the management access level for a user in the **Access Restriction** area. You must first enable management authorization by clicking the **Perform authorization for exec shell access** option on the **Configuration** > **Device Management** > **Users/AAA** > **AAA Access** > **Authorization** tab.

Choose one of the following options:

- **Full Access (ASDM, Telnet, SSH and console)**—If you configure authentication for management access using the local database, then this option lets the user use ASDM, SSH, Telnet, and the console port. If you also enable authentication, then the user can access global configuration mode.
  - **Privilege Level**—Sets the privilege level for ASDM and local command authorization. The range is 0 (lowest) to 15 (highest). Specify 15 to grant unrestricted admin access. The predefined ASDM roles use 15 for Admin, 5 for Read Only, and 3 for Monitor Only (which restricts the user to the Home and Monitoring panes).
- **CLI login prompt for SSH, Telnet and console (no ASDM access)**—If you configure authentication for management access using the local database, then this option lets the user use SSH, Telnet, and the console port. The user cannot use ASDM for configuration (if you configure HTTP authentication). ASDM monitoring is allowed. If you also configure enable authentication, then the user cannot access global configuration mode.
- **No ASDM, SSH, Telnet, or console access**—If you configure authentication for management access using the local database, then this option disallows the user from accessing any management access method for which you configured authentication (excluding the Serial option; serial access is allowed).

Step 7 (Optional) To enable public key authentication for SSH connections to the ASA on a per-user basis, see Configure HTTPS Access for ASDM, Other Clients, on page 1028.

Step 8 Click **VPN Policy** to configure VPN policy attributes for this user. See the VPN configuration guide.

Step 9 Click **Apply**.

The user is added to the local database, and the changes are saved to the running configuration.

Tip
You can search for specific text in each column of the **Configuration** > **Device Management > Users/AAA > User Accounts** pane. Enter the specific text that you want to locate in the **Find** box, then click the **Up** or **Down** arrow. You can also use the asterisk ("*") and question mark ("?") as wild card characters in the text search.

Test Local Database Authentication and Authorization

To determine whether the ASA can contact a local database and authenticate or authorize a user, perform the following steps:

Procedure

Step 1 From the **Configuration** > **Device Management** > **Users/AAA** > **AAA Server Groups** > **AAA Server Groups** table, click the server group in which the server resides.

Step 2 Click the server that you want to test from the **Servers in the Selected Group** table.

Step 3 Click **Test**.

The **Test AAA Server** dialog box appears for the selected server.

Step 4 Click the type of test that you want to perform—**Authentication** or **Authorization**.

Step 5 Enter a username.

Step 6 If you are testing authentication, enter the password for the username.

Step 7 Click **OK**.

The ASA sends an authentication or authorization test message to the server. If the test fails, ASDM displays an error message.

Monitoring the Local Database

See the following commands for monitoring the local database:

- **Monitoring > Properties > AAA Servers**

  This pane shows AAA server statistics.

- **Tools > Command Line Interface**

  This pane allows you to issue various non-interactive commands and view results.

History for the Local Database

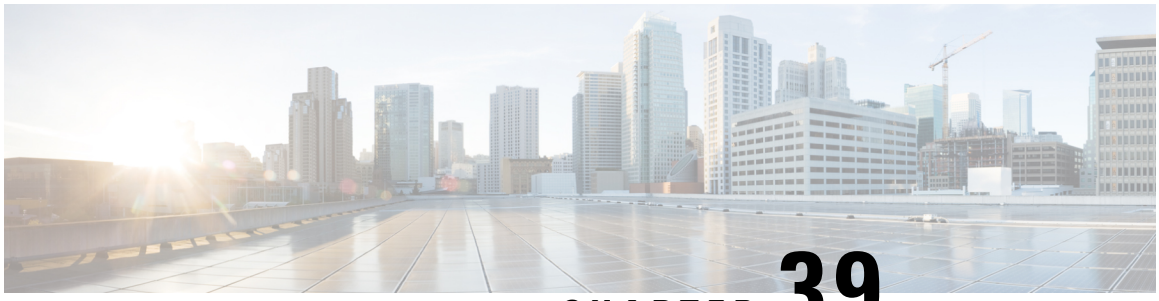
Table 46: History for the Local Database

Feature Name	Platform Releases	Description
Local database configuration for AAA	7.0(1)	Describes how to configure the local database for AAA use. We introduced the following screens: Configuration > Device Management > Users/AAA > AAA Server Groups Configuration > Device Management > Users/AAA > User Accounts.
Support for SSH public key authentication	9.1(2)	You can now enable public key authentication for SSH connections to the ASA on a per-user basis. You can specify a public key file (PKF) formatted key or a Base64 key. The PKF key can be up to 4096 bits. Use PKF format for keys that are too large to for the ASA support of the Base64 format (up to 2048 bits). We introduced the following screens: Configuration > Device Management > Users/AAA > User Accounts > Edit User Account > Public Key Authentication Configuration > Device Management > Users/AAA > User Accounts > Edit User Account > Public Key Using PKF *Also available in 8.4(4.1); PKF key format support is only in 9.1(2).*

Feature Name	Platform Releases	Description
Longer password support for local **username** and **enable** passwords (up to 127 characters)	9.6(1)	You can now create local **username** and **enable** passwords up to 127 characters (the former limit was 32). When you create a password longer than 32 characters, it is stored in the configuration using a PBKDF2 (Password-Based Key Derivation Function 2) hash. Shorter passwords continue to use the MD5-based hashing method. We modified the following screens: **Configuration** > **Device Setup** > **Device Name/Password** > **Enable Password** **Configuration** > **Device Management** > **Users/AAA** > **User Accounts** > **Add/Edit User Account** > **Identity**
SSH public key authentication improvements	9.6(2)	In earlier releases, you could enable SSH public key authentication without also enabling AAA SSH authentication with the Local user database . The configuration is now fixed so that you must explicitly enable AAA SSH authentication. To disallow users from using a password instead of the private key, you can now create a username without any password defined. We modifed the following screens: **Configuration** > **Device Management** > **Management Access** > **ASDM/HTTPS/Telnet/SSH** **Configuration** > **Device Management** > **Users/AAA** > **User Accounts** > **Add/Edit User Account**
PBKDF2 hashing for all local **username** and **enable** passwords	9.7(1)	Local **username** and **enable** passwords of all lengths are stored in the configuration using a PBKDF2 (Password-Based Key Derivation Function 2) hash. Previously, passwords 32 characters and shorter used the MD5-based hashing method. Already existing passwords continue to use the MD5-based hash unless you enter a new password. See the "Software and Configurations" chapter in the General Operations Configuration Guide for downgrading guidelines. We modified the following screens: **Configuration** > **Device Setup** > **Device Name/Password** > **Enable Password** **Configuration** > **Device Management** > **Users/AAA** > **User Accounts** > **Add/Edit User Account** > **Identity**

Feature Name	Platform Releases	Description
Separate authentication for users with SSH public key authentication and users with passwords	9.6(3)/9.8(1)	In releases prior to 9.6(2), you could enable SSH public key authentication (**ssh authentication**) without also explicitly enabling AAA SSH authentication with the Local user database (**aaa authentication ssh console LOCAL**). In 9.6(2), the ASA required you to explicitly enable AAA SSH authentication. In this release, you no longer have to explicitly enable AAA SSH authentication; when you configure the **ssh authentication** command for a user, local authentication is enabled by default for users with this type of authentication. Moreover, when you explicitly configure AAA SSH authentication, this configuration only applies for for usernames with *passwords*, and you can use any AAA server type (**aaa authentication ssh console radius_1**, for example). For example, some users can use public key authentication using the local database, and other users can use passwords with RADIUS. We did not modify any screens.
Stronger local user and enable password requirements	9.17(1)	For local users and the enable password, the following password requirements were added: • Password length—Minimum 8 characters. Formerly, the minimum was 3 characters. • Repetitive and sequential characters—Three or more consecutive sequential or repetitive ASCII characters are disallowed. For example, the following passwords will be rejected: • **abc**user1 • user**543** • user**aaaa** • user2**666** New/Modified screens: • **Configuration** > **Device Management** > **Users/AAA** > **User Accounts** • **Configuration** > **Device Setup** > **Device Name/Password**

Feature Name	Platform Releases	Description
Local user lockout changes	9.17(1)	The ASA can lock out local users after a configurable number of failed login attempts. This feature did not apply to users with privilege level 15. Also, a user would be locked out indefinitely until an admin unlocked their account. Now, users will be unlocked after 10 minutes unless an admin uses the **clear aaa local user lockout** command before then. Privilege level 15 users are also now affected by the lockout setting. New/Modified commands: **aaa local authentication attempts max-fail** , **show aaa local user**
SSH and Telnet password change prompt	9.17(1)	The first time a local user logs into the ASA using SSH or Telnet, they are prompted to change their password. They will also be prompted for the first login after an admin changes their password. If the ASA reloads, however, users will not be prompted even if it is their first login. New/Modified commands: **show aaa local user**



CHAPTER 39

RADIUS Servers for AAA

This chapter describes how to configure RADIUS servers for AAA.

- About RADIUS Servers for AAA, on page 981
- Guidelines for RADIUS Servers for AAA, on page 992
- Configure RADIUS Servers for AAA, on page 992
- Test RADIUS Server Authentication and Authorization, on page 998
- Monitoring RADIUS Servers for AAA, on page 998
- History for RADIUS Servers for AAA, on page 999

About RADIUS Servers for AAA

The ASA supports the following RFC-compliant RADIUS servers for AAA:

- Cisco Secure ACS 3.2, 4.0, 4.1, 4.2, and 5.x
- Cisco Identity Services Engine (ISE)
- RSA RADIUS in RSA Authentication Manager 5.2, 6.1, and 7.x
- Microsoft

Supported Authentication Methods

The ASA supports the following authentication methods with RADIUS servers:

- PAP—For all connection types.
- CHAP and MS-CHAPv1—For L2TP-over-IPsec connections.
- MS-CHAPv2—For L2TP-over-IPsec connections, and for regular IPsec remote access connections when the password management feature is enabled. You can also use MS-CHAPv2 with clientless connections.
- Authentication Proxy modes—For RADIUS-to Active-Directory, RADIUS-to-RSA/SDI, RADIUS-to-Token server, and RSA/SDI-to-RADIUS connections,



Note To enable MS-CHAPv2 as the protocol used between the ASA and the RADIUS server for a VPN connection, password management must be enabled in the tunnel group general attributes. Enabling password management generates an MS-CHAPv2 authentication request from the ASA to the RADIUS server. See the description of the **password-management** command for details.

If you use double authentication and enable password management in the tunnel group, then the primary and secondary authentication requests include MS-CHAPv2 request attributes. If a RADIUS server does not support MS-CHAPv2, then you can configure that server to send a non-MS-CHAPv2 authentication request by using the **no mschapv2-capable** command.

User Authorization of VPN Connections

The ASA can use RADIUS servers for user authorization of VPN remote access and firewall cut-through-proxy sessions using dynamic ACLs or ACL names per user. To implement dynamic ACLs, you must configure the RADIUS server to support them. When the user authenticates, the RADIUS server sends a downloadable ACL or ACL name to the ASA. Access to a given service is either permitted or denied by the ACL. The ASA deletes the ACL when the authentication session expires.

In addition to ACLs, the ASA supports many other attributes for authorization and setting of permissions for VPN remote access and firewall cut-through proxy sessions.

Supported Sets of RADIUS Attributes

The ASA supports the following sets of RADIUS attributes:

- Authentication attributes defined in RFC 2138 and 2865.
- Accounting attributes defined in RFC 2139 and 2866.
- RADIUS attributes for tunneled protocol support, defined in RFC 2868 and 6929.
- Cisco IOS Vendor-Specific Attributes (VSAs), identified by RADIUS vendor ID 9.
- Cisco VPN-related VSAs, identified by RADIUS vendor ID 3076.
- Microsoft VSAs, defined in RFC 2548.

Supported RADIUS Authorization Attributes

Authorization refers to the process of enforcing permissions or attributes. A RADIUS server defined as an authentication server enforces permissions or attributes if they are configured. These attributes have vendor ID 3076.

The following table lists the supported RADIUS attributes that can be used for user authorization.

Note RADIUS attribute names do not contain the cVPN3000 prefix. Cisco Secure ACS 4.x supports this new nomenclature, but attribute names in pre-4.0 ACS releases still include the cVPN3000 prefix. The ASAs enforce the RADIUS attributes based on attribute numeric ID, not attribute name.

All attributes listed in the following table are downstream attributes that are sent from the RADIUS server to the ASA except for the following attribute numbers: 146, 150, 151, and 152. These attribute numbers are upstream attributes that are sent from the ASA to the RADIUS server. RADIUS attributes 146 and 150 are sent from the ASA to the RADIUS server for authentication and authorization requests. All four previously listed attributes are sent from the ASA to the RADIUS server for accounting start, interim-update, and stop requests. Upstream RADIUS attributes 146, 150, 151, and 152 were introduced in Version 8.4(3).

Table 47: Supported RADIUS Authorization Attributes

Attribute Name	ASA	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	ACL ID
Access-List-Outbound	Y	87	String	Single	ACL ID
Address-Pools	Y	217	String	Single	Name of IP local pool
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, GN, SN, I, GENQ, DNQ, SER, use-entire-na
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 remo sessions: IPsec IKEv1, Secure Client SSL-TLS/DTLS/IKEv2, and Clientless SSL
Banner2	Y	36	String	Single	Banner string to display for Cisco VPN remo sessions: IPsec IKEv1, Secure Client SSL-TLS/DTLS/IKEv2, and Clientless SSL. T string is concatenated to the Banner1 string , if
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 C VPN 3 = Clientless SSL VPN 4 = Cut-Throu = L2TP/IPsec SSL VPN 6 = Secure Client II (IKEv2)

Attribute Name	ASA	Attr. No.	Syntax/Type	Single or Multi- Valued	Description or Value
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 cor 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
Group-Policy	Y	25	String	Single	Sets the group policy for the remote access VPN For Versions 8.2.x and later, use this attribute in IETF-Radius-Class. You can use one of the follo 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 detect 4 = 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

Attribute Name	ASA	Attr. No.	Syntax/Type	Single or Multi- Valued	Description or Value
IPsec-Authentication		13	Integer	Single	0 = None 1 = RADIUS 2 = LDAP (authoriza = NT Domain 4 = SDI 5 = Internal 6 = RAD Expiry 7 = 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
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 CP 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 name the client (1-255 characters).
IPsec-Split-Tunneling-Policy	Y	55	Integer	Single	0 = No split tunneling 1 = Split tunneling 2 = permitted
IPsec-Split-Tunnel-List	Y	27	String	Single	Specifies the name of the network or ACL th 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

Attribute Name	ASA	Attr. No.	Syntax/Type	Single or Multi- Valued	Description or 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
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 4 = 8 = Stateless-Required 15= 40/128-Encr/Stateles
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 Client Zone Labs 3 = NetworkICE 4 = Sygate 5 = Cisco (with Cisco Intrusion Prevention Security Agent
Required-Client-Firewall-Description	Y	47	String	Single	String

Attribute Name	ASA	Attr. No.	Syntax/Type	Single or Multi- Valued	Description or Value
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
Session Type	Y	151	Integer	Single	0 = None 1 = Secure Client SSL VPN 2 = Se IPSec VPN (IKEv2) 3 = Clientless SSL VPN Clientless Email Proxy 5 = Cisco VPN Clien = IKEv1 LAN-LAN 7 = IKEv2 LAN-LAN 8 = 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 ap 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 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

Attribute Name	ASA	Attr. No.	Syntax/Type	Single or Multi- Valued	Description or Value
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
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

Attribute Name	ASA	Attr. No.	Syntax/Type	Single or Multi- Valued	Description or Value
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= o 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
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, “Applicati 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 ap the domain name
WebVPN-Smart-Tunnel-Auto-Start	Y	138	Integer	Single	0 = Disabled 1 = Enabled 2 = Auto Start

Attribute Name	ASA	Attr. No.	Syntax/Type	Single or Multi- Valued	Description or Value
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
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

Supported IETF RADIUS Authorization Attributes

The following table lists the supported IETF RADIUS attributes.

Table 48: Supported IETF RADIUS Attributes

Attribute Name	ASA	Attr. No.	Syntax/Type	Single or Multi-Valued	Description or Value
IETF-Radius-Class	Y	25		Single	For Versions 8.2.x and later, we recommend that you use the Group-Policy attribute (VSA 3076, #25): • *group policy name* • OU=*group policy name* • OU=*group policy name*
IETF-Radius-Filter-Id	Y	11	String	Single	ACL name that is defined on the ASA, which applie only to full tunnel IPsec and SSL VPN clients.
IETF-Radius-Framed-IP-Address	Y	n/a	String	Single	An IP address
IETF-Radius-Framed-IP-Netmask	Y	n/a	String	Single	An IP address mask
IETF-Radius-Idle-Timeout	Y	28	Integer	Single	Seconds
IETF-Radius-Service-Type	Y	6	Integer	Single	Seconds. Possible Service Type values: • .Administrative—User is allowed access to the configure prompt. • .NAS-Prompt—User is allowed access to the ex prompt. • .remote-access—User is allowed network acces
IETF-Radius-Session-Timeout	Y	27	Integer	Single	Seconds

RADIUS Accounting Disconnect Reason Codes

These codes are returned if the ASA encounters a disconnect when sending packets:

Disconnect Reason Code
ACCT_DISC_USER_REQ = 1
ACCT_DISC_LOST_CARRIER = 2
ACCT_DISC_LOST_SERVICE = 3
ACCT_DISC_IDLE_TIMEOUT = 4
ACCT_DISC_SESS_TIMEOUT = 5
ACCT_DISC_ADMIN_RESET = 6
ACCT_DISC_ADMIN_REBOOT = 7

Disconnect Reason Code
ACCT_DISC_PORT_ERROR = 8
ACCT_DISC_NAS_ERROR = 9
ACCT_DISC_NAS_REQUEST = 10
ACCT_DISC_NAS_REBOOT = 11
ACCT_DISC_PORT_UNNEEDED = 12
ACCT_DISC_PORT_PREEMPTED = 13
ACCT_DISC_PORT_SUSPENDED = 14
ACCT_DISC_SERV_UNAVAIL = 15
ACCT_DISC_CALLBACK = 16
ACCT_DISC_USER_ERROR = 17
ACCT_DISC_HOST_REQUEST = 18
ACCT_DISC_ADMIN_SHUTDOWN = 19
ACCT_DISC_SA_EXPIRED = 21
ACCT_DISC_MAX_REASONS = 22

Guidelines for RADIUS Servers for AAA

This section describes the guidelines and limitations that you should check before configuring RADIUS servers for AAA.

- You can have up to 200 server groups in single mode or 4 server groups per context in multiple mode.
- Each group can have up to 16 servers in single mode or 8 servers in multiple mode.
- The maximum length of the RADIUS payload is 4096 bytes.

Configure RADIUS Servers for AAA

This section describes how to configure RADIUS servers for AAA.

Procedure

Step 1 Load the ASA attributes into the RADIUS server. The method that you use to load the attributes depends on which type of RADIUS server that you are using:

- If you are using Cisco ACS: the server already has these attributes integrated. You can skip this step.
- For RADIUS servers from other vendors (for example, Microsoft Internet Authentication Service): you must manually define each ASA attribute. To define an attribute, use the attribute name or number, type, value, and vendor code (3076).

Step 2 Configure RADIUS Server Groups, on page 993.

Step 3 Add a RADIUS Server to a Group, on page 995.

Step 4 (Optional) Add an Authentication Prompt, on page 997.

Configure RADIUS Server Groups

If you want to use an external RADIUS server for authentication, authorization, or accounting, you must first create at least one RADIUS server group per AAA protocol and add one or more servers to each group.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Users/AAA** > **AAA Server Groups**.

Step 2 Click **Add** in the **AAA Server Groups** area.

The **Add AAA Server Group** dialog box appears.

Step 3 Enter a name for the group in the **AAA Server Group** field.

Step 4 Choose the RADIUS server type from the **Protocol** drop-down list.

Step 5 Select the **Accounting Mode**.

- **Simultaneous**—Send accounting data to all servers in the group.
- **Single**—Send accounting data to only one server.

Step 6 Configure the method (**Reactivation Mode**) by which failed servers in a group are reactivated.

- **Depletion**, **Dead Time**—Reactivate failed servers only after all of the servers in the group are inactive. This is the default reactivation mode. Specify the amount of time, between 0 and 1440 minutes, that elapses between the disabling of the last server in the group and the subsequent reenabling of all servers. Deadtime applies only if you configure fallback to the local database; authentication is attempted locally until the deadtime elapses. The default is 10 minutes.
- **Timed**—Reactivate failed servers after 30 seconds of down time.

Step 7 In **Max Failed Attempts**, specify the maximum number of failed AAA transactions with a RADIUS server in the group before trying the next server.

The range is from 1 and 5. The default is 3.

If you configure a fallback method using the local database (for management access only), and all the servers in the group fail to respond, or their responses are invalid, then the group is considered to be unresponsive, and the fallback method is tried. The server group remains marked as unresponsive for a period of 10 minutes (if you use the default reactivation mode and dead time), so that additional AAA requests within that period

do not attempt to contact the server group, and the fallback method is used immediately. To change the unresponsive period from the default, see change the **Dead Time**.

If you do not have a fallback method, the ASA continues to retry the servers in the group.

Step 8 (Optional.) Enable the periodic generation of RADIUS interim-accounting-update messages by selecting the desired options.

These options are relevant only if you are using this server group for Secure Client or clientless SSL VPN.

- **Enable interim accounting update**—If you use this command without selecting the **Update Interval** option, the ASA sends interim-accounting-update messages only when a VPN tunnel connection is added to a clientless VPN session. When this happens the accounting update is generated in order to inform the RADIUS server of the newly assigned IP address.
- **Update Interval**—Enables the periodic generation and transmission of accounting records for every VPN session that is configured to send accounting records to the server group in question. You can change the interval, in hours, for sending these updates. The default is 24 hours, the range is 1 to 120.

Note
For server groups containing ISE servers, select both options. ISE maintains a directory of active sessions based on the accounting records that it receives from NAS devices like the ASA. However, if ISE does not receive any indication that the session is still active (accounting message or posture transactions) for a period of 5 days, it will remove the session record from its database. To ensure that long-lived VPN connections are not removed, configure the group to send periodic interim-accounting-update messages to ISE for all active sessions.

Step 9 (Optional.) If this group contains AD Agents or Cisco Directory Agent (CDA) servers only, select **Enable Active Directory Agent Mode**.

CDA or AD Agents are used in identity firewall, and are not full-featured RADIUS servers. If you select this option, you can use this group for identity firewall purposes only.

Step 10 (Optional) If you are using this server group for ISE Policy Enforcement in remote access VPN, configure the following options:

- **Enable dynamic authorization**—Enable the RADIUS Dynamic Authorization (ISE Change of Authorization, CoA) services for the AAA server group. When you use the server group in a VPN tunnel, the RADIUS server group will be registered for CoA notification and the ASA will listen to the port for the CoA policy updates from ISE. Enable dynamic authorization only if you are using this server group in a remote access VPN in conjunction with ISE.
- **Dynamic Authorization Port**—If you enable dynamic authorization, you can specify the listening port for RADIUS CoA requests. The default is 1700. The valid range is 1024 to 65535.
- **Use authorization only mode**—If you do not want to use ISE for authentication, enable authorize-only mode for the RADIUS server group. This indicates that when this server group is used for authorization, the RADIUS Access Request message will be built as an “Authorize Only” request as opposed to the configured password methods defined for the AAA server. If you do configure a common password for the RADIUS server, it will be ignored.

  For example, you would use authorize-only mode if you want to use certificates for authentication rather than this server group. You would still use this server group for authorization and accounting in the VPN tunnel.

Step 11 (Optional.) Configure the **VPN3K Compatibility Option** to specify whether or not a downloadable ACL received from a RADIUS packet should be merged with a Cisco AV pair ACL.

This option applies only to VPN connections. For VPN users, ACLs can be in the form of Cisco AV pair ACLs, downloadable ACLs, and an ACL that is configured on the ASA. This option determines whether or not the downloadable ACL and the AV pair ACL are merged, and does not apply to any ACLs configured on the ASA.

- **Do not merge** —Downloadable ACLs will not be merged with Cisco AV pair ACLs. If both an AV pair and a downloadable ACL are received, the AV pair has priority and is used. This is the default option.
- **Place the downloadable ACL after Cisco AV-pair ACL**
- **Place the downloadable ACL before Cisco AV-pair ACL**

Step 12 Click **OK**.

The **Add AAA Server Group** dialog box closes, and the new server group is added to the **AAA Server Groups** table.

Step 13 Click **Apply** to save the changes to the running configuration.

Add a RADIUS Server to a Group

To add a RADIUS server to a group, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Users/AAA** > **AAA Server Groups**, and in the **AAA Server Groups** area, click the server group to which you want to add a server.

Step 2 Click **Add** in the **Servers in the Selected Group** area (lower pane).

The **Add AAA Server Group** dialog box appears for the server group.

Step 3 Choose the interface name on which the authentication server resides.

Step 4 Add either a server name or IP address for the server that you are adding to the group.

Step 5 Specify the timeout value for connection attempts to the server.

Specify the timeout interval (1-300 seconds) for the server; the default is 10 seconds. For each AAA transaction the ASA retries connection attempts (based on the retry interval) until the timeout is reached. If the number of consecutive failed transactions reaches the maximum-failed-attempts limit specified in the AAA server group, the AAA server is deactivated and the ASA starts sending requests to another AAA server if it is configured.

Step 6 Select **Require message authenticator from RADIUS Server** to specify that the Message-Authenticator attribute must be present in RADIUS requests and responses.

The Message-Authenticator attribute is used to protect against Blast-RADIUS attacks. If you have upgraded your RADIUS server so it supports the message authenticator, you can enable this option to help protect against these attacks. When enabled, all requests and responses must have the message authenticator, or authentication will fail.

Step 7 Specify how you want the ASA to handle netmasks received in downloadable ACLs. Choose from the following options:

- **Detect automatically**—The ASA attempts to determine the type of netmask expression used. If the ASA detects a wildcard netmask expression, the ASA converts it to a standard netmask expression.

  Note
  Because some wildcard expressions are difficult to detect clearly, this setting may misinterpret a wildcard netmask expression as a standard netmask expression.

- **Standard**—The ASA assumes downloadable ACLs received from the RADIUS server contain only standard netmask expressions. No translation from wildcard netmask expressions is performed.
- **Wildcard**—The ASA assumes downloadable ACLs received from the RADIUS server contain only wildcard netmask expressions, and it converts them all to standard netmask expressions when the ACLs are downloaded.

Step 8 Specify a case-sensitive password that is common among users who access this RADIUS authorization server through this ASA. Be sure to provide this information to your RADIUS server administrator.

Note
For an authentication RADIUS server (rather than authorization), do not configure a common password.

If you leave this field blank, the username is the password for accessing this RADIUS authorization server.

Never use a RADIUS authorization server for authentication. Common passwords or usernames as passwords are less secure than assigning unique user passwords.

Although the password is required by the RADIUS protocol and the RADIUS server, users do not need to know it.

Step 9 If you use double authentication and enable password management in the tunnel group, then the primary and secondary authentication requests include MS-CHAPv2 request attributes. If a RADIUS server does not support MS-CHAPv2, then you can configure that server to send a non-MS-CHAPv2 authentication request by unchecking this check box.

Step 10 Specify the length of time, from 1 to 10 seconds, that the ASA waits between attempts to contact the server.

Note
For the RADIUS protocol, if the server responds with an ICMP Port Unreachable message, the retry-interval setting is ignored and the AAA server is immediately moved to the failed state. If this is the only server in the AAA group, it is reactivated and another request is sent to it. This is the intended behavior.

Step 11 Click **Simultaneous** or **Single**.

In Single mode, the ASA sends accounting data to only one server.

In Simultaneous mode, the ASA sends accounting data to all servers in the group.

Step 12 Specify the server port to be used for accounting of users. The default port is 1646.

Step 13 Specify the server port to be used for authentication of users. The default port is 1645.

Step 14 Specify the shared secret key used to authenticate the RADIUS server to the ASA. The server secret that you configure should match the one configured on the RADIUS server. If you do not know the server secret, ask the RADIUS server administrator. The maximum field length is 64 characters.

Step 15 Click **OK**.

The **Add AAA Server Group** dialog box closes, and the AAA server is added to the AAA server group.

Step 16 In the **AAA Server Groups** pane, click **Apply** to save the changes to the running configuration.

Add an Authentication Prompt

You can specify the AAA challenge text for HTTP, FTP, and Telnet access through the ASA when requiring user authentication from RADIUS servers. This text is primarily for cosmetic purposes and appears above the username and password prompts that users see when they log in. If you do not specify an authentication prompt, users see the following when authenticating with a RADIUS server:

Connection Type	Default Prompt
FTP	FTP authentication
HTTP	HTTP authentication
Telnet	None

To add an authentication prompt, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Users/AAA** > **Authentication Prompt**.

Step 2 Enter text in the **Prompt** field to add as a message to appear above the username and password prompts that users see when they log in.

The following table shows the allowed character limits for authentication prompts:

Application	Character Limit
Microsoft Internet Explorer	37
Telnet	235
FTP	235

Step 3 Add messages in the **User accepted message** and **User rejected message** fields.

If the user authentication occurs from Telnet, you can use the **User accepted message** and **User rejected message** options to display different status prompts to indicate that the authentication attempt is either accepted or rejected by the RADIUS server.

If the RADIUS server authenticates the user, the ASA displays the **User accepted message** text, if specified, to the user; otherwise, the ASA displays the **User rejected message** text, if specified. Authentication of HTTP and FTP sessions displays only the challenge text at the prompt. The **User accepted message** and **User rejected message** text are not displayed.

Step 4 Click **Apply** to save the changes to the running configuration.

Test RADIUS Server Authentication and Authorization

To determine whether the ASA can contact a RADIUS server and authenticate or authorize a user, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Users/AAA** > **AAA Server Groups**.

Step 2 Click the server group in which the server resides in the **AAA Server Groups** table.

Step 3 Click the server that you want to test in the **Servers in the Selected Group** table.

Step 4 Click **Test**.

The **Test AAA Server** dialog box appears for the selected server.

Step 5 Click the type of test that you want to perform—**Authentication** or **Authorization**.

Step 6 Enter a username.

Step 7 Enter the password for the username if you are testing authentication.

Step 8 Click **OK**.

The ASA sends an authentication or authorization test message to the server. If the test fails, an error message appears.

Monitoring RADIUS Servers for AAA

See the following commands for monitoring the status of RADIUS servers for AAA:

- **Monitoring > Properties > AAA Servers**

  This pane shows the RADIUS server running configuration.

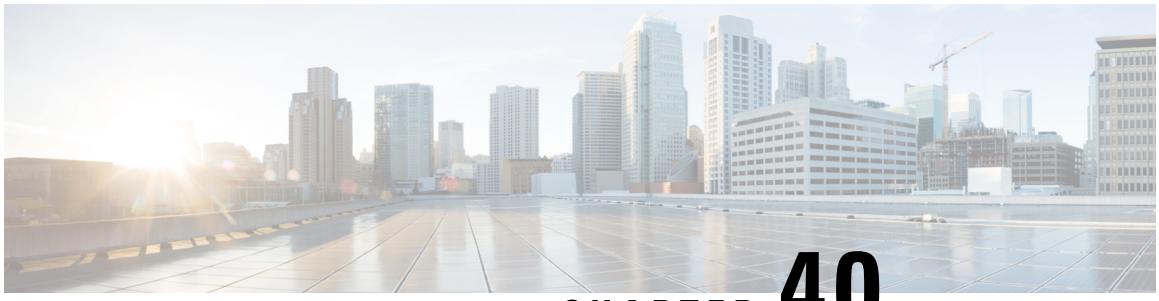
- **Tools > Command Line Interface**

  This pane allows you to issue various non-interactive commands and view results.

History for RADIUS Servers for AAA

Table 49: History for RADIUS Servers for AAA

Feature Name	Platform Releases	Description
RADIUS Servers for AAA	7.0(1)	Describes how to configure RADIUS servers for AAA. We introduced the following screens: Configuration > Device Management > Users/AAA > AAA Server Groups Configuration > Device Management > Users/AAA > Authentication Prompt.
Key vendor-specific attributes (VSAs) sent in RADIUS access request and accounting request packets from the ASA	8.4(3)	Four New VSAs—Tunnel Group Name (146) and Client Type (150) are sent in RADIUS access request packets from the ASA. Session Type (151) and Session Subtype (152) are sent in RADIUS accounting request packets from the ASA. All four attributes are sent for all accounting request packet types: Start, Interim-Update, and Stop. The RADIUS server (for example, ACS and ISE) can then enforce authorization and policy attributes or use them for accounting and billing purposes.
Increased limits for AAA server groups and servers per group.	9.13(1)	You can configure more AAA server groups. In single context mode, you can configure 200 AAA server groups (the former limit was 100). In multiple context mode, you can configure 8 (the former limit was 4). In addition, in multiple context mode, you can configure 8 servers per group (the former limit was 4 servers per group). The single context mode per-group limit of 16 remains unchanged. We modified the AAA screens to accept these new limits.
Support for the RADIUS Message-Authenticator attribute.	9.23(1)	The Message-Authenticator attribute is used to protect against Blast-RADIUS attacks. If you have upgraded your RADIUS server so it supports the message authenticator, you can enable this option to help protect against these attacks. When enabled, all requests and responses must have the message authenticator, or authentication will fail. We added the **Require message authenticator from RADIUS Server** option to the Add/Edit RADIUS Server dialog box.



CHAPTER 40

TACACS+ Servers for AAA

This chapter describes how to configure TACACS+ servers used in AAA.

- About TACACS+ Servers for AAA, on page 1001
- Guidelines for TACACS+ Servers for AAA, on page 1002
- Configure TACACS+ Servers, on page 1003
- Test TACACS+ Server Authentication and Authorization, on page 1006
- Monitoring TACACS+ Servers for AAA, on page 1006
- History for TACACS+ Servers for AAA, on page 1007

About TACACS+ Servers for AAA

The ASA supports TACACS+ server authentication with the following protocols: ASCII, PAP, CHAP, and MS-CHAPv1.

TACACS+ Attributes

The ASA provides support for TACACS+ attributes. TACACS+ attributes separate the functions of authentication, authorization, and accounting. The protocol supports two types of attributes: mandatory and optional. Both the server and client must understand a mandatory attribute, and the mandatory attribute must be applied to the user. An optional attribute may or may not be understood or used.

Note To use TACACS+ attributes, make sure that you have enabled AAA services on the NAS.

The following table lists supported TACACS+ authorization response attributes for cut-through-proxy connections.

Table 50: Supported TACACS+ Authorization Response Attributes

Attribute	Description
acl	Identifies a locally configured ACL to be applied to the connection.
idletime	Indicates the amount of inactivity in minutes that is allowed before the authenticated user session is terminated.

Attribute	Description
timeout	Specifies the absolute amount of time in minutes that authentication credentials remain active before the authenticated user session is terminated.

The following table lists supported TACACS+ accounting attributes.

.

Table 51: Supported TACACS+ Accounting Attributes

Attribute	Description
bytes_in	Specifies the number of input bytes transferred during this connection (stop records only).
bytes_out	Specifies the number of output bytes transferred during this connection (stop records only).
cmd	Defines the command executed (command accounting only).
disc-cause	Indicates the numeric code that identifies the reason for disconnecting (stop records only).
elapsed_time	Defines the elapsed time in seconds for the connection (stop records only).
foreign_ip	Specifies the IP address of the client for tunnel connections. Defines the address on the lowest security interface for cut-through-proxy connections.
local_ip	Specifies the IP address that the client connected to for tunnel connections. Defines the address on the highest security interface for cut-through-proxy connections.
NAS port	Contains a session ID for the connection.
packs_in	Specifies the number of input packets transferred during this connection.
packs_out	Specifies the number of output packets transferred during this connection.
priv-level	Set to the user privilege level for command accounting requests or to 1 otherwise.
rem_iddr	Indicates the IP address of the client.
service	Specifies the service used. Always set to “shell” for command accounting only.
task_id	Specifies a unique task ID for the accounting transaction.
username	Indicates the name of the user.

Guidelines for TACACS+ Servers for AAA

This section describes the guidelines and limitation that you should check before configuring TACACS+ servers for AAA.

IPv6

The AAA server can use either an IPv4 or IPv6 address.

Additional Guidelines

- You can have up to 200 server groups in single mode or 4 server groups per context in multiple mode.
- Each group can have up to 16 servers in single mode or 8 servers in multiple mode.
- For FPR1000, FPR2100, or FPR3100 Series that are running in ASA appliance mode, you must comply with these username conventions:
  - Must be Linux-valid usernames.
  - Must be lower-case only.
  - May include alphanumeric characters, period (.), or hyphen (-).
  - Must not include other special characters such as at sign (@) and slash (/).

Configure TACACS+ Servers

This section describes how to configure TACACS+ servers.

Procedure

Step 1 Configure TACACS+ Server Groups, on page 1003.
Step 2 Add a TACACS+ Server to a Group, on page 1004.
Step 3 (Optional) Add an Authentication Prompt, on page 1005.

Configure TACACS+ Server Groups

If you want to use a TACACS+ server for authentication, authorization, or accounting, you must first create at least one TACACS+ server group and add one or more servers to each group. You identify TACACS+ server groups by name.

To add a TACACS+ server group, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Users/AAA** > **AAA Server Groups**.
Step 2 Click **Add** in the **AAA Server Groups** area.

The **Add AAA Server Group** dialog box appears.

Step 3 Enter a name for the group in the **Server Group** field.

Step 4 Choose the **TACACS**+ server type from the **Protocol** drop-down list:

Step 5 Click **Simultaneous** or **Single** in the **Accounting Mode** field.

In Single mode, the ASA sends accounting data to only one server.

In Simultaneous mode, the ASA sends accounting data to all servers in the group.

Step 6 Click **Depletion** or **Timed** in the **Reactivation Mode** field.

In Depletion mode, failed servers are reactivated only after all of the servers in the group are inactive. In depletion mode, when a server is deactivated, it remains inactive until all other servers in the group are inactive. When and if this occurs, all servers in the group are reactivated. This approach minimizes the occurrence of connection delays due to failed servers.

In Timed mode, failed servers are reactivated after 30 seconds of down time.

Step 7 If you chose the Depletion reactivation mode, enter a time interval in the **Dead Time** field.

The dead time is the duration of time, in minutes, that elapses between the disabling of the last server in a group and the subsequent re-enabling of all servers. Deadtime applies only if you configure fallback to the local database; authentication is attempted locally until the deadtime elapses.

Step 8 Add the maximum number of failed AAA transactions with a server to allow.

This option sets the number of failed AAA transactions before declaring a nonresponsive server to be inactive.

Step 9 Click **OK**.

The **Add AAA Server Group** dialog box closes, and the new server group is added to the **AAA Server Groups** table.

Step 10 Click **Apply** to save the changes to the running configuration.

Add a TACACS+ Server to a Group

To add a TACACS+ server to a group, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Users/AAA** > **AAA Server Groups**.

Step 2 Click the server group to which you want to add a server.

Step 3 Click **Add** in the **Servers in the Selected Group** area.

The **Add AAA Server Group** dialog box appears for the server group.

Step 4 Choose the interface name on which the authentication server resides.

Step 5 Add either a server name or IP address for the server that you are adding to the group.

Step 6 Specify the timeout value for connection attempts to the server.

Specify the timeout interval (1-300 seconds) for the server; the default is 10 seconds. For each AAA transaction the ASA retries connection attempts (based on the retry interval) until the timeout is reached. If the number

of consecutive failed transactions reaches the maximum-failed-attempts limit specified in the AAA server group, the AAA server is deactivated and the ASA starts sending requests to another AAA server if it is configured.

Step 7 Specify the server port. The server port is either port number 139, or the TCP port number used by the ASA to communicate with the TACACS+ server.

Step 8 Specify the server secret key. The shared secret key used to authenticate the TACACS+ server to the ASA. The server secret that you configure here should match the one that is configured on the TACACS+ server. If you do not know the server secret, ask the TACACS+ server administrator. The maximum field length is 64 characters.

Step 9 Click **OK**.

The **Add AAA Server Group** dialog box closes, and the AAA server is added to the AAA server group.

Step 10 Click **Apply** to save the changes to the running configuration.

Add an Authentication Prompt

You can specify text to display to the user during the AAA authentication challenge process. You can specify the AAA challenge text for HTTP, FTP, and Telnet access through the ASA when requiring user authentication from TACACS+ servers. This text is primarily for cosmetic purposes and appears above the username and password prompts that users see when they log in.

If you do not specify an authentication prompt, users see the following when authenticating with a TACACS+ server:

Connection Type	Default Prompt
FTP	FTP authentication
HTTP	HTTP authentication
Telnet	None

To add an authentication prompt, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Users/AAA** > **Authentication Prompt**.

Step 2 Add text to appear above the username and password prompts that users see when they log in.

The following table shows the allowed character limits for authentication prompts:

Application	Character Limit for Authentication Prompt
Microsoft Internet Explorer	37
Telnet	235

Application	Character Limit for Authentication Prompt
FTP	235

Step 3 Add messages in the **User accepted message** and **User rejected message** fields.

If the user authentication occurs from Telnet, you can use the **User accepted message** and **User rejected message** options to display different status prompts to indicate that the authentication attempt is accepted or rejected by the AAA server.

If the AAA server authenticates the user, the ASA displays the **User accepted message** text, if specified, to the user; otherwise, the ASA displays the **User rejected message** text, if specified. Authentication of HTTP and FTP sessions displays only the challenge text at the prompt. The User accepted message and User rejected message text are not displayed.

Step 4 Click **Apply** to save the changes to the running configuration.

Test TACACS+ Server Authentication and Authorization

To determine whether the ASA can contact a TACACS+ server and authenticate or authorize a user, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Users/AAA** > **AAA Server Groups**.

Step 2 Click the server group in which the server resides.

Step 3 Click the server that you want to test.

Step 4 Click **Test**.

The **Test AAA Server** dialog box appears for the selected server.

Step 5 Click the type of test that you want to perform—**Authentication** or **Authorization**.

Step 6 Enter a username.

Step 7 If you are testing authentication, enter the password for the username.

Step 8 Click **OK**.

The ASA sends an authentication or authorization test message to the server. If the test fails, an error message appears.

Monitoring TACACS+ Servers for AAA

See the following commands for monitoring TACACS+ servers for AAA:

- **Monitoring > Properties > AAA Servers**

This pane shows the configured TACACS+ server statistics.

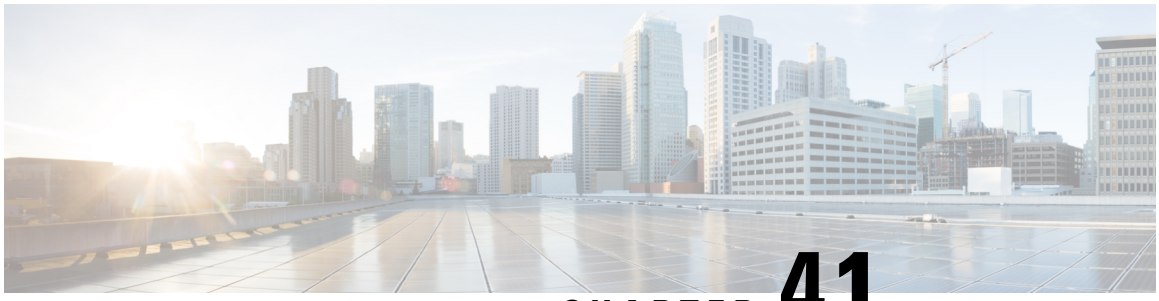
- **Tools > Command Line Interface**

  This pane allows you to issue various non-interactive commands and view results.

History for TACACS+ Servers for AAA

Table 52: History for TACACS+ Servers for AAA

Feature Name	Platform Releases	Description
TACACS+ Servers	7.0(1)	Describes how to configure TACACS+ servers for AAA. We introduced the following screens: Configuration > Device Management > Users/AAA > AAA Server Groups Configuration > Device Management > Users/AAA > Authentication Prompt.
TACACS+ servers with IPv6 addresses for AAA	9.7(1)	You can now use either an IPv4 or IPv6 address for the AAA server.
Increased limits for AAA server groups and servers per group.	9.13(1)	You can configure more AAA server groups. In single context mode, you can configure 200 AAA server groups (the former limit was 100). In multiple context mode, you can configure 8 (the former limit was 4). In addition, in multiple context mode, you can configure 8 servers per group (the former limit was 4 servers per group). The single context mode per-group limit of 16 remains unchanged. We modified the AAA screens to accept these new limits.



CHAPTER 41

LDAP Servers for AAA

This chapter describes how to configure LDAP servers used in AAA.

- About LDAP and the ASA, on page 1009
- Guidelines for LDAP Servers for AAA, on page 1012
- Configure LDAP Servers for AAA, on page 1013
- Test LDAP Server Authentication and Authorization, on page 1017
- Monitoring LDAP Servers for AAA, on page 1017
- History for LDAP Servers for AAA, on page 1018

About LDAP and the ASA

The ASA is compatible with the most LDAPv3 directory servers, including:

- Sun Microsystems JAVA System Directory Server, now part of Oracle Directory Server Enterprise Edition, and formerly named the Sun ONE Directory Server
- Microsoft Active Directory
- Novell
- OpenLDAP

By default, the ASA autodetects whether it is connected to Microsoft Active Directory, Sun LDAP, Novell, OpenLDAP, or a generic LDAPv3 directory server. However, if autodetection fails to determine the LDAP server type, you can manually configure it.

How Authentication Works with LDAP

During authentication, the ASA acts as a client proxy to the LDAP server for the user, and authenticates to the LDAP server in either plain text or by using the SASL protocol. By default, the ASA passes authentication parameters, usually a username and password, to the LDAP server in plain text.

The ASA supports only Digest-MD5 SASL mechanism. Using this mechanism, ASA responds to LDAP server with an MD5 value computed from the username and password.

When user LDAP authentication has succeeded, the LDAP server returns the attributes for the authenticated user. For VPN authentication, these attributes generally include authorization data that is applied to the VPN session. In this case, using LDAP accomplishes authentication and authorization in a single step.

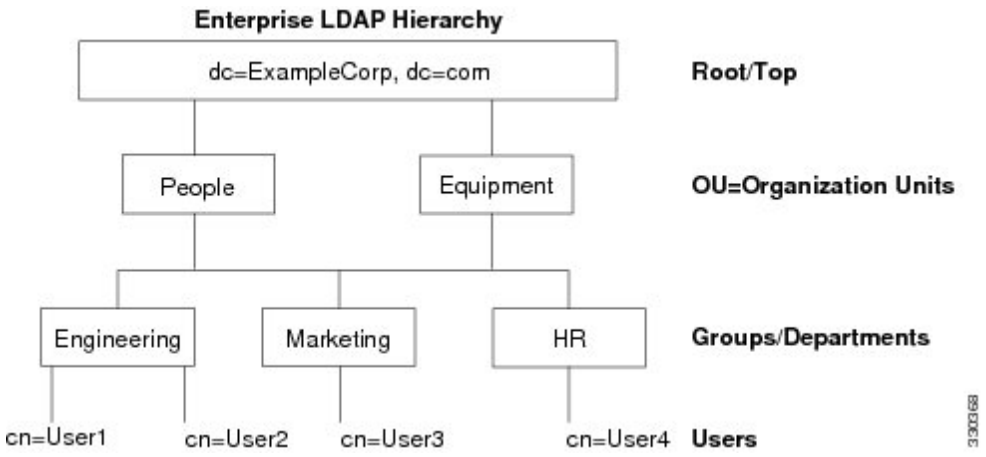
Note For more information about LDAP, see RFCs 1777, 2251, and 2849.

LDAP Hierarchy

Your LDAP configuration should reflect the logical hierarchy of your organization. For example, suppose an employee at your company, Example Corporation, is named Employee1. Employee1 works in the Engineering group. Your LDAP hierarchy could have one or many levels. You might decide to set up a single-level hierarchy in which Employee1 is considered a member of Example Corporation. Or you could set up a multi-level hierarchy in which Employee1 is considered to be a member of the department Engineering, which is a member of an organizational unit called People, which is itself a member of Example Corporation. See the following figure for an example of a multi-level hierarchy.

A multi-level hierarchy has more detail, but searches return results more quickly in a single-level hierarchy.

Figure 96: A Multi-Level LDAP Hierarchy



Search the LDAP Hierarchy

The ASA lets you tailor the search within the LDAP hierarchy. You configure the following three fields on the ASA to define where in the LDAP hierarchy that your search begins, the extent, and the type of information you are looking for. Together, these fields limit the search of the hierarchy to only the part that includes the user permissions.

- LDAP Base DN defines where in the LDAP hierarchy that the server should begin searching for user information when it receives an authorization request from the ASA.
- Search Scope defines the extent of the search in the LDAP hierarchy. The search proceeds this many levels in the hierarchy below the LDAP Base DN. You can choose to have the server search only the level immediately below it, or it can search the entire subtree. A single level search is quicker, but a subtree search is more extensive.
- Naming Attribute(s) defines the RDN that uniquely identifies an entry in the LDAP server. Common naming attributes can include cn (Common Name), sAMAccountName, and userPrincipalName.

The figure shows a sample LDAP hierarchy for Example Corporation. Given this hierarchy, you could define your search in different ways. The following table shows two sample search configurations.

In the first example configuration, when Employee1 establishes the IPsec tunnel with LDAP authorization required, the ASA sends a search request to the LDAP server, indicating it should search for Employee1 in the Engineering group. This search is quick.

In the second example configuration, the ASA sends a search request indicating that the server should search for Employee1 within Example Corporation. This search takes longer.

Table 53: Example Search Configurations

No.	LDAP Base DN	Search Scope	Naming Attribute	Result
1	group= Engineering,ou=People,dc=ExampleCorporation, dc=com	One Level	cn=Employee1	Quicker search
2	dc=ExampleCorporation,dc=com	Subtree	cn=Employee1	Longer search

Bind to an LDAP Server

The ASA uses the login DN and login password to establish trust (bind) with an LDAP server. When performing a Microsoft Active Directory read-only operation (such as authentication, authorization, or group search), the ASA can bind using a login DN with fewer privileges. For example, the login DN can be a user whose AD "Member Of" designation is part of Domain Users. For VPN password management operations, the login DN needs elevated privileges, and must be part of the Account Operators AD group.

The following is an example of a login DN:

```
cn=Binduser1,ou=Admins,ou=Users,dc=company_A,dc=com
```

The ASA supports the following authentication methods:

- Simple LDAP authentication with an unencrypted password on port 389
- Secure LDAP (LDAP-S) on port 636
- Simple Authentication and Security Layer (SASL) MD5

The ASA does not support anonymous authentication.



Note As an LDAP client, the ASA does not support the transmission of anonymous binds or requests.

LDAP Attribute Maps

The ASA can use an LDAP directory for authenticating users for:

- VPN remote access users
- Firewall network access/cut-through-proxy sessions

- Setting policy permissions (also called authorization attributes), such as ACLs, bookmark lists, DNS or WINS settings, and session timers.
- Setting the key attributes in a local group policy

The ASA uses LDAP attribute maps to translate native LDAP user attributes to ASA attributes. You can bind these attribute maps to LDAP servers or remove them. You can also show or clear attribute maps.

The LDAP attribute map does not support multi-valued attributes. For example, if a user is a member of several AD groups, and the LDAP attribute map matches more than one group, the value chosen is based on the alphabetization of the matched entries.

To use the attribute mapping features correctly, you need to understand LDAP attribute names and values, as well as the user-defined attribute names and values.

The names of frequently mapped LDAP attributes and the type of user-defined attributes that they would commonly be mapped to include the following:

- IETF-Radius-Class (Group_Policy in ASA version 8.2 and later)—Sets the group policy based on the directory department or user group (for example, Microsoft Active Directory memberOf) attribute value. The group policy attribute replaced the IETF-Radius-Class attribute with ASDM version 6.2/ASA version 8.2 or later.
- IETF-Radius-Filter-Id—Applies an access control list or ACL to VPN clients, IPsec, and SSL.
- IETF-Radius-Framed-IP-Address—Assigns a static IP address assigned to a VPN remote access client, IPsec, and SSL.
- Banner1—Displays a text banner when the VPN remote access user logs in.
- Tunneling-Protocols—Allows or denies the VPN remote access session based on the access type.



Note A single LDAP attribute map may contain one or many attributes. You can only map one LDAP attribute from a specific LDAP server.

Guidelines for LDAP Servers for AAA

This section includes the guidelines and limitations that you should check before configuring LDAP servers for AAA.

IPv6

The AAA server can use either an IPv4 or IPv6 address.

Additional Guidelines

- The DN configured on the ASA to access a Sun directory server must be able to access the default password policy on that server. We recommend using the directory administrator, or a user with directory administrator privileges, as the DN. Alternatively, you can place an ACL on the default password policy.
- You must configure LDAP over SSL to enable password management with Microsoft Active Directory and Sun servers.

- The ASA does not support password management with Novell, OpenLDAP, and other LDAPv3 directory servers.
- Beginning with Version 7.1(x), the ASA performs authentication and authorization using the native LDAP schema, and the Cisco schema is no longer needed.
- You can have up to 200 server groups in single mode or 4 server groups per context in multiple mode.
- Each group can have up to 16 servers in single mode or 8 servers in multiple mode.
- When a user logs in, the LDAP servers are accessed one at a time, starting with the first server that you specify in the configuration, until a server responds. If all servers in the group are unavailable, the ASA tries the local database if you configured it as a fallback method (management authentication and authorization only). If you do not have a fallback method, the ASA continues to try the LDAP servers.

Configure LDAP Servers for AAA

This section describes how to configure LDAP servers for AAA.

Procedure

Step 1 Configure LDAP attribute maps. See Configure LDAP Attribute Maps, on page 1013.

Step 2 Add an LDAP server group. See Configure LDAP Server Groups, on page 1014.

Step 3 Add a server to the group, then configure server parameters. See Add an LDAP Server to a Server Group , on page 1015.

Configure LDAP Attribute Maps

To configure LDAP attribute maps, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Remote Access VPN** > **AAA Local Users** > **LDAP Attribute Map** (for local users), or **Configuration** > **Device Management** > **Users/AAA** > **LDAP Attribute Map** (for all other users), then click **Add**.

The **Add LDAP Attribute Map** dialog box appears with the **Mapping of Attribute Name** tab active.

Step 2 Create a name for this attribute map.

Step 3 Add the name of one of the LDAP attributes to be mapped.

Step 4 Choose a Cisco attribute.

Step 5 Click **Add**.

Step 6 To map more attributes, repeat Steps 1 through 5.

Step 7 Click the **Mapping of Attribute Value** tab to map the value of any of the LDAP attributes to a new value in the mapped Cisco attribute.

Step 8 Click **Add** to display the **Add Mapping of Attribute Value** dialog box.

Step 9 Enter the value for this LDAP attribute that you expect to be returned from the LDAP server.

Step 10 Enter the value you want to use in the Cisco attribute when this LDAP attribute contains the previous LDAP attribute value.

Step 11 Click **Add**.

Step 12 To map more attribute values, repeat Steps 8 through 11.

Step 13 Click **OK** twice to close each dialog box.

Step 14 Click **Apply** to save the settings to the running configuration.

Configure LDAP Server Groups

To create and configure an LDAP server group, then add an LDAP server to that group, perform the following steps:

Before you begin

You must add an attribute map before you may add an LDAP server to an LDAP server group.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Users/AAA** > **AAA Server Groups**, or **Configuration** > **Remote Access VPN** > **AAA/Local Users** > **AAA Server Groups** for VPN users.

Step 2 Click **Add**.

The **Add AAA Server Group** dialog box appears.

Step 3 Enter a name for the AAA server group.

Step 4 Choose the LDAP server type from the **Protocol** drop-down list.

Step 5 Click the radio button for the reactivation mode you want to use (**Depletion** or **Timed**).

In Depletion mode, failed servers are reactivated only after all of the servers in the group are inactive.

In Timed mode, failed servers are reactivated after 30 seconds of down time.

a) If you chose the Depletion reactivation mode, enter a time interval in the **Dead Time** field.

The dead time is the duration of time, in minutes, that elapses between the disabling of the last server in a group and the subsequent re-enabling of all servers. Deadtime applies only if you configure fallback to the local database; authentication is attempted locally until the deadtime elapses.

Step 6 Add the maximum number of failed AAA transactions with the server to allow.

This option sets the number of failed connection attempts allowed before declaring a non-responsive server to be inactive.

Step 7 Click **OK**.

The **Add AAA Server Group** dialog box closes, and the new server group is added to the AAA server group.

Step 8 Click **Apply** to save the changes to the running configuration.

Add an LDAP Server to a Server Group

To add an LDAP server to a server group, perform the following steps:

Procedure

Step 1 Choose one of the following:

- **Configuration Remote Access VPN AAA/Local Users AAA Server Groups** for VPN users.
- **Configuration** >**Device Management** > **Users/AAA** > **AAA Server Groups**

Step 2 Select the server group to which you want to add a server, then click **Add**.

The **Add AAA Server** dialog box appears for the selected server group.

Step 3 Choose the name of the interface that connects to the LDAP server.

Step 4 Add either the server name or IP address of the LDAP server.

Step 5 Add a timeout value or keep the default. The timeout is the duration of time, in seconds, that the ASA waits for a response from the primary server before sending the request to the backup server.

Step 6 In the **LDAP Parameters for authentication/authorization** area, configure the following settings:

- **Enable LDAP over SSL** (also called secure LDAP or LDAP-S)—Check this check box to use SSL to secure communications between the ASA and the LDAP server.

  Note
  If you do not configure the SASL protocol, we strongly recommend that you secure LDAP communications with SSL.

- **Reference Identity Name**—Enter reference-identity name to validate LDAP server identity.
- **Server Port**—Enter TCP port number 389, the port which the ASA uses to access the LDAP server for simple (non-secure) authentication, or TCP port 636 for secure authentication (LDAP-S). All LDAP servers support authentication and authorization. Only Microsoft AD and Sun LDAP servers additionally provide a VPN remote access password management capability, which requires LDAP-S.
- **Server Type**—Specify the LDAP server type from the drop-down list. The available options include the following:
  - **Detect Automatically/Use Generic Type**
  - **Microsoft**
  - **Novell**
  - **OpenLDAP**
  - **Sun, now part of Oracle Directory Server Enterprise Edition**

- **Base DN**—Enter the base distinguished name (DN) or location in the LDAP hierarchy where the server should begin searching when it receives an LDAP request (for example, OU=people, dc=cisco, dc=com).
- **Scope**—Specify the extent of the search that the server should perform in the LDAP hierarchy when it receives an authorization request from the drop-down list. The following options are available:
  - **One Level**—Searches only one level beneath the Base DN. This option is quicker.
  - **All Levels**—Searches all levels beneath the Base DN (that is, searches the entire subtree hierarchy). This option takes more time.
- **Naming Attribute(s)**—Enter the relative distinguished name attribute(s) that uniquely identify an entry on the LDAP server. Common naming attributes are Common Name (CN), sAMAccountName, userPrincipalName, and User ID (uid).
- **Login DN and Login Password**—The ASA uses the login DN and login password to establish trust (bind) with an LDAP server. Specify the login password, which is the password for the login DN user account.
- **LDAP Attribute Map**—Select one of the attribute maps that you created for this LDAP server to use. These attribute maps map LDAP attribute names to Cisco attribute names and values.
- **SASL MD5 authentication**—This enables the MD5 mechanism of the SASL to authenticate communications between the ASA and the LDAP server.
- **LDAP Parameters for Group Search**—The fields in this area configure how the ASA requests AD groups.
  - **Group Base DN**—Specifies the location in the LDAP hierarchy to begin searching for the AD groups (that is, the list of memberOf enumerations). If this field is not configured, the ASA uses the base DN for AD group retrieval. ASDM uses the list of retrieved AD groups to define AAA selection criteria for dynamic access policies. For more information, see the **show ad-groups** command.
  - **Group Search Timeout**—Specify the maximum time to wait for a response from an AD server that was queried for available groups.
- **LDAP SSL Client Certificate/Client Identity Certificate Trustpoint**—If you enable LDAP over SSL, you can select the certificate trustpoint that the ASA client should present to the LDAP server for authentication. A trustpoint is required if you configure the LDAP server to authenticate client certificates. If you do not configure a certificate, the ASA does not present one when the LDAP server asks for it. If an LDAP server is configured to require a peer certificate, the secure LDAP session will not complete and authentication/authorization requests will fail.

Step 7 Click **OK**.

The **Add AAA Server** dialog box closes, and the AAA server is added to the AAA server group.

Step 8 Click **Apply** to save the changes to the running configuration.

Test LDAP Server Authentication and Authorization

To determine whether the ASA can contact an LDAP server and authenticate or authorize a user, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Users/AAA** > **AAA Server Groups**.

Step 2 Select the server group in which the server resides.

Step 3 Select the server that you want to test.

Step 4 Click **Test**.

The **Test AAA Server** dialog box appears for the selected server.

Step 5 Click the type of test that you want to perform—**Authentication** or **Authorization**.

Step 6 Enter a username.

Step 7 If you are testing authentication, enter the password for the username.

Step 8 Click **OK**.

The ASA sends either an authentication or authorization test message to the server. If the test fails, an error message appears.

Monitoring LDAP Servers for AAA

See the following commands for monitoring LDAP servers for AAA:

- **Monitoring > Properties > AAA Servers**

  This pane shows the configured AAA server statistics.

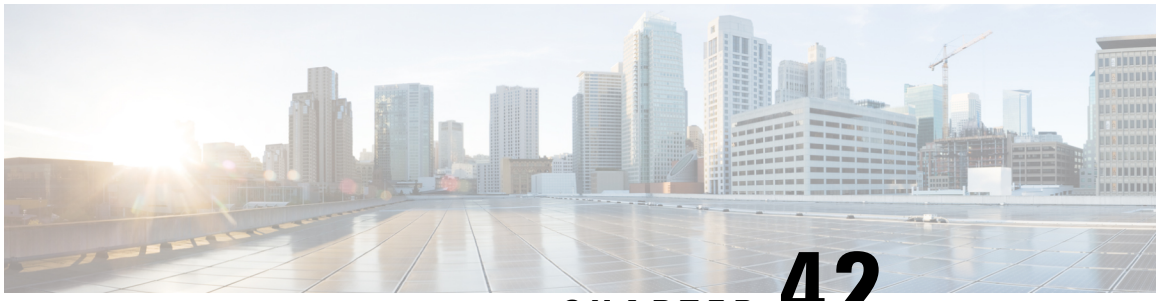
- **Tools > Command Line Interface**

  This pane allows you to issue various non-interactive commands and view results.

History for LDAP Servers for AAA

Table 54: History for AAA Servers

Feature Name	Platform Releases	Description
LDAP Servers for AAA	7.0(1)	LDAP Servers describe support for AAA and how to configure LDAP servers. We introduced the following screens: Configuration > Device Management > Users/AAA > AAA Server Groups Configuration > Remote Access VPN > AAA Local Users > LDAP Attribute Map.
LDAP servers with IPv6 addresses for AAA	9.7(1)	You can now use either an IPv4 or IPv6 address for the AAA server.
Increased limits for AAA server groups and servers per group.	9.13(1)	You can configure more AAA server groups. In single context mode, you can configure 200 AAA server groups (the former limit was 100). In multiple context mode, you can configure 8 (the former limit was 4). In addition, in multiple context mode, you can configure 8 servers per group (the former limit was 4 servers per group). The single context mode per-group limit of 16 remains unchanged. We modified the AAA screens to accept these new limits.
Mutual LDAPS authentication.	9.18(1)	You can configure a client certificate for the ASA to present to the LDAP server when it requests a certificate to authenticate. This feature applies when using LDAP over SSL. If an LDAP server is configured to require a peer certificate, the secure LDAP session will not complete and authentication/authorization requests will fail. We modified the following screen: **Configuration** > **Device Management** > **Users/AAA** > > **AAA Server Groups**, Add/Edit LDAP server.



CHAPTER 42

RSA SecurID Servers for AAA

The following topics explain how to configure RSA SecurID servers used in AAA. The RSA SecureID servers are also known as SDI servers, because SDI is the protocol used to communicate with them. You can use RSA SecurID servers for the authentication of management connections, network access, and VPN user access.

- About RSA SecurID Servers, on page 1019
- Guidelines for RSA SecurID Servers for AAA, on page 1019
- Configure RSA SecurID Servers for AAA, on page 1020
- Monitor RSA SecurID Servers for AAA, on page 1022
- History for RSA SecurID Servers for AAA, on page 1022

About RSA SecurID Servers

You can use RSA SecurID servers either directly for authentication, or indirectly, as a second factor for authentication. In the latter case, you would configure the relationship to the SecurID server between the SecurID server and your RADIUS server, and configure the ASA to use the RADIUS server.

But, if you want to directly authenticate against the SecurID server, you would create a AAA server group for the SDI protocol, which is the protocol used to communicate with these servers.

When you use SDI, you need only specify the primary SecurID server when you create the AAA server group. The ASA will retrieve the sdiconf.rec file, which lists all of the SecurID server replicas, when it first connects to the server. The ASA can then use these replicas for authentication if the primary server does not respond.

In addition, you must register the ASA as an authentication agent in the RSA Authentication Manager. Authentication attempts will fail until you register the ASA.

Guidelines for RSA SecurID Servers for AAA

- You can have up to 200 server groups in single mode or 8 server groups per context in multiple mode.
- Each group can have up to 16 servers in single mode or 8 servers in multiple mode. When a user logs in, the servers are accessed one at a time starting with the first server you specify in the configuration, until a server responds.

Configure RSA SecurID Servers for AAA

The following topics explain how to configure RSA SecurID server groups. You can then use these groups when configuring management access or VPNs.

Configure RSA SecurID AAA Server Groups

If you want to use direct communication with an RSA SecurID server for authentication, you must first create at least one SDI server group and add one or more servers to each group. If you are using the SecurID server in a proxy relationship with a RADIUS server, you do not need to configure an SDI AAA server group on the ASA.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Users/AAA** > **AAA Server Groups**.

Step 2 Click **Add** in the **AAA Server Groups** area.

The **Add AAA Server Group** dialog box appears.

Step 3 Enter a name for the group in the **Server Group** field.

Step 4 Choose the **SDI** server type from the **Protocol** drop-down list:

Step 5 Click **Depletion** or **Timed** in the **Reactivation Mode** field.

In Depletion mode, failed servers are reactivated only after all of the servers in the group are inactive. In depletion mode, when a server is deactivated, it remains inactive until all other servers in the group are inactive. When and if this occurs, all servers in the group are reactivated. This approach minimizes the occurrence of connection delays due to failed servers.

In Timed mode, failed servers are reactivated after 30 seconds of down time.

Step 6 If you chose the Depletion reactivation mode, enter a time interval in the **Dead Time** field.

The dead time is the duration of time, in minutes, that elapses between the disabling of the last server in a group and the subsequent re-enabling of all servers. Deadtime applies only if you configure fallback to the local database; authentication is attempted locally until the deadtime elapses.

Step 7 Specify the maximum number of failed AAA transactions with a AAA server in the group before trying the next server.

This option sets the number of failed AAA transactions before declaring a nonresponsive server to be inactive.

Step 8 Click **OK**.

Add RSA SecurID Servers to an SDI Server Group

Before you can use an SDI server group, you must add at least one RSA SecurID server to the group.

Servers in an SDI server group use the authentication and server management protocol (ACE) to communicate with the ASA.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Users/AAA** > **AAA Server Groups**.

Step 2 Select the server group to which you want to add a server.

Step 3 Click **Add** in the **Servers in the Selected Group** area.

The **Add AAA Server Group** dialog box appears for the server group.

Step 4 Choose the **Interface Name** through which the authentication server resides.

Step 5 Enter either the name or IP address for the server that you are adding to the group.

Step 6 Specify the timeout value for connection attempts to the server.

Specify the timeout interval (1-300 seconds) for the server; the default is 10 seconds. For each AAA transaction the ASA retries connection attempts (based on the retry interval) until the timeout is reached. If the number of consecutive failed transactions reaches the maximum-failed-attempts limit specified in the AAA server group, the AAA server is deactivated and the ASA starts sending requests to another AAA server if it is configured.

Step 7 Select the retry interval, which is the time the system waits before retrying a connection request. You can select from 1-10 seconds. The default is 10 seconds.

Step 8 Specify the server port. The server port is either the default port number 5500, or the TCP port number used by the ASA to communicate with the RSA SecurID server.

Step 9 Click **OK**.

Import the SDI Node Secret File

You can manually import the node-secret file that is generated by the RSA Authentication Manager (SecurID) server.

Procedure

Step 1 Export the node secret file from the RSA Authentication Manager server. For details, see the RSA Authentication Manager documentation.

Step 2 Choose **Configuration** > **Device Management** > **Users/AAA** > **AAA SDI**.

Step 3 Click **Upload**, select the unzipped node secret file that was exported from the RSA Authentication Manager, and upload it to the system.

Step 4 Under **Import Node Secret for SDI**, enter the following information:

- **Server IP**—The IP address or fully-qualified hostname of the RSA Authentication Manager server to which the node secret belongs.
- **Password**—The password used to protect the file when you exported it.

- **File Name**—Click **Browse** and select the unzipped node secret file that you uploaded.

Monitor RSA SecurID Servers for AAA

You can use the following commands to monitor and clear RSA SecurID-related information. Enter commands from the **Tools** > **Command Line Interface** window.

- **Monitoring** > **Properties** > **AAA Servers**

  This window shows the AAA server statistics.

- **show aaa-server**

  Shows the AAA server statistics. Use the **clear aaa-server statistics** command to clear the server statistics.

- **show running-config aaa-server**

  Shows the AAA servers that are configured for the system. Use the **clear configure aaa-server** command to remove the AAA server configuration.

- **show aaa sdi node-secrets**

  Shows which RSA SecurID servers have an imported node secret file. Use the **clear aaa sdi node-secret** command to remove a node secret file.

History for RSA SecurID Servers for AAA

Feature Name	Platform Releases	Description
SecurID Servers	7.2(1)	Support for SecurID servers for AAA for management authentication. SecurID was supported in previous releases for VPN authentication.
IPv6 addresses for AAA	9.7(1)	You can now use either an IPv4 or IPv6 address for the AAA server.
Increased limits for AAA server groups and servers per group.	9.13(1)	You can configure more AAA server groups. In single context mode, you can configure 200 AAA server groups (the former limit was 100). In multiple context mode, you can configure 8 (the former limit was 4). In addition, in multiple context mode, you can configure 8 servers per group (the former limit was 4 servers per group). The single context mode per-group limit of 16 remains unchanged. We modified the AAA screens to accept these new limits.

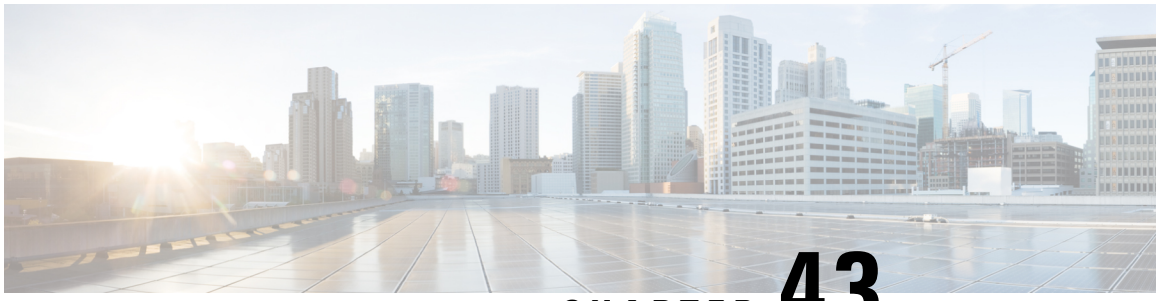
Feature Name	Platform Releases	Description
Manual import of node secret file from the RSA Authentication Manager for SDI AAA server groups.	9.15(1)	You can import the node secret file that you export from the RSA Authentication Manager for use with SDI AAA server groups. We added the following screen: **Configuration** > **Device Management** > **Users/AAA** > **AAA SDI**.



PART VII

System Administration

- Management Access, on page 1027
- Software and Configurations, on page 1069
- Response Automation for System Events, on page 1099
- Testing and Troubleshooting, on page 1105



CHAPTER 43

Management Access

This chapter describes how to access the ASA for system management through Telnet, SSH, and HTTPS (using ASDM), how to authenticate and authorize users, and how to create login banners.

- Configure Management Remote Access, on page 1027
- Configure AAA for System Administrators, on page 1043
- Monitoring Device Access, on page 1059
- History for Management Access, on page 1060

Configure Management Remote Access

This section describes how to configure ASA access for ASDM, Telnet, or SSH, and other management parameters such as a login banner.

Configure ASA Access for HTTPS, Telnet, or SSH

This section describes how to configure ASA access for HTTPS (including ASDM and CSM), Telnet, or SSH. See the following guidelines:

- To access the ASA interface for management access, you do not also need an access rule allowing the host IP address. You only need to configure management access according to the sections in this chapter. If, however, you configure HTTP redirect to redirect HTTP connections to HTTPS automatically, you must enable an access rule to allow HTTP; otherwise, the interface cannot listen to the HTTP port.
- Management access to an interface other than the one from which you entered the ASA is not supported. For example, if your management host is located on the outside interface, you can only initiate a management connection directly to the outside interface. The only exception to this rule is through a VPN connection. See Configure Management Access Over a VPN Tunnel, on page 1038.
- The ASA allows:
  - A maximum of 5 concurrent Telnet connections per context, if available, with a maximum of 100 connections divided among all contexts.
  - A maximum of 5 concurrent SSH connections per context, if available, with a maximum of 100 connections divided among all contexts.

- In single context mode, you can have a maximum 5 ASDM concurrent sessions. In multiple context mode, you can have a maximum of 5 concurrent ASDM sessions per context, with a maximum of 200 ASDM instances among all contexts.

  ASDM sessions use two HTTPS connections: one for monitoring that is always present, and one for making configuration changes that is present only when you make changes. For example, the multiple-context mode system limit of 200 ASDM sessions represents a limit of 400 HTTPS sessions.

- A maximum of 6 concurrent non-ASDM HTTPS sessions in single context mode or per context, if available, with a maximum or 100 HTTPS sessions among all contexts.

Configure HTTPS Access for ASDM, Other Clients

This section describes how to configure ASA access for HTTPS, including ASDM and CSM.

If you enable both SSL (**webvpn** > **enable** *interface*) and HTTPS access on the same interface, you can access Secure Client from **https://***ip_address* and ASDM from **https://***ip_address***/admin**, both on port 443. If you also enable authentication for HTTPS (Configure Authentication for CLI, ASDM, and enable command Access, on page 1045), then you must specify a different port for ASDM access.

Before you begin

- In multiple context mode, complete this procedure in the context execution space. To change from the system to a context configuration, in the **Configuration > Device List** pane, double-click the context name under the active device IP address.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Management Access** > **ASDM/HTTPS/Telnet/SSH**, and click **Add**.

The **Add Device Access Configuration** dialog box appears.

Step 2 Choose **ASDM/HTTPS**.

Step 3 Choose the management interface and set the host IP addresses allowed, and click **OK**.

Specify any named interface. For bridge groups, specify the bridge group member interface. For VPN management access only (see Configure Management Access Over a VPN Tunnel, on page 1038), specify the named BVI interface.

Step 4 To require certificate authentication, in the **Specify the interface requires client certificate to access ASDM** area, click **Add** to specify the interface and an optional certificate map that must be matched for successful authentication. See **Configuration** > **Remote Access VPN** > **Network (Client) Access** > **Advanced** > **IPSec** > **Certificate to Connection Map** > **Rules** to create the certificate map. For more information, see Configure ASDM Certificate Authentication, on page 1046.

Step 5 Configure **HTTP Settings**.

- **Enable HTTP Server**—Enable the HTTPS server.
- **Port Number**—Set the port number. The default is 443.

- **Idle Timeout**—Set the idle timeout for ASDM connections, from 1-1440 minutes. The default is 20 minutes. The ASA disconnects an ASDM connection that is idle for the set period of time.
- **Session Timeout**—Set the session timeout for ASDM sessions, from 1-1440 minutes. This timeout is disabled by default. The ASA disconnects an ASDM session that exceeds the set period of time.
- **Connection Session Timeout**—Set the idle timeout for all HTTPS connections, including ASDM, WebVPN, and other clients, from 10-86400 seconds. This timeout is disabled by default. The ASA disconnects a connection that is idle for the set period of time. If you set both the **Idle Timeout** and the **Connection Session Timeout**, the **Connection Session Timeout** takes precendence.

Step 6 Click **Apply**.

Step 7 (Optional) Allow non-browser-based HTTPS clients to access HTTPS services on the ASA. By default, ASDM, CSM, and REST API are allowed.

Many specialty clients (for example, python libraries, curl, and wget) do not support Cross-site request forgery (CSRF) token-based authentication, so you need to specifically allow these clients to use the ASA basic authentication method. For security purposes, you should only allow required clients.

a) Choose **Configuration** > **Device Management** > **Management Access** > **HTTP Non-Browser Client Support**, and click **Add**.

b) In the **User-Agent String from the HTTP Header** field, specify the client's User-Agent string in the HTTP header of the HTTP request.

You can specify the complete string or a partial string; partial strings must match the start of the User-Agent string. We recommend complete strings for better security. Note that the string is case-sensitive.

For example, **`curl`** will match the following User-Agent string:

```
curl/7.19.7 (x86_64-redhat-linux-gnu) libcurl/7.19.7 NSS/3.19.1 Basic
ECC zlib/1.2.3 libidn/1.18 libssh2/1.4.2
```

`curl` will *not* match the following User-Agent string:

```
abcd curl/7.19.7 (x86_64-redhat-linux-gnu) libcurl/7.19.7 NSS/3.19.1
Basic ECC zlib/1.2.3 libidn/1.18 libssh2/1.4.2
```

`CURL` will *not* match the following User-Agent string:

```
curl/7.19.7 (x86_64-redhat-linux-gnu) libcurl/7.19.7 NSS/3.19.1 Basic
ECC zlib/1.2.3 libidn/1.18 libssh2/1.4.2
```

Configure SSH Access

To configure SSH access to the ASA, enable the SSH server and identify the allowed IP addresses. To authenticate users, you can use the following methods:

- Username and password using either the local database or a AAA server.
- Username and public key using the local database.

SSH Guidelines

- To access the ASA interface for SSH access, you do not also need an access rule allowing the host IP address. You only need to configure SSH access according to this section.
- SSH access to an interface other than the one from which you entered the ASA is not supported. For example, if your SSH host is located on the outside interface, you can only initiate a management connection directly to the outside interface.
- The ASA allows a maximum of 5 concurrent SSH connections per context/single mode, with a maximum of 100 connections divided among all contexts. However, because configuration commands might obtain locks on resources being changed, you should make changes in one SSH session at a time to ensure all changes are applied correctly.
- SSH does not support:
  - EDDSA key pair
  - RSA key pair in FIPS mode
- To use the ASA **copy** command to copy a file to or from an SCP server, you have to:
  - Enable SSH access on the ASA for the SCP server subnet/host .
  - Generate a key pair (for physical ASAs only).
- The SSH default username is no longer supported. You can no longer connect to the ASA using SSH with the **pix** or **asa** username and the login password. To use SSH, you must configure AAA authentication by choosing **Configuration** > **Device Management** > **Users/AAA** > **AAA Access** > **Authentication**; then define a local user by choosing **Configuration** > **Device Management** > **Users/AAA**. If you want to use a AAA server for authentication instead of the local database, we recommend also configuring local authentication as a backup method.
- Only SSH Version 2 is supported.

Enable the SSH Server

Enable the SSH server and specify which IP addresses are allowed to connect. You can also configure other SSH server settings.

Before you begin

- In multiple context mode, complete this procedure in the context execution space. To change from the system to a context configuration, in the **Configuration > Device List** pane, double-click the context name under the active device IP address.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Management Access** > **ASDM/HTTPS/Telnet/SSH**, and click **Add**.

The **Add Device Access Configuration** dialog box appears.

Step 2 Choose **SSH**.

Step 3 Choose the management interface and set the host IP addresses allowed, and click **OK**.

Specify any named interface. For bridge groups, specify the bridge group member interface.

Step 4 (Optional) Configure **SSH Settings**.

- **SSH Timeout**—Set the timeout from 1 to 60 minutes. The default is 5 minutes. The default duration is too short in most cases, and should be increased until all pre-production testing and troubleshooting have been completed.
- **Key Exchange Hostkey**—By default, the ASA tries to use keys in the following order if they exist:ECDSA, and then RSA. If you explicitly choose the RSA key, then you must generate a key that is 2048 bits or higher. For upgrade compatibility, the ASA will use smaller RSA host keys only when the default host key setting is used.
- **DH Key Exchange** (Admin context only)—Click the applicable radio button to choose the Diffie-Hellman (DH) Key Exchange Group. If no DH group key-exchange method is specified, the DH group 14 SHA256 key-exchange method is used. For more information about using DH key-exchange methods, see RFC 4253. You can only set the key exchange in the Admin context; this value is used by all contexts.

Step 5 Click **Apply**.

Step 6 (Optional) Configure SSH cipher encryption and integrity algorithms:

a) Choose **Configuration** > **Device Management** > **Advanced** > **SSH Ciphers**.

b) Select **Encryption**, and click **Edit**.

c) From the SSH cipher security level drop-down list, choose one of the following levels.

Ciphers are used in the order they are listed. For pre-defined lists, they are listed from highest to lowest security.

- **All**—Specifies using all ciphers: 3des-cbc aes128-cbc aes192-cbc aes256-cbc aes128-ctr chacha20-poly1305@openssh.com aes192-ctr aes256-ctr
- **Custom**—Specifies a custom cipher encryption configuration string that you enter in the **Cipher algorithms/custom string** field, separated by colons.
- **Fips**—Specifies only FIPS-compliant ciphers: aes128-cbc aes256-cbc
- **High**—Specifies only high-strength ciphers: aes256-cbc chacha20-poly1305@openssh.com aes256-ctr
- **Low**—Specifies low, medium, and high strength ciphers: 3des-cbc aes128-cbc aes192-cbc aes256-cbc aes128-ctr aes192-ctr aes256-ctr
- **Medium**—Specifies the medium and high strength ciphers (the default): 3des-cbc aes128-cbc aes192-cbc aes256-cbc aes128-ctr aes192-ctr aes256-ctr

d) Select **Integrity**, and click **Edit**.

e) From the SSH cipher security level drop-down list, choose one of the following levels:

- **All**—Specifies using all ciphers: hmac-sha1 hmac-sha1-96 hmac-sha2-256 hmac-md5 hmac-md5-96
- **Custom**—Specifies a custom cipher encryption configuration string that you enter in the **Cipher algorithms/custom string** field, separated by colons.
- **Fips**—Specifies only FIPS-compliant ciphers: hmac-sha1 hmac-sha2-256

- **High**—Specifies only high-strength ciphers (the default): hmac-sha2-256
- **Low**—Specifies low, medium, and high strength ciphers: hmac-sha1 hmac-sha1-96 hmac-md5 hmac-md5-96
- **Medium**—Specifies the medium and high strength ciphers: hmac-sha1 hmac-sha1-96

Step 7 Enable the Secure Copy server.

a) Depending on your context mode:

- For single mode, choose **Configuration > Device Management > Management Access > File Access > Secure Copy (SCP)**.
- For multiple mode in the System, choose **Configuration > Device Management > Device Administration > Secure Copy**

b) Check the **Enable secure copy server** check box.

Step 8 Click **Apply**.

Examples

The following example shows an SCP session to the ASA. From a client on the external host, perform an SCP file transfer. For example, in Linux enter the following command:

scp -v -pw *password* [*path*/]*source_filename*
*username@asa_address***:**{**disk0**|**disk1**}**:/**[*path*/]*dest_filename*

The **-v** is for verbose, and if **-pw** is not specified, you will be prompted for a password.

Configure SSH for Password Access

Configure SSH authentication using a username and password.

Before you begin

- In multiple context mode, complete this procedure in the context execution space. To change from the system to a context configuration, in the **Configuration > Device List** pane, double-click the context name under the active device IP address.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Users/AAA** > **AAA Access** > **Authentication**.

Step 2 Check the **SSH** check box.

Step 3 Choose the **LOCAL** database or a AAA server from the **Server Group** drop-down list.

This setting does not affect local public key authentication for usernames that have a public key configured. The ASA implicitly uses the local database for public key authentication. This setting only affects local usernames with passwords. If you want to allow either public key authentication or password use by a local user, then you need to explicitly configure local authentication with this setting to allow password access.

Step 4 Click **Apply**.

Step 5 Create a user in the local database or on a AAA server that can be used for SSH access. See the following procedure to add a local username (recommended).

a) Choose **Configuration** > **Device Management** > **Users/AAA** > **User Accounts**, then click **Add**.

The **Add User Account-Identity** dialog box appears.

b) Enter a username and password, then confirm the password.

You can use public key authentication or password authentication for the same local user. Public key authentication is not supported using a AAA server.

c) Click **OK**, then click **Apply**.

Step 6 Generate a key pair (for physical ASAs only).

For the ASAv, the key pairs are automatically created after deployment. The ASAv only supports the RSA key.

a) Choose **Configuration** > **Device Management** > **Certificate Management** > **Identity Certificates**.
b) Click **Add** and click the **Add a new identity certificate** radio button.
c) Click **New**.
d) In the **Add Key Pair** dialog box, specify the type and size, and click **Generate Now**.

The default key pair used is ECDSA, and then RSA. For RSA, choose a size 2048 bits or higher..

Note
With CiscoSSH stack, only the configured key pair is shared by ASA at the time of establishing a connection.

You can then **Cancel** out of the certificate dialog box, because you only wanted to generate the key pair.

e) Click **Apply**.

Configure SSH for Public Key Access

Configure SSH authentication using a public key.

Before you begin

- In multiple context mode, complete this procedure in the context execution space. To change from the system to a context configuration, in the **Configuration > Device List** pane, double-click the context name under the active device IP address.

Procedure

Step 1 Create a user in the local database that can be used for SSH access.

a) Choose **Configuration** > **Device Management** > **Users/AAA** > **User Accounts**, then click **Add**.

The **Add User Account-Identity** dialog box appears.

b) Enter a username and password, then confirm the password.

You might want to create a user without a password if you want to force the user to use public key authentication instead of password authentication. If you configure public key authentication as well as a password, then the user can log in with either method if you explicitly configure AAA authentication in this procedure.

c) Enable public key authentication. Choose one of the following panes:

- **Public Key Authentication**—Paste in a Base64-encoded public key. You can generate the key using any SSH key generation software (such as ssh keygen) that can generate ssh-rsa, or ecdsa-sha2-nistp raw keys (with no certificates).When you view an existing key, the key is encrypted using a SHA-256 hash. If you need to copy and paste a hashed key, check the **Key is hashed** check box.

  To remove an authentication key, click **Delete Key** to display a confirmation dialog box. Click **Yes** to remove the authentication key, or click **No** to retain it.

- **Public Key Using PKF**—Check the **Specify a new PKF** key check box, and paste or import a public key file (PKF) formatted key, up to 4096 bits. Use this format for keys that are too large to paste in Base64 format. For example, you can generate a 4096-bit key using ssh keygen, then convert it to PKF, and import on this pane. When you view an existing key, the key is encrypted using a SHA-256 hash. If you need to copy and paste a hashed key, copy it from the **Public Key Authentication** pane, and paste it in that pane on the new ASA with the **Key is hashed** check box checked.

  To remove an authentication key, click **Delete Key** to display a confirmation dialog box. Click **Yes** to remove the authentication key, or click **No** to retain it.

d) Click **OK**, then click **Apply**.

Step 2 Generate a key pair (for physical ASAs only).

For the ASAv, the key pairs are automatically created after deployment. The ASAv only supports the RSA key.

a) Choose **Configuration** > **Device Management** > **Certificate Management** > **Identity Certificates**.
b) Click **Add** and click the **Add a new identity certificate** radio button.
c) Click **New**.
d) In the **Add Key Pair** dialog box, specify the type and size, and click **Generate Now**.

The default key pair used is ECDSA, and then RSA. For RSA, choose a size 2048 bits or higher..

Note
With CiscoSSH stack, only the configured key pair is shared by ASA at the time of establishing a connection.

You can then **Cancel** out of the certificate dialog box, because you only wanted to generate the key pair.

e) Click **Apply**.

Examples

The following example shows how to authenticate using a PKF formatted key:

The following example generates a shared key for SSH on a Linux or Macintosh system, and imports it to the ASA:

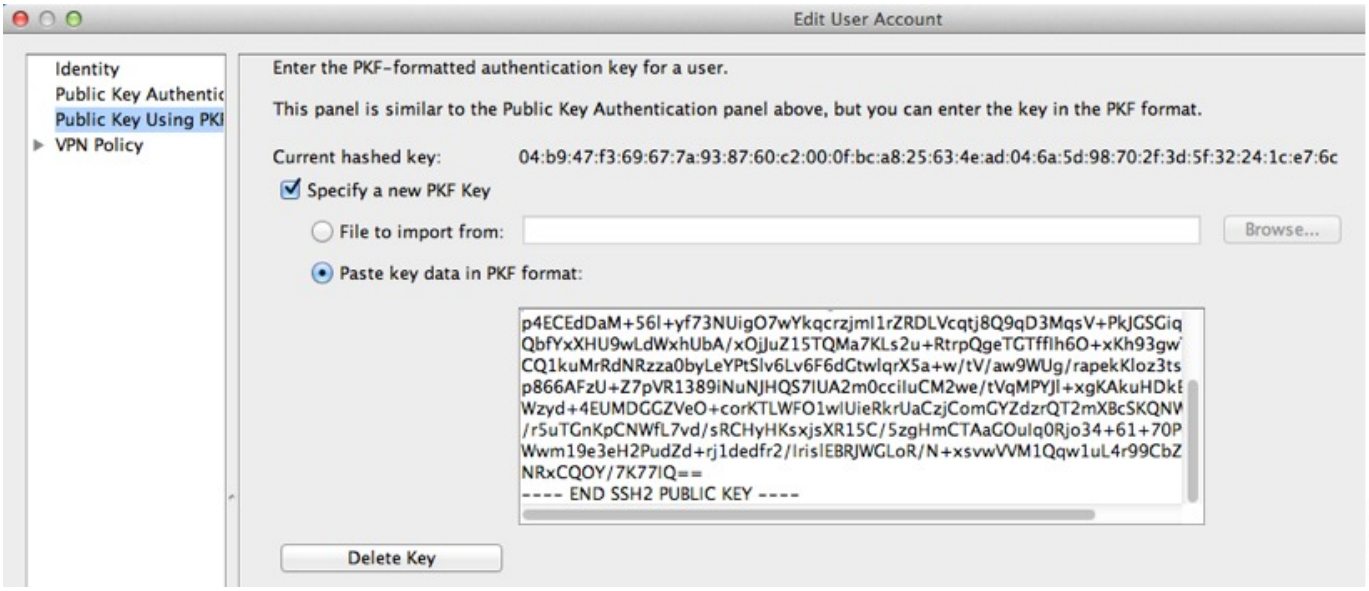
1. Generate the RSA public and private keys for 4096 bits on your computer:

```
jcrichton-mac:~ john$ ssh-keygen -b 4096
Generating public/private rsa key pair.
Enter file in which to save the key (/Users/john/.ssh/id_rsa):
/Users/john/.ssh/id_rsa already exists.
Overwrite (y/n)? y
Enter passphrase (empty for no passphrase): pa$$phrase
Enter same passphrase again: pa$$phrase
Your identification has been saved in /Users/john/.ssh/id_rsa.
Your public key has been saved in /Users/john/.ssh/id_rsa.pub.
The key fingerprint is:
c0:0a:a2:3c:99:fc:00:62:f1:ee:fa:f8:ef:70:c1:f9 john@jcrichton-mac
The key's randomart image is:
+--[ RSA 4096]----+
| .               |
|  o  .           |
|+...  o          |
|B.+.....         |
|.B ..+  S        |
|  =   o          |
|   + . E         |
|  o o            |
| ooooo           |
+-----------------+
```

2. Convert the key to PKF format:

```
jcrichton-mac:~ john$ cd .ssh
jcrichton-mac:.ssh john$ ssh-keygen -e -f id_rsa.pub
---- BEGIN SSH2 PUBLIC KEY ----
Comment: "4096-bit RSA, converted by john@jcrichton-mac from OpenSSH"
AAAAB3NzaC1yc2EAAAADAQABAAACAQDNUvkgza37lB/Q/fljpLAv1BbyAd5PJCJXh/U4LO
hleR/qgIROjpnDaS7Az8/+sjHmq0qXC5TXkzWihvRZbhefyPhPHCi0hIt4oUF2ZbXESA/8
jUT4ehXIUE7FrChffBBtbD4d9FkV8A2gwZCDJBxEM26ocbZCSTx9QC//wt6E/zRcdoqiJG
p4ECEdDaM+56l+yf73NUigO7wYkqcrzjmI1rZRDLVcqtj8Q9qD3MqsV+PkJGSGiqZwnyIl
QbfYxXHU9wLdWxhUbA/xOjJuZ15TQMa7KLs2u+RtrpQgeTGTffIh6O+xKh93gwTgzaZTK4
CQ1kuMrRdNRzza0byLeYPtSlv6Lv6F6dGtwlqrX5a+w/tV/aw9WUg/rapekKloz3tsPTDe
p866AFzU+Z7pVR1389iNuNJHQS7IUA2m0cciIuCM2we/tVqMPYJl+xgKAkuHDkBlMS4i8b
Wzyd+4EUMDGGZVeO+corKTLWFO1wIUieRkrUaCzjComGYZdzrQT2mXBcSKQNWlSCBpCHsk
/r5uTGnKpCNWfL7vd/sRCHyHKsxjsXR15C/5zgHmCTAaGOuIq0Rjo34+61+70PCtYXebxM
Wwm19e3eH2PudZd+rj1dedfr2/IrislEBRJWGLoR/N+xsvwVVM1Qqw1uL4r99CbZF9NghY
NRxCQOY/7K77IQ==
---- END SSH2 PUBLIC KEY ----
jcrichton-mac:.ssh john$
```

3. Copy the key to your clipboard.
4. In ASDM, choose **Configuration** > **Device Management** > **Users/AAA** > **User Accounts**, select the username and then click **Edit**. Click **Public Key Using PKF** and paste the key into the window:



5. Verify the user can SSH to the ASA. For the password, enter the SSH key password you specified when you created the key pair.

```
jcrichton-mac:.ssh john$ ssh test@10.86.118.5
The authenticity of host '10.86.118.5 (10.86.118.5)' can't be established.
RSA key fingerprint is 39:ca:ed:a8:75:5b:cc:8e:e2:1d:96:2b:93:b5:69:94.
Are you sure you want to continue connecting (yes/no)? yes
```

The following dialog box appears for you to enter your passphrase:



Meanwhile, in the terminal session:

```
Warning: Permanently added '10.86.118.5' (RSA) to the list of known hosts.
Identity added: /Users/john/.ssh/id_rsa (/Users/john/.ssh/id_rsa)
Type help or '?' for a list of available commands.
asa>
```

Configure Telnet Access

This section describes how to configure ASA access for Telnet. You cannot use Telnet to the lowest security interface unless you use Telnet inside a VPN tunnel.

Before you begin

- In multiple context mode, complete this procedure in the context execution space. To change from the system to a context configuration, in the **Configuration > Device List** pane, double-click the context name under the active device IP address.
- To gain access to the ASA CLI using Telnet, enter the login password. You must manually set the password before using Telnet.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Management Access** > **ASDM/HTTPS/Telnet/SSH**, and click **Add**.

The **Add Device Access Configuration** dialog box appears.

Step 2 Choose **Telnet**.

Step 3 Choose the management interface and set the host IP addresses allowed, and click **OK**.

Specify any named interface. For bridge groups, specify the bridge group member interface. For VPN management access only (see Configure Management Access Over a VPN Tunnel, on page 1038), specify the named BVI interface.

Step 4 (Optional) Set the **Telnet Timeout**. The default timeout value is 5 minutes.

Step 5 Click **Apply**.

Step 6 Set a login password before you can connect with Telnet; there is no default password.

a) Choose **Configuration** > **Device Setup** > **Device Name/Password**.
b) Check the **Change the password to access the console of the security appliance** check box in the **Telnet Password** area.
c) Enter the old password (leave this field blank for a new ASA), new password, then confirm the new password.
d) Click **Apply**.

Configure HTTP Redirect for ASDM Access or Clientless SSL VPN

You must use HTTPS to connect to the ASA using ASDM or clientless SSL VPN. For your convenience, you can redirect HTTP management connections to HTTPS. For example, by redirecting HTTP, you can enter either **http://10.1.8.4/admin/** or **https://10.1.8.4/admin/** and still arrive at the ASDM launch page at the HTTPS address.

You can redirect both IPv4 and IPv6 traffic.

Before you begin

Normally, you do not need an access rule allowing the host IP address. However, for HTTP redirect, you must enable an access rule to allow HTTP; otherwise, the interface cannot listen to the HTTP port.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **HTTP Redirect**.

The table shows the currently configured interfaces and whether redirection is enabled on an interface.

Step 2 Select the interface that you use for ASDM, and click **Edit**.

Step 3 Configure the following options in the **Edit HTTP/HTTPS Settings** dialog box:

- **Redirect HTTP to HTTPS**—Redirects HTTP requests to HTTPS.
- **HTTP Port**—Identifies the port from which the interface redirects HTTP connections. The default is 80.

Step 4 Click **OK**.

Configure Management Access Over a VPN Tunnel

If your VPN tunnel terminates on one interface, but you want to manage the ASA by accessing a different interface, you must identify that interface as a management-access interface. For example, if you enter the ASA from the outside interface, this feature lets you connect to the inside interface using ASDM, SSH, or Telnet; or you can ping the inside interface when entering from the outside interface.



Note This feature is not supported for SSH if you use the CiscoSSH stack, which is the default.



Note This feature is not supported for SNMP. For SNMP over VPN, we recommend enabling SNMP on a loopback interface. You don't need the management-access feature enabled to use SNMP on the loopback interface. Loopback also works for SSH.

VPN access to an interface other than the one from which you entered the ASA is not supported. For example, if your VPN access is located on the outside interface, you can only initiate a connection directly to the outside interface. You should enable VPN on the directly-accessible interface of the ASA and use name resolution so that you don’t have to remember multiple addresses.

Management access is available via the following VPN tunnel types: IPsec clients, IPsec Site-to-Site, Easy VPN, and the Secure Client SSL VPN.

Before you begin

- This feature is not supported on management-only interfaces.

- When you use a management-access interface and you configure identity NAT, you must configure NAT with the route lookup option. For more information see the "NAT and VPN Management Access" section in the *NAT Examples and Reference* chapter in the appropriate release of the ASA Firewall CLI Configuration Guide.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Management Access** > **Management Interface**.

Step 2 Choose the interface with the highest security (the inside interface) from the **Management Access Interface** drop-down list.

For Easy VPN and Site-to-Site tunnels, you can specify a named BVI (in routed mode).

Step 3 Click **Apply**.

The management interface is assigned, and the change is saved to the running configuration.

Change the Console Timeout

The console timeout sets how long a connection can remain in privileged EXEC mode or configuration mode; when the timeout is reached, the session drops into user EXEC mode. By default, the session does not time out. This setting does not affect how long you can remain connected to the console port, which never times out.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Management Access** > **Command Line (CLI)** > **Console Timeout**.

Step 2 Define a new timeout value in minutes, To specify an unlimited amount of time, enter **0**. The default value is 0.

Step 3 Click **Apply**.

The timeout value change is saved to the running configuration.

Customize a CLI Prompt

The ability to add information to a prompt allows you to see at-a-glance which ASA you are logged into when you have multiple modules. During a failover, this feature is useful when both ASAs have the same hostname.

In multiple context mode, you can view the extended prompt when you log in to the system execution space or the admin context. Within a non-admin context, you only see the default prompt, which is the hostname and the context name.

By default, the prompt shows the hostname of the ASA. In multiple context mode, the prompt also displays the context name. You can display the following items in the CLI prompt:

cluster-unit	Displays the cluster unit name. Each unit in a cluster can have a unique name.
context	(Multiple mode only) Displays the name of the current context.
domain	Displays the domain name.
hostname	Displays the hostname.
priority	Displays the failover priority as pri (primary) or sec (secondary).
state	Displays the traffic-passing state or role of the unit. For failover, the following values are displayed for the state keyword: • **act**—Failover is enabled, and the unit is actively passing traffic. • **stby**— Failover is enabled, and the unit is not passing traffic and is in a standby, failed, or other non-active state. • **actNoFailover**—Failover is not enabled, and the unit is actively passing traffic. • **stbyNoFailover**—Failover is not enabled, and the unit is not passing traffic. This might happen when there is an interface failure above the threshold on the standby unit. For clustering, the values for control and data are shown.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Management Access** > **Command Line (CLI)** > **CLI Prompt.**

Step 2 Do any of the following to customize the prompt:

- Click the attribute in the **Available Prompts** list, then click **Add**. You can add multiple attributes to the prompt. The attribute is moved from the **Available Prompts** list to the **Selected Prompts** list.
- Click the attribute in the **Selected Prompts** list, then click **Delete**. The attribute is moved from the **Selected Prompts** list to the **Available Prompts** list.
- Click the attribute in the **Selected Prompts** list and click **Move Up** or **Move Down** to change the order in which the attributes appear.

The prompt is changed and appears in the **CLI Prompt Preview** field.

Step 3 Click **Apply**.

The new prompt is saved to the running configuration.

Configure a Login Banner

You can configure a message to display when a user connects to the ASA, before a user logs in, or before a user enters privileged EXEC mode.

Before you begin

- 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 possible criminal consequences.
```

- After a banner has been added, Telnet or SSH sessions to the ASA may close if:
  - There is not enough system memory available to process the banner message(s).
  - A TCP write error occurs when trying to display banner message(s).
- See RFC 2196 for guidelines about banner messages.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Management Access** > **Command Line (CLI)** > **Banner**.

Step 2 Add your banner text to the field for the type of banner that you are creating for the CLI:

- The session (exec) banner appears when a user accesses privileged EXEC mode at the CLI.
- The login banner appears when a user logs in to the CLI.
- The message-of-the-day (motd) banner appears when a user first connects to the CLI.
- The ASDM banner appears when a user connects to ASDM, after user authentication. The user is given two options for dismissing the banner:
  - **Continue**—Dismiss the banner and complete login.
  - **Disconnect**—Dismiss the banner and terminate the connection.
- Only ASCII characters are allowed, including a new line (Enter), which counts as two characters.
- Do not use tabs in the banner, because they are not preserved in the CLI version.
- There is no length limit for banners other than those for RAM and flash memory.
- You can dynamically add the hostname or domain name of the ASA by including the strings **$(hostname)** and **$(domain)**.
- If you configure a banner in the system configuration, you can use that banner text within a context by using the **$(system)** string in the context configuration.

Step 3 Click **Apply**.

The new banner is saved to the running configuration.

Set a Management Session Quota

You can establish a maximum number of simultaneous ASDM, SSH, and Telnet sessions that are allowed on the ASA. If the maximum is reached, no additional sessions are allowed and a syslog message is generated. To prevent a system lockout, the management session quota mechanism cannot block a console session.

Note In multiple context mode, you cannot configure the number of ASDM sessions, where the maximum is fixed at 5 sessions.

Note If you also set a resource limit per context for the maximum administrative sessions (SSH, etc.), then the lower value will be used.

Before you begin

In multiple context mode, complete this procedure in the context execution space. To change from the System to a context configuration, in the **Configuration** > **Device List** pane, double-click the context name under the active device IP address.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Management Access** > **Management Session Quota**.

Step 2 Enter the maximum number of simultaneous sessions.

- **Aggregate**—Sets the aggregate number of sessions between 1 and 15. The default is 15.
- **HTTP Sessions**—Sets the maximum HTTPS (ASDM) sessions, between 1 and 5. The default is 5.
- **SSH Sessions**—Sets the maximum SSH sessions, between 1 and 5. The default is 5.
- **Telnet Sessions**—Sets the maximum Telnet sessions, between 1 and 5. The default is 5.
- **User Sessions**—Sets the maximum sessions per user, between 1 and 5. The default is 5.

Step 3 Click **Apply** to save the configuration changes.

Configure AAA for System Administrators

This section describes how to configure authentication, management authorization, and command authorization for system administrators.

Configure Management Authentication

Configure authentication for CLI and ASDM access.

About Management Authentication

How you log into the ASA depends on whether or not you enable authentication.

About SSH Authentication

See the following behavior for SSH access with and without authentication:

- No Authentication—SSH is not available without authentication.
- Authentication—When you enable SSH authentication, you enter the username and password as defined on the AAA server or local user database. For public key authentication, the ASA only supports the local database. If you configure SSH public key authentication, then the ASA uses the local database implicitly. You only need to explicitly configure SSH authentication when you use a username and password to log in. You access user EXEC mode.

About Telnet Authentication

See the following behavior for Telnet access with and without authentication:

- No Authentication—If you do not enable any authentication for Telnet, you do not enter a username; you enter the login password. There is no default password, so you must set one before you can Telnet to the ASA. You access user EXEC mode.
- Authentication—If you enable Telnet authentication, you enter the username and password as defined on the AAA server or local user database. You access user EXEC mode.

About ASDM Authentication

See the following behavior for ASDM access with and without authentication. You can also configure certificate authentication, with or without AAA authentication.

- No Authentication—By default, you can log into ASDM with a blank username and the enable password, which is blank by default. We suggest that you change the enable password as soon as possible so that it does not remain blank; see Set the Hostname, Domain Name, and the Enable and Telnet Passwords, on page 693. When you enter the **enable** command at the CLI for the first time, you are prompted to change the password; this behavior is not enforced when you log into ASDM. Note that if you enter a username and password at the login screen (instead of leaving the username blank), ASDM checks the local database for a match.
- Certificate Authentication—(Single, routed mode only) You can require that the user have a valid certificate. Enter the certificate username and password, and the ASA validates the certificate against the PKI trustpoint.

- AAA Authentication—When you enable ASDM (HTTPS) authentication, you enter the username and password as defined on the AAA server or local user database. You can no longer use ASDM with a blank username and the enable password.
- AAA Authentication plus Certificate Authentication—(Single, routed mode only) When you enable ASDM (HTTPS) authentication, you enter the username and password as defined on the AAA server or local user database. If the username and password are different for the certificate authentication, you are prompted to enter them as well. You can opt to pre-fill the username derived from your certificate.

About Serial Authentication

See the following behavior for access to the serial console port with and without authentication:

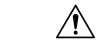
- No Authentication—If you do not enable any authentication for serial access, you do not enter a username or password. You access user EXEC mode.
- Authentication—If you enable authentication for serial access, you enter the username and password as defined on the AAA server or local user database. You access user EXEC mode.

About Enable Authentication

To enter privileged EXEC mode after logging in, enter the **enable** command. How this command works depends on whether or not you enable authentication:

- No Authentication—If you do not configure enable authentication, enter the system enable password when you enter the **enable** command, which is blank by default. The first time you enter the **enable** command, you are prompted to change it. However, if you do not use enable authentication, after you enter the **enable** command, you are no longer logged in as a particular user, which can affect user-based features such as command authorization. To maintain your username, use enable authentication.
- Authentication—If you configure enable authentication, the ASA prompts you for your username and password as defined on the AAA server or local user database. This feature is particularly useful when you perform command authorization, in which usernames are important in determining the commands that a user can enter.

For enable authentication using the local database, you can use the **login** command instead of the **enable** command. The **login** command maintains the username, but requires no configuration to turn on authentication.



Caution If you add users to the local database who can gain access to the CLI and whom you do not want to enter privileged EXEC mode, you should configure command authorization. Without command authorization, users can access privileged EXEC mode (and all commands) at the CLI using their own password if their privilege level is 2 or greater (2 is the default). Alternatively, you can discourage the login command by using a AAA server for authentication instead of the local database, or you can set all local users to level 1 so you can control who can use the system enable password to access privileged EXEC mode.

Sessions from the Host Operating System to the ASA

Some platforms support running the ASA as a separate application: for example the ASA on the Firepower 4100/9300. For sessions from the host operating system to the ASA, you can configure serial and Telnet authentication, depending on the type of connection.

In multiple context mode, you cannot configure any AAA commands in the system configuration. However, if you configure Telnet or serial authentication in the admin context, then authentication also applies to these sessions. The admin context AAA server or local user database is used in this instance.

Configure Authentication for CLI, ASDM, and enable command Access

Before you begin

- Configure Telnet, SSH, or HTTP access.
- For external authentication, configure a AAA server group. For local authentication, add users to the local database.
- HTTP management authentication does not support the SDI protocol for a AAA server group.
- This feature does not affect SSH public key authentication for local usernames with the **ssh authentication** command. The ASA implicitly uses the local database for public key authentication. This feature only affects usernames with passwords. If you want to allow either public key authentication or password use by a local user, then you need to explicitly configure local authentication with this procedure to allow password access.

Procedure

Step 1 To authenticate users who use the **enable** command, choose **Configuration** > **Device Management** > **Users/AAA** > **AAA Access** > **Authentication**, then configure the following settings:

a) Check the **Enable** check box.
b) Choose a server group name or the LOCAL database.
c) (Optional) If you chose a AAA server, you can configure the ASA to use the local database as a fallback method if the AAA server is unavailable. Check the **Use LOCAL when server group fails** check box. We recommend that you use the same username and password in the local database as the AAA server, because the ASA prompt does not give any indication of which method is being used.

Step 2 To authenticate users who access the CLI or ASDM, choose **Configuration** > **Device Management** > **Users/AAA** > **AAA Access** > **Authentication**, then configure the following settings:

a) Check one or more of the following check boxes:
   - **HTTP/ASDM**—Authenticates the ASDM client that accesses the ASA using HTTPS.
   - **Serial**—Authenticates users who access the ASA using the console port.
   - **SSH**—Authenticates users who access the ASA using SSH (password only; public key authentication implicitly uses the local database).
   - **Telnet**—Authenticates users who access the ASA using Telnet.

b) For each service that you checked, choose a server group name or the LOCAL database.
c) (Optional) If you chose a AAA server, you can configure the ASA to use the local database as a fallback method if the AAA server is unavailable. Check the **Use LOCAL when server group fails** check box. We recommend that you use the same username and password in the local database as the AAA server because the ASA prompt does not give any indication of which method is being used.

Step 3 Click **Apply**.

Configure ASDM Certificate Authentication

You can require certificate authentication, with or without AAA authentication. The ASA validates the certificate against the PKI trustpoint.

Before you begin

This feature is supported in single, routed mode only.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Management Access** > **ASDM/HTTPS/Telnet/SSH**.

Step 2 In the **Specify the interface requires client certificate to access ASDM** area, click **Add** to specify the interface and an optional certificate map that must be matched for successful authentication.

You configure certificate authentication for each interface, so that connections on a trusted/inside interface do not have to provide a certificate. See **Configuration** > **Site-to-Site VPN** > **Advanced** > **IPSec** > **Certificate to Connection Map** > **Rules** to create the certificate map.

Step 3 (Optional) To set the attribute used by ASDM to derive the username from the certificate, choose **Configuration** > **Device Management** > **Management Access** > **HTTP Certificate Rule**.

Choose one of the following methods:

- **Specify the Certificate Fields to be used**—Select a value from the **Primary Field** and the **Secondary Field** drop-down lists.
- **Use the entire DN as the username**
- **Use script to select username**—Click **Add** to add the script content.

Check the **Prefill Username** check box to pre-fill the username when prompted for authentication. If the username is different from the one you initially typed in, a new dialog box appears with the username pre-filled. You can then enter the password for authentication.

By default, ASDM uses CN OU attributes.

Step 4 Click **Apply**.

Control CLI and ASDM Access with Management Authorization

The ASA lets you distinguish between administrative and remote-access users when they authenticate. User role differentiation can prevent remote access VPN and network access users from establishing an administrative connection to the ASA.

Before you begin

RADIUS or LDAP (mapped) users

When users are authenticated through LDAP, the native LDAP attributes and their values can be mapped to ASA attributes to provide specific authorization features. Configure Cisco VSA CVPN3000-Privilege-Level with a value between 0 and 15. and then map the LDAP attributes to Cisco VAS CVPN3000-Privilege-Level.

The RADIUS IETF **service-type** attribute, when sent in an access-accept message as the result of a RADIUS authentication and authorization request, is used to designate which type of service is granted to the authenticated user

The RADIUS Cisco VSA **privilege-level** attribute (Vendor ID 3076, sub-ID 220), when sent in an access-accept message, is used to designate the level of privilege for the user.

TACACS+ users

Authorization is requested with "service=shell," and the server responds with PASS or FAIL.

Local users

Configure the **Access Restriction** option for a given username. By default, the access restriction is **Full Access**, which allows full access to any services specified by the **Authentication** tab options.

Management Authorization Attributes

See the following table for AAA server types and valid values for management authorization. The ASA uses these values to determine the level of management access.

Management Level	RADIUS/LDAP (Mapped) Attributes	TACACS+ Attributes	Local Database Attributes
Full Access—Allows full access to any services specified by the **Authentication** tab options	Service-Type 6 (Administrative), Privilege-Level 1	PASS, privilege level 1	admin
Partial Access—Allows access to the CLI or ASDM when you configure the **Authentication** tab options. However, if you configure **enable** authentication with the **Enable** option, then the CLI user cannot access privileged EXEC mode using the **enable** command.	Service-Type 7 (NAS prompt), Privilege-Level 2 and higher The Framed (2) and Login (1) service types are treated the same way.	PASS, privilege level 2 and higher	nas-prompt
No Access—Denies management access. The user cannot use any services specified by the **Authentication** tab options (excluding the **Serial** option; serial access is allowed). Remote access (IPsec and SSL) users can still authenticate and terminate their remote access sessions. All other service types (Voice, FAX, and so on) are treated the same way.	Service-Type 5 (Outbound)	FAIL	remote-access

Additional Guidelines

- Serial console access is not included in management authorization.

- You must also configure AAA authentication for management access to use this feature. See Configure Authentication for CLI, ASDM, and enable command Access, on page 1045.
- If you use external authentication, you must pre-configure a AAA server group before you enable this feature.
- HTTP authorization is supported in single, routed mode only.

Procedure

Step 1 To enable management authorization for HTTP sessions, choose **Configuration** > **Device Management** > **Users/AAA** > **AAA Access** > **Authorization**, and check the **HTTP** check box in the **Enable Authorization for ASA Command Access** Area.

Note
To configure ASA Command Access, see Configure Local Command Authorization, on page 1050.

Step 2 To enable management authorization for Telnet and SSH sessions, choose **Configuration** > **Device Management** > **Users/AAA** > **AAA Access** > **Authorization**, and check the **Enable** check box in the **Perform authorization for exec shell access** Area.

Step 3 Select either the **Remote** or **Local** radio buttons to specify the server to be used for authorization of exec shell access.

Step 4 To enable management authorization, check the **Allow privileged users to enter into EXEC mode on login** check box.

The **auto-enable** option allows users Full Access to be placed directly in privileged EXEC mode. Otherwise, users are placed in user EXEC mode.

Configure Command Authorization

If you want to control access to commands, the ASA lets you configure command authorization, where you can determine which commands that are available to a user. By default when you log in, you can access user EXEC mode, which offers only minimal commands. When you enter the **enable** command (or the **login** command when you use the local database), you can access privileged EXEC mode and advanced commands, including configuration commands.

You can use one of two command authorization methods:

- Local privilege levels
- TACACS+ server privilege levels

About Command Authorization

You can enable command authorization so only authorized users can enter commands.

Supported Command Authorization Methods

You can use one of two command authorization methods:

- Local privilege levels—Configure the command privilege levels on the ASA. When a local, RADIUS, or LDAP (if you map LDAP attributes to RADIUS attributes) user authenticates for CLI access, the ASA places that user in the privilege level that is defined by the local database, RADIUS, or LDAP server. The user can access commands at the assigned privilege level and below. Note that all users access user EXEC mode when they first log in (commands at level 0 or 1). The user needs to authenticate again with the **enable** command to access privileged EXEC mode (commands at level 2 or higher), or they can log in with the **login** command (local database only).



Note You can use local command authorization without any users in the local database and without CLI or **enable** authentication. Instead, when you enter the **enable** command, you enter the system enable password, and the ASA places you in level 15. You can then create enable passwords for every level, so that when you enter **enable** *n* (2 to 15), the ASA places you in level *n*. These levels are not used unless you enable local command authorization.

- TACACS+ server privilege levels—On the TACACS+ server, configure the commands that a user or group can use after authenticating for CLI access. Every command that a user enters at the CLI is validated with the TACACS+ server.

Security Contexts and Command Authorization

AAA settings are discrete per context, not shared among contexts.

When configuring command authorization, you must configure each security context separately. This configuration provides you the opportunity to enforce different command authorizations for different security contexts.

When switching between security contexts, administrators should be aware that the commands permitted for the username specified when they login may be different in the new context session or that command authorization may not be configured at all in the new context. Failure to understand that command authorizations may differ between security contexts could confuse an administrator.



Note The system execution space does not support AAA commands; therefore, command authorization is not available in the system execution space.

Command Privilege Levels

By default, the following commands are assigned to privilege level 0. All other commands are assigned to privilege level 15.

- **show checksum**
- **show curpriv**
- **enable**
- **help**
- **show history**
- **login**

- **logout**
- **pager**
- **show pager**
- **clear pager**
- **quit**
- **show version**

If you move any configure mode commands to a lower level than 15, be sure to move the **configure** command to that level as well, otherwise, the user cannot enter configuration mode.

Configure Local Command Authorization

Local command authorization lets you assign commands to one of 16 privilege levels (0 to 15). By default, each command is assigned either to privilege level 0 or 15. You can define each user to be at a specific privilege level, and each user can enter any command at the assigned privilege level or below. The ASA supports user privilege levels defined in the local database, a RADIUS server, or an LDAP server (if you map LDAP attributes to RADIUS attributes).

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Users/AAA** > **AAA Access** > **Authorization**.

Step 2 Check the **Enable authorization for ASA command access** > **Enable** check box.

Step 3 Choose **LOCAL** from the **Server Group** drop-down list.

Step 4 When you enable local command authorization, you have the option of manually assigning privilege levels to individual commands or groups of commands or enabling the predefined user account privileges.

- Click **Set ASDM Defined User Roles** to use predefined user account privileges.

  The **ASDM Defined User Roles Setup** dialog box appears. Click **Yes** to use the predefined user account privileges: **Admin** (privilege level 15, with full access to all CLI commands; **Read Only** (privilege level 5, with read-only access); and **Monitor Only** (privilege level 3, with access to the **Monitoring** section only).

- Click **Configure Command Privileges** to manually configure command levels.

  The **Command Privileges Setup** dialog box appears. You can view all commands by choosing **All Modes** from the **Command Mode** drop-down list, or you can choose a configuration mode to view the commands available in that mode. For example, if you choose context, you can view all commands available in context configuration mode. If a command can be entered in user EXEC or privileged EXEC mode as well as configuration mode, and the command performs different actions in each mode, you can set the privilege level for these modes separately.

  The **Variant** column displays show, clear, or cmd. You can set the privilege only for the show, clear, or configure form of the command. The configure form of the command is typically the form that causes a configuration change, either as the unmodified command (without the show or clear prefix) or as the no form.

  To change the level of a command, double-click it or click **Edit**. You can set the level between 0 and 15. You can only configure the privilege level of the main command. For example, you can configure

the level of all **aaa** commands, but not the level of the **aaa authentication** command and the **aaa authorization** command separately.

To change the level of all commands that appear, click **Select All**, then **Edit**.

Click **OK** to accept your changes.

Step 5 (Optional) Check the **Perform authorization for exec shell access** > **Enable** check box to enable AAA users for command authorization. Without this option, the ASA only supports privilege levels for local database users and defaults all other types of users to level 15.

This command also enables management authorization. See Control CLI and ASDM Access with Management Authorization, on page 1046.

Step 6 Click **Apply**.

The authorization settings are assigned, and the changes are saved to the running configuration.

Configure Commands on the TACACS+ Server

You can configure commands on a Cisco Secure Access Control Server (ACS) TACACS+ server as a shared profile component, for a group, or for individual users. For third-party TACACS+ servers, see your server documentation for more information about command authorization support.

See the following guidelines for configuring commands in Cisco Secure ACS Version 3.1; many of these guidelines also apply to third-party servers:

- The ASA sends the commands to be authorized as shell commands, so configure the commands on the TACACS+ server as shell commands.

Note Cisco Secure ACS might include a command type called “pix-shell.” Do not use this type for ASA command authorization.

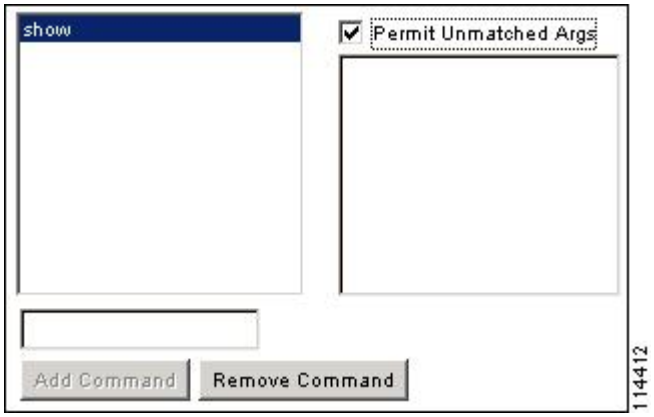
- The first word of the command is considered to be the main command. All additional words are considered to be arguments, which need to be preceded by **permit** or **deny**.

  For example, to allow the **show running-configuration aaa-server** command, add **show running-configuration** to the command field, and type **permit aaa-server** in the arguments field.

- You can permit all arguments of a command that you do not explicitly deny by checking the **Permit Unmatched Args** check box.

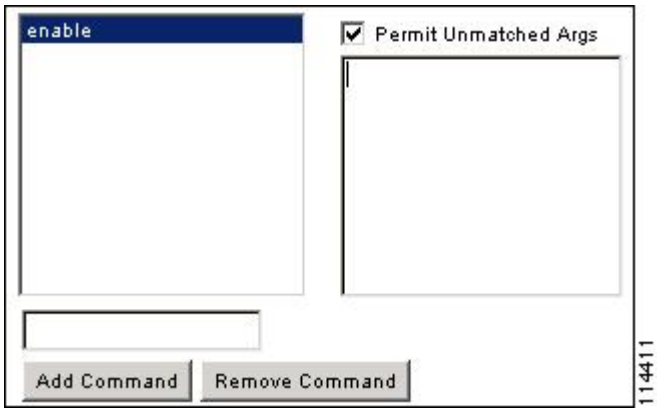
  For example, you can configure just the **show** command, then all the **show** commands are allowed. We recommend using this method so that you do not have to anticipate every variant of a command, including abbreviations and a question mark, which shows CLI usage (see the following figure).

Figure 97: Permitting All Related Commands



- For commands that are a single word, you *must* permit unmatched arguments, even if there are no arguments for the command, for example **enable** or **help** (see the following figure).

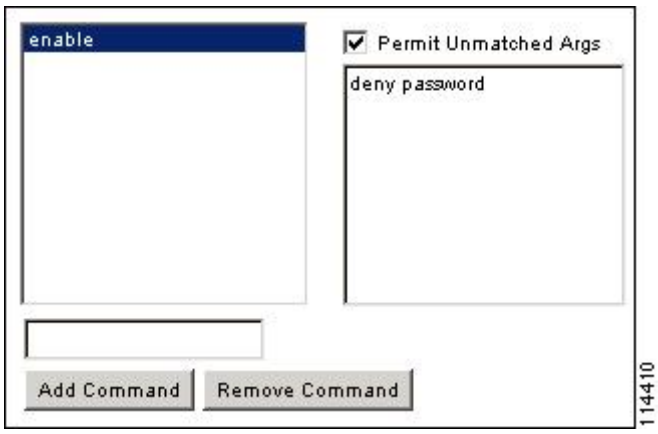
Figure 98: Permitting Single Word Commands



- To disallow some arguments, enter the arguments preceded by **deny**.

  For example, to allow **enable**, but not **enable password**, enter **enable** in the commands field, and **deny password** in the arguments field. Be sure to check the **Permit Unmatched Args** check box so that **enable** alone is still allowed (see the following figure).

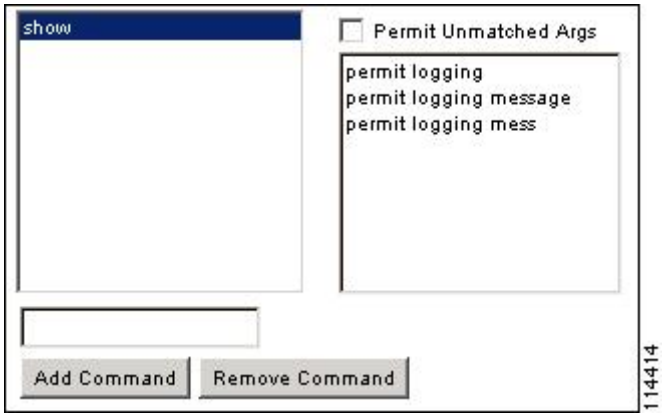
Figure 99: Disallowing Arguments



- When you abbreviate a command at the command line, the ASA expands the prefix and main command to the full text, but it sends additional arguments to the TACACS+ server as you enter them.

  For example, if you enter **sh log**, then the ASA sends the entire command to the TACACS+ server, **show logging**. However, if you enter **sh log mess**, then the ASA sends **show logging mess** to the TACACS+ server, and not the expanded command **show logging message**. You can configure multiple spellings of the same argument to anticipate abbreviations (see the following figure).

Figure 100: Specifying Abbreviations



- We recommend that you allow the following basic commands for all users:
  - **show checksum**
  - **show curpriv**
  - **enable**
  - **help**
  - **show history**
  - **login**
  - **logout**

- **pager**
- **show pager**
- **clear pager**
- **quit**
- **show version**

Configure TACACS+ Command Authorization

If you enable TACACS+ command authorization, and a user enters a command at the CLI, the ASA sends the command and username to the TACACS+ server to determine if the command is authorized.

Before you enable TACACS+ command authorization, be sure that you are logged into the ASA as a user that is defined on the TACACS+ server, and that you have the necessary command authorization to continue configuring the ASA. For example, you should log in as an admin user with all commands authorized. Otherwise, you could become unintentionally locked out.

Do not save your configuration until you are sure that it works the way you want. If you get locked out because of a mistake, you can usually recover access by restarting the ASA.

Be sure that your TACACS+ system is completely stable and reliable. The necessary level of reliability typically requires that you have a fully redundant TACACS+ server system and fully redundant connectivity to the ASA. For example, in your TACACS+ server pool, include one server connected to interface 1, and another to interface 2. You can also configure local command authorization as a fallback method if the TACACS+ server is unavailable.

To configure command authorization using a TACACS+ server, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Users/AAA** > **AAA Access** > **Authorization**.

Step 2 Check the **Enable authorization for command access** > **Enable** check box.

Step 3 Choose a AAA server group name from the **Server Group** drop-down list.

Step 4 (Optional) You can configure the ASA to use the local database as a fallback method if the AAA server is unavailable. To do so, check the **Use LOCAL when server group fails** check box. We recommend that you use the same username and password in the local database as the AAA server, because the ASA prompt does not give any indication which method is being used. Be sure to configure users in the local database and command privilege levels.

Step 5 Click **Apply**.

The command authorization settings are assigned, and the changes are saved to the running configuration.

Configure a Password Policy for Local Database Users

When you configure authentication for CLI or ASDM access using the local database, you can configure a password policy that requires a user to change their password after a specified amount of time and also requires password standards such as a minimum length and the minimum number of changed characters.

The password policy only applies to administrative users using the local database, and not to other types of traffic that can use the local database, such as VPN or AAA for network access, and not to users authenticated by a AAA server.

After you configure the password policy, when you change a password (either your own or another user's), the password policy applies to the new password. Any existing passwords are grandfathered in. The new policy applies to changing the password with the **User Accounts** pane as well as the **Change My Password** pane.

Before you begin

- Configure AAA authentication for CLI or ASDM access using the local database.
- Specify usernames in the local database.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Users/AAA** > **Password Policy**.

Step 2 Configure any mix of the following options:

- **Minimum Password Length**—Enter the minimum length for passwords. Valid values range from 3 to 64 characters. The recommended minimum password length is 8 characters.
- **Lifetime**—Enter the interval in days after which passwords expire for remote users (SSH, Telnet, HTTP); users at the console port are never locked out due to password expiration. Valid values are between 0 and 65536 days. The default value is 0 days, a value indicating that passwords will never expire.

  7 days before the password expires, a warning message appears. After the password expires, system access is denied to remote users. To gain access after expiration, do one of the following:

  - Have another administrator change your password.
  - Log in to the physical console port to change your password.

- **Minimum Number Of**—Specify the minimum of characters from the following types:
  - **Numeric Characters**—Enter the minimum number of numeric characters that passwords must have. Valid values are between 0 and 64 characters. The default value is 0.
  - **Lower Case Characters**—Enter the minimum number of lower case characters that passwords must have. Valid values range from 0 to 64 characters. The default value is 0.
  - **Upper Case Characters**—Enter the minimum number of upper case characters that passwords must have. Valid values range from 0 to 64 characters. The default value is 0.
  - **Special Characters**—Enter the minimum number of special characters that passwords must have. Valid values range from 0 to 64 characters. Special characters include the following: !, @, #, $, %, ^, &, *, '(‘ and ‘)’.The default value is 0.

- **Different Characters from Previous Password**—Enter the minimum number of characters that you must change between new and old passwords. Valid values are between 0 and 64 characters. The default value is 0. Character matching is position independent, meaning that new password characters are considered changed only if they do not appear anywhere in the current password.

- **Enable Reuse Interval**—You can prohibit the reuse of a password that matches previously used passwords, between 2 and 7 previous passwords. The previous passwords are stored in the configuration under each username in encrypted form using the **password-history** command; this command is not user-configurable.
- **Prevent Passwords from Matching Usernames**—Prohibit a password that matches a username.

Step 3 (Optional) Check the **Enable Password and Account Protection** check box to require users to change their password on the **Change My Password** pane instead of the **User Accounts** pane. The default setting is disabled: a user can use either method to change their password.

If you enable this feature and try to change your password on the **User Accounts** pane, the following error message is generated:

```
ERROR: Changing your own password is prohibited
```

Step 4 Click **Apply** to save the configuration settings.

Change Your Password

If you configure a password lifetime in the password policy, you need to change your password to a new one when the old password expires. This password change method is required if you enable password policy authentication. If password policy authentication is not enabled, then you can use this method, or you can change your user account directly.

To change your username password, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Users/AAA** > **Change Password**.
Step 2 Enter your old password.
Step 3 Enter your new password.
Step 4 Confirm your new password.
Step 5 Click **Make Change**.
Step 6 Click the **Save** icon to save your changes to the running configuration.

Enable and View the Login History

By default, the login history is saved for 90 days. You can disable this feature or change the duration, up to 365 days.

Before you begin

- The login history is only saved per unit; in failover and clustering environments, each unit maintains its own login history only.
- Login history data is not maintained over reloads.
- This feature applies to usernames in the local database or from a AAA server when you enable local AAA authentication for one or more of the CLI management methods (SSH, Telnet, serial console). ASDM logins are not saved in the history.

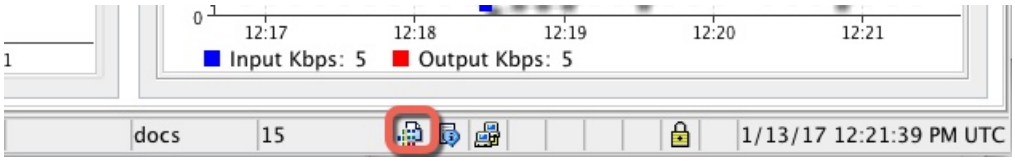
Procedure

Step 1 Choose **Configuration** > **Device Management** > **Users/AAA** > **Login History**.

Step 2 Check the **Configure login history reporting for administrators** check box. This feature is enabled by default.

Step 3 Set the **Duration** between 1 and 365 days. The default is 90.

Step 4 To view the login history, from any ASDM screen you can click on the **Login History** icon in the bottom **Status** bar:



The login history for all users displays in a dialog box.

Configure Management Access Accounting

You can send accounting messages to the TACACS+ accounting server when you enter any command other than **show** commands at the CLI. You can configure accounting when users log in, when they enter the **enable** command, or when they issue commands.

For command accounting, you can only use TACACS+ servers.

To configure management access and enable command accounting, perform the following steps:

Procedure

Step 1 To enable accounting of users when they enter the **enable** command, perform the following steps:

a) Choose **Configuration** > **Device Management** > **Users/AAA** > **AAA Access** > **Accounting** , then check the **Require accounting to allow accounting of user activity** > **Enable** check box.

b) Choose a RADIUS or TACACS+ server group name.

Step 2 To enable accounting of users when they access the ASA using Telnet, SSH, or the serial console, perform the following steps:

a) Check the **Serial**, **SSH**, and/or **Telnet** check boxes in the **Require accounting for the following types of connections** area.
b) Choose a RADIUS or TACACS+ server group name for each connection type.

Step 3 To configure command accounting, perform the following steps:

a) Check the **Enable** check box in the **Require accounting for the following types of connections** area.
b) Choose a TACACS+ server group name. RADIUS is not supported.

You can send accounting messages to the TACACS+ accounting server when you enter any command other than **show** commands at the CLI.

c) If you customize the command privilege level using the **Command Privilege Setup** dialog box, you can limit which commands the ASA accounts for by specifying a minimum privilege level in the **Privilege level** drop-down list. The ASA does not account for commands that are below the minimum privilege level.

Step 4 Click **Apply**.

The accounting settings are assigned, and the changes are saved to the running configuration.

Recover from a Lockout

In some circumstances, when you turn on command authorization or CLI authentication, you can be locked out of the ASA CLI. You can usually recover access by restarting the ASA. However, if you already saved your configuration, you might be locked out.

The following table lists the common lockout conditions and how you might recover from them.

Table 55: CLI Authentication and Command Authorization Lockout Scenarios

Feature	Lockout Condition	Description	Workaround: Single Mode	Workaround: Multiple Mode
Local CLI authentication	No users have been configured in the local database.	If you have no users in the local database, you cannot log in, and you cannot add any users.	Log in and reset the passwords and **aaa** commands.	Session into the ASA from the switch. From the system execution space, you can change to the context and add a user.

Feature	Lockout Condition	Description	Workaround: Single Mode	Workaround: Multiple Mode
TACACS+ command authorization TACACS+ CLI authentication RADIUS CLI authentication	The server is down or unreachable and you do not have the fallback method configured.	If the server is unreachable, then you cannot log in or enter any commands.	1. Log in and reset the passwords and AAA commands. 2. Configure the local database as a fallback method so you do not get locked out when the server is down.	1. If the server is unreachable because the network configuration is incorrect on the ASA, session into the ASA from the switch. From the system execution space, you can change to the context and reconfigure your network settings. 2. Configure the local database as a fallback method so that you do not get locked out when the server is down.
TACACS+ command authorization	You are logged in as a user without enough privileges or as a user that does not exist.	You enable command authorization, but then find that the user cannot enter any more commands.	Fix the TACACS+ server user account. If you do not have access to the TACACS+ server and you need to configure the ASA immediately, then log into the maintenance partition and reset the passwords and **aaa** commands.	Session into the ASA from the switch. From the system execution space, you can change to the context and complete the configuration changes. You can also disable command authorization until you fix the TACACS+ configuration.
Local command authorization	You are logged in as a user without enough privileges.	You enable command authorization, but then find that the user cannot enter any more commands.	Log in and reset the passwords and **aaa** commands.	Session into the ASA from the switch. From the system execution space, you can change to the context and change the user level.

Monitoring Device Access

- **Monitoring** > **Properties** > **Device Access** > **ASDM/HTTPS/Telnet/SSH Sessions**

  The top pane lists the connection types, session IDs, and IP addresses for users connected through ASDM, HTTPS, and Telnet sessions. To disconnect a specific session, click **Disconnect**.

  The bottom pane lists the clients, usernames, connection states, software versions, incoming encryption types, outgoing encryption types, incoming HMACs, outgoing HMACs, SSH session IDs, remaining rekey data, remaining rekey time, data-based rekeys, time-based rekeys, and the last rekey time. To disconnect a specific session, click **Disconnect**.

- **Monitoring** > **Properties** > **Device Access** > **Authenticated Users**

This pane lists the usernames, IP addresses, dynamic ACLs, inactivity timeouts (if any), and absolute timeouts for users who were authenticated by AAA servers.

- **Monitoring** > **Properties** > **Device Access** > **AAA Locked Out Users**

  This pane lists the usernames of locked-out AAA local users, the number of failed attempts to authenticate, and the times that users were locked out. To clear a specific user who has been locked out, click **Clear Selected Lockout**. To clear all users who have been locked out, click **Clear All Lockouts**.

- **Tools > Command Line Interface**

  This pane allows you to issue various non-interactive commands and view results.

History for Management Access

Table 56: History for Management Access

Feature Name	Platform Releases	Description
The ASA SSH stack was deprecated	9.23(1)	You can no longer use the ASA SSH stack. The Cisco SSH stack is now the only stack. New/Modified screens: • Single context mode: **Configuration** > **Device Management** > **Management Access** > **ASDM/HTTPS/Telnet/SSH** • Multiple context mode: **Configuration** > **Device Management** > **SSH Stack**
CiscoSSH stack now default	9.19(1)	The Cisco SSH stack is now used by default. New/Modified screens: • Single context mode: **Configuration** > **Device Management** > **Management Access** > **ASDM/HTTPS/Telnet/SSH** • Multiple context mode: **Configuration** > **Device Management** > **SSH Stack**
Loopback interface support for SSH and Telnet	9.18(2)	You can now add a loopback interface and use it for the following features: • SSH • Telnet New/Modified commands: **interface loopback**, **ssh**, **telnet** New/Modified screens: **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces** > **Add Loopback Interface** ASDM support was added in 7.19.

Feature Name	Platform Releases	Description
CiscoSSH stack	9.17(1)	The ASA uses a proprietary SSH stack for SSH connections. You can now choose to use the CiscoSSH stack instead, which is based on OpenSSH. The default stack continues to be the ASA stack. Cisco SSH supports: • FIPS compliance • Regular updates, including updates from Cisco and the open source community Note that the CiscoSSH stack does not support: • SSH to a different interface over VPN (management-access) • EdDSA key pair • RSA key pair in FIPS mode If you need these features, you should continue to use the ASA SSH stack. There is a small change to SCP functionality with the CiscoSSH stack: to use the ASA **copy** command to copy a file to or from an SCP server, you have to enable SSH access on the ASA for the SCP server subnet/host. New/Modified screens: • Single context mode: **Configuration** > **Device Management** > **Management Access** > **ASDM/HTTPS/Telnet/SSH** • Multiple context mode: **Configuration** > **Device Management** > **SSH Stack**
Local user lockout changes	9.17(1)	The ASA can lock out local users after a configurable number of failed login attempts. This feature did not apply to users with privilege level 15. Also, a user would be locked out indefinitely until an admin unlocked their account. Now, users will be unlocked after 10 minutes unless an admin uses the **clear aaa local user lockout** command before then. Privilege level 15 users are also now affected by the lockout setting. New/Modified commands: **aaa local authentication attempts max-fail** , **show aaa local user**
SSH and Telnet password change prompt	9.17(1)	The first time a local user logs into the ASA using SSH or Telnet, they are prompted to change their password. They will also be prompted for the first login after an admin changes their password. If the ASA reloads, however, users will not be prompted even if it is their first login. Note that any service that uses the local user database, such as VPN, will also have to use the new password if it was changed during an SSH or Telnet login. New/Modified commands: **show aaa local user**

Feature Name	Platform Releases	Description
SSH security improvements	9.16(1)	SSH now supports the following security improvements: • Host key format—**crypto key generate** {**eddsa** | **ecdsa**}. In addition to RSA, we added support for the EdDSA and ECDSA host keys. The ASA tries to use keys in the following order if they exist: EdDSA, ECDSA, and then RSA. If you explicitly configure the ASA to use the RSA key with the **ssh key-exchange hostkey rsa** command, you must generate a key that is 2048 bits or higher. For upgrade compatibility, the ASA will use smaller RSA host keys only when the default host key setting is used. RSA support will be removed in a later release. • Key exchange algorithms—**ssh key-exchange group** {**ecdh-sha2-nistp256** | **curve25519-sha256**} • Encryption algorithms—**ssh cipher encryption chacha20-poly1305@openssh.com** • SSH version 1 is no longer supported—The **ssh version** command is removed. New/Modified screens: • **Configuration** > **Device Management** > **Management Access** > **ASDM/HTTPS/Telnet/SSH** • **Configuration** > **Device Management** > **Certificate Management** > **Identity Certificates** • **Configuration** > **Device Management** > **Advanced** > **SSH Ciphers**
Management access for SNMP	9.14(2)	When configuring management access over a VPN tunnel, include the IP address of the outside interface in the crypto map access-list as part of the VPN configuration for secure SNMP polling over a site-to-site VPN.
HTTPS idle timeout setting	9.14(1)	You can now set the idle timeout for all HTTPS connections to the ASA, including ASDM, WebVPN, and other clients. Formerly, using the **http server idle-timeout** command, you could only set the ASDM idle timeout. If you set both timeouts, the new command takes precendence. New/Modified screen: **Configuration** > **Device Management** > **Management Access** > **ASDM/HTTPS/Telnet/SSH** > **HTTP Settings** > **Connection Idle Timeout** check box.
SSH encryption ciphers are now listed in order from highest to lowest security for pre-defined lists	9.13(1)	SSH encryption ciphers are now listed in order from highest security to lowest security for pre-defined lists (such as medium or high). In earlier releases, they were listed from lowest to highest, which meant that a low security cipher would be proposed before a high security cipher. New/Modified screens: **Configuration** > **Device Management** > **Advanced** > **SSH Ciphers**

Feature Name	Platform Releases	Description
Setting the SSH key exchange mode is restricted to the Admin context	9.12(2)	You must set the SSH key exchange in the Admin context; this setting is inherited by all other contexts. New/Modified screen: **Configuration** > **Device Management** > **Management Access** > **ASDM/HTTPS/Telnet/SSH** > **SSH Settings** > **DH Key Exchange**
enable password change now required on login	9.12(1)	The default **enable** password is blank. When you try to access privileged EXEC mode on the ASA, you are now required to change the password to a value of 3 characters or longer. You cannot keep it blank. The **no enable password** command is no longer supported. At the CLI, you can access privileged EXEC mode using the **enable** command, the **login** command (with a user at privilege level 2+), or an SSH or Telnet session when you enable **aaa authorization exec auto-enable**. All of these methods require you to set the enable password. This password change requirement is not enforced for ASDM logins. In ASDM, by default you can log in without a username and with the **enable** password. No modified screens.
Configurable limitation of admin sessions	9.12(1)	You can configure the maximum number of aggregate, per user, and per-protocol administrative sessions. Formerly, you could configure only the aggregate number of sessions. This feature does not affect console sessions. Note that in multiple context mode, you cannot configure the number of HTTPS sessions, where the maximum is fixed at 5 sessions. The **quota management-session** command is also no longer accepted in the system configuration, and is instead available in the context configuration. The maximum aggregate sessions is now 15; if you configured 0 (unlimited) or 16+, then when you upgrade, the value is changed to 15. New/Modified screens: **Configuration** > **Device Management** > **Management Access** > **Management Session Quota**
Notifications for administrative privilege level changes	9.12(1)	When you authenticate for enable access (**aaa authentication enable console**) or allow privileged EXEC access directly (**aaa authorization exec auto-enable**), then the ASA now notifies users if their assigned access level has changed since their last login. New/Modified screens: **Status** bar **>** **Login History** icon

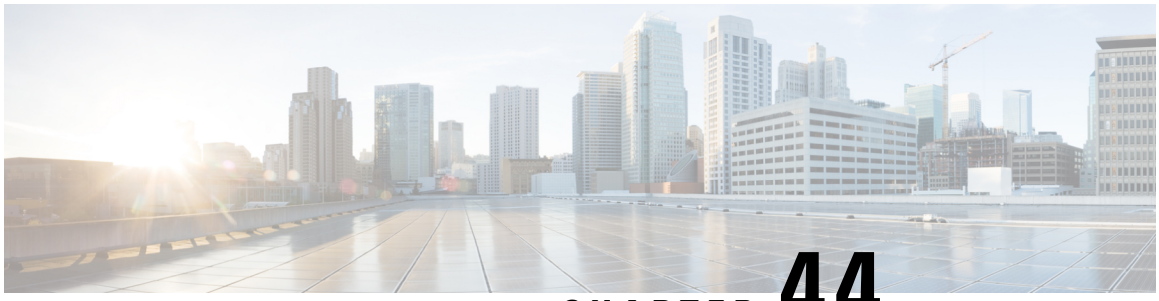
Feature Name	Platform Releases	Description
SSH stronger security	9.12(1)	See the following SSH security improvements: • Diffie-Hellman Group 14 SHA256 key exchange support. This setting is now the default. The former default was Group 1 SHA1. • HMAC-SHA256 integrity cipher support. The default is now the high security set of ciphers (hmac-sha2-256 only). The former default was the medium set. New/Modified screens: • **Configuration** > **Device Management** > **Management Access** > **ASDM/HTTPS/Telnet/SSH** • **Configuration** > **Device Management** > **Advanced** > **SSH Ciphers**
Allow non-browser-based HTTPS clients to access the ASA	9.12(1)	You can allow non-browser-based HTTPS clients to access HTTPS services on the ASA. By default, ASDM, CSM, and REST API are allowed. New/Modified screens. **Configuration** > **Device Management** > **Management Access** > **HTTP Non-Browser Client Support**
RSA key pair supports 3072-bit keys	9.9(2)	You can now set the modulus size to 3072. New or modified screen: **Configuration** > **Device Management** > **Certificate Management** > **Identity Certificates**
VPN management access on Bridged Virtual Interfaces (BVIs)	9.9(2)	You can now enable management services, such as **telnet**, **http**, and **ssh**, on a BVI if VPN **management-access** has been enabled on that BVI. For non-VPN management access, you should continue to configure these services on the bridge group member interfaces. New or Modified commands: **https**, **telnet**, **ssh**, **management-access**
SSH version 1 has been deprecated	9.9(1)	SSH version 1 has been deprecated, and will be removed in a future release. The default setting has changed from both SSH v1 and v2 to just SSH v2. New/Modified screens: • **Configuration** > **Device Management** > **Management Access** > **ASDM/HTTPS/Telnet/SSH**

Feature Name	Platform Releases	Description
Separate authentication for users with SSH public key authentication and users with passwords	9.6(3)/9.8(1)	In releases prior to 9.6(2), you could enable SSH public key authentication (**ssh authentication**) without also explicitly enabling AAA SSH authentication with the Local user database (**aaa authentication ssh console LOCAL**). In 9.6(2), the ASA required you to explicitly enable AAA SSH authentication. In this release, you no longer have to explicitly enable AAA SSH authentication; when you configure the **ssh authentication** command for a user, local authentication is enabled by default for users with this type of authentication. Moreover, when you explicitly configure AAA SSH authentication, this configuration only applies for for usernames with *passwords*, and you can use any AAA server type (**aaa authentication ssh console radius_1**, for example). For example, some users can use public key authentication using the local database, and other users can use passwords with RADIUS. We did not modify any screens.
Login history	9.8(1)	By default, the login history is saved for 90 days. You can disable this feature or change the duration, up to 365 days. This feature only applies to usernames in the local database when you enable local AAA authentication for one or more of the management methods (SSH, ASDM, Telnet, and so on). We introduced the following screen: **Configuration** > **Device Management** > **Users/AAA** > **Login History**
Password policy enforcement to prohibit the reuse of passwords, and prohibit use of a password matching a username	9.8(1)	You can now prohibit the reuse of previous passwords for up to 7 generations, and you can also prohibit the use of a password that matches a username. We modified the following screen: **Configuration** > **Device Management** > **Users/AAA** > **Password Policy**
ASA SSL Server mode matching for ASDM	9.6(2)	For an ASDM user who authenticates with a certificate, you can now require the certificate to match a certificate map. We modified the following screen: **Configuration** > **Device Management** > **Management Access** > **ASDM/HTTPS/Telnet/SSH**
SSH public key authentication improvements	9.6(2)	In earlier releases, you could enable SSH public key authentication without also enabling AAA SSH authentication with the Local user database . The configuration is now fixed so that you must explicitly enable AAA SSH authentication. To disallow users from using a password instead of the private key, you can now create a username without any password defined. We modifed the following screens: **Configuration** > **Device Management** > **Management Access** > **ASDM/HTTPS/Telnet/SSH** **Configuration** > **Device Management** > **Users/AAA** > **User Accounts** > **Add/Edit User Account**

Feature Name	Platform Releases	Description
ASDM management authorization	9.4(1)	You can now configure management authorization separately for HTTP access vs. Telnet and SSH access. We modified the following screen: **Configuration** > **Device Management** > **Users/AAA** > **AAA Access** > **Authorization**
ASDM username from certificate configuration	9.4(1)	When you enable ASDM certificate authentication, you can configure how ASDM extracts the username from the certificate; you can also enable pre-filling the username at the login prompt. We introduced the following screen: **Configuration > Device Management > Management Access > HTTP Certificate Rule.**
Improved one-time password authentication	9.2(1)	Administrators who have sufficient authorization privileges may enter privileged EXEC mode by entering their authentication credentials once. The **auto-enable** option was added to the **aaa authorization exec** command. We modified the following screen: Configuration > Device Management > Users/AAA > AAA Access > Authorization.
HTTP redirect support for IPV6	9.1(7)/9.6(1)	When you enable HTTP redirect to HTTPS for ASDM access or clientless SSL VPN, you can now redirect traffic sent an to IPv6 address. We added functionality to the following screen: **Configuration** > **Device Management** > **HTTP Redirect**
Configurable SSH encryption and integrity ciphers	9.1(7.4)/9.3(3.6)(1)	Users can select cipher modes when doing SSH encryption management and can configure HMAC and encryption for varying key exchange algorithms. You might want to change the ciphers to be more or less strict, depending on your application. Note that the performance of secure copy depends partly on the encryption cipher used. By default, the ASA negotiates one of the following algorithms in order: 3des-cbc aes128-cbc aes192-cbc aes256-cbc aes128-ctr aes192-ctr aes256-ctr. If the first algorithm proposed (3des-cbc) is chosen, then the performance is much slower than a more efficient algorithm such as aes128-cbc. To change the proposed ciphers, use **ssh cipher encryption custom aes128-cbc**, for example. We introduced the following screen: **Configuration** > **Device Management** > **Advanced** > **SSH Ciphers**
AES-CTR encryption for SSH	9.1(2)	The SSH server implementation in the ASA now supports AES-CTR mode encryption.
Improved SSH rekey interval	9.1(2)	An SSH connection is rekeyed after 60 minutes of connection time or 1 GB of data traffic. .
For the ASASM in multiple context mode, support for Telnet and virtual console authentication from the switch.	8.5(1)	Although connecting to the ASASM from the switch in multiple context mode connects to the system execution space, you can configure authentication in the admin context to govern those connections.

Feature Name	Platform Releases	Description
Support for administrator password policy when using the local database	8.4(4.1), 9.1(2)	When you configure authentication for CLI or ASDM access using the local database, you can configure a password policy that requires a user to change their password after a specified amount of time and also requires password standards such as a minimum length and the minimum number of changed characters. We introduced the following screen: Configuration > Device Management > Users/AAA > Password Policy.
Support for SSH public key authentication	8.4(4.1), 9.1(2)	You can enable public key authentication for SSH connections to the ASA on a per-user basis. You can specify a public key file (PKF) formatted key or a Base64 key. The PKF key can be up to 4096 bits. Use PKF format for keys that are too large to for the ASA support of the Base64 format (up to 2048 bits). We introduced the following screens: Configuration > Device Management > Users/AAA > User Accounts > Edit User Account > Public Key Authentication Configuration > Device Management > Users/AAA > User Accounts > Edit User Account > Public Key Using PKF. *PKF key format support is only in 9.1(2) and later.*
Support for Diffie-Hellman Group 14 for the SSH Key Exchange	8.4(4.1), 9.1(2)	Support for Diffie-Hellman Group 14 for SSH Key Exchange was added. Formerly, only Group 1 was supported. We modified the following screen: Configuration > Device Management > Management Access > ASDM/HTTPS/Telnet/SSH.
Support for a maximum number of management sessions	8.4(4.1), 9.1(2)	You can set the maximum number of simultaneous ASDM, SSH, and Telnet sessions. We introduced the following screen: Configuration > Device Management > Management Access > Management Session Quota.
Increased SSH security; the SSH default username is no longer supported.	8.4(2)	Starting in 8.4(2), you can no longer connect to the ASA using SSH with the pix or asa username and the login password. To use SSH, you must configure AAA authentication using the **aaa authentication ssh console LOCAL** command (CLI) or Configuration > Device Management > Users/AAA > AAA Access > Authentication (ASDM); then define a local user by entering the **username** command (CLI) or choosing Configuration > Device Management > Users/AAA > User Accounts (ASDM). If you want to use a AAA server for authentication instead of the local database, we recommend also configuring local authentication as a backup method.

Feature Name	Platform Releases	Description
Management Access	7.0(1)	We introduced this feature. We introduced the following screens: Configuration > Device Management > Management Access > ASDM/HTTPS/Telnet/SSH Configuration > Device Management > Management Access > Command Line (CLI) > Banner Configuration > Device Management > Management Access > CLI Prompt Configuration > Device Management > Management Access > ICMP Configuration > Device Management > Management Access > File Access > FTP Client Configuration > Device Management > Management Access > File Access > Secure Copy (SCP) Server Configuration > Device Management > Management Access > File Access > Mount-Points Configuration > Device Management > Users/AAA > AAA Access > Authentication Configuration > Device Management > Users/AAA > AAA Access > Authorization Configuration > Device Management > Users/AAA > AAA Access > > Accounting.



CHAPTER 44

Software and Configurations

This chapter describes how to manage the ASA software and configurations.

- Upgrade the Software, on page 1069
- Load an Image Using ROMMON (ISA 3000), on page 1069
- Upgrade the ROMMON Image: ISA 3000, on page 1071
- Downgrade Your Software, on page 1072
- Manage Files, on page 1077
- Set the ASA Image, ASDM, and Startup Configuration, on page 1084
- Back Up and Restore Configurations or Other Files, on page 1086
- Schedule a System Restart, on page 1092
- Hot Swap an SSD on the Secure Firewall 3100/4200, on page 1093
- Disable the USB Port, on page 1095
- History for Software and Configurations, on page 1096

Upgrade the Software

See the Cisco ASA Upgrade Guide for full upgrade procedures.

Load an Image Using ROMMON (ISA 3000)

To load a software image onto an ASA from the ROMMON mode using TFTP, perform the following steps.

Procedure

Step 1 Connect to the ASA console port according to the instructions in Access the ISA 3000 Console, on page 17.

Step 2 Power off the ASA, then power it on.

Step 3 During startup, press the **Escape** key when you are prompted to enter ROMMON mode.

Step 4 In ROMMOM mode, define the interface settings to the ASA, including the IP address, TFTP server address, gateway address, software image file, and port, as follows:

```
rommon #1> interface gigabitethernet0/0
```

```
rommon #2> address 10.86.118.4
rommon #3> server 10.86.118.21
rommon #4> gateway 10.86.118.21
rommon #5> file asa961-smp-k8.bin
```

Note
Be sure that the connection to the network already exists.

The **interface** command is ignored on the ASA 5506-X, ASA 5508-X, and ASA 5516-X, and ISA 3000 platforms, and you must perform TFTP recovery on these platforms from the Management 1/1 interface.

Step 5 Validate your settings:

```
rommon #6> set
ROMMON Variable Settings:
  ADDRESS=10.86.118.3
  SERVER=10.86.118.21
  GATEWAY=10.86.118.21
  PORT=GigabitEthernet0/0
  VLAN=untagged
  IMAGE=asa961-smp-k8.bin
  CONFIG=
  LINKTIMEOUT=20
  PKTTIMEOUT=4
  RETRY=20
```

Step 6 Ping the TFTP server:

```
rommon #7> ping server
Sending 20, 100-byte ICMP Echoes to server 10.86.118.21, timeout is 4 seconds:

Success rate is 100 percent (20/20)
```

Step 7 Save the network settings for future use:

```
rommon #8> sync
Updating NVRAM Parameters...
```

Step 8 Load the software image:

```
rommon #9> tftpdnld
ROMMON Variable Settings:
  ADDRESS=10.86.118.3
  SERVER=10.86.118.21
  GATEWAY=10.86.118.21
  PORT=GigabitEthernet0/0
  VLAN=untagged
  IMAGE=asa961-smp-k8.bin
  CONFIG=
  LINKTIMEOUT=20
  PKTTIMEOUT=4
  RETRY=20

tftp asa961-smp-k8.bin@10.86.118.21 via 10.86.118.21

Received 14450688 bytes
```

```
Launching TFTP Image...
Cisco ASA Security Appliance admin loader (3.0) #0: Mon Mar 5 16:00:07 MST 2016

Loading...
```

After the software image is successfully loaded, the ASA automatically exits ROMMON mode.

Step 9 Booting the ASA from ROMMON mode does not preserve the system image across reloads; you must still download the image to flash memory. See the Cisco ASA Upgrade Guide for full upgrade procedures.

Upgrade the ROMMON Image: ISA 3000

Follow these steps to upgrade the ROMMON image for the ISA 3000. For the ASA models, the ROMMON version on your system must be 1.1.8 or greater. We recommend that you upgrade to the latest version.

You can only upgrade to a new version; you cannot downgrade.



Caution The ISA 3000 ROMMON upgrade for 1.0.5 takes twice as long as previous ROMMON versions, approximately 15 minutes. **Do not** power cycle the device during the upgrade. If the upgrade is not complete within 30 minutes or it fails, contact Cisco technical support; **do not** power cycle or reset the device.

Before you begin

Obtain the new ROMMON image from Cisco.com, and put it on a server to copy to the ASA. The ASA supports FTP, TFTP, SCP, HTTP(S), and SMB servers. Download the image from:

- ISA 3000: https://software.cisco.com/download/home/286288493/type

Procedure

Step 1 Copy the ROMMON image to the ASA flash memory. This procedure shows an FTP copy; enter **copy ?** for the syntax for other server types.

copy ftp://[*username***:***password***@**]*server_ip***/asa5500-firmware-***xxxx***.SPA disk0:asa5500-firmware-***xxxx***.SPA**

Step 2 To see your current version, enter the **show module** command and look at the Fw Version in the output for Mod 1 in the MAC Address Range table:

```
ciscoasa# show module
[...]
Mod  MAC Address Range                 Hw Version   Fw Version   Sw Version
---- --------------------------------- ------------ ------------ ---------------
   1 7426.aceb.ccea to 7426.aceb.ccf2  0.3          1.1.5        9.4(1)
 sfr 7426.aceb.cce9 to 7426.aceb.cce9  N/A          N/A
```

Step 3 Upgrade the ROMMON image:

upgrade rommon disk0:asa5500-firmware-*xxxx*.SPA

Example:

```
ciscoasa# upgrade rommon disk0:asa5500-firmware-1108.SPA
Verifying file integrity of disk0:/asa5500-firmware-1108.SPA

Computed Hash   SHA2: d824bdeecee1308fc64427367fa559e9
                      eefe8f182491652ee4c05e6e751f7a4f
                      5cdea28540cf60acde3ab9b65ff55a9f
                      4e0cfb84b9e2317a856580576612f4af

Embedded Hash   SHA2: d824bdeecee1308fc64427367fa559e9
                      eefe8f182491652ee4c05e6e751f7a4f
                      5cdea28540cf60acde3ab9b65ff55a9f
                      4e0cfb84b9e2317a856580576612f4af


Digital signature successfully validated
File Name                     : disk0:/asa5500-firmware-1108.SPA
Image type                    : Release
    Signer Information
        Common Name           : abraxas
        Organization Unit     : NCS_Kenton_ASA
        Organization Name     : CiscoSystems
    Certificate Serial Number : 553156F4
    Hash Algorithm            : SHA2 512
    Signature Algorithm       : 2048-bit RSA
    Key Version               : A
Verification successful.
Proceed with reload? [confirm]
```

Step 4 Confirm to reload the ASA when you are prompted.

The ASA upgrades the ROMMON image, and then reloads the operating system.

Downgrade Your Software

In many cases, you can downgrade your ASA software and restore a backup configuration from the previous software version. The method of downgrading depends on your ASA platform.

Guidelines and Limitations for Downgrading

See the following guidelines before downgrading:

- **There is no official Zero Downtime Downgrade support for clustering**—However, in some cases, Zero Downtime Downgrading will work. See the following known issues for downgrading; note that there may be other issues that require you to reload your cluster units, thus causing downtime.
  - **Downgrade to a pre-9.9(1) release with clustering**—9.9(1) and later includes an improvement in the backup distribution. If you have 3 or more units in the cluster, you must perform the following steps:
    1. Remove all secondary units from the cluster (so the cluster consists only of the primary unit).

2. Downgrade 1 secondary unit, and rejoin it to the cluster.
3. Disable clustering on the primary unit; downgrade it, and rejoin the cluster.
4. Downgrade the remaining secondary units, and join them back to the cluster, one at a time.

- **Downgrade to a pre-9.9(1) release when you enable cluster site redundancy**—You should disable site redundancy if you want to downgrade (or if you want to add a pre-9.9(1) unit to a cluster). Otherwise, you will see side effects, for example, dummy forwarding flows on the unit running the old version.
- **Downgrade from 9.8(1) with clustering and crypto-map**—There is no Zero Downtime Downgrade support when downgrading from 9.8(1) when you have a crypto-map configured. You should clear the crypto-map configuration before downgrading, and then re-apply the configuration after the downgrade.
- **Downgrade from 9.8(1) with clustering unit health check set to .3 to .7 seconds**—If you downgrade your ASA software after setting the hold time to .3 - .7 (**health-check holdtime**), this setting will revert to the default of 3 seconds because the new setting is unsupported.
- **Downgrade from 9.5(2) or later to 9.5(1) or earlier with clustering (CSCuv82933)**—There is no Zero Downtime Downgrade support when downgrading from 9.5(2). You must reload all units at roughly the same time so that a new cluster is formed when the units come back online. If you wait to reload the units sequentially, then they will be unable to form a cluster.
- **Downgrade from 9.2(1) or later to 9.1 or earlier with clustering**—Zero Downtime Downgrade is not supported.

- **Downgrade issue from 9.22 or later**—If you disable the USB port using the usb-port disable command, but then downgrade to an earlier release, the port will remain disabled, and you cannot re-enable it without erasing the NVRAM (the FXOS local-mgmt **erase secure all** command).
- **Downgrade issue from 9.18 or later**—There is a behavior change in 9.18 where the **access-group** command will be listed before its **access-list** commands. If you downgrade, the **access-group** command will be rejected because it has not yet loaded the **access-list** commands. This outcome occurs even if you had previously enabled the **forward-reference enable** command, because that command is now removed. Before you downgrade, be sure to copy all **access-group** commands manually, and then after downgrading, re-enter them.
- **Downgrade from 9.10(1) for smart licensing**—Due to changes in the smart agent, if you downgrade, you must re-register your device to the Cisco Smart Software Manager. The new smart agent uses an encrypted file, so you need to re-register to use an unencrypted file required by the old smart agent.
- **Downgrade to 9.5 and earlier with passwords using PBKDF2 (Password-Based Key Derivation Function 2) hash**—Versions before 9.6 do not support PBKDF2 hashing. In 9.6(1), **enable** and **username** passwords longer than 32 characters use PBKDF2 hashing. In 9.7(1), new passwords of all lengths use PBKDF2 hashing (existing passwords continue to use MD5 hashing). If you downgrade, the **enable** password reverts to the default (which is blank). Usernames will not parse correctly, and the **username** commands will be removed. You must re-create your local users.
- **Downgrade from Version 9.5(2.200) for the ASA Virtual**—The ASA virtual does not retain the licensing registration state. You need to re-register with the **license smart register idtoken id_token force** command (for ASDM: see the **Configuration > Device Management > Licensing > Smart Licensing** page, and use the **Force registration** option); obtain the ID token from the Smart Software Manager.

- **VPN tunnels are replicated to the standby unit even if the standby unit is running a version of software that does not support the Ciphersuite that the original tunnel negotiated**—This scenario occurs when downgrading. In this case, disconnect your VPN connection and reconnect.
- **Downgrade ASA virtual with flexible permanent license reservation from Version 9.23(1) to earlier versions** —We recommend that you do not downgrade an ASA virtual with Version 9.23(1) and flexible permanent license reservation. The license registration status will become **Unregistered**.

Incompatible Configuration Removed After Downgrading

When you downgrade to an old version, commands that were introduced in later versions will be removed from the configuration. There is no automated way to check the configuration against the target version before you downgrade. You can view when new commands were added in ASA new features by release.

You can view rejected commands *after* you downgrade using the **show startup-config errors** command. If you can perform a downgrade on a lab device, you can preview the effects using this command before you perform the downgrade on a production device.

In some cases, the ASA migrates commands to new forms automatically when you upgrade, so depending on your version, even if you did not manually configure new commands, the downgrade could be affected by configuration migrations. We recommend that you have a backup of your old configuration that you can use when you downgrade. In the case of upgrading to 8.3, a backup is automatically created (<old_version>_startup_cfg.sav). Other migrations do not create back-ups. See the "Version-Specific Guidelines and Migrations" in the ASA Upgrade guide for more information about automatic command migrations that could affect downgrading.

See also known downgrade issues in Guidelines and Limitations for Downgrading, on page 1072.

For example, an ASA running version 9.8(2) includes the following commands:

```
access-list acl1 extended permit sctp 192.0.2.0 255.255.255.0 198.51.100.0 255.255.255.0
username test1 password $sha512$1234$abcdefghijklmnopqrstuvwxyz privilege 15
snmp-server user snmpuser1 snmpgroup1 v3 engineID abcdefghijklmnopqrstuvwxyz encrypted auth
 md5 12:ab:34 priv aes 128 12:ab:34
```

When you downgrade to 9.0(4), you will see the following errors on startup:

```
access-list acl1 extended permit sctp 192.0.2.0 255.255.255.0 198.51.100.0 255.255.255.0
                                  ^
ERROR: % Invalid input detected at '^' marker.

username test1 password $sha512$1234$abcdefghijklmnopqrstuvwxyz pbkdf2 privilege 15
                                                                ^
ERROR: % Invalid input detected at '^' marker.

snmp-server user snmpuser1 snmpgroup1 v3 engineID abcdefghijklmnopqrstuvwxyz encrypted auth
 md5 12:ab:34 priv aes 128 12:ab:34
                                       ^
ERROR: % Invalid input detected at '^' marker.
```

In this example, support for **sctp** in the **access-list extended** command was added in version 9.5(2), support for **pbkdf2** in the **username** command was added in version 9.6(1), and support for **engineID** in the **snmp-server user** command was added in version 9.5(3).

Downgrade the ASA Appliance

You can downgrade the ASA software version by setting the ASA version to the old version, restoring the backup configuration to the startup configuration, and then reloading. This procedure applies to the following models:

- Firepower 1000
- Secure Firewall 1200
- Firepower 2100
- Secure Firewall 3100
- Secure Firewall 4200

Before you begin

This procedure requires a backup configuration of the ASA before you upgraded, so you can restore the old configuration. If you do not restore the old configuration, you may have incompatible commands representing new or changed features. Any new commands will be rejected when you load the old software version.

Procedure

Step 1 Load the old ASA software version using the upgrade procedure in the ASA upgrade guide for standalone, failover, or clustering deployments. In this case, specify the old ASA version instead of a new version. **Important:** Do *not* reload the ASA yet.

Step 2 At the ASA CLI, copy the backup ASA configuration to the startup configuration. For failover, perform this step on the active unit. This step replicates the command to the standby unit.

copy *old_config_url* **startup-config**

It's important that you do not save the running configuration to the startup configuration using **write memory**; this command will overwrite your backup configuration.

Example:

```
ciscoasa# copy disk0:/9.13.1_cfg.sav startup-config
```

Step 3 Reload the ASA.

ASA CLI

reload

ASDM

Choose **Tools** > **System Reload**.

Downgrade the Firepower 4100/9300

You can downgrade the ASA software version by restoring the backup configuration to the startup configuration, setting the ASA version to the old version, and then reloading.

Before you begin

- This procedure requires a backup configuration of the ASA before you upgraded, so you can restore the old configuration. If you do not restore the old configuration, you may have incompatible commands representing new or changed features. Any new commands will be rejected when you load the old software version.
- Make sure the old ASA version is compatibile with the current FXOS version. If not, downgrade FXOS as the first step before you restore the old ASA configuration. Just make sure the downgraded FXOS is also compatible with the current ASA version (before you downgrade it). If you cannot achieve compatibility, we suggest you do not perform a downgrade.

Procedure

Step 1 At the ASA CLI, copy the backup ASA configuration to the startup configuration. For failover or clustering, perform this step on the active/control unit. This step replicates the command to the standby/data units.

copy *old_config_url* **startup-config**

It's important that you do not save the running configuration to the startup configuration using **write memory**; this command will overwrite your backup configuration.

Example:

```
ciscoasa# copy disk0:/9.8.4_cfg.sav startup-config
```

Step 2 In FXOS, use the chassis manager or FXOS CLI to use the old ASA software version using the upgrade procedure in the ASA upgrade guide for standalone, failover, or clustering deployments. In this case, specify the old ASA version instead of a new version.

Step 3 If you are also downgrading FXOS, use the chassis manager or FXOS CLI to set the old FXOS software version to be the current version using the upgrade procedure in the ASA upgrade guide for standalone, failover, or clustering deployments.

Downgrade the ISA 3000

The downgrade feature provides a shortcut for completing the following functions on ISA 3000 models:

- Clearing the boot image configuration (**clear configure boot**).
- Setting the boot image to be the old image (**boot system**).
- (Optional) Entering a new activation key (**activation-key**).
- Saving the running configuration to startup (**write memory**). This sets the BOOT environment variable to the old image, so when you reload, the old image is loaded.

- Copying the old configuration backup to the startup configuration (**copy** *old_config_url* **startup-config**).
- Reloading (**reload**).

Before you begin

- This procedure requires a backup configuration of the ASA before you upgraded, so you can restore the old configuration.

Procedure

Step 1 Choose **Tools** > **Downgrade Software** .

The Downgrade Software dialog box appears.

Step 2 For the **ASA Image**, click **Select Image File**.

The **Browse File Locations** dialog box appears.

Step 3 Click one of the following radio buttons:

- **Remote Server**—Choose ftp, smb, or http from the drop-down list, and type the path to the old image file.
- **Flash File System**—Click **Browse Flash** to choose the old image file on the local flash file system.

Step 4 For the **Configuration**, click **Browse Flash** to choose the pre-migration configuration file.

Step 5 (Optional) In the **Activation Key** field, enter the old activation key if you need to revert to a pre-8.3 activation key.

Step 6 Click **Downgrade**.

Manage Files

ASDM provides a set of file management tools to help you perform basic file management tasks. The File Management tool lets you view, move, copy, and delete files stored in flash memory, transfer files, and to manage files on remote storage devices (mount points).

Note In multiple context mode, this tool is only available in the system security context.

Configure File Access

The ASA can use an FTP client, secure copy client, or TFTP client. You can also configure the ASA as a secure copy server so you can use a secure copy client on your computer.

Configure the FTP Client Mode

The ASA can use FTP to upload or download image files or configuration files to or from an FTP server. In passive FTP, the client initiates both the control connection and the data connection. The server, which is the recipient of the data connection in passive mode, responds with the port number to which it is listening for the specific connection.

Procedure

Step 1 From the Configuration > Device Management > Management Access > File Access > FTP Client pane, check the **Specify FTP mode as passive** check box.

Step 2 Click **Apply**.

The FTP client configuration is changed and the change is saved to the running configuration.

Configure the ASA Secure Copy Client

You can configure SCP settings when the ASA acts as an SCP client using the **copy** command.

The performance of SCP depends partly on the encryption cipher used. By default, the ASA negotiates one of the following algorithms in order: 3des-cbc aes128-cbc aes192-cbc aes256-cbc aes128-ctr aes192-ctr aes256-ctr. If the first algorithm proposed (3des-cbc) is chosen, then the performance is much slower than a more efficient algorithm such as aes128-cbc. To change the proposed ciphers, use the **Configuration** > **Device Management** > **Advanced** > **SSH Ciphers** pane; for example, choose **Custom** and set it to aes128-cbc.

Before you begin

- The ASA license must have the strong encryption (3DES/AES) license to support SSH Version 2 connections.
- Unless otherwise specified, for multiple context mode, complete this procedure in the system execution space. If you are not already in the System configuration mode, in the Configuration > Device List pane, double-click **System** under the active device IP address.
- For the SCP server, enable SSH on the ASA according to Configure SSH Access, on page 1029.

Procedure

Step 1 Depending on your context mode:

- For single mode, choose **Configuration > Device Management > Management Access > File Access > Secure Copy (SCP)**.
- For multiple mode in the System, choose **Configuration > Device Management > Device Administration > Secure Copy**

Step 2 (Optional) The ASA stores the SSH host key for each SCP server to which it connects. You can manually add or delete servers and their keys from the ASA database if desired.

To add a key:

a) Click **Add** for a new server, or select the server from the Trusted SSH Hosts table, and click **Edit**.
b) For a new server, in the Host field, enter the server IP address.
c) Check the **Add public key for the trusted SSH host** check box.
d) Specify one of the following keys:

- Fingerprint—Enter the already hashed key; for example, a key that you copied from **show** command output.
- Key—Enter the public key or hashed value of the SSH host. The key string is the Base64 encoded RSA public key of the remote peer. You can obtain the public key value from an open SSH client; that is, from the .ssh/id_rsa.pub file. After you submit the Base64 encoded public key, that key is then hashed via SHA-256.

To delete a key, select the server from the Trusted SSH Hosts table, and click **Delete**.

Step 3 (Optional) To be informed when a new host key is detected, check the **Inform me when a new host key is detected** check box.

By default, this option is enabled. When this option is enabled, you are prompted to accept or reject the host key if it is not already stored on the ASA. When this option is disabled, the ASA accepts the host key automatically if it was not stored before.

Step 4 Click **Apply**.

Configure the ASA TFTP Client Path

TFTP is a simple client/server file transfer protocol, which is described in RFC 783 and RFC 1350 Rev. 2. You can configure the ASA as a TFTP client so that it can copy files to or from a TFTP server. In this way, you can back up and propagate configuration files to multiple ASAs.

This section lets you predefine the path to a TFTP server so you do not need to enter it in commands such as **copy** and **configure net**.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Management Access** > **File Access** > **TFTP Client**, and check the **Enable** check box.

Step 2 From the Interface Name drop-down list, choose the interface to use as a TFTP client.

Step 3 In the IP Address field, enter the IP address of the TFTP server on which configuration files will be saved.

Step 4 In the Path field, enter the path to the TFTP server on which configuration files will be saved.

For example: /tftpboot/asa/config3

Step 5 Click **Apply**.

Add Mount Points

You can add a CIFS or FTP mount point.

Add a CIFS Mount Point

To define a Common Internet File System (CIFS) mount point, perform the following steps.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Management Access** > **File Access** > **Mount-Points**, and click **Add** > **CIFS Mount Point** .

The Add CIFS Mount Point dialog box appears.

Step 2 Check the **Enable mount point** check box.

This option attaches the CIFS file system on the ASA to the UNIX file tree.

Step 3 In the Mount Point Name field, enter the name of an existing CIFS location.

Step 4 In the Server Name or IP Address field, enter the name or IP address of the server in which the mount point is located.

Step 5 In the Share Name field, enter the name of the folder on the CIFS server.

Step 6 In the NT Domain Name field, enter the name of the NT Domain in which the server resides.

Step 7 In the User Name field, enter the name of the user authorized for file system mounting on the server.

Step 8 In the Password field, enter the password for the user authorized for file system mounting on the server.

Step 9 In the Confirm Password field, reenter the password.

Step 10 Click **OK**.

The Add CIFS Mount Point dialog box closes.

Step 11 Click **Apply**.

Add an FTP Mount Point

For an FTP mount point, the FTP server must have a UNIX directory listing style. Microsoft FTP servers have a default of the MS-DOS directory listing style.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Management Access** > **File Access** > **Mount-Points**, and click **Add** > **FTP Mount Point**.

The Add FTP Mount Point dialog box appears.

Step 2 Check the **Enable** check box.

This option attaches the FTP file system on the ASA to the UNIX file tree.

Step 3 In the Mount Point Name field, enter the name of an existing FTP location.

Step 4 In the Server Name or IP Address field, enter the name or IP address of the server where the mount point is located.

Step 5 In the Mode field, click the radio button for the FTP mode (**Active** or **Passive**). When you choose Passive mode, the client initiates both the FTP control connection and the data connection. The server responds with the number of its listening port for this connection.

Step 6 In the Path to Mount field, enter the directory path name to the FTP file server.

Step 7 In the User Name field, enter the name of the user authorized for file system mounting on the server.

Step 8 In the Password field, enter the password for the user authorized for file system mounting on the server.

Step 9 In the Confirm Password field, reenter the password.

Step 10 Click **OK**.

The Add FTP Mount Point dialog box closes.

Step 11 Click **Apply**.

Access the File Management Tool

To use the file management tools, perform the following steps.

Procedure

Step 1 In the main ASDM application window, choose **Tools** > **File Management**.

The File Management dialog box appears.

- The Folders pane displays the available folders on disk.
- Flash Space shows the total amount of flash memory and how much memory is available.
- The Files area displays the following information about files in the selected folder:
  - Path
  - Filename
  - Size (bytes)
  - Time Modified
  - Status, which indicates whether a selected file is designated as a boot configuration file, boot image file, ASDM image file, SVC image file, CSD image file, or APCF image file.

Step 2 Click **View** to display the selected file in your browser.

Step 3 Click **Cut** to cut the selected file for pasting to another directory.

Step 4 Click **Copy** to copy the selected file for pasting to another directory.

Step 5 Click **Paste** to paste the copied file to the selected destination.

Step 6 Click **Delete** to remove the selected file from flash memory.

Step 7 Click **Rename** to rename a file.

Step 8 Click **New Directory** to create a new directory for storing files.

Step 9 Click **File Transfer** to open the File Transfer dialog box. See Transfer Files, on page 1082 for more information.

Step 10 Click **Mount Points** to open the Manage Mount Points dialog box. See Add Mount Points , on page 1080 for more information.

Transfer Files

The File Transfer tool lets you transfer files from either a local or remote location. You can transfer a local file on your computer or a flash file system to and from the ASA. You can transfer a remote file to and from the ASA using HTTP, HTTPS, TFTP, FTP, or SMB.



Note For the IPS SSP software module, before you download the IPS software to disk0, make sure at least 50% of the flash memory is free. When you install IPS, IPS reserves 50% of the internal flash memory for its file system.

Transfer Files Between Local PC and Flash

To transfer files between your local computer and a flash file system, perform the following steps.

Procedure

Step 1 In the main ASDM application window, choose **Tools** > **File Management**.

The File Management dialog box appears.

Step 2 Click the down arrow next to **File Transfer**, and then click **Between Local PC and Flash**.

The File Transfer dialog box appears.

Step 3 Select and *drag* the file(s) from either your local computer or the flash file system that you want to upload or download to the desired location. Alternatively, select the file(s) from either your local computer or the flash file system that you want to upload or download, and click the right arrow or left arrow to transfer the file(s) to the desired location.

Step 4 Click **Close** when you are done.

Transfer Files Between Remote Server and Flash

To transfer files between a remote server and a flash file system, perform the following steps.

Procedure

Step 1 In the main ASDM application window, choose **Tools** > **File Management**.

The File Management dialog box appears.

Step 2 Click the down arrow from the File Transfer drop-down list, and then click **Between Remote Server and Flash**.

The File Transfer dialog box appears.

Step 3 To transfer a file from a remote server, click the **Remote server** option.

Step 4 Define the source file to be transferred.

a) (Optional) Specify the interface through which the ASA communicates with the server. If you do not specify the interface, the ASA checks the management-only routing table; if there are no matches, it then checks the data routing table.

b) Choose the path to the location of the file, including the IP address of the server.

Note
File transfer supports IPv4 and IPv6 addresses.

c) Enter the type (if the path is FTP) or the port number (if the path is HTTP or HTTPS) of the remote server. Valid FTP types are the following:

- ap—ASCII files in passive mode
- an—ASCII files in non-passive mode
- ip—Binary image files in passive mode
- in—Binary image files in non-passive mode

Step 5 To transfer the file from the flash file system, click the **Flash file system** option.

Step 6 Enter the path to the location of the file or click **Browse Flash** to find the file location.

Step 7 In addition, you can copy a file from your startup configuration, running configuration, or an SMB file system through the CLI. For instructions about using the **copy** command, see the CLI configuration guide.

Step 8 Define the destination of the file to be transferred.

a) To transfer the file to the flash file system, choose the **Flash file system** option.

b) Enter the path to the location of the file or click **Browse Flash** to find the file location.

Step 9 To transfer a file to a remote server, choose the **Remote server** option.

a) (Optional) Specify the interface through which the ASA communicates with the server. If you do not specify the interface, the ASA checks the management-only routing table; if there are no matches, it then checks the data routing table.

b) Enter the path to the location of the file.

c) For FTP transfers, enter the type. Valid types are the following:

- ap—ASCII files in passive mode
- an—ASCII files in non-passive mode

- ip—Binary image files in passive mode
- in—Binary image files in non-passive mode

Step 10 Click **Transfer** to start the file transfer.

The Enter Username and Password dialog box appears.

Step 11 Enter the username, password, and domain (if required) for the remote server.

Step 12 Click **OK** to continue the file transfer.

The file transfer process might take a few minutes; make sure that you wait until it is finished.

Step 13 Click **Close** when the file transfer is finished.

Set the ASA Image, ASDM, and Startup Configuration

If you have more than one ASA or ASDM image, you should specify the image that you want to boot. If you do not set the image, the default boot image is used, and that image may not be the one intended. For the startup configuration, you can optionally specify a file in the visible file system instead of a hidden directory.

See the following model guidelines:

- Firepower 4100/9300 chassis—ASA upgrades are managed by FXOS; you cannot upgrade the ASA within the ASA operating system, so do not use this procedure for the ASA image. You can upgrade the ASA and FXOS separately from each other, and they are listed separately in the FXOS directory listing. The ASA package always includes ASDM.
- Firepower 1000, Secure Firewall 1200/3100/4200—The ASA, ASDM, and FXOS images are bundled together into a single package. Package updates are managed by the ASA using this procedure. Although these platforms use the ASA to identify the image to boot, the underlying mechanism is different from legacy ASAs. See the command description below for more information.
- ASDM for the models—ASDM can be upgraded from within the ASA operating system, so you do not need to only use the bundled ASDM image. For Firepower 4100/9300, ASDM images that you upload manually do not appear in the FXOS image list; you must manage ASDM images from the ASA.



Note When you upgrade the ASA bundle, the ASDM image in the bundle replaces the previous ASDM bundle image on the ASA because they have the same name (**asdm.bin**). But if you manually chose a different ASDM image that you uploaded (for example, **asdm-782.bin**), then you continue to use that image even after a bundle upgrade. To make sure that you are running a compatible version of ASDM, you should either upgrade ASDM before you upgrade the bundle, or you should reconfigure the ASA to use the bundled ASDM image (**asdm.bin**) just before upgrading the ASA bundle.

- ASA Virtual—The initial deployment ASA virtual package puts the ASA image in the read-only boot:/ partition. When you upgrade the ASA virtual, you specify a different image in flash memory. Note that if you later clear your configuration, then the ASA virtual will revert to loading the original deployment

image. The initial deployment ASA virtual package also includes an ASDM image that it places in flash memory. You can upgrade the ASDM image separately.

- **disk0:** is the internal memory. Other drive numbers represent external storage such as a USB drive, SSD, or SD card.

See the following default settings:

- ASA image:
  - Firepower 1000, Secure Firewall 1200/3100/4200—Boots the previously-running boot image.
  - ISA 3000—Boots the first application image that it finds in internal flash memory.
  - ASA Virtual—Boots the image in the read-only boot:/ partition that was created when you first deployed.
  - Firepower 4100/9300 chassis—The FXOS system determines which ASA image to boot. You cannot use this procedure to set the ASA image.
- ASDM image on all ASAs—Boots the first ASDM image that it finds in internal flash memory, or if one does not exist in this location, then in external flash memory.
- Startup configuration—By default, the ASA boots from a startup configuration that is a hidden file.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **System Image/Configuration** > **Boot Image/Configuration**.

Firepower 1000 Secure Firewall 1200/3100/4200: You can only add a single image. If you upgrade to a new image, then you must delete the previous image you set. When you apply this change, the system performs an action: the system validates and unpacks the image and copies it to the boot location (an internal location on disk0 managed by FXOS). The new image will load when you reload the ASA. If you change your mind prior to reloading, you can delete the **Boot Image Location** and reapply to remove the new image from the boot location, so the current image continues to run. You can even delete the original image file from the ASA flash memory after you apply this change, and the ASA will boot correctly from the boot location. Unlike other models, this command in the startup configuration does not affect the booting image. The last-loaded boot image will always run upon reload. You can only load images with the original filename from the Cisco download site. If you change the filename, it will not load.

ASA virtual and ISA 3000: You can specify up to four local binary image files for use as the startup image, and one image located on a TFTP server for the device to boot from. If you specify an image located on a TFTP server, it must be first in the list. If the device cannot reach the TFTP server to load the image, it tries to load the next image file in the list located in flash.

Step 2 Click **Add** in the Boot Image/Configuration pane.

Step 3 Browse to the image from which you want to boot. For a TFTP image, enter the TFTP URL in the File Name field. Click **OK**.

Step 4 Arrange the images in order by using the Move Up and Move Down buttons.

Step 5 (Optional) In the Boot Configuration File Path field, specify the startup configuration file by clicking **Browse Flash** and choosing the configuration. Click **OK**.

This feature is important for when you work with large configurations that do not fit in the hidden directory. If you save a large configuration and see the following error message, be sure to instead save the configuration to a new file using this command:

```
%Error writing. nvram:/startup-config (No space left on device:)
```

Step 6 In the ASDM Image File Path field, specify the ASDM image by clicking **Browse Flash** and choosing the image. Click **OK**.

Step 7 Click **Apply**.

Back Up and Restore Configurations or Other Files

We recommend that you make regular backups of your configuration and other system files to guard against system failure.

Perform a Complete System Backup or Restoration

These procedures describe how to back up and restore configurations and images to a zip file and transfer it to your local computer.

Before You Begin Backup or Restore

- You should have at least 300 MB of disk space available at the backup or restore location before you start a backup or restore.
- The ASA must be in single context mode.
- If you make any configuration changes during or after a backup, those changes will not be included in the backup. If you change a configuration after making the backup, then perform a restore, this configuration change will be overwritten. As a result, the ASA might behave differently.
- You can start only one backup or restore at a time.
- You can only restore a configuration to the same ASA version as when you performed the original backup. You cannot use the restore tool to migrate a configuration from one ASA version to another. If a configuration migration is required, the ASA automatically upgrades the resident startup configuration when it loads the new ASA OS.
- If you use clustering, you can only back up or restore the startup-configuration, running-configuration, and identity certificates. You must create and restore a backup separately for each unit.
- If you use failover, you must create and restore a backup separately for the active and standby units.
- If you set a master passphrase for the ASA, then you need that master passphrase to restore the backup configuration that you create with this procedure. If you do not know the master passphrase for the ASA, see Configure the Master Passphrase, on page 698 to learn how to reset it before continuing with the backup.
- If you import PKCS12 data (with the **crypto ca trustpoint** command) and the trustpoint uses RSA keys, the imported key pair is assigned the same name as the trustpoint. Because of this limitation, if you specify a different name for the trustpoint and its key pair after you have restored an ASDM configuration,

the startup configuration will be the same as the original configuration, but the running configuration will include a different key pair name. This means that if you use different names for the key pair and trustpoint, you cannot restore the original configuration. To work around this issue, make sure that you use the same name for the trustpoint and its key pair.

- You cannot back up using the CLI and restore using ASDM, or vice versa.
- Each backup file includes the following content:
  - Running-configuration
  - Startup-configuration
  - All security images

    Cisco Secure Desktop and Host Scan images

    Cisco Secure Desktop and Host Scan settings

    Secure Client (SVC) images and profiles

    Secure Client (SVC) customizations and transforms
  - Identity certificates (includes RSA key pairs tied to identity certificates; excludes standalone keys)
  - VPN pre-shared keys
  - SSL VPN configurations
  - Application Profile Custom Framework (APCF)
  - Bookmarks
  - Customizations
  - Dynamic Access Policy (DAP)
  - Plug-ins
  - Pre-fill scripts for connection profiles
  - Proxy Auto-config
  - Translation table
  - Web content
  - Version information

Back Up the System

This procedure describes how to perform a complete system backup.

Note If the backup process stalls, your ASDM may not have enough memory to load the configuration. You can monitor the Java console for the "java.lang.OutOfMemoryError" message to see if running out of memory is the problem. To increaes the ASDM memory, see Increase the ASDM Configuration Memory, on page 26.

Procedure

Step 1 Create a folder on your computer to store backup files so they will be easy to find in case you need to restore them later.

Step 2 Choose **Tools** > **Backup Configurations**.

The Backup Configurations dialog box appears. Click the down arrow in the **SSL VPN Configuration** area to view the backup options for SSL VPN configurations. By default, all configuration files are checked and will be backed up if they are available. If you want to back up all of the files in the list, go to Step 5.

Step 3 Uncheck the **Backup All** check box if you want to select the configurations to back up.

Step 4 Check the check box next to the option that you want to back up.

Step 5 Click **Browse Local to specify a directory and file name for the backup .zip file.**

Step 6 In the Select dialog box, choose the directory in which you want to store the backup file.

Step 7 Click **Select**. The path appears in the Backup File field.

Step 8 Enter the name of the destination backup file after the directory path. The backup file name must be between 3 and 232 characters long.

Step 9 Click **Backup**. The backup proceeds immediately unless you are backing up certificates or the ASA is using a master passphrase.

Step 10 If you have configured and enabled a master passphrase on your ASA, you receive a warning message with a suggestion to change the master passphrase, if you do not know it, before proceeding with the backup. Click Yes to proceed with the backup if you know the master passphrase. The backup proceeds immediately unless you are backing up identity certificates.

Step 11 If you are backing up an identity certificate, you are asked to enter a separate passphrase to be used for encoding the certificates in PKCS12 format. You can enter a passphrase or skip this step.

Note
Only identity certificates are backed up by this process.

- To encrypt certificates, enter and confirm your certificate passphrase in the Certificate Passphrase dialog box and click OK. You will need to remember the password you enter in this dialog box when restoring the certificates.
- Clicking **Cancel** skips the step and does not back up certificates.

After clicking OK or Cancel, the backup begins immediately.

Step 12 After the backup is complete, the status window closes and the Backup Statistics dialog box appears to provide success and failure messages.

Note
Backup “failure messages” are most likely caused by the lack of an existing configuration for the types indicated.

Step 13 Click **OK** to close the Backup Statistics dialog box.

Restore the Backup

You can specify configurations and images to restore from a zip tar.gz file on your local computer.

Note If the restore process stalls, your ASDM may not have enough memory to load the configuration. You can monitor the Java console for the "java.lang.OutOfMemoryError" message to see if running out of memory is the problem. To increaes the ASDM memory, see Increase the ASDM Configuration Memory, on page 26.

Procedure

Step 1 Choose **Tools** > **Restore Configurations**.

Step 2 In the Restore Configurations dialog box, click **Browse Local Directory**, choose the zip file on your local computer that contains the configuration to restore, then click **Select**. The path and the zip filename appear in the **Local File** field.

The zip file that you restore must be created by choosing the **Tools > Backup Configurations** option.

Step 3 Click **Next**. The second Restore Configuration dialog box appears. Check the check boxes next to the configurations that you want to restore. All available SSL VPN configurations are selected by default.

Step 4 Click **Restore**.

Step 5 If you specified a certificate passphrase with which to encrypt the certificates when you created the backup file, ASDM prompts you to enter the passphrase.

Step 6 If you chose to restore the running configuration, you are asked if you want to merge the running configuration, replace the running configuration, or skip this part of the restoration process.

- Merging configurations combines the current running configuration and the backed-up running configuration.
- Replacing the running configuration uses the backed-up running configuration only.
- Skipping the step does not restore the backed-up running configuration.

ASDM displays a status dialog box until the restore operation is finished.

Step 7 If you replaced or merged the running configuration, close ASDM and restart it. If you did not restore the running configuration or the running configuration, refresh the ASDM session for the changes to take effect.

Configure Automatic Backup and Restore (ISA 3000)

On the ISA 3000, you can configure automatic backups to a particular location every time you save your configuration.

Automatic restore lets you easily configure new devices with a complete configuration loaded on an SD flash memory card. Automatic restore is enabled in the default factory configuration.

Configure Automatic Backup (ISA 3000)

On the ISA 3000, you can configure automatic backups to a particular location every time you save your configuration.

Before you begin

This feature is only available on the ISA 3000.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Auto Backup & Restore Configuration**.

Step 2 Check or uncheck **Automate Backup Configuration** to enable or disable automatic backups.

If you enable automatic backups, when you save the configuration, the configuration is automatically saved to the backup location as well as to the startup configuration. The backup file has the name "auto-backup-asa.tgz".

Set the following parameters:

- **Interface**—Specifies the interface to reach the backup URL, if you specify off-device storage. If you do not specify the interface name, the ASA checks the management-only routing table; if there are no matches, it then checks the data routing table.
- **Location**—Specifies the storage medium to be used for backing up data. You can specify a URL or local storage. disk0 is the internal flash drive. disk1 is an optional USB memory stick on USB 1. disk2 is an optional USB memory stick on USB 2. And disk3 is the SD memory card. The default for automatic restore is disk3:.
- **Passphrase**—Sets the passphrase to secure the backed-up data. The default for automatic restore is "cisco".

Configure Automatic Restore (ISA 3000)

Automatic restore mode restores the system configuration on a device without any user intervention. For example, you insert an SD memory card containing a saved backup configuration into a new device and then power the device on. When the device comes up, it checks the SD card to decide if the system configuration needs to be restored. (The restoration is only initiated if the backup file has the "fingerprint" of a different device. The fingerprint of the backup file is updated to match the current device during a backup or restore operation. So if the device has already completed a restore, or if it has created its own backup, then the automatic restore is skipped.) If the fingerprint shows a restoration is required, the device replaces the system configuration (startup-config, running-config, SSL VPN configuration, and so on; see Back Up the System, on page 1087 for details about the contents of the backup). When the device finishes booting, it is running the saved configuration.

Automatic restore is enabled in the default factory configuration, so you can easily configure new devices with a complete configuration loaded on an SD memory card without having to perform any pre-configuration of the device.

Because the device needs to decide early in the boot process if the system configuration needs to be restored, it checks ROMMON variables to determine if the device is in automatic restore mode and to obtain the location of the backup configuration. The following ROMMON variables are used:

- **RESTORE_MODE =** {**auto** | **manual**}

  The default is **auto**.

- **RESTORE_LOCATION =** {**disk0:** | **disk1:** | **disk2:** | **disk3:**}

  The default is **disk3:**.

- **RESTORE_PASSPHRASE** = *key*

  The default is **cisco**.

To change the automatic restore settings, complete the following procedure.

Before you begin

- This feature is only available on the ISA 3000.
- If you use the default restore settings, you need an SD memory card installed (part number SD-IE-1GB=).
- If you need to restore the default configuration to ensure that automatic restore is enabled, use the **configure factory default** command. This command is only available in transparent firewall mode, so if you are in routed firewall mode, use the **firewall transparent** command first.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Auto Backup & Restore Configuration**.

Step 2 Check or uncheck **Automate Restore Configuration** to enable or disable automatic restore.

The name of the file that is restored is "auto-backup-asa.tgz". If you enable automatic restore, set the following parameters:

- **Location**—Specifies the storage medium to be used for restoring data. disk0 is the internal flash drive. disk1 is an optional USB memory stick on USB 1. disk2 is an optional USB memory stick on USB 2. And disk3 is the SD memory card. The default is disk3.
- **Passphrase**—Sets the passphrase to read the backed-up data. The default is "cisco".

Save the Running Configuration to a TFTP Server

This feature stores a copy of the current running configuration file on a TFTP server.

Procedure

Step 1 Choose **File** > **Save Running Configuration to TFTP Server**.

The Save Running Configuration to TFTP Server dialog box appears.

Step 2 Enter the TFTP server IP address and file path on the TFTP server in which the configuration file will be saved, and then click **Save Configuration**.

Note
To configure default TFTP settings, choose **Configuration > Device Management > Management Access > File Access > TFTP Client**. After you have configured this setting, the TFTP server IP address and file path on the TFTP server appear automatically in this dialog box.

Schedule a System Restart

The System Reload tool lets you schedule a system restart or cancel a pending restart.

Procedure

Step 1 Choose **Tools** > **System Reload**.

Step 2 In the Reload Scheduling area, define the following settings:

a) For the Configuration State, choose either to save or discard the running configuration at restart time.

b) For the Reload Start Time, choose from the following options:

- Click **Now** to perform an immediate restart.
- Click **Delay by** to delay the restart by a specified amount of time. Enter the time before the restart begins in hours and minutes or only minutes.
- Click **Schedule at** to schedule the restart to occur at a specific time and date. Enter the time of day the restart is to occur, and select the date of the scheduled restart.

c) In the Reload Message field, enter a message to send to open instances of ASDM at restart time.

d) Check the **On reload failure force immediate reload after** check box to show the amount of time elapsed in hours and minutes or only minutes before a restart is attempted again.

e) Click **Schedule Reload** to schedule the restart as configured.

The Reload Status area displays the status of the restart.

Step 3 Choose one of the following:

- Click **Cancel Reload** to stop a scheduled restart.
- Click **Refresh** to refresh the Reload Status display after a scheduled restart is finished.
- Click **Details** to display the results of a scheduled restart.

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 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.

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:

```
ciscoasa(config)# 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:

```
ciscoasa# 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:

ciscoasa# 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.

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:

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

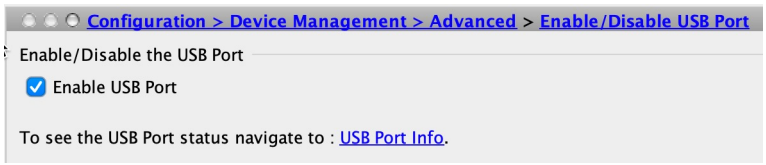
- Enabling or disabling the USB port requires a reload.
- If the USB port is disabled and you downgrade to a version that does not support this feature, the port will remain disabled, and you cannot re-enable it without erasing the NVRAM (the FXOS local-mgmt **erase secure all** command).
- If you perform a ROMMON **factory-reset** or FXOS local-mgmt **erase secure**, the USB port will be re-enabled.
- For high availability or clustering, disable or enable the USB port on the active unit or control node. The command is replicated to the other nodes. However, you need to reload each unit for the change to take effect.

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

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Advanced** > **Enable/Disable USB Port**.

Figure 101: Enable/Disable the USB Port



Step 2 Uncheck the **Enable USB Port** checkbox and click **Apply**.

Click **OK** on the Information dialog box.

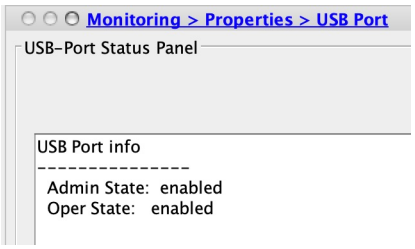
To re-enable the USB port, check the **Enable USB Port** checkbox and click **Apply**.

Step 3 Click **Save**.

Step 4 Choose **Tools** > **System Reload** to reload the device for your changes to take effect.

Step 5 Choose **Monitoring** > **Properties** > **USB Port** to view the port status.

Figure 102: USB Port Status



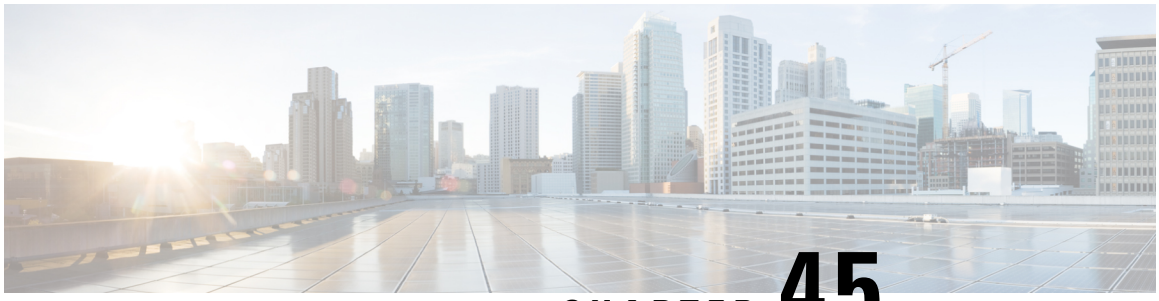
The Admin State shows the USB port configuration. The Oper State shows the current operation. For example, if you disable the USB port but do not reload, the Admin State will show disabled while the Oper State would will enabled.

History for Software and Configurations

Feature Name	Platform Releases	Feature Information
Secure Copy client and server	9.1(5)/9.2(1)	The ASA now supports the Secure Copy (SCP) client and server to transfer files to and from a SCP server. We modified the following screens: Tools > File Management > File Transfer > Between Remote Server and Flash **Configuration > Device Management > Management Access > File Access > Secure Copy (SCP) Server**

Feature Name	Platform Releases	Feature Information
Configurable SSH encryption and integrity ciphers	9.1(7)9.4(3)9.5(3)9.6(1)	Users can select cipher modes when doing SSH encryption management and can configure HMAC and encryption for varying key exchange algorithms. You might want to change the ciphers to be more or less strict, depending on your application. Note that the performance of secure copy depends partly on the encryption cipher used. By default, the ASA negotiates one of the following algorithms in order: 3des-cbc aes128-cbc aes192-cbc aes256-cbc aes128-ctr aes192-ctr aes256-ctr. If the first algorithm proposed (3des-cbc) is chosen, then the performance is much slower than a more efficient algorithm such as aes128-cbc. To change the proposed ciphers, use **ssh cipher encryption custom aes128-cbc**, for example. We introduced the following screen: **Configuration** > **Device Management** > **Advanced** > **SSH Ciphers**
Auto Update server certificate verification enabled by default	9.2(1)	The Auto Update server certificate verification is now enabled by default; for new configurations, you must explicitly disable certificate verification. If you are upgrading from an earlier release, and you did not enable certificate verification, then certificate verification is not enabled, and you see the following warning: `WARNING: The certificate provided by the auto-update servers will not be verified. In order to verify this certificate please use the verify-certificate option.` The configuration will be migrated to explicitly configure no verification. We modified the following screen: Configuration > Device Management > System/Image Configuration > Auto Update > Add Auto Update Server.
System backup and restore using the CLI	9.3(2)	You can now back up and restore complete system configurations, including images and certificates, using the CLI. We did not modify any ASDM screens.
Recovering and loading a new ASA 5506W-X image	9.4(1)	We now support the recovery and loading of a new ASA 5506W-X image. We did not modify any ASDM screens.

Feature Name	Platform Releases	Feature Information
Automatic Backup and Restore for the ISA 3000	9.7(1)	You can enable auto-backup and/or auto-restore functionality using pre-set parameters in the backup and restore commands. The use cases for these features include initial configuration from external media; device replacement; roll back to an operable state. We introduced the following screen: **Configuration** > **Device Management** > **Auto Backup & Restore Configuration**
CiscoSSH stack requires SSH access when using the SCP client	9.17(1)	If you use the CiscoSSH stack, to use the ASA **copy** command to copy a file to or from an SCP server, you have to enable SSH access on the ASA for the SCP server subnet/host
RAID support for SSDs on the Secure Firewall 3100	9.17(1)	The SSDs are self-encrypting drives (SEDs), and if you have 2 SSDs, they form a software RAID. New/Modified commands: **raid, show raid, show ssd**
Disable the USB port	9.22(1)	By default, the type-A USB port is enabled and could not be disabled. You can now disable USB port access for security purposes on the following models: • Firepower 1000 • Secure Firewall 1200 • Secure Firewall 3100 • Secure Firewall 4200 This setting is stored in firmware and requires a reload. Moreover, 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. **Note** This feature does not affect the USB console port, if present. New/Modified screens: . • **Configuration** > **Device Management** > **Advanced** > **Enable/Disable USB Port** • **Monitoring** > **Properties** > **USB Port** > **USB Port Info**



CHAPTER 45

Response Automation for System Events

This chapter describes how to configure the Embedded Event Manager (EEM).

- About the EEM, on page 1099
- Guidelines for the EEM, on page 1100
- Configure the EEM, on page 1101
- Monitoring the EEM, on page 1104
- History for the EEM, on page 1104

About the EEM

The EEM service enables you to debug problems and provides general purpose logging for troubleshooting. There are two components: events to which the EEM responds or listens, and event manager applets that define actions as well as the events to which the EEM responds. You may configure multiple event manager applets to respond to different events and perform different actions.

Supported Events

The EEM supports the following events:

- Syslog—The ASA uses syslog message IDs to identify syslog messages that trigger an event manager applet. You may configure multiple syslog events, but the syslog message IDs may not overlap within a single event manager applet.
- Timers—You may use timers to trigger events. You may configure each timer only once for each event manager applet. Each event manager applet may have up to three timers. The three types of timers are the following:
  - Watchdog (periodic) timers trigger an event manager applet after the specified time period following the completion of the applet actions and restart automatically.
  - Countdown (one-shot) timers trigger an event manager applet once after the specified time period and do not restart unless they are removed, then re-added.
  - Absolute (once-a-day) timers cause an event to occur once a day at a specified time, and restart automatically. The time-of-day format is in hh:mm:ss.

    You may configure only one timer event of each type for each event manager applet.

- None—The none event is triggered when you run an event manager applet manually using the CLI or ASDM.
- Crash—The crash event is triggered when the ASA crashes. In some scenarios, a force crash is triggered:

  If ASA is configured to reload on block depletion, and when ASA remains out of memory for a configured duration, it sends out syslog and collects troubleshoot data. ASA force crashes and triggers the reload process to release the memory block. In a HA setup, under such cases, failover is triggered. On a cluster setup, the node leaves the cluster.

  Regardless of the value of the **output** command, the **action** commands are directed to the crashinfo file. The output is generated before the **show tech** command.

Actions on Event Manager Applets

When an event manager applet is triggered, the actions on the event manager applet are performed. Each action has a number that is used to specify the sequence of the actions. The sequence number must be unique within an event manager applet. You may configure multiple actions for an event manager applet. The commands are typical CLI commands, such as **show blocks**.

Output Destinations

You may send the output from the actions to a specified location using the **output** command. Only one output value may be enabled at any one time. The default value is **output none**. This value discards any output from the **action** commands. The command runs in global configuration mode as a user with privilege level 15 (the highest). The command may not accept any input, because it is disabled. You may send the output of the **action** CLI commands to one of three locations:

- **None**, which is the default and discards the output
- **Console**, which sends the output to the ASA console
- **File**, which sends the output to a file. The following four file options are available:
  - **Create a unique file**, which creates a new, uniquely named file each time that an event manager applet is invoked
  - **Create/overwrite a file**, which overwrites a specified file each time that an event manager applet is invoked.
  - **Create/append to a file**, which appends to a specified file each time that an event manager applet is invoked. If the file does not yet exist, it is created.
  - **Create a set of files**, which creates a set of uniquely named files that are rotated each time that an event manager applet is invoked.

Guidelines for the EEM

This section describes guidelines and limitations that you should check before configuring the EEM.

Context Mode Guidelines

Not supported in multiple context mode.

Additional Guidelines

- During a crash, the state of the ASA is generally unknown. Some commands may not be safe to run during this condition.
- The name of an event manager applet may not contain spaces.
- You cannot modify the None event and Crashinfo event parameters.
- Performance may be affected because syslog messages are sent to the EEM for processing.
- The default output is **output none** for each event manager applet. To change this setting, you must enter a different output value.
- You may have only one output option defined for each event manager applet.

Configure the EEM

Configuring the EEM consists of the following tasks:

Procedure

Step 1 Create an Event Manager Applet and Configure Events, on page 1101.
Step 2 Configure an Action and Destinations for Output from an Action, on page 1102.
Step 3 Run an Event Manager Applet, on page 1103.
Step 4 Track Memory Allocation and Memory Usage, on page 1103.

Create an Event Manager Applet and Configure Events

To create an event manager applet and configure events, perform the following steps:

Procedure

Step 1 In ASDM, choose **Configuration** > **Device Management** > **Advanced** > **Embedded Event Manager**.
Step 2 Click **Add** to display the **Add Event Manager Applet** dialog box.
Step 3 Enter the name of the applet (without spaces) and describe what it does. The description may be up to 256 characters long. You may include spaces in description text if it is placed within quotes.
Step 4 Click **Add** in the **Events** area to display the **Add Event Manager Applet Event** dialog box.
Step 5 Choose the event type that you want to configure from the **Type** drop-down list. The available options are **crashinfo**, **None**, **Syslog**, **Once-a-day timer**, **One-shot timer**, and **Periodic** timer.

- **Syslog**: Enter a single syslog message or a range of syslog messages. If a syslog message occurs that matches the specified individual syslog message or range of syslog messages, an event manager applet is triggered. (Optional) Enter the number of times in the **occurrences** field that the syslog message must occur for an event manager applet to be invoked. The default is 1 occurrence every 0 seconds. Valid values are from 1 - 4294967295. (Optional) Enter the number of seconds in the **period** field within which the syslog messages must occur to invoke the action. This value limits how frequently an event manager applet is invoked to at most once in the configured period. Valid values are from 0 - 604800. A value of 0 means that no period is defined.
- **Periodic**: Enter the time period in seconds. The number of seconds may range from 1- 604800.
- **Once-a-day timer**: Enter the time of day in hh:mm:ss. The time range is from 00:00:00 (midnight) to 23:59:59.
- **One-shot timer**: Enter the time period in seconds. The number of seconds may range from 1- 604800.
- **None**: Choose this option to invoke an event manager applet manually.
- **crashinfo**: Choose this option to trigger a crash event when the ASA crashes.

Configure an Action and Destinations for Output from an Action

To configure an action and specific destinations for sending output from an action, perform the following steps:

Procedure

Step 1 Click **Add** to display the **Add Event Manager Applet** dialog box.

Step 2 Enter the name of the applet (without spaces) and describe what it does. The description may be up to 256 characters long.

Step 3 Click **Add** in the **Actions** area to display the **Add Event Manager Applet Action** dialog box.

Step 4 Enter the unique sequence number in the **Sequence #** field. Valid sequence numbers range from 0 - 4294967295.

Step 5 Enter the CLI command in the **CLI Command** field. The command runs in global configuration mode as a user with privilege level 15 (the highest). The command may not accept any input, because it is disabled.

Step 6 Click **OK** to close the **Add Event Manager Applet Action** dialog box.

The newly added action appears in the **Actions** list.

Step 7 Click **Add** to open the **Add Event Manager Applet** dialog box.

Step 8 Choose one of the available output destination options:

- Choose the **None** option from the **Output Location** drop-down list to discard any output from the **action** commands. This is the default setting.
- Choose the **Console** option from the **Output Location** drop-down list to sends the output of the **action** commands to the console.

  Note
  Running this command affects performance.

- Choose the **File** option from the **Output Location** drop-down list to send the output of the **action** commands to a new file for each event manager applet that is invoked. The **Create a unique file** option is automatically selected as the default.

  The filename has the format of eem-*applet*-*timestamp*.log, in which *applet* is the name of the event manager applet and *timestamp* is a dated time stamp in the format of YYYYMMDD-hhmmss.

- Choose the **File** option from the **Output Location** drop-down list, then choose the **Create a set of files** option from the drop-down list to create a set of rotated files.

  When a new file is to be written, the oldest file is deleted, and all subsequent files are renumbered before the first file is written. The newest file is indicated by 0, and the oldest file is indicated by the highest number. Valid values for the rotate value range from 2 - 100. The filename format is eem-*applet*-*x*.log, in which *applet* is the name of the applet, and *x* is the file number.

- Choose the **File** option from the **Output Location** drop-down list, then choose the **Create/overwrite a file option** from the drop-down list to write the **action** command output to a single file, which is overwritten every time.

- Choose the **FileFile** option from the **Output Location** drop-down list, then choose the **Create/append a file** option from the drop-down list to writes the **action** command output to a single file, but that file is appended to every time.

Step 9 Click **OK** to close the **Add Event Manager Applet** dialog box.

The specified output destination appears in the **Embedded Event Manager** pane.

Run an Event Manager Applet

To run an event manager applet, perform the following steps:

Procedure

Step 1 In the **Embedded Event Manager** pane, select an event manager applet from the list that has been configured with the **None** event.

Step 2 Click **Run**.

Track Memory Allocation and Memory Usage

To log memory allocation and memory usage, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Advanced** > **Embedded Event Manager**.

Step 2 Click **Add** to display the **Add Event Manager Applet** dialog box.

Step 3 Click **Add** again to display the **Add Event Manager Applet Event** dialog box.

Step 4 Choose **memory-logging-wrap** from the drop-down list.

Step 5 Click **OK** to add it to the **Events** list.

Step 6 Click **OK** again to add it to the **Applets** list.

Monitoring the EEM

See the following commands to monitor the EEM.

- **Monitoring > Properties > EEM Applets**

  This pane shows the list of EEM applets and their hit count value.

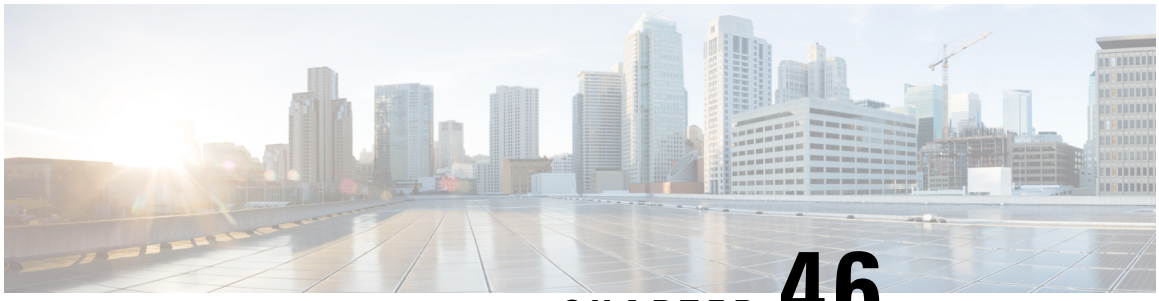
- **Tools > Command Line Interface**

  This pane allows you to issue various non-interactive commands and view results.

History for the EEM

Table 57: History for the EEM

Feature Name	Platform Releases	Description
Embedded Event Manager (EEM)	9.2(1)	The EEM service enables you to debug problems and provides general purpose logging for troubleshooting. There are two components: events to which the EEM responds or listens, and event manager applets that define actions as well as the events to which the EEM responds. You may configure multiple event manager applets to respond to different events and perform different actions. We introduced the following screens: **Configuration > Device Management > Advanced > Embedded Event Manager, Monitoring > Properties > EEM Applets**.
Memory tracking for the EEM	9.4(1)	We have added a new debugging feature to log memory allocations and memory usage, and to respond to memory logging wrap events. We modified the following screen: **Configuration > Device Management > Advanced > Embedded Event Manager > Add Event Manager Applet > Add Event Manager Applet Event**.



CHAPTER 46

Testing and Troubleshooting

This chapter describes how to troubleshoot the ASA and test basic connectivity.

- Recover Enable and Telnet Passwords, on page 1105
- Configure and Run Captures with the Packet Capture Wizard, on page 1109
- CPU Usage and Reporting, on page 1115
- Test Your Configuration, on page 1120
- Monitoring Performance and System Resources, on page 1128
- Monitoring Connections, on page 1130
- History for Testing and Troubleshooting , on page 1130

Recover Enable and Telnet Passwords

If you forget the enable or Telnet passwords, you can recover them for ASA virtual and ISA 3000 models. You must perform the task using the CLI.

Note For other platforms, you cannot recover lost passwords. You can only restore the factory default configuration, and reset the passwords to the default. For Firepower 4100/9300, see the FXOS configuration guide. For other models, see the FXOS troubleshooting guide.

Recover Passwords on the ISA 3000

To recover passwords for the ISA 3000 perform the following steps:

Procedure

Step 1 Connect to the ASA console port.

Step 2 Power off the ASA, then power it on.

Step 3 After startup, press the **Escape** key when you are prompted to enter ROMMON mode.

Step 4 To update the configuration register value, enter the following command:

```
rommon #1> confreg 0x41

You must reset or power cycle for new config to take effect
```

The ASA displays the current configuration register value and a list of configuration options. Record the current configuration register value, so you can restore it later.

```
Configuration Register: 0x00000041

Configuration Summary
 [ 0 ] password recovery
 [ 1 ] display break prompt
 [ 2 ] ignore system configuration
 [ 3 ] auto-boot image in disks
 [ 4 ] console baud: 9600
 boot: ...... auto-boot index 1 image in disks
```

Step 5 Reload the ASA by entering the following command:

```
rommon #2> boot
Launching BootLoader...
Boot configuration file contains 1 entry.

Loading disk0:/asa932-226-k8.bin... Booting...Loading...
```

The ASA loads the default configuration instead of the startup configuration.

Step 6 Access the privileged EXEC mode by entering the following command:

```
ciscoasa# enable
```

Step 7 When prompted for the password, press **Enter**.

The password is blank.

Step 8 Load the startup configuration by entering the following command:

```
ciscoasa# copy startup-config running-config
```

Step 9 Access the global configuration mode by entering the following command:

```
ciscoasa# configure terminal
```

Step 10 Change the passwords, as required, in the default configuration by entering the following commands:

```
ciscoasa(config)# password password
ciscoasa(config)# enable password password
ciscoasa(config)# username name password password
```

Step 11 Load the default configuration by entering the following command:

```
ciscoasa(config)# no config-register
```

The default configuration register value is 0x1. See the command reference for more information about the configuration register.

Step 12 Save the new passwords to the startup configuration by entering the following command:

```
ciscoasa(config)# copy running-config startup-config
```

Recover Passwords or Images on the ASA Virtual

To recover passwords or images on the ASA virtual, perform the following steps:

Procedure

Step 1 Copy the running configuration to a backup file on the ASA virtual:

copy running-config *filename*

Example:

```
ciscoasa# copy running-config backup.cfg
```

Step 2 Restart the ASA virtual:

reload

Step 3 From the GNU GRUB menu, press the down arrow, choose the <**filename**> **with no configuration load** option, then press **Enter**. The filename is the default boot image filename on the ASA virtual. The default boot image is never automatically booted through the **fallback** command. Then load the selected boot image.

```
GNU GRUB version 2.0(12)4
bootflash:/asa100123-20-smp-k8.bin
bootflash: /asa100123-20-smp-k8.bin with no configuration load
```

Example:

```
GNU GRUB version 2.0(12)4
bootflash: /asa100123-20-smp-k8.bin with no configuration load
```

Step 4 Copy the backup configuration file to the running configuration.

copy *filename* **running-config**

Example:

```
ciscoasa (config)# copy backup.cfg running-config
```

Step 5 Reset the password.

enable password *password*

Example:

```
ciscoasa(config)# enable password cisco123
```

Step 6 Save the new configuration.

write memory

Example:

```
ciscoasa(config)# write memory
```

Disable Password Recovery for ISA 3000 Hardware

Note You cannot disable password recovery on the ASA virtual, Secure Firewall models.

To disable password recovery to ensure that unauthorized users cannot use the password recovery mechanism to compromise the ASA, perform the following steps.

Before you begin

On the ASA, the **no service password-recovery** command prevents you from entering ROMMON mode with the configuration intact. When you enter ROMMON mode, the ASA prompts you to erase all Flash file systems. You cannot enter ROMMON mode without first performing this erasure. If you choose not to erase the Flash file system, the ASA reloads. Because password recovery depends on using ROMMON mode and maintaining the existing configuration, this erasure prevents you from recovering a password. However, disabling password recovery prevents unauthorized users from viewing the configuration or inserting different passwords. In this case, to restore the system to an operating state, load a new image and a backup configuration file, if available.

The **service password-recovery** command appears in the configuration file for information only. When you enter the command at the CLI prompt, the setting is saved in NVRAM. The only way to change the setting is to enter the command at the CLI prompt. Loading a new configuration with a different version of the command does not change the setting. If you disable password recovery when the ASA is configured to ignore the startup configuration at startup (in preparation for password recovery), then the ASA changes the setting to load the startup configuration as usual. If you use failover, and the standby unit is configured to ignore the startup configuration, then the same change is made to the configuration register when the **no service password-recovery** command replicates to the standby unit.

Procedure

Disable password recovery.

no service password-recovery

Example:

```
ciscoasa (config)# no service password-recovery
```

Configure and Run Captures with the Packet Capture Wizard

You can use the Packet Capture Wizard to configure and run captures for troubleshooting errors. The captures can use ACLs to limit the type of traffic captured, the source and destination addresses and ports, and one or more interfaces. The wizard runs one capture on each of the ingress and egress interfaces. You can save the captures on your PC to examine them in a packet analyzer.

Note This tool does not support clientless SSL VPN capture.

To configure and run captures, perform the following steps:

Procedure

Step 1 Choose **Wizards** > **Packet Capture Wizard**.

The **Overview of Packet Capture** screen appears, with a list of the tasks through which the wizard will guide you to complete. Those tasks include the following:

- Selecting an ingress interface.
- Selecting an egress interface.
- Setting the buffer parameters.
- Running the captures.
- Saving the captures to your PC (optional).

Step 2 Click **Next**.

In a clustering environment, the **Cluster Option** screen appears. Go to Step 3.

In a non-clustering environment, the **Ingress Traffic Selector** screen appears. Go to Step 4.

Step 3 Choose one of the following options in the **Cluster Option** screen for running a capture: **This device only** or **The whole cluster**, then click **Next** to display the **Ingress Selector** screen.

Step 4 Click the **Select Interface** radio button to capture packets on an interface.

In a clustering environment, to capture only the cluster control plane packets, select the **CP-Cluster** check box.

Step 5 Click the **Use backplane channel** radio button to capture packets on the ASA CX dataplane.

Step 6 Do one of the following in the **Packet Match Criteria** area:

- Click the **Select access list** radio button to specify the ACL to use for matching packets, then choose the ACL from the **Select ACL** drop-down list. Click **Manage** to display the **ACL Manager** pane to add a previously configured ACL to the current drop-down list. Choose an ACL, then click **OK**.

  The access list option is disabled when you enable switch packet capture. For details, see Ingress Traffic Selector, on page 1112.

- Click the **Specify Packet Parameters** radio button to specify packets parameters.

a) Do one of the following in the **ICMP Capture** drop-down list:

Note
The **ICMP Capture** field is populated only when you select **The whole cluster** as the cluster option in the previous window.

- Select **include-decrypted** to capture decrypted IPsec packets which contain both normal and decrypted traffic once they enter the firewall device.
- Select **persist** to capture persistent packets on cluter units.

Step 7 To continue, see Ingress Traffic Selector, on page 1112.

Step 8 Click **Next** to display the **Egress Traffic Selector** screen.

Step 9 Click the **Select Interface** radio button to capture packets on an interface.

In a clustering environment, to capture the cluster control plane packets, select the **CP-Cluster** check box.

Note
To know more details on the Egress Traffic Selector fields, see Egress Traffic Selector, on page 1113.

To know more details on the Egress Traffic Selector fields, see Egress Traffic Selector, on page 1113.

Step 10 Click **Next** to display the **Buffers & Captures** screen. To continue, see Buffers.

Step 11 Check the **Get capture every 10 seconds** check box in the **Capture Parameters** area to obtain the latest capture every 10 seconds automatically. By default, this capture uses the circular buffer.

Step 12 You specify the buffer size and packet size in the **Buffer Parameters** area. The buffer size is the maximum amount of memory that the capture can use to store packets. The packet size is the longest packet that the capture can hold. We recommend that you use the longest packet size to capture as much information as possible.

a) (Optional. Applicable only for Secure Firewall 3100 devices) Check the **Switch** check box to store captured switch packets.
b) Enter the packet size. The valid size ranges from 14 - 1522 bytes. For switch packet capture, the valid size ranges from 64 to 9006 bytes.
c) Enter the buffer size. The valid size ranges from 1534 - 33554432 bytes. For switch packet capture, the valid size ranges from 256 to 2048 bytes.
d) Check the **Use circular buffer** check box to store captured packets.

Note
When you choose this setting, if all the buffer storage is used, the capture starts overwriting the oldest packets.

Step 13 Click **Next** to display the **Summary** screen, which shows the cluster options for all units in the cluster (if you are using clustering), traffic selectors, and buffer parameters that you have entered. To continue, see Summary.

Step 14 Click **Next** to display the **Run Captures** screen, then click **Start** to begin capturing packets. Click **Stop** to end the capture. To continue, see Run Captures, on page 1114. If you are using clustering, go to Step 16.

Step 15 Click **Get Capture Buffer** to determine how much buffer space you have remaining. Click **Clear Buffer on Device** to remove the current content and allow room in the buffer to capture more packets.

Step 16 In a clustering environment, on the **Run Captures** screen, perform one or more of the following steps:

- Click **Get Cluster Capture Summary** to view a summary of packet capture information for all units in the cluster, followed by packet capture information for each unit.
- Click **Get Capture Buffer** to determine how much buffer space you have remaining in each unit of the cluster. The **Capture Buffer from Device** dialog box appears.
- Click **Clear Capture Buffer** to remove the current content for one or all of the units in a cluster and allow room in the buffer to capture more packets.

Step 17 Click **Save captures** to display the **Save Capture** dialog box. You have the option of saving either the ingress capture, the egress capture, or both. To continue, see Save Captures.

Step 18 Click **Save Ingress Capture** to display the **Save capture file** dialog box. Specify the storage location on your PC, then click **Save**.

Step 19 Click **Launch Network Sniffer Application** to start the packet analysis application specified in **Tools** > **Preferences** for analyzing the ingress capture.

Step 20 Click **Save Egress Capture** to display the **Save capture file** dialog box. Specify the storage location on your PC, then click **Save**.

Step 21 Click **Launch Network Sniffer Application** to start the packet analysis application specified in **Tools** > **Preferences** for analyzing the egress capture.

Step 22 Click **Close**, then click **Finish** to exit the wizard.

Guidelines for Packet Capture

Context Mode

- You can configure captures on the cluster control link within a context; only the packet that is associated with the context sent in the cluster control link is captured.
- You can only configure one capture for a shared VLAN; if you configure a capture in multiple contexts on the shared VLAN, then only the last capture that was configured is used.
- If you remove the last-configured (active) capture, no captures become active, even if you have previously configured a capture in another context; you must remove the capture and add it again to make it active.
- All traffic that enters the interface to which the capture is attached is captured, including traffic to other contexts on the shared VLAN. Therefore, if you enable a capture in Context A for a VLAN that is also used by Context B, both Context A and Context B ingress traffic are captured.

- For egress traffic, only the traffic of the context with the active capture is captured. The only exception is when you do not enable the ICMP inspection (therefore the ICMP traffic does not have a session in the accelerated path). In this case, both ingress and egress ICMP traffic for all contexts on the shared VLAN is captured.

Additional Guidelines

- If the ASA receives packets with an incorrectly formatted TCP header and drops them because of the *invalid-tcp-hdr-length* ASP drop reason, the **show capture** command output on the interface where those packets are received does not show those packets.
- You can only capture IP traffic; you cannot capture non-IP packets such as ARPs.
- For inline SGT tagged packets, captured packets contain an additional CMD header that your PCAP viewer might not understand.
- Packet captures include packets that the system modifies or injects into the connection due to inspection, NAT, TCP normalization, or other features that adjust the content of a packet.
- The trace of the lifespan of an injected virtual packet in a datapath does not exactly reflect how the datapath handles the physical packets. This difference depends on the software version, configuration, and type of the injected virtual packets. Following are configuration settings that might lead to the disparity:
  - at least 2 NAT statements for the same host exist.
  - forward and reverse flows of a connection having different protocols. For example, forward flow is UDP or TCP, reverse flow is ICMP.
  - ICMP error inspection being enabled.

Ingress Traffic Selector

To configure the ingress interface, source and destination hosts or networks, and the protocol for packet capture, perform the following steps:

Procedure

Step 1 Choose the ingress interface name from the drop-down list.

Step 2 Enter the ingress source host and network. Click the **Use backplane channel** radio button to capture packets on the ASA CX dataplane.

Step 3 Enter the ingress destination host and network.

Step 4 Enter the protocol type to capture. Available protocols are ah, eigrp, esp, gre, icmp, icmp6, igmp, igrp, ip, ipinip, nos, ospf, pcp, pim, snp, tcp, or udp.

a) Enter the ICMP type for ICMP only. Available types include all, alternate address, conversion-error, echo, echo-reply, information-reply, information-request, mask-reply, mask-request, mobile-redirect, parameter-problem, redirect, router-advertisement, router-solicitation, source-quench, time-exceeded, timestamp-reply, timestamp-request, traceroute, or unreachable.

b) Specify the source and destination port services for the TCP and UDP protocols only. Available options include the following:

- Choose **All Services** to include all services.
- Choose **Service Groups** to include a service group.

  To include a specific service, choose one of the following: aol, bgp, chargen, cifx, citrix-ica, ctiqbe, daytime, discard, domain, echo, exec, finger, ftp, ftp-data, gopher, h323, hostname, http, https, ident, imap4, irc, , klogin, kshell, ldap, ldaps, login, lotusnotes, lpd, netbios-ssn, nntp, pcanywhere-data, pim-auto-rp, pop2, pop3, pptp, rsh, rtsp, sip, smtp, sqlnet, ssh, sunrpc, tacacs, talk, telnet, uucp, or whois.

Step 5 To enable packet capture for the Cisco TrustSec service, in the **Security Group Tagging** area, check the **SGT number** check box and enter the security group tag number. Valid security group tag numbers range from 2 to 65519.

Step 6 (Optional. Applicable only for Secure Firewall 3100 devices and Secure Firewall 4200 model devices) To enable switch packet capture, in the **Switch Control** area, check the **Switch** check box and specify the inner VLAN and outer VLAN ranges (1 to 4096).

Note
When you enable switch packet capture, the access list option is disabled.

Step 7 (Optional. This option is applicable when you enable switch packet capture.) To configure the ingress traffic direction parameter for the packet capture of Secure Firewall 4200 model devices, in the **Direction Control** area, from the **Direction** drop-down, choose a direction.

Egress Traffic Selector

To configure the egress interface, source and destination hosts/networks, and source and destination port services for packet capture, perform the following steps:

Procedure

Step 1 Click the **Select Interface** radio button to capture packets on an interface. Click the **Use backplane channel** radio button to capture packets on the ASA CX dataplane.

Step 2 Choose the egress interface name from the drop-down list.

Step 3 Enter the egress source host and network.

Step 4 Enter the egress destination host and network.

The protocol type selected during the ingress configuration is already listed.

Step 5 (Optional. Applicable only for Secure Firewall 3100 devices and Secure Firewall 4200 model devices) If you have enabled switch packet capture, specify the inner VLAN and outer VLAN ranges (1 to 4096). To enable switch packet capture, see Ingress Traffic Selector, on page 1112.

Step 6 (Optional. This option is applicable when you enable switch packet capture.) To configure the egress traffic direction parameter for the packet capture of Secure Firewall 4200 model devices, in the **Direction Control** area, from the **Direction** drop-down, choose a direction.

Buffers

To configure the packet size, buffer size, and use of the circular buffer for packet capture, perform the following steps:

Procedure

Step 1 Enter the longest packet that the capture can hold. Use the longest size available to capture as much information as possible.

Step 2 Enter the maximum amount of memory that the capture can use to store packets.

Step 3 Use the circular buffer to store packets. When the circular buffer has used all of the buffer storage, the capture will overwrite the oldest packets first.

Summary

The **Summary** screen shows the cluster options (if you are using clustering), traffic selectors, and the buffer parameters for the packet capture selected in the previous wizard screens.

Run Captures

To start and stop the capture session, view the capture buffer, launch a network analyzer application, save packet captures, and clear the buffer, perform the following steps:

Procedure

Step 1 Click **Start** to begin the packet capture session on a selected interface.

Step 2 Click **Stop** to stop the packet capture session on a selected interface.

Step 3 Click **Get Capture Buffer** to obtain a snapshot of the captured packets on the interface.

Step 4 Click **Ingress** to show the capture buffer on the ingress interface.

Step 5 Click **Egress** to show the capture buffer on the egress interface.

Step 6 Click **Clear Buffer on Device** to clear the buffer on the device.

Step 7 Click **Launch Network Sniffer Application** to start the packet analysis application for analyzing the ingress capture or the egress capture specified in **Tools** > **Preferences**.

Step 8 Click **Save Captures** to save the ingress and egress captures in either ASCII or PCAP format.

Save Captures

To save the ingress and egress packet captures to ASCII or PCAP file format for further packet analysis, perform the following steps:

Procedure

Step 1 Click **ASCII** to save the capture buffer in ASCII format.

Step 2 Click **PCAP** to save the capture buffer in PCAP format.

Step 3 Click **Save ingress capture** to specify a file in which to save the ingress packet capture.

Step 4 Click **Save egress capture** to specify a file in which to save the egress packet capture.

CPU Usage and Reporting

The CPU Utilization report summarizes the percentage of the CPU used within the time specified. Typically, the core operates at approximately 30 to 40 percent of total CPU capacity during non-peak hours and approximately 60 to 70 percent capacity during peak hours.

vCPU Usage in the ASA Virtual

Use the **show cpu usage** command on the ASA virtual to display CPU utilization statistics. The ASA virtual vCPU usage shows the amount of vCPUs used for the data path, control point, and external processes.

The vCPU usage reported by the cloud service provider (such as VMware, Azure, OCI, and so on) includes the ASA virtual usage as described plus:

- ASA virtual idle time
- %SYS overhead used for the ASA virtual VM
- Overhead of moving packets between vSwitches, vNICs, and pNICs. This overhead can be quite significant.

CPU Usage Example

The following is an example in which the reported vCPU usage is substantially different:

- ASA Virtual reports: 40%
- DP: 35%
- External Processes: 5%
- vSphere reports: 95%
- ASA (as ASA virtual reports): 40%
- ASA idle polling: 10%

- Overhead: 45%

The overhead is used to perform hypervisor functions and to move packets between NICs and vNICs using the vSwitch.

Usage can exceed 100% because the ESXi server can use additional compute resources for overhead on behalf of the ASA virtual.

VMware CPU Usage Reporting

In vSphere, click the **VM Performance** tab, then click **Advanced** to display the **Chart Options** drop-down list, which shows vCPU usage for each state (%USER, %IDLE, %SYS, and so on) of the VM. This information is useful for understanding VMware's perspective on where CPU resources are being used.

On the ESXi server shell (you access the shell by using SSH to connect to the host), esxtop is available. Esxtop has a similar look and feel to the Linux **top** command and provides VM state information for vSphere performance, including the following:

- Details on vCPU, memory, and network usage
- vCPU usage for each state of each VM.
- Memory (type M while running) and network (type N while running), as well as statistics and the number of RX drops

ASA Virtual and vCenter Graphs

There are differences in the CPU % numbers between the ASA virtual and vCenter:

- The vCenter graph numbers are always higher than the ASA virtual numbers.
- vCenter calls it %CPU usage; the ASA virtual calls it %CPU utilization.

The terms "%CPU utilization" and "%CPU usage" mean different things:

- CPU utilization provides statistics for physical CPUs.
- CPU usage provides statistics for logical CPUs, which is based on CPU hyperthreading. But because only one vCPU is used, hyperthreading is not turned on.

vCenter calculates the CPU % usage as follows:

Amount of actively used virtual CPUs, specified as a percentage of the total available CPUs

This calculation is the host view of the CPU usage, not the guest operating system view, and is the average CPU utilization over all available virtual CPUs in the virtual machine.

For example, if a virtual machine with one virtual CPU is running on a host that has four physical CPUs and the CPU usage is 100%, the virtual machine is using one physical CPU completely. The virtual CPU usage calculation is Usage in MHz / number of virtual CPUs x core frequency

When you compare the usage in MHz, both the vCenter and ASA virtual numbers match. According to the vCenter graph, MHz % CPU usage is calculated as 60/(2499 x 1 vCPU) = 2.4

Amazon CloudWatch CPU Usage Reporting

You can view the metrics explorer to monitor resources by their tags and properties. Perform the following steps to view CPU utilization statistics for a specific instance:

Procedure

Step 1 Open the **CloudWatch** console and choose **Metrics** in the navigation pane.

Step 2 Select the **EC2** metric namespace and select **Per-instance Metrics** dimension.

Step 3 Enter **CPUUtilization** in the search field and press Enter. Select the row for the required instance to display a graph for the **CPUUtilization** metric for that instance.

See Amazon CloudWatch documentation for more information.

ASA Virtual and Amazon CloudWatch Graphs

The Amazon CloudWatch graph numbers are higher than the numbers because of the different ways in which CPU usage is calculated on the ASA virtual and the CloudWatch.

When ASA virtual is running in poll mode, each CPU runs a loop of lightweight commands instead of entering power saving mode or any other idle state. This improves performance by keeping each core active at all times instead of having to be turned on/off or have their clocks adjusted by the Intel power states.

Inside ASA virtual, this activity is understood to be idling behaviour, and the CPU usage is correctly calculated. However, on the Amazon CloudWatch, the idle behaviour appears like normal CPU activity because all CPU cycles have instructions to run, and this causes the CloudWatch to display a high CPU usage percentage (85-90%).

Azure CPU Usage Reporting

Perform the following steps to view CPU utilization % across all monitored VMs by using VM Insights from the Azure Monitor:

Procedure

Step 1 Go to the Azure portal, select **Monitor** and choose **Virtual Machines** in the **Solutions** section.

Step 2 Select the **Performance** tab to display the **CPU Utilization %** chart. This chart displays the top five machines with the highest average processor utilization.

Perform the following steps to view the CPU utilization % chart directly from a specific Azure VM:

Procedure

Step 1 Go to the Azure portal and select **Virtual Machines**.

Step 2 From the list of VMs, choose a VM.

Step 3 In the **Monitoring** section, select **Insights**.

Step 4 Select the **Performance** tab.

See How to chart performance with VM Insights for more information.

ASA Virtual and Azure Graphs

There are differences in the CPU % numbers between the ASA virtual and Azure. The Azure graph numbers are always higher than the ASA virtual numbers because Azure calculates the CPU % usage as the amount of actively used virtual CPUs, specified as a percentage of the total available CPUs.

This calculation is the host view of the CPU usage, not the guest operating system view, and is the average CPU utilization over all available virtual CPUs in the virtual machine.

For example, if a virtual machine with one virtual CPU is running on a host that has four physical CPUs and the CPU usage is 100%, the virtual machine is using one physical CPU completely. The virtual CPU usage calculation is Usage in MHz / number of virtual CPUs x core frequency.

Azure also rate limits the amount of CPU that is requested by the guest OS. Consider a scenario in which ASA virtual is reporting CPU usage at 40% and the hypervisor is reporting CPU usage at 90%. Now, if the ASA virtual needs more processing power, the CPU usage may go above 80% and the hypervisor may then report CPU usage at more than 95%. This leads to the hypervisor throttling the ASA virtual CPU even though the ASA virtual is just running a loop of lightweight commands in poll mode exhibiting idling behavior.

Hyper-V CPU Usage Reporting

In addition to viewing CPU, RAM, and disk space configuration information for available Cloud Servers, you can also view disk, I/O, and networking information. Use this information to help you decide which Cloud Server is right for your needs. You can view the available servers through either the command-line nova client or the **Cloud Control Panel** interface.

On the command line, run the following command:

```
nova flavor-list
```

All available server configurations are displayed. The list contains the following information:

- ID - The server configuration ID
- Name - The configuration name, labeled by RAM size and performance type
- Memory_MB - The amount of RAM for the configuration
- Disk - The size of the disk in GB (for general purpose Cloud Servers, the size of the system disk)
- Ephemeral - The size of the data disk

- Swap - The size of the swap space
- VCPUs - The number of virtual CPUs associated with the configuration
- RXTX_Factor - The amount of bandwidth, in Mbps, allocated to the PublicNet ports, ServiceNet ports, and isolated networks (cloud networks) attached to a server
- Is_Public - Not used

ASA Virtual and Hyper-V Graphs

There are differences in the CPU % numbers between the ASA Virtual and Hyper-V:

- The Hyper-V graph numbers are always higher than the ASA Virtual numbers.
- Hyper-V calls it %CPU usage; the ASA Virtual calls it %CPU utilization.

The terms “%CPU utilization” and “%CPU usage” mean different things:

- CPU utilization provides statistics for physical CPUs.
- CPU usage provides statistics for logical CPUs, which is based on CPU hyperthreading. But because only one vCPU is used, hyperthreading is not turned on.

Hyper-V calculates the CPU % usage as follows:

Amount of actively used virtual CPUs, specified as a percentage of the total available CPUs

This calculation is the host view of the CPU usage, not the guest operating system view, and is the average CPU utilization over all available virtual CPUs in the virtual machine.

For example, if a virtual machine with one virtual CPU is running on a host that has four physical CPUs and the CPU usage is 100%, the virtual machine is using one physical CPU completely. The virtual CPU usage calculation is Usage in MHz / number of virtual CPUs x core frequency



Note It is recommended to look into ASA Virtual reports to get an accurate CPU usage percentage.

OCI CPU Usage Reporting

You can view the CPU utilization % in OCI by using the compute instance metric: **oci_computeagent**. The CpuUtilization metric displays the activity level from the CPU and is expressed as a percentage of total time. Perform the following steps to view metric charts for a single compute instance:

Procedure

Step 1 Open the navigation menu and click **Instances** under **Compute**.

Step 2 Click an instance and click **Metrics** under **Resources**.

Step 3 Select **oci_computeagent** in the Metric namespace list.

See Compute Instance Metrics for more information.

ASA Virtual and OCI Graphs

The OCI graph numbers are always higher than the ASA virtual numbers because OCI calculates the CPU % usage as the amount of actively used virtual CPUs, specified as a percentage of the total available CPUs.

This calculation is the host view of the CPU usage, not the guest operating system view, and is the average CPU utilization over all available virtual CPUs in the virtual machine.

For example, if a virtual machine with one virtual CPU is running on a host that has four physical CPUs and the CPU usage is 100%, the virtual machine is using one physical CPU completely. The virtual CPU usage calculation is Usage in MHz / number of virtual CPUs x core frequency.

Test Your Configuration

This section describes how to test connectivity for the single mode ASA or for each security context, how to ping the ASA interfaces, and how to allow hosts on one interface to ping through to hosts on another interface.

Test Basic Connectivity: Pinging Addresses

Ping is a simple command that lets you determine if a particular address is alive and responsive. The following topics explain more about the command and what types of testing you can accomplish with it.

What You Can Test Using Ping

When you ping a device, a packet is sent to the device and the device returns a reply. This process enables network devices to discover, identify, and test each other.

You can use ping to do the following tests:

- Loopback testing of two interfaces—You can initiate a ping from one interface to another on the same ASA, as an external loopback test to verify basic “up” status and operation of each interface.
- Pinging to an ASA—You can ping an interface on another ASA to verify that it is up and responding.
- Pinging through an ASA—You can ping through an intermediate ASA by pinging a device on the other side of the ASA. The packets will pass through two of the intermediate ASA’s interfaces as they go in each direction. This action performs a basic test of the interfaces, operation, and response time of the intermediate unit.
- Pinging to test questionable operation of a network device—You can ping from an ASA interface to a network device that you suspect is functioning incorrectly. If the interface is configured correctly and an echo is not received, there might be problems with the device.
- Pinging to test intermediate communications—You can ping from an ASA interface to a network device that is known to be functioning correctly. If the echo is received, the correct operation of any intermediate devices and physical connectivity is confirmed.

Choosing Between ICMP and TCP Ping

The ASA includes the traditional ping, which sends ICMP Echo Request packets and gets Echo Reply packets in return. This is the standard tool and works well if all intervening network devices allow ICMP traffic. With ICMP ping, you can ping IPv4 or IPv6 addresses, or host names.

However, some networks prohibit ICMP. If this is true of your network, you can instead use TCP ping to test network connectivity. With TCP ping, the ping sends TCP SYN packets, and considers the ping a success if it receives a SYN-ACK in response. With TCP ping, you can ping IPv4 addresses or host names, but you cannot ping IPv6 addresses.

Keep in mind that a successful ICMP or TCP ping simply means that the address you are using is alive and responding to that specific type of traffic. This means that basic connectivity is working. Other policies running on a device could prevent specific types of traffic from successfully getting through a device.

Enable ICMP

By default, you can ping from a high security interface to a low security interface. You just need to enable ICMP inspection to allow returning traffic through. If you want to ping from low to high, then you need to apply an ACL to allow traffic.

When pinging an ASA interface, any ICMP rules applied to the interface must allow Echo Request and Echo Response packets. ICMP rules are optional: if you do not configure them, all ICMP traffic to an interface is allowed.

This procedure explains all of ICMP configuration you might need to complete to enable ICMP pinging of ASA interfaces, or for pinging through an ASA.

Procedure

Step 1 Ensure ICMP rules allow Echo Request/Echo Response.

ICMP rules are optional and apply to ICMP packets sent directly to an interface. If you do not apply ICMP rules, all ICMP access is allowed. In this case, no action is required.

However, if you do implement ICMP rules, ensure that you include rules that permit any address for the Echo and Echo-Reply messages on each interface. Configure ICMP rules on the **Configuration > Device Management > Management Access > ICMP** pane.

Step 2 Ensure access rules allow ICMP.

When pinging a host through an ASA, access rules must allow ICMP traffic to leave and return. The access rule must at least allow Echo Request/Echo Reply ICMP packets. You can add these rules as global rules.

If you do not have access rules, you will need to also allow the other type of traffic you want, because applying any access rules to an interface adds an implicit deny, so all other traffic will be dropped.

Configure access rules on the **Configuration > Firewall > Access Rules** pane. If you are simply adding the rules for testing purposes, you can delete them after completing the tests.

Step 3 Enable ICMP inspection.

ICMP inspection is needed when pinging through the ASA, as opposed to pinging an interface. Inspection allows returning traffic (that is, the Echo Reply packet) to return to the host that initiated the ping, and also ensures there is one response per packet, which prevents certain types of attack.

You can simply enable ICMP inspection in the default global inspection policy.

a) Choose **Configuration > Firewall > Service Policy Rules**.
b) Edit the **inspection_default** global rule.
c) On the **Rule Actions > Protocol Inspection** tab, select ICMP.
d) Click **OK**, then **Apply**.

Ping Hosts

To ping any device, you simply choose **Tools > Ping**, enter the IP address or host name of the destination you are pinging, and click **Ping**. For TCP ping, you select **TCP** and also include the destination port. That is usually the extent of any test you need to run.

Example output for a successful ping:

```
Sending 5, 100-byte ICMP Echos to out-pc, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

If the ping fails, the output indicates ? for each failed attempt, and the success rate is less than 100 percent (complete failure is 0 percent):

```
Sending 5, 100-byte ICMP Echos to 10.132.80.101, timeout is 2 seconds:
?????
Success rate is 0 percent (0/5)
```

However, you can also add parameters to control some aspects of the ping. Following are your basic options:

- ICMP ping—You can select the interface used for the source IP address; however, the egress interface is determined by a route lookup using the data routing table. You can ping IPv4 or IPv6 addresses or host names.
- TCP ping—You must also select the TCP port for the destination you are pinging. For example, **www.example.com 80** to ping the HTTP port. You can ping IPv4 addresses or host names, but not IPv6 addresses.

  You also have the option to specify the interface used for the source IP address; however, the egress interface is determined by a route lookup using the data routing table.

  Finally, you can specify how often to repeat the ping (the default is 5 times) or the timeout for each attempt (the default is 2 seconds).

Test ASA Connectivity Systematically

If you want to do a more systematic test of ASA connectivity, you can use the following general procedure.

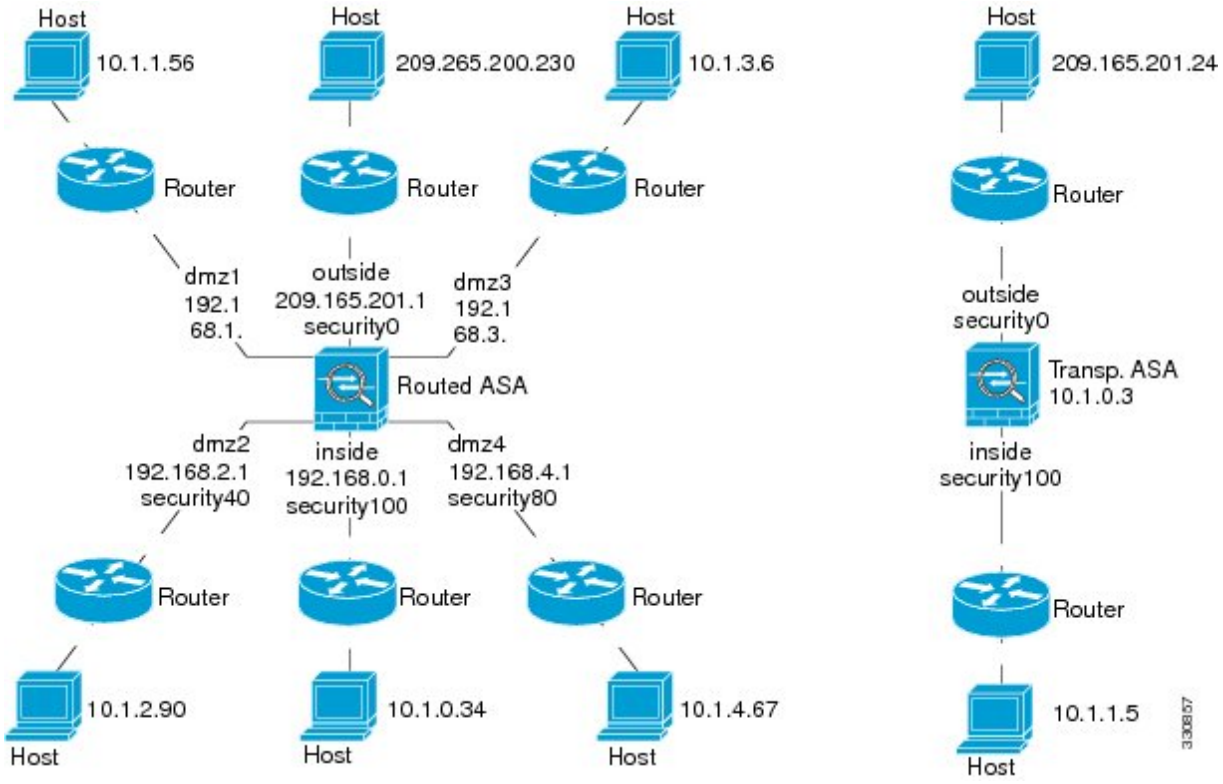
Before you begin

If you want to see the syslog messages mentioned in the procedure, enable logging (the **logging enable** command, or **Configuration > Device Management > Logging > Logging Setup** in ASDM).

Procedure

Step 1 Draw a diagram of your single-mode ASA or security context that shows the interface names, security levels, and IP addresses. The diagram should also include any directly connected routers and a host on the other side of the router from which you will ping the ASA.

Figure 103: Network Diagram with Interfaces, Routers, and Hosts



Step 2 Ping each ASA interface from the directly connected routers. For transparent mode, ping the BVI IP address. This test ensures that the ASA interfaces are active and that the interface configuration is correct.

A ping might fail if the ASA interface is not active, the interface configuration is incorrect, or if a switch between the ASA and a router is down (see the following figure). In this case, no debugging messages or syslog messages appear, because the packet never reaches the ASA.

Figure 104: Ping Failure at the ASA Interface

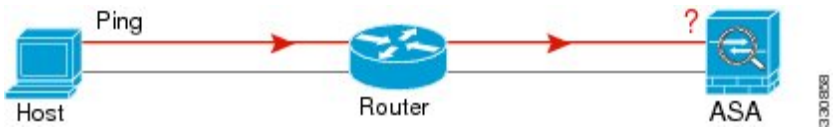
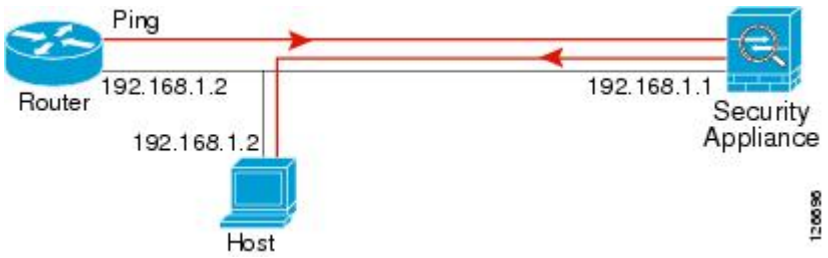


Figure 105: Ping Failure Because of IP Addressing Problems



If the ping reply does not return to the router, then a switch loop or redundant IP addresses might exist (see the following figure).

Step 3 Ping each ASA interface from a remote host. For transparent mode, ping the BVI IP address. This test checks whether the directly connected router can route the packet between the host and the ASA, and whether the ASA can correctly route the packet back to the host.

A ping might fail if the ASA does not have a return route to the host through the intermediate router (see the following figure). In this case, the debugging messages show that the ping was successful, but syslog message 110001 appears, indicating a routing failure has occurred.

Figure 106: Ping Failure Because the ASA Has No Return Route



Step 4 Ping from an ASA interface to a network device that you know is functioning correctly.

- If the ping is not received, a problem with the transmitting hardware or interface configuration may exist.
- If the ASA interface is configured correctly and it does not receive an echo reply from the "known good" device, problems with the interface hardware receiving function may exist. If a different interface with "known good" receiving capability can receive an echo after pinging the same "known good" device, the hardware receiving problem of the first interface is confirmed.

Step 5 Ping from the host or router through the source interface to another host or router on another interface. Repeat this step for as many interface pairs as you want to check. If you use NAT, this test shows that NAT is operating correctly.

If the ping succeeds, a syslog message appears to confirm the address translation for routed mode (305009 or 305011) and that an ICMP connection was established (302020). You can also enter either the **show xlate** or **show conns** command to view this information.

The ping might fail because NAT is not configured correctly. In this case, a syslog message appears, showing that the NAT failed (305005 or 305006). If the ping is from an outside host to an inside host, and you do not have a static translation, you get message 106010.

Figure 107: Ping Failure Because the ASA is Not Translating Addresses



Trace Routes to Hosts

If you are having problems sending traffic to an IP address, you can trace the route to the host to determine if there is a problem on the network path.

Procedure

Step 1 Make the ASA Visible on Trace Routes, on page 1125.

Step 2 Determine Packet Routes, on page 1126.

Make the ASA Visible on Trace Routes

By default, the ASA 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 ASA, and increase the rate limit on ICMP unreachable messages.

Procedure

Step 1 Decrement the TTL using a service policy.

a) Choose **Configuration > Firewall > Service Policy Rules**.
b) Add or edit a rule. For example, if you already have a rule to which you can add the option to decrement TTL, you do not need to create a new one.
c) Progress through the wizard to the Rule Actions page, applying the rule globally or to an interface, and specifying the traffic match. For example, you could create a global match any rule.
d) On the Rule Actions page, click the **Connection Settings** tab, and select **Decrement time to live for a connection**.
e) Click **OK** or **Finish**, then **Apply**.

Step 2 Increase the ICMP unreachable rate limit.

a) Choose **Configuration > Device Management > Management Access > ICMP**.
b) Increase the **IPv4 ICMP Unreachable Message Limits > Rate Limit** value at the bottom of the page. For example, increase it to 50.
c) Click **Apply**.

Determine Packet Routes

Use Traceroute to help you to determine the route that packets will take to their destination. A traceroute works by sending UDP packets or ICMPv6 echo to a destination on an invalid port. Because the port is not valid, the routers along the way to the destination respond with an ICMP or ICMPv6 Time Exceeded Message, and report that error to the ASA.

The traceroute shows the result of each probe sent. Every line of output corresponds to a TTL value in increasing order. The following table explains the output symbols.

Output Symbol	Description
*	No response was received for the probe within the timeout period.
U	No route to the destination.
nn msec	For each node, the round-trip time (in milliseconds) for the specified number of probes.
!N.	ICMP network unreachable. For ICMPv6, address is out of scope.
!H	ICMP host unreachable.
!P	ICMP unreachable. For ICMPv6, port not reachable.
!A	ICMP administratively prohibited.
?	Unknown ICMP error.

Procedure

Step 1 Choose **Tools** > **Traceroute**.

Step 2 Enter the destination hostname or IP address to which you are tracing the route. Configure a DNS server to use a host name.

Step 3 (Optional) Configure the characteristics of the trace. The defaults are appropriate in most cases.

- **Timeout**—How long to wait for a response before timing out. The default is 3 seconds.
- **Port**—The UDP port to use. The default is 33434.
- **Probe**—How many probes to send at each TTL level. The default is 3.
- **TTL**—The minimum and maximum time-to-live values for the probes. The minimum default is one, but you can set it to a higher value to suppress the display of known hops. The maximum default is 30. The traceroute terminates when the packet reaches the destination or when the maximum value is reached.
- **Specify source interface or IP address**—The interface to use as the source of the trace. You can specify the interface by name or by IP address. For IPv6, you cannot specify the source interface; you can only specify the source IP address. An IPv6 address is valid only if you enabled IPv6 on an ASA interface.In transparent mode, you must use the management address.
- **Reverse Resolve**—Whether to have the output display the names of hops encountered if DNS name resolution is configured. Deselect the option to show IP addresses only.
- **Use ICMP**—Whether to send ICMP probe packets instead of UDP probe packets.

Step 4 Click **Trace Route** to start the traceroute.

The **Traceroute Output** area displays detailed messages about the traceroute results.

Using the Packet Tracer to Test Policy Configuration

You can test your policy configuration by modeling a packet based on source and destination addressing and protocol characteristics. The trace does policy lookup to test access rules, NAT, routing, and so forth, to see if the packet would be permitted or denied.

By testing packets this way, you can see the results of your policies and test whether the types of traffic you want to allow or deny are handled as desired. Besides verifying your configuration, you can use the tracer to debug unexpected behavior, such as packets being denied when they should be allowed.

Procedure

Step 1 Choose **Tools** > **Packet Tracer**.

Step 2 Choose the source **Interface** for the packet trace.

Step 3 Specify the **Packet Type** for the packet trace. Available protocol types include: ICMP, IP, TCP, UDP, SCTP.

Step 4 (Optional.) If you want to trace a packet where the security group tag value is embedded in the Layer 2 CMD header (Trustsec), check **SGT number** and enter the security group tag number, 0-65533.

Step 5 (Transparent mode) If you want the packet tracer to enter a parent interface, which is later redirected to a subinterface, check **VLAN ID** and enter the ID, 1- 4096. VLAN ID is available only when the input interface is not a subinterface.

Step 6 (Transparent mode) Specify the **Destination MAC Address**.

Step 7 Specify the source and destination for the packets.

You can specify IPv4 or IPv6 addresses, fully-qualified domain names (FQDN), or security group names or tags if you use Cisco Trustsec. For the source address, you can also specify a username in the format Domain\username.

Step 8 Specify the protocol characteristics:

- ICMP—Enter the ICMP type, ICMP code (0-255), and optionally, the ICMP identifier.
- TCP/UDP/SCTP—Enter the source and destination port numbers.
- Raw IP—Enter the protocol number, 0-255.

Step 9 Use packet tracer to debug packets across cluster units. From the **Cluster Capture** drop-down list, select:

a) **decrypted**—Injects a decrypted packet in a VPN tunnel and also simulates a packet that comes across a VPN tunnel.
b) **persist**—Injects the packet you want to track across cluster units.
c) **bypass-checks**—Skips security checks like ACL, VPN filters, IPsec spoof, and uRPF.
d) **transmit**—Allows simulated packets to egress the ASA.

Step 10 Click **Start** to trace the packet.

The **Information Display Area** shows detailed messages about the results of the packet trace.

Monitoring Performance and System Resources

You can monitor a variety of system resources to identify performance or other potential problems.

Monitoring Performance

You can view ASA performance information in a graphical or tabular format.

Procedure

Step 1 Choose **Monitoring** > **Properties** > **Connection Graphs** > **Perfmon**.

Step 2 You can give the graph window a title by entering it in **Graph Window Title**, or you can choose an existing title.

Step 3 Select up to four entries from the Available Graphs list, then click **Add** to move them to the Selected Graphs list. The available options are the following:

- AAA Perfmon—Requests per second for authentication, authorization, and accounting requests.
- Inspection Perfmon—Packets per second for HTTP, FTP, and TCP inspection.
- Web Perfmon—Requests per second for URL access and URL server requests.
- Connections Perfmon—Connections per second for all connections, UDP connections, TCP connections, and TCP Intercept.
- Xlate Perfmon—NAT xlates per second.

Step 4 Click **Show Graphs**.

You can toggle each graph between graph and table views. You can also change how often the data refreshes, and export or print the data.

Monitoring Memory Blocks

You can view free and used memory blocks information in a graphical or tabular format.

Procedure

Step 1 Choose **Monitoring** > **Properties** > **System Resources Graphs** > **Blocks**.

Step 2 You can give the graph window a title by entering it in **Graph Window Title**, or you can choose an existing title.

Step 3 Select entries from the Available Graphs list, then click **Add** to move them to the Selected Graphs list. The available options are the following:

- Blocks Used—Displays the ASA used memory blocks.
- Blocks Free—Displays the ASA free memory blocks.

Step 4 Click **Show Graphs**.

You can toggle each graph between graph and table views. You can also change how often the data refreshes, and export or print the data.

Monitoring CPU

You can view CPU utilization.

Procedure

Step 1 Choose **Monitoring** > **Properties** > **System Resources Graphs** > **CPU**.

Step 2 You can give the graph window a title by entering it in **Graph Window Title**, or you can choose an existing title.

Step 3 Add CPU Utilization to the Selected Graphs list.

Step 4 Click **Show Graphs**.

You can toggle the graph between graph and table views. You can also change how often the data refreshes, and export or print the data.

Monitoring Memory

You can view memory utilization information in a graphical or tabular format.

Procedure

Step 1 Choose **Monitoring** > **Properties** > **System Resources Graphs** > **Memory**.

Step 2 You can give the graph window a title by entering it in **Graph Window Title**, or you can choose an existing title.

Step 3 Select entries from the Available Graphs list, then click **Add** to move them to the Selected Graphs list. The available options are the following:

- Free Memory—Displays the ASA free memory.

- Used Memory—Displays the ASA used memory.

Step 4 Click **Show Graphs**.

You can toggle each graph between graph and table views. You can also change how often the data refreshes, and export or print the data.

Monitoring Per-Process CPU Usage

You can monitor the processes that run on the CPU. You can obtain information about the percentage of CPU that is used by a certain process. CPU usage statistics are sorted in descending order to display the highest consumer at the top. Also included is information about the load on the CPU per process, at 5 seconds, 1 minute, and 5 minutes before the log time. This information is updated automatically every 5 seconds to provide real-time statistics. In ASDM, it is updated every 30 seconds.

To view CPU usage on a per-process basis, choose **Monitoring** > **Properties** > **Per-Process CPU Usage**.

You can stop the auto refresh, manually refresh the information, or save it to a file. You can also click **Configure CPU Usage Colors** to chose background and foreground colors based on usage percentages, to make it easier to scan for high-usage processes.

Monitoring Connections

To view current connections in a tabular format, in the ASDM main window, choose **Monitoring > Properties > Connections**. Information for each connection includes the protocol, source and destination address characteristics, idle time since the last packet was sent or received, and the amount of traffic in the connection.

History for Testing and Troubleshooting

Feature Name	Platform Releases	Description
IPv6 support for traceroute	9.7(1)	The **traceroute** command was modified to accept ar address. We modified the following screen: **Tools** > **Tracero**
Support for the packet tracer for bridge group member interfaces	9.7(1)	You can now use the packet tracer for bridge group interfaces. We added **VLAN ID** and **Destination MAC Addres** the packet-tracer screen:**Tools** > **Packet Tracer**
Manually start and stop packet captures	9.7(1)	You can now manually stop and start the capture. Added/Modified screens: **Wizards** > **Packet Capture Run Captures** Added/Modified options: **Start** button, **Stop** button

Feature Name	Platform Releases	Description
Enhanced packet tracer and packet capture capabilities	9.9(1)	The packet tracer has been enhanced with the follo • Trace a packet when it passes between clust • Allow simulated packets to egress the ASA • Bypass security checks for a similated pack • Treat a simulated packet as an IPsec/SSL dec The packet capture has been enhanced with the f features: • Capture packets after they are decrypted. • Capture traces and retain them in the persist New or modified screens: **Tools** > **Packet Tracer** We added **Cluster Capture** field to support thes **decrypted**, **persist**, **bypass-checks**, **transmit** We added two new options in the **Filter By** view **Sessions** drop-down list: **Origin** and **Origin-ID** **Monitoring** > **VPN** > **VPN Statistics** > **Packet Capture** We added **ICMP Capture** field in the Packet Ca screen:**Wizards** > **Packet Capture Wizard** We added two options **include-decrypted** and **per** ICMP Capture.
Packet capture support for matching IPv6 traffic without using an ACL	9.10(1)	If you use the **match** keyword for the **capture** co **any** keyword only matches IPv4 traffic. You can **any4** and **any6** keywords to capture either IPv4 o The **any** keyword continues to match only IPv4 t New/Modified commands: **capture match** No ASDM support.
New **debug telemetry** command for Forepower 9300/4100.	9.14(1)	If you use the **debug telemetry** command, debug related to telemetry are displayed. The debugs he the cause for errors when generating the telemetr No modified screens.

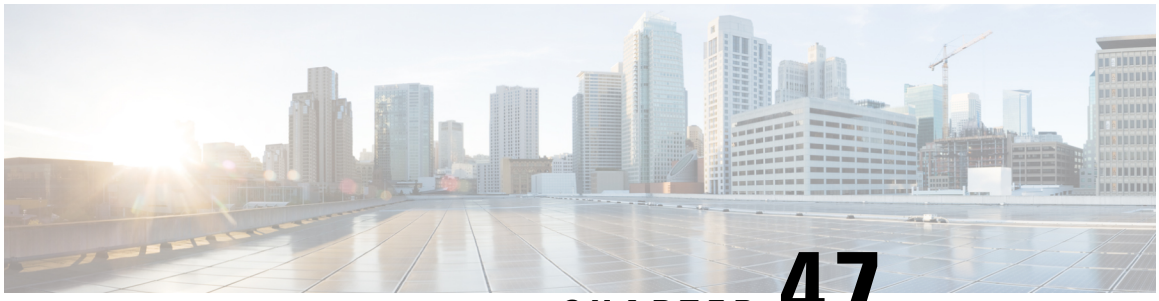
Feature Name	Platform Releases	Description
ping command changes	9.18(2)	To support pinging a loopback interface, the **ping** co now has changed behavior. If you specify the interfa command, the source IP address matches the specifie IP address, but the actual egress interface is determi route lookup using the data routing table. New/Modified commands: **ping**
Packet Capture for switches	9.20(1)	You can now configure to capture egress and ingress packets for a switch. This option is applicable only f Firewall 4200 model devices. New/Modified screens: **Wizards** > **Packet Capture Ingress Traffic Selector**and **Wizards** > **Packet Ca Wizard** > **Egress Traffic Selector**



PART VIII

Monitoring

- Logging, on page 1135
- SNMP, on page 1171
- Cisco Success Network and Telemetry Data, on page 1189
- Anonymous Reporting and Smart Call Home, on page 1199



CHAPTER 47

Logging

This chapter describes how to log system messages and use them for troubleshooting.

- About Logging, on page 1135
- Guidelines for Logging, on page 1143
- Configure Logging, on page 1145
- Monitoring the Logs, on page 1164
- History for Logging, on page 1167

About Logging

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.

The ASA system logs provide you with information for monitoring and troubleshooting the ASA. With the logging feature, you can do the following:

- Specify which syslog messages should be logged.
- Disable or change the severity level of a syslog message.
- Specify one or more locations where syslog messages should be sent, including:
  - An internal buffer
  - One or more syslog servers
  - ASDM
  - An SNMP management station
  - Specified e-mail addresses
  - Console
  - Telnet and SSH sessions.
- Configure and manage syslog messages in groups, such as by severity level or class of message.

- Specify whether or not a rate-limit is applied to syslog generation.
- Specify what happens to the contents of the internal log buffer when it becomes full: overwrite the buffer, send the buffer contents to an FTP server, or save the contents to internal flash memory.
- Filter syslog messages by locations, severity level, class, or a custom message list.

Logging in Multiple Context Mode

Each security context includes its own logging configuration and generates its own messages. If you log in to the system or admin context, and then change to another context, messages you view in your session are only those messages that are related to the current context.

Syslog messages that are generated in the system execution space, including failover messages, are viewed in the admin context along with messages generated in the admin context. You cannot configure logging or view any logging information in the system execution space.

You can configure the ASA to include the context name with each message, which helps you differentiate context messages that are sent to a single syslog server. This feature also helps you to determine which messages are from the admin context and which are from the system; messages that originate in the system execution space use a device ID of **system**, and messages that originate in the admin context use the name of the admin context as the device ID.

Syslog Message Analysis

The following are some examples of the type of information you can obtain from a review of various syslog messages:

- Connections that are allowed by ASA security policies. These messages help you spot holes that remain open in your security policies.
- Connections that are denied by ASA security policies. These messages show what types of activity are being directed toward your secured inside network.
- Using the ACE deny rate logging feature shows attacks that are occurring on your ASA.
- IDS activity messages can show attacks that have occurred.
- User authentication and command usage provide an audit trail of security policy changes.
- Bandwidth usage messages show each connection that was built and torn down as well as the duration and traffic volume used.
- Protocol usage messages show the protocols and port numbers used for each connection.
- Address translation audit trail messages record NAT or PAT connections being built or torn down, which are useful if you receive a report of malicious activity coming from inside your network to the outside world.

Syslog Message Format

Syslog messages are structured as follows:

```
[<PRI>] [Timestamp] [Device-ID] : %ASA-Class-Level-Message_number: Message_text
```

Field descriptions are as follows:

<PRI>	Priority value. When the logging EMBLEM is enabled, this value is displayed in the syslog message. Logging EMBLEM is compatible with UDP and not with TCP.
Timestamp	Date and time of the event is displayed. When logging of timestamps is enabled, and if the timestamp is configured to be in the RFC 5424 format, all timestamp in syslog messages display the time in UTC, as indicated by the RFC 5424 standard.
Device-ID	The device identifier string that was configured while enabling the logging device-id option through the user interface. If enabled, the device ID does not appear in EMBLEM-formatted syslog messages.
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 vpnc class denotes VPN Client.
ASA	The syslog message facility code for messages that are generated by the ASA. This value is always `ASA`.
Level	0 through 7. The level reflects the severity of the condition described by the syslog message—the lower the number, the more severe the condition.
Message_number	A unique six-digit number that identifies the syslog message. All messages are documented in the Cisco Secure Firewall ASA Series Syslog Messages guide.
Message_text	A text string that describes the condition. This portion of the syslog message sometimes includes IP addresses, port numbers, or usernames.

Example of a syslog message with logging EMBLEM, logging timestamp rfc5424, and device-id enabled.

```
<166>2018-06-27T12:17:46Z: %ASA-6-110002: Failed to locate egress interface for protocol
from src interface :src IP/src port to dest IP/dest port
```

Example of a syslog message with logging timestamp rfc5424 and device-id enabled.

```
2018-06-27T12:17:46Z asa : %ASA-6-110002: Failed to locate egress interface for protocol
from src interface :src IP/src port to dest IP/dest port
```

Syslog messages are structured as follows:

```
[<PRI>] [Timestamp] [Device-ID] : %ASA-Level-Message_number: Message_text
```

Field descriptions are as follows:

<PRI>	Priority value. When the logging EMBLEM is enabled, this value is displayed in the syslog message. Logging EMBLEM is compatible with UDP and not with TCP.
Timestamp	Date and time of the event is displayed. When logging of timestamps is enabled, and if the timestamp is configured to be in the RFC 5424 format, all timestamp in syslog messages display the time in UTC, as indicated by the RFC 5424 standard.
Device-ID	The device identifier string that was configured while enabling the logging device-id option through the user interface. If enabled, the device ID does not appear in EMBLEM-formatted syslog messages.
ASA	The syslog message facility code for messages that are generated by the ASA. This value is always `ASA`.

Level	0 through 7. The level reflects the severity of the condition described by the syslog message—the lower the number, the more severe the condition.
Message_number	A unique six-digit number that identifies the syslog message.
Message_text	A text string that describes the condition. This portion of the syslog message sometimes includes IP addresses, port numbers, or usernames.

All syslog messages that are generated by the device are documented in the Cisco Secure Firewall ASA Series Syslog Messages guide.

Example of a syslog message with logging EMBLEM, logging timestamp rfc5424, and device-id enabled.

```
<166>2018-06-27T12:17:46Z: %ASA-6-110002: Failed to locate egress interface for protocol
from src interface :src IP/src port to dest IP/dest port
```

Example of a syslog message with logging timestamp rfc5424 and device-id enabled.

```
2018-06-27T12:17:46Z asa : %ASA-6-110002: Failed to locate egress interface for protocol
from src interface :src IP/src port to dest IP/dest port
```

Severity Levels

The following table lists the syslog message severity levels. You can assign custom colors to each of the severity levels to make it easier to distinguish them in the ASDM log viewers. To configure syslog message color settings, either choose the **Tools > Preferences > Syslog** tab or, in the log viewer itself, click **Color Settings** on the toolbar.

Table 58: 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 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 ASA 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
- Syslog message severity level
- Syslog message class (equivalent to a functional area)

You customize these criteria by creating a message list that you can specify when you set the output destination. Alternatively, you can configure the ASA to send a particular message class to each type of output destination independently of the message list.

Syslog Message Classes

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 59: Syslog Message Classes and Associated Message ID Numbers

Class	Definition	Syslog Message ID Numbers
auth	User Authentication	109, 113
—	Access Lists	106
—	Application Firewall	415
—	Botnet Traffic Filtering	338
bridge	Transparent Firewall	110, 220
ca	PKI Certification Authority	717
citrix	Citrix Client	723
—	Clustering	747
—	Card Management	323
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	735
ha	Failover	101, 102, 103, 104, 105, 210, 311, 709
—	Identity-based Firewall	746
ids	Intrusion Detection System	400, 733
—	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	325
—	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	305
—	Network Access Point	713
np	Network Processor	319
—	NP SSL	725
ospf	OSPF Routing	318, 409, 503, 613
—	Password Encryption	742
—	Phone Proxy	337
rip	RIP Routing	107, 312
rm	Resource Manager	321
—	Smart Call Home	120
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
snmp	SNMP	212
—	ScanSafe	775
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
—	Threat Detection	733
tag-switching	Service Tag Switching	779
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	778
webfo	WebVPN Failover	721
webvpn	WebVPN and Secure Client	716

Sort Messages in the Log Viewers

You can sort messages in all ASDM log viewers (that is, the Real-Time Log Viewer, the Log Buffer Viewer, and the Latest ASDM Syslog Events Viewer). To sort tables by multiple columns, click the header of the first column that you want to sort by, then press and hold down the **Ctrl** key and at the same time, click the headers of the other column(s) that you want to include in the sort order. To sort messages chronologically, select both the date and time columns; otherwise, the messages are sorted only by date (regardless of the time) or only by time (regardless of the date).

When you sort messages in the Real-Time Log Viewer and in the Latest ASDM Syslog Events Viewer, the new messages that come in appear in the sorted order, instead of at the top, as they normally would be. That is, they are mixed in with the rest of the messages.

Custom Message Lists

Creating a custom message list is a flexible way to exercise control over which syslog messages are sent to which output destination. In a custom syslog message list, you specify groups of syslog messages using any or all of the following criteria:

- Severity level
- Message IDs
- Ranges of syslog message IDs
- Message class.

For example, you can use message lists to do the following:

- Select syslog messages with the severity levels of 1 and 2 and send them to one or more e-mail addresses.
- Select all syslog messages associated with a message class (such as ha) and save them to the internal buffer.

A message list can include multiple criteria for selecting messages. However, you must add each message selection criterion with a new command entry. It is possible to create a message list that includes overlapping message selection criteria. If two criteria in a message list select the same message, the message is logged only once.

Clustering

Syslog messages are an invaluable tool for accounting, monitoring, and troubleshooting in a clustering environment. Each ASA unit in the cluster (up to eight units are allowed) generates syslog messages independently; certain **logging** commands then enable you to control header fields, which include a time stamp and device ID. The syslog server uses the device ID to identify the syslog generator. You can use the **logging device-id** command to generate syslog messages with identical or different device IDs to make messages appear to come from the same or different units in the cluster.



Note To monitor syslog messages from units in a cluster, you must open an ASDM session to each of the units that you want to monitor.

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

- 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 ASA, you must specify a logging output destination. If you enable logging without specifying a logging output destination, the ASA 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. For example, to designate more than one syslog server as an output destination, specify separate entries in the **Syslog Server** pane for each syslog server.
- Sending syslogs over TCP is not supported on a standby device.

- 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. However, in multiple context mode, the limitation is 4 servers per context.
- The syslog server should be reachable through the ASA. 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 you use a custom message list to match only access list hits, the access list logs are not generated for access lists that have had their logging severity level increased to debugging (level 7). The default logging severity level is set to 6 for the **logging list** command. This default behavior is by design. When you explicitly change the logging severity level of the access list configuration to debugging, you must also change the logging configuration itself.

  The following is sample output from the **show running-config logging** command that does not include access list hits, because their logging severity level has been changed to debugging:

```
ciscoasa# show running-config logging
logging enable
logging timestamp
logging list test message 106100
logging buffered test
```

  The following is sample output from the **show running-config logging** command that does include access list hits:

```
ciscoasa# show running-config logging
logging enable
logging timestamp
logging buffered debugging
```

  In this case, the access list configuration does not change and the number of access list hits appears, as shown in the following example:

```
ciscoasa(config)# access-list global line 1 extended
permit icmp any host 4.2.2.2 log debugging interval 1 (hitcnt=7) 0xf36b5386
ciscoasa(config)# access-list global line 2 extended
permit tcp host 10.1.1.2 any eq www log informational interval 1 (hitcnt=18) 0xe7e7c3b8
ciscoasa(config)# access-list global line 3 extended
permit ip any any (hitcnt=543) 0x25f9e609
```

- When the ASA sends syslogs via TCP, the connection takes about one minute to initiate after the syslogd service restarts.
- The server certificate received from a Syslog Server must contain "ServAuth" in the Extended Key Usage field. This check will be done on non self-signed certificates only, self-signed certificates do not provide any value in this field.

Configure Logging

This section describes how to configure logging.

Enable Logging

To enable logging, perform the following steps:

Procedure

Step 1 In ASDM, choose one of the following:

- **Home** > **Latest ASDM Syslog Messages** > **Enable Logging**
- **Configuration** > **Device Management** > **Logging** > **Logging Setup**
- **Monitoring** > **Real-Time Log Viewer** > **Enable Logging**
- **Monitoring** > **Log Buffer** > **Enable Logging**

Step 2 Check the **Enable logging** check box to turn on logging.

Configure an Output Destination

To optimize syslog message usage for troubleshooting and performance monitoring, we recommend that you specify one or more locations where syslog messages should be sent, including an internal log buffer, one or more external syslog servers, ASDM, an SNMP management station, the console port, specified e-mail addresses, or Telnet and SSH sessions.

When you configure syslog logging on an interface with management-only access enabled, the dataplane related logs (syslog IDs 302015, 302014, 106023, and 304001) are dropped and does not reach the syslog server. The syslog messages are dropped because the datapath routing table does not have the management interface routing. Hence, ensure the interface that you are configuring has management-only access disabled

Send Syslog Messages to an External Syslog Server

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.

To send syslog messages to an external syslog server, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Logging** > **Logging Setup**.

Step 2 Check the **Enable logging** check box to turn on logging for the ASA.

Step 3 Check the **Enable logging on the failover standby unit** check box to turn on logging for the standby ASA, if available.

Step 4 Check the **Send debug messages as syslogs** check box to redirect all debugging trace output to system logs. The syslog message does not appear on the console if this option is enabled. Therefore, to view debugging messages, you must have logging enabled at the console and have it configured as the destination for the debugging syslog message number and severity level. The syslog message number to use is **711001**. The default severity level for this syslog message is debugging.

Step 5 Check the **Send syslogs in EMBLEM format** check box to enable EMBLEM format so that it is used for all logging destinations, except syslog servers.

Step 6 Specify the size of the internal log buffer to which syslog messages are saved if the logging buffer is enabled. When the buffer fills up, messages are overwritten unless you save the logs to an FTP server or to internal flash memory. The default buffer size is 4096 bytes. The range is 4096 to 1048576.

Step 7 To save the buffer content to the FTP server before it is overwritten, check the **Save Buffer To FTP Server** check box. To allow overwriting of the buffer content, uncheck this check box.

Step 8 Click **Configure FTP Settings** to identify the FTP server and configure the FTP parameters used to save the buffer content.

Step 9 Check the **Save Buffer To Flash** check box To save the buffer content to internal flash memory before it is overwritten.

Note
This option is only available in routed or transparent single mode.

Step 10 Click **Configure Flash Usage** to specify the maximum space to be used in internal flash memory for logging and the minimum free space to be preserved (in KB). Enabling this option creates a directory called “syslog” on the device disk on which messages are stored.

Note
This option is only available in single routed or transparent mode.

Step 11 Specify the queue size for system logs that are to be viewed in the ASA.

Configure FTP Settings

To specify the configuration for the FTP server that is used to save the log buffer content, perform the following steps:

Procedure

Step 1 Check the **Enable FTP client** check box to enable configuration of the FTP client.

Step 2 Specify the IP address of the FTP server.

Step 3 Specify the directory path on the FTP server to store the saved log buffer content.

Step 4 Specify the username to log in to the FTP server.

Step 5 Specify the password associated with the username to log in to the FTP server.

Step 6 Confirm the password, then click **OK**.

Configure Logging Flash Usage

To specify the limits for saving the log buffer content to internal flash memory, perform the following steps:

Procedure

Step 1 Specify the maximum amount of internal flash memory that can be used for logging (in KB).

Step 2 Specify the amount of internal flash memory that is preserved (in KB). When the internal flash memory approaches that limit, new logs are no longer saved.

Step 3 Click **OK** to close the **Configure Logging Flash Usage** dialog box.

Enable Secure Logging

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Logging** > **Syslog Server**.

Step 2 Select a syslog server for which you want to enable secure logging, then click **Edit**.

The **Edit Syslog Server** dialog box appears.

Step 3 Click the **TCP** radio button.

Secure logging does not support UDP; an error occurs if you try to use this protocol.

Step 4 Check the **Enable secure syslog with SSL/TLS** check box, then click **OK**.

Step 5 (Optional) Specify a **Reference Identity** object by name to enable RFC 6125 reference identity checks on the certificate received from the Syslog server.

See Configure Reference Identities, on page 741 for details on the reference identity object.

Generate Syslog Messages in EMBLEM Format to a Syslog Server

To generate syslog messages in EMBLEM format to a syslog server, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Logging** > **Syslog Server**.

Sending syslogs over IPv6 is supported.

Step 2 Click **Add** to add a new syslog server.

The **Add Syslog Server** dialog box appears.

Note
You can set up a maximum of four syslog servers per security context (up to a total of 16).

Step 3 Specify the number of messages that are allowed to be queued on the ASA when a syslog server is busy. A zero value means an unlimited number of messages may be queued.

Step 4 Check the **Allow user traffic to pass when TCP syslog server is down** check box to allow all traffic if any syslog server is down.

When the ASA is configured to send syslog messages to a TCP-connected syslog server, and if the syslog server fails, as a security protection, new connections through the ASA are blocked. To permit new connections, even when the syslog server is not operational, select this check box.

If you specify UDP, the ASA continues to allow new connections whether or not the syslog server is operational. Valid port values for either protocol are 1025 through 65535. The default UDP port is 514. The default TCP port is 1470.

Note
Sending syslogs over TCP is not supported on a standby ASA.

Generate Syslog Messages in EMBLEM Format to Other Output Destinations

To generate syslog messages in EMBLEM format to other output destinations, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Logging** > **Logging Setup**.

Step 2 Check the **Send syslogs in EMBLEM format** check box.

Add or Edit Syslog Server Settings

To add or edit syslog server settings, perform the following steps:

Procedure

Step 1 Choose the interface used to communicate with the syslog server from the drop-down list.

Step 2 Enter the IP address that is used to communicate with the syslog server.

Choose the protocol (either TCP or UDP) that is used by the syslog server to communicate with the ASA or ASASM. You can configure the ASA and ASASM to send data to a syslog server using either UDP or TCP. The default protocol is UDP if you do not specify a protocol.

Warning
If you specify TCP, when the ASA discovers syslog server failures, for security reasons, new connections through the ASA are blocked. To permit new connections despite syslog server failures, see Step 4 of Generate Syslog Messages in EMBLEM Format to a Syslog Server, on page 1147.

Step 3 Enter the port number used by the syslog server to communicate with the ASA or ASASM.

Step 4 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).

Step 5 Check the **Enable secure logging using SSL/TLS (TCP only)** check box to specify that the connection to the syslog server is secure through the use of SSL/TLS over TCP, and that the syslog message content is encrypted. You can optionally mention reference identity to validate the certificate based on the previously configured reference identity object. For more information, see Enable Secure Logging, on page 1147.

Step 6 Click **OK** to complete the configuration.

Send Syslog Messages to the Internal Log Buffer

You need to specify which syslog messages should be sent to the internal log buffer, which serves as a temporary storage location. New messages are appended to the end of the list. When the buffer is full, that is, when the buffer wraps, old messages are overwritten as new messages are generated, unless you configure the ASA to save the full buffer to another location.

To send syslog messages to the internal log buffer, perform the following steps:

Procedure

Step 1 Choose one of the following options to specify which syslog messages should be sent to the internal log buffer:

- **Home** > **Latest ASDM Syslog Messages** > **Configure ASDM Syslog Filters**
- **Configuration** > **Device Management** > **Logging** > **Logging Filters**

Step 2 Choose **Monitoring** > **Logging** > **Log Buffer** > **View**. Then choose **File** > **Clear Internal Log Buffer** in the **Log Buffer** pane to empty the internal log buffer.

Step 3 Choose **Configuration** > **Device Management** > **Logging** > **Logging Setup** to change the size of the internal log buffer. The default buffer size is 4 KB.

The ASA continue to save new messages to the internal log buffer and save the full log buffer content to internal flash memory. When saving the buffer content to another location, the ASA create log files with names that use the following time-stamp format:

```
LOG-YYYY-MM-DD-HHMMSS.TXT
```

where *YYYY* is the year, *MM* is the month, *DD* is the day of the month, and *HHMMSS* is the time in hours, minutes, and seconds.

Step 4 To save new messages to another location, choose one of the following options:

- Check the **Flash** check box to send new messages to internal flash memory, then click Configure **Flash Usage**. The **Configure Logging Flash Usage** dialog box appears.
  - **a.** Specify the maximum amount of flash memory in KB that you want to use for logging.
  - **b.** Specify the minimum amount of free space in KB that logging will preserve in flash memory.
  - **c.** Click **OK** to close this dialog box.
- Check the **FTP Server** check box to send new messages to an FTP server, then click **Configure FTP Settings**. The **Configure FTP Settings** dialog box appears.
  - **a.** Check the **Enable FTP Client** check box.
  - **b.** Enter the following information in the fields provided: FTP server IP address, path, username, and password.
  - **c.** Confirm the password, then click **OK** to close this dialog box.

Save an Internal Log Buffer to Flash

To save the internal log buffer to flash memory, perform the following steps:

Procedure

Step 1 Choose **File** > **Save Internal Log Buffer to Flash**.

The **Enter Log File Name** dialog box appears.

Step 2 Choose the first option to save the log buffer with the default filename, LOG-YYYY-MM-DD-hhmmss.txt.

Step 3 Choose the second option to specify a filename for the log buffer.

Step 4 Enter the filename for the log buffer, then click **OK**.

Change the Amount of Internal Flash Memory Available for Logs

To change the amount of internal flash memory available for logs, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Logging** > **Logging Setup**.

Step 2 Check the **Enable Logging** check box.

Step 3 Check the **Save Buffer to Flash** check box in the **Logging to Internal Buffer** area.

Step 4 Click **Configure Flash Usage**.

The **Configure Logging Flash Usage** dialog box appears.

Step 5 Enter the maximum amount of flash memory in KB allowed to be used for logging.

By default, the ASA can use up to 50 MB of internal flash memory for log data. The minimum amount of internal flash memory that must be free for the ASA to save log data is 3 MB. If a log file being saved to internal flash memory would cause the amount of free internal flash memory to fall below the configured minimum limit, the ASA deletes the oldest log files to ensure that the minimum amount of memory remains free after saving the new log file. If there are no files to delete or if, after all old files have been deleted, free memory is still below the limit, the ASA fails to save the new log file. The maximum limit of the flash-maximum-allocation value is 2 GB.

Step 6 Enter the minimum amount of free space in KB to be preserved for logging in flash memory.

Step 7 Click **OK** to close the **Configure Logging Flash Usage** dialog box.

View and Copy Logged Entries with the ASDM Java Console

Use the ASDM Java console to view and copy logged entries in a text format, which may help you troubleshoot ASDM errors.

To access the ASDM Java Console, perform the following steps:

Procedure

Step 1 Choose **Tools** > **ASDM Java Console**.

Step 2 Enter **m** in the console to show the virtual machine memory statistics.

Step 3 Enter **g** in the console to perform garbage collection.

Step 4 Open the Windows Task Manager and double-click the **asdm_launcher.exe** file to monitor memory usage.

Note
The maximum memory allocation allowed is 256 MB.

Send Syslog Messages to an E-mail Address

To send syslog messages to an e-mail address, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Logging** > **E-Mail Setup**.

Step 2 Specify the e-mail address that is used as the source address for syslog messages that are sent as e-mail messages.

Step 3 Click **Add** to enter a new e-mail address recipient of the specified syslog messages.

Step 4 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. The global filter specified in the **Logging Filters** pane is also applied to each e-mail recipient.

Step 5 Click **Edit** to modify an existing severity level of the syslog messages that are sent to this recipient.

Step 6 Click **OK** to close the **Add E-mail Recipient** dialog box.

Add or Edit E-Mail Recipients

To add or edit e-mail recipients and severity levels, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Logging** > **E-mail Setup**.

Step 2 Click **Add** or **Edit** to display the **Add/Edit E-Mail Recipient** dialog box.

Step 3 Enter the destination e-mail address, and choose the syslog severity level from the drop-down list. Severity levels are defined as follows:

- Emergency (level 0, system is unusable)

  Note
  Using a severity level of zero is not recommended.

- Alert (level 1, immediate action is needed)
- Critical (level 2, critical conditions)
- Error (level 3, error conditions)
- Warning (level 4, warning conditions)
- Notification (level 5, normal but significant conditions)
- Informational (level 6, informational messages only)
- Debugging (level 7, debugging messages only)

Note
The severity level used to filter messages for the destination e-mail address is the higher of the severity level specified in the **Add/Edit E-Mail Recipient** dialog box and the global filter set for all e-mail recipients in the **Logging Filters** pane.

Step 4 Click **OK** to close the **Add/Edit E-Mail Recipient** dialog box.

The added or revised entry appears in the **E-mail Recipients** pane.

Step 5 Click **Apply** to save your changes to the running configuration.

Configure the Remote SMTP Server

To configure the remote SMTP server to which e-mail alerts and notifications are sent in response to specific events, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Device Setup** > **Logging** > **SMTP**.

Step 2 Enter the IP address of the primary SMTP server.

Step 3 (Optional) Enter the IP address of the standby SMTP server, then click **Apply** to save your changes to the running configuration.

Send Syslog Messages to the Console Port

To send syslog messages to the console port, perform the following steps:

Procedure

Step 1 Choose one of the following options:

- **Home** > **Latest ASDM Syslog Messages** > **Configure ASDM Syslog Filters**
- **Configuration** > **Device Management** > **Logging** > **Logging Filters**

Step 2 Select the console in the **Logging Destination** column, then click **Edit**.

The **Edit Logging Filters** dialog box appears.

Step 3 Choose either syslogs from all event classes or syslogs from specific event classes to specify which syslog messages should be sent to the console port.

Send Syslog Messages to a Telnet or SSH Session

To send syslog messages to a Telnet or SSH session, perform the following steps:

Procedure

Step 1 Choose one of the following options:

- **Home** > **Latest ASDM Syslog Messages** > **Configure ASDM Syslog Filters**
- **Configuration** > **Device Management** > **Logging** > **Logging Filters**

Step 2 Select the **Telnet** and **SSH Sessions** in the **Logging Destination** column, then click **Edit**.

The **Edit Logging Filters** dialog box appears.

Step 3 Choose either syslogs from all event classes or syslogs from specific event classes to specify which syslog messages should be sent to a Telnet or an SSH session..

Step 4 Choose **Configuration** > **Device Management** > **Logging** > **Logging Setup** to enable logging for the current session only.

Step 5 Check the **Enable logging** check box, then click **Apply**.

Configure Syslog Messages

Configure Syslog Messaging

To configure syslog messaging, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Logging** > **Syslog Setup**.

Step 2 Choose a system log facility for syslog servers to use as a basis to file messages. The default is LOCAL(4)20, which is what most UNIX systems expect. However, because your network devices share eight available facilities, you might need to change this value for system logs.

Step 3 Check the **Include timestamp in syslogs** check box to add the date and time in each syslog message sent.

Use the **Timestamp Format** drop-down to select the legacy (mm:dd:yyyy hh:mm:ss) or RFC 5424 (yyyy:dd:mmTHH:mm:ssZ) format.

Step 4 Uncheck the **Hide username if its validity cannot be determined** check box to show invalid usernames in syslog messages for unsuccessful login attempts. The default setting is to hide usernames when the username is invalid or if the validity is unknown. If a user accidentally types a password instead of a username, for example, then it is more secure to hide the “username” in the resultant syslog message. You might want to show invalid usernames to help with troubleshooting login issues.

Step 5 Choose the information to be displayed in the **Syslog ID** table. Available options are as follows:

- Choose **Show all syslog IDs** to specify that the **Syslog ID** table should display the entire list of syslog message IDs.
- Choose **Show disabled syslog IDs** to specify that the **Syslog ID** table should display only those syslog message IDs that have been explicitly disabled.
- Choose **Show syslog IDs with changed logging** to specify that the **Syslog ID** table should display only those syslog message IDs with severity levels that have changed from their default values.
- Choose **Show syslog IDs that are disabled or with a changed logging level** to specify that the **Syslog ID** table should display only those syslog message IDs with severity levels that have been modified and the IDs of syslog messages that have been explicitly disabled.

Step 6 The **Syslog ID Setup Table** displays the list of syslog messages based on the setting in the Syslog ID Setup Table. Choose individual messages or ranges of message IDs that you want to modify. You can either disable the selected message IDs or modify their severity levels. To select more than one message ID in the list, click the first ID in the range and Shift-click the last ID in the range.

Step 7 Click **Advanced** to configure syslog messages to include a device ID.

Edit Syslog ID Settings

To change syslog message settings, perform the following steps:



Note The **Syslog ID**(s) field is display-only. The values that appear in this area are determined by the entries you chose in the **Syslog ID** table, located in the **Syslog Setup** pane.

Procedure

Step 1 Check the **Disable Message(s)** check box to disable messages for the syslog message ID(s) displayed in the **Syslog ID**(s) list.

Step 2 Choose the severity logging level of messages to be sent for the syslog message ID(s) displayed in the **Syslog ID**(s) list. Severity levels are defined as follows:

- Emergency (level 0, system is unusable)

  Note
  Using a severity level of zero is not recommended.

- Alert (level 1, immediate action is needed)
- Critical (level 2, critical conditions)
- Error (level 3, error conditions)
- Warning (level 4, warning conditions)
- Notification (level 5, normal but significant conditions)
- Informational (level 6, informational messages only)
- Debugging (level 7, debugging messages only)

Step 3 Click **OK** to close the **Edit Syslog ID Settings** dialog box.

Include a Device ID in Non-EMBLEM Formatted Syslog Messages

To include a device ID in non-EMBLEM formatted syslog messages, perform the following steps:

Procedure

Step 1 Check the **Enable syslog device ID** check box to specify that a device ID should be included in all non-EMBLEM formatted syslog messages.

Step 2 To specify which to use as the device ID, choose one of the following options:

- Hostname of the ASA
- Interface IP address

  Choose the interface name that corresponds to the selected IP address from the drop-down list.

  Check the **In an ASA cluster, always use control's IP address for the selected interface** check box if you are using clustering.
- String

  Specify an alphanumeric, user-defined string.
- ASA cluster name

Step 3 Click **OK** to close the **Advanced Syslog Configuration** dialog box.

Include the Date and Time in Syslog Messages

To include the date and time in syslog messages, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Logging** > **Syslog Setup**.

Step 2 Check the **Include timestamp in syslogs** check box in the **Syslog ID Setup** area.

Step 3 Click **Apply** to save your changes.

Disable a Syslog Message

To disable a specified syslog message, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Logging** > **Syslog Setup**.

Step 2 Select the syslog that you want to disable from the table, then click **Edit**.

The **Edit Syslog ID Settings** dialog box appears.

Step 3 Check the **Disable messages** check box, then click **OK**.

Change the Severity Level of a Syslog Message

To change the severity level of a syslog message, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Logging** > **Syslog Setup**.

Step 2 Select the syslog whose severity level you want to change from the table, then click **Edit**.

The **Edit Syslog ID Settings** dialog box appears.

Step 3 Choose the desired severity level from the **Logging Level** drop-down list, then click **OK**.

Block Syslog Messages on a Standby Unit

To block specific syslog messages from being generated on a standby unit, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Logging** > **Syslog Settings**.

Step 2 Choose a syslog ID in the table, then click **Edit**.

The **Edit Syslog ID Settings** dialog box appears.

Step 3 Check the **Disable messages on standby unit** check box to block syslog messages from being generated on a standby unit.

Step 4 Click **OK** to close this dialog box.

Include the Device ID in Non-EMBLEM Format Syslog Messages

To include the device ID in non-EMBLEM format syslog messages, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Logging** > **Syslog Setup** > **Advanced** > **Advanced Syslog Configuration.**

Step 2 Check the **Enable syslog device ID** check box.

Step 3 Click the **Hostname**, **Interface IP Address**, or **String** radio button in the **Device ID** area.

- If you chose the **Interface IP Address** option, make sure that the correct interface is selected in the drop-down list.
- If you chose the **String** option, enter the device ID in the **User-Defined ID** field. The string can include as many as 16 characters.

Note
If enabled, the device ID does not appear in EMBLEM-formatted syslog messages nor in SNMP traps.

Step 4 Click **OK** to close the **Advanced Syslog Configuration** dialog box.

Create a Custom Event List

You use the following three criteria to define an event list:

- Event Class
- Severity
- Message ID

To create a custom event list to send to a specific logging destination (for example, an SNMP server), perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Logging** > **Event Lists**.

Step 2 Click **Add** to display the **Add Event List** dialog box.

Step 3 Enter the name of the event list. No spaces are allowed.

Step 4 Click **Add** to display the **Add Class and Severity Filter** dialog box.

Step 5 Choose the event class from the drop-down list. Available event classes change according to the device mode that you are using.

Step 6 Choose the severity level from the drop-down list. Severity levels include the following:

- Emergency (level 0, system is unusable)

  Note
  Using a severity level of zero is not recommended.

- Alert (level 1, immediate action is needed)
- Critical (level 2, critical conditions)
- Error (level 3, error conditions)
- Warning (level 4, warning conditions)
- Notification (level 5, normal but significant conditions)
- Informational (level 6, informational messages only)
- Debugging (level 7, debugging messages only)

Step 7 Click **OK** to close the **Add Event List** dialog box.

Step 8 Click **Add** to display the **Add Syslog Message ID Filter** dialog box.

Step 9 Enter a syslog message ID or range of IDs (for example, 101001-199012) to include in the filter.

Step 10 Click **OK** to close the **Add Event List** dialog box.

The event of interest appears in the list.

Configure Logging Filters

Apply Message Filters to a Logging Destination

To apply message filters to a logging destination, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Logging** > **Logging Filters**.

Step 2 Choose the name of the logging destination to which you want to apply a filter. Available logging destinations are as follows:

- ASDM
- Console port
- E-Mail
- Internal buffer
- SNMP server
- Syslog server
- Telnet or SSH session

Included in this selection are the second column, Syslogs From All Event Classes, and the third column, Syslogs From Specific Event Classes. The second column lists the severity or the event class to use to filter messages for the logging destination, or whether logging is disabled for all event classes. The third column lists the event class to use to filter messages for that logging destination.

Step 3 Click **Edit** to display the **Edit Logging Filters** dialog box. To apply, edit, or disable filters, see Apply Logging Filters, on page 1159.

Apply Logging Filters

To apply filters, perform the following steps:

Procedure

Step 1 Choose the **Filter on severity** option to filter syslog messages according to their severity level.

Step 2 Choose the **Use event list** option to filter syslog messages according to an event list.

Step 3 Choose the **Disable logging from all event classes** option to disable all logging to the selected destination.

Step 4 Click **New** to add a new event list. To add a new event list, see Create a Custom Event List, on page 1158.

Step 5 Choose the event class from the drop-down list. Available event classes change according to the device mode that you are using.

Step 6 Choose the level of logging messages from the drop-down list. Severity levels include the following:

- Emergency (level 0, system is unusable)

  Note
  Using a severity level of zero is not recommended.

- Alert (level 1, immediate action is needed)
- Critical (level 2, critical conditions)
- Error (level 3, error conditions)
- Warning (level 4, warning conditions)
- Notification (level 5, normal but significant conditions)
- Informational (level 6, informational messages only)
- Debugging (level 7, debugging messages only)

Step 7 Click **Add** to add the event class and severity level, then click **OK**.

The selected logging destination for a filter appears at the top.

Add or Edit a Syslog Message ID Filter

To add or edit a syslog message ID filter, see Edit Syslog ID Settings, on page 1155.

Add or Edit a Message Class and Severity Filter

To add or edit a message class and severity level for filtering messages, perform the following steps:

Procedure

Step 1 Choose the event class from the drop-down list. Available event classes change according to the device mode that you are using.

Step 2 Choose the level of logging messages from the drop-down list. Severity levels include the following:

- Emergency (level 0, system is unusable)

  Note
  Using a severity level of zero is not recommended.

- Alert (level 1, immediate action is needed)
- Critical (level 2, critical conditions)
- Error (level 3, error conditions)

- Warning (level 4, warning conditions)
- Notification (level 5, normal but significant conditions)
- Informational (level 6, informational messages only)
- Debugging (level 7, debugging messages only)

Step 3 Click **OK** when you are done making selections.

Send All Syslog Messages in a Class to a Specified Output Destination

To send all syslog messages in a class to a specified output destination, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Logging** > **Logging Filters**.

Step 2 To override the configuration in the specified output destination, choose the output destination that you want to change, then click **Edit**.

The **Edit Logging Filters** dialog box appears.

Step 3 Revise the settings in either the **Syslogs from All Event Classes** or **Syslogs from Specific Event Classes** area, then click **OK** to close this dialog box.

For example, if you specify that messages at severity level 7 should go to the internal log buffer and that ha class messages at severity level 3 should go to the internal log buffer, then the latter configuration takes precedence.

To specify that a class should go to more than one destination, select a different filtering option for each output destination.

Limit the Rate of Syslog Message Generation

To limit the rate of syslog message generation, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Logging** > **Rate Limit**.

Step 2 Choose the logging level (message severity level) to which you want to assign rate limits. Severity levels are defined as follows:

- Emergency (level 0, system is unusable)
- Alert (level 1, immediate action is needed)
- Critical (level 2, critical conditions)

- Error (level 3, error conditions)
- Warning (level 4, warning conditions)
- Notification (level 5, normal but significant conditions)
- Informational (level 6, informational messages only)
- Debugging (level 7, debugging messages only)

Step 3 The **No of Messages** field displays the number of messages sent. The **Interval** (Seconds) field displays the interval, in seconds, that is used to limit how many messages at this logging level can be sent. Choose a logging level from the table and click **Edit** to display the **Edit Rate Limit for Syslog Logging Level** dialog box.

Step 4 To continue, see Assign or Change Rate Limits for Individual Syslog Messages, on page 1162.

Assign or Change Rate Limits for Individual Syslog Messages

To assign or change rate limits to individual syslog messages, perform the following steps:

Procedure

Step 1 To assign the rate limit of a specific syslog message, click **Add** to display the **Add Rate Limit for Syslog Message** dialog box.

Step 2 To continue, see Add or Edit the Rate Limit for a Syslog Message, on page 1162.

Step 3 To change the rate limit of a specific syslog message, click **Edit** to display the **Edit Rate Limit for Syslog Message** dialog box.

Step 4 To continue, see Edit the Rate Limit for a Syslog Severity Level, on page 1163.

Add or Edit the Rate Limit for a Syslog Message

To add or change the rate limit for a specific syslog message, perform the following steps:

Procedure

Step 1 To add a rate limit to a specific syslog message, click **Add** to display the **Add Rate Limit for Syslog Message** dialog box. To change a rate limit for a syslog message, click **Edit** to display the **Edit Rate Limit for Syslog Message** dialog box.

Step 2 Enter the message ID of the syslog message that you want to limit.

Step 3 Enter the maximum number of messages that can be sent in the specified time interval.

Step 4 Enter the amount of time, in seconds, that is used to limit the rate of the specified message, then click **OK**.

Note

To allow an unlimited number of messages, leave both the **Number of Messages** and **Time Interval** fields blank.

Edit the Rate Limit for a Syslog Severity Level

To change the rate limit of a specified syslog severity level, perform the following steps:

Procedure

Step 1 Enter the maximum number of messages at this severity level that can be sent.

Step 2 Enter the amount of time, in seconds, that is used to limit the rate of messages at this severity level, and click **OK**.

The selected message severity level appears.

Note
To allow an unlimited number of messages, leave both the **Number of Messages** and **Time Interval** fields blank.

Assign or Change Rate Limits for Dynamic Logging

You can assign rate limits for logging based on used resources (block size). By specifying a threshold value (percentage), the rate of syslog message generation is limited. You can further define the number of the messages permitted to be generated when the block size usage exceeds the threshold value.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Logging** > **Rate Limit**.

Step 2 Under **Rate Limits for Dynamic Logging**, specify the following:

- **Block**—Specify the percentage of free blocks that act as the threshold to trigger the dynamic rate-limit.
- **Message Limit**—Specify the number of messages allowed for the dynamic rate-limit. The default value is 10.

Step 3 Click **Apply**.

Step 4 To modify the saved values, enter the new values, and then click **Apply**.

Step 5 To disable the dynamic logging rate-limit, leave the fields blank.

Monitoring the Logs

See the following commands for monitoring logging status.

- **Monitoring** > **Logging** > **Log Buffer** > **View**

  This pane allows you to view the log buffer.

- **Monitoring** > **Logging** > **Real-Time Log Viewer** > **View**

  This pane allows you to view the real-time log.

- **Tools** > **Command Line Interface**

  This pane allows you to issue various non-interactive commands and view results.

- **Configuration** > **Firewall** > **Access Rules**

  This pane allows you to filter the live viewer of logging to that specific logs based on the search criteria (Rule Hex Id). To view the results, select the rule and click **Show Log**.

Filter Syslog Messages Through the Log Viewers

You can filter syslog messages based on one or multiple values that correspond to any column of the Real-Time Log Viewer and the Log Buffer Viewer.

To filter syslog messages through one of the log viewers, perform the following steps:

Procedure

Step 1 Choose one of the following options:

- **Monitoring** > **Logging** > **Real-Time Log Viewer** > **View**
- **Monitoring** > **Logging** > **Log Buffer** > **View**

Step 2 In either the **Real-Time Log Viewer** or the **Log Buffer Viewer** dialog box, click **Build Filter** on the toolbar.

Step 3 In the **Build Filter** dialog box, specify the filtering criteria to apply to syslog messages:

a) Choose one of the following three options in the **Date and Time** area: real-time, a specific time, or a time range. If you chose a specific time, indicate the time by entering the number and choosing hours or minutes from the drop-down list. If you chose a time range, click the drop-down arrow in the **Start Time** field to display a calendar. Choose a start date and a start time from the drop-down list, then click **OK**. Click the drop-down arrow in the **End Time** field to display a calendar. Choose an end date and an end time from the drop-down list, then click **OK**.

b) Enter a valid severity level in the **Severity** field. Alternatively, click the **Edit** icon on the right of the **Severity** field. Click the severity levels in the list on which you want to filter. To include severity levels 1-7, click **All**. Click **OK** to display these settings in the **Build Filter** dialog box. Click the **Info** icon on the right of the **Severity** field for additional information about the correct input format to use.

c) Enter a valid syslog ID in the **Syslog ID** field. Alternatively, click the **Edit** icon on the right of the **Syslog ID** field. Choose a condition on which to filter from the drop-down list, then click **Add**. Click **OK** to

display these settings in the **Build Filter** dialog box. Click the **Info** icon on the right of the **Syslog ID** field for additional information about the correct input format to use.

d) Enter a valid source IP address in the **Source IP Address** field, or click the **Edit** icon on the right of the **Source IP Address** field. Choose a single IP address or a specified range of IP addresses, then click **Add**. Check the **Do not include (exclude) this address or range** check box to exclude a specific IP address or range of IP addresses, Click **OK** to display these settings in the **Build Filter** dialog box. Click the **Info** icon on the right of the **Source IP Address** field for additional information about the correct input format to use.

e) Enter a valid source port in the **Source Port** field, or click the **Edit** icon on the right of the **Source Port** field. Choose a condition on which to filter from the drop-down list, then click **Add**. Click **OK** to display these settings in the **Build Filter** dialog box. Click the **Info** icon on the right of the **Source Port** field for additional information about the correct input format to use.

f) Enter a valid destination IP address in the **Destination IP Address** field, or click the **Edit** icon on the right of the **Destination IP Address** field. Choose a single IP address or a specified range of IP addresses, then click **Add**. Check the **Do not include (exclude) this address or range** check box o exclude a specific IP address or range of IP addresses. Click **OK** to display these settings in the **Build Filter** dialog box. Click the **Info** icon on the right of the **Destination IP Address** field for additional information about the correct input format to use.

g) Enter a valid destination port in the **Destination Port** field, or click the **Edit** icon on the right of the **Destination Port** field. Choose a condition on which to filter from the drop-down list, then click **Add**. Click **OK** to display these settings in the **Build Filter** dialog box. Click the **Info** icon on the right of the **Destination Port** field for additional information about the correct input format to use.

h) Enter filtering text for the **Description** field. The text may be any string of one or more characters, including a regular expression. However, semicolons are not valid characters, and this setting is case-sensitive. Multiple entries must be separated by commas.

i) Click **OK** to add the filter settings you have just specified to the **Filter By** drop-down list in the log viewers. The filter strings follow a specific format. The prefix FILTER: designates all custom filters that appear in the **Filter By** drop-down list. You may still type random text into this field.

The following table shows examples of the format used.

Build Filter Example	Filter String Format
Source IP = 192.168.1.1 or 0.0.0.0 Source Port = 67	FILTER: srcIP=192.168.1.1,0.0.0.0;srcPort=67;
Severity = Informational Destination IP = 1.1.1.1 through 1.1.1.10	FILTER: sev=6;dstIP=1.1.1.1-1.1.1.10;
Syslog ID not in the range 725001 through 725003	FILTER: sysID=!725001-725003;
Source IP = 1.1.1.1 Description = Built outbound	FILTER: srcIP=1.1.1.1;descr=Built outbound

Step 4 Choose one of the settings in the **Filter By** drop-down list to filter syslog messages, then click **Filter** on the toolbar. This setting also applies to all future syslog messages. Click **Show All** on the toolbar to clear all filters.

Note

You cannot save filters that you have specified with the **Build Filter** dialog box. These filters are valid only for the ASDM session during which they were created.

Edit Filtering Settings

To edit filtering settings that you created using the **Build Filter** dialog box, perform the following steps:

Procedure

Choose one of the following options:

- Revise a filter directly by entering the changes in the **Filter By** drop-down list.
- Choose a filter in the **Filter By** drop-down list, then click **Build Filter** to display the **Build Filter** dialog box. Click **Clear Filter** to remove the current filter settings and enter new ones. Otherwise, change the settings that appear, and click **OK**.

  Note
  These filter settings apply only to those defined in the **Build Filter** dialog box.

- Click **Show All** on the toolbar to stop filtering and show all syslog messages.

Issue Certain Commands Using the Log Viewers

You can issue the following commands using either of the log viewers: **ping**, **traceroute**, **whois**, and **dns lookup**.

To run any of these commands, perform the following steps:

Procedure

Step 1 Choose one of the following options:

- **Monitoring** > **Logging** > **Real-Time Log Viewer** > **View**
- **Monitoring** > **Logging** > **Log Buffer** > **View**

Step 2 Click **Tools** from the **Real-Time Log Viewer** or **Log Buffer** pane, then choose the command that you want to execute. Alternatively, you can right-click a specific syslog message that is listed to display a context menu, then choose the command that you want to execute.

The **Entering command** dialog box appears, with the command that you selected automatically showing in the drop-down list.

Step 3 Enter either the source or destination IP address of the selected syslog message in the **Address** field, then click **Go**.

The command output appears in the area provided.

Step 4 Click **Clear** to remove the output, and choose another command to execute from the drop-down list. Repeat Step 3, if necessary. Click **Close** when you are done.

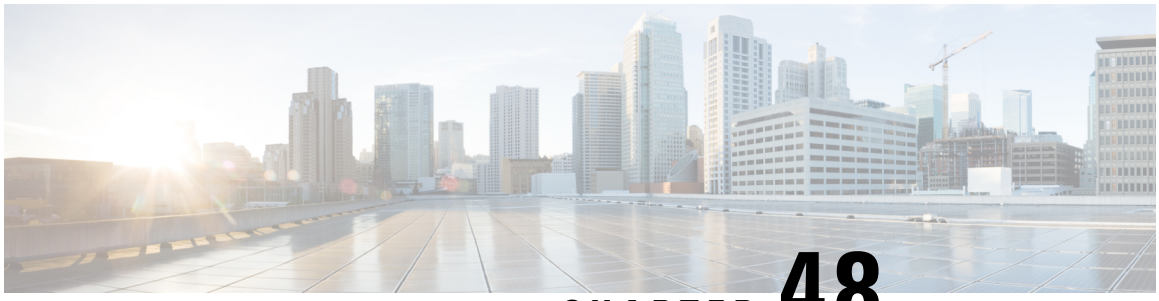
History for Logging

Table 60: History for Logging

Feature Name	Platform Releases	Description
Logging	7.0(1)	Provides ASA network logging information through various output destinations, and includes the option to view and save log files. We introduced the following screen: Configuration > Device Management > Logging > Logging Setup.
Rate limit	7.0(4)	Limits the rate at which syslog messages are generated. We modified the following screen: Configuration > Device Management > Logging > Rate Limit.
Logging list	7.2(1)	Creates a logging list to use in other commands to specify messages by various criteria (logging level, event class, and message IDs). We modified the following screen: Configuration > Device Management > Logging > Event Lists.
Secure logging	8.0(2)	Specifies that the connection to the remote logging host should use SSL/TLS. This option is valid only if the protocol selected is TCP. We modified the following screen: Configuration > Device Management > Logging > Syslog Server.
Logging class	8.0(4), 8.1(1)	Added support for the ipaa event class of logging messages. We modified the following screen: Configuration > Device Management > Logging > Logging Filters.
Logging class and saved logging buffers	8.2(1)	Added support for the dap event class of logging messages. Added support to clear the saved logging buffers (ASDM, internal, FTP, and flash). We modified the following screen: Configuration > Device Management > Logging > Logging Setup.
Password encryption	8.3(1)	Added support for password encryption.
Log viewers	8.3(1)	The source and destination IP addresses were added to the log viewers.

Feature Name	Platform Releases	Description
Enhanced logging and connection blocking	8.3(2)	When you configure a syslog server to use TCP, and the syslog server is unavailable, the ASA blocks new connections that generate syslog messages until the server becomes available again (for example, VPN, firewall, and cut-through-proxy connections). This feature has been enhanced to also block new connections when the logging queue on the ASA is full; connections resume when the logging queue is cleared. This feature was added for compliance with Common Criteria EAL4+. Unless required, we recommended allowing connections when syslog messages cannot be sent or received. To allow connections, continue to check the **Allow user traffic to pass when TCP syslog server is down** check box on the Configuration > Device Management > Logging > Syslog Servers pane. We introduced the following syslog messages: 414005, 414006, 414007, and 414008. We did not modify any ASDM screens.
Syslog message filtering and sorting	8.4(1)	Support has been added for the following: • Syslog message filtering based on multiple text strings that correspond to various columns • Creation of custom filters • Column sorting of messages. For detailed information, see the ASDM configuration guide. This feature interoperates with all ASA versions. We modified the following screens: Monitoring > Logging > Real-Time Log Viewer > View. Monitoring > Logging > Log Buffer Viewer > View.
Clustering	9.0(1)	Added support for syslog message generation in a clustering environment on the ASA 5580 and 5585-X. We modified the following screen: Configuration > Logging > Syslog Setup > Advanced > Advanced Syslog Configuration.
Blocking syslogs on a standby unit	9.4(1)	We added support for blocking the generation of specific syslog messages on the standby unit in a failover configuration. We modified the following screen: Configuration > Device Management > Logging > Syslog Setup.
Reference Identities for Secure Syslog Server connections	9.6(2)	TLS client processing now supports rules for verification of a server identity defined in RFC 6125, Section 6. Identity verification will be done during PKI validation for TLS connections to the Syslog Server. If the presented identity cannot be matched against the configured reference identity, the connection is not established. We modifed the following pages: **ASDM Configuration > Remote Access VPN > Advanced**, and **Configuration > Device Management > Logging > Syslog Servers -> Add or Edit.**

Feature Name	Platform Releases	Description
IPv6 address support for syslog servers	9.7(1)	You can now configure syslog servers with IPv6 addresses to record, send, and receive syslogs over TCP and UDP. We modified the following screen: **Configuration** > **Device Management** > **Logging** > **Syslog Servers** > **Add Syslog Server**
Logging class	9.12(1)	Added support for the BFD, BGP, interface, IPv6, Muticast, Object-Group-Search, PBR, routing, SLA class of logging messages. We modified the following screen: **Configuration** > **Device Management** > **Logging** > **Logging Filters**.
Loopback interface support for syslog	9.18(2)	You can now add a loopback interface and use it for syslog. New/Modified commands: **interface loopback**, **logging host** New/Modified screens: **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces** > **Add Loopback Interface** ASDM support was added in 7.19.
Rate limiting for SNMP syslogs	9.20(1)	If you do not set system-wide rate limiting, you can now configure rate limiting separately for syslogs sent to an SNMP server.



CHAPTER 48

SNMP

This chapter describes how to configure Simple Network Management Protocol (SNMP) to monitor ASA.

- About SNMP, on page 1171
- Guidelines for SNMP, on page 1174
- Configure SNMP, on page 1176
- Monitoring SNMP, on page 1182
- History for SNMP, on page 1183

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. The ASA 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 ASA interface lets you monitor the network devices through network management systems (NMSes), such as HP OpenView. The ASA support 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 ASA 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 ASA 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.



Note With intense workloads, deploying more than 10 NMS can impact the device's performance. To ensure device's stability and responsiveness, we recommend that you cautiously utilize NMS in conducting SNMP walk polling and in managing the trap traffic.

The ASA have an SNMP agent that notifies 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 ASA 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 61: SNMP Terminology

Term	Description
Agent	The SNMP server running on the ASA. 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, such as the ASA.
Object identifier (OID)	The system that identifies a device to its NMS and indicates to users the source of information monitored and displayed.
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.

SNMP Version 3 Overview

SNMP Version 3 provides security enhancements that are not available in SNMP Version 1 or Version 2c. SNMP Versions 1 and 2c transmit data between the SNMP server and SNMP agent in clear text. SNMP Version 3 adds authentication and privacy options to secure protocol operations. In addition, this version controls access to the SNMP agent and MIB objects through the User-based Security Model (USM) and View-based Access Control Model (VACM). The ASA also supports the creation of SNMP groups and users, as well as hosts, which is required to enable transport authentication and encryption for secure SNMP communications.

Security Models

For configuration purposes, the authentication and privacy options are grouped together into security models. Security models apply to users and groups, which are divided into the following three types:

- NoAuthPriv—No Authentication and No Privacy, which means that no security is applied to messages.
- AuthNoPriv—Authentication but No Privacy, which means that messages are authenticated.

- AuthPriv—Authentication and Privacy, which means that messages are authenticated and encrypted.

SNMP Groups

An SNMP group is an access control policy to which users can be added. Each SNMP group is configured with a security model, and is associated with an SNMP view. A user within an SNMP group must match the security model of the SNMP group. These parameters specify what type of authentication and privacy a user within an SNMP group uses. Each SNMP group name and security model pair must be unique.

SNMP Users

SNMP users have a specified username, a group to which the user belongs, authentication password, encryption password, and authentication and encryption algorithms to use. The authentication algorithm options are SHA-1, SHA-224, SHA-256 HMAC, and SHA-384. The encryption algorithm options are 3DES and AES (which is available in 128, 192, and 256 versions). When you create a user, you must associate it with an SNMP group. The user then inherits the security model of the group.



Note When configuring an SNMP v3 user account, ensure that the length of authentication algorithm is equal to or greater than the length of encryption algorithm.

SNMP Hosts

An SNMP host is an IP address to which SNMP notifications and traps are sent. To configure SNMP Version 3 hosts, along with the target IP address, you must configure a username, because traps are only sent to a configured user. SNMP target IP addresses and target parameter names must be unique on the ASA. Each SNMP host can have only one username associated with it. To receive SNMP traps, configure the SNMP NMS, and make sure that you configure the user credentials on the NMS to match the credentials for the ASA.



Note You can add up to 8192 hosts. However, only 128 of this number can be for traps.

Implementation Differences Between the ASA and Cisco IOS Software

The SNMP Version 3 implementation in the ASA differs from the SNMP Version 3 implementation in the Cisco IOS software in the following ways:

- The local-engine and remote-engine IDs are not configurable. The local engine ID is generated when the ASA starts or when a context is created.
- No support exists for view-based access control, which results in unrestricted MIB browsing.
- Support is restricted to the following MIBs: USM, VACM, FRAMEWORK, and TARGET.
- You must create users and groups with the correct security model.
- You must remove users, groups, and hosts in the correct sequence.
- Use of the snmp-server host command creates an ASA rule to allow incoming SNMP traffic.

SNMP Syslog Messaging

SNMP generates detailed syslog messages that are numbered 212*nnn*. Syslog messages indicate the status of SNMP requests, SNMP traps, SNMP channels, and SNMP responses from the ASA or ASASM to a specified host on a specified interface.

For detailed information about syslog messages, see the syslog messages guide.

Note SNMP polling fails if SNMP syslog messages exceed a high rate (approximately 4000 per second).

Application Services and Third-Party Tools

For information about SNMP support, see the following URL:

http://www.cisco.com/en/US/tech/tk648/tk362/tk605/tsd_technology_support_sub-protocol_home.html

For information about using third-party tools to walk SNMP Version 3 MIBs, see the following URL:

http://www.cisco.com/en/US/docs/security/asa/asa83/snmp/snmpv3_tools.html

Guidelines for SNMP

This section includes the guidelines and limitations that you should review before configuring SNMP.

Failover and Clustering Guidelines

- When using SNMPv3 with clustering or failover, if you add a new cluster unit after the initial cluster formation or you replace a failover unit, then SNMPv3 users are not replicated to the new unit. You must re-add the SNMPv3 users to the control/active unit to force the users to replicate to the new unit; or you can add the users directly on the new unit (SNMPv3 users and groups are an exception to the rule that you cannot enter configuration commands on a cluster data unit). Reconfigure each user by entering the **snmp-server user** *username group-name* **v3** command on the control/active unit or directly to the data/standby unit with the *priv-password* option and *auth-password* option in their unencrypted forms.

IPv6 Guidelines (All ASA Models)

SNMP can be configured over IPv6 transport so that an IPv6 host can perform SNMP queries and receive SNMP notifications from a device running IPv6 software. The SNMP agent and related MIBs have been enhanced to support IPv6 addressing.

Additional Guidelines

- Power supply traps are not issued for systems operating in Appliance mode.
- For the Firepower 2100 in Appliance mode, you cannot poll the hardware model and serial, and in ASA, traps are not generated for these details. Hence, configure SNMP on FXOS or Chassis Manager to poll Chassis Management IP instead of the interfaces on ASA instance.
- You must have Cisco Works for Windows or another SNMP MIB-II compliant browser to receive SNMP traps or browse a MIB.

- Management-access over a VPN tunnel is not supported with SNMP (the **management-access** command). For SNMP over VPN, we recommend enabling SNMP on a loopback interface. You don't need the management-access feature enabled to use SNMP on the loopback interface. Loopback also works for SSH.
- Does not support view-based access control, but the VACM MIB is available for browsing to determine default view settings.
- The ENTITY-MIB is not available in the non-admin context. Use the IF-MIB instead to perform queries in the non-admin context.
- The ENTITY-MIB is not available for the Firepower 9300. Instead, use CISCO-FIREPOWER-EQUIPMENT-MIB and CISCO-FIREPOWER-SM-MIB.
- On some devices, the order of interfaces (ifDescr) in the output of **snmpwalk** has been observed to change after a reboot. The ASA uses an algorithm to determine the ifIndex table that SNMP queries. When the ASA is booted up, the interfaces are added to the ifIndex table in the order loaded as the ASA reads the configuration. New interfaces added to the ASA are appended to the list of interfaces in the ifIndex table. As interfaces are added, removed, or renamed, it can affect the order of interfaces on reboot.
- When you provide an OID in the **snmpwalk** command, the snmpwalk tool queries all variables in the subtree that is below the specified OID and displays their values. Thus, to view a comprehensive output of the objects on the device, ensure to provide the OID in the **snmpwalk** command.
- Does not support SNMP Version 3 for the AIP SSM or AIP SSC.
- Does not support SNMP debugging.
- Does not support retrieval of ARP information.
- Does not support SNMP SET commands.
- When using NET-SNMP Version 5.4.2.1, only supports the encryption algorithm version of AES128. Does not support the encryption algorithm versions of AES256 or AES192.
- Changes to the existing configuration are rejected if the result places the SNMP feature in an inconsistent state.
- For SNMP Version 3, configuration must occur in the following order: group, user, host.
- For Secure Firewall models, the **snmpwalk** command polls FXOS mibs only from admin context.
- Before a group is deleted, you must ensure that all users associated with that group are deleted.
- Before a user is deleted, you must ensure that no hosts are configured that are associated with that username.
- If users have been configured to belong to a particular group with a certain security model, and if the security level of that group is changed, you must do the following in this sequence:
    - Remove the users from that group.
    - Change the group security level.
    - Add users that belong to the new group.
- The creation of custom views to restrict user access to a subset of MIB objects is not supported.
- All requests and traps are available in the default Read/Notify View only.

- The connection-limit-reached trap is generated in the admin context. To generate this trap. you must have at least one SNMP server host configured in the user context in which the connection limit has been reached.
- If the NMS cannot successfully request objects or is not correctly handling incoming traps from the ASA, performing a packet capture is the most useful method for determining the problem. Choose **Wizards > Packet Capture Wizard**, and follow the on-screen instructions.
- You can add up to 4000 hosts. However, only 128 of this number can be for traps.
- The total number of supported active polling destinations is 128.
- You can specify a network object to indicate the individual hosts that you want to add as a host group.
- You can associate more than one user with one host.
- You can specify overlapping network objects in different **host-group** commands. The values that you specify for the last host group take effect for the common set of hosts in the different network objects.
- If you delete a host group or hosts that overlap with other host groups, the hosts are set up again using the values that have been specified in the configured host groups.
- The values that the hosts acquire depend on the specified sequence that you use to run the commands.
- The limit on the message size that SNMP sends is 1472 bytes.
- The ASA supports an unlimited number of SNMP server trap hosts per context. The **show snmp-server host** command output displays only the active hosts that are polling the ASA, as well as the statically configured hosts.

Configure SNMP

This section describes how to configure SNMP.

Procedure

Step 1 Configure an SNMP management station to receive requests from the ASA.
Step 2 Configure SNMP traps.
Step 3 Configure SNMP Version 1 and 2c parameters or SNMP Version 3 parameters.

Configure an SNMP Management Station

To configure an SNMP management station, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Management Access** > **SNMP**. By default, the SNMP server is enabled.

Step 2 Click **Add** in the **SNMP Management Stations** pane.

The **Add SNMP Host Access Entry** dialog box appears.

Step 3 Choose the interface on which the SNMP host resides.

Step 4 Enter the SNMP host IP address.

Step 5 Enter the SNMP host UDP port, or keep the default, port 162.

Step 6 Add the SNMP host community string. If no community string is specified for a management station, the value set in the **Community String** (default) field on the **SNMP Management Stations** pane is used.

Step 7 Choose the SNMP version used by the SNMP host.

Step 8 If you have selected SNMP Version 3 in the previous step, choose the name of a configured user.

Step 9 To specify the method for communicating with this NMS, check either the **Poll** or **Trap** check box.

Step 10 Click **OK**.

The **Add SNMP Host Access Entry** dialog box closes.

Step 11 Click **Apply**.

The NMS is configured and changes are saved to the running configuration. For more information about SNMP Version 3 NMS tools, see the following URL:

http://www.cisco.com/en/US/docs/security/asa/asa82/snmp/snmpv3_tools.html

Configure SNMP Traps

To designate which traps that the SNMP agent generates and how they are collected and sent to NMSs, perform the following steps:



Note 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. For example, you can skip *Informational* syslog trap severity level.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Management Access** > **SNMP**.

Step 2 Click **Configure Traps**.

The **SNMP Trap Configuration** dialog box appears.

Step 3 Check the **SNMP Server traps configuration** check box.

The default configuration has all SNMP standard traps enabled. If you do not specify a trap type, the default is the **syslog** trap. The default SNMP traps continue to be enabled with the **syslog** trap. All other traps are disabled by default. To disable a trap, uncheck the applicable check box.

The traps are divided into the following categories:

a) **Standard SNMP Traps**, check all that apply.

Choose from **Critical CPU temperature**, **Chassis temperature**, and **Chassis Fan Failure**.

Note
The default configuration has all SNMP standard traps enabled.

b) **Environment Traps**, check all that apply.

Choose from **Authentication**, **Link up**, **Link down**, **Cold start**, and **Warm start**.

c) **Ikev2 Traps** check all that apply.

Choose from **Start** and **Stop**.

d) **Entity MIB Notifications**.

Check this item to receive notifications about field-replaceable units.

e) **IPsec Traps**, check all that apply.

Choose from **Start** and **Stop**.

f) **Remote Access Traps**.

Check this item to receive notifications when the number of sessions established exceeds the set threshold.

g) **Resource Traps**, check all that apply.

Choose from **Connection limit reached**, **Memory threshold reached**, and **Interface threshold reached**.

h) **NAT Traps**.

Check this item to receive notifications when IP packets are discarded by NAT because mapping space is not available.

i) **Syslog**.

Check **Enable syslog traps** to receive notifications when the number of sessions established exceeds the set threshold.

To configure the **syslog** trap severity level, choose **Configuration** > **Device Management** > **Logging** > **Logging Filters**

j) **CPU Utilization Traps**.

Check **CPU rising threshold reached** to receive notifications when the CPU usage is greater than the configured **CPU Utilization threshold** value for the configured **Monitoring interval**.

k) **SNMP interface threshold**.

Check **Configure threshold and interval** to receive notifications when the interface bandwidth utilization is greater than the configured **SNMP interface threshold** value.

Valid threshold values range from 30 to 99 percent. The default value is 70 percent.

l) **SNMP Memory threshold**.

Check **Configure memory threshold** to receive notifications when the CPU usage is greater than the configured threshold value for the **SNMP memory threshold** value.

When the used system context memory reaches 80 percent of the total system memory, the memory threshold trap is generated from the admin context. For all other user contexts, this trap is generated when the used memory reaches 80 percent of the total system memory in that particular context.

m) **Failover Traps**.

Check **Enable Failover related traps** to receive SNMP syslog traps for failover.

n) **Cluster Traps**.

Check **Enable cluster related traps** to receive SNMP syslog traps for cluster members.

o) **Peer-Flap Traps**.

Check **Enable bgp/ospf peer-flap related traps** to receive SNMP syslog traps for cluster peer MAC address flapping.

Step 4 Click **OK** to close the **SNMP Trap Configuration** dialog box.

Step 5 Click **Apply**.

The SNMP traps are configured and the changes are saved to the running configuration.

Configure Parameters for SNMP Version 1 or 2c

To configure parameters for SNMP Version 1 or 2c, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Management Access** > **SNMP.**

Step 2 Enter a default community string in the **Community String** (default) field if you are using SNMP Version 1 or 2c. Enter the password used by the SNMP NMSs when they send requests to the ASA. The SNMP community string is a shared secret among the SNMP NMSs and the network nodes being managed. The ASA uses the password to determine if the incoming SNMP request is valid. However, if SNMP monitoring is through the management interface instead of the diagnostic interface, polling takes place without ASA validating the community string. The password is a case-sensitive value up to 32 alphanumeric characters long. Spaces are not permitted. The default is public. SNMP Version 2c allows separate community strings to be set for each NMS. If no community string is configured for any NMS, the value set here is used by default.

Note
You should avoid the use of special characters (!, @, #, $, %, ^, &, *, \) in community strings. In general, using any special characters reserved for functions used by the operating system can cause unexpected results. For example, the backslash (\) is interpreted as an escape character and should not be used in the community string.

Step 3 Enter the name of the ASA system administrator. The text is case-sensitive and can be up to 127 alphabetic characters. Spaces are accepted, but multiple spaces are shortened to a single space.

Step 4 Enter the location of the ASA being managed by SNMP. The text is case-sensitive and can be up to 127 characters. Spaces are accepted, but multiple spaces are shortened to a single space.

Step 5 Enter the number of the ASA port that listens for SNMP requests from NMSes; or keep the default, port number 161.

Step 6 (Optional) Check the **Enable Global-Shared pool in the walk** checkbox to query free memory and used memory statistics through SNMP walk operations.

Important
When the ASA queries memory information, the CPU may be held by the SNMP process for too long before releasing the CPU to other processes. This can result in SNMP-related CPU hogs causing packet drops.

Step 7 Click **Add** in the **SNMP Host Access List** pane.

The **Add SNMP Host Access Entry** dialog box appears.

Step 8 Choose the interface name from which traps are sent from the drop-down list.

Step 9 Enter the IP address of the NMS or SNMP manager that can connect to the ASA.

Step 10 Enter the UDP port number. The default is 162.

Step 11 Choose the SNMP version that you are using from the drop-down list. If you choose Version 1 or Version 2c, you must enter the community string. If you choose Version 3, you must choose the username from the drop-down list.

The version specifies the SNMP version to use for traps and requests (polling). Communication with the server is allowed using the selected version only.

Step 12 Check the **Poll** check box in the **Server Poll/Trap Specification** area to limit the NMS to sending requests (polling) only. Check the **Trap** check box to limit the NMS to receiving traps only. You may check both check boxes to perform both functions of the SNMP host.

Step 13 Click **OK** to close the **Add SNMP Host Access Entry** dialog box.

The new host appears in the **SNMP Host Access List** pane.

Step 14 Click **Apply**.

SNMP parameters for Versions 1, 2c, or 3 are configured and the changes are saved to the running configuration.

Configure Parameters for SNMP Version 3

To configure parameters for SNMP Version 3, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Management Access** > **SNMP**.

Step 2 Click **Add** > **SNMP User**on the **SNMPv3 User/Group** tab in the **SNMPv3 Users pane** to add a configured user or a new user to a group. When you remove the last user in a group, ASDM deletes the group.

Note
After a user has been created, you cannot change the group to which the user belongs.

The **Add SNMP User Entry** dialog box appears.

Step 3 Choose the group to which the SNMP user belongs. The available groups are as follows:

- **Auth&Encryption**, in which users have authentication and encryption configured
- **Authentication_Only**, in which users have only authentication configured
- **No_Authentication**, in which users have neither authentication nor encryption configured

Note
You cannot change the group names.

Step 4 Click the **USM Model** tab to use the user security model (USM) groups.

Step 5 Click **Add**.

The **Add SNMP USM Entry** dialog box appears.

Step 6 Enter the group name.

Step 7 Choose the security level from the drop-down list. This setting allows you to assign a configured USM group as a security level to SNMPv3 users.

Step 8 Enter the name of a configured user or a new user. The username must be unique for the SNMP server group selected.

Step 9 Indicate the type of password you want to use by clicking one of the two radio buttons: **Encrypted** or **Clear Text**.

Step 10 Indicate the type of authentication you want to use by clicking one of the four radio buttons: **SHA**, **SHA224**, **SHA256**, or **SHA384**.

Step 11 Enter the password to use for authentication.

Step 12 Indicate the type of encryption you want to use by clicking one of these two radio buttons: **3DES** or **AES**.

Step 13 If you chose AES encryption, then choose the level of AES encryption to use: **128**, **192**, or **256**.

Step 14 Enter the password to use for encryption. The maximum number of alphanumeric characters allowed for this password is 64.

Step 15 Click **OK** to create a group (if this is the first user in that group), display this group in the **Group Name** drop-down list, and create a user for that group.

The **Add SNMP User Entry** dialog box closes.

Step 16 Click **Apply**.

SNMP parameters for Version 3 are configured, and the changes are saved to the running configuration.

Configure a Group of Users

To configure an SNMP user list with a group of specified users in it, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Management Access** > **SNMP.**

Step 2 Click **Add** > **SNMP User Group** on the **SNMPv3 User/Group** tab in the **SNMPv3 Users** pane to add a configured user group or a new user group. When you remove the last user in a group, ASDM deletes the group.

The **Add SNMP User Group** dialog box appears.

Step 3 Enter the user group name.

Step 4 Click the **Existing User/User Group** radio button to select an existing user or user group.

Step 5 Click the **Create new user** radio button to create a new user.

Step 6 Choose the group to which the SNMP user belongs. The available groups are as follows:

- **Auth&Encryption**, in which users have authentication and encryption configured
- **Authentication_Only**, in which users have only authentication configured
- **No_Authentication**, in which users have neither authentication nor encryption configured

Step 7 Enter the name of a configured user or a new user. The username must be unique for the SNMP server group selected.

Step 8 Indicate the type of password you want to use by clicking one of the two radio buttons: **Encrypted** or **Clear Text**.

Step 9 Indicate the type of authentication you want to use by clicking one of the four radio buttons: **SHA**, **SHA224**, **SHA256**, or **SHA384**.

Step 10 Enter the password to use for authentication.

Step 11 Confirm the password to use for authentication.

Step 12 Indicate the type of encryption you want to use by clicking one of these two radio buttons: **3DES** or **AES**.

Step 13 Enter the password to use for encryption. The maximum number of alphanumeric characters allowed for this password is 64.

Step 14 Confirm the password to use for encryption.

Step 15 Click **Add** to add the new user to the specified user group in the **Members in Group** pane. Click **Remove** to delete an existing user from the **Members in Group** pane.

Step 16 Click **OK** to create a new user for the specified user group.

The **Add SNMP User Group** dialog box closes.

Step 17 Click **Apply**.

SNMP parameters for Version 3 are configured, and the changes are saved to the running configuration.

Monitoring SNMP

See the following commands for monitoring SNMP. You can enter these commands using **Tools > Command Line Interface**.

- **show running-config snmp-server** [**default**]

  This command shows all SNMP server configuration information.

- **show running-config snmp-server group**

This command shows SNMP group configuration settings.

- **show running-config snmp-server host**

  This command shows configuration settings used by SNMP to control messages and notifications sent to remote hosts.

- **show running-config snmp-server host-group**

  This command shows SNMP host group configurations.

- **show running-config snmp-server user**

  This command shows SNMP user-based configuration settings.

- **show running-config snmp-server user-list**

  This command shows SNMP user list configurations.

- **show snmp-server engineid**

  This command shows the ID of the SNMP engine configured.

- **show snmp-server group**

  This command shows the names of configured SNMP groups. If the community string has already been configured, two extra groups appear by default in the output. This behavior is normal.

- **show snmp-server statistics**

  This command shows the configured characteristics of the SNMP server. To reset all SNMP counters to zero, use the **clear snmp-server statistics** command.

- **show snmp-server user**

  This command shows the configured characteristics of users.

History for SNMP

Table 62: History for SNMP

Feature Name	Version	Description
SNMP Versions 1 and 2c	7.0(1)	Provides ASA network monitoring and event information by transmitting data between the SNMP server and SNMP agent through the clear text community string. We modified the following screen: Configuration > Device Management > Management Access > SNMP.

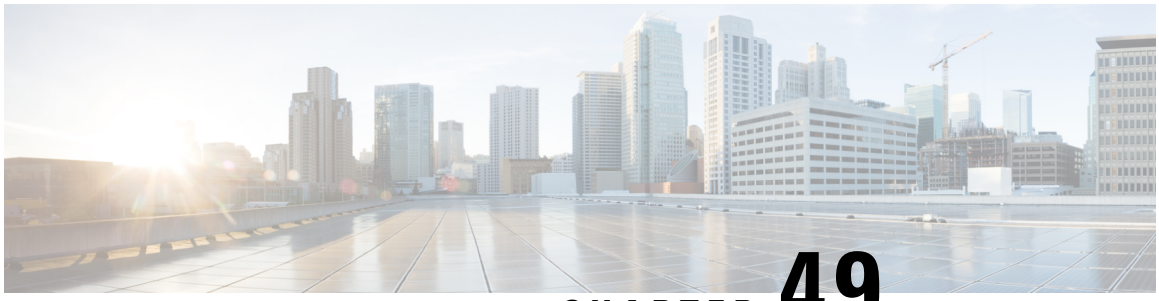
Feature Name	Version	Description
SNMP Version 3	8.2(1)	Provides 3DES or AES encryption and support for SNMP Version 3, the most secure form of the supported security models. This version allows you to configure users, groups, and hosts, as well as authentication characteristics by using the USM. In addition, this version allows access control to the agent and MIB objects and includes additional MIB support. We modified the following screen: Configuration > Device Management > Management Access > SNMP.
Password encryption	8.3(1)	Supports password encryption.
SNMP traps and MIBs	8.4(1)	Supports the following additional keywords: **connection-limit-reached**, **cpu threshold rising**, **entity cpu-temperature**, **entity fan-failure**, **entity power-supply**, **ikev2 stop | start**, **interface-threshold**, **memory-threshold**, **nat packet-discard**, **warmstart**. The entPhysicalTable reports entries for sensors, fans, power supplies, and related components. Supports the following additional MIBs: CISCO-ENTITY-SENSOR-EXT-MIB, CISCO-ENTITY-FRU-CONTROL-MIB, CISCO-PROCESS-MIB, CISCO-ENHANCED-MEMPOOL-MIB, CISCO-L4L7MODULE-RESOURCE-LIMIT-MIB, DISMAN-EVENT-MIB, DISMAN-EXPRESSION-MIB, ENTITY-SENSOR-MIB, NAT-MIB. Supports the following additional traps: ceSensorExtThresholdNotification, clrResourceLimitReached, cpmCPURisingThreshold, mteTriggerFired, natPacketDiscard, warmStart. We modified the following screen: Configuration > Device Management > Management Access > SNMP.
IF-MIB ifAlias OID support	8.2(5) / 8.4(2)	The ASA now supports the ifAlias OID. When you browse the IF-MIB, the ifAlias OID will be set to the value that has been set for the interface description.

Feature Name	Version	Description
ASA Services Module (ASASM)	8.5(1)	The ASASM supports all MIBs and traps that are present in 8.4(1), except for the following: Unsupported MIBs in 8.5(1): • CISCO-ENTITY-SENSOR-EXT-MIB (Only objects under the entPhySensorTable group are supported). • ENTITY-SENSOR-MIB (Only objects in the entPhySensorTable group are supported). • DISMAN-EXPRESSION-MIB (Only objects in the expExpressionTable, expObjectTable, and expValueTable groups are supported). Unsupported traps in 8.5(1): • ceSensorExtThresholdNotification (CISCO-ENTITY-SENSOR-EXT-MIB). This trap is only used for power supply failure, fan failure, and high CPU temperature events. • InterfacesBandwidthUtilization.
SNMP traps	8.6(1)	Supports the following additional keywords for the ASA 5512-X, 5515-X, 5525-X, 5545-X, and 5555-X: **entity power-supply-presence**, **entity power-supply-failure**, **entity chassis-temperature**, **entity chassis-fan-failure, entity power-supply-temperature**. We modified the following command: **snmp-server enable traps**.
VPN-related MIBs	9.0(1)	An updated version of the CISCO-IPSEC-FLOW-MONITOR-MIB.my MIB has been implemented to support the next generation encryption feature. The following MIBs have been enabled for the ASASM: • ALTIGA-GLOBAL-REG.my • ALTIGA-LBSSF-STATS-MIB.my • ALTIGA-MIB.my • ALTIGA-SSL-STATS-MIB.my • CISCO-IPSEC-FLOW-MONITOR-MIB.my • CISCO-REMOTE-ACCESS-MONITOR-MIB.my
Cisco TrustSec MIB	9.0(1)	Support for the following MIB was added: CISCO-TRUSTSEC-SXP-MIB.
SNMP OIDs	9.1(1)	Five new SNMP Physical Vendor Type OIDs have been added to support the ASA 5512-X, 5515-X, 5525-X, 5545-X, and 5555-X.
NAT MIB	9.1(2)	Added the cnatAddrBindNumberOfEntries and cnatAddrBindSessionCount OIDs to support the xlate_count and max_xlate_count entries, which are the equivalent to allowing polling using the **show xlate count** command.

Feature Name	Version	Description
SNMP hosts, host groups, and user lists	9.1(5)	You can now add up to 4000 hosts. The number of supported active polling destinations is 128. You can specify a network object to indicate the individual hosts that you want to add as a host group. You can associate more than one user with one host. We modified the following screen: Configuration > Device Management > Management Access > SNMP.
SNMP message size	9.2(1)	The limit on the message size that SNMP sends has been increased to 1472 bytes.
SNMP OIDs and MIBs	9.2(1)	The ASA now supports the cpmCPUTotal5minRev OID. The ASA virtual has been added as a new product to the SNMP sysObjectID OID and entPhysicalVendorType OID. The CISCO-PRODUCTS-MIB and CISCO-ENTITY-VENDORTYPE-OID-MIB have been updated to support the new ASA virtual platform. A new SNMP MIB for monitoring VPN shared license usage has been added.
SNMP OIDs and MIBs	9.3(1)	CISCO-REMOTE-ACCESS-MONITOR-MIB (OID 1.3.6.1.4.1.9.9.392) support has been added for the ASASM.
SNMP MIBs and traps	9.3(2)	The CISCO-PRODUCTS-MIB and CISCO-ENTITY-VENDORTYPE-OID-MIB have been updated to support the ASA 5506-X. The ASA 5506-X has been added as new products to the SNMP sysObjectID OID and entPhysicalVendorType OID tables. The ASA now supports the CISCO-CONFIG-MAN-MIB, which enables you to do the following: • Know which commands have been entered for a specific configuration. • Notify the NMS when a change has occurred in the running configuration. • Track the time stamps associated with the last time that the running configuration was changed or saved. • Track other changes to commands, such as terminal details and command sources. We modified the following screen: Configuration > Device Management > Management Access > SNMP > Configure Traps > SNMP Trap Configuration.
SNMP MIBs and traps	9.4(1)	The ASA 5506W-X, ASA 5506H-X, ASA 5508-X, and ASA 5516-X have been added as a new product to the SNMP sysObjectID OID and entPhysicalVendorType OID tables.

Feature Name	Version	Description
Unlimited SNMP server trap hosts per context	9.4(1)	The ASA supports unlimited SNMP server trap hosts per context. The **show snmp-server host** command output displays only the active hosts that are polling the ASA, as well as the statically configured hosts. We did not modify any ASDM screens.
Added support for ISA 3000	9.4(1.225)	The ISA 3000 family of products is now supported for SNMP. We added new OIDs for this platform. The **snmp-server enable traps entity** command has been modified to include a new variable *l1-bypass-status*. This enables hardware bypass status change. We did not modify any ASDM screens.
Support for the cempMemPoolTable in the CISCO-ENHANCED-MEMPOOL-MIB	9.6(1)	The cempMemPoolTable of the CISCO-ENHANCED-MEMPOOL-MIB is now supported. This is a table of memory pool monitoring entries for all physical entities on a managed system. **Note** The CISCO-ENHANCED-MEMPOOL-MIB uses 64-bit counters and supports reporting of memory on platforms with more than 4GB of RAM.
Support for E2E Transparent Clock Mode MIBs for the Precision Time Protocol (PTP)	9.7(1)	MIBs corresponding to E2E Transparent Clock mode are now supported. **Note** Only SNMP get, bulkget, getnext, and walk operations are supported.
SNMP over IPv6	9.9(2)	The ASA now supports SNMP over IPv6, including communicating with SNMP servers over IPv6, allowing the execution of queries and traps over IPv6, and supporting IPv6 addresses for existing MIBs. We added the following new SNMP IPv6 MIB objects as described in RFC 8096. • ipv6InterfaceTable (OID: 1.3.6.1.2.1.4.30)—Contains per-interface IPv6-specific information. • ipAddressPrefixTable (OID:1.3.6.1.2.1.4.32)—Includes all the prefixes learned by this entity. • ipAddressTable (OID: 1.3.6.1.2.1.4.34)—Contains addressing information relevant to the entity's interfaces. • ipNetToPhysicalTable (OID: 1.3.6.1.2.1.4.35)—Contains the mapping from IP addresses to physical addresses. New or modified screen: **Configuration** > **Device Management** > **Management Access** > **SNMP**
Support to enable and disable the results for free memory and used memory statistics during SNMP walk operations	9.10(1)	To avoid overutilization of CPU resources, you can enable and disable the query of free memory and used memory statistics collected through SNMP walk operations. We did not modify any ASDM screens.

Feature Name	Version	Description
Support to enable and disable the results for free memory and used memory statistics during SNMP walk operations	9.12(1)	To avoid overutilization of CPU resources, you can enable and disable the query of free memory and used memory statistics collected through SNMP walk operations. New or modified screen: **Configuration** > **Device Management** > **Management Access** > **SNMP**
SNMPv3 Authentication	9.14(1)	You can now use SHA-256 HMAC for user authentication. New/Modified screens: **Configuration** > **Device Management** > **Management Access** > **SNMP**
For Failover pairs in 9.14(1)+, the ASA no longer shares SNMP client engine data with its peer.	9.14(1)	The ASA no longer shares SNMP client engine data with its peer.
SNMP polling over site-to-site VPN	9.14(2)	For secure SNMP polling over a site-to-site VPN, include the IP address of the outside interface in the crypto map access-list as part of the VPN configuration.
Support for the CISCO-MEMORY-POOL-MIB OIDs is deprecated	9.15(1)	The CISCO-MEMORY-POOL-MIB OIDs (ciscoMemoryPoolUsed, ciscoMemoryPoolFree) are deprecated for systems that use 64-bit counters. The cempMemPoolTable of the CISCO-ENHANCED-MEMPOOL-MIB provides memory pool monitoring entries for systems that use 64-bit counters.
SNMPv3 Authentication	9.16(1)	You can now use SHA-224 and SHA-384 for user authentication. You can no longer use MD5 for user authentication. You can no longer use DES for encryption. New/Modified screens: **Configuration** > **Device Management** > **Management Access** > **SNMP**
Loopback interface support for SNMP	9.18(2)	You can now add a loopback interface and use it for SNMP. New/Modified commands: **interface loopback**, **snmp-server host** New/Modified screens: **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces** > **Add Loopback Interface** ASDM support was added in 7.19.
SNMP MIBs and traps	9.20(1)	The Secure Firewall 4200 model devices (FPR4215, FPR4225, FPR4245) have been added as a new product to the SNMP sysObjectID OID and entPhysicalVendorType OID tables. SNMP support for the two EPM cards(4X200G and 2X100G) of these Secure Firewall 4200 Series devices was added.



CHAPTER 49

Cisco Success Network and Telemetry Data

This chapter describes about Cisco Success Network and how to enable it on ASA. It also lists the telemetry data points that are sent to the Security Service Engine(SSE) cloud.

- About Cisco Success Network , on page 1189
- Enable or Disable Cisco Success Network , on page 1190
- View ASA Telemetry Data , on page 1191
- Cisco Success Network - Telemetry Data, on page 1191

About Cisco Success Network

Cisco Success Network is user-enabled cloud service that establishes a secured connection with the Security Service Exchange (SSE) cloud to stream ASA usage information and statistics. Streaming telemetry provides a mechanism to transmit ASA usage and other details in structured format (JSON) to remote management stations for the following benefits:

- To inform you of extra technical support services and monitoring that are available for your product.
- To help Cisco improve its products.

By default, the Cisco Success Network is enabled on the Firepower 4100/9300 platforms that hosts ASA devices (at the blade level). However, for the telemetry data to be transmitted, you must enable the configuration on FXOS at chassis level (see Cisco Firepower 4100/9300 FXOS CLI Configuration Guide) or enable the Cisco Success Network on the chassis manager (see Cisco Firepower 4100/9300 FXOS Firepower Chassis Manager Configuration Guide) ASA allows you to disable the telemetry service at any point in time.

The telemetry data that is collected on your ASA devices includes CPU, memory, disk, bandwidth, and license usage, configured feature list, cluster/failover information, and the alike. Refer Cisco Success Network - Telemetry Data, on page 1191.

Supported Platforms and Required Configurations

- Supported on FP9300/4100 platforms with ASA version 9.13.1 or above running on it.
- Requires FXOS version 2.7.1 or above to connect with the cloud.
- The SSE connector on FXOS must be connected to the SSE cloud. This connection is established by enabling and registering the smart license with smart licensing backend. The SSE connector on FXOS is automatically registered to the SSE cloud by registering smart license.

- The Cisco Success Network configuration must be enabled on chassis manager.
- The telemetry configuration must be enabled on ASA.

How Does ASA Telemetry Data Reach the SSE Cloud

Cisco Success Network is supported on Firepower 4100/9300 platforms in ASA 9.13(1) by default. The FXOS service manager sends telemetry request daily to the ASA application running on the platform. The ASA engine, based on the configuration and connectivity status, sends the telemetry data either in standalone mode or cluster mode to FXOS. That is, if the telemetry support is enabled in ASA and SSE connector status is connected, the telemetry thread pulls the needed information from various sources such as system or platform or device APIs, license APIs, CPU APIs, memory APIs, disk APIs, smart call home feature APIs, and so on. However, if the telemetry support is disabled in ASA or the SSE connector status is disconnected, ASA sends a response to FXOS (appAgent) indicating the telemetry configuration status and does not send any telemetry data.

FXOS has only one SSE connector instance running on it. When it gets registered with the SSE cloud, it is considered as one device and SSE infra assigns FXOS with one device ID. Any telemetry report that is sent through the SSE connector is categorized under the same device ID. Therefore, FXOS aggregates the telemetry report from each ASA into a single report. Other contents such as smart license account information are added to the report. FXOS then sends the final report to the SSE cloud. The telemetry data is saved in the SSE data exchange (DEX) and available for the Cisco IT team to use.

Enable or Disable Cisco Success Network

Before you begin

- Enable and register smart license on FXOS.
- Enable telemetry support on FXOS at the chassis level (see Cisco Firepower 4100/9300 FXOS CLI Configuration Guide) or enable the Cisco Success Network on the chassis manager (see Cisco Firepower 4100/9300 FXOS Firepower Chassis Manager Configuration Guide).

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Telemetry**.

The **Enable Cisco Success Network** checkbox is selected by default.

Step 2 Ensure the Cisco Success Network is enabled by checking the **Enable Cisco Success Network** check box.

Step 3 To disable the Cisco Success Network, clear the **Enable Cisco Success Network** check box.

Step 4 Click **Apply**.

What to do next

- You can view the telemetry configuration and activities log or the telemetry data. See View ASA Telemetry Data , on page 1191

- To view a sample of telemetry data and the data fields, see Cisco Success Network - Telemetry Data, on page 1191

View ASA Telemetry Data

Before you begin

- Enable the telemetry service on ASA. See Enable or Disable Cisco Success Network , on page 1190

Procedure

Step 1 Choose **Monitoring** > **Properties** > **Telemetry**.

Step 2 Under **Telemetry**, click the relevant option:

- **History**—To view the past 100 events related to telemetry configuration and activities.
- **Sample**—To view the instantly generated telemetry data in JSON format.
- **Last-report**—To view the latest telemetry data sent to FXOS in JSON format.

Step 3 Click **Refresh** to view the report.

Cisco Success Network - Telemetry Data

Cisco Success Network is supported on Firepower 4100/9300 platforms by default. The FXOS service manager sends telemetry request daily to the ASA engine running on the platform. The ASA engine, on receiving the request, based on the connectivity status, sends the telemetry data either in standalone mode or cluster mode to FXOS. Following tables provide information on the telemetry data points, its description, and sample values.

Table 63: Device Info

Data Point	Description	Example Value
Device Model	Device model	Cisco Adaptive Security Appliance
Serial Number	Serial number of the device	FCH183771EZ
System Time	System uptime	11658000
Platform	Hardware	FPR9K-SM-24
Deployment Mode	Deployment type	Native
Security context mode	Single/Multiple	Single

Table 64: Versions Info

Data Point	Description	Example Value
Version Global Variable	ASA version	9.13.1.5
Device Manager Version	Device manager version	7.10.1

Table 65: License Info

Data Point	Description	Example Value
Smart License Global Variable	Activated licenses	regid.2015-01.com.cisco.ASA - SSP-STRONG-ENCRYPTION, 1.0_555507e9-85f8-4e41-96de-860b59f10bbe

Table 66: Platform Info

Data Point	Description	Example Value
CPU	CPU usage in past 5 minutes	fiveSecondsPercentage: 0.2000000, oneMinutePercentage: 0, fiveMinutesPercentage: 0
Memory	Memory usage	freeMemoryInBytes: 225854966384, usedMemoryInBytes: 17798281616, totalMemoryInBytes: 243653248000
Disk	Disk usage	freeGB: 21.237285, usedGB: 0.238805, totalGB: 21.476090
Bandwidth	Bandwidth usage	receivedPktsPerSec: 3, receivedBytesPerSec: 212, transmittedPktsPerSec: 3, transmittedBytesPerSec: 399

Table 67: Feature Info

Data Point	Description	Example Value
Feature List	Enabled feature list	name: cluster status: enabled

Table 68: Cluster Info

Data Point	Description	Example Value
Cluster Info	Cluster information	clusterGroupName : ssp-cluster interfaceMode : spanned unitName : unit-3-3 unitState : SLAVE otherMembers : items : memberName : unit-2-1 memberState : MASTER memberSerialNum : FCH183771BA

Table 69: Failover Info

Data Point	Description	Example Value
Failover	Failover information	myRole: Primary, peerRole: Secondary, myState: active, peerState: standby, peerSerialNum: FCH183770EZ

Table 70: Login Info

Data Point	Description	Example Value
Login	Login history	loginTimes: 2 times in last 2 days, lastSuccessfulLogin: 12:25:36 PDT Mar 11 2019

ASA Telemetry Data Sample

Following is an example of the telemetry data that are sent from ASA in JSON format. When service manager receives this input, it aggregates the data from all ASAs and adds necessary headers/fields before sending to the SSE connector. The headers/fields include “version”, “metadata”, “payload” with “recordedAt”, “recordType”, “recordVersion”, and ASA telemetry data, "smartLicenseProductInstanceIdentifier", "smartLicenseVirtualAccountName", and alike.

```
{
  "version": "1.0",
  "metadata": {
    "topic": "ASA.telemetry",
    "contentType": "application/json"
```

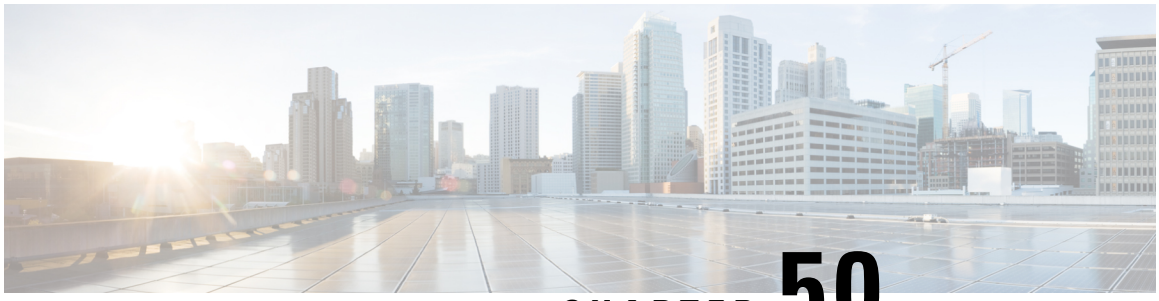
```
  },
  "payload": {
    "recordType": "CST_ASA",
    "recordVersion": "1.0",
    "recordedAt": 1557363423705,
    "SSP": {
      "SSPdeviceInfo": {
        "deviceModel": "Cisco Firepower FP9300 Security Appliance",
        "serialNumber": "JMX2235L01J",
        "smartLicenseProductInstanceIdentifier": "f85a5bb0-xxxx-xxxx-xxxx-xxxxxxxxx",
        "smartLicenseVirtualAccountName": "SSP-general",
        "systemUptime": 198599,
        "udiProductIdentifier": "FPR-C9300-AC"
      },
      "versions": {
        "items": [
          {
            "type": "package_version",
            "version": "92.7(1.342g)"
          }
        ]
      }
    },
    "asaDevices": {
      "items": [
        {
    "deviceInfo": {
        "deviceModel": "Cisco Adaptive Security Appliance",
        "serialNumber": "AANNNXXX",
        "systemUptime": 285,
        "udiProductIdentifier": "FPR9K-SM-36",
        "deploymentType": "Native",
        "securityContextMode": "Single"
    },
    "versions": {
        "items": [
            {
                "type": "asa_version",
                "version": "201.4(1)82"
            },
            {
                "type": "device_mgr_version",
                "version": "7.12(1)44"
            }
        ]
    },
    "licenseActivated": {
        "items": [
            {
                "type": "Strong encryption",
                "tag":
"regid.2015-01.com.cisco.ASA-SSP-STRONG-ENCRYPTION,1.0_xxxxxxx-xxxx-xxxx-96de-860b59f10bbe",

                "count": 1
            },
            {
                "type": "Carrier",
                "tag":
"regid.2015-01.com.cisco.ASA-SSP-MOBILE-SP,1.0_xxxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxx",
                "count": 1
            }
        ]
    },
    "CPUUsage": {
```

```
    "fiveSecondsPercentage": 0,
    "oneMinutePercentage": 0,
    "fiveMinutesPercentage": 0
},
"memoryUsage": {
    "freeMemoryInBytes": 99545662064,
    "usedMemoryInBytes": 20545378704,
    "totalMemoryInBytes": 120091040768
},
"diskUsage": {
    "freeGB": 21.237027,
    "usedGB": 0.239063,
    "totalGB": 21.476090
},
"bandwidthUsage": {
    "receivedPktsPerSec": 3,
    "receivedBytesPerSec": 268,
    "transmittedPktsPerSec": 4,
    "transmittedBytesPerSec": 461
},
"featureStatus": {
    "items": [
        {
            "name": "call-home",
            "status": "enabled"
        },
        {
            "name": "cluster",
            "status": "enabled"
        },
        {
            "name": "firewall_user_authentication",
            "status": "enabled"
        },
        {
            "name": "inspection-dns",
            "status": "enabled"
        },
        {
            "name": "inspection-esmtp",
            "status": "enabled"
        },
        {
            "name": "inspection-ftp",
            "status": "enabled"
        },
        {
            "name": "inspection-netbios",
            "status": "enabled"
        },
        {
            "name": "inspection-rsh",
            "status": "enabled"
        },
        {
            "name": "inspection-sip",
            "status": "enabled"
        },
        {
            "name": "inspection-sqlnet",
            "status": "enabled"
        },
        {
            "name": "inspection-sunrpc",
```

```
                "status": "enabled"
            },
            {
                "name": "inspection-tftp",
                "status": "enabled"
            },
            {
                "name": "inspection-xdmcp",
                "status": "enabled"
            },
            {
                "name": "logging-console",
                "status": "informational"
            },
            {
                "name": "management-mode",
                "status": "normal"
            },
            {
                "name": "sctp-engine",
                "status": "enabled"
            },
            {
                "name": "threat_detection_basic_threat",
                "status": "enabled"
            },
            {
                "name": "threat_detection_stat_access_list",
                "status": "enabled"
            },
            {
                "name": "webvpn-activex-relay",
                "status": "enabled"
            },
            {
                "name": "webvpn-dtls",
                "status": "enabled"
            }
        ]
    },
    "clusterInfo": {
        "clusterGroupName": "ssp-cluster",
        "interfaceMode": "spanned",
        "unitName": "unit-3-3",
        "unitState": "SLAVE",
        "otherMembers": {
            "items": [
                {
                    "memberName": "unit-2-1",
                    "memberState": "MASTER",
                    "memberSerialNum": "FCH183771BA"
                },
                {
                    "memberName": "unit-2-3",
                    "memberState": "SLAVE",
                    "memberSerialNum": "FLM1949C6JR"
                },
                {
                    "memberName": "unit-2-2",
                    "memberState": "SLAVE",
                    "memberSerialNum": "xxxxxxxx"
                },
                {
                    "memberName": "unit-3-2",
```

```
                    "memberState": "SLAVE",
                    "memberSerialNum": "xxxxxxxx"
                },
                {
                    "memberName": "unit-3-1",
                    "memberState": "SLAVE",
                    "memberSerialNum": "xxxxxxxx"
                }
            ]
        }
    },
    "loginHistory": {
        "loginTimes": "1 times in last 1 days",
        "lastSuccessfulLogin": "12:25:36 PDT Mar 11 2019"
    }
}
```



CHAPTER 50

Anonymous Reporting and Smart Call Home

This chapter describes how to configure the Anonymous Reporting and Smart Call Home services.

- About Anonymous Reporting, on page 1199
- About Smart Call Home, on page 1200
- Guidelines for Anonymous Reporting and Smart Call Home, on page 1201
- Configure Anonymous Reporting and Smart Call Home, on page 1202
- Monitoring Anonymous Reporting and Smart Call Home, on page 1206
- History for Anonymous Reporting and Smart Call Home, on page 1207

About Anonymous Reporting

You can help to improve the ASA platform by enabling Anonymous Reporting, which allows Cisco to securely receive minimal error and health information from the device. If you enable the feature, your customer identity will remain anonymous, and no identifying information will be sent.

Enabling Anonymous Reporting creates a trust point and installs a certificate. A CA certificate is required for your ASA to validate the server certificate present on the Smart Call Home web server and to form the HTTPS session so that your ASA can send messages securely. Cisco imports a certificate that is predefined in the software. If you decide to enable Anonymous Reporting, a certificate is installed on the ASA with a hardcoded trust point name: _SmartCallHome_ServerCA. When you enable Anonymous Reporting, this trust point is created, the appropriate certificate is installed, and you receive a message about this action. The certificate then appears in your configuration.

If the appropriate certificate already exists in your configuration when you enable Anonymous Reporting, no trust point is created, and no certificate is installed.

Note When you enable Anonymous Reporting, you acknowledge your consent to transfer the specified data to Cisco or to vendors operating on Cisco's behalf (including countries outside of the U.S.). Cisco maintains the privacy of all customers. For information about Cisco's treatment of personal information, see the Cisco Privacy Statement at the following URL: http://www.cisco.com/web/siteassets/legal/privacy.html

When the ASA configures Smart Call Home anonymous reporting in the background, the ASA automatically creates a trustpoint containing the certificate of the CA that issues the Call Home server certificate. The ASA now supports validation of the certificate if the issuing hierarchy of the server certificate changes, without the need for customer involvement to make certificate hierarchy changes. You can also automatically import the trustpool certificates so that ASA renews the certificate hierarchy without any manual intervention.

When you upgrade ASA 9.14(2.14), the trust point configuration automatically changes from CallHome_ServerCA to CallHome_ServerCA2.

DNS Requirement

A DNS server must be configured correctly for the ASA to reach the Cisco Smart Call Home server and send messages to Cisco. Because it is possible that the ASA resides in a private network and does not have access to the public network, Cisco verifies your DNS configuration and then configures it for you, if necessary, by doing the following:

1. Performing a DNS lookup for all DNS servers configured.
2. Getting the DNS server from the DHCP server by sending DHCPINFORM messages on the highest security-level interface.
3. Using the Cisco DNS servers for lookup.
4. Randomly using a static IP addresses for tools.cisco.com.

These tasks are performed without changing the current configuration. (For example, the DNS server that was learned from DHCP will not be added to the configuration.)

If there is no DNS server configured, and the ASA cannot reach the Cisco Smart Call Home Server, Cisco generates a syslog message with the warning severity level for each Smart Call Home message that is sent to remind you to configure DNS correctly.

See the syslog messages guide for information about syslog messages.

About Smart Call Home

When fully configured, Smart Call Home detects issues at your site and reports them back to Cisco or through other user-defined channels (such as e-mail or directly to you), often before you know that these issues exist. Depending on the seriousness of these problems, Cisco responds to your system configuration issues, product end-of-life announcements, security advisory issues, and so on by providing the following services:

- Identifying issues quickly with continuous monitoring, real-time proactive alerts, and detailed diagnostics.
- Making you aware of potential problems through Smart Call Home notifications, in which a service request has been opened, with all diagnostic data attached.
- Resolving critical problems faster with direct, automatic access to experts in Cisco TAC.

- Using staff resources more efficiently by reducing troubleshooting time.
- Generating service requests to Cisco TAC automatically (if you have a service contract), routed to the appropriate support team, which provides detailed diagnostic information that speeds problem resolution.

The Smart Call Home Portal offers quick access to required information that enables you to do the following:

- Review all Smart Call Home messages, diagnostics, and recommendations in one place.
- Check service request status.
- View the most up-to-date inventory and configuration information for all Smart Call Home-enabled devices.

Guidelines for Anonymous Reporting and Smart Call Home

This section includes the guidelines and limitation that you should review before configuring Anonymous reporting and Smart Call Home.

Anonymous Reporting Guidelines

- DNS must be configured.
- If an Anonymous Reporting message cannot be sent on the first try, the ASA retries two more times before dropping the message.
- Anonymous Reporting may coexist with other Smart Call Home configurations without changing the existing configuration. For example, if Smart Call Home is disabled before enabling Anonymous Reporting, it remains disabled, even after Anonymous Reporting has been enabled.
- If Anonymous Reporting is enabled, you cannot remove the trust point, and when Anonymous Reporting is disabled, the trust point remains. If Anonymous Reporting is disabled, you can remove the trust point, but disabling Anonymous Reporting does not cause the trust point to be removed.
- If you are using a multiple context mode configuration, the **dns**, **interface**, and **trustpoint** commands are in the admin context, and the **call-home** commands are in the system context.
- You can automate the update of the trustpool bundle at periodic intervals so that Smart Call Home can remain active if the self-signed certificate of the CA server changes. This trustpool auto renewal feature is not supported under multi-context deployments.

Smart Call Home Guidelines

- In multiple context mode, the subscribe-to-alert-group snapshot periodic command is divided into two commands: one to obtain information from the system configuration and one to obtain information from the user context.
- The Smart Call Home back-end server can accept messages in XML format only.
- A Smart Call Home message is sent to Cisco to report important cluster events if you have enabled clustering and configured Smart Call Home to subscribe to the diagnostic alert group with a critical severity level. A Smart Call Home clustering message is sent for only the following events:
    - When a unit joins the cluster

- When a unit leaves the cluster
- When a cluster unit becomes the cluster control unit
- When a secondary unit fails in the cluster

Each message that is sent includes the following information:

- The active cluster member count
- The output of the **show cluster info** command and the **show cluster history** command on the cluster control unit

Configure Anonymous Reporting and Smart Call Home

While Anonymous Reporting is part of the Smart Call Home service and allows Cisco to anonymously receive minimal error and health information from your device, the Smart Call Home service provides customized support of your system health, enabling Cisco TAC to monitor your devices and open a case when there is an issue, often before you know the issue has occurred.

You can have both services configured on your system at the same time, although configuring the Smart Call Home service provides the same functionality as Anonymous Reporting, plus customized services.

Configure Anonymous Reporting

To configure Anonymous Reporting, perform the following steps:

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Smart Call Home**.

Step 2 Check the **Enable Anonymous Reporting** check box.

Step 3 Click **Test Connection** to ensure that your system is able to send messages.

ASDM returns a success or error message to notify you of test results.

Step 4 Click **Apply** to save the configuration and enable Anonymous Reporting.

Configure Smart Call Home

To configure the Smart Call Home service, system setup, and alert subscription profiles, perform the following steps.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **Smart Call Home**.

Step 2 Check the **Enable Registered Smart Call Home** check box to enable Smart Call Home and register your ASA with Cisco TAC.

Step 3 Double-click **Advanced** System Setup. This area consists of three panes. Each pane can be expanded or collapsed by double-clicking the title row.

a) You can set up mail servers in the **Mail Servers** pane, through which Smart Call Home messages are delivered to e-mail subscribers.

b) You can enter the information of the person to contact in the **Contact Information** pane for the ASA that appears in Smart Call Home messages. This pane includes the following information:

- The name of the contact person.
- The contact phone number.
- The postal address of the contact person.
- The e-mail address of the contact.
- The “from” e-mail address in Smart Call Home e-mail.
- The “reply-to” e-mail address in Smart Call Home e-mail.
- The customer ID.
- The site ID.
- The contract ID.

c) You can adjust alert control parameters in the **Alert Control** pane. This pane includes the **Alert Group Status** pane, which lists the status (enabled or disabled) of the following alert groups:

- The diagnostics alert group.
- The configuration alert group.
- The environmental alert group.
- The inventory alert group.
- The snapshot alert group.
- The syslog alert group.
- The telemetry alert group.
- The threat alert group.
- The maximum number of Smart Call Home messages processed per minute.
- The “from” e-mail address in Smart Call Home e-mail.

Step 4 Double-click **Alert Subscription Profiles**. Each named subscription profile identifies subscribers and alert groups of interest.

a) Click **Add** or **Edit** to display the **Subscription Profile Editor**, in which you can create a new subscription profile or edit an existing subscription profile.

b) Click **Delete** to remove the selected profile.

c) Check the **Active** check box to send a Smart Call Home message of the selected subscription profile to subscribers.

Step 5 Click **Add** or **Edit** to display the **Add** or **Edit Alert Subscription Profile** dialog box.

a) The **Name** field is read-only and cannot be edited.
b) Check the **Enable this subscription profile** check box to enable or disable this particular profile.
c) Click either the **HTTP** or **Email** radio button in the **Alert Delivery Method** area.
d) Enter the e-mail address or web address in the **Subscribers** field.
e) Specify a **Reference Identity** object by name to enable RFC 6125 reference identity checks on the certificate received from the Syslog server.

See Configure Reference Identities, on page 741 for details on the reference identity object.

Step 6 The **Alert Dispatch** area lets the administrator specify which type of Smart Call Home information to send to subscribers and under what conditions. There are two types of alerts, time-based and event-based, chosen according to how the alert is triggered. The following alert groups are time-based: Configuration, Inventory, Snapshot, and Telemetry. The following alert groups are event-based: Diagnostic, Environmental, Syslog, and Threat.

Step 7 The **Message Parameters** area lets you adjust parameters that control messages sent to the subscriber, including the preferred message format and the maximum message size.

Step 8 For time-based alerts, click **Add** or **Edit** in the **Alert Dispatch** area to display the **Add** or **Edit Configuration Alert Dispatch Condition** dialog box.

a) Specify the frequency in the **Alert Dispatch Frequency** area in which to send the information to subscribers:

- For a monthly subscription, specify the day of the month, as well as the time of the day to send the information. If they are not specified, the ASA chooses appropriate values for them.
- For a weekly subscription, specify the day of the week, as well as the time of the day to send the information. If they are not specified, the ASA chooses appropriate values for them.
- For a daily subscription, specify the time of the day to send the information. If it is not specified, the ASA chooses an appropriate value for it.
- For an hourly subscription, specify the minute of the hour to send the information. If it is not specified, the ASA chooses an appropriate value for it. Hourly subscriptions are applicable to the snapshot and telemetry alert groups only.

b) Click the **Basic** or **Detailed** radio button to provide the desired level of information to subscribers.
c) Click **OK** to save the configuration.

Step 9 For diagnostic, environment, and threat event-based alerts, click **Add** or **Edit** in the **Alert Dispatch** area to display the **Create** or **Edit Diagnostic Alert Dispatch Condition** dialog box.

Step 10 Specify the event severity that triggers dispatch of the alert to subscribers in the **Event Severity** drop-down list, and then click **OK**.

Step 11 For inventory time-based alerts, click **Add** or **Edit** in the **Alert Dispatch** area to display the **Create** or **Edit Inventory Alert Dispatch Condition** dialog box.

Step 12 Specify how often to dispatch alerts to subscribers in the **Alert Dispatch Frequency** drop-down list, and then click **OK**.

Step 13 For snapshot time-based alerts, click **Add** or **Edit** in the **Alert Dispatch** area to display the **Create** or **Edit Snapshot Alert Dispatch Condition** dialog box.

a) Specify the frequency in the **Alert Dispatch Frequency** area in which to send the information to subscribers:

- For a monthly subscription, specify the day of the month, as well as the time of the day to send the information. If they are not specified, the ASA chooses appropriate values for them.
- For a weekly subscription, specify the day of the week, as well as the time of the day to send the information. If they are not specified, the ASA chooses appropriate values for them.
- For a daily subscription, specify the time of the day to send the information. If it is not specified, the ASA chooses an appropriate value for it.
- For an hourly subscription, specify the minute of the hour to send the information. If it is not specified, the ASA chooses an appropriate value for it. Hourly subscriptions are applicable to the snapshot and telemetry alert groups only.
- For an interval subscription, specify how often, in minutes, the formation is sent to the subscribers. This requirement is applicable to the snapshot alert group only.

b) Click **OK** to save the configuration.

Step 14 For syslog event-based alerts, click **Add** or **Edit** in the **Alert Dispatch** area to display the **Create** or **Edit Syslog Alert Dispatch Condition** dialog box.

a) Check the **Specify the event severity which triggers the dispatch of alert to subscribers** check box, and choose the event severity from the drop-down list.
b) Check the **Specify the message IDs of syslogs which trigger the dispatch of alert to subscribers** check box.
c) Specify the syslog message IDs that trigger dispatch of the alert to subscribers according to the on-screen instructions.
d) Click **OK** to save the configuration.

Step 15 For telemetry event-based alerts, click **Add** or **Edit** in the **Alert Dispatch** area to display the **Create** or **Edit Telemetry Alert Dispatch Condition** dialog box.

a) Specify the frequency in the **Alert Dispatch Frequency** area in which to send the information to subscribers:
   - For a monthly subscription, specify the day of the month, as well as the time of the day to send the information. If they are not specified, the ASA chooses appropriate values for them.
   - For a weekly subscription, specify the day of the week, as well as the time of the day to send the information. If they are not specified, the ASA chooses appropriate values for them.
   - For a daily subscription, specify the time of the day to send the information. If it is not specified, the ASA chooses an appropriate value for it.
   - For an hourly subscription, specify the minute of the hour to send the information. If it is not specified, the ASA chooses an appropriate value for it. Hourly subscriptions are applicable to the snapshot and telemetry alert groups only.

b) Click **OK** to save the configuration.

Step 16 Click **Test** to determine if the configured alerts are operating correctly.

Configure Auto Import of Trustpool Certificates

Smart licensing uses the Smart Call Home infrastructure. When the ASA configures Smart Call Home anonymous reporting in the background, the ASA automatically creates a trustpoint containing the certificate of the CA that issued the Call Home server certificate. The ASA now supports validation of the certificate if the issuing hierarchy of the server certificate changes, without the need for customer involvement to adjust certificate hierarchy changes. You can automate the update of the trustpool bundle at periodic intervals so that Smart Call Home can remain active if the self-signed certificate of the CA server changes. This feature is not supported under multi-context deployments.

Automatic import of trustpool certificate bundles requires you to specify the URL that ASA uses to download and import the bundle. Use the following command so the import happens daily at a regular interval with the default Cisco URL and default time of 22 hours:

```
ciscoasa(config-ca-trustpool)# auto-import-url Default
```

You can also enable auto import with a custom URL with the following command:

```
ciscoasa(config-ca-trustpool)# auto-import url http://www.thawte.com
```

To give you more flexibility to set downloads during off peak hours or other convenient times, enter the following command which enables the import with a custom time:

```
ciscoasa(config-ca-trustpool)# auto-import time 23:23:23
```

Setting the automatic import with both a custom URL and custom time requires the following command:

```
ciscoasa(config-ca-trustpool)# auto-import time 23:23:23 url http://www.thawte.com
```

Monitoring Anonymous Reporting and Smart Call Home

See the following commands for monitoring Anonymous Reporting and Smart Call Home services. You can enter these commands using **Tools > Command Line Interface**.

- **show call-home detail**

  This command shows the current Smart Call Home detail configuration.

- **show call-home mail-server status**

  This command shows the current mail server status.

- **show call-home profile** {profile name | **all**}

  This command shows the configuration of Smart Call Home profiles.

- **show call-home registered-module status** [**all**]

  This command shows the registered module status.

- **show call-home statistics**

  This command shows call-home detail status.

- **show call-home**

  This command shows the current Smart Call Home configuration.

- **show running-config call-home**

  This command shows the current Smart Call Home running configuration.

- **show smart-call-home alert-group**

  This command shows the current status of Smart Call Home alert groups.

- **show running-config all**

  This command shows details about the Anonymous Reporting user profile.

History for Anonymous Reporting and Smart Call Home

Table 71: History for Anonymous Reporting and Smart Call Home

Feature Name	Platform Releases	Description
Smart Call Home	8.2(2)	The Smart Call Home service offers proactive diagnostics and real-time alerts on the ASA, and provides higher network availability and increased operational efficiency. We introduced the following screen: Configuration > Device Management > Smart Call Home.
Anonymous Reporting	9.0(1)	You can help to improve the ASA platform by enabling Anonymous Reporting, which allows Cisco to securely receive minimal error and health information from a device. We modified the following screen: Configuration > Device Management > Smart Call Home.
Smart Call Home	9.1(2)	The **show local-host** command was changed to the **show local-host | include interface** command for telemetry alert group reporting.
Smart Call Home	9.1(3)	A Smart Call Home message is sent to Cisco to report important cluster events if you have enabled clustering and configured Smart Call Home to subscribe to the Diagnostic alert group with a Critical severity level. A Smart Call Home clustering message is sent for only the following three events: • When a unit joins the cluster • When a unit leaves the cluster • When a cluster unit becomes the cluster control unit Each message that is sent includes the following information: • The active cluster member count • The output of the **show cluster info** command and the **show cluster history** command on the cluster control unit

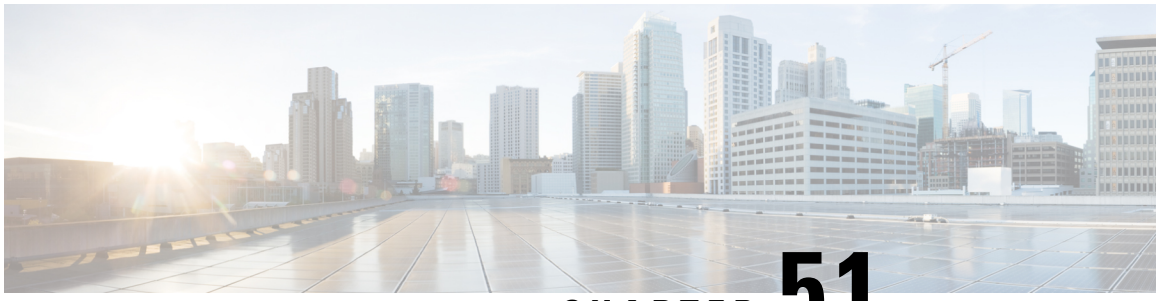
Feature Name	Platform Releases	Description
Reference Identities for Secure Smart Call Home Server connections	9.6(2)	TLS client processing now supports rules for verification of a server identity defined in RFC 6125, Section 6. Identity verification will be done during PKI validation for TLS connections to the Smart Call Home Server. If the presented identity cannot be matched against the configured reference identity, the connection is not established. We modifed the following page: **Configuration > Device Management > Smart Call Home**.



PART IX

Reference

- Addresses, Protocols, and Ports, on page 1211



CHAPTER 51

Addresses, Protocols, and Ports

This chapter provides a quick reference for IP addresses, protocols, and applications.

- IPv4 Addresses and Subnet Masks, on page 1211
- IPv6 Addresses, on page 1215
- Protocols and Applications, on page 1220
- TCP and UDP Ports, on page 1221
- Local Ports and Protocols, on page 1224
- ICMP Types, on page 1226

IPv4 Addresses and Subnet Masks

This section describes how to use IPv4 addresses in ASA. An IPv4 address is a 32-bit number written in dotted-decimal notation: four 8-bit fields (octets) converted from binary to decimal numbers, separated by dots. The first part of an IP address identifies the network on which the host resides, while the second part identifies the particular host on the given network. The network number field is called the network prefix. All hosts on a given network share the same network prefix but must have a unique host number. In classful IP, the class of the address determines the boundary between the network prefix and the host number.

Classes

IP host addresses are divided into three different address classes: Class A, Class B, and Class C. Each class fixes the boundary between the network prefix and the host number at a different point within the 32-bit address. Class D addresses are reserved for multicast IP.

- Class A addresses (1.xxx.xxx.xxx through 126.xxx.xxx.xxx) use only the first octet as the network prefix.
- Class B addresses (128.0.xxx.xxx through 191.255.xxx.xxx) use the first two octets as the network prefix.
- Class C addresses (192.0.0.xxx through 223.255.255.xxx) use the first three octets as the network prefix.

Because Class A addresses have 16,777,214 host addresses, and Class B addresses 65,534 hosts, you can use subnet masking to break these huge networks into smaller subnets.

Private Networks

If you need large numbers of addresses on your network, and they do not need to be routed on the Internet, you can use private IP addresses that the Internet Assigned Numbers Authority (IANA) recommends (see RFC 1918). The following address ranges are designated as private networks 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

Subnet Masks

A subnet mask lets you convert a single Class A, B, or C network into multiple networks. With a subnet mask, you can create an extended network prefix that adds bits from the host number to the network prefix. For example, a Class C network prefix always consists of the first three octets of the IP address. But a Class C extended network prefix uses part of the fourth octet as well.

Subnet masking is easy to understand if you use binary notation instead of dotted decimal. The bits in the subnet mask have a one-to-one correspondence with the Internet address:

- The bits are set to 1 if the corresponding bit in the IP address is part of the extended network prefix.
- The bits are set to 0 if the bit is part of the host number.

Example 1: If you have the Class B address 129.10.0.0 and you want to use the entire third octet as part of the extended network prefix instead of the host number, then you must specify a subnet mask of 11111111.11111111.11111111.00000000. This subnet mask converts the Class B address into the equivalent of a Class C address, where the host number consists of the last octet only.

Example 2: If you want to use only part of the third octet for the extended network prefix, then you must specify a subnet mask like 11111111.11111111.11111000.00000000, which uses only 5 bits of the third octet for the extended network prefix.

You can write a subnet mask as a dotted-decimal mask or as a /*bits* ("slash *bits*") mask. In Example 1, for a dotted-decimal mask, you convert each binary octet into a decimal number: 255.255.255.0. For a /*bits* mask, you add the number of 1s: /24. In Example 2, the decimal number is 255.255.248.0 and the /bits is /21.

You can also supernet multiple Class C networks into a larger network by using part of the third octet for the extended network prefix. For example, 192.168.0.0/20.

Determine the Subnet Mask

See the following table to determine the subnet mask based on how many hosts you want.



Note The first and last number of a subnet are reserved, except for /32, which identifies a single host.

Table 72: Hosts, Bits, and Dotted-Decimal Masks

Hosts	/Bits Mask	Dotted-Decimal Mask
16,777,216	/8	255.0.0.0 Class A Network
65,536	/16	255.255.0.0 Class B Network
32,768	/17	255.255.128.0
16,384	/18	255.255.192.0
8192	/19	255.255.224.0
4096	/20	255.255.240.0
2048	/21	255.255.248.0
1024	/22	255.255.252.0
512	/23	255.255.254.0
256	/24	255.255.255.0 Class C Network
128	/25	255.255.255.128
64	/26	255.255.255.192
32	/27	255.255.255.224
16	/28	255.255.255.240
8	/29	255.255.255.248
4	/30	255.255.255.252
Do not use	/31	255.255.255.254
1	/32	255.255.255.255 Single Host Address

Determine the Address to Use with the Subnet Mask

The following sections describe how to determine the network address to use with a subnet mask for a Class C-size and a Class B-size network.

Class C-Size Network Address

For a network between 2 and 254 hosts, the fourth octet falls on a multiple of the number of host addresses, starting with 0. For example, The following table shows the 8-host subnets (/29) of 192.168.0.x.



Note The first and last address of a subnet are reserved. In the first subnet example, you cannot use 192.168.0.0 or 192.168.0.7.

Table 73: Class C-Size Network Address

Subnet with Mask /29 (255.255.255.248)	Address Range
192.168.0.0	192.168.0.0 to 192.168.0.7
192.168.0.8	192.168.0.8 to 192.168.0.15
192.168.0.16	192.168.0.16 to 192.168.0.31
—	—
192.168.0.248	192.168.0.248 to 192.168.0.255

Class B-Size Network Address

To determine the network address to use with the subnet mask for a network with between 254 and 65,534 hosts, you need to determine the value of the third octet for each possible extended network prefix. For example, you might want to subnet an address like 10.1.*x*.0, where the first two octets are fixed because they are used in the extended network prefix, and the fourth octet is 0 because all bits are used for the host number.

To determine the value of the third octet, follow these steps:

1. Calculate how many subnets you can make from the network by dividing 65,536 (the total number of addresses using the third and fourth octet) by the number of host addresses you want.

   For example, 65,536 divided by 4096 hosts equals 16. Therefore, there are 16 subnets of 4096 addresses each in a Class B-size network.

2. Determine the multiple of the third octet value by dividing 256 (the number of values for the third octet) by the number of subnets:

   In this example, 256/16 = 16.

   The third octet falls on a multiple of 16, starting with 0.

The following table shows the 16 subnets of the network 10.1.



Note The first and last address of a subnet are reserved. In the first subnet example, you cannot use 10.1.0.0 or 10.1.15.255.

Table 74: Subnets of Network

Subnet with Mask /20 (255.255.240.0)	Address Range
10.1.0.0	10.1.0.0 to 10.1.15.255
10.1.16.0	10.1.16.0 to 10.1.31.255
10.1.32.0	10.1.32.0 to 10.1.47.255
—	—
10.1.240.0	10.1.240.0 to 10.1.255.255

IPv6 Addresses

IPv6 is the next generation of the Internet Protocol after IPv4. It provides an expanded address space, a simplified header format, improved support for extensions and options, flow labeling capability, and authentication and privacy capabilities. IPv6 is described in RFC 2460. The IPv6 addressing architecture is described in RFC 3513.

This section describes the IPv6 address format and architecture.

IPv6 Address Format

IPv6 addresses are represented as a series of eight 16-bit hexadecimal fields separated by colons (:) in the format: x:x:x:x:x:x:x:x. The following are two examples of IPv6 addresses:

- 2001:0DB8:7654:3210:FEDC:BA98:7654:3210
- 2001:0DB8:0000:0000:0008:0800:200C:417A



Note The hexadecimal letters in IPv6 addresses are not case-sensitive.

You do not need to include the leading zeros in an individual field of the address, but each field must contain at least one digit. So the example address 2001:0DB8:0000:0000:0008:0800:200C:417A can be shortened to 2001:0DB8:0:0:8:800:200C:417A by removing the leading zeros from the third through sixth fields from the left. The fields that contained all zeros (the third and fourth fields from the left) were shortened to a single zero. The fifth field from the left had the three leading zeros removed, leaving a single 8 in that field, and the sixth field from the left had the one leading zero removed, leaving 800 in that field.

It is common for IPv6 addresses to contain several consecutive hexadecimal fields of zeros. You can use two colons (::) to compress consecutive fields of zeros at the beginning, middle, or end of an IPv6 address (the colons represent the successive hexadecimal fields of zeros). The following table shows several examples of address compression for different types of IPv6 address.

Table 75: IPv6 Address Compression Examples

Address Type	Standard Form	Compressed Form
Unicast	2001:0DB8:0:0:0:BA98:0:3210	2001:0DB8::BA98:0:3210
Multicast	FF01:0:0:0:0:0:0:101	FF01::101
Loopback	0:0:0:0:0:0:0:1	::1
Unspecified	0:0:0:0:0:0:0:0	::



Note Two colons (::) can be used only once in an IPv6 address to represent successive fields of zeros.

An alternative form of the IPv6 format is often used when dealing with an environment that contains both IPv4 and IPv6 addresses. This alternative has the format x:x:x:x:x:x:y.y.y.y, where x represent the hexadecimal values for the six high-order parts of the IPv6 address and y represent decimal values for the 32-bit IPv4 part of the address (which takes the place of the remaining two 16-bit parts of the IPv6 address). For example, the IPv4 address 192.168.1.1 could be represented as the IPv6 address 0:0:0:0:0:0:FFFF:192.168.1.1 or ::FFFF:192.168.1.1.

IPv6 Address Types

The following are the three main types of IPv6 addresses:

- **Unicast**—A unicast address is an identifier for a single interface. A packet sent to a unicast address is delivered to the interface identified by that address. An interface may have more than one unicast address assigned to it.
- **Multicast**—A multicast address is an identifier for a set of interfaces. A packet sent to a multicast address is delivered to all addresses identified by that address.
- **Anycast**—An anycast address is an identifier for a set of interfaces. Unlike a multicast address, a packet sent to an anycast address is only delivered to the “nearest” interface, as determined by the measure of distances for the routing protocol.



Note There are no broadcast addresses in IPv6. Multicast addresses provide the broadcast functionality.

Unicast Addresses

This section describes IPv6 unicast addresses. Unicast addresses identify an interface on a network node.

Global Address

The general format of an IPv6 global unicast address is a global routing prefix followed by a subnet ID followed by an interface ID. The global routing prefix can be any prefix not reserved by another IPv6 address type.

All global unicast addresses, other than those that start with binary 000, have a 64-bit interface ID in the Modified EUI-64 format.

Global unicast address that start with the binary 000 do not have any constraints on the size or structure of the interface ID portion of the address. One example of this type of address is an IPv6 address with an embedded IPv4 address.

Site-Local Address

Site-local addresses are used for addressing within a site. They can be used to address an entire site without using a globally unique prefix. Site-local addresses have the prefix FEC0::/10, followed by a 54-bit subnet ID, and end with a 64-bit interface ID in the modified EUI-64 format.

Site-local routers do not forward any packets that have a site-local address for a source or destination outside of the site. Therefore, site-local addresses can be considered private addresses.

Link-Local Address

All interfaces are required to have at least one link-local address. You can configure multiple IPv6 addresses per interfaces, but only one link-local address.

A link-local address is an IPv6 unicast address that can be automatically configured on any interface using the link-local prefix FE80::/10 and the interface identifier in modified EUI-64 format. Link-local addresses are used in the neighbor discovery protocol and the stateless autoconfiguration process. Nodes with a link-local address can communicate; they do not need a site-local or globally unique address to communicate.

Routers do not forward any packets that have a link-local address for a source or destination. Therefore, link-local addresses can be considered private addresses.

IPv4-Compatible IPv6 Addresses

There are two types of IPv6 addresses that can contain IPv4 addresses.

The first type is the IPv4-compatibly IPv6 address. The IPv6 transition mechanisms include a technique for hosts and routers to dynamically tunnel IPv6 packets over IPv4 routing infrastructure. IPv6 nodes that use this technique are assigned special IPv6 unicast addresses that carry a global IPv4 address in the low-order 32 bits. This type of address is termed an IPv4-compatible IPv6 address and has the format ::y.y.y.y, where y.y.y.y is an IPv4 unicast address.



Note The IPv4 address used in the IPv4-compatible IPv6 address must be a globally unique IPv4 unicast address.

The second type of IPv6 address, which holds an embedded IPv4 address, is called the IPv4-mapped IPv6 address. This address type is used to represent the addresses of IPv4 nodes as IPv6 addresses. This type of address has the format ::FFFF:y.y.y.y, where y.y.y.y is an IPv4 unicast address.

Unspecified Address

The unspecified address, 0:0:0:0:0:0:0:0, indicates the absence of an IPv6 address. For example, a newly initialized node on an IPv6 network may use the unspecified address as the source address in its packets until it receives its IPv6 address.



Note The IPv6 unspecified address cannot be assigned to an interface. The unspecified IPv6 addresses must not be used as destination addresses in IPv6 packets or the IPv6 routing header.

Loopback Address

The loopback address, 0:0:0:0:0:0:0:1, may be used by a node to send an IPv6 packet to itself. The loopback address in IPv6 functions the same as the loopback address in IPv4 (127.0.0.1).



Note The IPv6 loopback address cannot be assigned to a physical interface. A packet that has the IPv6 loopback address as its source or destination address must remain within the node that created the packet. IPv6 routers do not forward packets that have the IPv6 loopback address as their source or destination address.

Interface Identifiers

Interface identifiers in IPv6 unicast addresses are used to identify the interfaces on a link. They need to be unique within a subnet prefix. In many cases, the interface identifier is derived from the interface link-layer address. The same interface identifier may be used on multiple interfaces of a single node, as long as those interfaces are attached to different subnets.

For all unicast addresses, except those that start with the binary 000, the interface identifier is required to be 64 bits long and to be constructed in the Modified EUI-64 format. The Modified EUI-64 format is created from the 48-bit MAC address by inverting the universal/local bit in the address and by inserting the hexadecimal number FFFE between the upper three bytes and lower three bytes of the of the MAC address.

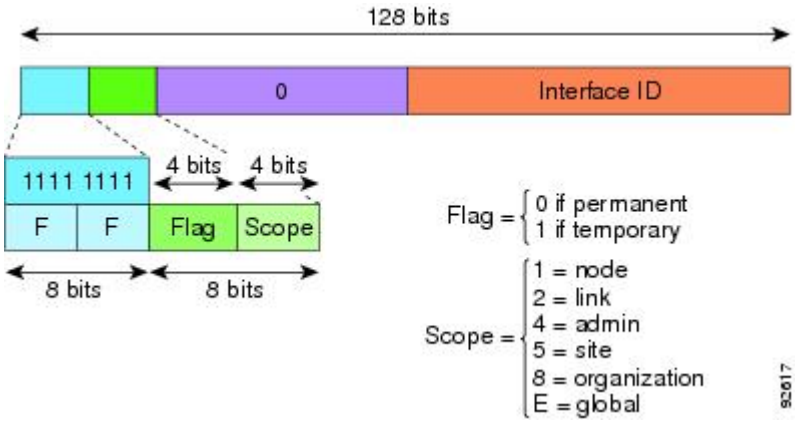
For example, and interface with the MAC address of 00E0.b601.3B7A would have a 64-bit interface ID of 02E0:B6FF:FE01:3B7A.

Multicast Address

An IPv6 multicast address is an identifier for a group of interfaces, typically on different nodes. A packet sent to a multicast address is delivered to all interfaces identified by the multicast address. An interface may belong to any number of multicast groups.

An IPv6 multicast address has a prefix of FF00::/8 (1111 1111). The octet following the prefix defines the type and scope of the multicast address. A permanently assigned (well known) multicast address has a flag parameter equal to 0; a temporary (transient) multicast address has a flag parameter equal to 1. A multicast address that has the scope of a node, link, site, or organization, or a global scope has a scope parameter of 1, 2, 5, 8, or E, respectively. For example, a multicast address with the prefix FF02::/16 is a permanent multicast address with a link scope. The following figure shows the format of the IPv6 multicast address.

Figure 108: IPv6 Multicast Address Format



IPv6 nodes (hosts and routers) are required to join the following multicast groups:

- The All Nodes multicast addresses:
  - FF01:: (interface-local)
  - FF02:: (link-local)
- The Solicited-Node Address for each IPv6 unicast and anycast address on the node: FF02:0:0:0:0:1:FFXX:XXXX/104, where XX:XXXX is the low-order 24-bits of the unicast or anycast address.

Note Solicited-Node addresses are used in Neighbor Solicitation messages.

IPv6 routers are required to join the following multicast groups:

- FF01::2 (interface-local)
- FF02::2 (link-local)
- FF05::2 (site-local)

Multicast address should not be used as source addresses in IPv6 packets.

Note There are no broadcast addresses in IPv6. IPv6 multicast addresses are used instead of broadcast addresses.

Anycast Address

The IPv6 anycast address is a unicast address that is assigned to more than one interface (typically belonging to different nodes). A packet that is routed to an anycast address is routed to the nearest interface having that address, the nearness being determined by the routing protocol in effect.

Anycast addresses are allocated from the unicast address space. An anycast address is simply a unicast address that has been assigned to more than one interface, and the interfaces must be configured to recognize the address as an anycast address.

The following restrictions apply to anycast addresses:

- An anycast address cannot be used as the source address for an IPv6 packet.
- An anycast address cannot be assigned to an IPv6 host; it can only be assigned to an IPv6 router.

Note Anycast addresses are not supported on the ASA.

Required Addresses

IPv6 hosts must, at a minimum, be configured with the following addresses (either automatically or manually):

- A link-local address for each interface
- The loopback address
- The All-Nodes multicast addresses
- A Solicited-Node multicast address for each unicast or anycast address

IPv6 routers must, at a minimum, be configured with the following addresses (either automatically or manually):

- The required host addresses
- The Subnet-Router anycast addresses for all interfaces for which it is configured to act as a router

- The All-Routers multicast addresses

IPv6 Address Prefixes

An IPv6 address prefix, in the format ipv6-prefix/prefix-length, can be used to represent bit-wise contiguous blocks of the entire address space. The IPv6-prefix must be in the form documented in RFC 2373 where the address is specified in hexadecimal using 16-bit values between colons. The prefix length is a decimal value that indicates how many of the high-order contiguous bits of the address comprise the prefix (the network portion of the address). For example, 2001:0DB8:8086:6502::/32 is a valid IPv6 prefix.

The IPv6 prefix identifies the type of IPv6 address. The following table shows the prefixes for each IPv6 address type.

Table 76: IPv6 Address Type Prefixes

Address Type	Binary Prefix	IPv6 Notation
Unspecified	000...0 (128 bits)	::/128
Loopback	000...1 (128 bits)	::1/128
Multicast	11111111	FF00::/8
Link-Local (unicast)	1111111010	FE80::/10
Site-Local (unicast)	1111111111	FEC0::/10
Global (unicast)	All other addresses.	
Anycast	Taken from the unicast address space.	

Protocols and Applications

The following table lists the protocol literal values and port numbers; either can be entered in ASA commands.

Table 77: Protocol Literal Values

Literal	Value	Description
ah	51	Authentication Header for IPv6, RFC 1826.
eigrp	88	Enhanced Interior Gateway Routing Protocol.
esp	50	Encapsulated Security Payload for IPv6, RFC 1827.
gre	47	Generic Routing Encapsulation.
icmp	1	Internet Control Message Protocol, RFC 792.
icmp6	58	Internet Control Message Protocol for IPv6, RFC 2463.
igmp	2	Internet Group Management Protocol, RFC 1112.

Literal	Value	Description
igrp	9	Interior Gateway Routing Protocol.
ip	0	Internet Protocol.
ipinip	4	IP-in-IP encapsulation.
ipsec	50	IP Security. Entering the ipsec protocol literal is equivalent to entering the esp protocol literal.
nos	94	Network Operating System (Novell's NetWare).
ospf	89	Open Shortest Path First routing protocol, RFC 1247.
pcp	108	Payload Compression Protocol.
pim	103	Protocol Independent Multicast.
pptp	47	Point-to-Point Tunneling Protocol. Entering the pptp protocol literal is equivalent to entering the gre protocol literal.
snp	109	Sitara Networks Protocol.
tcp	6	Transmission Control Protocol, RFC 793.
udp	17	User Datagram Protocol, RFC 768.

You can view protocol numbers online at the IANA website:

http://www.iana.org/assignments/protocol-numbers

TCP and UDP Ports

The following table lists the literal values and port numbers; either can be entered in ASA commands. See the following caveats:

- The ASA uses port 1521 for SQL*Net. This is the default port used by Oracle for SQL*Net. This value, however, does not agree with IANA port assignments.
- The ASA listens for RADIUS on ports 1645 and 1646. If your RADIUS server uses the standard ports 1812 and 1813, you can configure the ASA to listen to those ports using the **authentication-port** and **accounting-port** commands.
- To assign a port for DNS access, use the **domain** literal value, not **dns**. If you use **dns**, the ASA assumes you meant to use the **dnsix** literal value.

You can view port numbers online at the IANA website:

http://www.iana.org/assignments/port-numbers

Table 78: Port Literal Values

Literal	TCP or UDP?	Value	Description
aol	TCP	5190	America Online
bgp	TCP	179	Border Gateway Protocol, RFC 1163
biff	UDP	512	Used by mail system to notify users that new mail is received
bootpc	UDP	68	Bootstrap Protocol Client
bootps	UDP	67	Bootstrap Protocol Server
chargen	TCP	19	Character Generator
cifs	TCP, UDP	3020	Common Internet File System
citrix-ica	TCP	1494	Citrix Independent Computing Architecture (ICA) protocol
cmd	TCP	514	Similar to exec except that cmd has automatic authentication
ctiqbe	TCP	2748	Computer Telephony Interface Quick Buffer Encoding
daytime	TCP	13	Day time, RFC 867
discard	TCP, UDP	9	Discard
dnsix	UDP	195	DNSIX Session Management Module Audit Redirector
domain	TCP, UDP	53	DNS
echo	TCP, UDP	7	Echo
exec	TCP	512	Remote process execution
finger	TCP	79	Finger
ftp	TCP	21	File Transfer Protocol (control port)
ftp-data	TCP	20	File Transfer Protocol (data port)
gopher	TCP	70	Gopher
h323	TCP	1720	H.323 call signaling
hostname	TCP	101	NIC Host Name Server
http	TCP, UDP	80	World Wide Web HTTP
https	TCP	443	HTTP over SSL
ident	TCP	113	Ident authentication service

Literal	TCP or UDP?	Value	Description
imap4	TCP	143	Internet Message Access Protocol, version 4
irc	TCP	194	Internet Relay Chat protocol
isakmp	UDP	500	Internet Security Association and Key Management Protocol
klogin	TCP	543	KLOGIN
kshell	TCP	544	Korn Shell
ldap	TCP	389	Lightweight Directory Access Protocol
ldaps	TCP	636	Lightweight Directory Access Protocol (SSL)
login	TCP	513	Remote login
lotusnotes	TCP	1352	IBM Lotus Notes
lpd	TCP	515	Line Printer Daemon - printer spooler
mobile-ip	UDP	434	Mobile IP-Agent
nameserver	UDP	42	Host Name Server
netbios-dgm	UDP	138	NetBIOS Datagram Service
netbios-ns	UDP	137	NetBIOS Name Service
netbios-ssn	TCP	139	NetBIOS Session Service
nfs	TCP, UDP	2049	Network File System - Sun Microsystems
nntp	TCP	119	Network News Transfer Protocol
ntp	UDP	123	Network Time Protocol
pcanywhere-data	TCP	5631	pcAnywhere data
pcanywhere-status	UDP	5632	pcAnywhere status
pim-auto-rp	TCP, UDP	496	Protocol Independent Multicast, reverse path flooding, dense mode
pop2	TCP	109	Post Office Protocol - Version 2
pop3	TCP	110	Post Office Protocol - Version 3
pptp	TCP	1723	Point-to-Point Tunneling Protocol
radius	UDP	1645	Remote Authentication Dial-In User Service
radius-acct	UDP	1646	Remote Authentication Dial-In User Service (accounting)

Literal	TCP or UDP?	Value	Description
rip	UDP	520	Routing Information Protocol
rsh	TCP	514	Remote Shell
rtsp	TCP	554	Real Time Streaming Protocol
secureid-udp	UDP	5510	SecureID over UDP
sip	TCP, UDP	5060	Session Initiation Protocol
smtp	TCP	25	Simple Mail Transport Protocol
snmp	UDP	161	Simple Network Management Protocol
snmptrap	UDP	162	Simple Network Management Protocol - Trap
sqlnet	TCP	1521	Structured Query Language Network
ssh	TCP	22	Secure Shell
sunrpc	TCP, UDP	111	Sun Remote Procedure Call
syslog	UDP	514	System Log
tacacs	TCP, UDP	49	Terminal Access Controller Access Control System Plus
talk	TCP, UDP	517	Talk
telnet	TCP	23	RFC 854 Telnet
tftp	UDP	69	Trivial File Transfer Protocol
time	UDP	37	Time
uucp	TCP	540	UNIX-to-UNIX Copy Program
vxlan	UDP	4789	Virtual eXtensible Local Area Network (VXLAN)
who	UDP	513	Who
whois	TCP	43	Who Is
www	TCP, UDP	80	World Wide Web
xdmcp	UDP	177	X Display Manager Control Protocol

Local Ports and Protocols

The following table lists the protocols, TCP ports, and UDP ports that the ASA may open to process traffic destined to the ASA. Unless you enable the features and services listed in this table, the ASA does *not* open any local protocols or any TCP or UDP ports. You must configure a feature or service for the ASA to open

the default listening protocol or port. In many cases you can configure ports other than the default port when you enable a feature or service.

Table 79: Protocols and Ports Opened by Features and Services

Feature or Service	Protocol	Port Number	Comments
DHCP	UDP	67,68	—
Failover Control	105	N/A	—
HTTP	TCP	80	—
HTTPS	TCP	443	—
ICMP	1	N/A	—
IGMP	2	N/A	Protocol only open on destination IP address 224.0.0.1
ISAKMP/IKE	UDP	500	Configurable.
IPsec (ESP)	50	N/A	—
IPsec over UDP (NAT-T)	UDP	4500	—
IPsec over TCP (CTCP)	TCP	—	No default port is used. You must specify the port number when configuring IPsec over TCP.
NTP	UDP	123	—
OSPF	89	N/A	Protocol only open on destination IP address 224.0.0.5 and 224.0.0.6
PIM	103	N/A	Protocol only open on destination IP address 224.0.0.13
RIP	UDP	520	—
RIPv2	UDP	520	Port only open on destination IP address 224.0.0.9
SNMP	UDP	161	Configurable.
SSH	TCP	22	—
Stateful Update	8 (non-secure) 9 (secure)	N/A	—
Telnet	TCP	23	—
VPN Load Balancing	UDP	9023	Configurable.
VPN Individual User Authentication Proxy	UDP	1645, 1646	Port accessible only over VPN tunnel.

ICMP Types

The following table lists the ICMP type numbers and names that you can enter in ASA commands.

Table 80: ICMP Types

ICMP Number	ICMP Name
0	echo-reply
3	unreachable
4	source-quench
5	redirect
6	alternate-address
8	echo
9	router-advertisement
10	router-solicitation
11	time-exceeded
12	parameter-problem
13	timestamp-request
14	timestamp-reply
15	information-request
16	information-reply
17	mask-request
18	mask-reply
30	traceroute
31	conversion-error
32	mobile-redirect