



Cisco Secure Network Analytics Online Help 7.6.0



TOC

English (United States)	2
Glossary	1756

TOC

Domain Configuration: Add a Domain	75
Creating a Data Store Domain for Non-Data Store Users by Adding a Data Store Domain and Synchronizing	76
Creating a Data Store Domain for Non-Data Store Users by Importing an Existing Domain (Optional)	77
Related Topics	78
Domain Configuration: Synchronizing Domains	79
Before You Begin	79
Procedure	79
Removing a Domain Synchronization Target Domain	79
Using Dashboard Information to View Network Security Status	81
How to find this page	81
What you can view on this page	81
Global Time Interval	82
View the following components	83
What you can do next	90
Related Topics	91
Analytics: Alerts Summary	92
Open Alerts Summary	92
Filter the Alerts Table	92
View the Alerts Table	93
Edit Alerts Table Entries	94
View Related Configuration Pages	95
Download a .CSV File	95
Related Tasks	95
Related Topics	95
Analytics: Observations Dashboard	97
Open Observations Dashboard	97

Observations Highlights Overview	97
View Observations Highlights	98
View Descriptions for the Observations Highlights tables	98
Related Topics	105
ISE ANC Policy Assignments Report	106
How to find this page	106
View the ISE ANC Policy Assignments List	106
Related Topic	108
Analyzing Flows	109
Accessing Saved Flow Searches	110
How to find this page	110
What you can do on this page	110
Accessing Saved Flow Search Results	111
How to find this page	111
What you can do on this page	111
Flow Search: Overview	112
Open Flow Search	112
Flow Search: Table of Contents	112
Flow Search: Concepts You Need to Understand	113
Related Topics	113
Flow Search: How a Unique Flow is Defined	114
Example	114
Related Topics	115
Flow Search: Query Logic Used Between and Within Sections	116
Sections	116
Attributes	116
Related Topics	117
Flow Search: Search Logic Used in a Flow Search	118
Related Topics	119
Flow Search: IPv4 and IPv6 Address Notation	120

Examples	120
Related Topics	121
Flow Search: IPv4 Address Notation	122
Related Topics	123
Flow Search: IPv6 Address Notation	124
Related Topics	124
Flow Search: CIDR Notation	125
Related Topics	127
Flow Search: Entering Search Criteria Overview	129
Related Tasks	129
Related Topics	129
Flow Search: Guidelines For Entering Search Criteria	130
Alpha and Alphanumeric Search Criteria	130
Numeric Search Criteria	130
Related Tasks	132
Related Topics	132
Flow Search: General Search Criteria	133
Related Tasks	134
Related Topics	134
Flow Search: Subject Search Criteria	135
MAIN OPTIONS	135
ADVANCED OPTIONS	135
Related Tasks	136
Related Topics	136
Flow Search: Connection Search Criteria	137
MAIN OPTIONS	137
ADVANCED OPTIONS	137
Related Tasks	139
Related Topics	140
Flow Search: Peer Search Criteria	141

MAIN OPTIONS	141
ADVANCED OPTIONS	141
Related Tasks	142
Related Topics	142
Flow Search: Run a Flow Search Overview	143
Related Tasks	143
Related Topics	143
Flow Search: Data on Which a Flow Search Filters	144
Related Tasks	145
Related Topics	145
Flow Search: Run a Flow Search	146
Related Tasks	146
Related Topics	146
Flow Search: Run a Flow Search From a Menu	148
Default criteria that is used	149
Single IP address	149
Row of data	149
Graph	149
Related Tasks	150
Related Topics	150
Flow Search: Run a Flow Search From Saved Searches	151
Related Tasks	151
Related Topics	151
Flow Search Results: Overview	152
Open Flow Search Results	152
Flow Search Results: Table of Contents	152
Flow Search Results: View Flow Search Results Overview	153
Related Tasks	153
Related Topics	153
Flow Search Results: Navigate Results Page	154

Related Tasks	155
Related Topics	156
Flow Search Results: Filter Results	157
Alphanumeric Data	157
Numeric Data	157
Time-Related Data	158
Related Tasks	159
Related Topics	159
Flow Search Results: View Results	160
Related Tasks	160
Related Topics	160
Flow Search Results: Column Definitions	161
Related Tasks	179
Related Topics	179
Flow Search Results: What You Can Do Next Overview	180
Related Topics	180
Flow Search Results: Run a New Flow Search	181
Related Topics	181
Flow Search Results: Save Flow Search Results	182
Related Topics	182
Flow Search Results: Save Flow Search	183
Related Topics	183
Flow Search Results: Delete Results	184
Related Topics	184
Flow Search Results: Download the CSV File	185
Related Topics	185
Flow Search Results: Generate Network Diagrams	186
Related Topics	186
Investigating Flow Proxy Records	187
What you can view on this page	187

What you can do next	187
Related topics	187
Building a Top Report Search	188
Why you use this page	188
How to find this page	189
Concepts you need to understand to use this page	191
Examples	194
What you can do on this page	198
Search Type	199
Time Range	199
Search Name	199
Host IP Address or Range	200
Host Groups	200
Applications	200
Port/Protocol	200
Direction	201
Host IP Address or Range	201
Host Groups	201
Subject Orientation	201
Flow Collector Name	202
Records Returned	202
Interfaces	202
Order By	202
Performance Options	202
Default criteria that is used	205
Single IP Address	205
Row of data	206
Graph	206
What you can do next	206
Related topics	207

Viewing Top Report Results	208
How to find this page	208
Concepts you need to understand to use this page	210
Viewing information on this page	211
What you can do next	253
Related Topics	254
Determining the Existence of a Host	255
How to find this page	255
What you can do on this page	255
What you can do next	256
Related topics	257
Investigating Copyright Infringement	258
How to find this page	258
What you can do on this page	258
What you can view on this page	260
What you can do after results have been returned	261
Cisco Security Analytics and Logging (On Premises) Overview	262
Opening Security Analytics and Logging (OnPrem)	262
Configuring the Deployment	262
Security Analytics and Logging (OnPrem): Use the Event Viewer	263
Editing	263
Refining	263
Related Tasks	264
Related Topics	264
Security Analytics and Logging (OnPrem): Create a View	265
Related Tasks	265
Related Topics	265
Security Analytics and Logging (OnPrem): Update a View	266
Related Tasks	266
Related Topics	266

Security Analytics and Logging (OnPrem): Delete a View	267
Related Tasks	267
Related Topics	267
Security Analytics and Logging (OnPrem): Resources	268
Related Topics	268
Managing Jobs	269
How to find this page	269
What you can do on this page	269
What you can do next	273
Related Topics	273
Accessing Saved Flow Searches	275
How to find this page	275
What you can do on this page	275
Accessing Saved Flow Search Results	276
How to find this page	276
What you can do on this page	276
Investigating Hosts	277
How to find this page	277
Concepts you need to understand to use this page	278
What you can view on this page	281
What you can do next	281
Related topics	282
Investigating a Single Host	283
How to find this page	283
What you can view on this page	283
What you can do next	289
Related topics	290
Investigating Host Groups	291
How to find this page	291
What you can view on this page	292

What you can do next	299
Investigating Users	300
How to find this page	300
Concepts you need to understand to use this page	300
What you can view on this page	303
What you can do next	303
Related topics	303
Investigating a Single User	304
How to find this page	304
What you can do next	305
Related topics	305
Interface Monitoring: Overview	306
Open Response Management	306
Interface Monitoring: Table of Contents	306
Interface Monitoring: Filter Results Overview	307
Opening the Filter	307
What To Do Next	307
Related Topics	307
Interface Monitoring: Filter Results	308
What To Do Next	309
Related Topics	309
Interface Monitoring: Analyze Interface Status Data	310
Related Topics	311
Interface Monitoring: View Interface Traffic Details Overview	312
Opening the Interface Traffic Details	312
What To Do Next	312
Related Tasks	312
Related Topics	312
Interface Monitoring: View Interface Traffic Details	313
What To Do Next	313

Related Tasks	314
Related Topics	314
Interface Monitoring: View Top Application Traffic	315
What To Do Next	315
Related Tasks	316
Related Topics	316
Interface Monitoring: View the Packets Graph	317
Opening the Packets Graph	317
Viewing the Packets Graph	317
What To Do Next	317
Related Tasks	318
Related Topics	318
Interface Monitoring: View the Utilization Graph	319
Opening the Utilization Graph	319
Viewing the Utilization Graph	319
What To Do Next	319
Related Tasks	319
Related Topics	319
Interface Monitoring: Reports Overview	320
Related Tasks	320
Related Topics	320
Interface Monitoring: Report Filter Options	321
What To Do Next	322
Related Topics	322
Interface Monitoring: Run a Flow Search Report	323
Related Tasks	323
Related Topics	323
Interface Monitoring: Run a Top Report	325
Related Tasks	326
Related Topics	326

Analytics: Roles	327
Open the Roles page	327
Set the time frame	327
View the results	328
Export results	328
Related Topics	328
Report Builder	330
Opening Report Builder	330
Configuring Dashboards on Home Page	331
Best Practices	331
Failover	332
Related Tasks	332
Related Topic	332
Reports	333
Related Tasks	333
Related Topics	333
Report Templates	334
Related Tasks	336
Related Topics	337
Alarms Report	338
Parameters	338
Columns	339
Related Tasks	341
Related Topics	341
Alerts Report	342
Requirements	342
Parameters	342
Columns	342
Related Tasks	345
Related Topics	345

DSCP Status Report	346
Parameters	346
Columns	346
Related Tasks	348
Related Topics	348
Endpoint Traffic (NVM) Report	349
Requirements	349
Parameters	349
Report Results	349
Show More Results	350
Filter Data	350
Columns	350
Related Tasks	353
Related Topics	353
Firewall Log Collection Trend Report	354
Requirements	354
Parameters	354
Report Results	355
Related Tasks	355
Related Topics	355
Firewall Log Database Ingest Trend Report	356
Requirements	356
Parameters	356
Report Results	357
Related Tasks	357
Related Topics	357
Flow Collection Status Report	358
Parameters	358
Report Results	358
Common Scenarios	358

Columns	359
Configuration Information	361
Related Tasks	361
Related Topics	361
Flow Collection Trend by Exporter Report	362
Parameters	362
Report Results	362
Related Tasks	363
Related Topics	363
Flow Collection Trend by Flow Collector Report	364
Parameters	364
Report Results	364
Related Tasks	365
Related Topics	365
Flow Database Ingest Trend Report	366
Requirements	366
Parameters	366
Report Results	367
Related Tasks	367
Related Topics	367
Flow Rate Usage Report	368
Parameters	368
Report Results	368
Columns	368
Related Tasks	369
Related Topics	369
Host Group Application Traffic Report	370
Parameters	370
Columns	370
Related Tasks	371

Related Topics	372
Host Group Flow Traffic Report	373
Parameters	373
Columns	374
Related Tasks	375
Related Topics	375
Interface Application Traffic Report	376
Parameters	376
Columns	376
Related Tasks	377
Related Topics	377
Interface Network Traffic Report	378
Parameters	378
Columns	378
Related Tasks	379
Related Topics	380
Interface Service Traffic Report	381
Parameters	381
Columns	381
Related Tasks	382
Related Topics	382
Network and Server Performance Report	383
Parameters	383
Columns	384
Related Tasks	385
Related Topics	386
NVM Collection Trend Report	387
Requirements	387
Parameters	387
Report Results	388

Related Tasks	388
Related Topics	388
NVM Database Ingest Trend Report	389
Requirements	389
Parameters	389
Report Results	390
Related Tasks	390
Related Topics	390
Security Events Report	391
Parameters	391
Columns	391
Related Tasks	393
Related Topics	393
System Alarms Report	394
Parameters	394
Columns	394
Related Tasks	395
Related Topics	395
TrustSec Analytics Report	396
Parameters	396
Report Results	397
Time Range and Next Update	397
Security Group Tags (SGTs)	397
Colors and Icons	398
Security Group Activity	398
Policy Information	399
Related Tasks	400
Related Topics	400
TrustSec Policy Analytics Report	402
Parameters	402

Report Results	403
Time Range and Next Update	403
Security Group Tags (SGTs)	404
Colors and Icons	404
Security Group Activity	405
Policy Compliance	406
Related Tasks	408
Related Topics	408
Zeek Database Ingest Trend Report	409
Requirements	409
Parameters	409
Report Results	410
Related Tasks	410
Related Topics	410
Zeek Log Collection Trend Report	411
Requirements	411
Parameters	411
Report Results	412
Related Tasks	412
Related Topics	412
Zeek Logs Report	413
Requirements	413
Parameters	413
Report Results	414
Related Tasks	414
Related Topics	414
Creating a Report	415
Before You Start	415
Creating a Report	415
Related Tasks	416

Related Topics	417
Building Queries	418
Subject, Connection, Peer	418
Parameters	418
Example	419
Orientation	419
IPv4 IP Addresses	419
IPv6 IP Addresses	421
Related Tasks	421
Related Topics	421
Editing a Report	422
Before You Start	422
Editing a Report	422
Related Tasks	423
Related Topics	424
Deleting a Report	425
Related Tasks	425
Related Topics	425
Saving a Report	426
Related Tasks	426
Related Topics	426
Opening a Report	427
Related Tasks	427
Related Topics	427
Creating a TrustSec Report	428
Related Tasks	429
Related Topics	429
Investigating TrustSec Report Details	430
Related Tasks	430
Related Topics	430

Hiding/Showing Security Group Tags	431
Related Tasks	431
Related Topics	431
Investigating Possible Policy Violations	432
Related Tasks	434
Related Topics	434
Reviewing Colors and Icons	435
Related Tasks	436
Related Topics	436
Reviewing the Default Policy	437
Related Tasks	437
Related Topics	437
Reviewing Security Group Activity	438
Related Tasks	440
Related Topics	440
Running a Flow Search	441
Related Tasks	441
Related Topics	442
Running Top Reports	443
Related Tasks	443
Related Topics	443
Investigating the Details	445
Charts	445
Tables	445
Related Tasks	445
Related Topics	445
Adding a New Chart Type	446
Related Tasks	446
Related Topics	446
Filtering the Chart	447

Related Tasks	447
Related Topics	447
Reviewing Chart Details	448
Related Tasks	448
Related Topics	448
Running a Flow Search	449
Related Tasks	449
Related Topics	449
Running Top Reports	450
Related Tasks	450
Related Topics	450
Zooming In/Out	451
Related Tasks	451
Related Topics	451
Adjusting the Column Width	452
Related Tasks	452
Related Topics	452
Hiding/Showing Columns	453
Related Tasks	453
Related Topics	453
Moving a Column	454
Related Tasks	454
Related Topics	454
Sorting Table Columns	455
Related Tasks	455
Related Topics	455
Filtering the Endpoint Traffic (NVM) Report	456
Entering a Search Query	456
Filtering by Parameter	457
Related Tasks	457

Related Topics	457
Filtering a Table	458
Entering a Search Query	458
Filtering by Parameter	458
Related Tasks	459
Related Topics	459
Sorting Table Columns	460
Related Tasks	460
Related Topics	460
Dashboards	461
Related Tasks	461
Related Topics	461
Creating a Report Dashboard	462
Before You Start	462
Creating a Report Dashboard	462
Related Tasks	464
Related Topics	464
Creating a Network Insights Dashboard	466
Before You Start	466
Creating a Network Insights Dashboard	466
Related Tasks	467
Related Topics	468
Editing a Dashboard Layout	469
Related Tasks	469
Related Topics	469
Schedule History	470
Related Tasks	470
Related Topics	470
Scheduling a Report	471
Requirements	471

Supported Reports	471
Schedule a Report	471
Related Tasks	472
Related Topics	472
Deleting a Scheduled Report	473
Related Tasks	473
Related Topics	473
Exporting a Report	474
Related Tasks	474
Related Topics	474
Exporting an Endpoint Traffic (NVM) Report	475
Related Tasks	475
Related Topics	476
Visibility Assessment: Overview	477
Open Visibility Assessment	477
Visibility Assessment: Table of Contents	477
Visibility Assessment: Critical Summary Data Overview	478
Related Tasks	479
Related Topics	479
Security Events Versus Hits	480
Relates Topics	480
Visibility Assessment: Change the Time Range	482
Relates Topics	482
Visibility Assessment Reset All Data	483
Relates Topics	483
Visibility Assessment: Filter Columns	484
Clearing entries	484
Relates Topics	484
Visibility Assessment: Sort Columns	485
Relates Topics	485

Visibility Assessment: View Internal Network Scanners Data	486
What data does the Summary contain?	486
How the summary is sorted	486
Additional details you can view	486
Related Tasks	487
Related Topics	487
Visibility Assessment: View Remote Access Breach Data	488
What data does the Summary contain?	488
How the summary is sorted	488
Additional details you can view	488
Related Tasks	489
Related Topics	489
Visibility Assessment: View Possible Malware Data	490
What data does the Summary contain?	490
How the summary is sorted	490
Additional details you can view	490
Related Tasks	491
Related Topics	491
Visibility Assessment: View Vulnerable Protocol Servers Data	492
What data does the Summary contain?	492
How the summary is sorted	492
Additional details you can view	492
Related Tasks	493
Related Topics	493
Visibility Assessment: View DNS Risk Data	494
What data does the Summary contain?	494
How the summary is sorted	494
Additional details you can view	494
Related Tasks	495
Related Topics	495

Visibility Assessment: Internal Monitored Network Data Overview	496
What to do Next	496
Related Topics	496
Visibility Assessment: View Internal Monitored Network Data	497
Related Tasks	498
Related Topics	498
Visibility Assessment Report: Overview	499
What to do Next	499
Related Topics	499
View the Visibility Assessment Report	500
Related Topics	500
Visibility Assessment: Host Report Overview	501
Related Tasks	501
Related Topics	501
Visibility Assessment Host Report: Classify a Host	502
Related Tasks	502
Related Topics	502
Visibility Assessment Host Report: View Security Events for non-Top 10 Hosts	503
Related Tasks	503
Related Topics	503
Visibility Assessment Host Report: View Host Details	504
Related Tasks	504
Related Topics	504
Visibility Assessment: High Risk Countries Overview	505
Related Tasks	505
Related Topics	505
Visibility Assessment: View High Risk Countries Traffic	506
Related Tasks	506
Related Topics	506

Visibility Assessment: Configure the High Risk Countries Map	507
Related Tasks	507
Related Topics	507
Network Diagrams Overview	508
Who can access network diagrams?	508
Users most likely to use network diagrams	509
Open Network Diagrams	509
Network Diagrams: Table of Contents	509
Use Cases	510
Case 1	510
Case 2	510
Case 3	510
Related Tasks	510
Related Topics	510
Viewing Video Tutorials	511
Related Topics	511
Default Network Diagrams	512
Related Topics	512
Creating Network Diagrams Overview	513
Related Topics	513
1 Beginning Steps	514
What to do Next	514
Related Topics	514
2 Adding Host Groups	515
If you delete a host group in Secure Network Analytics	515
What to do Next	515
Related Topics	515
3 Creating Relationship Policies Overview	516
Related Tasks	516
Related Topics	516

Creating Relationship Policies Between Different Host Groups	517
What to do Next	518
Related Tasks	518
Related Topics	518
Creating Relationship Policies Within a Single Host Group	519
What to do Next	519
Related Tasks	519
Related Topics	519
(Optional) Adding One or More Additional Events to a Relationship Policy	520
Related Topics	520
(Optional) Enabling Events for a Relationship Policy to See Event Alarms for the Policy	521
Related Topics	521
(Optional) Adding Elements to Your Diagram	522
Related Topics	522
(Optional) Adding Images to Your Diagram Overview	523
Related Topics	523
(Optional) Adding Existing Icons or Pictures to a Diagram	524
Related Topics	524
(Optional) Uploading New Icons or Pictures	525
Related Topics	525
Interpreting Network Diagrams Overview	526
Related Topics	526
Filtering Your View	527
Related Tasks	527
Related Topics	527
Viewing Relationship Policy Traffic and Relationship Event Alarms	528
Traffic Volume	528
Alarms	528
Related Tasks	528

Related Topics	528
Viewing Core and Custom Security Event Alarms	530
Related Tasks	530
Related Topics	530
Viewing Preferences	531
Related Tasks	532
Related Topics	532
What You Can Pivot to From a Network Diagram Overview	533
Related Topics	533
Pivoting to Policies to Which a Host Belongs	534
Related Tasks	534
Related Topics	534
Pivoting to Relationship Policies Associated With an Edge	535
Related Tasks	535
Related Topics	535
Pivoting to Related Flow Traffic	536
Related Tasks	536
Related Topics	536
Pivoting to the Alarm List	537
View alarms associated with a relationship policy	537
View alarms associated with a host group	537
Related Tasks	537
Related Topics	537
Pivoting to Top Reports	539
Related Tasks	539
Related Topics	539
Pivoting to the Host Group Report	540
Related Tasks	540
Related Topics	540
Generating Network Diagrams From Flow Search Results Overview	541

Related Tasks	541
Related Topics	541
Generating Network Diagrams From Flow Search Results	542
Related Topics	542
Editing Relationship Policies Linked to a Network Diagram Overview	543
Related Tasks	543
Related Topics	543
Accessing Relationship Policies	544
Related Tasks	544
Related Topics	544
Enabling and Disabling Alarms	545
Related Tasks	545
Related Topics	545
Adding and Deleting Events	546
Add an event	546
Delete an event	546
Related Tasks	546
Related Topics	546
Editing Service and Application Filters	547
Related Tasks	547
Related Topics	547
Editing Event Details	548
Related Tasks	548
Related Topics	548
Working with Existing Network Diagrams Overview	549
Related Tasks	549
Related Topics	549
Opening a Network Diagram	550
Related Tasks	550
Related Topics	550

Editing Components in a Network Diagram	551
Related Tasks	551
Related Topics	551
Deleting Network Diagrams	552
Related Tasks	552
Related Topics	552
Duplicating Network Diagrams	554
Related Tasks	554
Related Topics	554
Importing Diagrams	555
Related Tasks	555
Related Topics	556
Exporting Network Diagrams	557
Related Tasks	557
Related Topics	557
Viewing Release Notes and the Secure Network Analytics Apps Compatibility Matrix	558
Related Topics	558
Managing and Configuring Host Groups	559
How to find this page	559
Concepts you need to understand to use this page	560
Examples	567
What you can do on this page	574
Host Group Tree	574
Host Group Name	575
Parent Host Group	575
Description	575
IP Addresses and Ranges	575
Advanced Options	575
Related topics	582

Configuring Host Group Automation	583
Related topics	588
About Alarm Severities	589
Access the Alarm Severity page	589
Clearing entries	592
Related topics	593
Alarm List: Overview	594
Alarm List: Exporter/Interface Alarms	596
Related Topics	598
Alarm List: Flow Collector System Alarms	599
Related Topics	602
Alarm List: Host Alarms	603
Related Topics	603
Alarm List: Host Group Relationship Alarms	604
Related Topics	607
Alarm List: Manager System Alarms	608
Related Topics	615
Alarm List: UDP Director Alarms	616
Related Topics	617
Alarm Suppression: Overview	618
Alarm Suppression: Table of Contents	618
Alarm Suppression: Alarms You Can Suppress	619
Related Tasks	625
Related Topics	625
Alarm Suppression: Rule Configuration	626
Rule Properties	626
Related Tasks	628
Related Topics	628
Alarm Suppression: Attribute Details	629
Related Tasks	632

Related Topics	632
Alarm Suppression: API Examples	633
Related Tasks	633
Related Topics	633
Security Event List	634
Addr (Address) Scan/tcp	635
Addr (Address) Scan/udp	640
Bad Flag ACK (Acknowledge) **	645
Bad Flag All **	649
Bad Flag NoFlg **	653
Bad Flag Rsvd (Reserved)	657
Bad Flag RST (Reset) **	661
Bad Flag SYN (Synchronize) FIN (Finish) **	665
Bad Flag URG (Urgent) **	669
Beaconing Host	673
Bot Command & Control Server	678
Bot Infected Host - Attempted C&C Activity	680
Bot Infected Host - Successful C&C Activity	682
Brute Force Login	684
Connection from Bogon Address Attempted	688
Connection from Bogon Address Successful	691
Connection from Tor Attempted	693
Connection from Tor Successful	694
Connection to Bogon Address Attempted	695
Connection to Bogon Address Successful	698
Connection to Tor Attempted	700
Connection to Tor Successful	702
Fake Application Detected	704
Flow Denied	709
Frag Packet Too Long **	711

Frag First Packet Too Short **	715
Frag Sizes Differ **	719
Half Open Attack	723
High File Sharing Index	726
High SMB Peers	730
High Total Traffic	734
High Traffic	736
High Volume Email	738
ICMP Comm (Communication) Admin **	742
ICMP Dest (Destination) Host Admin **	743
ICMP Dest (Destination) Host Unk (Unknown)**	744
ICMP Dest (Destination) Net Admin **	745
ICMP Dest (Destination) Net Unk (Unknown) **	746
ICMP Flood	747
ICMP Frag (Fragmentation) Needed **	751
ICMP Host Precedence **	752
ICMP Host Unreach **	754
ICMP Host Unreach TOS (Type of Service) **	755
ICMP Net Unreach **	756
ICMP Net Unreach TOS **	757
ICMP Port Unreach **	758
ICMP Precedence Cutoff **	759
ICMP Proto (Protocol) Unreach **	760
ICMP Received	761
ICMP Src (Source) Host Isolated **	763
ICMP Src (Source) Route Failed **	764
ICMP Timeout	765
Inside Tor Entry Detected	769
Inside Tor Exit Detected	771
Low Traffic	773

MAC Address Violation	774
Mail Rejects	775
Mail Relay	779
Max Flows Initiated	783
Max Flows Served	785
New Flows Initiated	787
New Flows Served	789
New Host Active	791
Packet Flood	792
Ping	797
Ping Oversized Packet	799
Ping Scan	804
Port Scan	809
Reset/tcp	813
Reset/udp	818
Scanner Talking	823
Slow Connection Flood	827
Spam Source	832
Src=Des (Source=Destination)	836
SSH Reverse Shell	838
Stealth Scan/tcp	842
Stealth Scan/udp	846
Suspect Data Hoarding	850
Suspect Data Loss	854
Suspect Long Flow	858
Suspect Quiet Long Flow	860
Suspect UDP Activity	862
SYN Flood	866
SYNs Received	870
Talks to Phantoms	872

Target Data Hoarding	877
Timeout/tcp	881
Timeout/udp	885
Touched	889
Trapped Host	890
UDP Flood	892
UDP Received	897
Watch Host Active	899
Watch Port Active	901
Worm Activity	903
Worm Propagation	908
Investigating Today's Active Alarms	912
How to find this page	912
Concepts you need to understand to use this page	912
What you can do on this page	913
What you can do next	914
Start date and time	916
End date and time	917
Related topics	918
Using Alarm Dashboard to Assess Potential Threats	919
How to find this page	919
What you can view on this page	919
Application Traffic	922
Summary of Sources	922
Alarms Table	923
Application Traffic	923
Summary of Targets	924
Alarms Table	924
What you can do next	927
Related Topics	930

Investigating Event Alarms	931
How to find this page	931
What you can view on this page	931
What you can do next	932
Related topics	935
Examining Details of Security Events	936
How to find this page	936
Option 1	937
Option 2	937
What you can view on this page	937
What you can do next	938
Related topics	939
Investigating Relationship Policy Alarms	940
How to find this page	940
What you can view on this page	940
What you can do next	941
Start date and time	942
End date and time	943
Related topics	944
System Notifications: Overview	945
System Notifications: Table of Contents	945
System Notifications: Notification Categories	946
General	946
Smart Licensing	946
Password Expiry	946
System Alarms	946
Related Topics	946
System Alarm Severity Descriptions	947
Related Topics	947
System Notifications: Notifications Panel Overview	948

Open the Notifications Panel	948
Related Tasks	948
Related Topics	948
System Notifications: View the Notifications Panel	949
What to do Next	950
Related Tasks	950
Related Topics	950
System Notifications: Alarm Detail Panel Overview	951
Open the Alert Details Panel	951
Related Tasks	951
Related Topics	951
System Notifications: View the Alarm Detail Panel	952
What to do Next	953
Related Tasks	953
Related Topics	953
System Notifications: Investigate System Notifications	954
Related Tasks	954
Related Topics	954
Secure Network Analytics Policies	955
Related Topics	957
Managing Custom Policies (Custom Security Events)	958
How to find this page	958
What you can do on this page	960
Filtering policies with multiple entries	963
Clearing entries	963
What you can do next	964
Related topics	965
Configuring Custom Policies (Custom Security Events)	966
How to find this page	966
Concepts you need to understand to use this page	967

Subject only	968
Subject or peer	968
Flow connection	969
Logic Terms	969
Rules vs. Values	970
IPv4 and IPv6 address formats	980
Examples	982
What you can do on this page	986
Query logic used for rules	988
When you delete a host group, service, or application	989
What you can do next	996
Related topics	996
Relationship Policy Management: Overview	998
Open Relationship Policy Management	998
Relationship Policy Management: Table of Contents	999
Relationship Policy Management: Use Context Menus	1000
Related Topics	1000
Relationship Policy Management: Filter Columns	1001
Filtering with multiple entries	1001
Editing a policy with filtered criteria	1001
Clearing entries	1002
Related Topics	1002
Relationship Policy Management: Sort Columns	1003
Related Topics	1003
Relationship Policy Management: What You Can Do On This Page Overview ..	1004
Related Topics	1004
Relationship Policy Management: View Current Policies	1005
What is a network diagram?	1005
Related Tasks	1006
Related Topics	1006

Relationship Policy Management: View or Edit Event Details of a Relationship Policy	1007
Related Tasks	1008
Related Topics	1009
Relationship Policy Management: Find All Relationship Policies Associated With a Host Group	1010
Related Tasks	1010
Related Topics	1010
Relationship Policy Management: Create a Policy	1011
Related Tasks	1011
Related Topics	1011
Relationship Policy Management: Edit a Policy	1012
Related Tasks	1012
Related Topics	1012
Relationship Policy Configuration: Overview	1013
Open Relationship Policy Configuration	1013
Relationship Configuration: Table of Contents	1013
Relationship Policy Configuration: Filter Columns	1014
Filtering with multiple entries	1014
Editing a policy with filtered criteria	1014
Clearing entries	1015
Related Topics	1015
Relationship Policy Configuration: Sort Columns	1016
Related Topics	1016
Relationship Policy Configuration: Concepts You Need to Understand	1017
Related Topics	1017
Relationship Policy Configuration: What You Can Do On This Page Overview	1018
Related Topics	1018
Relationship Policy Configuration: Create or Edit a Policy	1019
Related Topics	1021

Relationship Policy Configuration: Add an Event to a Relationship Policy	1022
Related Topics	1023
Relationship Policy Configuration: View or Edit Event Setting Details of a Relationship Policy	1025
Related Topics	1026
Relationship Policy Configuration: Enable or Disable an Event	1028
Related Topics	1028
Relationship Policy Configuration: Delete a Relationship Policy	1029
Related Topics	1029
Relationship Policy Configuration: Delete an Event from a Relationship Policy	1030
Related Topics	1030
Relationship Policy Configuration: Duplicate a Relationship Policy	1031
Related Topics	1031
Relationship Policy Configuration: Synchronize Your Flow Collector with the Manager	1032
Related Topics	1032
Managing Core Policies	1033
How to find this page	1033
Concepts you need to understand to use this page	1035
Default Policy, Inside Hosts	1047
Policy type	1048
"When Source is Host" and "When Source is Target" settings	1048
Scenario 1	1049
If the settings for the Host Policy are... ..	1049
When Host Is Source	1049
When Host Is Target	1049
And the settings for Role Policy 1 are... ..	1049
When Host Is Source	1049
When Host Is Target	1049
And the settings for Role Policy 2 are... ..	1049

When Host Is Source	1049
When Host Is Target	1049
Then the effective policy for the source host is... ..	1050
Then the effective policy for the target host is... ..	1050
Scenario 2	1050
If the settings for Role Policy 1 are... ..	1050
When Host Is Source	1050
When Host Is Target	1050
And the settings for Role Policy 2 are... ..	1050
When Host Is Source	1050
When Host Is Target	1050
Then the effective policy for the source host is... ..	1050
Then the effective policy for the target host is... ..	1050
Scenario 3	1050
If the settings for Role Policy 1 are... ..	1050
When Host Is Source	1050
When Host Is Target	1050
And the settings for Role Policy 2 are... ..	1051
When Host Is Source	1051
When Host Is Target	1051
Then the effective policy for the source host is... ..	1051
Then the effective policy for the target host is... ..	1051
Scenario 4	1051
If the settings for Role Policy 1 are... ..	1051
When Host Is Source	1051
When Host Is Target	1051
And the settings for Role Policy 2 are... ..	1051
When Host Is Source	1051
When Host Is Target	1051
Then the effective policy for the source host is... ..	1051

Then the effective policy for the target host is...	1051
Scenario 5	1052
If the settings for Role Policy 1 are...	1052
When Host Is Source	1052
When Host Is Target	1052
And the settings for Role Policy 2 are...	1052
When Host Is Source	1052
When Host Is Target	1052
Then the effective policy for the source host is...	1052
Then the effective policy for the target host is...	1052
Scenario 6	1052
If the settings for Role Policy 1 are...	1052
When Host Is Source	1052
When Host Is Target	1052
And the settings for Role Policy 2 are...	1052
When Host Is Source	1052
When Host Is Target	1052
Then the effective policy for the source host is...	1053
Then the effective policy for the target host is...	1053
What you can do on this page	1053
Filtering with multiple entries	1062
Editing a policy with filtered criteria	1063
Clearing entries	1063
What you can do next	1063
Related topics	1064
Configuring Host Policies	1065
How to find this page	1065
Concepts you need to understand to use this page	1066
Examples	1068
What you can do on this page	1072

Default Policy, Inside Hosts	1077
Filtering with multiple entries	1078
Editing a policy with filtered criteria	1078
Clearing entries	1079
What you can do next	1079
Related topics	1080
Configuring Role Policies	1081
How to find this page	1081
Concepts you need to understand to use this page	1082
Examples	1084
What you can do on this page	1088
Default Policy, Inside Hosts	1094
Filtering with multiple entries	1095
Editing a policy with filtered criteria	1095
Clearing entries	1095
What you can do next	1096
Related topics	1097
Editing Default Policies	1098
How to find this page	1098
What you can do on this page	1098
What you can do next	1099
Related topics	1099
Response Management: Overview	1101
Open Response Management	1102
Response Management: Table of Contents	1102
Response Management: Sort Columns	1103
Related Tasks	1103
Related Topics	1103
Response Management: IPv4 and IPv6 Address Notation	1104
Examples	1104

Response Management: IPv4 Address Notation	1106
Response Management: IPv6 Address Notation	1108
Response Management: Workflow	1109
What to do Next	1109
Related Tasks	1109
Related Topics	1109
Alarm Variables and Rule Types	1110
Related topics	1117
Alarm Variables Table	1118
Related topics	1125
Alert Variables Table	1126
Related topics	1129
Response Management Rules: Overview	1130
What to do Next	1130
Related Topics	1130
Response Management: Alarm Types	1131
Related Tasks	1131
Related Topics	1131
Response Management: Add Rules	1132
What to do Next	1133
Related Topics	1133
Alarm Tiering	1134
Related Tasks	1134
Related Topics	1134
Response Management: Add Condition Sets	1135
What to do Next	1139
Related Tasks	1139
Related Topics	1139
Response Management Conditions: Overview	1140
What to do Next	1140

Related Tasks	1140
Related Topics	1140
Response Management: Condition Types	1141
Alarms	1141
Alerts	1142
Related Tasks	1144
Related Topics	1144
Response Management: Add the Main Parent Operator and Conditions	1145
What to do Next	1146
Related Tasks	1146
Related Topics	1146
Response Management: Add Sub Conditions	1147
What to do Next	1148
Related Tasks	1148
Related Topics	1148
Response Management Actions: Overview	1149
What to do Next	1149
Related Topics	1149
Response Management: Action Types	1150
Related Tasks	1158
Related Topics	1158
Response Management: Add Actions to a Rule	1159
System performance regarding actions	1159
What to do Next	1160
Related Topics	1160
Response Management: Syslog Formats Overview	1161
What to do Next	1161
Related Topics	1161
Response Management: Add a Syslog Format to a Rule	1162
What to do Next	1163

Related Topics	1163
Facility Codes for Event Messages	1164
Related Tasks	1165
Related Topics	1165
Severity Levels for Event Messages	1166
Related Tasks	1166
Related Topics	1166
Response Management: Assign Actions Overview	1167
What to do Next	1167
Related Topics	1167
Response Management: Assign Actions to a Rule	1168
Related Topics	1168
Response Management: Rules, Actions, and Syslog Formats Overview	1169
Related Tasks	1169
Related Topics	1169
Response Management: Edit a Rule, Action, or Syslog Format	1170
Related Tasks	1170
Related Topics	1170
Response Management: Delete a Rule, Action, or Syslog Format	1171
Related Tasks	1171
Related Topics	1171
Response Management: Duplicate a Rule or Action	1172
Related Tasks	1172
Related Topics	1172
Classify Hosts Inside the Network: Overview	1173
How to find this page	1173
Classifying Hosts Inside the Network: Table of Contents	1173
Classify Hosts Inside the Network: Concepts You Need to Understand	1174
Related Topics	1174
Classify Hosts Inside the Network: Viewing Information	1175

Related Topics	1176
Classify Hosts Inside the Network: Classify Hosts Overview	1177
Related Topics	1177
Classify Hosts Inside the Network: Classify Multiple Hosts	1178
What To Do Next	1178
Related Topics	1178
Classify Hosts Inside the Network: Classify a Single Host	1179
What To Do Next	1179
Related Topics	1180
Classify Hosts Inside the Network: Apply a Policy	1181
Related Topics	1181
Analytics: Overview	1182
Enable Analytics	1183
Disable Analytics	1183
Analytics: Table of Contents	1183
Analytics: System Requirements	1184
Related Topics	1184
Analytics: APIs	1185
Related Topics	1185
Analytics: Alerts and Observations Overview	1186
Analytics: Alerts and Observations from Detections Findings	1187
Related Topics	1187
Analytics: Alert Prerequisite Chart	1188
Related Topics	1195
Analytics: Alert Descriptions	1196
Related Topics	1217
Analytics: Observation Descriptions	1218
Related Topics	1225
Analytics: Alerts Summary	1226
Open Alerts Summary	1226

Filter the Alerts Table	1226
View the Alerts Table	1227
Edit Alerts Table Entries	1228
View Related Configuration Pages	1229
Download a .CSV File	1229
Related Tasks	1229
Related Topics	1229
Analytics: Alerts Workflow	1231
Related Tasks	1231
Related Topics	1231
Analytics: Alerts FAQs	1232
Why are certain alerts disabled?	1232
What does an alert's status mean?	1232
How does a subnet's sensitivity affect alerts?	1232
Related Topics	1233
Analytics: Investigating Alerts	1234
Source Entities	1237
External Entities	1237
Related Tasks	1239
Related Topics	1239
Analytics: Alert Details	1240
Open Alert Details	1240
View Alert Type Details	1240
View Alert Rule Details	1240
View descriptions for the Supporting Observations table	1241
Other Pages You Can Access from the Alert Details Page	1248
Alert Priorities Configuration	1248
Flow Search Results	1248
Device Report	1248
Alerts Summary	1248

Observations by Device	1249
Enter Comments for an Alert	1249
Related Tasks	1249
Related Topics	1249
Analytics: Post Alerts to Cisco XDR	1250
Analytics: Device Report	1251
Open Device Report	1251
Device Report Overview	1251
Device Outline	1251
Alerts	1251
History	1252
Edit Time Range	1252
Summary	1252
Traffic	1253
IPs	1254
Related Topics	1255
Analytics: Roles	1256
Open the Roles page	1256
Set the time frame	1256
View the results	1257
Export results	1257
Related Topics	1257
Analytics: Observations Dashboard	1259
Open Observations Dashboard	1259
Observations Highlights Overview	1259
View Observations Highlights	1260
View Descriptions for the Observations Highlights tables	1260
Related Topics	1267
Analytics: Observation Types	1268
Open Observation Types	1268

View Observation Types	1268
Related Topics	1268
Analytics: Observations by Device	1269
Open Observations by Device	1269
View Observations by Device	1269
Related Topics	1270
Analytics: Selected Observations	1271
Open Selected Observations	1271
View Selected Observations	1271
View descriptions for the Selected Observations table	1271
Related Topics	1278
Analytics: Configuring Priorities	1279
Open Alert Priorities Configuration	1279
Configure Alert Priorities	1279
Related Topics	1280
Analytics: Configuring Expiration	1282
Open Alerts Expiration Configuration	1282
Configure Alert Expiration Days	1282
Related Topics	1282
Analytics: Configuring Country Watchlist	1283
Open Alerts Country Watchlist Configuration	1283
View Watched Countries	1283
Related Topics	1283
Troubleshooting	1284
Analytics jobs are lagging	1284
The secondary Manager has been promoted to primary Manager	1284
Your system is under-resourced for the amount of data being ingested	1284
Related Topics	1284
Detections Pack Management	1285
Review the Detections in the Pack	1285

Activate a Pack	1285
Operationalize Detections Pack	1286
Related Tasks	1286
Related Topics	1286
Configuring Services	1287
Clearing entries	1288
Custom Applications: Overview	1290
Open Custom Applications	1290
Custom Applications: Table of Contents	1290
Custom Applications: View Applications	1291
Related Topics	1291
Custom Applications: Add an Application	1292
Related Topics	1292
Custom Applications: Edit and Delete an Application	1293
Related Topics	1293
Domain Configuration: Overview	1294
Opening Domain Properties	1294
Deleting a Desktop Client Domain	1294
Domain Configuration: Table of Contents	1294
Domain Configuration: Configure Domain Settings	1295
Related Topics	1297
Domain Configuration: Add a Domain	1298
Creating a Data Store Domain for Non-Data Store Users by Adding a Data Store Domain and Synchronizing	1299
Creating a Data Store Domain for Non-Data Store Users by Importing an Existing Domain (Optional)	1300
Related Topics	1301
Domain Configuration: Edit a Domain	1302
Related Topics	1302
Domain Configuration: Delete a Domain	1303

Before you Begin	1303
Delete a Domain	1303
Related Topics	1303
Domain Configuration: Synchronizing Domains	1305
Before You Begin	1305
Procedure	1305
Removing a Domain Synchronization Target Domain	1305
Configure Flow Collectors: Overview	1307
Open Flow Collector Configuration	1307
Flow Collector Configuration: Table of Contents	1307
Configure Flow Collectors: Main Tab	1308
Open Flow Collector Configuration: Main Tab	1308
Global fields	1308
Data Collection	1308
System Alarms	1308
Store Flow Interface Data	1309
Alarm Configurations	1309
Related Topics	1311
Configure Flow Collectors: Advanced Tab	1312
Open Flow Collector Configuration: Advanced Tab	1312
Broadcast List	1312
Ignore List	1312
Watchlist	1312
Synchronize	1312
Flow Collector Security Thresholds	1313
Related Topics	1316
Exporters and Interfaces: Overview	1317
Open Exporters and Interfaces	1317
Exporters and Interfaces: Table of Contents	1317
Exporters and Interfaces: Filter Columns	1318

Clearing entries	1318
Related Tasks	1318
Related Topics	1318
Exporters and Interfaces: Sort Columns	1319
Related Tasks	1319
Related Topics	1319
Exporters and Interfaces: Manage Exporters Overview	1320
Related Topics	1320
Exporters and Interfaces: View Exporter Configurations	1321
Related Tasks	1322
Related Topics	1322
Exporters and Interfaces: Add an Exporter Configuration	1323
Related Tasks	1324
Related Topics	1324
Exporters and Interfaces: Edit a Single Exporter Configuration	1325
Related Tasks	1325
Related Topics	1325
Exporters and Interfaces: Bulk Edit Multiple Exporter Configurations	1326
Related Tasks	1326
Related Topics	1326
Exporters and Interfaces: Delete an Exporter Configuration	1327
Related Tasks	1327
Related Topics	1327
Exporters and Interfaces: Import a Single Exporter Configuration	1328
Related Tasks	1328
Related Topics	1328
Exporters and Interfaces: Export a Single Exporter Configuration	1329
Related Tasks	1329
Related Topics	1329
Exporters and Interfaces: Export all Exporter Configurations	1330

Related Tasks	1330
Related Topics	1330
Exporters and Interfaces: Manage Interfaces Overview	1331
Related Topics	1331
Exporters and Interfaces: View Interface Configurations	1332
Related Tasks	1334
Related Topics	1334
Exporters and Interfaces: Add an Interface Configuration	1335
Related Tasks	1337
Related Topics	1337
Exporters and Interfaces: Edit an Interface Configuration	1338
Related Tasks	1338
Related Topics	1338
Exporters and Interfaces: Bulk Edit Multiple Interface Configurations	1339
Related Tasks	1339
Related Topics	1339
Exporters and Interfaces: Delete an Interface Configuration	1340
Related Tasks	1340
Related Topics	1340
Exporter SNMP Profiles: Overview	1341
Open Exporter SNMP Profiles	1341
Exporter SNMP Profiles: Table of Contents	1341
Exporter SNMP Profiles: Filter Columns	1342
Clearing entries	1342
Related Tasks	1342
Related Topics	1342
Exporter SNMP Profiles: Sort Columns	1343
Related Tasks	1343
Related Topics	1343
Exporter SNMP Profiles: Manage SNMP Profile Configurations Overview	1344

Related Topics	1344
Exporter SNMP Profiles: View Exporter SNMP Configurations	1345
Related Topics	1346
Exporter SNMP Profiles: Add an Exporter SNMP Configuration	1347
Related Topics	1349
Exporter SNMP Profiles: Edit an Exporter SNMP Configuration	1350
Related Topics	1350
Exporter SNMP Profiles: Delete an Exporter SNMP Configuration	1351
Related Topics	1351
SNMP Configuration Fields for Custom Option	1352
Related Topics	1354
About Central Management	1355
Open Central Management	1355
Open Individual Appliances	1355
Open an Appliance through Central Management	1356
Open Appliance Administration through Direct Login	1356
Related Topics	1356
Central Management Inventory	1357
Appliance Status	1357
Actions Menu	1359
Related Tasks	1361
Related Topics	1361
Adding an Appliance to Central Management	1362
Add an Appliance to Central Management	1363
Related Tasks	1363
Related Topics	1363
Removing an Appliance from Central Management	1364
Troubleshooting	1364
Resetting Factory Defaults	1364
Related Tasks	1365

Related Topics	1365
	1366
Appliance Configuration	1367
Best Practices	1367
Open the Appliance Configuration Page	1367
Appliance Tab	1367
Network Services Tab	1368
General Tab	1368
External Services	1368
AIDE	1369
Open AIDE	1370
Enable AIDE	1370
Review AIDE Audit Logs	1370
Audit Log Destination	1372
Open Audit Log Destination	1372
1. Configure Audit Log Destination	1372
2. Enable SyslogCompliance	1373
Backup Configuration Encryption	1375
Open Backup Configuration Encryption	1375
Configure Encryption for Backups	1375
Cisco Security Cloud Control	1377
DoDIN Notifications	1378
Open DoDIN Notifications	1378
Configure and Enable DoDIN Notifications	1378
Compliance Mode	1381
FIPS	1381
Common Criteria	1381
Open Compliance Mode	1382
1. Prepare for Configuration	1382
2. Replace the Appliance Identity Certificates	1382

3. Enable FIPS Encryption Libraries	1383
Customer Success Metrics	1384
DNS Server and DNS Cache	1385
Open DNS Server and DNS Cache	1385
Add a New DNS Server	1385
Delete a DNS Server IP address	1385
Configure the Global DNS Time-Out Value	1385
Configure the DNS Cache Size	1386
Clear the Cache	1386
Host Naming	1387
Review Host Naming	1387
Internet Proxy	1388
Open Internet Proxy	1388
Enable Internet Proxy	1388
Local Resolution	1390
Open Local Resolution	1390
Add a New Host	1390
Delete a Host	1390
Network Interfaces	1391
Review Network Interfaces	1391
NTP Server	1392
1. Configure Authentication	1393
2. Confirm Authentication	1393
Opening Message	1395
Open Opening Message	1395
Create an Opening Message	1395
Password Policy	1396
Open Password Policy	1396
Define the Password Policy	1396
Protected Sessions Time-Out	1399

Open Protected Sessions Time-Out	1399
Set the Protected Sessions Time-Out	1399
Session Settings	1401
Open Session Settings	1401
Define Session Settings	1401
Related Tasks	1402
SMTP Configuration	1403
Open SMTP Configuration	1403
Configure SMTP for Response Management	1403
SNMP Agent	1405
Open SNMP Agent	1405
SSH	1413
Open SSH	1413
Enable SSH	1413
Disable SSH	1413
SSL/TLS Appliance Identity and Additional SSL/TLS Client Identities	1414
Review Certificates	1414
Certificate Expiration	1415
Threat Feed	1416
Cisco Talos IP Blocklist Integration	1416
Failover Configuration	1416
Enabling the Threat Feed	1416
1. License and Configure the Manager	1416
2. Enable the Threat Feed	1417
Review Alarms and Security Events	1417
Trust Store	1418
Review Trust Store Certificates	1418
Appliance Support	1419
Open the Appliance Support Page	1419
Audit Logs	1420

Open Audit Logs	1420
Review the Audit Log Data	1420
Sort by Column	1421
Filter Search Criteria	1421
Filter by Date	1421
Filter by Name, Location, Category	1422
Audit Log Destination	1423
Open Audit Log Destination	1423
1. Configure Audit Log Destination	1423
2. Enable SyslogCompliance	1424
Backup Configuration Files	1426
Create a Backup Configuration File	1428
Download a Backup Configuration File	1429
Restore an Appliance Backup Configuration File	1430
Review a Restoration Failure	1432
Diagnostics Packs	1433
Data Store Tab	1434
Opening the Data Store Tab	1434
Viewing the Data Store Database Status	1434
Starting a Data Node	1434
Stopping a Data Node	1435
Reviewing Last Action Results	1435
Related Topics	1435
Changing the Data Store Database Passwords	1436
Related Tasks	1436
Related Topics	1436
Adding a Data Store to a Non-Data Store Deployment and Transitioning Your Flow Collectors	1437
Data Store - Database Retention	1438
Database Fullness Chart	1438

Per Telemetry Contribution Chart	1438
Daily Storage	1438
Oldest Data in Data Store	1438
Opening the Data Store - Database Retention Tab	1439
Changing the Flow Interface Data Storage Period	1439
Related Topics	1439
Data Store - Database Update Status	1440
Opening the Data Store - Database Update Status Tab	1440
Monitoring the Database Update Status	1440
Related Topics	1441
Update Manager	1442
Open Update Manager	1442
Updates	1442
Updating Data Nodes	1442
Patches	1443
Types of Patches	1443
CIMC Firmware Version Common Update Patch	1443
Cisco Bundles Common Update Patch	1443
Detections Packs	1444
Related Topics	1444
App Manager	1445
App Compatibility Notice	1445
How to find this page	1446
What should you do on this page	1446
Access help	1449
Related Topics	1449
Smart Software Licensing	1450
Open Smart Software Licensing	1450
Product Instance Name	1450
Evaluation and Registration	1450

Transport Settings	1450
Database	1451
Opening the Database Tab	1451
Changing your Flow Collector Database Password	1451
Related Tasks	1451
User Management: Users Tab Overview	1452
Open the Users Tab	1452
Related Topics	1452
User Management: User Configuration: Table of Contents	1452
Users Tab: Use the Menu	1453
Related Topics	1453
Users Tab: Filter Columns	1454
Filtering with multiple entries	1454
Clearing entries	1454
Related Topics	1454
Users Tab: Manage Passwords	1455
Related Topics	1457
Users Tab: Enable and Disable User Accounts	1458
Related Tasks	1458
Related Topics	1458
Users Tab: Edit User Accounts	1459
Related Tasks	1459
Related Topics	1459
Users Tab: Delete User Accounts	1460
Related Tasks	1460
Related Topics	1460
Data Role Management: Data Roles Tab Overview	1461
Open the Data Roles Tab	1462
Related Topics	1462
Data Role Management: Data Roles Tab Overview: Table of Contents	1462

Data Roles Tab: Use Context Menus	1463
Related Topics	1463
Data Roles Tab: Filter Columns	1464
Filtering with multiple entries	1464
Editing a role with filtered criteria	1464
Clearing entries	1464
Related Topics	1464
Data Roles Tab: Sort Columns	1465
Related Topics	1465
Data Roles Tab: Edit Data Roles	1466
Related Topics	1466
Data Roles Tab: Delete Data Roles	1467
Related Topics	1467
Data Roles Tab: Duplicate Data Roles	1468
Related Topics	1468
Data Role Management: Data Role Configuration Overview	1469
Open the Data Roles Page	1469
What to do Next	1469
Related Topics	1469
Data Role Management: Data Role Configuration: Table of Contents	1470
Data Role Management: Create Data Roles	1471
Related Topics	1472
Managing Authentication and Authorization	1473
Authentication Services	1473
Authorization Services	1473
Related topics	1476
Configuring Authentication and Authorization	1477
Authentication Services	1477
Authorization Services	1477
RADIUS Message-Authenticator	1478

Support Details	1490
1. Prepare for Configuration	1491
2. Create an SSO Authentication Service	1491
Select the SSO Authentication Service	1491
Enter Your Basic Information	1492
Enter Your Identity Provider Information	1492
Disable Requested Authentication Context	1492
Add Your General Configuration Information	1493
3. Add a Local SSO User	1493
4. Test SAML Login	1494
Troubleshooting	1494
Related topics	1495
Session Limits	1496
Accessing Session Limits	1496
Configuring Session Limits	1496
Related Tasks	1496
Related Topics	1496
User Management: Sysadmin Password	1497
Open the Sysadmin Password Tab	1497
Change Sysadmin Password	1497
Related Topics	1497
User Management: User Configuration Overview	1498
Open the User page	1498
Related Topics	1499
User Management: User Configuration: Table of Contents	1499
User Configuration: Create User Accounts	1500
Related Tasks	1501
Related Topics	1501
User Configuration: Assign Data Roles to Users	1502
Related Tasks	1503

Related Topics	1503
User Configuration: View Web Role Permissions	1504
Related Tasks	1504
Related Topics	1504
User Configuration: Assign Web Roles to Users	1505
Related Tasks	1506
Related Topics	1506
User Configuration: Create Desktop Client Roles	1507
Related Tasks	1507
Related Topics	1508
User Configuration: Assign Desktop Client Roles to Users	1509
Related Tasks	1510
Related Topics	1510
Data Retention: Overview	1511
Open Data Retention	1511
Data Retention: Table of Contents	1512
Data Retention: Sort Columns	1513
Related Topics	1513
Data Retention: View Details	1514
Manager application configuration	1514
Data retention	1514
Related Topics	1514
Data Retention: Edit Data Retention Values	1515
Related Topics	1515
Default Data Retention Values	1516
Related Topics	1519
Data Retention: Configure Disk Space	1520
Related Topics	1520
DSCP Configuration: Overview	1521
Open DSCP Configurations	1521

DSCP Configurations: Table of Contents	1521
DSCP Configuration: Sort Columns	1522
Related Topics	1522
DSCP Configuration: Manage DSCP Configurations Overview	1523
Related Topics	1523
DSCP Configuration: View Configuration Details	1524
Manager application configuration	1524
DSCP configuration	1524
Related Tasks	1525
Related Topics	1525
DSCP Configuration: Add a DSCP Configuration	1526
Related Tasks	1526
Related Topics	1526
DSCP Configuration: Edit a DSCP Configuration	1527
Related Tasks	1527
Related Topics	1527
DSCP Configuration: Delete a DSCP Configuration	1528
Related Tasks	1528
Related Topics	1528
Failover Configuration	1529
Open Failover Configuration	1529
Configuring Failover	1529
Primary and Secondary Roles	1530
Redundant Site	1531
Requirements	1531
Adding Certificates to Trust Stores	1531
Trust Store Requirements	1531
Certificate Chain	1532
Uploading Certificates to the Trust Store	1532
1. Download the Appliance Identity Certificates	1532

2. Add Certificates to the Manager Trust Stores	1532
Open Site Redundancy Configuration	1533
Configuring a Redundant Site	1533
Disabling a Redundant Site	1534
Troubleshooting	1535
Interfaces and Exporters Cleaner	1536
Open the Interfaces and Exporters Cleaner	1536
Configure the Interfaces and Exporters Cleaner	1536
Configuring the UDP Director	1537
How to find this page	1537
What you can do on this page	1537
Related topics	1539
Forwarding Rules: Overview	1540
Open Forwarding Rules	1540
Forwarding Rules: Table of Contents	1540
Forwarding Rules: Add a Rule	1541
Related topics	1542
Forwarding Rules: Edit a Rule	1543
Related topics	1543
Forwarding Rules: Delete a Rule	1544
Related topics	1544
Forwarding Rules: Import a Rule	1545
Related topics	1545
Forwarding Rules: Export a Rule	1546
Related topics	1546
Managing External Lookup	1547
Related topics	1547
Configuring External Lookup	1548
Why you use this page	1548
How to find this page	1548

What you can do on this page	1548
Example 1	1549
Example 2	1551
What you can do next	1552
Related topics	1552
Performing an External Lookup	1553
Why you use this feature	1553
How to use this feature	1553
What you can do next	1553
Related Topics	1554
Managing Cisco® ISE or ISE-PIC	1555
Related topics	1555
Applying Mitigation Actions to a Host Using Cisco® ISE	1556
How to find Mitigation Actions links	1556
What you can do on this page	1556
Apply ANC Policies to a Host	1556
Removing ANC Policies from a Host	1557
What you can do next	1557
Related Topic	1557
Configuring Cisco® ISE or ISE-PIC	1558
What you can do next	1558
Related topics	1558
Viewing Mitigation Actions for a Host (Host History)	1559
How to find this page	1559
What you can view on this page	1559
What you can do next	1560
Related Topic	1560
ISE ANC Policy Assignments Report	1561
How to find this page	1561
View the ISE ANC Policy Assignments List	1561

Related Topic	1563
Configuring Secure Network Analytics to Use Active Directory	1564
How to find this page	1564
What you can do on this page	1565
What you can do next	1565
Managing the Cisco XDR Integration	1566
Access the Cisco XDR Integration	1566
Configure the Cisco XDR Integration	1566
Send Alerts and Observations to Cisco XDR	1566
Secure Network Analytics Help	1566
Get answers now	1566
Resources	1566
About the Cisco Secure Network Analytics Data Store	1569
About Cisco Secure Network Analytics Flow Collector NetFlow	1571
About NetFlow Version 9	1573
About the Cisco Secure Network Analytics Flow Collector sFlow	1576
About the Cisco Secure Network Analytics Flow Sensor	1578
Performance/Latency Metrics	1578
About the Cisco Secure Network Analytics Flow Sensor Virtual Edition	1580
About Secure Network Analytics Manager	1581
Related Topics	1581
Tasks You Can Perform	1582
Tasks you can perform only in the Desktop Client	1583
About the Cisco Secure Network Analytics UDP Director	1584
A note about load balancers	1585
Integrations	1586
App Manager	1588
App Compatibility Notice	1588
How to find this page	1589
What should you do on this page	1589

Access help	1592
Related Topics	1592
Home	1593
System	1593
Engine Status	1593
Services	1595
Network Interfaces	1595
Disk Usage	1595
Procedures	1596
UDP Director Engine Status	1596
Services	1596
Network Interfaces	1597
Field Descriptions	1597
Manager Docker Services	1603
Flow Collector Docker Services	1604
UDP Director and Flow Sensor Docker Services	1605
Related Topics	1609
Network Interface Configuration	1610
Status Report	1611
Field Descriptions	1611
Naming and DNS	1613
Network Host and IP Lookup	1613
Naming and DNS	1614
Network Host and IP Lookup	1614
System Time and NTP	1615
Services	1616
Alpha OIDs	1617
SSL Certificate	1621
Certificate Authority Certificates	1622
Forwarding Rules	1623

Enable the High Availability Cluster	1624
Field Descriptions	1624
Remote File System	1629
Proxy Ingest	1630
Procedures	1631
Add Proxy Server	1631
Add an Entry to Exclude	1631
Global Settings	1632
License Management	1633
Add/Edit/Delete Users	1634
Display this page	1634
Add a User	1634
Edit a User	1634
Delete a User	1635
Change Password	1636
Display the Change Password Page	1636
Change the Admin Password	1636
Backup/Restore Database (Appliance Admin Interface)	1639
Time	1640
Starting and Stopping the Backup	1640
Backup Display Information	1640
Back up a Database	1640
1. Disable SNMP Polling for a Manager	1641
2. Trim the Flow Collector Database	1641
1. Review your Database Storage Statistics	1642
2. Trim the Interface Details	1643
3. Trim Flow Details and CI Event Data	1643
3. Back up the Databases	1644
4. Re-enable SNMP Polling in the Manager	1645
Advanced Settings	1646

Procedures	1646
Advanced Settings for Flow Sensors	1652
Procedures	1652
Enable	1652
Disable	1652
Procedures	1653
Enter or change advanced settings	1653
Procedures	1655
Enable	1655
Disable	1655
Procedures	1655
Enable	1655
Disable	1656
Procedures	1656
Enable	1656
Assign the monitoring interface an IP addresses	1657
Add parameter information to flowsensor.xml	1657
Disable	1658
Procedures	1658
Enable	1658
Disable	1658
Procedures	1659
Enable	1659
Disable	1659
Procedures	1659
Enable	1659
Disable	1659
Procedures	1660
Specify which mode the Flow Sensor 4240 uses	1660
Specify which mode the Flow Sensor 4300 uses	1660

Procedures	1660
To select the type of cache mode	1660
Updating the IP Mode by Logging into the Flow Sensor	1661
Updating the IP Mode by Removing and Re-Adding the Flow Sensor to Central Management	1662
Database Storage Statistics	1664
Field Descriptions	1664
Backup/Restore Database (Appliance Admin Interface)	1670
Time	1671
Starting and Stopping the Backup	1671
Backup Display Information	1671
Back up a Database	1671
1. Disable SNMP Polling for a Manager	1672
2. Trim the Flow Collector Database	1672
1. Review your Database Storage Statistics	1673
2. Trim the Interface Details	1673
3. Trim Flow Details and CI Event Data	1674
3. Back up the Databases	1675
4. Re-enable SNMP Polling in the Manager	1676
Browse Files	1677
Procedures	1677
Packet Capture	1678
Generate a Packet Capture	1678
Tcpdump Expressions	1680
Tcpdump Template Examples	1680
Tcpdump Expressions	1681
Advanced BPF Filters	1688
Update Manager	1695
Open Update Manager	1695
Updates	1695


Updating Data Nodes	1695
Patches	1696
Types of Patches	1696
CIMC Firmware Version Common Update Patch	1696
Cisco Bundles Common Update Patch	1697
Detections Packs	1697
Related Topics	1697
Backup/Restore Configuration	1698
Diagnostics Pack	1699
Audit Log	1700
Operations	1701
About Central Management	1702
Open Central Management	1702
Open Individual Appliances	1702
Open an Appliance through Central Management	1703
Open Appliance Administration through Direct Login	1703
Related Topics	1703
About the Desktop Client	1704
Smart Software Licensing	1706
Open Smart Software Licensing	1706
Product Instance Name	1706
Evaluation and Registration	1706
Transport Settings	1707
Update Manager	1708
Open Update Manager	1708
Updates	1708
Updating Data Nodes	1708
Patches	1709
Types of Patches	1709
CIMC Firmware Version Common Update Patch	1709

Cisco Bundles Common Update Patch	1710
Detections Packs	1710
Related Topics	1710
Navigating the Manager	1711
System Navigation: Table of Contents	1711
Compatibility	1712
Supported Browsers	1712
Supported Operating Systems	1712
Use Context Menus	1713
Install the Desktop Client Overview	1714
What to do Next	1714
Install the Desktop Client Using Mac OS	1715
Change the memory size	1716
Install the Desktop Client Using Windows	1718
Change the memory size	1719
Use Filters	1721
Choose a Theme	1725
Before you log in	1725
After you log in	1725
Select a Language	1726
Before you log in	1726
After you log in	1726
Log Out	1727
Change Your Profile Name	1728
Change Your Password	1729
View System Notifications	1730
Search for a Host, IP Range, or User	1731
Secure Network Analytics Help	1731
Get answers now	1731
Resources	1731

Useful Information	1734
Search Tips	1734
Related Topics	1734
Access Help	1735
Search Tips	1735
Related Topics	1735
Desktop Client Help	1736
Glossary	1737
Secure Network Analytics Help	1737
Get answers now	1737
Resources	1737
Glossary	1741
Index	1752

Domain Configuration: Add a Domain

1. From the main menu, choose **[Current domain name] > Add Domain**.



2. Configure the following fields:

Domain Name: The name to be assigned to the domain. This name is shown on the Host Group tree.

Select Method: Select one of the methods described in the table below to designate which host group structure you want to use for the domain you are adding.

If you want to make this a Data Store domain, make sure you check the **Configure as a Data Store domain** check box.

If you select this method...	Then...
Default	Secure Network Analytics creates the domain with the default host group structure but without any Flow Collectors.
Import from File	Secure Network Analytics creates the domain and uses the appropriate configuration, based on the specific domain content you exported (host group, domain, or both). • XML files containing the domain configuration are not backwards compatible. These files are only compatible within the same system version number (for example, from Flow Collector v7.0 to Manager v7.0). • You can also import the entire host group configuration using the Host Group Management page.

If you select this method...	Then...
	• If you need to import interface groups in the Network Devices branch of the Host Group tree from another domain, use this option. You must first export the groupings as an XML file to your local drive. • None of the Flow Collectors contained in the XML file is imported.
 If you add a Flow Collector to an existing domain, that domain's specific configuration (policy, alarm severity, services, exporter SNMP, etc.) is applied to this Flow Collector.	

3. If you want to add a Data Store domain, check the **Configure as a Data Store domain** check box. A Data Store domain is for Secure Network Analytics systems that are using a Data Store and a Non-Data Store domain is for Secure Network Analytics systems that are not using a Data Store.
4. To save your configurations, click **Add**.
5. Complete the settings on this page. (See next section.)

Creating a Data Store Domain for Non-Data Store Users by Adding a Data Store Domain and Synchronizing

If you are in the process of transitioning a Non-Data Store Flow Collector to a Data Store Flow Collector and want to keep your data synchronized between your Non-Data Store domain and your newly created Data Store domain, follow the process outlined in [Domain Configuration: Synchronizing Domains](#). For more information on transitioning, refer to the [Secure Network Analytics System Configuration Guide](#).

Follow the steps below to add a new Data Store that you will be synchronizing with your Non-Data Store domain.

1. From the main menu, choose **[Current domain name] > Add Domain**.
2. Configure the following fields:

Domain Name: The name to be assigned to the domain. This name is shown on the Host Group tree.

Select Method: Select the Default method.

If you want to make this a Data Store domain, make sure you check the **Configure as a Data Store domain** check box.

3. To save your configurations, click **Add**.
4. From the main menu, choose **[Current domain name] > the Non Data Store domain that you want to synchronize with your Data Store domain**.
5. From the main menu, choose **Configure > System > Domain Properties**.
6. Select the **Edit** button.
7. Choose the Data Store domain that you want to synchronize this domain with in the **Target Domain to Synchronize** drop-down menu.
8. Click the **Save** button to save your changes.
9. Click the Synchronize button  that appears next to your Non-Data Store domain in the main menu to synchronize your domains.

Creating a Data Store Domain for Non-Data Store Users by Importing an Existing Domain (Optional)

If you are currently on a Non-Data Store domain and you want to add a Data Store domain to your Secure Network Analytics system for a future expansion into Data Store, you can do so by importing a Non-Data Store configuration into a new Data Store domain. When you import a existing domain, you won't have to re-configure items like alarms, host groups, and so on. Importing from an existing domain is like creating a new domain but with an existing configuration.

If the domain is newly created, then you will have to re-configure your Secure Network Analytics settings.

Follow the steps below to add a new Data Store domain and import all of its configuration from your Non-Data Store domain.

1. From the main menu, choose **[Current domain name] > the Non- Data Store domain that you want to synchronize with your Data Store domain**.
2. From the main menu, choose **Configure > System > Domain Properties**.
3. Make sure the **Export All configuration** radio button is selected. Refer to the [Domain Configuration: Configure Domain Settings](#) topic to view a list of the data that is exported.
4. Click the **Export** button to download the XML file.
5. From the main menu, choose **[Current domain name] > Add Domain**.
6. Enter a name for your new domain in the **Domain Name** field.

7. From the **Select Method** drop-down menu, select the **Import from File** option.
8. Select the XML file you downloaded in [step 4](#).
9. Check the **Configure as a Data Store domain** check box.
10. To add your new domain, click **Add**.

Related Topics

[Domain Configuration: Overview](#)

[Domain Configuration: Synchronizing Domains](#)

[Managing and Configuring Host Groups](#)

[Investigating Host Groups](#)

Domain Configuration: Synchronizing Domains

If you are in the process of transitioning a Non-Data Store Flow Collector to a Data Store Flow Collector, you may want to keep your configurations and tuning synchronized between your Non-Data Store domain and your Data Store domain. This topic describes the process for synchronizing your Non-Data Store domain with its associated Data Store domain.

Before You Begin

Ensure that you have already created a Data Store domain that you will be synchronizing with your Non-Data Store domain. If you have already followed the process described in [Adding a Data Store to a Non-Data Store Deployment and Transitioning Your Flow Collectors](#), your Data Store domain should already be created. Refer to [Domain Configuration: Add a Domain](#) for instructions on adding a domain.

Procedure

Follow these steps to synchronize your Non-Data Store domain (Source) with your Data Store domain (Target).

1. From the main menu, choose **[Current domain name] > the Non Data Store domain that you want to synchronize with your Data Store domain**.
2. From the main menu, choose **Configure > Domain Properties**.
3. Select the **Edit** button.
4. Choose the Data Store domain that you want to synchronize this domain with in the **Target Domain to Synchronize** drop-down menu.



You can only synchronize your target Data Store domain with one source Non-Data Store domain. If you attempt to synchronize your target Data Store domain with more than one source Non-Data Store domain, you will receive the following error, "This target domain is already in use by another Non-Data Store domain."

5. Click the **Save** button to save your changes.

Removing a Domain Synchronization Target Domain

Follow the steps below to remove a target domain.

1. From the main menu, choose **[Current domain name] > the Non Data Store domain that you want to synchronize with your Data Store domain**.

2. From the main menu, choose **Configure > Domain Properties**.
3. Select the **Edit** button.
4. Click the **Clear Target Domain** button.
5. Click the **Save** button to save your changes.

Using Dashboard Information to View Network Security Status

Use this page to view multiple dashboards on your Home page. Quickly see the default and customized reports on the dashboards.

How to find this page

Do one of the following to access this page

- When you log in to the Manager, the Security Insight Dashboard opens as the Home page.
- You can also access it by doing the following from any page:
 - From the main menu, choose **Monitor** > **Security Insights Dashboard**.

What you can view on this page

Viewing Security Dashboard

Choose the **Security** tab, to view the default **Security Insights Dashboard**. Use this page to quickly see the events that have triggered alarms.

Some information provides data for the past seven days, and some information concerns data for the last 24 hours.



Make sure to toggle to **Alerts view** to see the Alerts report. This page defaults to the **Alarms view**, refer to [Alarms View and Alerts View](#) for more information.

View the dashboard components that appear across the top of the page, which provide information about alarms, applications, and flows, provide both current and trend information.

You can do the following with the dashboard components:

- Print or download the chart or graph in a component. Several components provide a menu of commands to use if you want to perform either of these functions.
- To have a detailed view of information in a chart or table, click the  (**Fullscreen**) icon in the component's title bar to open the component in a separate window.



- Unless otherwise noted, all hosts listed are inside hosts.

- The information you see may be affected by point of view settings. If you do not have access to view certain types of information, those components are not displayed on this dashboard.

Viewing Pinned Dashboards

You can pin and view the dashboards to your Home page. You can navigate only pinned dashboards on the Home page as tabs beside the **Security** tab.

To pin a custom dashboard, do the following:

1. From the main menu, choose **Report > Report Builder > Dashboards**.
2. Choose a saved dashboard and click the  (**Pin**) icon . For more details, refer to [Report Builder](#).



If you have at least one custom dashboard pinned to a Home page you can modify your dashboards view by clicking **Customize** button when on any of your pinned dashboards tab.

Global Time Interval

There is a Global Time Interval of default **Last 24 Hours** applied to the pinned dashboards and the reports in the dashboards.

To change the time interval for the dashboard, click the **Last 24 Hours** and choose from the drop-down. The changes will be applied to all the reports in the opened dashboard. You can modify the **Time Range** for individual reports in the opened dashboard.



The Global Time Interval will be reset to default **Last 24 Hours** whenever you navigate over the dashboards tabs on the Home page.

The error  icon will appear on a report if:

- There is any report in the dashboard which is configured to be different from the Global Time Interval, regardless of what time range was selected in the Report Builder.
- The selected Time Range doesn't apply to that report.

View the following components

Alarms View and Alerts View

The following topics provide details about the **Alarms view** display:

- [Alarming Hosts](#)
- [Alarms by Type](#)
- [Flow Collection Trend](#)
- [Sorting Criteria](#)
- [Today's Alarms](#)
- [Top Alarming Hosts](#)
- [Top Applications](#)

When you toggle to **Alerts view**, the [Alerts Report](#) displays.

Alarming Hosts

Alarming hosts are shown across the top of the dashboard in various alarm categories (such as Concern Index or Recon). An alarm category is a “bucket” toward which a defined list of security events contributes. When network activity meets or exceeds a defined set of criteria specified for this alarm category, it triggers an alarm. Alarm categories contain only host alarms. (A host alarm is a security event that has triggered an alarm.) To learn more about alarm severity levels, see [About Alarm Severities](#).

Refer to the following guidelines regarding alarming hosts:

- Although the following "always bad" security events do not contribute index points to an alarm category, any of these immediately increases the alarm count for a host and assigns it a severity level of major (which causes both the number and the alarm trend to turn red).
 - All custom security events
 - Any security event contributing to the Policy Violation alarm (Note that these security events may contribute to other alarms as well, but they cause only the Policy Violation alarm to turn black.)
- If you reboot your Flow Collector, it deletes all alarm history. Since the alarming host widgets (which display the number of hosts receiving alarms since the last reset hour for a specific category) on the Security Insight dashboard don't update until the next reset hour, you may see a discrepancy between these values and the alarm

values on the Host List. Note that if you replace your Flow Collector, the new Flow Collector retains the alarm history from the old Flow Collector, and the values on the Security Insight dashboard and the Host List remain in sync.

- If you disable a custom event that had resulted in any Policy Violation alarms, the Host List may show fewer alarming hosts. Enabling the custom event may or may not restore the count.

The alarm categories at the top of this page actively receive alarms that affect this host, and they display the following data:

- The number of hosts receiving alarms (since the last reset hour) for that category. If no hosts are receiving alarms that contribute to the alarm category, the container displays a zero. If you click the number (value cannot be 0), the [Host List](#) page opens, filtered by that category.
- The color represents the current highest severity for an alarm occurring for any host within that category. If the highest severity is major, the color of the corresponding number turns red; if it is minor, it turns orange. If no alarms are contributing to an alarm category, its number turns gray.
- The 7-Day Trend chart's color always matches the number's color. For details about a specific day, place your cursor over the graph; a tooltip displays the date and the number of hosts receiving alarms (for a specific day) that contribute to that alarm category.


The alarm categories are as follows:

Name of Alarm Category	Description
Anomaly	Tracks events that indicate that hosts are behaving abnormally or generating traffic that is unusual, but isn't consistent with another category of activity.
C&C (Command & Control)	Indicates the existence of infected servers or hosts in your network attempting to contact a C&C server.
Concern Index	Indicates that the host's concern index has either exceeded the CI threshold or rapidly increased.


Name of Alarm Category	Description
Data Hoarding	Indicates that a source or target host within a network has downloaded an unusual amount of data from one or more hosts.
Exfiltration	Indicates that an abnormal amount of data has been transferred to inside or outside hosts.
Exploitation	Tracks direct attempts by hosts to compromise each other, such as through worm propagation and brute force password cracking.
DDoS Source	Indicates that a host has been identified as the source of a DDoS attack.
DDoS Target	Indicates that a host has been identified as the target of a DDoS attack.
Policy Violation	Indicates that the subject is exhibiting behavior that violates normal network policies.
Recon	Indicates the presence of unauthorized and potentially malicious scans using TCP or UDP and being run against your organization's hosts. These scans, referred to as "reconnaissance," are early indicators of attacks against your network, and the scans may come from outside or inside your organization.
Target Index	Indicates that the target IP has been the recipient of more than an acceptable number of scans or other malicious attacks.

Example 1:


For instances in which no hosts have received any contributing alarms since the last reset hour nor at any time during the last seven days, no trend chart is displayed.

**Example 2:**

For instances in which no hosts have received any contributing alarms since the last reset hour but at least one host has received a contributing alarm during the last seven days, the trend chart is dark gray.

**Example 3:**

For instances in which at least one host has received a contributing alarm since the last reset hour, the number is either red or orange, depending on whether the highest alarm severity assigned is major or minor. The 7-Day Trend chart's color is the same color as the number.



Top Alarming Hosts

This component displays the top seven inside hosts, sorted by alarm severity (represented by alarm category containers), that have been active on your network since the last reset hour. If a host is a member of two or more host groups and receives more than one alarm, Secure Network Analytics displays only one of the host groups of which this host is a member. The order of the hosts on this list may vary slightly from that of the Host List if the Manager and Flow Collector have not yet synced.

To see the alarm category name and percentage of threshold (or alarm count for the “always bad” alarms) reached by a host for a particular alarm category, hover your cursor over the applicable alarm category container.

From this component you can access the following information:

- The [Host List](#), which displays all hosts that have had active alarms since the last reset hour. To view this data, click the View All Hosts link at the bottom of the component. The Host List opens, sorted by severity.
- The [Host Report](#) for the host which you selected. You can view information about a single host’s activity as far back as the last seven days. To view this data, click the applicable host IP address.
- The [Alarm dashboard](#), which displays all alarms contributing to a specific alarm category for a specific host. To view this data, click the alarm category container for the host you want to investigate.

 To learn more about alarm severity levels, see [About Alarm Severities](#).

Sorting Criteria

To learn more about alarm severity levels, see [About Alarm Severities](#).

1. The primary sort is by hosts with alarms that have one or more contributing critical, or ["always bad," security events](#). These security events are so serious that it does not matter how many index points they are worth or how much they have contributed to the percentage of threshold reached for a particular alarm. Hosts with these alarms are always at the top of the list.
2. The secondary sort is by hosts based on severity of alarm.
3. The third level of sorting is by hosts with alarms that contain the highest percentage of threshold reached.

Example:

- A host has three alarms each with a medium level of severity and each at 40% of threshold, for a total of 120% of threshold.
- A second host has one alarm with a high severity level (not caused by an "always bad" security event) and which is at 80% of threshold.
- A third host has one alarm with a medium severity level and which is at 80% of threshold.

In the sort list, the alarm associated with the second host has a higher severity level, so that host will be at the top of the list. The first host, which has the three medium severity level alarms, will be next, because this host possesses the highest percentage of threshold reached. The third host will be the last one listed because it has a lower percentage of threshold reached than the first host.

Alarms by Type

This bar chart displays the number of host alarms and host group relationship alarms that have occurred over the past seven days. Each type of alarm is represented by a different color, which is identified in the legend beneath the component's title. When you hover your cursor over a bar, it displays the number of alarms for the alarm type you are hovering over that have occurred on the day indicated. It also displays the total number of all alarms, regardless of type, that occurred on that day. Click an alarm name in the legend to remove its data from displaying on the graph. Click the alarm name again to restore its data on the graph. You can use this chart to focus on the volume and frequency of a particular type of alarm's occurrence.

When you click a section on the chart representing an alarm category, the Alarm List opens. You can use this page to investigate alarms that have been active (but are not currently active) on your network since the last archive hour. It lists each alarm

occurrence and details such as alarm duration, hosts or host groups affected, whether or not it has been acknowledged, etc. This bar chart corresponds to the Today's Alarms graphic and shows the distribution of alarms.

Today's Alarms

This component displays in graphic form the host alarms and host group relationship alarms that are active on your system for the current day (i.e., for the 24-hour period since the last archive hour, which by default occurs at 12:00 a.m.). If you click the pie slice that corresponds to the applicable alarm, the Alarm List opens. You can use this page to investigate alarms that have been active (but are not currently active) on your network since the last archive hour. It lists each alarm occurrence and details such as alarm duration, hosts or host groups affected, whether or not it has been acknowledged, etc. This graphic corresponds to the current day bar on the Alarms by Type component and shows the distribution of alarms.

Flow Collection Trend

This component displays a line graph of the rate of traffic in flows per minute for a 24-hour period based on the archive hour. The chart aggregates data for all of the Flow Collectors on your system and provides visual information about trends and spikes in the flow of traffic on your network.

If you are accustomed to using the Desktop Client, you may be familiar with the Flow Collection Trend chart, which uses a bar chart to represent trend information for a single Flow Collector.

If you use multiple Flow Collectors, you will see a differently colored line for each Flow Collector. Below the graph is a legend that uses color-coded blocks for each Flow Collector. If you want to reduce the number of Flow Collectors on the graph, click the name of the Flow Collector you do not want shown in the graph. Its trend line is removed, and the graph rescales or resizes as necessary. Click the Flow Collector name in the legend again to restore its trend line in the graph.

If you move your cursor across the graph, a small information box shows the number of flows per second for each Flow Collector represented in the graph for the date and time (also displayed in the graph) where the cursor is pointed. If you have suppressed a line on the graph by clicking a Flow Collector's name in the legend, that Flow Collector's flow data is not displayed in the information box.

Top Applications

This component displays the top 20 applications used for inbound and outbound traffic for a domain. The percentages refer to the flows for the specific application out of the total inbound or outbound traffic. Since each half of the chart may represent between one and twenty applications, you should not expect the percentages to total 100 percent. You can use the data to see the distribution of applications for inbound and outbound traffic. When you examine the data, you need to evaluate whether the type of application used is valid for the direction of traffic and whether its percentage of use is reasonable for your system.

What you can do next

Investigate hosts

Use this workflow to further investigate a host that is contributing to an alarm. Follow the process listed below to view information about hosts on your network.

1. View [list of hosts](#) on your network.
2. Select a [specific host](#) for which you want to view details.

You also have the option of using the search box to search for an IP address or a range of IP addresses.

Investigate users

Use this workflow to further investigate a user who is contributing to an alarm. Follow the process listed below.

1. View a [list of users](#) on your network.
2. Select a [specific user](#) for whom you want to view details.

You also have the option of using the search box to search for one or more user names. If you want to search for users, use the search box drop-down arrow to select User.



- The name cannot contain any of these characters: | + = ? " < > () : ;
- The search box is not case sensitive.
- You can type partial name(s) with wild cards. The wild card can be in any position, such as srh*, *doe. You can use more than one wild card in each name.

Investigate security events

Use this workflow to investigate the types and numbers of security events contributing to an alarm on a particular host. Follow the process listed below.

1. View a [list of hosts](#) on your network.
2. Select a [specific host](#) for which you want to view details.
3. Select a specific alarm to view the [security events](#) that are triggering for that host.

Investigate alarms

Use this workflow to investigate the types and numbers of alarms.

In the Alarms by Type or Current Alarms component, click an alarm type to display the Alarm List filtered for that alarm type.

Related Topics

[Navigating the Manager](#)

[Analytics: Alert Prerequisite Chart](#)

[Choose a Theme](#)

Analytics: Alerts Summary

The Alerts Summary displays the alerts generated by the system according to your filter settings. It generates these alerts, representing potential malicious activity, based on an analysis of various information about your network, including the following:

- Monitored entities roles, and the observations logged for those entities
- Alert type priority
- IP scanner rules



Use Host Group Management to add known valid IP scanners to the default Network Scanners group (Host Group ID: 48). These devices should be listed by single comma separated IP or CIDR notation. By adding known scanners, you will prevent alerts from occurring when these devices scan external or internal hosts.

Open Alerts Summary

From the main menu, choose **Monitor > Alerts**.

OR

In the Supporting Observations section on the Alert Details page, from the drop-down list for the desired device, choose **Alerts**.

The Alerts Summary opens with a list of all the alerts related to that device.



For more information on how to use this page, see [Analytics: Investigating Alerts](#)

Filter the Alerts Table

Refer to the following information to learn how to filter for various settings.

Search - In the Search field, type the entry by which you want to filter the table. You can filter by any content that would potentially be displayed in any of the table columns, such as alert type, source type, time, etc. When finished, click **Apply** to the right of the Time Range fields.

Time Range - To filter by date and time:

- Enter the start date and time in the **Start** text box or click the calendar icon at the corner of the text box to select the start date and time.
- Enter the end date and time in the **End** text box or click the calendar icon at the

corner of the text box to select the end date and time.

- Click **Apply** . To revert, click **Reset**.

Set alerts by status – By clicking the desired option in this field, you can choose to view All alerts or only Open, Snoozed, Unpublished, or Closed alerts. By default, the table displays alerts that are assigned the All status.



When you open, re-open, or close an alert that matches a Response Management rule, all associated actions occur (email, syslog, webhook), informing you that the alert has been opened, re-opened, or closed, depending on the action you took.

- Snoozed alerts: When you close an alert (see [Edit Alerts Table Entries](#)), the Close Alert dialog opens. In this dialog you can specify if you want to snooze the alert for a certain period of time.
- Unpublished alerts: Unpublished alerts are newer alerts that are still being evaluated for their optimal level of efficacy. They are Off by default. Unpublished alerts work just as other alerts do (for instance, you can snooze and close them). Additionally, changing the priority of the alert does not affect whether or not an alert is published or unpublished. However, they will behave inaccurately in single node deployments.

To know which alerts are unpublished, see [Analytics: Alert Descriptions](#).

View the Alerts Table

The following information is displayed in the Alerts table.

Field	Description
Alert	The alert type that was generated. • To access the Alert Details page for an alert, click the alert. For information about the Alert Details page, see Analytics: Alert Details. • To see the status of an alert, hover over the  (Information) icon that is beside the desired alert. Use the Filter drop-down list to filter by a specific alert type.
Source	The source entity that caused the alert to be generated. To view details for an alert, hover over the  15+ (Alert) icon or

	the  (Eye) icon that is beside the IP address. (If the source is associated with a minimum of one open alert, the Alert icon is displayed. If the source is not associated with an open alert, the Eye icon is displayed.) Click the drop-down list to access the following options: - Choose Device (Network) to access the Device Report filtered by the source. For information about the Device page, see Analytics: Device Report. - Choose Alerts to access the Alerts Summary that contains a list of all the alerts related to that source. For information about the Alerts Summary, see Analytics: Alerts Summary. - Choose Observations to access the Observations by Device page that contains a list of all observations related to the source. For information about the Observations by Device page, see Analytics: Observations by Device. - Choose Flow Analysis to access the Flow Search Results page that contains flow information related to the source. For information about the Flow Search Results page, see Flow Search Results: Overview.
Time	The time the alert was last updated.
Tags	A tag is any type of label or text that you add to an alert so that it can be filtered by that label or text. Use the Tags drop-down list to filter by a specific tag.

Edit Alerts Table Entries

You can assign tags to specific alerts as well as change the status of alerts by using the drop-down lists in the Actions field (located immediately above the Alert table).

1. In the Alerts table, check the checkbox next to each alert for which you want to make one or more changes. Note that if you check multiple alerts, the changes you make will be uniformly applied to all of the alerts you selected.
2. In the Actions field, make your selection(s) from one of the following drop-down lists, one at a time:
 - Assign Tag
 - Change Status

You can change the status to Open, Close, or Unsnnooze. When you close an alert, the Close Alert dialog opens. In this dialog you can specify if you want to snooze the alert for a certain period of time. If you snooze an alert and then decide you want to reopen it, choose the Unsnnooze option for that alert.

To choose more than one option at a time from a particular drop-down list, press the **Shift key** while clicking each option. To de-select a choice, click it again.

3. To save your options, click outside the drop-down list on any white space on the page.

The applicable entries in the table change to reflect your selections.

View Related Configuration Pages

Click the  (**Related Config Links**) icon and then click the appropriate option to access either the Alert Priorities Configuration page or the Alerts Country Watchlist Configuration page.

Download a .CSV File

To download a .csv file containing either all available observations or just the current filtered view (if you have already filtered the table), click the  (**Download CSV**) icon located in the top right corner of the page and click the appropriate option.



You can download a maximum of 65,000 alert records from the Alerts Summary.

Related Tasks

[Analytics: Investigating Alerts](#)

[Analytics: Alert Details](#)

Related Topics

[Analytics: Alerts Workflow](#)

[Analytics: Alerts FAQs](#)

[Analytics: Device Report](#)

[Analytics: Overview](#)

Analytics: Observations Dashboard



- For information about the Observation Types page, see [Analytics: Observation Types](#).
- For information about the Observations by Device page, see [Analytics: Observations by Device](#).
- For information about the Selected Observations page, see [Analytics: Selected Observations](#).

Open Observations Dashboard

From the main menu, choose **Monitor > Observations**.

The Observations Dashboard opens.

Observations Highlights Overview

Observations are facts about an entity's behavior on the network, such as a heartbeat connection with an external IP address, an interaction with an entity on a watchlist, or a remote access session established with another entity. Observations on their own do not carry meaning beyond the fact of what they represent. A typical customer may have many thousands of observations and a few alerts.

Alerts are generated from combinations of observations; observations in isolation do not necessarily constitute malicious behavior. The recent observations, by themselves, do not necessarily mean there is malicious behavior on your network. Review your alerts for a better picture of the potential malicious behavior.

As the system inspects your traffic, it logs observations, or facts, about the entities on your network. Because observations are logged per entity, your network may generate more observations than can be reasonably reviewed. The system presents a subset of the most notable observations logged for your network. You can review and filter these to gain a better understanding of the types of behavior that may result in alerts being generated.



A high number of observations does not necessarily correspond with a high number of alerts, or even any alerts. For example, an entity with many observations may be passing a large amount of varied traffic, but Dynamic Entity Modeling may have identified this behavior as normal and expected for this entity. Similarly, a low number of observations does not necessarily correspond with a lack of alerts. For example, if the only activity detected for an entity is a continual attempt to log into a server with improper credentials, this may



generate a relatively small number of observations, and an alert noting the multiple failed login attempts.

View Observations Highlights

You can drill down and view all observations of that type. If you drill down, the system opens a new tab on this page that displays those observations. If you select a different observation type to drill down into, the system updates that new tab with these observations.

- To view all observations of a specific type, click the  (**Arrow Right**) icon next to the applicable observation type.

The Selected Observations page opens, which by default displays a table that contains all observations of the type you selected that have occurred within the last 24 hours.

For more information about the Selected Observations page, see [Analytics: Selected Observations](#).

- To download a .csv file containing all observations for a specific type, click the  (**Download CSV**) icon located at the top right corner of the associated table.

View Descriptions for the Observations Highlights tables



The fields displayed in these tables vary depending on the observation each table is associated with.

Option Name	Description
Affected Resource	The affected resource.
Affected Resource Type	The affected resource type.
Anomaly	The type of anomaly detected.
Bytes In	Amount of traffic (in bytes) that has been received by the device for a specific point in time.
Bytes Out	Amount of traffic (in bytes) that has been sent from the

Option Name	Description
	device for a specific point in time.
CIDR Range	The approximate CIDR notation scanning range (the actual range may be smaller).
Connected Device	The device to which the endpoint or source established a connection.
Connected IP	The IP with which this source communicated.
Connected Ports	The port with which this port communicated.
Corresponding Ports	The ports that were used in the communication.
Data Sink IP	The IP address of the external device to which data was uploaded.
Data Sink Profile	The Local profile of the observation source (this device) when uploading to the data sink.
Data Source	The internal device from which the data was downloaded.
Data Source IP	The IP address of the internal device from which the data was downloaded.
Data Source Profile	The local profile of the observation source (this device) when downloading from the data source.
Device	The associated endpoint or source. Click the drop-down list to access the following options: - Choose Device (Network) to access the Device Report filtered by the source. For information about the Device page, see Analytics: Device Report. - Choose Alerts to access the Alerts Summary that contains a list of all the alerts related to that source.

Option Name	Description
	For information about the Alerts Summary, see Analytics: Alerts Summary. - Choose Observations to access the Observations by Device page that contains a list of all observations related to the source. For information about the Observations by Device page, see Analytics: Observations by Device. - Choose Flow Analysis to access the Flow Search Results page that contains flow information related to the source. For information about the Flow Search Results page, see Flow Search Results: Overview.
Domain/URL	The domain/URL based on NGFW connection events or passive DNS.
Download Sec	The length of time required for the download to complete.
Download Speed bps	The download speed in bps.
Download Start	The time the download began for the external data sink.
External IP	The external IP address.
Failed Attempts	The number of times an entity has attempted to establish a connection to a device.
Heartbeat period (Seconds)	The time between heartbeats.
History Length (Days)	The number of days of history that was used to calculate the normal set.
Internal Port Set	The type of port set. For example, "connected internal" is the set of connected ports used for internal connections.

Option Name	Description
Last Active	The time that an observation was last active.
Local Device	The local device involved in the communication.
Local Port	The port over which the endpoint or source connected to the device.
Lookback Days	The number of days of history that was used to calculate the normal set.
Lost Port Sets	Ports that are no longer used on this date.
Matching Watchlists	If the watchlist is domain-based, the matching domain names are listed here.
Metric	The metric for this outlier. For example, an outlier for internal "Bytes In" indicates that the internal network traffic (where there is no internet) to the device has spiked.
New Connections	New connections on this date that weren't in the lookback period.
New Port Set	Ports that were used on this date that weren't used in the lookback period.
New Profile	A new device profile that differs from previous behavior.
Normal Connection Set	The connections found in the lookback period.
Normal Ports Set	The ports that were used in the lookback period.
Number of Heartbeats	The number of times the server was connected during this observed event.
Packets In	The packets received by the source.
Packets Out	The packets sent from the source.

Option Name	Description
Port	The source port used in the observed event.
Port Ranges	The device-scanned ports included in this range, and possibly others. Common targets may include web server targets.
Probability	The probability that you would see this outlier.
Profile	The role(s) associated with the endpoint or source to which the device connected.
Public Facing IP	A public IP address that was discovered on a watchlist.
Remote Device	The device this source communicated with. Click the drop-down list to access the following options: - Choose Device (Network) to access the Device Report filtered by the source. For information about the Device page, see Analytics: Device Report. - Choose Alerts to access the Alerts Summary that contains a list of all the alerts related to that source. For information about the Alerts Summary, see Analytics: Alerts Summary. - Choose Observations to access the Observations by Device page that contains a list of all observations related to the source. For information about the Observations by Device page, see Analytics: Observations by Device. - Choose Flow Analysis to access the Flow Search Results page that contains flow information related to the source. For information about the Flow Search Results page, see Flow Search Results: Overview.

Option Name	Description
Remote IP	The IP address with which this source communicated.
Remote Port	The port with which this source communicated.
Resource	The affected resource.
Retrieve Sessions	Provides a link to execute a supporting query on the Flow Search page. Use it to explore telemetry data related to the components of the observation, such as the observed time range, source and peer communication, observed port or protocol, and so on. Use the Flow Search for detailed session analysis.
Sample Size	The number of historical samples used in this calculation.
Scanned Device	The IP address or host name of the device that was scanned.
Scanned Ports	The device-scanned ports included in this range, and possibly others.
Scanner Device	The IP address or host name of the device that performed the scan.
Severity	The severity of the reported event.
Source	The associated endpoint or source. Click the drop-down list to access the following options: - Choose Device (Network) to access the Device Report filtered by the source. For information about the Device page, see Analytics: Device Report. - Choose Alerts to access the Alerts Summary that contains a list of all the alerts related to that source. For information about the Alerts Summary, see Analytics: Alerts Summary.

Option Name	Description
	- Choose Observations to access the Observations by Device page that contains a list of all observations related to the source. For information about the Observations by Device page, see Analytics: Observations by Device. - Choose Flow Analysis to access the Flow Search Results page that contains flow information related to the source. For information about the Flow Search Results page, see Flow Search Results: Overview.
Source IP	The IP of the source at the time of the observed event.
Source Port	The port of the source at the time of the observed event.
Suspect Connections	New connections to unusual external sources or that are on unusual ports.
Time	The time the observation occurred.
Time Window	The length of time during which the event was shared.
Transfer Size	The number of bytes transferred.
Upload Start	The start time of the upload to the external data sink.
Upload Sec	The duration of the upload.
Upload Speed bps	The upload speed in bps.
Violating Port	The port used by the source.
Violating Protocol	The network protocol used by the source (for example, TCP).
Violating Type	The type of violation.

Related Topics

[Analytics: Observation Types](#)

[Analytics: Observations by Device](#)

[Analytics: Selected Observations](#)

[Analytics: Overview](#)

ISE ANC Policy Assignments Report

Use this page to see a list of the hosts to which the ISE (Identity Services Engine) ANC (Adaptive Network Control) policy was applied for the last 30 days.

The advantage of configuring and applying ANC policies in [Response Management](#) is so that Secure Network Analytics can automatically assign an ANC policy to a host associated with alarms you have designated in a rule. You don't have to manually apply ANC policies from the [ANC Policies panel](#). (Note that you can only apply ANC policies from this panel, not configure them.) For more information, refer to the [ISE Configuration Guide](#).



To configure and apply ANC policies, you must have [configured Cisco® ISE](#) for your Secure Network Analytics system.

How to find this page

From the main menu, choose **Monitor > Integrations > ISE ANC Assignments**.

View the ISE ANC Policy Assignments List

The report lists up to a maximum of 1000 rows, with each row representing a single host. If a policy has been applied to a host more than once in the last 30 days, then you will see only the last occurrence in this report.

Report column definitions

Column	Description
Edit check box	To assign an ANC policy to several hosts at one time, check the check box beside each host to which you want to assign the policy. From the ANC Policy drop-down list, choose the applicable policy. These are policies that have been created in the ISE management interface. • Only policies that are available for ISE clusters where applying an ANC policy is possible are included in this list. • When applying an ANC policy to several hosts at one

Column	Description
	time, the hosts must belong to the same ISE cluster. When you choose multiple hosts that don't all belong to the same ISE cluster, you can only clear an ANC policy (by choosing the None policy).
Host IP	The host to which an ANC policy was applied. Click this entry to access the Host Report and check the history of the associated ANC policy for the host. The history shows each change that someone made to the policy.
ISE Cluster	The ISE cluster to which the host belongs.
MAC Address	The MAC address of the host. The MAC address is the unique identifier of a device.
Assignment Mode	Indicates the mode via which the policy was applied. Automatic Indicates that the policy was applied by a rule configured in Response Management. Manual Indicates that a user applied the policy from the ANC Policies panel (accessed via the Host Report).
Requested By	- If the policy was applied by a rule in Response Management, the entry is <i>Response Management</i>. - If a user applied the policy, the name of the user is displayed.
Time	The date and time the policy was applied.
Requested ANC Policy	The ANC policy that was requested at the time/date that is displayed in the Time column. If you assign a different policy using the context menu in the Assign ANC Policy column, this policy is then displayed in the Requested ANC Policy column.
Effective	Click Retrieve to see the ANC Policy that the host is

Column	Description
ANC Policy	currently assigned to in the ISE management interface.
Assign ANC Policy	To assign a different policy to a single host, click the context menu in this column and choose the ANC policy (these are policies that have been created in the ISE management interface). The policy is then displayed in the Requested ANC Policy column. Only policies that are available for ISE clusters where applying an ANC policy is possible are included in this list.

Related Topic

[Response Management: Overview](#)

[Applying Mitigation Actions to a Host Using Cisco® ISE](#)

[Viewing Mitigation Actions for a Host \(Host History\)](#)

Analyzing Flows

While monitoring or investigating activity on your network, one or more of the following may occur:

- You have determined that a specific host has been compromised. You want to retrieve information about conversations to and from that host, so you can identify the host suspected of causing the compromise.
- You see a high spike in traffic and want to analyze the data and determine the cause of the spike.
- An alarm has been triggered, and you need to determine if it is a threat to your network.

The flow analysis process allows you to make these determinations so you can secure your network:

- Access the Flow Search page and define settings to focus on the information you need.

If you have used the flow table in the Desktop Client, you are accustomed to opening a separate filter with multiple pages. In the Manager, you access the settings, or filters, on the Flow Search page. You can run the search or save it to run later. In addition, you can initiate a flow search for a single host or user from the [Host Report](#) or the [User Entity View](#) page.

- Analyze the information displayed.

Accessing Saved Flow Searches

Use this page to view a list of saved flow searches for the Manager. You can choose a search to edit, delete, run, or clone (i.e., copy).

How to find this page

From the main menu, choose **investigate > Search Management > Saved Searches**.

The Saved Searches page opens.

What you can do on this page

You can perform the following actions on this page.

To...	Click...
Review, run, edit, or save the flow or top report search under a different name	The applicable name in the Search Name column.
Modify an existing flow or top report search that you had previously saved	The ⋮ (Ellipsis) icon > Edit in the Actions column.
Permanently remove a flow or top report search you no longer need	The Ellipsis icon > Delete in the Actions column.
Run an existing flow or top report search that you had previously saved (the search remains listed on this page unless you delete it)	The Ellipsis icon > Run in the Actions column.
Copy an existing flow or top report search	The Ellipsis icon > Clone in the Actions column. <i>The flow search page opens, displaying a search with the same settings as the original search. You can change the settings now or edit them later. You will have the option to review the cloned search and save it using the same name or a different name.</i>

Accessing Saved Flow Search Results

Use this page to open and view results of a previous saved flow search. Each row on the page represents a different set of results.

- You will be able to download flow search results to a .CSV file when 20,000 or more records are returned for a query.

How to find this page

From the main menu, choose **Investigate > Search Management > Saved Results**.

The Saved Results page opens.

What you can do on this page

- View saved flow search results from the last 14 days.
- To permanently remove the results for a specific search, in the Actions column, click **Delete**.

Flow Search: Overview



For information about building a top report search, see [Building a Top Report Search](#).

The flow search page allows you to define a search with the following commonly used settings:

- Time range
- IP address or range
- Port/protocol
- Host groups

The flow search is useful as a way of monitoring the behavior of a host or range of hosts at a high level. If you need more detail or want to perform an in-depth forensic investigation, you can then use the Advanced Options section.

Open Flow Search

Do one of the following to access this page

- Go to the Flow Search page (from the main menu, choose **Investigate > Flow Search**).
- Go to the Saved Searches page (from the main menu, choose **investigate > Search Management > Saved Searches**). In the Actions column for the applicable saved flow search, choose **Edit**.
- Click the **⋮ (Ellipsis)** icon next to the applicable IP address, in the Actions column, or on a point in a graph. Note that if you are pivoting from the Traffic by Peer Host Group graph on the Host Report page or the Top Host Groups by Traffic graph on the Host Group Report page, you must click a host group, a column, or the line between two host groups.

The Flow Search page opens.

Flow Search: Table of Contents

Flow Search: Concepts You Need to Understand

Before running a flow search, refer to the following topics for concepts you should first understand:

- [Flow Search: How a Unique Flow is Defined](#)
- [Flow Search: Query Logic Used Between and Within Sections](#)
- [Flow Search: Search Logic Used in a Flow Search](#)

Related Topics

[Flow Search: Entering Search Criteria Overview](#)

[Flow Search: Run a Flow Search Overview](#)

[Flow Search: Overview](#)

Flow Search: How a Unique Flow is Defined

A Secure Network Analytics flow is an audit trail of connectivity between two endpoints (it is not the same as a TCP socket or firewall session). It is based on source IP, destination IP/port, and timestamp.

Duration of a given flow requires only a single packet with the same client-to-server relationship on the same service port. If communication between the client and server continues to send bytes over the same server port, this communication is combined into the same flow conversation, even if there are multiple TCP sessions within that flow.

If any of the following occurs, Secure Network Analytics ends that flow, and any additional communication are included in a new flow.


- Communication that previously made up the flow between two endpoints stops for over five minutes
- If the source IP, destination IP, or destination port changes for a specific communication occurrence

Example

Communication is occurring between client IP 1.1.1.1 and client IP 2.2.2.2 over the following ports/protocols:

- 80/tcp, which is one flow
- 443/tcp, which is a separate flow

If the two endpoints stop communicating over port/protocol 80/tcp for five minutes or longer, that flow is over. However, the communication occurring over port/protocol 443/tcp may still be ongoing.



Related Topics

[Flow Search: Concepts You Need to Understand](#)

[Flow Search: Entering Search Criteria Overview](#)

[Flow Search: Run a Flow Search Overview](#)

[Flow Search: Overview](#)

Flow Search: Query Logic Used Between and Within Sections

The flow search builder relies on Boolean logic. Refer to the following for an explanation.

Sections

- The flow search builder contains three sections, Subject, Connection, and Peer. These sections are ANDed together, so a search must return results that match the search criteria in all of them. The Connection section returns results matching the server OR client with the given connection search criteria.
- The orientation field specifies if the Subject sections are for client, server, or either. For example, if you select server, the search returns results that match the Subject section as the server and the Peer section as the client. If you select either, the search returns results matching the Subject section as the client and Peer as the server UNIONed with Subject section as the server and Peer as the client.

Attributes

- Each section contains attributes, such as Host, User, Devices, and so on. Each of these attributes functions as a set. If you enter search criteria in multiple sets, the sets are ANDed together within the section.
- Within an attribute (Host, User, etc.), you can add multiple sets of search criteria. By default, these are designated as "includes" and are ORed together. The search returns results that match one, some, or all of these sets. For example, if you want to search for three IP addresses, you would simply add three IP addresses within that attribute. The search returns results matching any of the three IP addresses.
- Within an attribute, you can also specify that one or more sets of search criteria be excluded. These individual "excludes" are ORed together (just as the individual "includes" are ORed together). However, each of these "exclude" sets functions as an AND NOT with the other "include" sets within that one attribute. Note that If a user is excluded from the Subject in a query, it is automatically excluded from the Peer, and vice versa.

Example:

You want to run a search on one IP address in the Subject section and two IP addresses in the Peer section.

1. In the Subject section, in the Host IP Address or Range attribute, enter the IP address you want included in the text field.
2. In the Peer section, in the Host IP Address or Range attribute, enter the two IP addresses in the text field.

If the orientation is...	Then the search returns results that match the following...
Server	The Subject IP address is sending flows to the first OR the second Peer IP address.
Client	The first Peer IP address OR the second Peer IP address is sending flows to the Subject IP address.
Either	The Subject IP address is sending flows to the first OR the second Peer IP address AND The first Peer IP address OR the second Peer IP address is sending flows to the Subject IP address.

Related Topics

[Flow Search: Concepts You Need to Understand](#)

[Flow Search: Entering Search Criteria Overview](#)

[Flow Search: Run a Flow Search Overview](#)

[Flow Search: Overview](#)

Flow Search: Search Logic Used in a Flow Search

- If you include certain criteria, marked by an  (**Information**) icon on the page, in your flow search or run a flow search with a Max Records Returned of 20,000 or more (whose results can only be viewed by downloading them to a CSV file), you cannot cancel the flow search before it is complete or see what percentage of the job has run thus far. However, when the job has finished, a blue *Completed* pill is displayed on the Flow Search Results page.
- When a job is in progress, a banner at the top of the page displays the message *Your query is in progress*. A blue *In Progress* pill is also displayed. While the job is running, the message *In Progress* is displayed on the Jobs > [Job Management](#) page in the Current Jobs table. When the job has finished, a blue *Completed* pill is displayed on the Flow Search Results page, and on the Job Management page, the job is displayed in the Finished Jobs table with a status of *Completed* (only if there were results).
- Note that the search uses 2,000 flow records as the default. When that number of flow records has been returned, the search may stop, no matter how much time has elapsed relative to the time interval you specified.
- For a flow search request with a Max Records Returned of 10,000 or less, both the search and the results are retained on the Job Management page for 48 hours from the time the job completed.
- For a flow search request with a Max Records Returned of 20,000 or more (whose results can only be viewed by downloading them to a CSV file), both the search and the results are retained on the Job Management page for 7 days from the time the job completed.
- You can perform a flow search for up to 400,000 flows in one search.
- Flow searches with a Max Records Returned of 10,000 or less, as well as all top report jobs, run in a different queue from that of flow searches with a Max Records Returned of 20,000 or more.
- You can run a maximum of four jobs at the same time if the Max Records Returned for the flow search is 10,000 or less. It does not matter what the Records Returned value is for a top report job.
- You can run only one flow search at a time if the Max Records Returned is 20,000 or more.

- If you start additional top report jobs or flow searches during this time (for either of these two categories of flow searches), they are placed in their respective queue with a Pending status.
- You can cancel a flow search with a Pending status.

For more information about flow search guidelines, see the "Manage Flow Search Jobs" section in the [Managing Jobs](#) topic.

Related Topics

[Flow Search: Concepts You Need to Understand](#)

[Flow Search: Entering Search Criteria Overview](#)

[Flow Search: Run a Flow Search Overview](#)

[Flow Search: Overview](#)

Flow Search: IPv4 and IPv6 Address Notation

You can use either IPv4 or IPv6 addresses for hosts or host groups. The notation forms that must be used for each type are listed in the applicable table. You must separate the IP addresses by a comma, comma space, space, or return.

The Manager displays IPv6 addresses and IPv4 addresses simultaneously. Of the advantages provided by IPv6, the fact that it supports a larger address space is probably the most well known. IPv6 supports 128-bit addresses, whereas IPv4 supports only 32-bit addresses.

The Manager continues to support IPv4. The Flow Collectors can monitor IPv6 traffic, but they are not IPv6 addressable at this time.

Examples

- An IPv6 address consists of eight blocks of hexadecimal digits separated by colons, as shown in this example (where the x's represent hexadecimal digits):

Example: xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx

- You can replace blocks of 4-digit zeros with double colons (e.g., ::), but only once per IPv6 address. (Leading zeros can also be omitted.) Consider the following address (where the x's represent hexadecimal digits):

Example: xxxx:xxxx:0000:0000:0000:0000:0000:xxxx

This address can be rewritten like this:

Example: xxxx:xxxx::xxxx.

- Like an IPv4 address, an IPv6 address may include a prefix length that specifies how many left-most bits of the address make up the prefix. In IPv4 addresses, the prefix is used to determine the subnet mask (a.k.a., CIDR notation). In IPv6 addresses, it identifies the subnet that an interface belongs to, and routers use it for forwarding.

Example: xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx/40

- Take care when compressing an IPv6 address that contains a prefix, because if not done correctly, it can result in an incorrect interpretation when the address is uncompressed. You can write an IPv4 address as an IPv6 address, like this:

Example: xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:192.168.0.2

- If the first 6 blocks are zeros, you can compress the address like this:

Example: ::192.168.0.2

Related Topics

[Flow Search: IPv4 Address Notation](#)

[Flow Search: IPv6 Address Notation](#)

[Flow Search: CIDR Notation](#)

[Flow Search: Overview](#)

Flow Search: IPv4 Address Notation

If you enter this type of IP address...	You must use this formatting...
Single IP address	10.52.1.55 Includes only the host at 10.52.1.55.
Trailing dot subnet	10.52. Includes any IP address from 10.52.0.0 to 10.52.255.255. Note: You must include the trailing period (.) at the end.
Classless Inter-Domain Routing (CIDR) notation (i.e., using a “/” to indicate the masked bits)	10.52.1.0/24 Includes any IP address from 10.52.1.0 to 10.52.1.255. Note: Use this notation if you prefer to enter the IP address using a “mask.”
Range of networks	10-11. Includes any IP address from 10.0.0.0 to 11.255.255.255 10.52-53. Includes any IP address from 10.52.0.0 to 10.52.255.255, and from 10.53.0.0 to 10.53.255.255 10.52-55.3. Includes any IP address from 10.52.3.0 to 10.52.3.255; from 10.53.3.0 to 10.53.3.255; from 10.54.3.0 to 10.54.3.255; and from 10.55.3.0 to 10.55.3.255 Note: Be sure to include the trailing period (.) at the end.
Range of hosts	10.52.1.0-10 Includes any IP address from 10.52.1.0 to 10.52.1.10 10.52.0.0-10.52.255.255 Includes any IP address from 10.52.0.0 to 10.52.255.255

If you enter this type of IP address...	You must use this formatting...
Comma-separated list	10.52.1.10,10.52.1.50,10.100.1.20 Includes only these three hosts.

Related Topics

[Flow Search: IPv4 and IPv6 Address Notation](#)

[Flow Search: IPv6 Address Notation](#)

[Flow Search: CIDR Notation](#)

[Flow Search: Overview](#)

Flow Search: IPv6 Address Notation

If you enter this type of IP address...	You must use this formatting...
Single IP address	2001:0DB8:0000:0056:0000:ABCD:EF12:3456 2001:DB8:0:56:0:ABCD:EF12:3456 2001:DB8::56:0:ABCD:EF12:3456 2001:DB80:0:56::ABCD:EF12:3456 2001:DB80:0:56::ABCD:239.18.52.86
Global Routing Prefix subnet	2001:DB8:0:56::/64
Range of hosts	2001:DB8:0:56:ABCD:EF12:3456:1-10 2001:DB8:0:56:ABCD:EF12:3456:1 - 2001:DB8:0:56:ABCD:EF12:3456:10

Related Topics

[Flow Search: IPv4 and IPv6 Address Notation](#)

[Flow Search: IPv4 Address Notation](#)

[Flow Search: CIDR Notation](#)

[Flow Search: Overview](#)

Flow Search: CIDR Notation

Use the following table to understand classless inter-domain routing (CIDR) notation:

Class	CIDR Length	Mask	Number of Addresses
A	/1	128.0.0.0	2,147,483,648
	/2	192.0.0.0	1,073,741,824
	/3	224.0.0.0	536,870,912
	/4	240.0.0.0	268,435,456
	/5	248.0.0.0	134,217,728
	/6	252.0.0.0	67,108,864
	/7	254.0.0.0	33,554,432
	/8	255.0.0.0	16,777,216
	/9	255.128.0.0	8,388,608
	/10	255.192	4,194,

Cl as s	CI DR Le ngt h	Mask	Numbe r of Addres ses
		.0.0	304
	/11	255.224 .0.0	2,097, 152
	/12	255.240 .0.0	1,048, 576
	/13	255.248 .0.0	524,28 8
	/14	255.252 .0.0	262,14 4
	/15	255.254 .0.0	131,07 2
B	/16	255.255 .0.0	65,536
	/17	255.255 .128.0	32,768
	/18	255.255 .192.0	16,384
	/19	255.255 .224.0	8,192
	/20	255.255 .240.0	4,096
	/21	255.255 .248.0	2,048

Cl as s	CI DR Le ngt h	Mask	Numbe r of Addres ses
	/22	255.255 .252.0	1,024
	/23	255.255 .254.0	512
C	/24	255.255 .255.0	256
	/25	255.255 .255.128	128
	/26	255.255 .255.192	64
	/27	255.255 .255.224	32
	/28	255.255 .255.240	16
	/29	255.255 .255.248	8
	/30	255.255 .255.252	4
	/31	255.255 .255.254	2

Related Topics

[Flow Search: IPv4 and IPv6 Address Notation](#)

[Flow Search: IPv4 Address Notation](#)

[Flow Search: IPv6 Address Notation](#)

Flow Search: Overview

Flow Search: Entering Search Criteria Overview

At the top of this page you can enter general criteria that applies to the subject, connection, and peer, which includes search type, time range, search name, and max records returned.

Below this you can enter the IP address for the subject, peer, or both, as well as the port/protocol for the connection.

The Advanced Options section allows you to define a search using more specific settings than those in the main section. You might, for example, use the advanced section as part of an in-depth forensic investigation of suspicious behavior involving a specific host or for a specific time interval. If you are conducting a forensic investigation, you should use this section to fine tune settings so the search returns the data that is most relevant. You may need to set up a series of sequential searches to make sure the results cover the time span you need.

Related Tasks

[Flow Search: Guidelines For Entering Search Criteria](#)

[Flow Search: General Search Criteria](#)

[Flow Search: Subject Search Criteria](#)

[Flow Search: Connection Search Criteria](#)

[Flow Search: Peer Search Criteria](#)

Related Topics

[Flow Search: Concepts You Need to Understand](#)

[Flow Search: Run a Flow Search Overview](#)

[Flow Search: Overview](#)

Flow Search: Guidelines For Entering Search Criteria

The search criteria, sorted by Subject, Connection, and Peer, are displayed at the top of the page and within the associated section on the page in a green text box.



To restore the fields on this page to the default search values (last five minutes, 2,000 max records), click **Restore Defaults**.

Alpha and Alphanumeric Search Criteria

Use these guidelines when entering alpha or alphanumeric search criteria on this page.

- Use a space or press **Enter** to advance to another line to separate multiple entries in a field. You cannot use a space as a delimiter in the Payload field because many of the valid values you can enter for this field contain a space.
- Type an exclamation point (!) before a search criteria to exclude that item from the search. The excluded search criteria are displayed in a black text box at the top of the page.
- Click **X** on a text box to remove a filter from your search.
- If you included certain criteria, marked by an **Information** icon on the page, in your flow search or run a flow search with a Max Records Returned of 20,000 or more (whose results can only be viewed by downloading them to a CSV file), you cannot cancel a flow search before it is complete or see what percentage of the job has run thus far. However, when the job has finished, a blue *Completed* pill is displayed on the Flow Search Results page.
- If you are using a MAC address to define a device, use the full address. Do not enter a partial address.
- To use the Cisco TrustSec® tags, you must have previously created Security Group Tags for your Cisco devices in your network. Use these guidelines when entering alpha or alphanumeric search criteria on this page.

Numeric Search Criteria

Use these guidelines when entering numeric search criteria on this page:

- In fields that require numerical values, such as Total Bytes, use the following symbols. If you try to enter a numeric value without one of these symbols, you see an error message.

- Greater than: >
- Greater than or equal to: >=
- Less than: <
- Less than or equal to: <=
- Range: -
- For numeric entries, if you want to search for results using two entries where the results overlap each other to some degree, search on a range of values instead of using the "less than" symbol (<) and the greater than symbol (>).

Example: If you want to search for flows that contain values greater than 300M and less than 600M for subject bytes, enter 300M-600M in the Subject: Bytes field. Do not enter >300M <600M, as the results include flows that are greater than 300M EVEN IF they are greater than 600M, and flows that are less than 600M EVEN IF they are less than 300M (therefore every flow would be a match). So in this scenario, it would be as if you did not even enter additional filter criteria, because the flow search would include the same results as the original search.

- You can enter either the number of units (e.g., >40) or the number of units followed by the applicable abbreviation (e.g., >40M).
- You can enter the number of units (e.g., 1) with the unit of measurement immediately after it (no space).

Example: 10M, not 10 M

- In fields that use bytes (not bps, which is bits per second) as the unit of measure, 1K = 1024
- When you enter a number over 1023, Secure Network Analytics changes that number to the nearest unit after you have clicked away from that field.

Example: 1024 changes to 1K; 4,000,000,000 changes to 3.73G

- In fields that do not use bytes as the unit of measure, 1K = 1000
- When you enter a number over 999, Secure Network Analytics changes that number to the nearest unit after you have clicked away from that field.

Example: 1000 changes to 1K; 1024 changes to 1.02K; 2,000,000,000 = 2 G

- Do not use these special characters in any field in the query builder: () % ' "

Related Tasks

[Flow Search: General Search Criteria](#)

[Flow Search: Subject Search Criteria](#)

[Flow Search: Connection Search Criteria](#)

[Flow Search: Peer Search Criteria](#)

Related Topics

[Flow Search: Entering Search Criteria Overview](#)

[Flow Search: Concepts You Need to Understand](#)

[Flow Search: Run a Flow Search Overview](#)

[Flow Search: Overview](#)

Flow Search: General Search Criteria

Search Type

Select **Flow** to run a flow search.

Time Range

Use the Time Range drop-down list to choose your desired setting for the time and date you want your flow search to cover. Click in the field to see a drop-down list of available choices.

If you want to define a custom time range, you can do one of the following:

- Paste or type the entry into the Time Range field using this format: MM/DD/YYYY HH:MM AM - MM/DD/YYYY HH:MM PM.
- Choose **Custom** from the drop-down list. Two calendars open. Today's date will be highlighted in blue in whichever calendar contains this date.
 1. First, click the date you want to use as the start date (if you want to use today's date as the start date, you must still click it even though it is highlighted). To move to another month in either calendar, use the arrow.
 2. After you have clicked the start date, find the desired end date (which must be a later date than the start date) and click it. The days between those two dates will be highlighted in light blue.
 - If you need to change your selection after you have selected both dates, you will need to start over. (Each time you click, it toggles between selecting a start date and an end date.)
 3. Choose the desired times and click **Apply** to accept. Time is displayed according to your browser's time zone.

Search Name

A default search name is pre-populated in the Search Name field; however, you can edit this entry.

Max Records Returned

You can specify the number of results you want returned from the drop-down list. The default number is 2,000. If you select 10,000 or lower, the results automatically load to the Flow Search Results page. If you want to export them to a comma-separated values (CSV) file, click **Export**.

If you select a number that is over 10,000, the results are created in a downloadable CSV file. To download the results, do one of the following:

- On the Job Management page, in the row of the applicable job, click the **Ellipsis** icon in the Actions column and select **Download CSV**.
- After the flow search has finished and the Flow Search Results page opens, click **Download CSV File** (located at the top of the page).

Related Tasks

[Flow Search: General Search Criteria](#)

[Flow Search: Subject Search Criteria](#)

[Flow Search: Connection Search Criteria](#)

[Flow Search: Peer Search Criteria](#)

Related Topics

[Flow Search: Entering Search Criteria Overview](#)

[Flow Search: Concepts You Need to Understand](#)

[Flow Search: Run a Flow Search Overview](#)

[Flow Search: Overview](#)

Flow Search: Subject Search Criteria

Use this section to define characteristics of the subject.

MAIN OPTIONS

Host IP Address or Range

Note the following guidelines when defining IP addresses:

- Valid formats include IPv4, IPv6, or CIDR notation (192.168.10/16).
- You can enter an IP address or a range (192.168.10-100).
- You can use a wild card at the end of a partial IP address (e.g., 10.0.3*, 10.0.35.*).

Host Groups

Note the following guidelines when selecting host groups:

- Click **Select** under Host Groups to open the Host Group Selector panel.
- Use the include and exclude tabs to choose hosts to include or exclude in your search. The selected hosts are displayed at the top of the panel.
- You can use the Search box in the Host Group Selector panel to quickly search for a host group.
- Click **Apply** to include these hosts in your flow search. Click **Cancel** to discard your selections.

ADVANCED OPTIONS

Bytes

The total number of estimated payload bytes sent by the subject. Ethernet header bytes are subtracted in addition to bytes for MPLS and VLAN tags if they are present in the Netflow/IPFIX record for each packet.

Port/Protocol

When you know that a particular IP address has been involved in a conversation over a given port, but you do not know if this host was the subject or peer, complete the Port/Protocol field in the Connection section. In this situation, Secure Network Analytics ORs together the criteria in these two fields and returns flows when the port/protocol you entered was used by either a subject or peer that matches the IP address you entered (or one of the IP addresses if you entered a range). The table that displays these results indicate if it was the subject or the peer who used the port/protocol you entered.

TrustSec ID

In order to search for TrustSec information or to create a custom event using it, you must have already implemented Cisco TrustSec and have created Security Tag Groups (SGT) on your network.

Orientation

Specify if you want to return results for instances where the subject is the client, server, or either (the client is considered to be the host that initiated the session).

Related Tasks

[Flow Search: Subject Search Criteria](#)

[Flow Search: General Search Criteria](#)

[Flow Search: Connection Search Criteria](#)

[Flow Search: Peer Search Criteria](#)

Related Topics

[Flow Search: Entering Search Criteria Overview](#)

[Flow Search: Concepts You Need to Understand](#)

[Flow Search: Run a Flow Search Overview](#)

[Flow Search: Overview](#)

Flow Search: Connection Search Criteria

Use this section to define settings for the connection between the subject and peer.

MAIN OPTIONS

Applications

Note the following guidelines when selecting applications:

- Click **Select** under Applications to open the Application Selector panel.
- Use the include and exclude tabs to choose applications to include or exclude in your search. The selected applications are displayed at the top of the panel.
- You can use the Search box in the Application Selector panel to quickly search for an application.
- The Applications list includes both custom and default applications.
- Click **Apply** to include these applications in your flow search. Click **Cancel** to discard your selections.

Port/Protocol

When you know that a particular IP address has been involved in a conversation over a given port, but you do not know if this host was the subject or peer, complete the Port/Protocol field in the Connection section. In this situation, Secure Network Analytics ORs together the criteria in these two fields and returns flows when the port/protocol you entered was used by either a subject or peer that matches the IP address you entered (or one of the IP addresses if you entered a range). The table that displays these results indicate if it was the subject or the peer who used the port/protocol you entered.

ADVANCED OPTIONS

Flow Direction

Select the direction of traffic for which you want the search to return results: All, Bidirectional, or Unidirectional. Traffic is defined as one or more packets that traverse an exporter's observation domain in the network. A packet consists of header bytes plus 0 or more payload bytes. Packets have a direction associated with them based on the source and destination IP address fields in the header bytes. Exporters observing traffic sense the direction and many provide the directional information in their exports. The Flow Collector creates unidirectional flows when it collects traffic exports in one direction and bidirectional flows when it collects traffic exports in the reverse direction.

Total Bytes

The total estimated payload bytes transmitted by both the subject and peer. Ethernet header bytes are subtracted in addition to bytes for MPLS and VLAN tags if they are present in the Netflow/IPFIX record for each packet.

Payload

You cannot use a space as a delimiter because many of the valid values you can enter for this field contain a space. You can use a comma or press **Enter** to advance to another line to separate entries.

Flow Collector Name

Select the Flow Collector(s) from which you want the data collected.

Protocols

Note the following when choosing protocols:

- Click either the **Include** or **Exclude** radio button, depending on whether you want to include or exclude protocols. You can choose only to either include or exclude protocols; you cannot do both.
- Click inside the **Protocols** field to see the list of protocols from which you can choose. In the Protocols field, included protocols will be displayed in green and excluded protocols will be displayed in gray.
- To delete an entry, click the **X** beside the applicable protocol.

Include Interface Data

If you want to view interface data related to the search parameters returned for a flow search, you must do the following:

On the Flow Search page

- Select the Include Interface Data check box.
- If you do not select the Filter By Flow Action check box, all flows are included in the flow search. See the Filter By Flow Action entry below for more information.

On the Flow Search Results page

- Ensure that the Subject Interfaces and Peer Interfaces columns are displayed on this page. If they are not, click **Manage Column** and select the Subject Interfaces and Peer Interfaces check boxes and then click **Set** to accept your selections.
- Click the arrow in the row containing the flow result you want to view. The row expands to display a General tab and an Interfaces tab. Click the Interfaces tab to view the interface details.

Filter By Flow Action

You must first select the Include Interface Data check box before you can select this check box. If you select *Permitted and Denied*, the event queries for any individual flow that is tagged as both *Permitted* and *Denied*. In the case of one flow traversing multiple routers, your network configuration may permit the flow to cross one router and deny it from crossing another.

Exporters & Interfaces

Click **Select** to open the Exporters & Interfaces panel. Select the exporters and/or interfaces by which you want to filter the flow search. When you select an exporter, all the interfaces under that exporter are automatically selected. However, if you expand the exporter, you see that the interface check boxes are not actually checked. Therefore, if you want to select only some interfaces under an exporter, clear the exporter check box and then select the desired interfaces.

Encryption

These fields contain "--" on the Flow Search Results page if you do not have a Cisco encrypted traffic analytics-enabled switch and router. Certain encryption fields may also contain "--" if the flow record does not have information for that particular field. For more information about Secure Network Analytics and encrypted traffic analytics, see the [encrypted traffic analytics white paper](#) and the [encrypted traffic analytics deployment guides](#).

Note the following guidelines when selecting encryption fields:

- Click **Select** to open the Encryption panel.
- To view the list of available search criteria, click a field to open the drop-down box.
- You can type in a field to quickly search for a specific option.
- As you select encryption fields, the available options in each drop-down box is filtered down to possible values to get accurate results from your search. For example, if you select Diffie-Hellman (DH) in the Encryption Key Exchange field, the Encryption Authentication Algorithm field limits your options to RSA, Anon, or DSS.
- Click **Done** to include these encryption fields in your flow search. Click **Cancel** to discard your selections.

Related Tasks

[Flow Search: Connection Search Criteria](#)

[Flow Search: General Search Criteria](#)

[Flow Search: Subject Search Criteria](#)

[Flow Search: Peer Search Criteria](#)

Related Topics

[Flow Search: Entering Search Criteria Overview](#)

[Flow Search: Concepts You Need to Understand](#)

[Flow Search: Run a Flow Search Overview](#)

[Flow Search: Overview](#)

Flow Search: Peer Search Criteria

Use this section to define characteristics of the peer (i.e., the other entity) involved in the flow.

MAIN OPTIONS

Host IP Address or Range

Note the following guidelines when defining IP addresses:

- Valid formats include IPv4, IPv6, or CIDR notation (192.168.10/16).
- You can enter an IP address or a range (192.168.10-100).

Host Groups

Note the following guidelines when selecting host groups:

- Click **Select** under Host Groups to open the Host Group Selector panel.
- Use the include and exclude tabs to choose hosts to include or exclude in your search. The selected hosts are displayed at the top of the panel.
- You can use the Search box in the Host Group Selector panel to quickly search for a host group.
- Click **Apply** to include these hosts in your flow search. Click **Cancel** to discard your selections.

ADVANCED OPTIONS

Bytes

The total number of estimated payload bytes sent by the peer. Ethernet header bytes are subtracted in addition to bytes for MPLS and VLAN tags if they are present in the Netflow/IPFIX record for each packet.

Port/Protocol

When you know that a particular IP address has been involved in a conversation over a given port, but you do not know if this host was the subject or peer, complete the Port/Protocol field in the Connection section. In this situation, Secure Network Analytics ORs together the criteria in these two fields and returns flows when the port/protocol you entered was used by either a subject or peer that matches the IP address you entered (or one of the IP addresses if you entered a range). The table that displays these results indicate if it was the subject or the peer who used the port/protocol you entered.

TrustSec ID

In order to search for TrustSec information or to create a custom event using it, you must have already implemented Cisco TrustSec and have created Security Tag Groups (SGT) on your network.

Related Tasks

[Flow Search: Peer Search Criteria](#)

[Flow Search: General Search Criteria](#)

[Flow Search: Subject Search Criteria](#)

[Flow Search: Connection Search Criteria](#)

Related Topics

[Flow Search: Entering Search Criteria Overview](#)

[Flow Search: Concepts You Need to Understand](#)

[Flow Search: Run a Flow Search Overview](#)

[Flow Search: Overview](#)

Flow Search: Run a Flow Search Overview

There are multiple locations from which to run a flow search other than from the Flow Search page. Additionally, depending on the location from which you run a flow search, the flow search filters on different data.

For more information, refer to the following topics:

Related Tasks

- [Flow Search: Data on Which a Flow Search Filters](#)
[Flow Search: Run a Flow Search](#)
- [Flow Search: Run a Flow Search From a Menu](#)
- [Flow Search: Run a Flow Search From Saved Searches](#)

Related Topics

- [Flow Search: Entering Search Criteria Overview](#)
- [Flow Search: Concepts You Need to Understand](#)
- [Flow Search: Overview](#)

Flow Search: Data on Which a Flow Search Filters

This table lists the data that each flow search filters on when run from various locations.

If you run the flow search from the ...	The flow search filters on the ...
Alarm dashboard	Source IP, Target IP, and First Active to current time. If the source or target hosts are multiple hosts, then the Source IP and Target IP do not send information to the query page. If the alarm is still active, then the report filters on First Active.
Alarm List	Source IP, Target IP, and First Active to current time. If the source or target hosts are multiple hosts, then the Source IP and Target IP do not send information to the query page. If the alarm is still active, then the report filters on First Active.
Host Group Report	• Summary section - Host group and time range selected from the graph. • Top Host Groups by Traffic section - Two host groups selected from the graph and the last 12 hours. • Top Applications section - Host group, application, and time range selected from the graph. • Top Alarming Hosts section - IP address and last five minutes.
Host List	First Sent to Last Sent and IP Address. If the First Sent and Last Sent are blank, the report filters on the last five minutes.
Host Report	• Single IP address - IP Address selected and last five minutes. • Traffic by Peer Host Group section - Two host groups selected from the graph and the last 12 hours.
Host Search	First Sent to Last Sent and IP Address. If First Sent to Last Sent is greater than seven days, then the report filters on the seven days prior to Last Sent.

Security Event Details page	Source IP, Target IP and First Active to Last Active. If the source or target hosts are multiple hosts, then the Source IP and Target IP do not send information to the query page. If the event is still active, then the report filters on First Active to current time.
Security Insight dashboard	Top Alarming Hosts section - IP address and last five minutes.

Related Tasks

- [Flow Search: Data on Which a Flow Search Filters](#)
- [Flow Search: Run a Flow Search From a Menu](#)
- [Flow Search: Run a Flow Search From Saved Searches](#)

Related Topics

- [Flow Search: Run a Flow Search Overview](#)
- [Flow Search: Entering Search Criteria Overview](#)
- [Flow Search: Concepts You Need to Understand](#)
- [Flow Search: Overview](#)

Flow Search: Run a Flow Search

To run a flow search from the Flow Search page, complete the following steps:

1. From the main menu, choose **Investigate > Flow Search**).
2. After configuring your settings, click **Search**.

The Flow Search Results page opens.

When a job is in progress, a banner at the top of the page displays the message *Your query is in progress*. A blue *In Progress* pill is also displayed. While the job is running, the message *In Progress* is displayed on the Jobs > [Job Management](#) page in the Current Jobs table. When the job has finished, a blue *Completed* pill is displayed on the Flow Search Results page, and on the Job Management page, the job is displayed in the Finished Jobs table with a status of *Completed* (only if there were results).

In both Data Store and Non-Data Store deployments, you can cancel a flow search that has a status of *Pending* or *In Progress*. If you cancel a flow search that is in progress, you can still view the partial results that were returned. On the Job Management page, click the name of the job to open the Flow Results page, which displays the results that were returned before you canceled the job.

The job name for the flow search consists of the date and time that the flow was generated, so it is possible that job names for different flow search jobs could be duplicated.

3. (Optional) To save the results, click **Save Results**.
4. (Optional) To save the search criteria, click **Save Search**.



See the "Manage Flow Search Jobs" section in the [Managing Jobs](#) topic for more information about flow search guidelines.

Related Tasks

- [Flow Search: Data on Which a Flow Search Filters](#)
- [Flow Search: Run a Flow Search From a Menu](#)
- [Flow Search: Run a Flow Search From Saved Searches](#)

Related Topics

- [Flow Search: Run a Flow Search Overview](#)
- [Flow Search: Entering Search Criteria Overview](#)

- [Flow Search: Concepts You Need to Understand](#)
- [Flow Search: Overview](#)

Flow Search: Run a Flow Search From a Menu

You can run a flow search by accessing the menu (represented by the **Ellipsis** icon) located on most pages throughout the Manager WebApp. By default, the results use the criteria that is contextually related to the IP address or the row you have chosen from which to run the flow search. You can edit the search criteria if you desire.

To access the menu, do one of the following:

- Click the **Ellipsis** icon located next to the applicable IP address.
- Click the **Ellipsis** icon in the Actions column of a data table or configuration table.
- Click a point in a graph. Exceptions are as follows:
- For the Traffic by Peer Host Group graph on the Host Report page and the Top Host Groups by Traffic graph on the Host Group Report page, you must click a host group, a column, or the line between two host groups.

You can run a flow search from the menu for flows that

- are associated with a specific host (but not multiple hosts), host group, security event, alarm, or application
- have occurred between two host groups

To run a flow search...	From the menu, click...
With the default criteria	View Flows. The Flow Search Results page opens with pre-populated fields containing contextual information related to the IP address or row from which you are running the flow search. The flow search begins to run automatically.
With criteria you have edited	Edit (located next to the View Flows option), make the necessary edits on the Flow Search page, and click Search.

To run a flow search...	From the menu, click...
	The Flow Search Results page opens.

Default criteria that is used

Single IP address

When you run a flow search from a menu associated with a single IP, the following default criteria is used:

- IP address.
- The IP address is the Subject.
- Orientation is *Either*.
- Time frame is for last five minutes.

Row of data

When you run a flow search from a menu associated with a row of data, the following default criteria is used (the data in the row dictates which default criteria is used):

- Subject/peer, source/target, or client/server .
- By default the subject, source, or client in the row is the Subject on the Flow Search page. However, if any of these fields contain multiple values, then the peer, target, or server becomes the default Subject on the Flow Search page.
- Time frame of the row data.

Graph

When you run a flow search from a point within a graph, the following default criteria is used:

- Orientation is *Either*.
- The time frame is based on where you initiated the flow search from within the graph. (For information about this, see [Flow Search: Run a Flow Search From a Menu](#) in this topic.
- For the Traffic by Peer Host Group graph on the Host Report page and the Top Host Groups by Traffic graph on the Host Group Report page, the following applies:
 - The current host group (displayed in the middle of the graph) is the Subject.
 - The host group at the other end of the line selected is the Peer.

- [Flow Search: Run a Flow Search From a Menu](#)

Related Tasks

- [Flow Search: Data on Which a Flow Search Filters](#)
- [Flow Search: Run a Flow Search From a Menu](#)
- [Flow Search: Run a Flow Search From Saved Searches](#)

Related Topics

- [Flow Search: Run a Flow Search Overview](#)
[Flow Search: Entering Search Criteria Overview](#)
- [Flow Search: Concepts You Need to Understand](#)
- [Flow Search: Overview](#)

Flow Search: Run a Flow Search From Saved Searches

1. From the main menu, click **Investigate > Search Management > Saved Searches**.
2. Click the saved search you want to run.
The Flow Search page opens.
3. Configure your settings and click **Search**.

Related Tasks

- [Flow Search: Data on Which a Flow Search Filters](#)
- [Flow Search: Run a Flow Search From Saved Searches](#)
- [Flow Search: Run a Flow Search From a Menu](#)

Related Topics

- [Flow Search: Run a Flow Search Overview](#)
[Flow Search: Entering Search Criteria Overview](#)
- [Flow Search: Concepts You Need to Understand](#)
- [Flow Search: Overview](#)

Flow Search Results: Overview

 For information about top report results, refer to [Viewing Top Report Results](#).

Use this page to view and analyze detailed information about a flow between two end points for an alarm, security event, host, or user you have been investigating. The Flow Table includes data for hosts that have translated IP addresses (i.e., if your router has Network Address Translation [NAT] enabled). Flow results are pulled from all Flow Collectors on your network, unless specific Flow Collectors are selected in the Advanced Connection Options section.

Open Flow Search Results

Do one of the following to access this page

- Run a flow search.
- From the main menu, choose **Investigate > Search Management > Saved Results**.
- On the Job Management page, click **a job name**.

Flow Search Results: Table of Contents

Flow Search Results: View Flow Search Results Overview

Refer to the following for flow search retention times:

- For a flow search request with a Max Records Returned of 10,000 or less, both the search and the results are retained for 48 hours on the Job Management page from the time the job completed.
- For a flow search request with a Max Records Returned of 20,000 or more (whose results can only be viewed by downloading them to a .CSV file), both the search and the results are retained for 7 days on the Job Management page from the time the job completed.

Related Tasks

[Flow Search Results: Navigate Results Page](#)

[Flow Search Results: Filter Results](#)

[Flow Search Results: View Results](#)

[Flow Search Results: Column Definitions](#)

Related Topics

[Flow Search Results: What You Can Do Next Overview](#)

[Flow Search Results: Overview](#)

Flow Search Results: Navigate Results Page

When navigating this page, refer to the following guidelines:

- To select columns to display or suppress (permitted only in the tabular view), click **Manage Columns** in the table. The Flow Search Table Columns panel opens. These selections continue to be displayed in future Top Report results until you remove them.
- If you want the Actions Column to display on this page, you must select the Actions check box in the Flow Search Table Columns panel. The Actions option is located at the top of the panel in the All section. Click **Set** at the bottom of the panel to accept your selection.
- To view detailed results for a flow, click the arrow in the row containing the flow result you want to view.
- The General tab contains general information (e.g., packets, bytes, and payload) about the associated flow result. To view detailed proxy records, click the **View URL Data** link located at the top of the Subject section within this tab.
- The Interfaces tab contains detailed information about the interfaces associated with the flow result. Click the exporter name to access the Host Report for the IP address associated with this flow search.
- To enable the Interfaces tab to be displayed on this page, you must select the Include Interface Data check box on the Flow Search page.
- To produce your own reports or if you need to share the flow results with others in your organization, you can export the results to a comma-separated values (.CSV) file. Do one of the following:
 - Click **Export > All** in the table to include data from all of the available columns (even if you have not selected all of them to be displayed on the Flow Search Results page). Click **Manage Columns** in the Flow Search Results table to see a list of all the columns that are available on which you can filter the results.
 - Click **Export > Visible** in the table to include data from only the columns on which you filtered the associated flow search results.
- To view the summary of the values for the following fields returned in the results for a flow search, click **Summary**.
- The export function exports the tabular view. When you export a job, local time is used. However, the time displayed in the job file name on the Jobs Management page uses UTC to reflect the time that the flow search or top report job was created on the server.

- To generate a network diagram, click **More > View diagram**.
- By default, the flow results are displayed in Tabular View (refer to the first image below). This view is indicated by the  (**Tabular View**) icon (the icon with three horizontal lines), located in the Flow Search Results table. To switch to Conversational View (refer to the second image below), click the  (**Conversational View**) icon (the icon with the grid), located to the right of the **List View** icon. The icon that has a white background indicates the active view.

The default sort is by Total Traffic Summary in descending order. The direction of the arrows indicates client/server orientation of the traffic flow.

Tabular View

Flow Search Results (2,000)

Edit Search Last 30 Days (Time Range) 2,000 (Max Records) Save Search Save Results Start New Search

Subject: **Either (Orientation)** Delete Search Completed

Connection: **All (Flow Direction)**

Manage Columns Summary Export 

Start	Duration	Subject IP Add...	Subject Port/Pr...	Subject Host G...	Subject Bytes	Application	Total Bytes	Peer IP Address	Peer Port/Prot...	Peer Host Grou...	Peer Bytes	Ac
Ex. 06/09/...	Ex. <=50min4s	Ex. 10.10.10.1	Ex. 57100/UDP	Ex. "catch All"	Ex. <=50M	Ex. "Corporate"	Ex. <=50M	Ex. 10.255.25...	Ex. 2055/UDP	Ex. "Catch All"	Ex. <=50M	
Jan 25, 2023 8:40:26 AM (3hr 41min 57s ago)	3hr 37min 33s	199.76.147.190	47360/TCP	United States	522.67 M	Amazon Cloud	522.68 M	52.46.133.3 ...	443/TCP	United States	12.73 K	...
Jan 25, 2023 11:53:50 AM (28min 33s ago)	24min 9s	10.83.179.14 ...	53664/UDP	Catch All	234.37 M	HTTPS	319.94 M	10.90.86.104 ...	443/UDP	Catch All	85.58 M	...

Conversational View

Flow Search Results (2,000)

Edit Search Last 30 Days (Time Range) 2,000 (Max Records) Save Search Save Results Start New Search

Subject: **Either (Orientation)** Delete Search Completed

Connection: **All (Flow Direction)**

Manage Columns Summary Export 

Duration	Subject	Subject Port/Protocol	Traffic Summary	Peer Port/Protocol	Peer	Actions
Start: Jan 25, 2023 8:40:26 AM End: Jan 25, 2023 12:17:59 PM Duration: 3hours 37minutes 33seconds	 199.76.147.190 ... View URL Data United States	47360/TCP	522.67 MB 1.3 Mpackets → Amazon Cloud ← 12.73 KB 30 packets	443/TCP	 52.46.133.3 ... United States	...
Start: Jan 25, 2023 11:53:50 AM End: Jan 25, 2023 12:17:59 PM Duration: 24minutes 9seconds	 10.83.179.14 ... View URL Data RFC 1918 98.01.97.ab:27.93	53664/UDP	234.37 MB 201.04 Kpackets → HTTPS ←	443/UDP	 10.90.86.104 ... lwrlabvpn.lancope.ciscolabs.com	...

Related Tasks

Flow Search Results: Filter Results

[Flow Search Results: View Results](#)

[Flow Search Results: Column Definitions](#)

Related Topics

[Flow Search Results: View Flow Search Results Overview](#)

[Flow Search Results: What You Can Do Next Overview](#)

[Flow Search Results: Overview](#)

Flow Search Results: Filter Results

To filter results, type an entry in the applicable field (you do not have to press Enter to begin the filtering). You can filter on multiple fields at one time. It does not matter in which order you filter the fields. Every time you filter on additional or different criteria, the flow search filters off the original flow search results. Across fields (columns), the criteria is ANDED together (the search returns results that match all the filter criteria). Within fields (columns), criteria is OR'd together (the search returns results that match one, some, or all the filter criteria).

Alphanumeric Data

- When entering multiple entries into a text field, leave a space between each entry: *874/TCP 47178/TCP*
- The results include any portion of an alphanumeric entry. If you want an exact match, you must enclose the entry within quotation marks.

Examples:

- If you enter *HTTP* in the Application field, the results include all flows that contain HTTP, HTTPS, HTTP (unclassified), and HTTPS (unclassified).
 - If you want flows containing only HTTP as a value for the Application field, you must enter *"HTTP"*.
-
- To exclude an item from the results, preface it with an exclamation point (!):
!874/TCP

If the item you want to exclude contains a space and you need an exact match, you must enclose it in quotation marks: *!"Catch All"*. If you do not, the results exclude the first part (Catch) and include the second part (All).

Numeric Data

- For numeric entries, if you want to search for results using two entries where the results overlap each other to some degree, search on a range of values instead of using the "less than" symbol (<) and the greater than symbol (>).
- If you want to search for flows that contain values greater than 300M and less than 600M for subject bytes, enter 300M-600M in the Subject Bytes field. Do not enter >300M <600M, as the results include flows that are greater than 300M EVEN IF they

are greater than 600M, and flows that are less than 600M EVEN IF they are less than 300M (therefore every flow would be a match).

So in this scenario, it would be as if you did not even enter additional filter criteria, because the flow search would include the same results as the original search.

- When entering multiple entries into a text field, leave a space between each entry: *874/TCP 47178/TCP*.
- Numeric entries must be exact. Results do not include portions of entries.
- To exclude an item from the results, type an exclamation point (!) before the item.
- You can use the following units of measure: K, M, G, and T. Do not leave a space between the number and unit of measure (e.g., enter *10M*, not *10 M*).
- Use hr, min, and s for hour(s), minute(s), and second(s), respectively. Do not leave any spaces within a single entry (e.g., enter *3hr15min0s*). You do not have to search on all of these time units; you can search on a few or one, if desired.
- You can use the following symbols. Do not leave a space between the number and symbol.
- Greater than: >
- Greater than or equal to: >=
- Less than: <
- Less than or equal to: <=
- Range: -

Time-Related Data

To edit the time range, you can do one of the following:

- Paste or type the entry into the Start or End field using this format: MM/DD/YYYY HH:MM AM - MM/DD/YYYY HH:MM PM.
- Click inside the **Start** field to choose a range of start days and times for flow data. If you want to choose a range of end days and times for flow data, click inside the **End** field (if you don't see the End field, click **Manage Columns** on the Flow Search Results table, check the **End** column check box, and click **Set**).

Two calendars open. Today's date will be highlighted in blue in whichever calendar contains this date.

1. First, click the date you want to use as the start date of the range (if you want to use today's date as the start date, you must still click it even though it is highlighted). To move to another month in either calendar, use the arrow.

2. After you have clicked the start date, find the desired end date of the range (which must be a later date than the start date) and click it. The days between those two dates will be highlighted in light blue.
 - If you need to change your selection after you have selected both dates, you will need to start over. (Each time you click, it toggles between selecting a start date and an end date.)
3. Choose the desired times and then click **Apply** to accept. Time is displayed according to your browser's time zone.

Related Tasks

[Flow Search Results: Navigate Results Page](#)

[Flow Search Results: View Results](#)

[Flow Search Results: Column Definitions](#)

Related Topics

[Flow Search Results: View Flow Search Results Overview](#)

[Flow Search Results: What You Can Do Next Overview](#)

[Flow Search Results: Overview](#)

Flow Search Results: View Results

On this page you can retrieve answers to the following questions:

- Which host is the source of the data?
- Was any data uploaded, and if so, how much and to which host?
- Where are the hosts involved in the flow located? Are these locations expected for your network or for this user?
- Which ports did the flow use?
- What routers or switches did the flow pass through? Is there any information indicating Round Trip Time (RTT), Server Response Time (SRT), or Time to Live (TTL) has been degraded, thus affecting network performance?
- What applications and services were involved in the flow?
- How long did the connection last? Is the length of time unusual for your network, this host, or this user?
- Does either host have a Network Translation Address (NAT) address, indicating an IP address has been changed?
- Which payload was seen by a Flow Sensor?
- Which users are involved in the flow? If identity data is available, it is displayed in the Search Subject column, the Peer column, or both. This information is displayed only in the Conversational view; however, if you export the Tabular view to a .CSV file, the identity data is displayed in the .CSV file.
- Which URL was seen by a Flow Sensor?

Related Tasks

[Flow Search Results: Navigate Results Page](#)

[Flow Search Results: Filter Results](#)

[Flow Search Results: Column Definitions](#)

Related Topics

[Flow Search Results: View Flow Search Results Overview](#)

[Flow Search Results: What You Can Do Next Overview](#)

[Flow Search Results: Overview](#)

Flow Search Results: Column Definitions

For column definitions, refer to the following table.

Column	Description
Actions	Click the ... (Ellipsis) icon to perform any of the following actions: • Run another flow search with the same parameters (click View Flows). • Run another flow search but with one or more different parameters (click View Flows > Edit). • Run a top report query (click Top Reports > [top report name]). • Run a host report for either the source IP or target IP (click External Lookup > Host Report).
Appliance	The Flow Collector that observed the data.
Application	The application that the corresponding flow was carrying as determined by the configuration of the application definitions.

Column	Description
Application (Flow Sensor)	The application ID that the Flow Sensor observed for this flow.
Application (NBAR)	The application ID that an NBAR-enabled device has observed for this flow. This is a fixed list provided by Cisco®.
Application (PacketShaper)	The application ID that a PacketShaper device has observed for this flow. This is a fixed list provided by BlueCoat.
Application (Palo Alto Networks)	The application name that a Palo Alto Networks firewall detected for this flow.
Byte Rate	The total transfer rate, in bytes per second, of data exchanged between the client host and the server host. This field is blank if a Cisco ASA firewall is the only device reporting the flow.
Domain	The domain in which the Flow Collector exists.
Duration	The length of time for which a flow was active.

Column	Description
Encryption (Key Exchange)	The encryption method used in the Key Exchange between two endpoints (for example, Diffie-Hellman, RSA) .
Encryption (TLS/SSL Version)	The TLS/SSL Version used in the flow encryption between the two endpoints (for example, 1.0, 1.1, 1.3).
Encryption Authentication Algorithm	The Algorithm used in the Encryption Authentication between the two endpoints (for example, HMAC-SHA-1, HMAC-MD5).
Encryption Algorithm and Key Length	The Algorithm used in the Encryption between the two endpoints (for example, DES-CBC, 3DES-CBC, Triple-DES, Blowfish, AES-CBC).
Encryption MAC	The Message Authentication Code algorithm used in data integrity checks (for example, SHA256, SHA384, MD5).
End	The time (YYYY-MM-DD HH:MM:SS) the last packet in the flow was transmitted, followed by

Column	Description
	the duration.
Flow Action	The status of the last event from a device such as a firewall, as follows: • Denied – All connections in the flow were blocked by device. • Permitted – All connections in the flow were accepted by a device. • Denied & Permitted – Some connections in the flow were accepted by a device, but other connections in the flow were blocked by a device. This field is empty if no corresponding device has contributed to the flow. You can see the flow actions per interface or interface group on the Quick View for Flow: Interfaces Tab.
Flow ID	Every flow record in the Flow Collector has a unique ID. The ID is useful for debugging and

Column	Description
	troubleshooting.
MPLS Label	The Multi-Protocol Label Switching label.
Packet Rate	The total transfer rate, in packets per second, of data exchanged between the client and the server. This field is blank if a Cisco ASA firewall is the only device reporting the flow.
Peer ASN	The peer AS number as listed in the IANA regional registries.
Peer ASN Assignment	The peer AS handle as assigned by the IANA.
Peer Byte Rate	The Bytes per second rate of the peer.
Peer Byte Ratio	The ratio, in percent, of server bytes to total peer bytes
Peer Bytes	The total number of estimated payload bytes sent by the peer. Ethernet header bytes are subtracted in addition to bytes for MPLS and VLAN tags if they are present in the

Column	Description
	Netflow/IPFIX record for each packet.
Peer File Hash	The SHA256 hash of the process image on disk associated with the peer flow.
Peer FIN Packets	The number of TCP packets, initiated by the peer, with the FIN flag enabled (i.e., finished sending data).
Peer Hostname	The hostname of the peer.
Peer Host Groups	The host groups to which the peer belongs.
Peer Interfaces	The IP address of up to ten exporters used by the peer, followed by the SNMP Interface Index (in parentheses) of the interface or interface group on the router or switch closest to the peer. The exporters are listed in descending order of TTL if that information is available. The process of retrieving data lengthens when you include subject interfaces in your flow query.

Column	Description
Peer IP Address	The IP address and name (if available) assigned to the peer. The IP address is displayed if the DNS resolution was not successful.
Peer Location	The geographic location of the peer.
Peer MAC Address	The MAC address of the peer.
Peer MAC Vendor	The ethernet interface card vendor for the peer.
Peer NAT	The Network Translation IP Address of the peer.
Peer NAT Hostname	The Network Translation Hostname of the peer.
Peer NAT Port	The Network Translation Port of the peer.
Peer Orientation	The client/server orientation of the peer related to the subject
Peer Packet Rate	The rate, in packets per second, that the peer has sent packets to the subject.
Peer Packets	The number of TCP packets, initiated by the peer.

Column	Description
Peer Parent File Hash	The SHA256 hash of the process image on disk of the parent process associated with the peer flow.
Peer Parent Process Name	The name of the parent process of the process associated with the peer flow.
Peer Payload	Bytes in the first packet Payload field of the peer.
Peer Port/Protocol	The port that the corresponding flow was using for the peer. The protocol that the corresponding flow was carrying for the peer.
Peer Process Account	The user account/authority that created the process associated with the peer flow.
Peer Process Name	The name of the process associated with the peer flow.
Peer RST Packets	The number of TCP packets, initiated by the server, with the RST (i.e., reset) flag enabled (i.e., breaking the connection).

Column	Description
Peer SYN Packets	The number of TCP packets, initiated by the server, with the SYN flag enabled and the ACK flag disabled (i.e., a server is establishing a connection with a client).
Peer SYN/ACK Packets	The number of TCP packets, initiated by the server, with the SYN and ACK flags enabled (i.e., the connection is established).
Peer TrustSec ID	The Security Group Tag (SGT) assigned to the peer via the Identity Services Engine (ISE) or Netflow/IPFIX.
Peer TrustSec Name	The Security Group Name (SGN) assigned to the peer via the Identity Services Engine (ISE) or Netflow/IPFIX.
Peer User	The name of the peer. An identity device such as the Cisco Identity Services Engine (ISE), Palo Alto Networks firewall, or Cisco ASA must be installed on your system for data to display in this column.

Column	Description
	If a user name is not available at the moment an alarm occurs, the engine cannot capture it, and the user name will not be displayed in this column. However, the engine may capture the user name after the flow begins and the alarm occurs, and then insert it in the Flow Table.
Protocol	Generally displays UDP or TCP, unless the flow is using FTP, in which case TCP-FTP displays.
RTT Average	The Average Round-Trip Time, or the average amount of time (in milliseconds) required for all the TCP connections to occur in the flow.
RTT Maximum	The Maximum Round-Trip Time, or the greatest amount of time (in milliseconds) required for all the TCP connections to occur in the flow.
RTT Minimum	The Minimum Round-

Column	Description
	Trip Time, or the least amount of time (in milliseconds) required for all the TCP connections to occur in the flow.
Service	The common name of the service used by the flow, if available.
SRT Average	The Average Server Response Time, or the average amount of time (in milliseconds) between the first client request and the first server response for all the TCP connections in the flow.
SRT Maximum	The Maximum Server Response Time, or the greatest amount of time (in milliseconds) between the first client request and the first server response among all the TCP connections in the flow.
SRT Minimum	The Minimum Server Response Time, or the least amount of time (in milliseconds) between the first client request and the first server

Column	Description
	response among all the TCP connections in the flow.
Start	The time (YYYY-MM-DD HH:MM:SS) the first packet in the flow was transmitted. The dates and times for flows correspond to the location of the Flow Collector.
Subject ASN	The subject AS number as listed in the IANA regional registries.
Subject ASN Assignment	The subject AS handle as assigned by the IANA.
Subject Byte Rate	The Bytes per second rate of the subject.
Subject Byte Ratio	The ratio, in percent, of server bytes to total subject bytes.
Subject Bytes	The total number of estimated payload bytes sent by the subject. Ethernet header bytes are subtracted in addition to bytes for MPLS and VLAN tags if they are present in the Netflow/IPFIX record for each packet.

Column	Description
Subject File Hash	The SHA256 hash of the process image on disk of the parent process associated with the subject flow.
Subject FIN Packets	The number of TCP packets, initiated by the subject, with the FIN flag enabled (i.e., finished sending data).
Subject Hostname	The name of the subject.
Subject Host Groups	The host groups to which the subject belongs.
Subject Interfaces	The IP address of up to ten exporters used by the subject, followed by the SNMP Interface Index (in parentheses) of the interface or interface group on the router or switch closest to the subject. The exporters are listed in descending order of TTL if that information is available. The process of retrieving data lengthens when you include subject interfaces in your flow query.
Subject IP Address	The IP address and name (if available)

Column	Description
	assigned to the subject. The IP address is displayed if the DNS resolution was not successful.
Subject Location	The geographic location of the subject.
Subject MAC Address	The MAC address of the subject.
Subject MAC Vendor	The ethernet interface card vendor for the subject.
Subject NAT	The Network Translation IP Address of the subject.
Subject NAT Hostname	The Network Translation Hostname of the subject.
Subject NAT Port	The Network Translation Port of the subject.
Subject Orientation	The client/server orientation of the subject.
Subject Packet Rate	The rate, in packets per second, that the subject has sent packets to the peer.
Subject Packets	The number of TCP packets, initiated by the subject.

Column	Description
Subject Parent File Hash	The SHA256 hash of the process image on disk of the parent process associated with the subject flow.
Subject Parent Process Name	The name of the parent process of the process associated with the subject flow.
Subject Payload	Bytes in the first packet Payload field of the subject.
Subject Port/Protocol	The port that the corresponding flow was using for the subject. The protocol that the corresponding flow was carrying for the subject.
Subject Process Account	The user account/authority that created the process associated with the subject flow.
Subject Process Name	The name of the process associated with the subject flow.
Subject RST Packets	The number of TCP packets, initiated by the subject, with the RST (i.e., reset) flag enabled (i.e., breaking the

Column	Description
	connection).
Subject SYN Packets	The number of TCP packets, initiated by the subject, with the SYN flag enabled and the ACK flag disabled (i.e., a client is establishing a connection with a server).
Subject SYN/ACK Packets	The number of TCP packets, initiated by the subject, with the SYN and ACK flags enabled (i.e., the connection is established).
Subject TrustSec ID	The Security Group Tag (SGT) assigned to the subject via the Identity Services Engine (ISE) or Netflow/IPFIX.
Subject TrustSec Name	The Security Group Name (SGN) assigned to the subject via the Identity Services Engine (ISE) or Netflow/IPFIX.
Subject User	The name of the subject. An identity device such as the Cisco® Identity Services Engine (ISE), Palo Alto Networks firewall, or Cisco ASA must be installed on your

Column	Description
	system for data to display in this column. If a user name is not available at the moment an alarm occurs, the engine cannot capture it, and the user name will not be displayed in this column. However, the engine may capture the user name after the flow begins and the alarm occurs, and then insert it in the Flow Table.
TCP Connections	The number of TCP connections (i.e., three-way handshakes, from the first SYN to the last ACK) that occur during the flow.
TCP Retransmissions	The number of TCP packets that were retransmitted during the flow. Retransmission usually occurs because packets are either damaged or lost.
TCP Retransmission Ratio	The percentage of TCP packets that were retransmitted during the flow. Retransmission usually occurs because

Column	Description
	packets are either damaged or lost.
Total Bytes	The total estimated payload bytes transmitted by both the client and server. Ethernet header bytes are subtracted in addition to bytes for MPLS and VLAN tags if they are present in the Netflow/IPFIX record for each packet.
Total Packets	The total number of packets exchanged between the client and the server. This field is blank if a Cisco ASA firewall is the only device reporting the flow.
Total Traffic (Bps)	The total transfer rate, in bits per second, of data exchanged between the client and the server. This field is blank if a Cisco ASA firewall is the only device reporting the flow.
VLAN ID	A number identifying a virtual LAN.

Related Tasks

[Flow Search Results: Navigate Results Page](#)

[Flow Search Results: Filter Results](#)

[Flow Search Results: View Results](#)

Related Topics

[Flow Search Results: View Flow Search Results Overview](#)

[Flow Search Results: What You Can Do Next Overview](#)

[Flow Search Results: Overview](#)

Flow Search Results: What You Can Do Next Overview

You can complete the following actions from the Flow Search Results page:

- [Flow Search Results: Run a New Flow Search](#)
- [Flow Search Results: Save Flow Search Results](#)
- [Flow Search Results: Save Flow Search](#)
- [Flow Search Results: Delete Results](#)
- [Flow Search Results: Download the CSV File](#)

Related Topics

[Flow Search Results: View Flow Search Results Overview](#)

[Flow Search Results: Overview](#)

Flow Search Results: Run a New Flow Search

From the Flow Search Results, complete the following steps:

1. Click **Start New Search**.

The Flow Search page opens.

2. If necessary, edit the desired criteria and then click **Search**. When a job is in progress, a banner at the top of the page displays the message *Your query is in progress*. A blue *In Progress* pill is also displayed. While the job is running, the message *In Progress* is displayed on the Jobs > [Job Management](#) page in the Current Jobs table. When the job has finished, a blue *Completed* pill is displayed on the Flow Search Results page, and on the Job Management page, the job is displayed in the Finished Jobs table with a status of *Completed* (only if there were results).

Related Topics

[Flow Search Results: What You Can Do Next Overview](#)

[Flow Search Results: View Flow Search Results Overview](#)

[Flow Search Results: Overview](#)

Flow Search Results: Save Flow Search Results

To save flow search results, complete the following steps:

1. Run a flow search from the Flow Search page (from the main menu, choose **Investigate > Flow Search**).
2. When the Flow Search Results page opens, click **Save Results**.
3. Enter a name in the Results Name field and, if desired, enter a description in the Description field.

Go to the Save Results page to see this saved result set.

Related Topics

[Flow Search Results: What You Can Do Next Overview](#)

[Flow Search Results: View Flow Search Results Overview](#)

[Flow Search Results: Overview](#)

Flow Search Results: Save Flow Search

To save flow search criteria, complete the following steps:

1. Run a flow search from the Flow Search page (from the main menu, choose **Investigate > Flow Search**).
2. When the Flow Search Results page opens, click **Save Search**.
3. Enter a name in the Search Name field and, if desired, enter a description in the Description field.

Go to the Save Searches page to see this saved result set.

You can save multiple flow searches with the same name, or you can give each one a unique name. You can use any special character in the name. Saved searches are displayed on the Analyze > [Saved Searches](#) page.

Related Topics

[Flow Search Results: What You Can Do Next Overview](#)

[Flow Search Results: View Flow Search Results Overview](#)

[Flow Search Results: Overview](#)

Flow Search Results: Delete Results

On the Job Management page, in the Actions column for the applicable row, do one of the following:

- If the results for the flow search was created in a downloadable CSV file, you must first click the *** (**Ellipsis**) icon to open the context menu and then click **Delete**.
- For all other flow searches, only a Delete button is displayed, so simply click **Delete**.

Related Topics

[Flow Search Results: What You Can Do Next Overview](#)

[Flow Search Results: View Flow Search Results Overview](#)

[Flow Search Results: Overview](#)

Flow Search Results: Download the CSV File

On the Job Management page, in the row of the applicable job, click the **⋮ (Ellipsis)** icon in the Actions column and select **Download CSV**.

- The flow search job naming conventions used for the zipped and unzipped files are as follows:
- flowAnalysis-[date]-[time].csv,zip, with the time (in current local time) reflecting when the download was initiated.
- flowAnalysis-[date]-[time].csv, with the time (in UTC) reflecting when the flow search job was created on the server.

Related Topics

[Flow Search Results: What You Can Do Next Overview](#)

[Flow Search Results: View Flow Search Results Overview](#)

[Flow Search Results: Overview](#)

Flow Search Results: Generate Network Diagrams

You can generate a network diagram that displays traffic based on a flow search and then monitor the network diagram to see if any alarms are raised for the associated traffic.

To generate a network diagram, complete the following steps.

1. On the Flow Search Results page, click **More > View Diagram** in the table.
 1. Secure Network Analytics generates the network diagram. A progress indicator is displayed during this process.
 2. After the diagram is finished generating, it opens in Edit mode in Network Diagrams so that you can make any further changes if necessary. The diagram contains the deepest child host group of the parent host group that is involved in the flow search.
2. (Optional) If you decide that you want to cancel the generation of the network diagram for any reason (e.g., you want to edit the flow configuration settings), simply click **Cancel** (located inside the progress indicator).
3. In order for Secure Network Analytics to create the associated policies and begin to monitor the traffic, you must click **Save diagram** after the network diagram is generated.

Secure Network Analytics saves the diagram as a snapshot view (the data in the diagram remains static).

Related Topics

[Flow Search Results: What You Can Do Next Overview](#)

[Flow Search Results: View Flow Search Results Overview](#)

[Flow Search Results: Overview](#)

Investigating Flow Proxy Records

Use this page to investigate a host's flow proxy records, which have extensive flow data about a host's transactions between it and the proxy server.

How to find this page

1. From the main menu, choose **Investigate > Flow Search**.
2. Enter the parameters for the flow search and click **Search**.
The Flow Search Results page opens.
3. From the Flow Search Results page, click the ellipsis beside the relevant host IP address.
4. Click **View URL Data**.

What you can view on this page

View the following information on this page:

- Proxy Session Duration
- Source IP/Port
- Proxy IP/Port
- Traffic summary
- Destination IP/Port
- URL host
- URL
- Proxy user name

What you can do next

Determine if anything in the proxy records has caused the alarm or anomaly in traffic. If so, escalate it according to proper procedures.

Related topics

[About Alarm Categories](#)

[Security Event List](#)

Building a Top Report Search

Why you use this page

Use a top report search to do the following

Run this top report...	To determine which...
Top Applications	Applications have been used the most within a specified time period by the searched-on host. You could run this report to help determine what role a segment of the network is playing. For example, if the most used application for the Host or Host Group is SQL, you can determine if it is most likely a database cluster.
Top Conversations	Conversations containing the largest number of bytes the searched-on host has been involved with within a specified period of time. You could run the Top Hosts report to identify systems that are utilizing the most bandwidth, and then from there run the Top Conversations report to see which conversations contained the most data transfer.
Top Hosts	Hosts the searched-on host has communicated with the most within a specified time period of time. You could run this report to identify the servers that are running on a particular port, or to identify systems that are utilizing the most bandwidth.
Top Peers	Peers the searched-on host has communicated with the most within a specified time period of time. If users wanted to identify, for example, who China is connecting with the most in their network, they could run the Top Peers report.
Top Ports	Ports have been used the most within a specified time period by the searched-on host. For Network Segmentation and maintaining appropriate firewall rules, customers need visibility on which servers are in their data centers and to what ports those servers are listening. In this case, users could run the Top Ports report to identify the ports that have been used within a specific time frame. From the report results, users could set or modify their firewall rules.

Top Protocols	Protocols have been used the most within a specified time period by the searched-on host. Use this report to identify which protocols are being used within a network.
Top Services	Services have been used the most within a specified time period by the searched-on host. If there is a giant spike in traffic and you suspect a DDoS attack, you could run a Top Services report. If DNS is the top service, then you could determine if there is an application-based DDoS attack occurring.

How to find this page

Do one of the following to access this page

You can access this page by doing any of the following:

- Go to the Flow Search page (from the main menu, choose **Investigate > Flow Search**).
- Go to the Saved Searches page (from the main menu, choose **investigate > Search Management > Saved Searches**).
- Click the **⋮ (Ellipsis)** icon next to the applicable IP address, in the Actions column, or on a point in a graph. Note that if you are pivoting from the Traffic by Peer Host Group graph on the Host Report page or the Top Host Groups by Traffic graph on the Host Group Report page, you must click a host group, a column, or the line between two host groups.

The Flow Search page opens. Select a top report type in the Search Type field.

Data that each top report filters on when run from various locations

If you run the report from the ...	The report filters on the ...
Alarm dashboard	Source IP, Target IP, and First Active to current time. If the source or target hosts are multiple hosts, then the Source IP and Target IP do not send information to the query page. If the alarm is still active, then the report filters on First Active.

Alarm List	Source IP, Target IP, and First Active to current time. If the source or target hosts are multiple hosts, then the Source IP and Target IP do not send information to the query page. If the alarm is still active, then the report filters on First Active.
Host Group Report	• Summary section - Host group and time range selected from the graph. • Top Host Groups by Traffic section - Two host groups selected from the graph and the last 12 hours. • Top Applications section - Host group, application, and time range selected from the graph. • Top Alarming Hosts section - IP address and last five minutes.
Host List	First Sent to Last Sent and IP Address. If the First Sent and Last Sent are blank, the report filters on the last five minutes.
Host Report	• Single IP address - IP Address selected and last five minutes. • Traffic by Peer Host Group section - Two host groups selected from the graph and the last 12 hours.
Host Search	First Sent to Last Sent and IP Address. If First Sent to Last Sent is greater than seven days, then the report filters on the seven days prior to Last Sent.
Security Event Details page	Source IP, Target IP and First Active to Last Active. If the source or target hosts are multiple hosts, then the Source IP and Target IP do not send information to the query page. If the event is still active, then the report filters on First Active to current time.
Security Insight dashboard	Top Alarming Hosts section - IP address and last five minutes.

Concepts you need to understand to use this page

Which query logic does the system use between and within sections

The top report search builder relies on Boolean logic. Refer to the following for an explanation.

- **Sections:**
 - The top report search builder contains three sections, Subject, Connection, and Peer. These sections are ANDed together, so a search must return results that match the search criteria in all of them. The Connection section returns results matching the server OR client with the given connection search criteria.
 - The orientation field specifies if the Subject sections are for client, server, or either. For example, if you select server, the search returns results that match the Subject section as the server and the Peer section as the client. If you select either, the search returns results matching the Subject section as the client and Peer as the server UNIONed with Subject section as the server and Peer as the client.
- **Attributes:**
 - Each section contains attributes, such as Host, User, Devices, and so on. Each of these attributes functions as a set. If you enter search criteria in multiple sets, the sets are ANDed together within the section.
 - Within an attribute (Host, User, etc.), you can add multiple sets of search criteria. By default, these are designated as "includes" and are ORed together. The search returns results that match one, some, or all of these sets. For example, if you want to search for three IP addresses, you would simply add three IP addresses within that attribute. The search returns results matching any of the three IP addresses.
 - Within an attribute, you can also specify that one or more sets of search criteria be excluded. These individual "excludes" are ORed together (just as the individual "includes" are ORed together). However, each of these "exclude" sets functions as an AND NOT with the other "include" sets within that one attribute. Note that If a user is excluded from the Subject in a query, it is automatically excluded from the Peer, and vice versa.

Example:

You want to run a search on one IP address in the Subject section and two IP addresses in the Peer section.

- In the Subject section, in the Host IP Address or Range attribute, enter the IP address you want included in the text field.
- In the Peer section, in the Host IP Address or Range attribute, enter the two IP addresses in the text field.

If the orientation is server, the search returns results that match the following:

- The Subject IP address is sending flows to the first OR the second Peer IP address.

If the orientation is client, the search returns results that match the following:

- The first Peer IP address OR the second Peer IP address is sending flows to the Subject IP address.

If the orientation is either, the search returns results that match the following:

- The Subject IP address is sending flows to the first OR the second Peer IP address.

UNION

- The first Peer IP address OR the second Peer IP address is sending flows to the Subject IP address.

IPv4 IP address

If you enter this type of IP address...	You must use this formatting...
Single IP address	10.52.1.55 Includes only the host at 10.52.1.55.
Trailing dot subnet	10.52. Includes any IP address from 10.52.0.0 to 10.52.255.255.

If you enter this type of IP address...	You must use this formatting...
	Note: You must include the trailing period (.) at the end.
Classless Inter-Domain Routing (CIDR) notation (i.e., using a “/” to indicate the masked bits)	10.52.1.0/24 Includes any IP address from 10.52.1.0 to 10.52.1.255. Note: Use this notation if you prefer to enter the IP address using a “mask.”
Range of networks	10-11. Includes any IP address from 10.0.0.0 to 11.255.255.255 10.52-53. Includes any IP address from 10.52.0.0 to 10.52.255.255, and from 10.53.0.0 to 10.53.255.255 10.52-55.3. Includes any IP address from 10.52.3.0 to 10.52.3.255; from 10.53.3.0 to 10.53.3.255; from 10.54.3.0 to 10.54.3.255; and from 10.55.3.0 to 10.55.3.255 Note: Be sure to include the trailing period (.) at the end.
Range of hosts	10.52.1.0-10 Includes any IP address from 10.52.1.0 to 10.52.1.10 10.52.0.0-10.52.255.255 Includes any IP address from 10.52.0.0 to 10.52.255.255
Comma-separated list	10.52.1.10,10.52.1.50,10.100.1.20 Includes only these three hosts.

IPv6 IP address

If you enter this type of IP address...	You must use this formatting...
Single IP address	2001:0DB8:0000:0056:0000:ABCD:EF12:3456 2001:DB8:0:56:0:ABCD:EF12:3456 2001:DB8::56:0:ABCD:EF12:3456 2001:DB80:0:56::ABCD:EF12:3456 2001:DB80:0:56::ABCD:239.18.52.86
Global Routing Prefix subnet	2001:DB8:0:56::/64
Range of hosts	2001:DB8:0:56:ABCD:EF12:3456:1-10 2001:DB8:0:56:ABCD:EF12:3456:1- 2001:DB8:0:56:ABCD:EF12:3456:10

About IPv4 and IPv6 addresses

You can use either IPv4 or IPv6 addresses for hosts or host groups. The notation forms that must be used for each type are listed in the applicable table. You must separate the IP addresses by a comma, comma space, space, or return.

The Manager displays IPv6 addresses and IPv4 addresses simultaneously. Of the advantages provided by IPv6, the fact that it supports a larger address space is probably the most well known. IPv6 supports 128-bit addresses, whereas IPv4 supports only 32-bit addresses.

The Manager continues to support IPv4. The Flow Collectors can monitor IPv6 traffic, but they are not IPv6 addressable at this time.

Examples

- An IPv6 address consists of eight blocks of hexadecimal digits separated by colons, as shown in this example (where the x's represent hexadecimal digits):

Example: xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx

- You can replace blocks of 4-digit zeros with double colons (e.g., ::), but only once per IPv6 address. (Leading zeros can also be omitted.) Consider the following address (where the x's represent hexadecimal digits):

Example: xxxx:xxxx:0000:0000:0000:0000:0000:xxxx

This address can be rewritten like this:

Example: xxxx:xxxx::xxxx.

- Like an IPv4 address, an IPv6 address may include a prefix length that specifies how many left-most bits of the address make up the prefix. In IPv4 addresses, the prefix is used to determine the subnet mask (a.k.a., CIDR notation). In IPv6 addresses, it identifies the subnet that an interface belongs to, and routers use it for forwarding.

Example: xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx/40

- Take care when compressing an IPv6 address that contains a prefix, because if not done correctly, it can result in an incorrect interpretation when the address is uncompressed. You can write an IPv4 address as an IPv6 address, like this:

Example: xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:192.168.0.2

- If the first 6 blocks are zeros, you can compress the address like this:

Example: ::192.168.0.2

CIDR notation

Use the following table to understand classless inter-domain routing (CIDR) notation:

Class	CIDR Length	Mask	Number of Addresses
A	/1	128.0.0.0	2,147,483,648
	/2	192.0.0.0	1,073,

Cl as s	CI DR Le ngt h	Mask	Numbe r of Addres ses
		0	741,824
	/3	224.0.0.0	536,870,912
	/4	240.0.0.0	268,435,456
	/5	248.0.0.0	134,217,728
	/6	252.0.0.0	67,108,864
	/7	254.0.0.0	33,554,432
	/8	255.0.0.0	16,777,216
	/9	255.128.0.0	8,388,608
	/10	255.192.0.0	4,194,304
	/11	255.224.0.0	2,097,152
	/12	255.240.0.0	1,048,576
	/13	255.248	524,28

Cl as s	CI DR Le ngt h	Mask	Numbe r of Addres ses
		.0.0	8
	/14	255.252 .0.0	262,14 4
	/15	255.254 .0.0	131,07 2
B	/16	255.255 .0.0	65,536
	/17	255.255 .128.0	32,768
	/18	255.255 .192.0	16,384
	/19	255.255 .224.0	8,192
	/20	255.255 .240.0	4,096
	/21	255.255 .248.0	2,048
	/22	255.255 .252.0	1,024
	/23	255.255 .254.0	512
C	/24	255.255 .255.0	256

Class	CIDR Length	Mask	Number of Addresses
	/25	255.255. .255.128	128
	/26	255.255. .255.192	64
	/27	255.255. .255.224	32
	/28	255.255. .255.240	16
	/29	255.255. .255.248	8
	/30	255.255. .255.252	4
	/31	255.255. .255.254	2

What you can do on this page

Use these guidelines when entering data on this page

Note the following guidelines when entering alphanumeric, numeric, and time data into text fields when building your query:

- In fields that require numerical values, such as Total Bytes, use the following symbols. If you try to enter a numeric value without one of these symbols, you receive an error message.
 - Greater than : >
 - Greater than or equal to: >=
 - Less than: <

- Less than or equal to: <=
- Range: -
- Do not use these special characters in any field in the query builder: () % '
- If you are using a MAC address to define a device, use the full address. Do not enter a partial address.
- The Applications list includes both custom and default applications. A search field is available so you can easily find an application name.
- To use the Cisco TrustSec® tags, you must have previously created Security Group Tags for your Cisco devices in your network.

General Section options

Search Type

Select the specific top report for which you want to run a search.

Time Range

The search defaults to Today (since the last reset hour). If you want to enter a range of dates, click **Custom** in the Time Range drop-down list to enable the From and To fields. Then enter the dates and time for the range. The hour setting is based on 24 hours. To set the hour and minutes, click the field and then slide the vertical bar; the earlier times are to the left, and the later times are to the right. The time is converted to UTC time.

If you have accessed the Top Peers report from the Search Results table on the Host Search page, the report filters on the applicable criteria that pertains to the 7 days prior to the Last Sent (End Date).

Example: If you run a Top Host Report where the First Sent (Start Date) value is August 1 at 10:00 am and the Last Sent (End Date) value is August 22 at 5:20 pm, the results contain the 50 hosts that the searched-on host had communicated the most with from August 15 at 5:20 pm through August 22 at 5:20 pm (a total of seven days).

Search Name

A default search name is pre-populated in the Search Name field; however, you can edit this entry.

Subject section options

Use this section to define characteristics of the subject.

Host IP Address or Range

Note the following guidelines when defining IP addresses:

- Valid formats include IPv4, IPv6, or CIDR notation (192.168.10/16).
- You can enter an IP address or a range (192.168.10-100).
- You can use a wild card at the end of a partial IP address (e.g., 10.0.3*, 10.0.35.*).

Host Groups

Note the following guidelines when selecting host groups:

- Click **Select** under Host Groups to open the Host Group Selector panel.
- Use the include and exclude tabs to choose hosts to include or exclude in your search. The selected hosts are displayed at the top of the panel.
- You can use the Search box in the Host Group Selector panel to quickly search for a host group.
- Click **Apply** to include these hosts in your flow search. Click **Cancel** to discard your selections.

Connection section options

Use this section to define settings for the connection between the subject and peer.

Applications

Note the following guidelines when selecting applications:

- Click **Select** under Applications to open the Application Selector panel.
- Use the include and exclude tabs to choose applications to include or exclude in your search. The selected applications are displayed at the top of the panel.
- You can use the Search box in the Application Selector panel to quickly search for an application.
- Click **Apply** to include these applications in your top report search. Click **Cancel** to discard your selections.

Port/Protocol

When you know that a particular IP address has been involved in a conversation over a given port, but you do not know if this host was the subject or peer, complete the Port/Protocol field in the Connection section. In this situation, Secure Network Analytics ORs together the criteria in these two fields and returns flows when the port/protocol you entered was used by either a subject or peer that matches the IP address you entered (or

one of the IP addresses if you entered a range). The table that displays these results indicates whether it was the subject or the peer who used the port/protocol you entered.

Direction

Specify the type of traffic you want included in the results: Total, Inbound, Outbound, or Within.

Peer section options

Use this section to define characteristics of the peer (i.e., the other entity) involved in the flow.

Host IP Address or Range

Note the following guidelines when defining IP addresses:

- Valid formats include IPv4, IPv6, or CIDR notation (192.168.10/16).
- You can enter an IP address or a range (192.168.10-100).
- You can use a wild card at the end of a partial IP address (e.g., 10.0.3*, 10.0.35.*).

Host Groups

Note the following guidelines when selecting host groups:

- Click **Select** under Host Groups to open the Host Group Selector panel.
- Use the include and exclude tabs to choose hosts to include or exclude in your search. The selected hosts are displayed at the top of the panel.
- You can use the Search box in the Host Group Selector panel to quickly search for a host group.
- Click **Apply** to include these hosts in your top report search. Click **Cancel** to discard your selections.

Advanced options

The Advanced Options section allows you to define a search using more specific settings than those in the main section. You might, for example, use the advanced section as part of an in-depth forensic investigation of suspicious behavior involving a specific host. You can set the following criteria:

Subject Orientation

Specify if you want the criteria you set on this page to return results if they match the client, server, or either one.

Flow Collector Name

Select the Flow Collector(s) from which you want the data collected.

Records Returned

Designate the number of records you want returned. The maximum number that can be returned is 200,000.

Interfaces

Click **Select** to open the Interfaces panel. Select the interfaces (listed under their associated exporter) by which you want to filter the top report search. You can select a maximum of 50 interfaces.

Order By

Designate the attributes by which you want this report filtered. You can select bytes, packets, flows, or TCP connection.

Performance Options

When you select one or more interfaces in the Advanced Options section, the Performance Options settings are disabled. Refer to the following table.

Performance Option	Description
Standard	Use the Standard option to return - The most commonly required information for each specific report. "Other" records, which are records that are grouped together that meet the query criteria but exceed the number of records that are requested to be returned.
Advanced	Use the Advanced option to enable the return of all available information in the results. These options enable you to easily omit the data from a group of columns to decrease the time it takes to complete a large query. The Advanced sub-options listed below enable you to easily omit

	the data from a group of columns to decrease the time it takes to complete a large query.	
	Excludes bps/pps values	Increases the speed of the query by omitting the bps and pps values for the 95th percentile, maximum, minimum, and average data.
	Excludes Other records	"Other" records are records that are grouped together that meet the query criteria but exceed the number of records that are requested to be returned.
	Excludes Counts	Excludes the flows, hosts, peers, host client flows and peer server flows counts from the top report results.

Run the top report search

Click **Search**. The Top [Report Name] Search Results page opens. When a job is in progress, a banner at the top of the page displays the message *Your query is in progress*. A blue *In Progress* pill is displayed. While the job is running, the message *In Progress* is displayed on the Jobs > [Job Management](#) page in the Current Jobs table. When the job has finished, a blue *Completed* pill is displayed on the Top [Report Name] Search Results page, and on the Job Management page, the job is displayed in the Finished Jobs table with a status of *Completed* (only if there were results).

In both Data Store and Non-Data Store deployments, you can cancel a top report search that has a status of *Pending* or *In Progress*.

All top report jobs, as well as flow searches with a Max Records Returned of 10,000 or less, run in a different queue from that of flow searches with a Max Records Returned of 20,000 or more.

- You can run a maximum of four jobs at the same time if the Max Records Returned for the flow search is 10,000 or less. It does not matter what the Records Returned value is for a top report job.
- If you start additional top report jobs or flow searches during this time (for either of these two categories of flow searches), they are placed in their respective queue with a Pending status.

If your top report query fails, it may be because the maximum number of queries that can run simultaneously has been reached and your query timed out before the other queries finished running. Wait until the queries currently running have finished, and then rerun your top report.

For more information about top report search guidelines, see the "Manage Top Report Search Jobs" section in the [Managing Jobs](#) topic.

How to run a top report from other pages

You can also run a top report from the Saved Searches page or by clicking the **Ellipsis** icon located next to some IP addresses on certain pages.

To run a top report from the Saved Searches page

1. From the main menu, click **Analyze > Saved Searches**.
2. On the Saved Searches page, click the arrow in the Run column for the saved search you want to run.

The [Top Report Name] Search Results page opens.

To run a top report from a menu

You can run a top report by accessing the menu (represented by the **Ellipsis** icon) located on most pages throughout the Web App. By default, the results use the criteria that is contextually related to the IP address or the row you have chosen from which to run the top report. You can edit the search criteria if you desire.

To access the menu, do one of the following:

- Click the **Ellipsis** icon located next to the applicable IP address.
- Click the **Ellipsis** icon in the Actions column of a data table or configuration table.

- Click a point in a graph. Exceptions are as follows:
- For the Traffic by Peer Host Group graph on the Host Report page and the Top Host Groups by Traffic graph on the Host Group Report page, you must click a host group, a column, or the line between two host groups.

You can run a top report from the menu for information that is associated with

- a specific host (but not multiple hosts), host group, security event, alarm, or application
- activity that has occurred between two host groups

To run a top report...	From the menu, click...
With the default criteria	Top Reports > [Report Name]. The Flow Search Results page opens with pre-populated fields containing contextual information related to the IP address or row from which you are running the flow search.
With criteria you have edited	Top Reports > Edit (for the applicable report), make the necessary edits on the Flow Search page, and click Search. The Flow Search Results page opens.

Default criteria that is used

Single IP Address

When you run a top report search from a menu associated with a single IP, the following default criteria is used:

- IP address.
- The IP address is the Subject.

- Orientation is *Either*.
- Time frame is for last five minutes.

Row of data

When you run a top report search from a menu associated with a row of data, the following default criteria is used (the data in the row dictates which default criteria is used):

- IP address, subject/peer, source/target, or client/server.
- By default the subject, source, or client in the row is the Subject on the [Top Report Name] Search page. However, if any of these fields contain multiple values, then the peer, target, or server becomes the default Subject on the [Top Report Name] Search page.
- Time frame of the row data.

Graph

When you run a top report search from a point within a graph, the following default criteria is used:

- Orientation is *Either*.
- The time frame is based on where you initiated the flow search from within the graph. For information about this, see the table in the "How to find this page" section in [Viewing Top Report Results](#).
- For the Traffic by Peer Host Group graph on the Host Report page and Top Host Groups by Traffic graph on the Host Group Report page, the following applies:
 - The current host group (displayed in the middle of the graph) is the Subject.
 - The host group at the other end of the line selected is the Peer.

What you can do next

Save the top report search

Click **Save**.

You can save multiple top report searches with the same name, or you can give each one a unique name. You can use any special character in the name. Saved searches are listed on the Analyze > [Saved Searches](#) page.

Run a saved search

1. Click **Load Saved Search > Manage Saved Searches**.
2. Select the saved search you want to run.

Related topics

[Accessing Saved Flow Search Results](#)

[Accessing Saved Flow Searches](#)

[Managing Jobs](#)

Viewing Top Report Results

Top reports provide a way for you to view pertinent information quickly so that you can more effectively investigate anomalies in your network. Top Report results are pulled from all Flow Collectors on your network, unless specific Flow Collectors are selected in the Advanced Options section.

How to find this page

Do any of the following to access this page

- Run a top report
- From the main menu, choose **Investigate > Search Management > Saved Results**.
- On the Job Management page, click a job name.

The results and any saved queries are retained for 48 hours. The total figures displayed at the bottom of the page are the totals for the entire search, not just for that particular page.


Data that each top report filters on

If you run the report from the ...	The report filters on the ...
Alarm dashboard	Source IP, Target IP, and First Active to current time. If the source or target hosts are multiple hosts, then the Source IP and Target IP do not send information to the query page. If the alarm is still active, then the report filters on First Active.
Alarm List	Source IP, Target IP, and First Active to current time. If the source or target hosts are multiple hosts, then the Source IP and Target IP do not send information to the query page. If the alarm is still active, then the report filters on First Active.
Host Group Report	• Summary section - Host group and time range selected from the graph. • Top Host Groups by Traffic section - Two host groups selected from the graph and the last 12 hours.


	• Top Applications section - Host group, application, and time range selected from the graph. • Top Alarming Hosts section - IP address and last five minutes.
Host List	First Sent to Last Sent and IP Address. If the First Sent and Last Sent are blank, the report filters on the last five minutes.
Host Report	• Single IP address - IP Address selected and last five minutes. • Traffic by Peer Host Group section - Two host groups selected from the graph and the last 12 hours.
Host Search	First Sent to Last Sent and IP Address. If First Sent to Last Sent is greater than seven days, then the report filters on the seven days prior to Last Sent.
Security Event Details page	Source IP, Target IP and First Active to Last Active. If the source or target hosts are multiple hosts, then the Source IP and Target IP do not send information to the query page. If the event is still active, then the report filters on First Active to current time.
Security Insight dashboard	Top Alarming Hosts section - IP address and last five minutes.

Concepts you need to understand to use this page

How is data calculated for host groups, IP range/list, VM servers, or VMs



How is data calculated for interfaces



Viewing information on this page

How to navigate this page

- To select columns to display or suppress in the tabular view, click **Manage Columns**. These selections continue to be displayed in future Top Report results until you remove them.
- To produce your own reports or if you need to share the flow results with others in your organization, you can export the results to a comma-separated values (.CSV) file. Do one of the following:
 - Click **Export > All** in the table to include data from all of the available columns (even if you have not selected all of them to be displayed on the Flow Search Results page). Click **Manage Columns** in the Flow Search Results table to see a list of all the columns that are available on which you can filter the results.
 - Click **Export > Visible** in the table to include data from only the columns on which you filtered the associated flow search results.

The export function exports the tabular view. When you export a job, local time is used. However, the time displayed in the job file name on the Jobs Management page uses UTC to reflect the time that the flow search or top report job was created on the server.

Report column definitions

You may see differences in the number of total flows from one Top report to another. The Top Hosts, Top Peers, and Top Conversations reports will display approximately double the total number of flows as those shown on the other Top reports or Flow Table. This is because each flow contributes two records for these three reports.

When you are evaluating the number of flows in these reports, the ratio of flows per application, host, etc., is more important than the raw number of flows.



Subject and peer do not denote orientation of the flow unless you have specified the orientation when building the flow (designating the subject as either client or server).

Top Applications

Column	Description
% of Bytes	The utilization percentage for the traffic used by the corresponding application. The percentage is displayed next to a graphic (for which the color never changes) that illustrates the percentage. If you chose to filter the report by bytes, this column will contain data.
95th Packet Rate (pps)	The highest rate of packet transmission (in packets per second) observed on the corresponding date, not including the top 5% of the peak rates. The 95th Percentile is used by firms that perform "burstable billing," where the top 5% of bandwidth peaks are cut off. With this method, infrequent peaks are ignored. For more information, refer to: http://en.wikipedia.org/wiki/Burstable_billing
95th Traffic (bps)	The highest rate of traffic (in bits per second) observed on the corresponding date, not including the top 5% of the peak rates.

Column	Description
	The 95th Percentile is used by firms that perform "burstable billing," where the top 5% of bandwidth peaks are cut off. With this method, infrequent peaks are ignored. For more information, refer to: http://en.wikipedia.org/wiki/Burstable_billing
Application	The application that the corresponding flow was carrying as determined by the configuration of the application definitions.
Average Packet Rate (pps)	The average rate of traffic (in packets per second) observed on the corresponding date.
Average Traffic (bps)	The average rate of traffic (in bits per second) observed on the corresponding date.
Bytes	The number of bytes used by the corresponding application.
Flow Collector Name	The name of the Flow Collector that observed the data.
Flows	The number of flows used by the corresponding application.
Host Bytes	The number of bytes for the hosts selected as the subject (on the Search page) used by the corresponding application.
Host Bytes Ratio	The ratio, in percent, of host bytes to the bytes used by the corresponding application. The percentage is displayed next to a bar graph (for which the color never changes) that illustrates the percentage. If you don't select a host, the displayed entry is 50% (since application traffic consists of both hosts and

Column	Description
	peers.)
Host Client Bytes	The number of bytes used by client hosts selected as the subject (on the Search page) used by the corresponding application.
Host Client Flows	The number of flows used by client hosts selected as the subject (on the Search page) used by the corresponding application.
Host Client Packets	The number of packets used by client hosts selected as the subject (on the Search page) used by the corresponding application.
Host Client TCP Connections *	The number of TCP connections (i.e., three-way handshakes, from the first SYN to the last ACK) used by client hosts selected as the subject (on the Search page) used by the corresponding application.
Host Clients	The number of client hosts selected as the subject (on the Search page) used by the corresponding application.
Host Packets	The number of packets for the hosts selected as the subject (on the Search page) used by the corresponding application.
Host Role	The role of the hosts selected as the subject (on the Search page) used by the corresponding application. The role can be one of the following: • Client • Server • Either
Host Server Bytes	The number of bytes used by server hosts selected as

Column	Description
	the subject (on the Search page) used by the corresponding application.
Host Server Flows	The number of flows used by server hosts selected as the subject (on the Search page) used by the corresponding application.
Host Server Packets	The number of packets used by server hosts selected as the subject (on the Search page) used by the corresponding application.
Host Sever TCP Connections *	The number of TCP connections (i.e., three-way handshakes, from the first SYN to the last ACK) used by server hosts selected as the subject (on the Search page) used by the corresponding application.
Host Servers	The number of hosts selected as the subject (on the Search page), acting as servers, used by the corresponding application.
Hosts	The number of hosts selected as the subject (on the Search page). These are usually the hosts on your network's side of the conversation.
Maximum Packet Rate (Pps)	The peak rate of traffic (in packets per second) used by the corresponding application.
Maximum Traffic (Bps)	The peak rate of traffic (in bits per second) used by the corresponding application.
Minimum Packet Rate (Pps)	The smallest rate of traffic (in packets per second) used by the corresponding application.
Minimum Traffic (Bps)	The smallest rate of traffic (in bits per second) used by the corresponding application.

Column	Description
Packets	The number of packets used by the corresponding application.
Peer Byte Ratio	The ratio, in percent, of peer bytes to the bytes used by the corresponding application. The percentage is displayed next to a graphic (for which the color never changes) that illustrates the percentage.
Peer Bytes	The number of bytes used by hosts selected as the peer (on the Search page) used by the corresponding application.
Peer Client Bytes	The number of bytes used by client hosts among the peer hosts (as selected on the Search page) used by the corresponding application.
Peer Client Flows	The number of flows used by client hosts among the peer hosts (as selected on the Search page) used by the corresponding application.
Peer Client Packets	The number of packets used by client hosts among the peer hosts (as selected on the Search page) used by the corresponding application.
Peer Client TCP Connections *	The number of TCP connections (i.e., three-way handshakes, from the first SYN to the last ACK) used by client hosts among the peer hosts (as selected on the Search page) used by the corresponding application.
Peer Clients	The number of client hosts among the peer hosts (as selected on the Search page) used by the corresponding application.
Peer Packets	The number of packets for the peer hosts (as selected on the Search page) used by the corresponding

Column	Description
	application.
Peer Role	The role of the peer host used by the corresponding application. The role can be one of the following: • Client • Server • Client and Server
Peer Server Bytes	The number of bytes used by server hosts among the peer hosts (as selected on the Search page) used by the corresponding application.
Peer Server Flows	The number of flows used by server hosts among the peer hosts (as selected on the Search page) used by the corresponding application.
Peer Server Packets	The number of packets used by server hosts among the peer hosts used by the corresponding application.
Peer Server TCP Connection *	The number of TCP connections (i.e., three-way handshakes, from the first SYN to the last ACK) used by server hosts among the peer hosts (as selected on the Search page) used by the corresponding application.
Peer Servers	The number of peer hosts (as selected on the Search page), acting as servers, used by the corresponding application.
Peers	The number of peer hosts (as selected on the Search page) associated with the corresponding application. These are usually the hosts on the other side of the conversation.
Records	The number of flow records sent in the selected

Column	Description
	direction associated with the corresponding application.
TCP Connections *	The number of TCP connections (i.e., three-way handshakes, from the first SYN to the last ACK) completed by the corresponding application.

* Not available when the document is filtered on an interface or interface group.

Top Conversations

Column	Description
% of Bytes	The utilization percentage for the conversation occurring most often in the traffic. The percentage is displayed next to a graphic (for which the color never changes) that illustrates the percentage. If you chose to filter the report by bytes, this column will contain data.
95th Packet Rate (pps)	The highest rate of packet transmission (in packets per second) observed on the corresponding date, not including the top 5% of the peak rates. The 95th Percentile is used by firms that perform "burstable billing," where the top 5% of bandwidth peaks are cut off. With this method, infrequent peaks are ignored. For more information, refer to: http://en.wikipedia.org/wiki/Burstable_billing
95th Traffic (bps)	The highest rate of traffic (in bits per second) observed on the corresponding date, not including the top 5% of the peak rates. The 95th Percentile is used by firms that perform "burstable billing," where the top 5% of bandwidth

Column	Description
	peaks are cut off. With this method, infrequent peaks are ignored. For more information, refer to: http://en.wikipedia.org/wiki/Burstable_billing
Actions	Click the ⋯ (Ellipsis) icon to perform any of the following actions: • Run another flow search with the same parameters (click View Flows). • Run another flow search but with one or more different parameters (click View Flows > Edit). • Run a top report query (click Top Reports > [top report name]). • Run a host report for either the source IP or target IP (click External Lookup > Host Report).
Average Packet Rate (pps)	The average rate of traffic (in packets per second) observed on the corresponding date.
Average Traffic (bps)	The average rate of traffic (in bits per second) observed on the corresponding date.
Bytes	The number of bytes used in the corresponding conversation.
Flow Collector Name	The name of the Flow Collector that observed the data.
Flows	The number of flows used in the corresponding conversation.
Host Bytes	The number of bytes for the hosts selected as the subject (on the Search page) used in the corresponding conversation.

Column	Description
Host Bytes Ratio	The ratio, in percent, of host bytes to the bytes used by the corresponding application. The percentage is displayed next to a bar graph (for which the color never changes) that illustrates the percentage. If you don't select a host, the displayed entry is 50% (since conversation traffic consists of both hosts and peers.)
Host Client Bytes	The number of bytes used by client hosts selected as the subject (on the Search page) used in the corresponding conversation.
Host Client Flows	The number of flows used by client hosts selected as the subject (on the Search page) used in the corresponding conversation.
Host Client Packets	The number of packets used by client hosts selected as the subject (on the Search page) used in the corresponding conversation.
Host Client TCP Connections *	The number of TCP connections (i.e., three-way handshakes, from the first SYN to the last ACK) used by client hosts selected as the subject (on the Search page) used in the corresponding conversation.
Host Clients	The number of client hosts selected as the subject (on the Search page) used in the corresponding conversation.
Host Country	The country in which the host selected as the subject (on the Search page) in the corresponding conversation is located.
Host Groups	The host group(s) to which the host selected as the subject (on the Search page) belongs.

Column	Description
Host IP Address	The IP address and name (if available) assigned to the host, selected as the subject (on the Search page), participating in the corresponding conversation. The IP address is displayed if the DNS resolution was not successful.
Host Name	The name assigned to the host selected as the subject (on the Search page), if available.
Host Packets	The number of packets for the hosts selected as the subject (on the Search page) used in the corresponding conversation.
Host Role	The role of the hosts selected as the subject (on the Search page) used in the corresponding conversation. The role can be one of the following: • Client • Server • Either
Host Server Bytes	The number of bytes used by server hosts selected as the subject (on the Search page) used in the corresponding conversation.
Host Server Flows	The number of flows used by server hosts selected as the subject (on the Search page) used in the corresponding conversation.
Host Server Packets	The number of packets used by server hosts selected as the subject (on the Search page) used in the corresponding conversation.
Host Sever TCP Connections *	The number of TCP connections (i.e., three-way handshakes, from the first SYN to the last ACK) used by server hosts selected as the subject (on the Search

Column	Description
	page) used in the corresponding conversation.
Host Servers	The number of hosts selected as the subject (on the Search page), acting as servers, used in the corresponding conversation.
Hosts	The number of hosts selected as the subject (on the Search page) using the corresponding application. These are usually the hosts on your network's side of the conversation.
Maximum Packet Rate (Pps)	The peak rate of traffic (in packets per second) used in the corresponding conversation.
Maximum Traffic (Bps)	The peak rate of traffic (in bits per second) used in the corresponding conversation.
Minimum Packet Rate (Pps)	The smallest rate of traffic (in packets per second) used in the corresponding conversation.
Minimum Traffic (Bps)	The smallest rate of traffic (in bits per second) used in the corresponding conversation.
Packets	The number of packets used in the corresponding conversation.
Peer Byte Ratio	The ratio, in percent, of peer bytes to the bytes used in the corresponding conversation. The percentage is displayed next to a graphic (for which the color never changes) that illustrates the percentage.
Peer Bytes	The number of bytes used by hosts selected as the peer (on the Search page) used in the corresponding conversation.
Peer Client Bytes	The number of bytes used by client hosts among the

Column	Description
	peer hosts (as selected on the Search page) used in the corresponding conversation.
Peer Client Flows	The number of flows used by client hosts among the peer hosts (as selected on the Search page) used in the corresponding conversation.
Peer Client Packets	The number of packets used by client hosts among the peer hosts (as selected on the Search page) used in the corresponding conversation.
Peer Client TCP Connections *	The number of TCP connections (i.e., three-way handshakes, from the first SYN to the last ACK) used by client hosts among the peer hosts (as selected on the Search page) used in the corresponding conversation.
Peer Clients	The number of client hosts among the peer hosts (as selected on the Search page) used in the corresponding conversation.
Peer Country	The country in which the peer host selected as the subject (on the Search page) in the corresponding conversation is located.
Peer Host Groups	The host group(s) to which the peer host selected as the subject (on the Search page) belongs.
Peer IP Address	The IP address and name (if available) assigned to the peer host, selected as the subject (on the Search page), participating in the corresponding conversation. The IP address is displayed if the DNS resolution was not successful.
Peer Name	The name assigned to the peer host selected as the subject (on the Search page), if available.

Column	Description
Peer Packets	The number of packets for the peer hosts (as selected on the Search page) used in the corresponding conversation.
Peer Role	The role of the peer host using the corresponding application. The role can be one of the following: • Client • Server • Client and Server
Peer Server Bytes	The number of bytes used by server hosts among the peer hosts (as selected on the Search page) used in the corresponding conversation.
Peer Server Flows	The number of flows used by server hosts among the peer hosts (as selected on the Search page) used in the corresponding conversation.
Peer Server Packets	The number of packets used by server hosts among the peer hosts used in the corresponding conversation.
Peer Server TCP Connection *	The number of TCP connections (i.e., three-way handshakes, from the first SYN to the last ACK) used by server hosts among the peer hosts (as selected on the Search page) used in the corresponding conversation.
Peer Servers	The number of peer hosts (as selected on the Search page), acting as servers, used in the corresponding conversation.
Peers	The number of peer hosts (as selected on the Search page) associated with the corresponding conversation. These are usually the hosts on the other

Column	Description
	side of the conversation.
Port	The port that the corresponding conversation used.
Records	The number of flow records sent in the selected direction used in the corresponding conversation.
TCP Connections *	The number of TCP connections (i.e., three-way handshakes, from the first SYN to the last ACK) completed in the corresponding conversation.

* Not available when the document is filtered on an interface or interface group.

Top Hosts

Column	Description
% of Bytes	The utilization percentage for the corresponding host. The percentage is displayed next to a graphic (for which the color never changes) that illustrates the percentage. If you chose to filter the report by bytes, this column will contain data.
95th Packet Rate (pps)	The highest rate of packet transmission (in packets per second) observed for the corresponding host, not including the top 5% of the peak rates. The 95th Percentile is used by firms that perform "burstable billing," where the top 5% of bandwidth peaks are cut off. With this method, infrequent peaks are ignored. For more information, refer to: http://en.wikipedia.org/wiki/Burstable_billing
95th Traffic (bps)	The highest rate of traffic (in bits per second) observed on the corresponding date, not including the

Column	Description
	top 5% of the peak rates. The 95th Percentile is used by firms that perform "burstable billing," where the top 5% of bandwidth peaks are cut off. With this method, infrequent peaks are ignored. For more information, refer to: http://en.wikipedia.org/wiki/Burstable_billing
Average Packet Rate (pps)	The average rate of traffic (in packets per second) observed on the corresponding date.
Average Traffic (bps)	The average rate of traffic (in bits per second) observed on the corresponding date.
Bytes	The number of bytes used by all hosts selected as the subject and as the peer (on the Search page).
Flow Collector Name	The name of the Flow Collector that observed the data.
Flows	The number of flows used by all hosts selected as the subject and as the peer (on the Search page).
Host Bytes	The number of bytes for the hosts selected as the subject (on the Search page).
Host Bytes Ratio	The ratio, in percent, of host bytes to the bytes used. The percentage is displayed next to a bar graph (for which the color never changes) that illustrates the percentage. If you don't select a host, the displayed entry is 50% (since flow traffic consists of both hosts and peers.)
Host Client Bytes	The number of bytes used by client hosts selected as the subject (on the Search page).

Column	Description
Host Client Flows	The number of flows used by client hosts selected as the subject (on the Search page).
Host Client Packets	The number of packets used by client hosts selected as the subject (on the Search page).
Host Client TCP Connections *	The number of TCP connections (i.e., three-way handshakes, from the first SYN to the last ACK) used by client hosts selected as the subject (on the Search page).
Host Clients	The number of client hosts selected as the subject (on the Search page).
Host Country	The country in which the host selected as the subject (on the Search page) is located.
Host Groups	The host group(s) to which the host selected as the subject (on the Search page) belongs.
Host IP Address	The IP address and name (if available) assigned to the host selected as the subject (on the Search page) . The IP address is displayed if the DNS resolution was not successful.
Host Name	The name assigned to the host selected as the subject (on the Search page), if available.
Host Packets	The number of packets for the hosts selected as the subject (on the Search page).
Host Role	The role of the hosts selected as the subject (on the Search page). The role can be one of the following: • Client • Server

Column	Description
	• Either
Host Server Bytes	The number of bytes used by server hosts selected as the subject (on the Search page).
Host Server Flows	The number of flows used by server hosts selected as the subject (on the Search page).
Host Server Packets	The number of packets used by server hosts selected as the subject (on the Search page).
Host Sever TCP Connections *	The number of TCP connections (i.e., three-way handshakes, from the first SYN to the last ACK) used by server hosts selected as the subject (on the Search page).
Host Servers	The number of hosts selected as the subject (on the Search page), acting as servers.
Hosts	The number of hosts selected as the subject (on the Search page) used in the corresponding flow. These are usually the hosts on your network's side of the flow.
Maximum Packet Rate (Pps)	The peak rate of traffic (in packets per second) used by the hosts selected as the subject (on the Search page).
Maximum Traffic (Bps)	The peak rate of traffic (in bits per second) used by the hosts selected as the subject (on the Search page).
Minimum Packet Rate (Pps)	The smallest rate of traffic (in packets per second) used by the hosts selected as the subject (on the Search page).

Column	Description
Minimum Traffic (Bps)	The smallest rate of traffic (in bits per second) used by the hosts selected as the subject (on the Search page).
Packets	The number of packets used by the hosts selected as the subject (on the Search page).
Peer Byte Ratio	The ratio, in percent, of peer bytes to the bytes used by the hosts selected as the subject (on the Search page). The percentage is displayed next to a graphic (for which the color never changes) that illustrates the percentage.
Peer Bytes	The number of bytes used by hosts selected as the peer (on the Search page).
Peer Client Bytes	The number of bytes used by client hosts among the peer hosts (as selected on the Search page).
Peer Client Flows	The number of flows used by client hosts among the peer hosts (as selected on the Search page).
Peer Client Packets	The number of packets used by client hosts among the peer hosts (as selected on the Search page).
Peer Client TCP Connections *	The number of TCP connections (i.e., three-way handshakes, from the first SYN to the last ACK) used by client hosts among the peer hosts (as selected on the Search page).
Peer Clients	The number of client hosts among the peer hosts (as selected on the Search page).
Peer Packets	The number of packets for the peer hosts (as selected on the Search page).

Column	Description
Peer Role	The role of the peer host. The role can be one of the following: • Client • Server • Client and Server
Peer Server Bytes	The number of bytes used by server hosts among the peer hosts (as selected on the Search page).
Peer Server Flows	The number of flows used by server hosts among the peer hosts (as selected on the Search page).
Peer Server Packets	The number of packets used by server hosts among the peer hosts.
Peer Server TCP Connection *	The number of TCP connections (i.e., three-way handshakes, from the first SYN to the last ACK) used by server hosts among the peer hosts (as selected on the Search page).
Peer Servers	The number of peer hosts (as selected on the Search page), acting as servers.
Peers	The number of active hosts among the peer hosts (as selected on the Search page).
Records	The number of flow records sent in the selected direction by the hosts selected as the subject (on the Search page).
TCP Connections *	The number of TCP connections (i.e., three-way handshakes, from the first SYN to the last ACK) completed by the hosts selected as the subject (on the Search page).

* Not available when the document is filtered on an interface or interface group.

Top Peers

Column	Description
% of Bytes	The utilization percentage for the corresponding peer. The percentage is displayed next to a graphic (for which the color never changes) that illustrates the percentage. If you chose to filter the report by bytes, this column will contain data.
95th Packet Rate (pps)	The highest rate of packet transmission (in packets per second) observed for the corresponding peer, not including the top 5% of the peak rates. The 95th Percentile is used by firms that perform "burstable billing," where the top 5% of bandwidth peaks are cut off. With this method, infrequent peaks are ignored. For more information, refer to: http://en.wikipedia.org/wiki/Burstable_billing
95th Traffic (bps)	The highest rate of traffic (in bits per second) observed on the corresponding date, not including the top 5% of the peak rates. The 95th Percentile is used by firms that perform "burstable billing," where the top 5% of bandwidth peaks are cut off. With this method, infrequent peaks are ignored. For more information, refer to: http://en.wikipedia.org/wiki/Burstable_billing
Average Packet Rate (pps)	The average rate of traffic (in packets per second) observed on the corresponding date.
Average Traffic (bps)	The average rate of traffic (in bits per second) observed on the corresponding date.
Bytes	The number of bytes used by all hosts selected as the

Column	Description
	subject and as the peer (on the Search page).
Flow Collector Name	The name of the Flow Collector that observed the data.
Flows	The number of flows used by all hosts selected as the subject and as the peer (on the Search page).
Host Bytes	The number of bytes for the hosts selected as the subject (on the Search page).
Host Bytes Ratio	The ratio, in percent, of host bytes to the bytes used. The percentage is displayed next to a bar graph (for which the color never changes) that illustrates the percentage. If you don't select a host, the displayed entry is 50% (since flow traffic consists of both hosts and peers.)
Host Client Bytes	The number of bytes used by client hosts selected as the subject (on the Search page).
Host Client Flows	The number of flows used by client hosts selected as the subject (on the Search page).
Host Client Packets	The number of packets used by client hosts selected as the subject (on the Search page).
Host Client TCP Connections *	The number of TCP connections (i.e., three-way handshakes, from the first SYN to the last ACK) used by client hosts selected as the subject (on the Search page).
Host Clients	The number of client hosts selected as the subject (on the Search page).
Host Packets	The number of packets for the hosts selected as the

Column	Description
	subject (on the Search page).
Host Role	The role of the hosts selected as the subject (on the Search page). The role can be one of the following: • Client • Server • Either
Host Server Bytes	The number of bytes used by server hosts selected as the subject (on the Search page).
Host Server Flows	The number of flows used by server hosts selected as the subject (on the Search page).
Host Server Packets	The number of packets used by server hosts selected as the subject (on the Search page).
Host Sever TCP Connections *	The number of TCP connections (i.e., three-way handshakes, from the first SYN to the last ACK) used by server hosts selected as the subject (on the Search page).
Host Servers	The number of hosts selected as the subject (on the Search page), acting as servers.
Hosts	The number of hosts selected as the subject (on the Search page) used in the corresponding flow. These are usually the hosts on your network's side of the flow.
Maximum Packet Rate (Pps)	The peak rate of traffic (in packets per second) used by the hosts selected as the subject (on the Search page).
Maximum Traffic	The peak rate of traffic (in bits per second) used by

Column	Description
(Bps)	the hosts selected as the subject (on the Search page).
Minimum Packet Rate (Pps)	The smallest rate of traffic (in packets per second) used by the hosts selected as the subject (on the Search page).
Minimum Traffic (Bps)	The smallest rate of traffic (in bits per second) used by the hosts selected as the subject (on the Search page).
Packets	The number of packets used by the hosts selected as the subject (on the Search page).
Peer Byte Ratio	The ratio, in percent, of peer bytes to the bytes used by the hosts selected as the subject (on the Search page). The percentage is displayed next to a graphic (for which the color never changes) that illustrates the percentage.
Peer Bytes	The number of bytes used by hosts selected as the peer (on the Search page).
Peer Client Bytes	The number of bytes used by client hosts among the peer hosts (as selected on the Search page).
Peer Client Flows	The number of flows used by client hosts among the peer hosts (as selected on the Search page).
Peer Client Packets	The number of packets used by client hosts among the peer hosts (as selected on the Search page).
Peer Client TCP Connections *	The number of TCP connections (i.e., three-way handshakes, from the first SYN to the last ACK) used by client hosts among the peer hosts (as selected on the Search page).

Column	Description
Peer Clients	The number of client hosts among the peer hosts (as selected on the Search page).
Peer Host Groups	The host group(s) to which the peer host selected as the subject (on the Search page) belongs.
Peer IP Address	The IP address and name (if available) assigned to the top host among the peer hosts (as selected on the Search page). The IP address is displayed if the DNS resolution was not successful.
Peer Name	The name assigned to the peer host selected as the subject (on the Search page), if available.
Peer Packets	The number of packets for the peer hosts (as selected on the Search page).
Peer Role	The role of the peer host. The role can be one of the following: • Client • Server • Client and Server
Peer Server Bytes	The number of bytes used by server hosts among the peer hosts (as selected on the Search page).
Peer Server Flows	The number of flows used by server hosts among the peer hosts (as selected on the Search page).
Peer Server Packets	The number of packets used by server hosts among the peer hosts.
Peer Server TCP Connection *	The number of TCP connections (i.e., three-way handshakes, from the first SYN to the last ACK) used by server hosts among the peer hosts (as selected on

Column	Description
	the Search page).
Peer Servers	The number of peer hosts (as selected on the Search page), acting as servers.
Peers	The number of active hosts among the peer hosts (as selected on the Search page).
Records	The number of flow records sent in the selected direction by the hosts selected as the subject (on the Search page).
TCP Connections *	The number of TCP connections (i.e., three-way handshakes, from the first SYN to the last ACK) completed by the hosts selected as the subject (on the Search page).

* Not available when the document is filtered on an interface or interface group.

Top Ports

Column	Description
% of Bytes	The utilization percentage for the traffic using the corresponding port. The percentage is displayed next to a graphic (for which the color never changes) that illustrates the percentage. If you chose to filter the report by bytes, this column will contain data.
95th Packet Rate (pps)	The highest rate of packet transmission (in packets per second) observed on the corresponding date, not including the top 5% of the peak rates. The 95th Percentile is used by firms that perform "burstable billing," where the top 5% of bandwidth

Column	Description
	peaks are cut off. With this method, infrequent peaks are ignored. For more information, refer to: http://en.wikipedia.org/wiki/Burstable_billing
95th Traffic (bps)	The highest rate of traffic (in bits per second) observed on the corresponding date, not including the top 5% of the peak rates. The 95th Percentile is used by firms that perform "burstable billing," where the top 5% of bandwidth peaks are cut off. With this method, infrequent peaks are ignored. For more information, refer to: http://en.wikipedia.org/wiki/Burstable_billing
Average Packet Rate (pps)	The average rate of traffic (in packets per second) observed on the corresponding date.
Average Traffic (bps)	The average rate of traffic (in bits per second) observed on the corresponding date.
Bytes	The number of bytes using the corresponding port.
Flow Collector Name	The name of the Flow Collector that observed the data.
Flows	The number of flows using the corresponding port.
Host Bytes	The number of bytes for the hosts selected as the subject (on the Search page) using the corresponding port.
Host Bytes Ratio	The ratio, in percent, of host bytes to the bytes using the corresponding port. The percentage is displayed next to a bar graph (for which the color never changes) that illustrates the percentage. If you don't select a host, the displayed entry is 50%

Column	Description
	(since application traffic consists of both hosts and peers.)
Host Client Bytes	The number of bytes used by client hosts selected as the subject (on the Search page) using the corresponding port.
Host Client Flows	The number of flows used by client hosts selected as the subject (on the Search page) using the corresponding port.
Host Client Packets	The number of packets used by client hosts selected as the subject (on the Search page) using the corresponding port.
Host Client TCP Connections *	The number of TCP connections (i.e., three-way handshakes, from the first SYN to the last ACK) used by client hosts selected as the subject (on the Search page) using the corresponding port.
Host Clients	The number of client hosts selected as the subject (on the Search page) using the corresponding port.
Host Packets	The number of packets for the hosts selected as the subject (on the Search page) using the corresponding port.
Host Role	The role of the hosts selected as the subject (on the Search page) using the corresponding port. The role can be one of the following: • Client • Server • Either
Host Server Bytes	The number of bytes used by server hosts selected as

Column	Description
	the subject (on the Search page) using the corresponding port.
Host Server Flows	The number of flows used by server hosts selected as the subject (on the Search page) using the corresponding port.
Host Server Packets	The number of packets used by server hosts selected as the subject (on the Search page) using the corresponding port.
Host Sever TCP Connections *	The number of TCP connections (i.e., three-way handshakes, from the first SYN to the last ACK) used by server hosts selected as the subject (on the Search page) using the corresponding port.
Host Servers	The number of hosts selected as the subject (on the Search page), acting as servers, using the corresponding port.
Hosts	The number of hosts selected as the subject (on the Search page). These are usually the hosts on your network's side of the conversation.
Maximum Packet Rate (Pps)	The peak rate of traffic (in packets per second) using the corresponding port.
Maximum Traffic (Bps)	The peak rate of traffic (in bits per second) using the corresponding port.
Minimum Packet Rate (Pps)	The smallest rate of traffic (in packets per second) using the corresponding port.
Minimum Traffic (Bps)	The smallest rate of traffic (in bits per second) using the corresponding port.

Column	Description
Packets	The number of packets using the corresponding port.
Peer Byte Ratio	The ratio, in percent, of peer bytes to the bytes using the corresponding port. The percentage is displayed next to a graphic (for which the color never changes) that illustrates the percentage.
Peer Bytes	The number of bytes used by hosts selected as the peer (on the Search page) using the corresponding port.
Peer Client Bytes	The number of bytes used by client hosts among the peer hosts (as selected on the Search page) using the corresponding port.
Peer Client Flows	The number of flows used by client hosts among the peer hosts (as selected on the Search page) using the corresponding port.
Peer Client Packets	The number of packets used by client hosts among the peer hosts (as selected on the Search page) using the corresponding port.
Peer Client TCP Connections *	The number of TCP connections (i.e., three-way handshakes, from the first SYN to the last ACK) used by client hosts among the peer hosts (as selected on the Search page) using the corresponding port.
Peer Clients	The number of client hosts among the peer hosts (as selected on the Search page) using the corresponding port.
Peer Packets	The number of packets for the peer hosts (as selected on the Search page) using the corresponding port.
Peer Role	The role of the peer host using the corresponding

Column	Description
	port. The role can be one of the following: • Client • Server • Client and Server
Peer Server Bytes	The number of bytes used by server hosts among the peer hosts (as selected on the Search page) using the corresponding port.
Peer Server Flows	The number of flows used by server hosts among the peer hosts (as selected on the Search page) using the corresponding port.
Peer Server Packets	The number of packets used by server hosts among the peer hosts using the corresponding port.
Peer Server TCP Connection *	The number of TCP connections (i.e., three-way handshakes, from the first SYN to the last ACK) used by server hosts among the peer hosts (as selected on the Search page) using the corresponding port.
Peer Servers	The number of peer hosts (as selected on the Search page), acting as servers, using the corresponding port.
Peers	The number of peer hosts (as selected on the Search page) associated with the corresponding port. These are usually the hosts on the other side of the conversation.
Port	The port that the corresponding flow used.
Records	The number of flow records sent through the corresponding port in the selected direction.

Column	Description
TCP Connections *	The number of TCP connections (i.e., three-way handshakes, from the first SYN to the last ACK) sent through the corresponding port.

* Not available when the document is filtered on an interface or interface group .

Top Protocols

Column	Description
% of Bytes	The utilization percentage for the traffic using the corresponding protocol. The percentage is displayed next to a graphic (for which the color never changes) that illustrates the percentage. If you chose to filter the report by bytes, this column will contain data.
95th Packet Rate (pps)	The highest rate of packet transmission (in packets per second) observed on the corresponding date, not including the top 5% of the peak rates. The 95th Percentile is used by firms that perform "burstable billing," where the top 5% of bandwidth peaks are cut off. With this method, infrequent peaks are ignored. For more information, refer to: http://en.wikipedia.org/wiki/Burstable_billing
95th Traffic (bps)	The highest rate of traffic (in bits per second) observed on the corresponding date, not including the top 5% of the peak rates. The 95th Percentile is used by firms that perform "burstable billing," where the top 5% of bandwidth peaks are cut off. With this method, infrequent peaks are ignored. For more information, refer to:

Column	Description
	http://en.wikipedia.org/wiki/Burstable_billing
Average Packet Rate (pps)	The average rate of traffic (in packets per second) observed on the corresponding date.
Average Traffic (bps)	The average rate of traffic (in bits per second) observed on the corresponding date.
Bytes	The number of bytes using the corresponding protocol.
Flow Collector Name	The name of the Flow Collector that observed the data.
Flows	The number of flows using the corresponding protocol.
Host Bytes	The number of bytes for the hosts selected as the subject (on the Search page) using the corresponding protocol.
Host Bytes Ratio	The ratio, in percent, of host bytes to the bytes using the corresponding protocol. The percentage is displayed next to a bar graph (for which the color never changes) that illustrates the percentage. If you don't select a host, the displayed entry is 50% (since application traffic consists of both hosts and peers.)
Host Client Bytes	The number of bytes used by client hosts selected as the subject (on the Search page) using the corresponding protocol.
Host Client Flows	The number of flows used by client hosts selected as the subject (on the Search page) using the corresponding protocol.

Column	Description
Host Client Packets	The number of packets used by client hosts selected as the subject (on the Search page) using the corresponding protocol.
Host Client TCP Connections *	The number of TCP connections (i.e., three-way handshakes, from the first SYN to the last ACK) used by client hosts selected as the subject (on the Search page) using the corresponding protocol.
Host Clients	The number of client hosts selected as the subject (on the Search page) using the corresponding protocol.
Host Packets	The number of packets for the hosts selected as the subject (on the Search page) using the corresponding protocol.
Host Role	The role of the hosts selected as the subject (on the Search page) using the corresponding protocol. The role can be one of the following: • Client • Server • Either
Host Server Bytes	The number of bytes used by server hosts selected as the subject (on the Search page) using the corresponding protocol.
Host Server Flows	The number of flows used by server hosts selected as the subject (on the Search page) using the corresponding protocol.
Host Server Packets	The number of packets used by server hosts selected as the subject (on the Search page) using the corresponding protocol.

Column	Description
Host Sever TCP Connections *	The number of TCP connections (i.e., three-way handshakes, from the first SYN to the last ACK) used by server hosts selected as the subject (on the Search page) using the corresponding protocol.
Host Servers	The number of hosts selected as the subject (on the Search page), acting as servers, using the corresponding protocol.
Hosts	The number of hosts selected as the subject (on the Search page). These are usually the hosts on your network's side of the conversation.
Maximum Packet Rate (Pps)	The peak rate of traffic (in packets per second) using the corresponding protocol.
Maximum Traffic (Bps)	The peak rate of traffic (in bits per second) using the corresponding protocol.
Minimum Packet Rate (Pps)	The smallest rate of traffic (in packets per second) using the corresponding protocol.
Minimum Traffic (Bps)	The smallest rate of traffic (in bits per second) using the corresponding protocol.
Packets	The number of packets used by the corresponding application.
Peer Byte Ratio	The ratio, in percent, of peer bytes to the bytes using the corresponding protocol. The percentage is displayed next to a graphic (for which the color never changes) that illustrates the percentage.
Peer Bytes	The number of bytes used by hosts selected as the peer (on the Search page) using the corresponding protocol.

Column	Description
Peer Client Bytes	The number of bytes used by client hosts among the peer hosts (as selected on the Search page) using the corresponding protocol.
Peer Client Flows	The number of flows used by client hosts among the peer hosts (as selected on the Search page) using the corresponding protocol.
Peer Client Packets	The number of packets used by client hosts among the peer hosts (as selected on the Search page) using the corresponding protocol.
Peer Client TCP Connections *	The number of TCP connections (i.e., three-way handshakes, from the first SYN to the last ACK) used by client hosts among the peer hosts (as selected on the Search page) using the corresponding protocol.
Peer Clients	The number of client hosts among the peer hosts (as selected on the Search page) using the corresponding protocol.
Peer Packets	The number of packets for the peer hosts (as selected on the Search page) using the corresponding protocol.
Peer Role	The role of the peer host using the corresponding protocol. The role can be one of the following: • Client • Server • Client and Server
Peer Server Bytes	The number of bytes used by server hosts among the peer hosts (as selected on the Search page) using the corresponding protocol.

Column	Description
Peer Server Flows	The number of flows used by server hosts among the peer hosts (as selected on the Search page) using the corresponding protocol.
Peer Server Packets	The number of packets used by server hosts among the peer hosts using the corresponding protocol.
Peer Server TCP Connection *	The number of TCP connections (i.e., three-way handshakes, from the first SYN to the last ACK) used by server hosts among the peer hosts (as selected on the Search page) using the corresponding protocol.
Peer Servers	The number of peer hosts (as selected on the Search page), acting as servers, using the corresponding protocol.
Peers	The number of peer hosts (as selected on the Search page) using the corresponding protocol. These are usually the hosts on the other side of the conversation.
Protocol	The protocol that the corresponding flow used.
Records	The number of flow records sent in the selected direction and using the corresponding protocol.
TCP Connections *	The number of TCP connections (i.e., three-way handshakes, from the first SYN to the last ACK) using the corresponding protocol.

* Not available when the document is filtered on an interface or interface group .

Top Services

Column	Description
% of Bytes	The utilization percentage for the traffic using the

Column	Description
	corresponding service. The percentage is displayed next to a graphic (for which the color never changes) that illustrates the percentage. If you chose to filter the report by bytes, this column will contain data.
95th Packet Rate (pps)	The highest rate of packet transmission (in packets per second) observed on the corresponding date, not including the top 5% of the peak rates. The 95th Percentile is used by firms that perform "burstable billing," where the top 5% of bandwidth peaks are cut off. With this method, infrequent peaks are ignored. For more information, refer to: http://en.wikipedia.org/wiki/Burstable_billing
95th Traffic (bps)	The highest rate of traffic (in bits per second) observed on the corresponding date, not including the top 5% of the peak rates. The 95th Percentile is used by firms that perform "burstable billing," where the top 5% of bandwidth peaks are cut off. With this method, infrequent peaks are ignored. For more information, refer to: http://en.wikipedia.org/wiki/Burstable_billing
Average Packet Rate (pps)	The average rate of traffic (in packets per second) observed on the corresponding date.
Average Traffic (bps)	The average rate of traffic (in bits per second) observed on the corresponding date.
Bytes	The number of bytes using the corresponding service.
Flow Collector Name	The name of the Flow Collector that observed the data.

Column	Description
Flows	The number of flows using the corresponding service.
Host Bytes	The number of bytes for the hosts selected as the subject (on the Search page) using the corresponding service.
Host Bytes Ratio	The ratio, in percent, of host bytes to the bytes using the corresponding service. The percentage is displayed next to a bar graph (for which the color never changes) that illustrates the percentage. If you don't select a host, the displayed entry is 50% (since application traffic consists of both hosts and peers.)
Host Client Bytes	The number of bytes used by client hosts selected as the subject (on the Search page) using the corresponding service.
Host Client Flows	The number of flows used by client hosts selected as the subject (on the Search page) using the corresponding service.
Host Client Packets	The number of packets used by client hosts selected as the subject (on the Search page) using the corresponding service.
Host Client TCP Connections *	The number of TCP connections (i.e., three-way handshakes, from the first SYN to the last ACK) used by client hosts selected as the subject (on the Search page) using the corresponding service.
Host Clients	The number of client hosts selected as the subject (on the Search page) using the corresponding service.
Host Packets	The number of packets for the hosts selected as the subject (on the Search page) using the corresponding

Column	Description
	service.
Host Role	The role of the hosts selected as the subject (on the Search page) using the corresponding service. The role can be one of the following: • Client • Server • Either
Host Server Bytes	The number of bytes used by server hosts selected as the subject (on the Search page) using the corresponding service.
Host Server Flows	The number of flows used by server hosts selected as the subject (on the Search page) using the corresponding service.
Host Server Packets	The number of packets used by server hosts selected as the subject (on the Search page) using the corresponding service.
Host Sever TCP Connections *	The number of TCP connections (i.e., three-way handshakes, from the first SYN to the last ACK) used by server hosts selected as the subject (on the Search page) using the corresponding service.
Host Servers	The number of hosts selected as the subject (on the Search page), acting as servers, using the corresponding service.
Hosts	The number of hosts selected as the subject (on the Search page). These are usually the hosts on your network's side of the conversation.
Maximum Packet Rate (Pps)	The peak rate of traffic (in packets per second) using the corresponding service.

Column	Description
Maximum Traffic (Bps)	The peak rate of traffic (in bits per second) using the corresponding service.
Minimum Packet Rate (Pps)	The smallest rate of traffic (in packets per second) using the corresponding service.
Minimum Traffic (Bps)	The smallest rate of traffic (in bits per second) using the corresponding service.
Packets	The number of packets using the corresponding service.
Peer Byte Ratio	The ratio, in percent, of peer bytes to the bytes using the corresponding service. The percentage is displayed next to a graphic (for which the color never changes) that illustrates the percentage.
Peer Bytes	The number of bytes used by hosts selected as the peer (on the Search page) using the corresponding service.
Peer Client Bytes	The number of bytes used by client hosts among the peer hosts (as selected on the Search page) using the corresponding service.
Peer Client Flows	The number of flows used by client hosts among the peer hosts (as selected on the Search page) using the corresponding service.
Peer Client Packets	The number of packets used by client hosts among the peer hosts (as selected on the Search page) using the corresponding service.
Peer Client TCP Connections *	The number of TCP connections (i.e., three-way handshakes, from the first SYN to the last ACK) used by client hosts among the peer hosts (as selected on the Search page) using the corresponding service.

Column	Description
Peer Clients	The number of client hosts among the peer hosts (as selected on the Search page) using the corresponding service.
Peer Packets	The number of packets for the peer hosts (as selected on the Search page) using the corresponding service.
Peer Role	The role of the peer host using the corresponding service. The role can be one of the following: • Client • Server • Client and Server
Peer Server Bytes	The number of bytes used by server hosts among the peer hosts (as selected on the Search page) using the corresponding service.
Peer Server Flows	The number of flows used by server hosts among the peer hosts (as selected on the Search page) using the corresponding service.
Peer Server Packets	The number of packets used by server hosts among the peer hosts using the corresponding service.
Peer Server TCP Connection *	The number of TCP connections (i.e., three-way handshakes, from the first SYN to the last ACK) used by server hosts among the peer hosts (as selected on the Search page) using the corresponding service.
Peer Servers	The number of peer hosts (as selected on the Search page), acting as servers, using the corresponding service.
Peers	The number of peer hosts (as selected on the Search page) associated with the corresponding service. These are usually the hosts on the other side of the

Column	Description
	conversation.
Records	The number of flow records sent in the selected direction associated with the corresponding service.
Service	The service that the corresponding flow used.
TCP Connections *	The number of TCP connections (i.e., three-way handshakes, from the first SYN to the last ACK) completed by the corresponding service.

* Not available when the document is filtered on an interface or interface group.

What you can do next

Start a new top report search

1. Click **Start New Search**.

The Top [Report Name] Search page opens.

2. If necessary, edit the desired criteria and then click **Search**. The Top [Report Name] Search Results page opens. When a job is in progress, a banner at the top of the page displays the message *Your query is in progress*. A blue *In Progress* pill is also displayed. While the job is running, the message *In Progress* is displayed on the Jobs > [Job Management](#) page in the Current Jobs table. When the job has finished, a blue *Completed* pill is displayed on the Top [Report Name] Search Results page, and on the Job Management page, the job is displayed in the Finished Jobs table with a status of *Completed* (only if there were results).

Save the results

Click **Save Results**.

You can save multiple top report searches with the same name, or you can give each one a unique name. You can use any special character in the name. Saved results are displayed on the Analyze > [Saved Results](#) page. Any saved results are retained until you delete them.

Save the top report search

Click **Save**.

You can save multiple top report searches with the same name, or you can give each one a unique name. You can use any special character in the name. Saved searches are displayed on the Analyze > [Saved Searches](#) page.

Run a saved search

1. Click **Start New Search**.

The Top [Report Name] Search page opens.

2. Click **Load Saved Search > Manage Saved Searches** and select the saved search you want to run.

Delete results

1. Access the Job Management page.
2. Click the menu in the Actions Column.
3. Click **Delete**.

Related Topics

[Accessing Saved Flow Search Results](#)

[Accessing Saved Flow Searches](#)

[Managing Jobs](#)

Determining the Existence of a Host

Use this page to determine if a specific IP address has ever been seen on your network since the installation of Secure Network Analytics on your network. This functionality is available beginning in v6.7

If the Host Aggregate Traffic table on the Manager has not received data from the flow aggregator service on the Flow Collector for more than one hour, a warning note is displayed denoting this. This note will also display the date and time the First Sent and Last Sent information was last updated. You will be requested to contact the [Cisco Support](#) team to restart the service.

How to find this page

Do the following to access this page

1. On any page in the Web App, from the main menu, choose **Investigate > Host Search**.
2. In the IP Address field, enter the IP address(es) for which you want to search.
3. Click **Search**. You can search for a maximum of 256 IP addresses per search and can perform multiple searches if necessary. Refer to [Entering IP Addresses](#) for acceptable conventions to use for entry of multiple IP addresses.

IPv6 IP addresses are not supported.

The results are displayed in the Search Results table. If you perform a search on an IP address for which there is currently an entry in the Search Results table, that entry will be updated; no additional entry will be created. All results are retained for 14 days.

What you can do on this page

View data

On this page the Search Results table is displayed, which contains the IP addresses for which searches have been performed. If no results are found, *Never* and *None* are populated in the applicable columns in the row.

View the data in the following fields within the table:

- Time of Search—The time the search was performed.
- Search Subject—The IP address that was searched.

- **First Sent (Start Date)** –The date and time the IP address first communicated on the network.
- **Last Sent (End Date)** –The date and time the IP address last communicated on the network.
- **Total Bytes**–The total number of bytes sent and received from the IP address.
 - The number in this field is not an exact calculation of the total bytes for the flow; it is an approximate value that is meant to be used to determine the relative size of bytes sent in the flow.
- **Flow Collectors**–The Flow Collectors which have seen the IP address. This column will be blank if a particular Flow Collector is no longer available.

What you can do next

Run a flow search

- To run a flow search associated with a single IP address, click the menu next to the applicable IP address to access the relevant options.
- To run a flow search associated with a row of data, click the menu in the Actions column to access the relevant options.

Run a top report search

- To [run a top report](#) search associated with a single IP address, click the menu next to the applicable IP address to access the relevant options.
- To run a top report search associated with a row of data, click the menu in the Actions column to access the relevant options.

Investigate the activity of a specific host

Investigate the activity of a specific host, To do this, click that host's IP address. The [Host Report](#) for that host opens.

Search for another IP address or addresses

Search for another IP address or addresses. To perform another search, type your entry in the IP Address field and click **Search**.

If the Host Aggregate Traffic table on the Manager has not received data from the flow aggregator service on the Flow Collector for more than one hour, this could indicate that the service to populate data is down.

Related topics

Investigating a Single Host

Investigating Copyright Infringement

Use this page to identify the host IP address that was used to download or distribute copyrighted material, such as documents or videos.

You can perform three tasks with this page:

- Build and run a copyright infringement query.
- View the results of a query. The results of these queries provide you with information to assist your investigation and allow you to take corrective action.
- View a list of previously run queries. The information is retained for 14 days and is then automatically deleted.

How to find this page

From the main menu, choose **Investigate > Copyright Infringement**.

What you can do on this page

Build a copyright infringement query

Build a copyright infringement query by completing the following steps.

1. Obtain the information you need to enter into the applicable fields from the third-party notification sent to your company that a copyright infringement has occurred. The table below describes each item.

Field Name	Description
IP Address	The IP address shown on the copyright infringement notice. This is the Network Address Translation (NAT) IP address of your company's firewall or proxy server, which is used by internal hosts to access the internet.
Port	The port number associated with the NAT version of the IP address; this is also provided on the copyright infringement notice.
Date & Time	The date and time shown on the copyright infringement notice (ensure you enter the time

Field Name	Description
	zone indicated in the notice) that indicates when the download occurred. The query searches for all flows that have occurred starting from five minutes before the date and time entered here to five minutes afterward.
Time Zone	Enter the time zone as indicated on the copyright infringement notice.

2. After defining the settings, click **Search** to run the query. A dialog opens displaying the query result(s).

The query result(s) provides the following information:

Field Name	Description
Search Subject and Subject Port	The host and port number of your company's firewall or proxy server. They should be the same as the ones shown in the copyright infringement notice. You can initiate a packet query by clicking the ... (Ellipsis) icon that is displayed beside the Search Subject IP address.
Date & Time	The date and time specified for the query. This time is given in the time zone that your browser is set to.
Inside Host and Inside Port	The internal host and port number that originated the download request. You can initiate a packet query by clicking the Ellipsis icon that is displayed beside the Inside Host IP address.
Outside Host and Outside Port	The external host and port number where the infringement information came from.

Field Name	Description
	You can initiate a packet query by clicking the Ellipsis icon that is displayed beside the Outside Host IP address.
Application	The application that the flow was carrying as determined by the application definitions configured in the Desktop Client.
Total Bytes	The total number of bytes that were transferred in the flow.

3. In the row containing the record you want to view, click **Select**. The dialog closes and the query result is displayed in the Search Results table on the Copyright Infringement page.

Only one record returned from a query can be displayed at one time in the Search Results table on the Copyright Infringement page.

4. If multiple results were returned and you want to view another record, do the following:
 - a. In the row for the applicable query, click the ellipsis in the Actions column and then click **Other client IPs**.
The dialog displaying multiple records opens.
 - b. Repeat step 3.

What you can view on this page

View the following information on this page:

- Review the record returned by the query. Do the translated host/port, date and time, and total bytes match the corresponding information on the copyright infringement notice? If so, refer to the suggested steps in the next section.
- If a query returns a record that shows the client host, server host, and port information is undefined or missing, then in the Search Results table on the Copyright Infringement page the following results occur:
 - The Inside Host entry will be *Not Found*.
 - The Inside Port, Outside Host, Outside Port, and Application fields will be

blank.

- The Total Bytes entry will be *0B*.

What you can do after results have been returned

Do the following as necessary:

1. Conduct a detailed investigation by Running a flow search for flows that occurred within the 10-minute time range (that was specified in the Date & Time field) for the inside Host and Inside Port. To do this, in the Actions column, click the ellipsis to open the menu and click **Flows**.
2. View other records (if multiple records were returned). To do this, in the row for the applicable query, click the ellipsis in the Actions column and then click **Other Client IPs**.
3. After you have conducted your investigation and, if necessary, warned the offending user, follow your company's policy about notifying the third-party monitoring organization about the outcome.
4. If you no longer need to retain the search results, delete them. To do this, click the menu in the Actions column and click **Delete**.

Cisco Security Analytics and Logging (On Premises) Overview

Security Analytics and Logging (OnPrem) allows you to store and view data from another Cisco product deployment, such as a Cisco Secure Firewall (formerly Firepower) appliance deployment.

In the case of the Secure Firewall deployment, you can export your Secure Firewall Security Events and data plane events from your Cisco Secure Firewall Threat Defense (formerly Firepower Threat Defense) devices, whether managed by a Cisco Secure Firewall Management Center (formerly Firepower Management Center) or other platforms, to Security Analytics and Logging (OnPrem) to store that information. Security Analytics and Logging (OnPrem) also supports events that can be exported from your ASA and NGIPS devices.

If you are on a domain with a Manager and Flow Collector (Non-Data Store domain), you will receive an error indicating no data returned.



Security Analytics and Logging (OnPrem) is only supported on Data Store (Manager, Flow Collector, and Data Node) deployments. Refer to the [Cisco Security Analytics and Logging \(On Premises\): Firewall Event Integration Guide](#) for more information.

Opening Security Analytics and Logging (OnPrem)

- From the main menu, choose **Investigate > Security Analytics and Logging (OnPrem)**.
- Cross-launch from the Cisco Secure Firewall Management Center.
For more information on cross-launch, refer to the *Investigate Events Using Cross-launch* section in the [Cisco Security Analytics and Logging \(On Premises\): Firewall Event Integration Guide](#).

Configuring the Deployment

Refer to the [Cisco Security Analytics and Logging \(On Premises\): Firewall Event Integration Guide](#).

Security Analytics and Logging (OnPrem): Use the Event Viewer

The Event Graph and the table below it provide information about connection, security, intrusion, file, and malware events. The Event Graph is a representation of the data which displays in the table.

Editing

To edit the events which display, do the following:

- Click the **Time range** drop-down menu to select a predefined range, or select **Custom** to choose the **Start date** and **End date**.
- Check the **Security Events Only** check-box to include only events related to a possible security issue
- Click **Apply Filters to Total** to update the Event Graph and Table with the selected parameters.
- Click **Cancel** to stop running the query.
- Click **Refresh Time** to update the **End date** to the current time.



Make sure you have at least one predefined time range before using the **Refresh Time** option.

Refining

If you'd like to further refine how the Event Graph and table display, use the following options:

- **Columns:** To select columns to display or suppress, click **Columns** in the table. The Show/Hide Columns panel opens.
- **Summary:** To view the summary of the values for the following fields returned in the results for a flow search, click **Summary** in the table.
- **Export (CSV):** If you need to share the Event results with others in your organization, click **Export (CSV)** in the table to export the first 10,000 events to a comma-separated values (.CSV) file.
- **Icons:**
 - To minimize the Event Graph, click the  (Expand) icon in the table.
 - To show the Event Graph, click the  (Expand) icon.

- To expand the event information, click the ► (**Triangle Right**) icon for the event.
- To filter based on data in the table, click the ⋮ (**Ellipsis**) icon when hovering over the applicable field.

When you filter data in the table, the graph also updates. If you highlight areas of the Event Graph to view specific events and choose a **Start date** and **End date**, the graph and the table populate automatically.

Initially, you'll see the default fields used to filter and display event results. Make sure to use the **Columns** option if you'd like to show or hide fields from the list.



Note: The **Event Type** field is unique in that you can either choose from the drop-down to make a selection, or you can enter a specific Event ID (for ASA events).

Related Tasks

[Security Analytics and Logging \(OnPrem\): Create a View](#)

[Security Analytics and Logging \(OnPrem\): Update a View](#)

[Security Analytics and Logging \(OnPrem\): Delete a View](#)

Related Topics

[Cisco Security Analytics and Logging \(On Premises\) Overview](#)

[Security Analytics and Logging \(OnPrem\): Resources](#)

Security Analytics and Logging (OnPrem): Create a View

Use the following instructions to create a new view with your selected options for the Event Viewer.



Custom time ranges will not be saved in a view due to event storage limits. The **Time range** option will default to **Last Hour**.

1. Click the **Views** drop-down.
2. Select **Create a New View**.
3. Enter a name for the view.
4. Click **Enter**.

Related Tasks

[Security Analytics and Logging \(OnPrem\): Use the Event Viewer](#)

[Security Analytics and Logging \(OnPrem\): Update a View](#)

[Security Analytics and Logging \(OnPrem\): Delete a View](#)

Related Topics

[Cisco Security Analytics and Logging \(On Premises\) Overview](#)

[Security Analytics and Logging \(OnPrem\): Resources](#)

Security Analytics and Logging (OnPrem): Update a View

Use the following instructions to update a saved view.

1. Click the **Views** drop-down menu.
2. Select a saved view.
3. Edit the options for the Event Viewer.

To edit the events shown, refer to [Security Analytics and Logging \(OnPrem\): Use the Event Viewer](#).



Custom time ranges will not be saved in a view due to event storage limits. The **Time range** option will default to **Last Hour**.

4. Click **Update View** to save your changes.

To discard your changes, click **Clear Changes**.

Related Tasks

[Security Analytics and Logging \(OnPrem\): Use the Event Viewer](#)

[Security Analytics and Logging \(OnPrem\): Create a View](#)

[Security Analytics and Logging \(OnPrem\): Delete a View](#)

Related Topics

[Cisco Security Analytics and Logging \(On Premises\) Overview](#)

[Security Analytics and Logging \(OnPrem\): Resources](#)

Security Analytics and Logging (OnPrem): Delete a View

Use the following instructions to delete a saved view.

1. Click the **Views** drop-down.
2. Click the  (**Delete**) icon next to the view you want to delete.

Related Tasks

[Security Analytics and Logging \(OnPrem\): Use the Event Viewer](#)

[Security Analytics and Logging \(OnPrem\): Create a View](#)

[Security Analytics and Logging \(OnPrem\): Update a View](#)

Related Topics

[Cisco Security Analytics and Logging \(On Premises\) Overview](#)

[Security Analytics and Logging \(OnPrem\): Resources](#)

Security Analytics and Logging (OnPrem): Resources

For additional information, refer to the following:

- [Getting Started with Security Analytics and Logging \(OnPrem\)](#)
- [Security Analytics and Logging \(OnPrem\): Firewall Event Integration Guide](#)

Related Topics

[Cisco Security Analytics and Logging \(On Premises\) Overview](#)

Managing Jobs

Use this page to view, cancel, or delete search results such as flow search, ANC queries, packet queries, etc.

Depending on the specific information you are looking for, you can view any of the following

- Name of the person who created a job
- Job status
- Dates and times a job started and ended

How to find this page

Do the following to access this page

From the main menu, choose **Investigate > Search Management > Job Management**.

What you can do on this page

Perform tasks common to all jobs

- View details (e.g., job owner, start time) about a job.
- When a job completes, it is moved to the Finished Jobs table.
- Cancel a job that is pending or in progress.
 - If you cancel a pending job, it is displayed in the Finished Jobs table with a status of *Canceled*.
 - If you cancel a job that is in progress, it is displayed in the Finished Jobs table with a status of *Canceled*.
- Open a completed or canceled job from the Finished Jobs table. The job is only available for a limited amount of time. Refer to the Data Retention section below for more information.
- Delete a job that has completed or that has been canceled. You may need to delete completed jobs to make space in the database for new jobs. Refer to the Data Retention section below for more information.
 - If you delete a job for which you have already saved results, you will not be able to access those saved results.

- Before deleting a job that has saved results, you will see a confirmation dialog. Click **No** to retain the saved results. The job will continue to be listed in the Finished Jobs table, and you can then select another job to delete.
- For all jobs that can be exported:
 - When you export a job, UTC is used for time-related data.

Manage flow search jobs

- If you run Flow Table searches on the Desktop Client, those are not listed on this page in the Manager.
- Open a flow search results job that is in progress from the Current Jobs table. When you do this, you will see partial results.
- The only jobs you can open at present are flow search results. You may see other types of jobs listed, such as the completed [network classification](#) query, but you cannot open those completed jobs. Depending on the type of job, it may download a separate file.
- When a job is in progress, a banner at the top of the page displays the message *Your query is in progress*. A blue *In Progress* pill is also displayed. While the job is running, the message *In Progress* is displayed on the Jobs > [Job Management](#) page in the Current Jobs table. When the job has finished, a blue *Completed* pill is displayed on the Flow Search Results page, and on the Job Management page, the job is displayed in the Finished Jobs table with a status of *Completed* (only if there were results).
- Note that the search uses approximately 10,000 flow records per Flow Collector as the default. When the default number of flow records has been returned, the search may stop, no matter how much time has elapsed relative to the time interval you set.
- On the Flow Search page, you can specify a time interval for the data you want to be retrieved. Remember that when the default number of 10,000 flow records has been returned, the search may stop, no matter how much time has elapsed relative to the time interval you set.
- If you are conducting a forensic investigation, and the time interval is an important part of your investigation, then you should use the Flow Search page to fine tune settings, so the search returns the data that is most relevant. You may need to set up a series of sequential searches to make sure the results cover the time span you need.

- Other types of searches, such as displaying security events from the Security Insight dashboard or listing hosts or users, may not run if the maximum number of flow searches are in progress.
- When a search completes, it is moved to the Finished Jobs table; the next pending search begins running, and its status changes to In Progress.
- In both Data Store and Non-Data Store deployments, you can cancel a flow search that has a status of *Pending* or *In Progress*. If you cancel a flow search that is in progress, you can still view the partial results that were returned. On the Job Management page, click the name of the job to open the Flow Results page, which displays the results that were returned before you canceled the job.
- If you included certain criteria, marked by an  (**Information**) icon on the page, in your flow search, or run a flow search with a Max Records Returned of 20,000 or more (whose results can only be viewed by downloading them to a CSV file), you cannot cancel a flow search before it is complete or see what percentage of the job has run thus far. However, when the job has finished, a blue *Completed* pill is displayed on the Flow Search Results page.
- For a flow search request with a Max Records Returned of 10,000 or less, both the search and the results are retained on the Job Management page for 48 hours from the time the job completed.
- For a flow search request with a Max Records Returned of 20,000 or more (whose results can only be viewed by downloading them to a .CSV file), both the search and the results are retained on the Job Management page for 7 days from the time the job completed.
- You can perform a flow search for up to 400,000 flows in one search.
- Flow searches with a Max Records Returned of 10,000 or less, as well as all top report jobs, run in a different queue from that of flow searches with a Max Records Returned of 20,000 or more.
 - You can run a maximum of four jobs at the same time if the Max Records Returned for the flow search is 10,000 or less. It does not matter what the Records Returned value is for a top report job.
 - You can run only one flow search at a time if the Max Records Returned is 20,000 or more.
 - If you start additional top report jobs or flow searches during this time (for either of these two categories of flow searches), they will be placed in their respective queue with a Pending status.

- The flow search job naming conventions used for the downloadable .CSV zipped and unzipped files are as follows:
 - flowAnalysis-[date]-[time].csv,zip, with the time (in current local time) reflecting when the download was initiated.
 - flowAnalysis-[date]-[time].csv, with the time (in UTC) reflecting when the file was created on the server.

The completion time is shown in the Finished Jobs table in the End Date/Time column. Both the jobs and the results are then deleted automatically.

You cannot change the retention periods. Saved jobs or results, however, are not deleted automatically.

- Possible job statuses are In Progress, Completed, Pending, Canceled, and Failed.
- If you must delete a job that has saved results from a flow search, but you want to use those results in the future, you can export them to a .CSV file. To do so, access the Saved Results page and click the name of the flow search results. On the Flows Search Results page, click **Export**.

Manage top report jobs

- All top report jobs, as well as flow searches with a Max Records Returned of 10,000 or less, run in a different queue from that of flow searches with a Max Records Returned of 20,000 or more.
 - You can run a maximum of four jobs at the same time if the Max Records Returned for the flow search is 10,000 or less. It does not matter what the Records Returned value is for a top report job.
 - If you start additional top report jobs or flow searches during this time (for either of these two categories of flow searches), they will be placed in their respective queue with a Pending status.
- In both Data Store and Non-Data Store deployments, you can cancel a top report search that has a status of *Pending* or *In Progress*.
- Completed jobs are listed in the Finished Jobs table; the searches and their results are retained for 48 hours from the time the job completed. The completion time is shown in the Finished Jobs table in the End Date/Time column. Both the jobs and the results are then deleted automatically after 48 hours. You cannot change the 48-hour retention period.
- When a search completes, it is moved to the Finished Jobs table; the next pending job begins running, and its status changes to In Progress.

- Other types of jobs, such as displaying security events from the Security Insight dashboard or listing hosts or users, may not run if the maximum number of top reports are in progress.
- Possible job statuses are In Progress, Completed, Pending, and Failed.

Respond to failed jobs

- In the Finished Jobs table on this page, if the Status column displays Failed, then there may be a problem with your system's database that you need to correct; examples include a lost connection or a database that is full.
- If you see the Failed status and a banner on this page advising you that the database is nearly full, then you cannot run any additional queries until you delete one or more jobs listed in the Finished Jobs table.
- If the job data is important, then you need to save the search and/or results or export the results to a .CSV file.

Respond to data retention rules

- If you want to reuse flow searches or refer to their results again, click the name in the Job Name column to open the [Flow Search Results](#) page. From there, you can [save the search](#), [save the results](#), or export the results to a .CSV file.
- You cannot open Assisted Network Classification (ANC) queries from the Job Management page. To view those results, from the main menu, click **Configure > DETECTION Network Scanners** to open the [Network Scanners](#) page. If the ANC query produced results, they will be displayed on that page.

What you can do next

Do the following as necessary

- To view available results of a job in progress, click the name of the job in the Current Jobs table.
- To view the results of a completed job, click the appropriate name in the Job Name column in the Finished Jobs table.
- To view results that were returned before a job was canceled, look for Canceled in the Status column in the Finished Jobs table. Then click the corresponding name in the Job Name column.

Related Topics

Classifying Hosts Inside the Network

Flow Search Results: Overview

Viewing Packet Query Results

Accessing Saved Flow Searches

Use this page to view a list of saved flow searches for the Manager. You can choose a search to edit, delete, run, or clone (i.e., copy).

How to find this page

From the main menu, choose **investigate > Search Management > Saved Searches**.

The Saved Searches page opens.

What you can do on this page

You can perform the following actions on this page.

To...	Click...
Review, run, edit, or save the flow or top report search under a different name	The applicable name in the Search Name column.
Modify an existing flow or top report search that you had previously saved	The ⋮ (Ellipsis) icon > Edit in the Actions column.
Permanently remove a flow or top report search you no longer need	The Ellipsis icon > Delete in the Actions column.
Run an existing flow or top report search that you had previously saved (the search remains listed on this page unless you delete it)	The Ellipsis icon > Run in the Actions column.
Copy an existing flow or top report search	The Ellipsis icon > Clone in the Actions column. <i>The flow search page opens, displaying a search with the same settings as the original search. You can change the settings now or edit them later. You will have the option to review the cloned search and save it using the same name or a different name.</i>

Accessing Saved Flow Search Results

Use this page to open and view results of a previous saved flow search. Each row on the page represents a different set of results.

- You will be able to download flow search results to a .CSV file when 20,000 or more records are returned for a query.

How to find this page

From the main menu, choose **Investigate > Search Management > Saved Results**.

The Saved Results page opens.

What you can do on this page

- View saved flow search results from the last 14 days.
- To permanently remove the results for a specific search, in the Actions column, click **Delete**.

Investigating Hosts

Use this page (Host List) to see which inside hosts have been active on your network since the last archive hour, the host group to which they belong, their location, and the categories of alarms the hosts are triggering.

- If you access this page by clicking a host group in the Top Alarming Host Groups component on the Security Insight Dashboard, this page is filtered by that host group.
- If the Host Aggregate Traffic table on the Manager has not received data from the flow aggregator service on the Flow Collector for more than one hour, a warning note is displayed denoting this. This note also displays the date and time the First Sent and Last Sent information was last updated.

 Contact **Cisco Support** to restart the service.

- If you reboot your Flow Collector, it deletes all alarm history. Since the alarming host widgets (which display the number of hosts receiving alarms since the last reset hour for a specific category) on the Security Insight Dashboard don't update until the next reset hour, you may see a discrepancy between these values and the alarm values on the Host List. Note that if you replace your Flow Collector, the new Flow Collector retains the alarm history from the old Flow Collector, and the values on the Security Insight Dashboard and the Host List remain in sync.
- If you disable a custom event that had resulted in any Policy Violation alarms, the Host List may show fewer alarming hosts. Enabling the custom event may or may not restore the count.

How to find this page

Do one of the following to access this page

Do any of the following to access the Host List:

- From the main menu, choose **Investigate > Assets > Hosts**.
- On the Security Insight Dashboard, click a number in an alarm category at the top of the page for an alarm you want to investigate. The Host List is filtered by that category.
- On the Host Group Report, click a number in an alarm category at the top of the page for an alarm you want to investigate. The Host List is filtered by that category and the applicable host group.

Concepts you need to understand to use this page

Sorting criteria

 To learn more about alarm severity levels, see [About Alarm Severities](#).

This list is sorted using the following criteria:

- By default, the list of hosts is sorted by the severity of the alarm categories, so the hosts with the worst alarms are at the top of the list. The order is based on the aggregate, or combination, of alarm categories for each host. The list is sorted in descending order, i.e., the host with the highest levels of severity in alarm categories is at the top of the list. Note the severity sorting link above the table. It is disabled by default, but it becomes enabled if you manually change the sorting. Use it to return to aggregate sorting.
- You can sort on any column in the list that has sort arrows in the column heading. In this case, sorting is based on comparing severity and percent of threshold. Note that if you choose to sort on a specific column, such as Host Name, the "Sorted by overall severity" link becomes enabled.
- If you click the "Sorted by overall severity" link, the list of hosts sorts again according to the default criteria described below. Using the filters in the left pane does not affect the "Sorted by overall severity" link.

Secure Network Analytics determines the sort hierarchy as follows:

1. The primary sort is by hosts with alarms that have one or more contributing critical, or ["always bad," security events](#). These security events are so serious that it does not matter how many index points they are worth or how much they have contributed to the percentage of threshold reached for a particular alarm. Hosts with these alarms are always at the top of the list.
2. The secondary sort is by hosts based on severity of alarm.
3. The third level of sorting is by hosts with alarms that contain the highest percentage of threshold reached.

Example:

- A host has three alarms each with a medium level of severity and each at 40% of threshold, for a total of 120% of threshold.
- A second host has one alarm with a high severity level (not caused

by an "always bad" security event) and which is at 80% of threshold.

- A third host has one alarm with a medium severity level and which is at 80% of threshold.

In the sort list, the alarm associated with the second host has a higher severity level, so that host will be at the top of the list. The first host, which has the three medium severity level alarms, will be next, because this host possesses the highest percentage of threshold reached. The third host will be the last one listed because it has a lower percentage of threshold reached than the first host.

Alarm value determination

The information you see on the Host List page depends on how you accessed this page.

- When you access the Host List page from an alarm category on the Security Insight Dashboard or the Host Group Report, the following parameters apply:
 - The number of alarming hosts that is displayed at the top of the Host List page in parentheses beside the Hosts heading represents only the number of alarming hosts for the selected category from which you pivoted from any of the aforementioned locations; hosts that are still accumulating points prior to alarming are NOT included in this number.
 - In the Filter Results By: Alarms section on the Host List page, the number in parentheses beside each alarm category represents only the number of alarming hosts for the selected category pivoted from; hosts that are still accumulating points prior to alarming are NOT included in this number.
 - In the Hosts table on the Host List page, only alarming hosts are displayed in the table; hosts that are still accumulating points prior to alarming are NOT included in this table.
- When you access the Host List page by performing one of the following options, the parameters listed below apply:
 - Click **Monitor > Hosts** from the main menu.
 - Click **View All Hosts** from within the Top Alarming Hosts section on the Security Insight Dashboard or on the Host Group Report.
 - The number of alarming hosts that is displayed at the top of the Host List page in parentheses beside the Hosts heading represents the hosts that are alarming AND the hosts that are still accumulating points prior to alarming for the selected category from which you pivoted from any of

the aforementioned locations.

- In the Filter Results By: Alarms section on the Host List page, the number in parentheses beside each alarm category represents only the number of alarming hosts for the selected category pivoted from; hosts that are still accumulating points prior to alarming are NOT included in this number.
- In the Hosts table on the Host List page, alarming hosts AND hosts that are still accumulating points prior to alarming for the selected category from which you pivoted are displayed in this table.



If a user filters further by a particular alarm category on this page, only values representing alarming hosts are shown.

Refer to the following parameters when viewing information on this page

- The Host List displays a maximum of 2000 hosts.
- If alarms are occurring for a host group, the column for each affected alarm category displays one of the following colors:
 - Black - Critical, or “always bad.”
 - Red - Major
 - Orange - Minor
 - Yellow - Trivial
- The color of the column indicates activity that has triggered security events for a specific host, but does not indicate the type of security event generating the alarm or the volume of alarms. To view the [security events](#) contributing to a specific alarm category for a specific host, click the alarm category column on the row for that host. (You can only do this if the alarm category is red.)
- The numerical value in each cell shows the percent of threshold for a specific alarm category and host.
- You should always investigate alarms with the severity level of Critical.
- A host must have the alarm category check box enabled in its effective policy to have a severity color assigned to it in the alarm category columns.
- If a host does not have any points accumulating towards an alarm category, it does not have a severity color in any of the alarm category columns on this page.

- Although the Watch Host Active alarm is an "always bad" security event, it displays as red (instead of black) on the Host List.

For more details about alarm severities, see [About Alarm Severities](#).

What you can view on this page

View the following information on this page

- [Severity levels](#) of alarm categories to see which hosts have generated the most severe alarms.
- The hosts generating the most alarms to decide which one(s) to investigate.
- The latest date and time the host communicated in Secure Network Analytics (Last Active column). You can view data as far back as seven days prior to the current day.
- The host groups associated with the most suspicious hosts.
- The countries (i.e., locations) associated with the most suspicious hosts.



Use the [filter](#) options to focus on the specific information you consider important.

What you can do next

Run a flow search

- To run a flow search associated with a single IP address, click the **⋮ (Ellipsis)** icon next to the applicable IP address to access the relevant options.
- To run a flow search associated with a row of data, in the Actions column, click the Ellipsis icon to access the relevant options.

Run a top report search

- To [run a top report](#) search associated with a single IP address, click the **Ellipsis** icon next to the applicable IP address to access the relevant options.
- To run a top report search associated with a row of data, in the Actions column, click the Ellipsis icon to access the relevant options.

Investigate the activity of a specific host

Investigate the activity of a specific host. To do this, click that host's IP address. The [Host Report](#) for that host opens.

Related topics

[About Alarm Categories](#)

[Security Event List](#)

[Investigating a Single Host](#)

[Investigating Users](#)

Investigating a Single Host

Use this page (Host Report) to view information about a single host's activity as far back as the last seven days to determine if there is a security risk associated with that host. The information shown pertains to an individual host. You can also use this page to assign a host to one or more host groups.

This page has comparable information to the Desktop Client's Host Snapshot and Quick View for Flow.

How to find this page

Do the following to access this page

1. From the main menu, choose **Investigate > Assets > Hosts**.

The Host List opens.

2. Click a host or click the  (**Search**) icon and enter an IP address or range in the Search field.

The Host Report opens from which you can view detailed information about a specific host or session.

What you can view on this page

Use the information on this page to answer these questions

- Who is this host?
 - Use information in the Host Summary component to quickly identify high-level, important information about the host, its location, the groups to which it belongs, and the policy assigned to it.
- Who is this host contacting?
 - Based on the Traffic by Peer Host Group, are there any host groups that raise concerns or questions? Is the volume of traffic reasonable?
- What is this host doing that I should be concerned about?
 - Are any of the alarm categories at the top of the page showing a volume of alarms that causes concern? If so, click the number of active alarms for the category you want to investigate. The Alarm dashboard displays security events for that category filtered by this host.

- Does the Alarms by Type bar chart display any spikes or trends that you need to investigate?
- In the Users & Sessions component, is the user valid for that device? Are the start and end times and session duration reasonable for this host?
- In the Application Traffic component, which applications is this host using the most? Which ones are transferring the most data? How many are being used to contact inside hosts, and how many are being used to contact outside hosts? Do the results represent normal trends for your system and this host?
- Do the 7-day and 24-hour trend lines for application traffic between this host and inside or outside hosts show spikes or dangerous trends that you need to investigate? Do you notice absence of activity that seems unusual for the type of host and certain types of applications?

Example: You notice a Web server that has no activity for the HTTP application for several days. This would be very unusual, so you check the server to determine if it is functioning normally or if there may be a security problem.

Components on this page

The following components are available on the Host Report and provide different types of information to assist your analysis or investigation.

Alarm Categories

An alarm category is a “bucket” toward which a defined list of security events contributes. When network activity meets or exceeds a defined set of criteria specified for this alarm category, it triggers an alarm. Alarm categories contain only host alarms. (A host alarm is a security event that has triggered an alarm.)

The alarm categories at the top of this page actively receiving alarms that affect this host. To investigate alarms that are currently contributing to a particular alarm category for this host, click the number in the corresponding alarm category to the [Alarm dashboard](#).

The alarm categories are as follows:

Name of Alarm Category	Description
Anomaly	Tracks events that indicate that hosts are behaving abnormally or generating traffic that is unusual, but isn't consistent with another

Name of Alarm Category	Description
	category of activity.
C&C (Command & Control)	Indicates the existence of infected servers or hosts in your network attempting to contact a C&C server.
Concern Index	Indicates that the host's concern index has either exceeded the CI threshold or rapidly increased.
Data Hoarding	Indicates that a source or target host within a network has downloaded an unusual amount of data from one or more hosts.
Exfiltration	Indicates that an abnormal amount of data has been transferred to inside or outside hosts.
Exploitation	Tracks direct attempts by hosts to compromise each other, such as through worm propagation and brute force password cracking.
DDoS Source	Indicates that a host has been identified as the source of a DDoS attack.
DDoS Target	Indicates that a host has been identified as the target of a DDoS attack.
Policy Violation	Indicates that the subject is exhibiting behavior that violates normal network policies.
Recon	Indicates the presence of unauthorized and potentially malicious scans using TCP or UDP and being run against your organization's hosts. These scans, referred to as "reconnaissance," are early indicators of attacks against your network, and the scans may come from outside or inside your organization.

Name of Alarm Category	Description
Target Index	Indicates that the target IP has been the recipient of more than an acceptable number of scans or other malicious attacks.

Host Summary

Information that provides the following details about the host:

- Host IP address
- Host's status
- Name assigned to host
- Host group(s) to which the host is assigned
- Host's location, such as country
- When the host was first and last seen by Secure Network Analytics
- Date and time when the corresponding Flow Collector last observed traffic for the host.
- Policy applied to the host
- Last active media access control (MAC) address for the host.

This MAC address is provided by a Flow Collector (NetFlow).

- ISE ANC policies applied to the host

Click **Edit** to set which ANC policies apply to the host. For details, see [Applying Mitigation Actions to a Host Using Cisco® ISE](#).

The following buttons are available in this section:

- Flows – Run a report but make edits to the parameters first.
- Classify – Assign this host to one or more host groups. If the host group you want to use is not shown, you must add it through the [Desktop Client](#).
- History – View a log of all mitigation actions that have been requested for a host. For details, see [Viewing Mitigation Actions for a Host \(Host History\)](#).

Traffic by Peer Host Group

This component displays aggregated data for server or client peers belonging to specific host groups that are communicating with this host (displayed in the middle of the component). Host groups inside your network are listed on the left, and host groups outside your network are listed on the right. The top eight host groups based on volume are listed separately. If there are more than eight, the remainder are grouped together as Other, for a total of nine groups. The thickness of the line between this host and a host group indicates the volume of traffic. When you hover your cursor over a line connecting two host groups, a tooltip shows the volume of traffic transferred between these two host groups.

If you click a host group, a column, or the line between two host groups, you can access the menu (represented by the **⋮** [Ellipsis] icon) from which you can run a flow search or [top report](#) to view the traffic associated with the two host groups to conduct a deeper analysis into specific areas of interest.

Alarms by Type

Number of alarms that have occurred over the past seven days. Each type of alarm is represented by a different color, which is identified in the legend beneath the component's title. When you click a type of alarm in the legend, that alarm is suppressed in the bars on the chart, and the Event Count scale changes accordingly. Clicking the alarm in the legend again restores it in the bars on the chart. You can use this chart to focus on the volume and frequency of a particular type of alarm's occurrence affecting this host. When you click a section on the chart representing an alarm category, the [Alarm dashboard](#) page opens, displaying information about each occurrence of the alarm on a specific date filtered by this host.

Top Security Events

Top ten security events that have been active on your network since the last reset hour displayed in two different tables. Information in these tables provides insight into what is causing the security events to occur. Use the data to diagnose and triage specific security events. Secure Network Analytics collects and displays security event information regardless of whether or not the events have caused an alarm to fire.

Click the Source tab to view security events where the host is the source, and click the Target tab to view security events where the host is the target. On the Source tab, the security events are sorted by volume of Concern Index (CI) points for this host that have

been active on your network since the last reset hour. On the Target tab, the security events are sorted by volume of Target Index (TI) points for this host that have been active on your network since the last reset hour. The number of active security events for the source and for the target is displayed in parentheses on the associated tab.

You can view the following information in this table:

- View the target hosts associated with the security events.
 - Are there any that might be misclassified?
 - Does an external source combined with other information about the security event indicate that an attack might be imminent. If so, you might want to perform a packet query. See Perform a packet query further down in this topic.
- View the date and time the security events became active. Is there a pattern that might indicate scanning activity).
- View the type of security event and the number of times the event occurred (displayed in the Count column).
- To view all the security events associated with this host that have occurred since the last reset hour, click **View More Security Events** at the bottom of the table. The Security Events Details page is displayed.
- Click the arrow beside the applicable security event to expand it and view the following information:
 - Information indicating why the security event was triggered. Depending on the event, some details that may be provided are the number of concern index point this event has accumulated since the last reset hour as well as the tolerance level for which it has been set.
 - Detailed information is not provided for every security event type, since the details of many events are the event itself (e.g., Ping Scan).
- A description of the security event.

Users & Sessions

User and session information associated with the host, including:

- Information about the device used during the last 24 hours. The devices may or may not be currently active.
- User name associated with the session using that host.
- Start and end time for the session.



You must have an identity management device such as Cisco Identity Services Engine (ISE), Cisco Identity Services Engine Passive Identity Connector (ISE-PIC), or Cisco ASA installed on your network to display information in this component.

Application Traffic

Applications used by this host and their volume of traffic; you can view information about internal or external traffic by selecting either the Internal or External tab. The % column provides information to show the percent of total traffic carried by each application. For each application, you can also view 7-day and 24-hour trends. For details about a specific day or time, place your cursor over the graph; a tooltip displays the date, time, and amount of traffic transferred.

What you can do next

Investigate the activity of a specific host

- To view the [security events](#) contributing to a specific alarm category for this host, click the number in the applicable alarm category at the top of the page.
- To view detailed information about a [specific user](#) associated with this host, click the user name in the Users & Sessions section.

Assign this host to one or more host groups

1. Click **Classify** in the Host Summary component.

The Host Groups Selector appears on the right and displays a tree structure with host groups. The tree is similar to the Enterprise tree in the Desktop Client. If the host group you want to use is not shown, you must add it through the Desktop Client.

If a host is already assigned to one or more host groups, including Catch All, the buttons for those groups are selected. The host group names are shown in the Host Summary component on this page.

2. Expand the applicable host group branch to find the host group you want to use.

Keep in mind that certain host groups are restricted and have properties that you cannot change. You cannot select one of the restricted host groups to classify hosts, because that affects the IP addresses associated with the host group. Host groups that you cannot change are:

- Inside Hosts, Catch All
- Outside Hosts, Countries and all its sub-branches
- Command &Control Servers
- Bogon
- Tor

In the Inside Hosts group, you can choose a sub-branch host group and use it to classify hosts. If you need to add a sub-branch host group, use the Enterprise tree in the Desktop Client. For more information, refer to *About Host Groups and Policies* in the Desktop Client Help.

 You can use the Search box to quickly search for a host group.

3. In the expanded tree, click the name of the host group to which you want to assign this host. You can select multiple host groups, however, a host may not be in both an Inside and Outside host group.

Each button you select displays a check.

4. Click **Apply**.

A message banner appears under the page title to confirm that you have successfully classified the host. The host group names are shown in the Host Summary component on this page.

Related topics

About Alarm Categories

Security Event List

Investigating Host Groups

Use this page to view detailed information about a host group.

How to find this page

Do the following to access this page

You can access this page by doing one of the following:

- From the main menu, choose **Investigate > Assets > Host Groups**.
- Click a link that is contained on multiple pages within the Web App.

On the...	Click the Host Group name that is displayed as a link in the...
Host List	Host Groups column.
Host Report - Host Summary Section	Host Groups section.
Alarm List	Source Host Groups and Target Host Groups columns.
Alarm Dashboard	Source Host Groups and Target Host Groups columns.
Security Event Details page	Source Host Groups and Target Host Groups columns.
	Subject Host Groups or Peer Host Groups columns.
Flow Search Results page	Make sure to use the tabular view, which is the default view. To change to the conversational view, click the  (Conversational View) icon (the icon with the grid), located to the right of the  (Tabular View) icon (the icon with three horizontal lines). The icon that has a white background indicates the active view.



If there isn't a host group for a flow, or if there are multiple host groups for a flow, then no link will be available. The entry in the applicable column is *Multiple Host Groups*.

What you can view on this page

View the following components to assist your analysis or investigation. To view the report for another host group, click **Change Host Group**.

Active alarming hosts within a host group

Although the following "always bad" security events don't always contribute index points to an alarm category, one of these will immediately increase the alarm count for a host and assign a severity level of "major" (which causes both the number and the alarm trend to turn red):

- All custom security events
- Any security event contributing to the Policy Violation alarm (Note that these security events may contribute to other alarms as well, but they cause only the Policy Violation alarm to turn black.)

If you reboot your Flow Collector, it deletes all alarm history. Since the alarming host widgets (which display the number of hosts receiving alarms since the last reset hour for a specific category) on the Security Insight Dashboard don't update until the next reset hour, you may see a discrepancy between these values and the alarm values on the Host List.

If you replace your Flow Collector, the new Flow Collector retains the alarm history from the old Flow Collector, and the values on the Security Insight Dashboard and the Host List remain in sync.

If you disable a custom event that had resulted in any Policy Violation alarms, the Host List may show fewer alarming hosts. Enabling the custom event may or may not restore the count.


Active alarming hosts within the corresponding host group are shown across the top of the page in various alarm categories (such as Concern Index or Recon) and display the following data:

- The number of member hosts of the applicable host group receiving alarms (since the last reset hour) for that category. If no member hosts are receiving alarms that contribute to the alarm category, the container displays a zero. If you click the number (value cannot be 0), the [Host List](#) page opens, filtered by that category.


- The color represents the current highest severity for an alarm occurring for any member host of the applicable host group within that category. If the highest severity is major, the color of the corresponding number turns red; if it is minor, it turns orange. If no alarms are contributing to an alarm category, its number turns gray. To learn more about alarm severity levels, see [About Alarm Severities](#).
- The 7-Day Trend chart's color always matches the number's color. For details about a specific day, place your cursor over the graph; a tooltip displays the date and the number of member hosts of the applicable host group receiving alarms (for a specific day) that contribute to that alarm category.

Example 1:

For instances in which no member hosts of the applicable host group have received any contributing alarms (since the last reset hour nor at any time during the last seven days), no trend chart is displayed.


**Example 2:**

For instances in which no member hosts of the applicable host group have received any contributing alarms (since the last reset hour) but at least one member host has received a contributing alarm during the last seven days, the trend chart is dark gray.



Example 3:

For instances in which at least one member host of the applicable host group has received a contributing alarm since the last reset hour, the number is either red or orange, depending on whether the highest alarm severity assigned is major or minor. To learn more about alarm severity levels, see [About Alarm Severities](#)). The 7-Day Trend chart's color is the same color as the number.



The alarm categories are as follows:

Name of Alarm Category	Description
Anomaly	Tracks events that indicate that hosts are behaving abnormally or generating traffic that is unusual, but isn't consistent with another category of activity.
C&C (Command & Control)	Indicates the existence of infected servers or hosts in your network attempting to contact a C&C server.
Concern Index	Indicates that the host's concern index has either exceeded the CI threshold or rapidly increased.
Data Hoarding	Indicates that a source or target host within a network has downloaded an unusual amount of data from one or more hosts.
Exfiltration	Indicates that an abnormal amount of data has

Name of Alarm Category	Description
	been transferred to inside or outside hosts.
Exploitation	Tracks direct attempts by hosts to compromise each other, such as through worm propagation and brute force password cracking.
DDoS Source	Indicates that a host has been identified as the source of a DDoS attack.
DDoS Target	Indicates that a host has been identified as the target of a DDoS attack.
Policy Violation	Indicates that the subject is exhibiting behavior that violates normal network policies.
Recon	Indicates the presence of unauthorized and potentially malicious scans using TCP or UDP and being run against your organization's hosts. These scans, referred to as "reconnaissance," are early indicators of attacks against your network, and the scans may come from outside or inside your organization.
Target Index	Indicates that the target IP has been the recipient of more than an acceptable number of scans or other malicious attacks.

Summary

This component displays the following:

- The location of the Host group in the Enterprise Tree.
- The traffic data currently monitored in the Host group. This section shows the total data monitored since the last reset hour and the total data separated by the amount of internal, inbound, and outbound traffic for that host group.
- The total traffic for the last seven days by hour in the Host Group. If you hover your cursor over a point in time on the Total Traffic graph, the date, time, and the amount of traffic is displayed. If you click a point in the graph, you can access the context

menu, from which you can run a flow search associated with particular host groups or [run a top report](#) to conduct a deeper analysis into specific areas of interest.



The number of byte counts and number of flow records for child host groups may not be equal to the byte counts and number of flow records of their parent group. If a flow goes from one child group to another child group, it is only counted once in the parent group, as it is a single flow.

Top Host Groups By Traffic

This component displays the top 10 Inside and Outside host groups with which the current host group (displayed in the middle of the graph) has communicated within the last 12 hours. The height and width of the columns is significant:

- **Height:** The height of the columns shown for the top inside host group and top outside host group is proportional to the amount of data generated by each of those two host groups.
- **Width:** The width of the lines connecting the current host group (in the middle of the graph) to each of the top ten inside and outside host groups is proportional to the amount of traffic sent between each pair of host groups. When you hover your cursor over a line connecting two host groups, the total amount of traffic transferred between these two host groups is displayed.

At the bottom of this component you can see the number of additional host groups over 10 (outside of both the Inside and Outside host groups) with which the current host group has communicated as well as the percentage of total traffic they represent. This allows you to view the percentage of total traffic that the top 10 Inside and Outside host groups represent.

If you click a host group, a column, or the line between two host groups, you can access the context menu from which you can run a flow search or [top report](#) to view the traffic associated with the two host groups to conduct a deeper analysis into specific areas of interest.

The title bar includes the last time this component was updated. When you access this page, the Secure Network Analytics system determines whether this component has been loaded within the last hour. If this has not occurred, the system queries the Flow Collector (sFlow) to populate the graph. If the component has been loaded within the last hour, the cached data is displayed. When you click **Update**, the system queries the Flow Collector (sFlow) and populates the graph with the latest data.

Top Alarming Hosts

This component displays the top seven alarming hosts in the Host group, sorted by alarm severity (represented by alarm category containers), that have been active on your network since the last reset hour.

By default, the top seven alarming hosts in the Inside Host Group is displayed, but you can change the host group by clicking **Change Host Group**. If a host is a member of two or more host groups and receives more than one alarm, Secure Network Analytics displays only one of the host groups of which this host is a member.

For information about how this list is sorted alarm severity levels, see [About Alarm Severities](#).



The order of the host groups on this list may vary slightly from that of the Host List if the Manager and Flow Collector have not yet synced.

To see the alarm category name and percentage of threshold (or alarm count for the “always bad” alarms) reached by a host (within a host group) for a particular alarm category, hover your cursor over the applicable alarm category container.

From this component you can access the following information:

- The [Host List](#), which displays all hosts that have had active alarms since the last reset hour. To view this data, click the View All Hosts link at the bottom of the component. The Host List opens, sorted by severity.
- The [Host Report](#) for the host which you selected. You can view information about a single host’s activity as far back as the last seven days. To view this data, click the applicable host IP address.
- The [Alarm Dashboard](#), which displays all alarms contributing to a specific alarm category for a specific host. To view this data, click the alarm category container for the host you want to investigate.
- Flows associated with a host. To do this, click the context menu beside the applicable host to access the relevant options. See for more information.
- Pertinent information from top reports so that you can conduct a deeper analysis into specific areas of interest. To do this, click the context menu beside the applicable host to access the relevant options. For more information, see [Building a Top Report Search](#).

Sorting criteria

1. The primary sort is by hosts with alarms that have one or more contributing critical, or ["always bad," security events](#).



These security events are so serious that it does not matter how many index points they are worth or how much they have contributed to the percentage of threshold reached for a particular alarm. Hosts with these alarms are always at the top of the list.

2. The secondary sort is by hosts based on severity of alarm.
3. The third level of sorting is by hosts with alarms that contain the highest percentage of threshold reached.

Example:

- A host group has three alarms each with a medium level of severity and each at 40% of threshold, for a total of 120% of threshold.
- A second host group has one alarm with a high severity level (not caused by an "always bad" security event) and which is at 80% of threshold.
- A third host group has one alarm with a medium severity level and which is at 80% of threshold.

In the sort list, the alarm associated with the second host group has a higher severity level, so that host group will be at the top of the list. The first host group, which has the three medium severity level alarms, will be next, because this host group possesses the highest percentage of threshold reached. The third host group will be the last one listed because it has a lower percentage of threshold reached than the first host group.

To learn more about alarm severity levels, see [About Alarm Severities](#).

Top Applications

This component displays the top nine applications and the sum of all other applications, displayed under Other, currently monitored in the Host group for the last 12 hours in five minute intervals. The top application graphs are separated by inbound, outbound, and internal traffic. The legend at the bottom of this component displays the top nine

applications:

- Click an application name in the legend to remove the application data from displaying on the graph.
- Click the Application name again to restore its data on the graph.

If you hover your cursor over a point in time on the Top Applications graph, the application name, start time, traffic bytes, and traffic direction is displayed. If you click a point in the graph, you can access the context menu from which you can run a flow search or [top report](#) for this host group to conduct a deeper analysis into specific areas of interest.

What you can do next

Do the following as necessary:

- Access the [Host Group Management page](#) to view or update host group configurations. To do this, from the Actions drop-down menu, select **Host Group Management**.
- To view flows associated with a particular host, host group, or application, run a flow search.
- To view top reports that will give more in-depth information about specific areas of interest, [run a top report search](#).



This page is not available on the secondary Manager.

Investigating Users

Use this page (User List) to see which users have been active on your network since the last archive hour, as well as information about their sessions, location, and current alarms.

The information is retrieved from an [identity management device](#) such as Cisco® Identity Services Engine (ISE), Cisco Identity Services Engine Passive Identity Connector (ISE-PIC), or Cisco ASA. As part of a security investigation, this device can identify users and the devices they are using on the network. For this feature to work properly, the identity management device must be configured separately.



You must have one of the previously mentioned identity management devices configured for use with Secure Network Analytics; otherwise, this page does not display any data.

How to find this page

Do the following to access this page

From the main menu, choose **Investigate > Assets > Users**.

The User List opens.

Concepts you need to understand to use this page

User Alarm Index Calculation for Multiple Sessions

A user may have multiple active sessions across different devices (for example, a laptop and a mobile phone). In such cases, the appliance considers all current active sessions for that user across all their associated devices to calculate alarm indexes and severity scores. This ensures that the displayed alarm indexes and severity levels accurately reflect the combined security posture from all active user sessions.

Sorting criteria

This list is sorted using the following criteria

- By default, the list of users is sorted by the severity of the alarm categories, so the users with the worst alarms are at the top of the list. The order is based on the aggregate, or combination, of alarm categories for each user. The list is sorted in descending order, i.e., the user with the highest levels of severity in alarm categories is at the top of the list. Note that the severity sorting link above the table.

It is disabled by default, but it becomes enabled if you manually change the sorting. Use it to return to aggregate sorting.

- You can sort on any column in the list that has sort arrows in the column heading. In this case, sorting is based on comparing severity and percent of threshold. Note that if you choose to sort on a specific column, such as Last Session Start, the "Sorted by overall severity" link becomes enabled.
- If you click the "Sorted by overall severity" link, the list of users sorts again according to the default criteria described below. Using the filters in the left pane does not affect the sorting.

Secure Network Analytics determines the sort hierarchy as follows:

1. The primary sort is by users with alarms that have one or more contributing critical, or ["always bad," security events](#). These security events are so serious that it does not matter how many Concern Index (CI) points they are worth or how much they have contributed to the percentage of threshold reached for a particular alarm. Users who have hosts with these alarms are always at the top of the list.
2. The secondary sort is by users based on severity of alarm.
3. The third level of sorting is by users with alarms that contain the highest percentage of threshold reached.

Example:

- A user has three alarms each with a medium level of severity and each at 40% of threshold, for a total of 120% of threshold.
- A second user has one alarm with a high severity level (not caused by an "always bad" security event) and which is at 80% of threshold.
- A third user has one alarm with a medium severity level and which is at 80% of threshold.

In the sort list, the alarm associated with the second user has a higher severity level, so that user will be at the top of the list. The first user, who has the three medium severity level alarms, will be next, because this user possesses the highest percentage of threshold reached. The third user will be the last one listed because it has a lower percentage of threshold reached than the first user.



To learn more about alarm severity levels, see [About Alarm Severities](#).

Refer to the following parameters when viewing information on this page

- The Users list displays a maximum of **2000 users**.

The page shows only the users who had an active session associated with an **Inside Host** in the **Top 4000** based on the **Total Security Index in the last 24 hours**.

- If alarms are occurring for a user's last active host, the column for each affected alarm category displays one of the following colors:
 - **Black** - Critical, or "always bad"
 - **Red** - Major
 - **Orange** - Minor
 - **Yellow** - Trivial
- **If colors are not shown**, enable the alarm category for the host associated with the user in its **Effective Policy**. Refer to [Managing Core Policies](#) for details.
- The color of the column indicates activity that has triggered security events for a host used by a specific user, but does not indicate the type of security event generating the alarm or the volume of alarms.
- If the alarm category is **Red**, you can click the **Alarm Category** column for the user you want to examine to view the [Alarm List](#) page.
- You should always investigate alarms with the severity level of **Critical**.
- The numerical value in each cell shows the percent of threshold for a specific alarm category and a user.
- **If a user's host does not have any points** accumulating towards an alarm category, the user does not have a severity color in any of the alarm category columns on this page.

 To learn more about alarm severities, see [About Alarm Severities](#).

What you can view on this page

View the following information on this page

- Severity levels of alarms to see which users have generated the most severe alarms.
- The user names generating the most alarms to ensure they are authorized users.
- The number of sessions to ensure they represent a reasonable number.
- The countries (i.e., locations) associated with the highest levels of alarms.
- The number and relative proximity of countries within the last 24 hours.
- Session start and end times to see if they represent suspicious activity.
- The number of devices to ensure they represent a reasonable number.



Use the filter options to focus on the specific information you consider important.

What you can do next

Do the following as necessary:

- To investigate the activity of a specific user, click that user's name to open the [User Report](#)
- To investigate an alarm category for a specific user, click the alarm category column in that row.

Related topics

[About Alarm Categories](#)

[Security Event List](#)

[Investigating a Single Host](#)

[Investigating Hosts](#)

Investigating a Single User

Use this page (User Report) to view information about a single user's activity. The information shown is based on the last 24 hours.

This page has comparable information to the Desktop Client's User Snapshot.

The information is retrieved from an [identity management device](#) such as Cisco Identity Services Engine (ISE), Cisco Identity Services Engine Passive Identity Connector (ISE-PIC), or Cisco ASA. As part of a security investigation, this information can identify users and the devices they are using on the network. For this feature to work properly, the identity management device must be configured separately.



You must have one of the previously mentioned identity management devices configured for use with Secure Network Analytics; otherwise, this page does not display any data.

How to find this page

Do the following to access this page

1. From the main menu, choose **Investigate > Assets > Users**.

The User List opens.

2. Click a user or click the  (**Search**) icon and enter a user's name in the Search field.

The User Report opens from which you can view more detailed information.

Follow these guidelines when entering a user name:

- The name cannot contain any of these characters: | + = ? " < > () : ;
- The search box is not case sensitive.
- You can type partial name(s) with wild cards. The wild card can be in any position, such as srh*, *doe. You can use more than one wild card in each name.

What you can view on this page

View the following information on this page



- If you disable an alarm in the Desktop Client, that alarm does not appear on this dashboard.

- You must [configure Secure Network Analytics to use Active Directory](#) to display the user information on this page. You must have admin privileges to configure integration with Active Directory.

- Alarms that display the [severity level](#) for hosts associated with a particular user. The percentage value displayed within the colored portion at the top shows the percentage of threshold that has been reached by the host with the highest severity level, based on all hosts associated with this user. The title at the top of this container changes color to indicate high, medium, or low levels of severity. If no alarms exist, the container is blue. The color reflects the highest severity level of alarms for the host on which the user has an active session.
- Contact information for user associated with a host you are investigating. The information is pulled from Active Directory, so you can quickly evaluate if the employee is a valid user for the host, session, devices, locations, and applications. If you need to contact the user or the user's manager, that information is provided.
- Host, user name, and session information associated with the user.
- Applications (both server and client) used by this user to determine which applications he or she is using the most, how many are being used to contact inside hosts, and how many are being used to contact outside hosts.
- Devices used by this user to determine which devices are being used the most.

What you can do next

To investigate the activity of a specific user, do the following as necessary:

- To view detailed information about a [specific host](#) associated with this user, click the host's IP address. The Host Report for that host opens.

Related topics

[About Alarm Categories](#)

[Security Event List](#)

Interface Monitoring: Overview

Use this page to view the inbound and outbound interface traffic for a domain since the last reset hour. You can analyze the data to assess the possibility of an attack or other threats against your network and devices.

The interface status results are sorted by the Current Utilization value in descending order.

Open Response Management

From the main menu, choose **Investigate > Assets > Interfaces**.

The Interfaces page is displayed.



To view the help topic for the Interface Monitoring home page (the Interface Status page), see [Interface Monitoring: Analyze Interface Status Data](#).

Interface Monitoring: Table of Contents

Interface Monitoring: Filter Results Overview

You can filter the results by choosing to include or exclude specific Flow Collectors and exporters.

Opening the Filter

To open the Flow Collectors & Exporters Filter panel, click **Filter by Device**.

The Include tab opens by default.

What To Do Next

[Interface Monitoring: Filter Results](#)

Related Topics

[Interface Monitoring: Analyze Interface Status Data](#)

[Interface Monitoring: View Interface Traffic Details Overview](#)

[Interface Monitoring: Reports Overview](#)

[Interface Monitoring: Overview](#)

Interface Monitoring: Filter Results

Note the following information when using the filter:

- At the top of the page, each included filter criterion is displayed in a green text box and each excluded filter criterion is displayed in a black text box.
- If you hover your cursor over an exporter in the Filter panel, the exporter name and IP address are displayed. If the exporter is a Flow Sensor or a Virtual Flow Sensor, the names of all alarms that have occurred for that exporter since the last reset hour are displayed as well.
- To [run a flow search or top report](#), click a point in the graph and then click the **Ellipsis** icon.

Note the following guidelines when selecting Flow Collectors and exporters:

- To search for a Flow Collector, enter the Flow Collector name in the Search field and press **Enter**. The associated Flow Collector name becomes red and italicized.
- To search for an exporter, enter either the exporter name or IP address in the Search field and press **Enter**. The associated exporter name becomes red and italicized.
- You can select multiple Flow Collectors as well as exporters that are on different Flow Collectors.
- If you select a Flow Collector, all exporters listed under that Flow Collector are automatically included in the results (note that you will not see check marks in the exporters' check boxes in this instance). For any exporter that you do not want to include, click **Exclude** at the top of the Filter panel to open the Exclude tab and select the applicable exporters. Click **View** at the bottom of the panel to accept your selection(s).
- You can select exporters without having to select the associated Flow Collector. To do this, click the arrow beside the Flow Collector. The list of exporters on this Flow Collector is displayed. Select the exporters that you want to include in the results. For any exporter that you do not want to include, click **Exclude** at the top of the Filter panel to open the Exclude tab and select the applicable exporters. Click **View** at the bottom of the panel to accept your selection (s).
- To include all network devices in the results, select the All Network Devices check box.
- The color of an exporter corresponds to its highest alarm severity since the last reset hour.

The following icons distinguish the various types of exporters.

Icon	Icon Name
	Exporter
	Firewall
	Flow Sensor
	Virtual Flow Sensor

What To Do Next

[Interface Monitoring: Analyze Interface Status Data](#)

Related Topics

[Interface Monitoring: Filter Results Overview](#)

[Interface Monitoring: View Interface Traffic Details Overview](#)

[Interface Monitoring: Reports Overview](#)

[Interface Monitoring: Overview](#)

Interface Monitoring: Analyze Interface Status Data

Keep the following questions in mind as you analyze interface status data:

- Which Flow Collector is this?
- Which exporter is sending the data?
- Which interface is the source of the data?
- What is the current percent of available bandwidth (Current Utilization) consumed by inbound or outbound traffic (direction is indicated by the Direction column) on the corresponding interface? This value is preceded by a bar that shows the following:
 - Threshold – represented by a thin line that intersects each rectangular bar.



The bar changes color when the threshold is approached, as described in the following table.

Bar Color	% of Threshold Reached
Empty Graphic	0% of configured threshold.
Green	0% to 50% of configured threshold.
Yellow	51% to 75% of configured threshold.
Orange	76% to 99% of configured threshold.
Red	100% of configured threshold or greater.

- At which rate is the interface currently receiving or sending (direction is indicated by the Direction column) traffic?

- What is the peak percent of available bandwidth (Maximum Utilization) that has been consumed by inbound or outbound traffic (direction is indicated by the Direction column) on the corresponding interface? This data uses the same logic as that used for the Current Utilization entries, as described earlier.
- What is the peak rate (in bits per second) at which traffic has been received or sent (direction is indicated by the Direction column) by the interface since the reset hour?
- At which speed is traffic being sent?
- Which exporters have high severity levels assigned to them? To determine this, click **Filter by Device** to open the Filter panel. (For more information about filtering, see [To filter the results](#) in this topic) Once you filter for these exporters, you can view more data that pertains to each of these exporters in the Interface Status table. Expand the row for an exporter to see more details, such as the Top Application Traffic graph. Click a point in the graph to run the Top Applications report to see which applications have a high amount of traffic. For information about running a top report, see "To run a top report from an IP address menu" in [Building a Top Report Search](#)

If you hover over a particular point in the graph, you can view the following information

- Current utilization
- Maximum utilization
- Average utilization
- Threshold percentage
- Speed of traffic

Related Topics

[Interface Monitoring: Filter Results Overview](#)

[Interface Monitoring: View Interface Traffic Details Overview](#)

[Interface Monitoring: Overview](#)

Interface Monitoring: View Interface Traffic Details Overview

You can view the interface traffic details associated with a particular entry in the Interface Status table.

Opening the Interface Traffic Details

To view the interface traffic details associated with a particular entry in the Interface Status table, click the arrow at the beginning of the row to expand the row and view the related data.

What To Do Next

[Interface Monitoring: View Interface Traffic Details](#)

[Interface Monitoring: View Top Application Traffic](#)

[Interface Monitoring: View the Packets Graph](#)

[Interface Monitoring: View the Utilization Graph](#)

Related Tasks

[Interface Monitoring: Filter Results](#)

Related Topics

[Interface Monitoring: Overview](#)

Interface Monitoring: View Interface Traffic Details

The following information is provided in this section and represents activity since the last reset hour:

- Description
- Total bytes
- Peak Inbound (peak rate at which inbound traffic has been received)
- Peak Outbound (peak rate at which outbound traffic has been sent)
- Average Packet Rate
- Current Packet Rate
- Maximum Packet Rate
- Alarms that have been triggered since the last reset hour. The colored circle next to an alarm name represents the highest severity for that alarm since the last reset hour. The numeral denotes the number of times the alarm has been triggered since the last reset hour.

Icon Color	Alarm Severity
Red	Critical
Orange	Major
Yellow	Minor
Dark Blue	Trivial
Light Blue	Informational

To run a flow search or top report, click the **...** (**Ellipsis**) icon next to the applicable Interface or Exporter IP to access the relevant options.

What To Do Next

[Interface Monitoring: View Top Application Traffic](#)

[Interface Monitoring: View the Packets Graph](#)

[Interface Monitoring: View the Utilization Graph](#)

Related Tasks

[Interface Monitoring: Filter Results](#)

Related Topics

[Interface Monitoring: View Interface Traffic Details Overview](#)

[Interface Monitoring: Overview](#)

Interface Monitoring: View Top Application Traffic

The Top Application Traffic graph (which is displayed by default) shows the following information for the selected time frame (which by default is for the last six hours):

- The top ten applications that constitute the Inbound and outbound traffic for the selected interface, based on volume. If there are more than ten applications, the remaining applications are grouped together as Other.
 - Inbound traffic is blue and located above the horizontal base line.
 - Outbound traffic is green and located below the horizontal base line.
- The left vertical axis represents the rate at which an application is using the bandwidth.
- The right vertical axis represents the percentage of bandwidth that an application has used.
- For details about a specific day or time, place your cursor on a point in the graph; a tooltip displays the application type, date, time, traffic direction, and rate at which the application is using bandwidth. If you have selected the Utilization graph, then you will also see the percentage of utilization.

You can do the following in this graph:

- To zoom in on an area, press down and drag your cursor across the area of interest.
- Click an application name in the legend to remove the application data from displaying on the graph. Click the Application name again to restore its data on the graph.
- To run a flow search or top report, click a point in the graph and make your selection from the pop-up menu.
- To view historical graph data, click the drop-down list to the right of the Utilization check box and select the desired time frame. The time span (length of time between the FROM and TO fields) must be 24 hours or less.
- The number of retention days for data may be different depending on what you have configured on the Manager Data Retention page. (To access this page, go to the Desktop Client, and in the Enterprise tree, right-click the **(Manager)** icon and then select **Configuration > Properties > Data Retention**.)

What To Do Next

[Interface Monitoring: View Interface Traffic Details](#)

[Interface Monitoring: View the Packets Graph](#)

[Interface Monitoring: View the Utilization Graph](#)

Related Tasks

[Interface Monitoring: Filter Results](#)

Related Topics

[Interface Monitoring: View Interface Traffic Details Overview](#)

[Interface Monitoring: Overview](#)

Interface Monitoring: View the Packets Graph

Opening the Packets Graph

Click **Packets** to display the Packets graph.

Viewing the Packets Graph

You can view the following information for the selected time frame (which by default is for the last six hours):

- The rate at which inbound and outbound packets are being transferred.
- Inbound packet information is shown in blue and located above the horizontal base line.
- Outbound packet information is shown in green and located below the horizontal base line.
- The left vertical axis represents the rate at which the packets are being transferred.
- For details about a specific day or time, place your cursor on a point in the graph; a tooltip displays the date, time, and rate at which packets are being sent to or sent by (depending on whether you click in the Inbound or Outbound section) the interface.

You can do the following in this graph:

- To zoom in on an area, press down and drag your cursor across the area of interest.
- When you click a packet type name in the legend, its corresponding data is suppressed on the graph. Click the name again to restore the data.
- To run a flow search or top report, click a point in the graph and then click the **Ellipsis** icon.
- To view historical graph data, click the drop-down list to the right of the Utilization check box and select the desired time frame. The time span (length of time between the FROM and TO fields) must be 24 hours or less.
- The number of retention days for data may be different depending on what you have configured on the Manager Data Retention page. (To access this page, go to the Desktop Client, and in the Enterprise tree, right-click the **Manager** icon and then select **Configuration > Properties > Data Retention**.)

What To Do Next

[Interface Monitoring: View Interface Traffic Details](#)

[Interface Monitoring: View Top Application Traffic](#)

[Interface Monitoring: View the Utilization Graph](#)

[Related Tasks](#)

[Interface Monitoring: Filter Results](#)

[Related Topics](#)

[Interface Monitoring: View Interface Traffic Details Overview](#)

[Interface Monitoring: Overview](#)

Interface Monitoring: View the Utilization Graph

Opening the Utilization Graph

Click the Utilization check box to display an overlay of the Utilization graph on either the Application or Packets graph (whichever is displayed). This gives you the ability to compare utilization data with either top application or packets data.

Viewing the Utilization Graph

You can view the following information when you place your cursor on a point in the graph:

- In the Application graph, the tooltip displays the date, time, traffic direction, utilization percentage for either the inbound or outbound traffic (depending on where your cursor is in the graph) for the specific time, and the threshold value for the interface.
- In the Packets graph, the tooltip displays the date, time, traffic direction, utilization percentage for either the inbound or outbound traffic (depending on where your cursor is in the graph) for the specific time, and the threshold value for the interface.

What To Do Next

[Interface Monitoring: View Interface Traffic Details](#)

[Interface Monitoring: View Top Application Traffic](#)

[Interface Monitoring: View the Packets Graph](#)

Related Tasks

[Interface Monitoring: Filter Results](#)

Related Topics

[Interface Monitoring: View Interface Traffic Details Overview](#)

[Interface Monitoring: Overview](#)

Interface Monitoring: Reports Overview

Run a flow search or top report to perform a high-level forensic investigation.

Related Tasks

[Interface Monitoring: Run a Flow Search Report](#)

[Interface Monitoring: Run a Top Report](#)

Related Topics

[Interface Monitoring: Report Filter Options](#)

[Interface Monitoring: Overview](#)

Interface Monitoring: Report Filter Options

When you run a flow search or a top report to perform a high-level forensic investigation, depending from where you run the report, it will filter on different values:

If you run a flow search or top report from the...	The flow search or top report filters on the following values...
Filter panel, Exporter column, or Interface column	Exporter or Interface: The selected exporter or interface (whichever applies). Time frame: Last five minutes Orientation: Either. Flow direction: All.
Top Application Traffic graph	Interface: The selected interface. Application: The selected application. Time frame: Based on the retention value that is set for your Manager. For more information, refer to the "Manager Properties: Data Retention Page" topic in the Desktop Client help. Orientation: Either. Flow direction: All.
Packets or Utilization graphs	Interface: The selected interface. Time frame: Based on the retention value that is set for your Manager. For more information, refer to the "Manager Properties: Data Retention Page" topic in the Desktop Client help. Orientation: Either. Flow direction: All.

What To Do Next

[Interface Monitoring: Run a Flow Search Report](#)

[Interface Monitoring: Run a Top Report](#)

Related Topics

[Interface Monitoring: Reports Overview](#)

[Interface Monitoring: Overview](#)

Interface Monitoring: Run a Flow Search Report

Run a flow search filtered by an exporter

1. In the Exporter column, click the **Ellipsis** icon.

If you are editing the flow search before running it and have selected the Edit option using one of the above-mentioned locations, note that under the Exporters & Interfaces heading in the Connection column on the Flow Search page, the IP address of the exporter for which you are running the flow search will be displayed in a green text box to indicate that the search will be filtered by this exporter (if you make no change to this setting).

2. Click **View Flows** to run the report with the default settings, or click **Edit** if you want to edit any settings before running the flow search.

Run a flow search filtered by an interface

1. Do one of the following:

- In the Interface column, click the **Ellipsis** icon.
- In the Details section that displays when you expand a row, click a point in the graph and then click the **Ellipsis** icon.

If you are editing the flow search before running it and have selected the Edit option using one of the above-mentioned locations, note that under the Exporters & Interfaces heading in the Connection column on the Flow Search page, the IP address of the interface for which you are running the flow search will be displayed in a green text box to indicate that the search will be filtered by this interface (if you make no change to this setting).

2. Click **View Flows** to run the report with the default settings, or click **Edit** if you want to edit any settings before running the flow search.

Related Tasks

[Interface Monitoring: Run a Top Report](#)

Related Topics

[Interface Monitoring: Reports Overview](#)

[Interface Monitoring: Report Filter Options](#)

[Interface Monitoring: Overview](#)

Interface Monitoring: Run a Top Report

1. To run a flow search filtered by an interface from the Interface Status page, do one of the following:

- In the Interface column, click the **Ellipsis** icon.
- In the Details section that displays when you expand a row, click a point in the graph and then click the **Ellipsis** icon.

If you are editing the top report before running it, note that under the Interfaces heading in the Advanced Options section on the [Top Report Name] Search page, the IP address of the interface for which you are running the top report will be displayed in a green text box to indicate that the top report will be filtered by this interface (if you make no change to this setting).

2. Click **Top Report > [report name]** to run the report with the default settings, or click **Edit** if you want to edit any settings before running the top report.

You can run the following top reports:

Run this top report...	To determine which...
Top Applications	Applications have been used the most within a specified time period by the searched-on host. You could run this report to help determine what role a segment of the network is playing. For example, if the most used application for the Host or Host Group is SQL, you can determine if it is most likely a database cluster.
Top Conversations	Conversations containing the largest number of bytes the searched-on host has been involved with within a specified period of time. You could run the Top Hosts report to identify systems that are utilizing the most bandwidth, and then from there run the Top Conversations report to see which conversations contained the most data transfer.
Top Hosts	Hosts the searched-on host has communicated with the most within a specified time period of time. You could run this report to identify the servers that are running on a particular port, or to identify systems that are utilizing the most bandwidth.
Top Interfaces	

Top Peers	Peers the searched-on host has communicated with the most within a specified time period of time. If users wanted to identify, for example, who China is connecting with the most in their network, they could run the Top Peers report.
Top Ports	Ports have been used the most within a specified time period by the searched-on host. For Network Segmentation and maintaining appropriate firewall rules, customers need visibility on which servers are in their data centers and to what ports those servers are listening. In this case, users could run the Top Ports report to identify the ports that have been used within a specific time frame. From the report results, users could set or modify their firewall rules.
Top Protocols	Protocols have been used the most within a specified time period by the searched-on host. Use this report to identify which protocols are being used within a network.
Top Services	Services have been used the most within a specified time period by the searched-on host. If there is a giant spike in traffic and you suspect a DDoS attack, you could run a Top Services report. If DNS is the top service, then you could determine if there is an application-based DDoS attack occurring.

Related Tasks

[Interface Monitoring: Run a Flow Search Report](#)

Related Topics

[Interface Monitoring: Reports Overview](#)

[Interface Monitoring: Report Filter Options](#)

[Interface Monitoring: Overview](#)

Analytics: Roles

Use this page to view a list of all roles (in the Active Roles table), along with at least one associated device for each role, that were active on your system for a specific period of time. The entries that are displayed in the Active Roles table depend on the telemetry types your system is ingesting. You can also view a list of all roles that are currently inactive on your system (the Inactive Roles table). The default time frame for both tables is the last 7 days, though you can set the time frame to a maximum of 90 days.

- You can see the Roles page only if you've enabled Analytics.
- The following roles will always be inactive in Cisco Secure Network Analytics (and therefore will always be listed in the Inactive Roles table):



- AWS
- Azure
- GCP
- Any other Cloud telemetry-related roles

Open the Roles page

From the main menu, choose **Investigate > Assets > Roles**.

The Roles page opens.



After you enable Analytics, you need to wait 24 hours for all jobs to run before the Roles page displays information.

Set the time frame

The time frame for the results defaults to the last 7 days. To configure a different time frame, complete the following steps:

1. Click either **Start** or **End** text boxes.
The calendar opens.
2. Choose your desired start and end date from the calendar.
3. Click **Apply** to confirm and set the selected date range.
Click **Reset**, to revert to the default date settings.

View the results

To view the matching devices for a role, check the check box of the role listed in the Active Roles table and the matching sources and role names for this role are added to a table that opens to the right.

Each time you check the check box of the role to view its matching sources and role names, the matching sources and role names for previous selections continue to be displayed. To remove matching sources and role names for a role, uncheck the check box of the role listed in the Active Roles table.

To view details about a matching source, hover your cursor over the  (**Alert**) icon or the  (**Eye**) icon that is beside the applicable source. (If the source is associated with a minimum of one open alert, the **Alert** icon is displayed. If the source is not associated with an open alert, the **Eye** icon is displayed.)

If you click the menu to the right of a matching source IP address, you can choose from several options to access the following pages:

Choose...	To access this page...
Device	Device Report. In the Roles section you will see the roles associated with the device, and in the Profiles section you will see the processes or protocols associated with the roles.
Alerts	Alerts Summary
Observations	Observations by Device
Flow Analysis	Flow Search Results

Export results

To export the results on this page to a CSV file, click the  (**Download CSV**) icon.

Related Topics

[Analytics: Alerts Workflow](#)

[Analytics: Alerts FAQs](#)

[Analytics: Investigating Alerts](#)

[Analytics: Alert Details](#)

[Analytics: Alerts Summary](#)

[Analytics: Overview](#)

Report Builder

Use Report Builder to create and customize your reports and dashboards. We've provided templates for building your reports and parameters for defining your search criteria.

Report results are based on the following:

- **Configuration:** Certain reports require a Data Store domain or additional configuration. For details, refer to [Report Templates](#).
- **Data:** It may take 24 hours or more to collect data. For details, refer to [Report Templates](#).
- **Data Roles:** The available data is determined by your assigned user data role. For more information, refer to [Data Role Management: Data Role Configuration Overview](#).

Whether you run a report on a routine basis, or only when needed to investigate an issue, you can review the details by editing the query and/or changing the chart or table view.

Opening Report Builder

1. Log in to your Manager.
2. From the main menu, choose **Report > Report Builder**.

Reports

All reports (saved by you or other users) are listed on this page. Click the report name to open a saved report, or add a new report as follows:

- [Report Templates](#)
- [Creating a Report](#)
- [Creating a TrustSec Report](#)
- [Opening a Report](#)
- [Editing a Report](#)
- [Deleting a Report](#)

Dashboards

All dashboards (saved by you or other users) are listed on this page.

Use dashboards to group your critical reports or create your own Network Insights Dashboard. You can also quickly refresh the report data on your dashboards.

Configuring Dashboards on Home Page

- **Pin or Unpin:** To pin or unpin the dashboard from your Home page, click the  (**Pin**) icon on a saved dashboard.



The **Security** dashboard tab is pinned to the Home page by default and cannot be unpinned.

- **Customize:** To organize the dashboards on the Home page, click **Customize** button, when on the pinned dashboard tab.
 - Drag and drop to reorder your dashboard tabs.
 - Pin or unpin dashboards from your Home page.
 - Add more dashboards to your Home page.

Click the dashboard name to open a saved dashboard, or add a new dashboard as follows:

- [Creating a Report Dashboard](#)
- [Creating a Network Insights Dashboard](#)
- [Opening a Report](#)
- [Editing a Report](#)
- [Deleting a Report](#)
- [Sorting Table Columns](#)

Schedule History

All scheduled reports (saved by you or other users) are listed on this page.

If your report supports scheduling, you can create a schedule and email delivery list. For details, review the following:

- [Scheduling a Report](#)
- [Deleting a Scheduled Report](#)

Best Practices

To run reports efficiently, review the following:

- **Limit Total Reports:** Whether you are creating or editing a report, limit the total number of open reports.
- **Time Range Parameter:** If the report template includes custom time ranges, choose short time ranges. This will help maximize performance.
- **Include/Exclude Parameters:** If you select **Include** for a parameter (such as Applications), add at least 1 parameter to the field. Otherwise, the report will search all data in that category, and it will take a long time to run and use a large amount of resources.

Failover

If you have configured your Managers as a [failover pair](#), you can run reports on the secondary Manager. However, make sure you manage your reports on the primary Manager. Editing, saving, and deleting reports are not available on the secondary Manager.

Related Tasks

[Investigating the Details](#)

[Investigating TrustSec Report Details](#)

[Saving a Report](#)

[Exporting a Report](#)

[Scheduling a Report](#)

Related Topic

[Report Templates](#)

Reports

All reports (saved by you or other users) are listed on the **Report > Report Builder > Reports** page. For details, refer to [Report Builder](#).

Click the report name to open a saved report, or add a new report as follows:

- [Report Templates](#)
- [Creating a Report](#)
- [Creating a TrustSec Report](#)

Related Tasks

[Editing a Report](#)

[Exporting a Report](#)

[Investigating the Details](#)

[Investigating TrustSec Report Details](#)

[Scheduling a Report](#)

Related Topics

[Report Builder](#)

[Report Templates](#)

Report Templates

We've provided templates to build your reports, and the summary is shown here. For more information about each template, available parameters, and report results, select the report name.

Name (click for more information)	Description	Data Store Domain Required
Alarms Report	View alarms for a selected IP address.	
Alerts Report	View alerts from Analytics.	yes
Dashboard	Create a custom dashboard with multiple reports that are customized for your needs.	
DSCP Status Report	View traffic, bandwidth, and utilization for a selected interface.	
Endpoint Traffic (NVM) Report	View endpoint traffic from your Network Visibility Module (NVM).	yes
Firewall Log Collection Trend Report	View the collection trend for Security Analytics and Logging (OnPrem) firewall events received by your Flow Collectors.	yes
Firewall Log Database Ingest Trend Report	View the Security Analytics and Logging (OnPrem) firewall records written to your Data Store.	yes
Flow Collection Status Report	View errors or performance issues on exporters for a selected IP address.	
Flow Collection Trend by Exporter Report	View the collection trend for flows received by selected Flow Collectors and exporters. The results are shown by exporter.	
Flow Collection Trend by Flow Collector Report	View the collection trend for flows received by your Flow Collectors.	

Flow Database Ingest Trend Report	View the flow records written to your Flow Collector databases or Data Store.	
Flow Rate Usage Report	View how the Flow Rate license (FPS) is being utilized across different telemetry types.	
Host Group Application Traffic Report	View application traffic for a selected host group.	
Host Group Flow Traffic Report	View flow traffic between host groups for selected applications, services, or protocols.	
Interface Application Traffic Report	View application traffic for a selected interface.	
Interface Network Traffic Report	View network traffic for a selected interface.	
Interface Service Traffic Report	View service traffic for a selected interface.	
Network and Server Performance Report	View Round Trip Time and Server Response Time between hosts or host groups. To use this report, you need a Flow Sensor or exporter that generates this data.	
Network Insights Dashboard	The Network Insights Dashboard is a template from which you can create your own network dashboard that, by default, contains the following reports: • Host Group Application Traffic Report • Host Group Flow Traffic Report • Network and Server Performance Report	

	• Flow Collection Trend by Flow Collector Report • Flow Collection Trend by Exporter Report • Firewall Log Collection Trend Report • NVM Collection Trend Report	
NVM Collection Trend Report	View the collection trend for Network Visibility Module (NVM) flows received by your Flow Collectors.	yes
NVM Database Ingest Trend Report	View the Network Visibility Module (NVM) records written to your Data Store.	yes
Security Events Report	View security events for a selected IP address.	
System Alarms Report	View system alarms over time.	
TrustSec Analytics Report	View traffic volume between Security Group Tags (SGTs) and gain insights into exact application flows between SGTs.	
TrustSec Policy Analytics Report	View policy compliance, including possible violations of the ISE TrustSec policy, for selected security groups based on observed traffic analytics.	
Zeek Database Ingest Trend Report	View Zeek logging events received by your Flow Collectors.	yes
Zeek Log Collection Trend Report	View the collection trend for Zeek flows received by your Flow Collectors.	yes
Zeek Logs Report	View Zeek logging events received by your Flow Collectors.	yes

Related Tasks

Creating a Report

[Creating a TrustSec Report](#)

[Related Topics](#)

[Report Builder](#)

[Cisco Support](#)

Alarms Report

Use this report to review a summary of the active and inactive alarms. You can investigate alarms on each exporter for the Flow Collector you select.

Parameters

Parameter	Description
Active or First Start Active and/or Last Active Time	Select Active to review alarms that were active in the last 10 minutes. To select a time range, select First Start Active and/or Last Active Time , and select preset ranges or enter a custom range.
Flow Collector	Select a Secure Network Analytics Flow Collector in your network.
Host IP Address or Range	You can enter the following: • Multiple IP Addresses: For each address, type the IP address and press Enter. • IP Range: Type <IP address > - <range>. For example, 10.52.1.0-10 includes any IP address from 10.52.1.0 to 10.52.1.10
Host Group	Select 1 host group. Click Select , and check a host group check box.
Orientation	Specify if you want to return results for instances where the subject is the client, server, or either. (The client is the host that initiated the session.)
Alarms	Select Include or Exclude . There are two ways to add alarms to your query: • Click the field, and select an alarm name from the menu. • Type an alarm name in the field, and select it from the menu.

Columns

We've listed all of the columns that are available for this report in alphabetical order here. To hide columns or change the order, refer to [Investigating the Details](#).

Column Name	Description
Active	Indicates whether or not the alarm is active.
Active Time Start	The date and time of the alarm's first occurrence or the alarm condition became active.
Alarm	Name of the alarm category or individual alarm. This column can show multiple alarms.
Alarm ID	We assign a unique ID number to each alarm.
Appliance	The Flow Collector that received the data or reported the alarm.
Details	Additional information about the alarm.
Duration	The amount of time (in minutes) that the alarm was active.
Last Active Time	The status and duration of the alarm. If the alarm is inactive, it shows the date and time that the alarm was most recently active. It also shows the length of time that has passed since the host was last active.
Policy	The host policy or host group relationship policy that triggered the alarm.
Severity	The severity level assigned to the alarm: • Informational • Trivial • Minor • Major • Critical

Source	The host IP address that caused the alarm, system alarm, or security event.
Source Country	The country of the source host that caused the alarm or security event.
Source Device Type	The source host's device type as detected by Cisco Identity Services Engine.
Source Host Group Path	The hierarchy of all parent host groups that include the source host.
Source Host Groups	The host groups that include the source IP address.
Source MAC Address	The MAC address of the source host device and the company name associated with the host MAC address. This information is not available if the user is connected through a VPN. To access this data, add Cisco Identity Services Engine as a device.
Source User Name	The user name associated with the source host in the monitored network. To see this data, install an identity device such as Cisco Identity Services Engine, Palo Alto Networks firewall, or Cisco ASA.
Target	The host IP address that is showing the suspicious activity. The target is the victim of the condition that caused the alarm.
Target Country	The country of the target host. For alarms, note the following: • The Manager uses the ISO 3166 standard for country names. • This column also identifies whether the host is local or private (RFC 1918).

Target Device Type	The target host's device type as detected by Cisco Identity Services Engine.
Target Host Group Path	The hierarchy of all parent host groups that include the target host.
Target Host Groups	The host groups that include the target host.
Target Host Name	The name assigned to the target host, if available.
Target MAC Address	The MAC address of the target host device and the company name associated with the host MAC address. This information is not available if the user is connected through a VPN. To access this data, add Cisco Identity Services Engine as a device.
Target User Name	The user name associated with the target host in the monitored network. To see this data, install an identity device such as Cisco Identity Services Engine, Palo Alto Networks firewall, or Cisco ASA.

Related Tasks

[Creating a Report](#)

[Editing a Report](#)

[Investigating the Details](#)

[Saving a Report](#)

[Exporting a Report](#)

Related Topics

[Report Builder](#)

[Building Queries](#)

[Report Templates](#)

Alerts Report



Make sure Analytics is enabled. From the main menu, choose **Configure > Detection > Analytics**, then click *Analytics On* Analytics On .

Use the Alerts report to view and evaluate alerts from Analytics:

- What alert types?
- When did the alerts occur?

This report shows the alerts that have been triggered during the time range you specify. The report provides an overview of the alerts, including the associated Host IPs and devices.



Observations are not shown on this report. To view observations, go to **Monitor > Observations** from the main menu.

Requirements

- To receive data on this report, you need Secure Network Analytics with a Data Store deployment and Analytics enabled. For information and instructions, refer to the [Appliance Installation Guide \(Hardware or Virtual Edition\) and the System Configuration Guide](#).
- To run this report, select a Data Store domain. To change domains, refer to Select a Domain.

Parameters

Parameter	More Information
Time Range	If you choose Custom , select a short time range for maximum performance. If you enter a long time range, the report may take a long time to query the data.

Columns

We've listed all of the columns that are available for this report in alphabetical order here. To hide columns or change the order, refer to [Investigating the Details](#).

Column Name	Description
Alert	Name of the alert.
Created Time	Date and time of the first observation.
Description	Description of the alert.
Last Updated	Date and time of the latest observation.
MITRE ATT&CK Tactics	MITRE ATT&CK tactics, if applicable.
Natural Time	Provides a reference to relative time, such as 1 hour ago or 5 days ago, for example.
Priority	Alert type priority as configured through Configure > Alerts from the main menu.
Publish Time	Exact time stamp when an

	alert is generated and available within Secure Network Analytics.
Resolved	Indicates if the alert has been resolved: True (alert has been closed) or False (alert is still open).
Source	Host IP address that caused the alert.
Source ID	A unique identifier that indicates the source associated with the alert.
Tags	User-defined tags that are associated with the alert.
Text	Additional information about the alert, typically provides the alert name, the IP address, and

	the URL for further investigation.

Related Tasks

[Creating a Report](#)

[Editing a Report](#)

[Investigating the Details](#)

[Saving a Report](#)

[Exporting a Report](#)

Related Topics

[Analytics: Alert Prerequisite Chart](#)

[Analytics: Alert Details](#)

[Analytics: Alerts Summary](#)

[Report Builder](#)

[Report Templates](#)

DSCP Status Report

Use this report to review Differentiated Services Code Point (DSCP) status, which is useful for standard network information and reviewing the health of your network. For example, run this report and review the data to determine if your network needs more bandwidth.

- **Managing DSCP:** To adjust your configuration settings, refer to [DSCP Configuration: Manage DSCP Configurations Overview](#).
- **Quality of Service:** For more information, refer to [Implementing Quality of Service Policies with DSCP](#).

Parameters

Parameter	Description
Time Range	If you choose Custom , select a short time range for maximum performance. If you enter a long time range, the report may take a long time to query the data.
Flow Collector	Select a Secure Network Analytics Flow Collector in your network.
Exporter	Select an exporter.
Interface	Select an interface.

Columns

We've listed all of the columns that are available for this report in alphabetical order here. To hide columns or change the order, refer to [Investigating the Details](#).

Column Name	Description
Allocated Bandwidth Inbound	The allocated rate of bandwidth for inbound traffic (bits per second).
Allocated Bandwidth Outbound	The allocated rate of bandwidth for outbound traffic (in bits per second).

Current Allocated Utilization Inbound	The current percent of the allocated configuration consumed by inbound traffic for the DSCP interface.
Current Allocated Utilization Inbound & Outbound	The current percent of available bandwidth consumed by inbound traffic for the DSCP interface.
Current Allocated Utilization Outbound	The current percent of the allocated configuration that is consumed by outbound traffic for the DSCP interface.
Current Traffic Inbound	The rate the interface is receiving inbound traffic (in bits per second).
Current Traffic Outbound	The rate the interface is receiving outbound traffic (in bits per second).
DSCP	The plain-text description associated with the differentiated services code point (DSCP) and the DSCP value that is encoded into the IP packet header. This value determines the Per-Hop Behavior (PHB). If the DSCP value is zero (0), the traffic is classified for the default PHB because it did not meet the requirements of any other traffic class. The default PHB typically has best-effort forwarding characteristics.
Exporter	The IP address of the router or switch that is sending the data.
Flow Collector	The Flow Collector that is receiving the data.
Interface	The name of the interface that provided the data. If you configured the Manager with the exporter SNMP settings, this column shows data provided by the exporter. If not, the Manager assigns a name.
Interface Speed In	The speed (line speed) for inbound traffic on the interface (in bits per second).
Interface Speed Out	The speed (line speed) for outbound traffic on the interface

	(in bits per second).
Maximum Traffic Inbound	The peak rate at which inbound traffic (in bits per second) has been received by the interface since the archive hour.
Maximum Traffic Outbound	The peak rate at which outbound traffic (in bits per second) has been sent by the interface since the archive hour.
Maximum Utilization Inbound & Outbound	The value from the Utilization Summary Inbound or Utilization Summary Outbound column, whichever is greater.

Related Tasks

[Creating a Report](#)

[Editing a Report](#)

[Investigating the Details](#)

[Saving a Report](#)

[Exporting a Report](#)

Related Topics

[Report Builder](#)

[Report Templates](#)

Endpoint Traffic (NVM) Report

Use this report to review endpoint traffic from your Network Visibility Module (NVM). We collect user, device, application, location, and destination data so you can investigate what users are doing while they are on or off the network. You can investigate and monitor details such as:

- Users running malware
- Users running unapproved applications
- Users sending/receiving data of an unusual amount or frequency

Requirements

- To receive data on this report, you need Secure Network Analytics with a Data Store deployment. For information and instructions, refer to the [Appliance Installation Guide \(Hardware or Virtual Edition\)](#) and the [System Configuration Guide](#).
- Make sure your Flow Collector is configured to receive data from your Network Visibility Module. For instructions, refer to the [Network Visibility Module \(NVM\) Configuration Guide](#).

Parameters

Parameter	Description
Time Range	If you choose Custom , select a short time range for maximum performance. If you enter a long time range, the report may take a long time to query the data.
Host IP Address or Range	You can enter the following: • Multiple IP Addresses: For each address, type the IP address and press Enter. • IP Range: Type <IP address > - <range>. For example, 10.52.1.0-10 includes any IP address from 10.52.1.0 to 10.52.1.10


Report Results

Any time you run a new query in the Endpoint Traffic (NVM) report, it shows the most recent 300 records based on the [end time](#). To see more results, review the following:

- [Show More Results](#)
- [Filter Data](#)
- [Columns](#)

Show More Results

To show more results, scroll down. The number of records updates as you scroll through the report.



Endpoint Traffic (NVM)

Showing 300 Matching Results

scroll to see more results

End Time	Source IP Address	Source Port
9/10/2020 11:03:49 PM	9/10/2020 11:03:50 PM	10.0.0.1 5000
9/10/2020 11:03:50 PM	9/10/2020 11:03:50 PM	10.0.0.1 5001
9/10/2020 11:03:49 PM	9/10/2020 11:03:50 PM	10.0.0.1 6000

Filter Data

You can also filter the current records, or filter and query all data in the database. For instructions, refer to [Filtering the Endpoint Traffic \(NVM\) Report](#).

Columns

When you first run the report, we show a default set of columns. You can keep this view or choose specific columns to show or hide.

We've listed all of the columns that are available for this report in alphabetical order here. To hide columns or change the order, refer to [Investigating the Details](#).

Column Name	Description
Agent Version	The software version of the agent/client software.
Bytes In	The total number of incoming bytes on the flow at layer 4.
Bytes Out	The total number of outgoing bytes on the flow at layer 4.

Destination IP Address	The IP address of the host that received the communication.
Destination Port	The port number where the host received the communication.
DNS Suffix	The DNS suffix associated with the interface for the flow.
End Time	The date and time the flow ended.
Host Name	The FQDN (host name) that resolved the destination IP on the flow.
Interface Details List	The network interface details, which include service set identifier (SSID), state (trusted or untrusted), or if it is configured as a split tunnel.
Interface Index	The network interface index as reported by the operating system.
Interface Mac Address	The network interface MAC address.
Interface Name	The network interface name as reported by the operating system.
Interface Type	The network interface type (for example: wired, wireless, cellular, or VPN). We provide this information as a numeric value.
Interface UID	The unique identifier for the network interface metadata.
Last Time	The date and time the flow ended.
Module Hash List	The list of SHA256 hashes of modules associated with the Module Name List.
Module Name List	The list of names of modules hosted by the process that generated the flow.
Operating System Edition	The operating system edition of the endpoint device.

Operating System Name	The operating system name of the endpoint device.
Operating System Version	The operating system version on the endpoint device.
Parent Process Account	The account that initiated the parent process.
Parent Process Account Type	The parent process account type.
Parent Process Args (Arguments)	The command line arguments associated with the parent process.
Parent Process Hash	The SHA256 hash of the parent process.
Parent Process ID	The process identifier of the parent process.
Parent Process Name	The parent name of the process associated with the flow. For example: iexplorer.exe, chrome.exe, etc.
Parent Process Path	The complete file path of the parent process.
Process Account	The account that initiated the process associated to the flow.
Process Account Type	The process account type.
Process Args (Arguments)	The process command line arguments associated with the flow.
Process Hash	The SHA256 hash of the process associated with the flow.
Process ID	The unique identifier of the process.
Process Name	The process name associated with the flow. For example: chrome.exe.
Process Path	The complete file path of the process.
Protocol	The protocol used in the flow communication.
Source IP Address	The IP address of the host sending the communication. It

	can be an inside or outside host.
Source Port	The port number that received the communication.
Start Time	The date and time the flow started.
System Manufacturer	The endpoint device manufacturer name.
System Type	The endpoint device system type.
UDID	The unique identifier of the endpoint device.
User	The user logged in to the physical device. If this field is empty, the user was not logged in to the device or the user logged in remotely (via RDP, SSH, etc.).
User Account Type	The account type of the logged-in user.
Virtual Station Name	The endpoint device name.

Related Tasks

[Creating a Report](#)

[Editing a Report](#)

[Filtering the Endpoint Traffic \(NVM\) Report](#)

[Investigating the Details](#)

[Saving a Report](#)

[Exporting an Endpoint Traffic \(NVM\) Report](#)

Related Topics

[Report Builder](#)

[Building Queries](#)

[Report Templates](#)

Firewall Log Collection Trend Report

Use the Firewall Log Collection Trend report to view the collection trend for Cisco Security Analytics and Logging (On Premises) firewall events received by your Flow Collectors, which can help you evaluate:

- Are all Flow Collectors receiving flows?
- What is the quantity of flows received?
- Were there any interruptions?

This report is important for troubleshooting and capacity planning as follows:

- **Initial Deployments:** As you deploy new Flow Collectors, run this report to verify they are ingesting the expected amount of telemetry.
- **Troubleshooting:** Review the [report results](#) to confirm your Flow Collectors are receiving Security Analytics and Logging (OnPrem) flows, review the collection rate, and find anomalies in telemetry ingest for an individual Flow Collector or multiple Flow Collectors. Also, compare the current collection rates with historical rates to identify trends in your data.

To confirm your Data Store is receiving Security Analytics and Logging (OnPrem) flows, run the [Firewall Log Database Ingest Trend Report](#).

- **Capacity Planning:** Review the [report results](#) to evaluate the telemetry ingest trends over time. Use the growth trends you observe to plan for future expansion and ensure your Flow Collectors are not overloaded based on their specifications.

Requirements

- To receive data on this report, you need Secure Network Analytics with a Data Store deployment. For information and instructions, refer to the [Appliance Installation Guide \(Hardware or Virtual Edition\)](#) and the [System Configuration Guide](#).
- Make sure your Flow Collector is configured to receive data from Security Analytics and Logging (OnPrem). For instructions, refer to the [Cisco Security Analytics and Logging \(On Premises\): Firewall Event Integration Guide](#).
- To run this report, select a Data Store domain. For more information, refer to Select a Domain.

Parameters

Parameter	More Information

Time Range	If you choose Custom , select a short time range for maximum performance. If you enter a long time range, the report may take a long time to query the data.

Report Results

The Firewall Log Collection Trend report shows your query results by Flow Collector.

- **Filter:** Click individual Flow Collectors in the legend. You can also change the y-axis criteria to events per second (EPS) or bytes.
- **Zoom In:** To focus on a section of the chart, [zoom in](#) to see the details.

For more information, refer to [Investigating the Details](#).

Related Tasks

[Creating a Report](#)

[Editing a Report](#)

[Investigating the Details](#)

[Saving a Report](#)

[Exporting a Report](#)

Related Topics

[Report Builder](#)

[Report Templates](#)

Firewall Log Database Ingest Trend Report

Use the Firewall Log Database Ingest Trend report to view Cisco Security Analytics and Logging (On Premises) firewall records written to your Data Store, which can help you evaluate:

- Is the Data Store receiving flows?
- What is the quantity of flows received?
- Were there any interruptions?

This report is important for troubleshooting and capacity planning as follows:

- **Initial Deployments:** As you deploy a new Data Store, run this report to verify it is ingesting the expected amount of telemetry.
- **Troubleshooting:** Review the [report results](#) to confirm your Data Store is receiving Security Analytics and Logging (OnPrem) flows and find anomalies in telemetry ingest. Also, compare the current collection rates with historical rates to identify trends in your data.

To confirm your Flow Collectors are receiving Security Analytics and Logging (OnPrem) flows, run the **Firewall Log Collection Trend Report**.

- **Capacity Planning:** Review the [report results](#) to evaluate the telemetry ingest trends over time. Use the growth trends you observe to plan for future expansion and ensure your Data Store is not overloaded based on the specifications.

Requirements

- To receive data on this report, you need Secure Network Analytics with a Data Store deployment. For information and instructions, refer to the [Appliance Installation Guide \(Hardware or Virtual Edition\)](#) and the [System Configuration Guide](#).
- Make sure your Flow Collector is configured to receive data from Security Analytics and Logging (OnPrem). For instructions, refer to the [Cisco Security Analytics and Logging \(On Premises\): Firewall Event Integration Guide](#).
- To run this report, select a Data Store domain. For more information, refer to [Select a Domain](#).

Parameters

Parameter	More Information

Time Range	If you choose Custom , select a short time range for maximum performance. If you enter a long time range, the report may take a long time to query the data.

Report Results

The Firewall Log Database Ingest Trend report shows your query results by Flow Collector.

- **Filter:** Click individual Flow Collectors in the legend.
- **Zoom In:** To focus on a section of the chart, [zoom in](#) to see the details.

For more information, refer to [Investigating the Details](#).

Related Tasks

[Creating a Report](#)

[Editing a Report](#)

[Investigating the Details](#)

[Saving a Report](#)

[Exporting a Report](#)

Related Topics

[Report Builder](#)

[Report Templates](#)

Flow Collection Status Report

Use this report to check for errors and performance issues on your Flow Collector. You can investigate issues by reviewing the status and data on each exporter for the Flow Collector you select.

Parameters

Parameter	Description
Date	Select Today, or select a specific date to search.
Flow Collector	Select a Secure Network Analytics Flow Collector in your network.

Report Results

You can identify errors in your data by reviewing each [column](#) and [Common Scenarios](#).

- **Error:** If **Error** is shown, hover your cursor over it for more information.
- **Configuration Information:** If you see errors in your data, you may need to adjust your configuration settings. Refer to [Configuration Information](#) for details.

Status	Exporter	Flow Type	Average Flow Rate (fps)	Average NetFlow Traffic (bps)
Error	10.1	NETFLOW_V9	846	560.78K
		NETFLOW_V9	83	271.94K
		NETFLOW_V9	82	41.81K
Error	10.1	NETFLOW_V9	69	51.15K

Common Scenarios

Review the following data in the report:

- **Exporter Status:** If Error is shown, hover your cursor over it for more information. If the error message indicates an **Unrealistic** error, correct the exporter configuration. For example, if there is an Unrealistic Start Time and Stop Time, where the Start Time is after the Stop Time, configure your exporter and Flow Collector so they are

using the same NTP server.

- **Utilization:** If Utilization Inbound and Utilization Outbound are incorrect, check your network speed configuration settings on your Flow Collector port.
- **Duplicate Flows:** If the exporter is processing duplicate flows, and it is not a backup device, review your configuration.
- **Longest Duration Export:** Confirm the range is 60 to 120 seconds. If the range is exceeded, there may be a backlog of data, and the Flow Collector is not processing it fast enough. To correct the configuration, decrease the exporter **Timeout** setting.
- **SNMP Status:** If Failed is shown, it means that the Manager cannot poll the exporter.

Columns

We've listed all of the columns that are available for this report in alphabetical order here. To hide columns or change the order, refer to [Investigating the Details](#).

Column Name	Description
Active Today	The exporter activity status. If Yes is shown, data has been received from the exporter since the end of the last hour.
Average Flow Rate	The average rate that flows have been received from the exporter at the beginning of each hour.
Average NetFlow Traffic	The average rate that NetFlow packets have been received from the exporter at the beginning of each hour.
Current Flow Rate	The rate that flows are being received from the exporter.
Current NetFlow Traffic	The rate that NetFlow packets have been received from the exporter. This is based on the time period when you last ran the report.
Duplicate Flows	The total number of packets received from the exporter at the beginning of each hour that has already been received.
Duplicate Flows %	The percentage of data received from the exporter that has already received.

Exporter	The IP address of the router or switch that is sending the data.
Flow Type	The version of NetFlow data that the exporter is sending.
Interface Count	The number of interfaces on the exporter that is sending data to the Flow Collector.
Last Export	The date and time that the most recent data was received from the exporter.
Longest Duration Export	The total length of time (in seconds) that the flow with the longest duration was active from the first packet to the last packet. This value is updated at the beginning of each hour and capped at 1,800 seconds (30 minutes). If 1,800 is shown, then flows of longer duration may have been found based on the flow start and end time stamps received in the NetFlow data.
Maximum Flow Rate	The peak rate of flows received from the exporter at the beginning of each hour.
Maximum NetFlow traffic	The total number of NetFlow packets received from the exporter at the beginning of each hour.
NetFlow Packets	The total number of NetFlow packets received from the exporter at the beginning of each hour.
SNMP Status	The status of the Manager to query the exporter using SNMP. For Flow Sensors, the status is n/a.
SNMP Status Date	The date of the last time the Manager tried to query the exporter using SNMP.
SNMP Status Error	There is an error in the Manager query to the exporter using SNMP.
Status	The status of the data.

Total Flows	The total number of flows received from the exporter at the beginning of each hour.

Configuration Information

If you see errors in your data, you may need to adjust your configuration settings. Refer to the following:

- **Exporter Configuration:** Select **Configure > System > Exporters > Exporter SNMP Profiles** tab.
- **Flow Collector Configuration:** Select **Configure > System > Flow Collectors**.

Related Tasks

[Creating a Report](#)

[Editing a Report](#)

[Investigating the Details](#)

[Saving a Report](#)

[Exporting a Report](#)

Related Topics

[Report Builder](#)

[Report Templates](#)

Flow Collection Trend by Exporter Report

Use the Flow Collection Trend by Exporter report to view the collection trend for flows received by the Flow Collectors and exporters you select, which can help you evaluate:

- Are all Flow Collectors and/or exporters receiving flows?
- What is the quantity of flows received?
- Were there any interruptions?

This report is important for troubleshooting and capacity planning as follows:

- **Initial Deployments:** As you deploy new exporters, run this report to verify the Flow Collectors are ingesting the expected amount of telemetry.
- **Troubleshooting:** Review the [report results](#) to find anomalies in telemetry ingest for an individual exporter or multiple exporters. Also, compare the current collection rates with historical rates to identify trends in your data.

To view the data by Flow Collectors only, run the [Flow Collection Trend by Flow Collector Report](#).

- **Capacity Planning:** Review the [report results](#) to evaluate the telemetry ingest trends over time. Use the growth trends you observe to plan for future expansion and ensure your exporters and Flow Collectors are not overloaded based on their specifications.

Parameters

Parameter	More Information
Time Range	If you choose Custom , select a short time range for maximum performance. If you enter a long time range, the report may take a long time to query the data.
Flow Collectors, Exporters	Select Flow Collectors and/or exporters in your domain.

Report Results

You can filter your results by selecting specific exporters. We've listed all of the columns that are available for this report in alphabetical order here. To filter, hide columns, or change the order, refer to [Investigating the Details](#).

Column Name	Description
Date/Time	The date and time the data occurred.
Exporter	The IP address of the router or switch that is sending the data.
Flow Collector	The Flow Collector that is receiving the data.
Flow Rate	The rate that the Flow Collector is receiving flows from the exporters and Flow Sensors.
Flow Count	The total number of flows received from the exporter since the archive hour.

Related Tasks

[Creating a Report](#)

[Editing a Report](#)

[Investigating the Details](#)

[Saving a Report](#)

[Exporting a Report](#)

Related Topics

[Report Builder](#)

[Report Templates](#)

Flow Collection Trend by Flow Collector Report

Use the Flow Collection Trend by Flow Collector report to view the collection trend for flows received by your Flow Collectors, which can help you evaluate:

- Are all Flow Collectors receiving flows?
- What is the quantity of flows received?
- Were there any interruptions?

This report is important for troubleshooting and capacity planning as follows:

- **Initial Deployments:** As you deploy new Flow Collectors, run this report to verify they are ingesting the expected amount of telemetry.
- **Troubleshooting:** Review the [report results](#) to find anomalies in telemetry ingest for an individual Flow Collector or multiple Flow Collectors. Also, compare the current collection rates with historical rates to identify trends in your data.

To confirm your Flow Collectors and exporters are receiving flows, run the [Flow Collection Trend by Exporter Report](#).

- **Capacity Planning:** Review the [report results](#) to evaluate the telemetry ingest trends over time. Use the growth trends you observe to plan for future expansion and ensure your Flow Collectors are not overloaded based on their specifications.

Parameters

Parameter	More Information
Time Range	If you choose Custom , select a short time range for maximum performance. If you enter a long time range, the report may take a long time to query the data.

Report Results

You can filter your results by selecting specific Flow Collectors. We've listed all of the columns that are available for this report in alphabetical order here. To filter, hide columns, or change the order, refer to [Investigating the Details](#).

Column Name	Description
Date/Time	The date and time the data occurred.
Flow Collector	The Flow Collector that is receiving the data.
Flow Rate	The rate that the Flow Collector is receiving flows from the exporters and Flow Sensors.
Flow Count	The total number of flows received from the exporter since the archive hour.

Related Tasks

[Creating a Report](#)

[Editing a Report](#)

[Investigating the Details](#)

[Saving a Report](#)

[Exporting a Report](#)

Related Topics

[Report Builder](#)

[Report Templates](#)

Flow Database Ingest Trend Report

Use the Flow Database Ingest Trend report to view flow records written to your Data Store or Flow Collector database, which can help you evaluate:

Data Store (Data Store domain required)

- Is the Data Store receiving flows?
- What is the quantity of flows received?
- Were there any interruptions?

Flow Collector Database (Non-Data Store domain required)

- Is the Flow Collector database receiving flows?
- What is the quantity of flows received?
- Were there any interruptions?

This report is important for troubleshooting and capacity planning as follows:

- **Initial Deployments:** As you deploy a new Data Store or Flow Collector database, run this report to verify it is ingesting the expected amount of telemetry.
- **Troubleshooting:** Review the [report results](#) to confirm your Data Store or Flow Collector database is receiving flows and find anomalies in telemetry ingest. Also, compare the current collection rates with historical rates to identify trends in your data.
- **Capacity Planning:** Review the [report results](#) to evaluate the telemetry ingest trends over time. Use the growth trends you observe to plan for future expansion and ensure your Data Store or Flow Collector database is not overloaded based on the specifications.

Requirements

- To receive Data Store data on this report, you need Secure Network Analytics with a Data Store deployment. For information and instructions, refer to the [Appliance Installation Guide \(Hardware or Virtual Edition\)](#) and the [System Configuration Guide](#).
- To change domains, refer to Select a Domain.

Parameters

Parameter	More Information

Time Range	If you choose Custom , select a short time range for maximum performance. If you enter a long time range, the report may take a long time to query the data.

Report Results

The Flow Database Ingest Trend report shows your query results by Flow Collector.

- **Filter:** Click individual Flow Collectors in the legend.
- **Zoom In:** To focus on a section of the chart, [zoom in](#) to see the details.

For more information, refer to [Investigating the Details](#).

Related Tasks

[Creating a Report](#)

[Editing a Report](#)

[Investigating the Details](#)

[Saving a Report](#)

[Exporting a Report](#)

Related Topics

[Report Builder](#)

[Report Templates](#)

Flow Rate Usage Report

Use the Flow Rate Usage Report to track how your Flow Rate license (FPS) is being used across different telemetry types, including NetFlow, Security Analytics and Logging (OnPrem), NVM, and Zeek.

- This report consolidates telemetry usage data into a single interface, making it easier for you to track your consumption, understand your telemetry mix, and ensure you remain within the bounds of your purchased FPS license.
- Whether you're troubleshooting a sudden spike in usage or planning for future capacity, the detailed insights provided by the report allows you to make informed decisions about your telemetry ingestion strategy.
- Secure Network Analytics reports the 95th percentile of the daily Flow Rate usage to your Smart Account. This value helps you monitor your usage to compare it with your licensed limits.

Parameters

Parameter	Description
Day Range	Select a number of days or choose Custom to define the start date and end date. This report provides historical data, with the most recent information covering the previous day's activity.

Report Results

The Flow Rate Usage report shows your query results.

- **Filter:** Click individual telemetry type in the legend.
- **Zoom In:** To focus on a section of the chart, [zoom in](#) to see the details.

Each bar represents an aggregation of actual FPS over a specific time window, determined by the selected Day Range. The 95th FPS value remains constant for each day within the chosen range. The 95th FPS represents the flow rate reported to Smart Licensing as your daily consumption.

Columns

We've listed all of the columns that are available for this report in alphabetical order here. To hide columns or change the order, refer to [Investigating the Details](#).

Column Name	Description
Date/Time	The date and time the data occurred.
FPS	The rate that the Flow Collector is receiving telemetry from its source(s).
95th FPS	95th percentile flow rate reported to Smart Licensing as the your consumption for the day.
Maximum FPS	Peak flow rate recorded during the day.
Telemetry	Sources of network data collected, such as NetFlow, NVM, Zeek, and Security Analytics and Logging (OnPrem)

Related Tasks

[Creating a Report](#)

[Editing a Report](#)

[Investigating the Details](#)

[Saving a Report](#)

[Exporting a Report](#)

Related Topics

[Report Builder](#)

[Report Templates](#)

Host Group Application Traffic Report

Use this report to review the application traffic for the host group you select. You can choose to include or exclude applications.

This is a good report to monitor data on a daily basis and see a broad overview. If you see a spike in the data, you can focus on the details and determine if there is an issue. You can also use this report to confirm that traffic is not showing any issues over time, which would indicate a problem (such as a virus) has been resolved.

Most issues are typically both application and service issues. To review application and service traffic, refer to [Host Group Flow Traffic Report](#).

Parameters

Parameter	Description
Time Range	If you choose Custom , select a short time range for maximum performance. If you enter a long time range, the report may take a long time to query the data.
Host Group	Select 1 host group. Click Select , and check a host group check box.
Applications	Select Include or Exclude. There are two ways to add applications to your query: - Click the field, and select an application name from the menu. - Type an application name in the field, and select it from the menu. We include default applications, and you can add custom applications at Configure > System > Applications.

Columns

We've listed all of the columns that are available for this report in alphabetical order here. To hide columns or change the order, refer to [Investigating the Details](#).

Column Name	Description

Time	The time the data was recorded.
Application	The selected application. We include default applications, and you can add custom applications for your network at Configure > System > Applications. Several applications are defined by default as unclassified to indicate they are using the traditional ports/protocols for specific applications. This list also includes Undefined TCP and Undefined UDP applications, with default ranges that encompass the entire possible range for the TCP and UDP protocols.
Traffic Within (bps)	The rate of all traffic using the application (during the 5 minute time interval shown in the Time column) that is: - received by hosts in the selected host group (and its subordinate host groups) from other hosts within this same host group (and its subordinate host groups) - sent by host in the selected host group (and its subordinate host groups) to other hosts within this same host group (and its subordinate host groups)
Traffic Inbound (bps)	The total rate of traffic received by the host group (and its subordinate host groups) during the 5 minute time interval shown in the Time column.
Traffic Outbound (bps)	The total rate of traffic sent by the host group (and its subordinate host groups) during the 5 minute time interval shown in the Time column.

Related Tasks

[Creating a Report](#)

[Editing a Report](#)

[Investigating the Details](#)

[Saving a Report](#)

[Exporting a Report](#)

Related Topics

[Report Builder](#)

[Building Queries](#)

[Report Templates](#)

Host Group Flow Traffic Report

Use this report to review host group flow traffic for a selected host group. Refine your search by including or excluding specific host groups. You can also include or exclude applications, services, and protocols.

If you see a spike in the data, focus on the details and determine if there is an issue. You can also confirm that traffic is not showing any issues over time, which would indicate a problem (such as a virus) has been resolved.

To focus on application traffic by host group, refer to [Host Group Application Traffic Report](#).

Parameters

Parameter	Description
Time Range	If you choose Custom , select a short time range for maximum performance. If you enter a long time range, the report may take a long time to query the data.
Flow Collector	Select a Secure Network Analytics Flow Collector in your network.
Host IP Address or Range	You can enter the following: • Multiple IP Addresses: For each address, type the IP address and press Enter. • IP Range: Type <IP address > - <range>. For example, 10.52.1.0-10 includes any IP address from 10.52.1.0 to 10.52.1.10
Host Group	Select 1 host group. Click Select , and check a host group check box.
Orientation	Specify if you want to return results for instances where the subject is the client, server, or either. (The client is the host that initiated the session.)
Applications	Select Include or Exclude . There are two ways to add applications to your query:

	- Click the field, and select an application name from the menu. - Type an application name in the field, and select it from the menu. We include default applications, and you can add custom applications at Configure > System > Applications.
Services	Select Include or Exclude. There are two ways to add services to your query: - Click the field, and select a service name from the menu. - Type a service name in the field, and select it from the menu. We include default services, and you can add custom services at Configure > System > Services.
Protocols	Select Include or Exclude. There are two ways to add protocols to your query: - Click the field, and select a protocol name from the menu. - Type a protocol name in the field, and select it from the menu.

Columns

We've listed all of the columns that are available for this report in alphabetical order here. To hide columns or change the order, refer to [Investigating the Details](#).

Column Name	Description
Date/Time	The date and time the data occurred.
Traffic Inbound	The total rate of traffic received by the host group (and its subordinate host groups) during the 1 minute time interval shown in the Date/Time column.

Traffic Outbound	The total rate of traffic sent by the host group (and its subordinate host groups) during the 1 minute time interval shown in the Date/Time column.

Related Tasks

[Creating a Report](#)

[Editing a Report](#)

[Investigating the Details](#)

[Saving a Report](#)

[Exporting a Report](#)

Related Topics

[Report Builder](#)

[Building Queries](#)

[Report Templates](#)

Interface Application Traffic Report

Use this report to review the application traffic for the exporter and interface you select.

Most issues are typically both application and service issues. To review interface service traffic, refer to [Interface Service Traffic Report](#).

Parameters

Parameter	Description
Time Range	If you choose Custom , select a short time range for maximum performance. If you enter a long time range, the report may take a long time to query the data.
Flow Collector	Select a Secure Network Analytics Flow Collector in your network.
Exporter	Select an exporter.
Interface	Select an interface.
Applications	Select Include or Exclude. There are two ways to add applications to your query: - Click the field, and select an application name from the menu. - Type an application name in the field, and select it from the menu. We include default applications, and you can add custom applications at Configure > System > Applications.

Columns

We've listed all of the columns that are available for this report in alphabetical order here. To hide columns or change the order, refer to [Investigating the Details](#).

Column Name	Description
Application	The selected application. We include default applications,

	and you can add custom applications for your network at Configure > System > Applications. Several applications are defined by default as unclassified to indicate they are using the traditional ports/protocols for specific applications. This list also includes Undefined TCP and Undefined UDP applications, with default ranges that encompass the entire possible range for the TCP and UDP protocols.
Date/Time	The date and time the data occurred.
Traffic Inbound	The total rate of traffic received by the host group (and its subordinate host groups) during the 1 minute time interval shown in the Date/Time column.
Traffic Outbound	The total rate of traffic sent by the host group (and its subordinate host groups) during the 1 minute time interval shown in the Date/Time column.

Related Tasks

[Creating a Report](#)

[Editing a Report](#)

[Investigating the Details](#)

[Saving a Report](#)

[Exporting a Report](#)

Related Topics

[Report Builder](#)

[Building Queries](#)

[Report Templates](#)

Interface Network Traffic Report

Use this report to review the network traffic for the exporter and interface you select.

Parameters

Parameter	Description
Time Range	If you choose Custom , select a short time range for maximum performance. If you enter a long time range, the report may take a long time to query the data.
Flow Collector	Select a Secure Network Analytics Flow Collector in your network.
Exporter	Select an exporter.
Interface	Select an interface.

Columns

We've listed all of the columns that are available for this report in alphabetical order here. To hide columns or change the order, refer to [Investigating the Details](#).

Column Name	Description
Date/Time	The date and time the data occurred.
Interface Inbound Threshold	The threshold of the inbound exporter interface speed that triggers the Interface Utilization Exceeded In alarm. To enable or disable alarms, refer to Alarm List: Exporter/Interface Alarms . To view the alarms, go to Investigate > Assets > Interfaces or Investigate > Flow Search > Network Diagrams .
Interface Outbound Threshold	The threshold of the outbound exporter interface speed that triggers the Interface Utilization Exceeded Out alarm. To enable or disable alarms, refer to Alarm List: Exporter/Interface Alarms . To view the alarms, go to Investigate > Assets > Interfaces or Investigate > Flow Search > Network Diagrams .

	Search > Network Diagrams.
Packets Inbound	The rate that inbound packets are being transferred. The default speed is 1G. However, the speed depends on your network card.
Packets Outbound	The rate that outbound packets are being transferred. The default speed is 1G. However, the speed depends on your network card.
Speed Inbound	The speed that inbound traffic flows are being transferred.
Speed Outbound	The speed that outbound traffic flows are being transferred.
Traffic Inbound	The total rate of traffic received by the host group (and its subordinate host groups) during the 1 minute time interval shown in the Date/Time column.
Traffic Inbound Threshold (bps)	The threshold of the inbound traffic on the interface in bytes per second.
Traffic Outbound	The total rate of traffic sent by the host group (and its subordinate host groups) during the 1 minute time interval shown in the Date/Time column.
Traffic Outbound Threshold (bps)	The threshold of the outbound traffic on the interface in bytes per second.
Utilization Inbound %	The percentage of inbound traffic utilization on the selected interface of the selected exporter.
Utilization Outbound %	The percentage of outbound traffic utilization on the selected interface of the selected exporter.

Related Tasks

[Creating a Report](#)

[Editing a Report](#)

[Investigating the Details](#)

[Saving a Report](#)

[Exporting a Report](#)

[Related Topics](#)

[Report Builder](#)

[Building Queries](#)

[Report Templates](#)

Interface Service Traffic Report

Use this report to review the service traffic for the exporter and interface you select.

Most issues are typically both application and service issues. To review interface application traffic, refer to [Interface Application Traffic Report](#).

Parameters

Parameter	Description
Time Range	If you choose Custom , select a short time range for maximum performance. If you enter a long time range, the report may take a long time to query the data.
Flow Collector	Select a Secure Network Analytics Flow Collector in your network.
Exporter	Select an exporter.
Interface	Select an interface.
Services	We include default services, and you can add custom applications for your network at Configure > System > Services . For example, you may want to add a service and include it in the File Sharing Index.

Columns

We've listed all of the columns that are available for this report in alphabetical order here. To hide columns or change the order, refer to [Investigating the Details](#).

Column Name	Description
Date/Time	The date and time the data occurred.
Service	The service reporting the data.
Traffic Inbound	The total rate of traffic received by the host group (and its subordinate host groups) during the 1 minute time interval shown in the Date/Time column.

Traffic Outbound	The total rate of traffic sent by the host group (and its subordinate host groups) during the 1 minute time interval shown in the Date/Time column.

Related Tasks

[Creating a Report](#)

[Editing a Report](#)

[Investigating the Details](#)

[Saving a Report](#)

[Exporting a Report](#)

Related Topics

[Report Builder](#)

[Building Queries](#)

[Report Templates](#)

Network and Server Performance Report

Use this report to review performance on your Flow Collectors, Flow Sensors, and exporters. For example, if you see that the review round trip time (RTT) is high or increasing, it could indicate an outage.

Requirements: To run this report, you need a Flow Sensor and exporter in your Secure Network Analytics network with round trip time (RTT) and server response time (SRT) data.

Parameters

Parameter	Description
Time Range	If you choose Custom , select a short time range for maximum performance. If you enter a long time range, the report may take a long time to query the data.
Flow Collector	Select a Secure Network Analytics Flow Collector in your network.
Host IP Address or Range	You can enter the following: • Multiple IP Addresses: For each address, type the IP address and press Enter. • IP Range: Type <IP address > - <range>. For example, 10.52.1.0-10 includes any IP address from 10.52.1.0 to 10.52.1.10
Host Group	Select 1 host group. Click Select , and check a host group check box.
Orientation	Specify if you want to return results for instances where the subject is the client, server, or either. (The client is the host that initiated the session.)
Applications	Select Include or Exclude . There are two ways to add applications to your query: • Click the field, and select an application name from the menu.

	Type an application name in the field, and select it from the menu. We include default applications, and you can add custom applications at Configure > System > Applications.
Services	Select Include or Exclude. There are two ways to add services to your query: - Click the field, and select a service name from the menu. - Type a service name in the field, and select it from the menu. We include default services, and you can add custom services at Configure > System > Services.
Protocols	Select Include or Exclude. There are two ways to add protocols to your query: - Click the field, and select a protocol name from the menu. - Type a protocol name in the field, and select it from the menu.

Columns

We've listed all of the columns that are available for this report in alphabetical order here. To hide columns or change the order, refer to [Investigating the Details](#).

Column Name	Description
Average Round Trip	The average amount of time (in milliseconds) required for all of the TCP connections to occur in each flow for each minute the flow is active.
Average Server Response	The average amount of time (in milliseconds) between the first client request and the first server data response for all of the TCP connections in each flow for each minute the

	flow is active.
Device	The Flow Collector that is receiving the data.
Max Round Trip	The average of the least amount of time (in milliseconds) required for all of the TCP connections to occur in each flow for each minute the flow is active.
Max Server Response	The average of the least amount of time (in milliseconds) required between the first client data set and the first server data response for all of the TCP connections in each flow for each minute the flow is active.
Min Round Trip	The average of the least amount of time (in milliseconds) required for all of the TCP connections to occur in each flow for each minute the flow is active.
Min Server Response	The average of the least amount of time (in milliseconds) required between the first client data set and the first server data response for all of the TCP connections in each flow for each minute the flow is active.
Retransmit Rate	The number of TCP packets that were retransmitted during the flow. This usually occurs because packets are either damaged or lost.
Retransmits	The percentage of TCP packets that were retransmitted during the flow. This usually occurs because packets are either damaged or lost.
RTT Flow Count	The number of flows containing Round Trip Time data.
SRT Flow Count	The number of flows containing Server Response Time data.
Time	The time the data was recorded.
Total Packets	The total number of packets received from the exporter.

Related Tasks

Creating a Report

[Editing a Report](#)

[Investigating the Details](#)

[Saving a Report](#)

[Exporting a Report](#)

[Related Topics](#)

[Report Builder](#)

[Building Queries](#)

[Report Templates](#)

NVM Collection Trend Report

Use the NVM Collection Trend report to view the collection trend for Network Visibility Module (NVM) flows received by your Flow Collectors, which can help you evaluate:

- Are all Flow Collectors receiving flows?
- What is the quantity of flows received?
- Were there any interruptions?

This report is important for troubleshooting and capacity planning as follows:

- **Initial Deployments:** As you deploy new Flow Collectors, run this report to verify they are ingesting the expected amount of telemetry.
- **Troubleshooting:** Review the [report results](#) to confirm your Flow Collectors are receiving NVM flows, review the collection rate, and find anomalies in telemetry ingest for an individual Flow Collector or multiple Flow Collectors. Also, compare the current collection rates with historical rates to identify trends in your data.

To confirm your Data Store is receiving NVM flows, run the [NVM Database Ingest Trend Report](#).

- **Capacity Planning:** Review the [report results](#) to evaluate the telemetry ingest trends over time. Use the growth trends you observe to plan for future expansion and ensure your Flow Collectors are not overloaded based on their specifications.

Requirements

- To receive data on this report, you need Secure Network Analytics with a Data Store deployment. For information and instructions, refer to the [Appliance Installation Guide \(Hardware or Virtual Edition\)](#) and the [System Configuration Guide](#).
- Make sure your Flow Collector is configured to receive data from your Network Visibility Module. For instructions, refer to the [Network Visibility Module \(NVM\) Configuration Guide](#).
- To run this report, select a Data Store domain. To change domains, refer to Select a Domain.

Parameters

Parameter	More Information
Time Range	If you choose Custom , select a short time range for

maximum performance. If you enter a long time range, the report may take a long time to query the data.

Report Results

The NVM Collection Trend report shows your query results by Flow Collector.

- **Filter:** Click individual Flow Collectors in the legend.
- **Zoom In:** To focus on a section of the chart, [zoom in](#) to see the details.

For more information, refer to [Investigating the Details](#).

Related Tasks

[Creating a Report](#)

[Editing a Report](#)

[Investigating the Details](#)

[Saving a Report](#)

[Exporting a Report](#)

Related Topics

[Report Builder](#)

[Report Templates](#)

NVM Database Ingest Trend Report

Use the NVM Database Ingest Trend report to view Network Visibility Module (NVM) records written to your Data Store, which can help you evaluate:

- Is the Data Store receiving flows?
- What is the quantity of flows received?
- Were there any interruptions?

This report is important for troubleshooting and capacity planning as follows:

- **Initial Deployments:** As you deploy a new Data Store, run this report to verify it is ingesting the expected amount of telemetry.
- **Troubleshooting:** Review the [report results](#) to confirm your Data Store is receiving NVM flows and find anomalies in telemetry ingest. Also, compare the current collection rates with historical rates to identify trends in your data.

To confirm your Flow Collectors are receiving NVM flows, run the [NVM Collection Trend Report](#).

- **Capacity Planning:** Review the [report results](#) to evaluate the telemetry ingest trends over time. Use the growth trends you observe to plan for future expansion and ensure your Data Store is not overloaded based on the specifications.

Requirements

- To receive data on this report, you need Secure Network Analytics with a Data Store deployment. For information and instructions, refer to the [Appliance Installation Guide \(Hardware or Virtual Edition\)](#) and the [System Configuration Guide](#).
- Make sure your Flow Collector is configured to receive data from your Network Visibility Module. For instructions, refer to the [Network Visibility Module \(NVM\) Configuration Guide](#).
- To run this report, select a Data Store domain. To change domains, refer to [Select a Domain](#).

Parameters

Parameter	More Information
Time Range	If you choose Custom , select a short time range for maximum performance. If you enter a long time range, the report may take a long time to query the data.

Report Results

The NVM Database Ingest Trend report shows your query results by Flow Collector.

- **Filter:** Click individual Flow Collectors in the legend.
- **Zoom In:** To focus on a section of the chart, [zoom in](#) to see the details.

For more information, refer to [Investigating the Details](#).

Related Tasks

[Creating a Report](#)

[Editing a Report](#)

[Investigating the Details](#)

[Saving a Report](#)

[Exporting a Report](#)

Related Topics

[Report Builder](#)

[Report Templates](#)

Security Events Report

Use this report to review a summary of all security events for the time period you select. For more information about security events, refer to [Investigating Behavior](#).

Parameters

Parameter	Description
Active or First Start Active and/or Last Active Time	Select Active to review security events that were active in the last 10 minutes. To select a time range, select First Start Active and/or Last Active Time , and select preset ranges or enter a custom range.
Flow Collector	Select a Secure Network Analytics Flow Collector in your network.
Host IP Address or Range	You can enter the following: • Multiple IP Addresses: For each address, type the IP address and press Enter. • IP Range: Type <IP address > - <range>. For example, 10.52.1.0-10 includes any IP address from 10.52.1.0 to 10.52.1.10
Host Group	Select 1 host group. Click Select , and check a host group check box.
Security Events	There are two ways to add a security event to your query: • Click the field, and select a security event name from the menu. • Type a security event name in the field, and select it from the menu.

Columns

We've listed all of the columns that are available for this report in alphabetical order here. To hide columns or change the order, refer to [Investigating the Details](#).

Column Name	Description
Appliance	The Flow Collector that received the data or reported the alarm.
Concern Index	The source host's current concern index.
Duration	The amount of time (in minutes) that the security event was active.
Hit Count	The number of times all of the listed security events have occurred. This is the total for all security event types.
ID	A unique identifier for the security event. We use it for debugging and troubleshooting.
Last Active Time	The date and time of the most recent event.
n/a	There is no data available.
Protocol	The TCP or UDP protocol used in the alarm or event.
Security Events	The security event identified in the data. For more information about the security event, go to the Security Insight Dashboard. Refer to Security Event List
Source Country	The country of the source host that caused the alarm or security event.
Source Host	The host IP address and name (if available) that shows the security event. It can be an inside or outside host.
Source Host Group Path (Outside Hosts > Countries)	The hierarchy of all parent host groups that include the source host.
Source Host Groups	The host groups that include the source IP address.
Start Active Time	The date and time of the security event's first occurrence or the alarm condition became active.

Target Country	The country of the target host. For alarms, note the following: • The Manager uses the ISO 3166 standard for country names. • This column also identifies whether the host is local or private (RFC 1918).
Target Host	The Target host IP address and name (if available) that shows the security event. It can be an inside or outside host. • Multiple Hosts: For security events that target the IP address of 0.0.0.0 • Subnet x.x.x.x: For security events that target an IP address with a zero in the last octet (for example: 10.0.1.1.0).
Target Host Name	The name assigned to the target host, if available.
Target Port	The port of the target host you queried.

Related Tasks

[Creating a Report](#)

[Editing a Report](#)

[Investigating the Details](#)

[Saving a Report](#)

[Exporting a Report](#)

Related Topics

[Report Builder](#)

[Building Queries](#)

[Report Templates](#)

System Alarms Report

Use this report to review a summary of the active system alarms. You can investigate all system alarms or choose specific alarms for your query.

For more information about system alarms, refer to [System Notifications: Overview](#).

Parameters

Parameter	Description
Active or First Start Active and/or Last Active Time	Select Active to review system alarms that were active in the last 10 minutes. To select a time range, select First Start Active and/or Last Active Time , and select preset ranges or enter a custom range.
System Alarms	Click the field and select a system alarm. You can add more than one system alarm to your query.

Columns

We've listed all of the columns that are available for this report in alphabetical order here. To hide columns or change the order, refer to [Investigating the Details](#).

Column Name	Description
Active	Indicates whether or not the system alarm is active.
Alarm	Name of the system alarm category or individual system alarm. This column can show multiple system alarms.
Alarm ID	We assign a unique ID number to each alarm.
Appliance	The appliance that received the data or reported the system alarm.
Date/Time	The date and time the data occurred.
Device	The device that is receiving the data.

Device Type	The type of device receiving the data.
Details	Additional information about the alarm.
Duration	The amount of time (in minutes) that the system alarm was active.
Severity	The severity level assigned to the alarm: • Informational • Trivial • Minor • Major • Critical
Source	The host IP address that caused the alarm, system alarm, or security event.

Related Tasks

[Creating a Report](#)

[Editing a Report](#)

[Investigating the Details](#)

[Saving a Report](#)

[Exporting a Report](#)

Related Topics

[Report Builder](#)

[Building Queries](#)

[Report Templates](#)

TrustSec Analytics Report

The TrustSec Analytics report shows the traffic volume between security group tags (SGTs) and details about the application flows between them. Use this report to gain insight into the communication on your network and confirm if your ISE policies are being enforced.

The TrustSec Analytics report shows the top 30 security group tags (SGTs) with the highest amount of traffic. You can review activity details for each pair interaction, including the traffic direction and volume, time range, and policy information. You can also refine your report by choosing which security group tags you want to show. For more information, refer to [Report Results](#).



To review policy compliance and possible policy violations for specific security group tags, refer to [TrustSec Policy Analytics Report](#).

- **Configuration:** To configure ISE with Secure Network Analytics, refer to [Managing Cisco® ISE or ISE-PIC](#). To configure TrustSec policies and security group tags (SGTs), refer to the instructions in your ISE Help.
- **Security Group Tag Maximum:** We can process up to 300 security group tags in traffic. If your network exceeds this amount, the report data may be incomplete.
- **Run Time:** Data collection starts after the reset hour (or archive hour). The first time you run and save the report, it may take up to 24 hours for the report to be available for review. It takes less time to view the first TrustSec Analytics report if you've already saved the [TrustSec Policy Analytics Report](#).
 - To configure the reset hour, go to **Configure > System > Domain Properties**.
 - If you select a preset day range in the parameters, the report automatically starts a new day (24 hours) of data collection after the reset hour.

Parameters

Parameter	Description
Day Range	Select a number of days or choose Custom to define the start date and end date. If you select a preset day range, the report will automatically begin a new day (24 hours) of data collection after the reset hour (or archive hour). To

	configure the reset hour, go to Configure > System > Domain Properties .
ISE Cluster	Select a Cisco Identity Services Engine cluster in your network that you want to use to show egress policy information in the report.
Egress Policy Matrix	Select an egress policy matrix. To show TrustSec policy information for your domain, make sure you select an egress policy matrix.

Report Results

When you run the TrustSec Analytics report, it shows the following information:

- [Time Range and Next Update](#)
- [Security Group Tags \(SGTs\)](#)
- [Colors and Icons](#)
- [Security Group Activity](#)
- [Policy Information](#)

Time Range and Next Update

Heading	Description
Time Range	The time range for the data shown in the report.
Next Update	If you selected a preset day range for your report, the report automatically starts a new day (24 hours) of data collection after the reset hour. Next Update shows the next time the report will be updated after the reset hour. Custom: If you selected a custom day range, the report is not updated automatically. The range remains fixed. Refer to Parameters for details.

Security Group Tags (SGTs)

The first time you open the report, it shows the top 30 security group tags with the highest traffic volume. You can keep this view or choose specific tags to show or hide. Refer to [Hiding/Showing Security Group Tags](#).

Colors and Icons


The report colors and icons indicate the traffic status and policy information. Refer to the report legend for brief information. For details, refer to the following:

Color or Icon	Legend Information	Description
 Blue	Traffic with Custom Policy	There was traffic between security groups, and there is a custom policy enabled for the security group tag pair.
 Green	Traffic	There was traffic in one direction or both directions between the security groups.
 Gray	No Traffic	There was no traffic between the security groups.
 Red	Denied Traffic	There was traffic, but the ISE policy is configured to deny that traffic.
	Policy Disabled	There is an ISE policy defined for the security group tag pair, but it is disabled.
	Policy Enabled	There is an ISE policy defined and enabled for the security group tag pair.
	Policy Monitor Mode	There is an ISE policy defined with Monitor Mode enabled for the security group tag pair. If Monitor Mode is configured for all policies, Monitor Mode is shown in the report heading.

Security Group Activity

If a security group is shown on the **Source** axis (y), traffic was detected **to** the security group shown on the **Destination** axis (x), and the traffic may be in both directions.

- **Activity:** To investigate traffic volume, direction, and policy information, click the cell where the security group tags intersect. For details, refer to [Reviewing Security Group Activity](#).
- **ID Numbers:** If there is an ID number shown instead of a security group name, the security group has not been defined in ISE, but it was detected in traffic. To configure the security group tag, refer to your ISE Help.
- **Unknown:** If a source or destination is Unknown, the traffic was detected but not tagged as a specific security group tag. To configure the security group tag, refer to your ISE Help.



Policy Information


The report [colors and icons](#) indicate the current policy type and status for each security group tag pair.

- **Default Policy:** To review the default policy information for the egress policy matrix, click the **Default Policy** check box. If a security group tag pair does not have a specific policy defined, the default policy is applied.
- **Policy Information:** To review policy details, click the cell where the security group tags intersect. For details, refer to [Reviewing Security Group Activity](#).
- **Policy Compliance:** To review policy compliance and possible policy violations for specific security group tags, run the [TrustSec Policy Analytics Report](#). Policy compliance and violations are not shown in the TrustSec Analytics report.

- **Monitor Mode:** If an ISE egress policy has Monitor Mode enabled, activity is analyzed and logged based on the policy configuration, but the policy is not enforced. In this case, your TrustSec Analytics report results may show more red cells with warning icons because offending traffic is identified and monitored but not denied.

The TrustSec Analytics report shows if Monitor Mode is enabled for a policy or all policies as follows:

- **Monitor Specific SGTs:** If Monitor Mode is applied to an individual policy, the (👁️) Monitor Mode icon is shown in the cell. Refer to [Colors and Icons](#) for details.
- **Monitor All:** If Monitor Mode is configured for all policies, **Monitor Mode** is shown in the report heading (as shown in the following example).



Related Tasks

[Creating a TrustSec Report](#)

[Exporting a Report](#)

[Investigating TrustSec Report Details](#)

[Saving a Report](#)

Related Topics

[Report Builder](#)

Report Templates

TrustSec Policy Analytics Report

Use this report to identify possible policy violations, misconfigurations, or deployment issues for the Cisco Identity Services Engine egress policy matrix you select.

To run this report, you will also select the security group tags (in the report parameters) that you want to investigate. We analyze the flows between security groups to determine if the traffic complies with current policies, which are based on single security group access control lists (SGACL). For details, refer to [Policy Compliance](#).

The colors and icons on the report show if there was a possible policy violation, unknown policy, custom policy, or denied traffic. The Security Group Activity panel shows the details of the security group activity and policy compliance. For more information, refer to [Report Results](#).

- **Configuration:** To configure ISE with Secure Network Analytics, refer to [Managing Cisco® ISE or ISE-PIC](#). To configure TrustSec policies and security group tags (SGTs), refer to the instructions in your ISE Help.
- **Security Group Tag Maximum:** We can process up to 300 security group tags in traffic. If your network exceeds this amount, the report data may be incomplete.
- **Report Maximum:** You can create and save only one TrustSec Policy Analytics report across all domains and users. After the first report is created, go to your saved reports to open it. If the report is not there, it may be saved to another domain.
- **Run Time:** Data collection starts after the reset hour (or archive hour). The first time you run and save the report, it may take up to 24 hours for the report to be available for review. It takes less time to view the first TrustSec Policy Analytics report if you've already saved the [TrustSec Analytics Report](#).
 - To configure the reset hour, go to **Configure > System > Domain Properties**.
 - If you select a preset day range in the parameters, the report automatically starts a new day (24 hours) of data collection after the reset hour.

Parameters

Parameter	Description
Day Range	Select a number of days or choose Custom to define the start date and end date. If you select a preset day range, the report will

	automatically begin a new day (24 hours) of data collection after the reset hour (or archive hour). To configure the reset hour, go to Configure > System > Domain Properties .
ISE Cluster	Select a Cisco Identity Services Engine cluster in your network that you want to use to show egress policy information in the report.
Egress Policy Matrix	Select an egress policy matrix. To show TrustSec policy information for your domain, make sure you select an egress policy matrix.
ISE Security Group Tag	Click Select , and check a security group tag check box.

Report Results

When you run the TrustSec Policy Analytics report, it shows the following information:

- **Time Range and Next Update**
- **Security Group Tags (SGTs)**
- **Colors and Icons**
- **Security Group Activity**
- **Policy Compliance**

Time Range and Next Update

Heading	Description
Time Range	The time range for the data shown in the report.
Next Update	If you selected a preset day range for your report, the report automatically starts a new day (24 hours) of data collection after the reset hour. Next Update shows the next time the report will be updated after the reset hour. Custom: If you selected a custom day range, the report is not updated automatically. The range remains fixed. Refer to Parameters for details.

Security Group Tags (SGTs)

The first time you open the report, it shows up to 30 security group tags in alphabetical order.

Colors and Icons

The report colors and icons indicate the traffic status and policy information. Refer to the report legend for brief information. For details, refer to the following:


Color or Icon	Legend Information	Description
 Blue	Traffic with Custom Policy	There was traffic between security groups, and there is a custom policy enabled for the security group tag pair.
 Green	Traffic	There was traffic in one direction or both directions between the security groups.
 Gray	No Traffic	There was no traffic between the security groups.
 Red	Denied Traffic	There was traffic, but the ISE policy is configured to deny that traffic.
	Policy Analysis Pending	Because the security group pair was just added to the report, it was not analyzed. It will be analyzed in the next update.
	Policy Disabled	There is an ISE policy defined for the security group tag pair, but it is disabled.
	Policy Enabled	There is an ISE policy defined and enabled for the security group tag pair.
	Policy Monitor Mode	There is an ISE policy defined with Monitor Mode enabled for the

		security group tag pair. If Monitor Mode is configured for all policies, Monitor Mode is shown in the report heading.
	Possible Policy Violation	There was traffic in both directions between security groups that does not comply with the defined policy.
	Unknown	It is unknown whether the policy was honored or violated in observed traffic, the ISE policy configuration is not supported on this report, or we encountered a processing error.

Security Group Activity

If a security group is shown on the **Source** axis (y), traffic was detected **to** the security group shown on the **Destination** axis (x), and the traffic may be in both directions.

- **Activity:** To investigate traffic volume, direction, and policy information, click the cell where the security group tags intersect. For details, refer to [Reviewing Security Group Activity](#).
- **ID Numbers:** If there is an ID number shown instead of a security group name, the security group has not been defined in ISE, but it was detected in traffic. To configure the security group tag, refer to your ISE Help.
- **Unknown:** If a source or destination is Unknown, the traffic was detected but not tagged as a specific security group tag. To configure the security group tag, refer to your ISE Help.



Policy Compliance

The report [colors and icons](#) indicate the current policy type and status for each security group tag pair.

- **Policy Analysis:** We currently analyze policies with 1 security group access control list (SGACL) maximum that:

- **permits all traffic, except for specific ports and protocols**

For example:

deny icmp

deny tcp dst eq 3389

deny tcp dst eq 1433

deny tcp dst eq 1521

deny tcp dst eq 445

deny tcp dst eq 137

deny udp dst eq 161

permit ip

OR

- **denies all traffic, except for specific ports and protocols**

For example:


```
permit tcp dst eq 22
permit tcp dst eq 80
permit tcp dst eq 443
permit tcp dst eq 3389
permit icmp
deny ip
```

Not Supported: We do not currently analyze policies that include:

- multiple security group access control lists
 - IPv6 access control lists
 - more than 12 lines in a security group access control list
 - port ranges
 - named aliases for ports (for example: http instead of 80)
 - keywords
- **Default Policy:** To review the default policy information for the egress policy matrix, click the **Default Policy** check box. If a security group tag pair does not have a specific policy defined, the default policy is applied.
 - **Policy Information and Compliance:** To review policy details, click the cell where the security group tags intersect and review the information on the Security Group Activity panel. If there are possible policy violations, the details are also shown on the panel. For more information, refer to [Investigating Possible Policy Violations](#).
 - **Monitor Mode:** If an ISE egress policy has Monitor Mode enabled, activity is analyzed and logged based on the policy configuration, but the policy is not enforced. In this case, your TrustSec Policy Analytics report results may show more red cells with warning icons because offending traffic is identified and monitored but not denied.

The TrustSec Policy Analytics report shows if Monitor Mode is enabled for a policy or all policies as follows:

- **Monitor Specific SGTs:** If Monitor Mode is applied to an individual policy, the (👁) Monitor Mode icon is shown in the cell. Refer to [Colors and Icons](#) for details.
- **Monitor All:** If Monitor Mode is configured for all policies, **Monitor Mode** is shown in the report heading (as shown in the following example).



Related Tasks

- Creating a TrustSec Report
- Exporting a Report
- Investigating TrustSec Report Details
- Saving a Report

Related Topics

- Report Builder
- Report Templates

Zeek Database Ingest Trend Report

 Data Store domain is required for Zeek telemetry.

Use the Zeek Database Ingest Trend report to view the Zeek conn.log events written to your Data Store, which can help you evaluate the following:

- Is the Data Store receiving Zeek conn.log events?
- What is the quantity of the Zeek conn.log events received?
- Were there any interruptions?

This report is important for troubleshooting and capacity planning as follows:

- **Initial Deployments:** As you deploy a new Data Store, run this report to verify it is ingesting the expected amount of Zeek telemetry.
- **Troubleshooting:** Review the [report results](#) to confirm your Data Store is receiving Zeek conn.log events and find anomalies in Zeek telemetry ingest. Also, compare the current collection rates with historical rates to identify trends in your data. To confirm you are receiving Zeek conn.log events, run the **Zeek Log Collection Trend Report**.
- **Capacity Planning:** Review the [report results](#) to evaluate Zeek telemetry ingest trends over time. Use the growth trends you observe to plan for future expansion and ensure your Data Store is not overloaded based on the specifications.

Requirements

- To receive Data Store data on this report, you need Secure Network Analytics with a Data Store deployment. For information and instructions, refer to the [Appliance Installation Guide \(Hardware or Virtual Edition\)](#) and the [System Configuration Guide](#).
- To change domains, refer to Select a Domain.

Parameters

Parameter	More Information
Time Range	If you choose Custom , select a short time range for maximum performance. If you enter a long time range, the report may take a long time to query the data.

Report Results

The Flow Database Ingest Trend report shows your query results by Flow Collector.

- **Filter:** Click individual Flow Collectors in the legend.
- **Zoom In:** To focus on a section of the chart, [zoom in](#) to see the details.

For more information, refer to [Investigating the Details](#). For more information about Zeek telemetry for Secure Network Analytics, refer to the [Zeek Configuration Guide](#).

Related Tasks

[Creating a Report](#)

[Editing a Report](#)

[Investigating the Details](#)

[Saving a Report](#)

[Exporting a Report](#)

Related Topics

[Report Builder](#)

[Report Templates](#)

Zeek Log Collection Trend Report

Use the Zeek Collection Trend report to view the collection trend for the Zeek log syslog events count and volume received by your Flow Collectors, which can help you evaluate:

- Are all Flow Collectors configured for Zeek logs collection receiving syslog events?
- What is the quantity of events received?
- Were there any interruptions?

This report is important for troubleshooting and capacity planning as follows:

- **Initial Deployments:** As you deploy new Flow Collectors, run this report to verify they are ingesting the expected amount of telemetry.
- **Troubleshooting:** Review the [report results](#) to confirm your Flow Collectors are receiving Zeek log syslog events, review the collection rate, and find anomalies in telemetry ingest for an individual Flow Collector or multiple Flow Collectors. Also, compare the current collection rates with historical rates to identify trends in your data.
- **Capacity Planning:** Review the [report results](#) to evaluate the telemetry ingest trends over time. Use the growth trends you observe to plan for future expansion and ensure your Flow Collectors are not overloaded based on their specifications.

Requirements

- To receive data on this report, you need Secure Network Analytics with a Data Store deployment. For information and instructions, refer to the [Appliance Installation Guide \(Hardware or Virtual Edition\)](#) and the [System Configuration Guide](#).
- Make sure your Flow Collector is configured to receive data from Zeek. For instructions, refer to the [System Configuration Guide](#).
- To run this report, select a Data Store domain. To change domains, refer to Select a Domain.

Parameters

Parameter	More Information
Time Range	If you choose Custom , select a short time range for maximum performance. If you enter a long time range, the report may take a long time to query the data.

Report Results

The Zeek Collection Trend report shows your query results by Flow Collector.

- **Filter:** Click individual Flow Collectors in the legend.
- **Zoom In:** To focus on a section of the chart, [zoom in](#) to see the details.



You can switch between Events Count Per Period and Events Bytes Per Period to review volume by count or the amount of data for the Zeek log syslog events received.

For more information, refer to [Investigating the Details](#). For more information about Zeek telemetry for Secure Network Analytics, refer to the [Zeek Configuration Guide](#).

Related Tasks

[Creating a Report](#)

[Editing a Report](#)

[Investigating the Details](#)

[Saving a Report](#)

[Exporting a Report](#)

Related Topics

[Report Builder](#)

[Report Templates](#)

Zeek Logs Report

Use this report to review the Zeek telemetry logging events for a specific Zeek log type for a Flow Collector.

Requirements

- To receive data on this report, you need Secure Network Analytics with a Data Store deployment. For information and instructions, refer to the [Appliance Installation Guide \(Hardware or Virtual Edition\)](#) and the [System Configuration Guide](#).
- Make sure Analytics is enabled. From the main menu, choose **Configure** > **Detection** > **Analytics**, then click *Analytics On* Analytics On .
- Make sure your Flow Collector is configured to receive data from Zeek. For instructions, refer to the [System Configuration Guide](#).
- To run this report, select a Data Store domain. To change domains, refer to Select a Domain.

Parameters

Parameter	More Information
Flow Collector	Select a Secure Network Analytics Flow Collector in your network.
Max Records	Select the maximum number of records. The limit is 10,000 records.
Time Range	If you choose Custom , select a short time range for maximum performance. If you enter a long time range, the report may take a long time to query the data.
Zeek Log Type	Select a Zeek log type.



Selecting a log other than conn.log in the Zeek Log Type field may cause the report to run long, but it must run to completion.

Report Results

The Zeek Logs report shows your query results by Flow Collector. For more information about Zeek telemetry for Secure Network Analytics, refer to the [Zeek Configuration Guide](#).

Related Tasks

[Creating a Report](#)

[Editing a Report](#)

[Investigating the Details](#)

[Saving a Report](#)

[Exporting a Report](#)

Related Topics

[Report Builder](#)

[Report Templates](#)

Creating a Report

Use the following instructions to select a report template and choose parameters to create your report.

Before You Start

Note the following:

- **Report Details:** For information about each report template, parameters, and results, refer to [Report Templates](#).
- **Domains:** Some of the reports are exclusive to a Data Store domain or Non-Data Store domain. To run a report with Data Store data, make sure you have a Data Store domain selected. For more information, refer to [Report Templates](#).
- **Building Queries:** For more information about subject, connection, and peer parameters, refer to [Building Queries](#).
- **TrustSec Reports:** To create a [TrustSec Analytics Report](#) or [TrustSec Policy Analytics Report](#), follow the instructions in [Creating a TrustSec Report](#).
- **Dashboard Reports:** To create Dashboard reports, follow the instructions in [Creating a Report Dashboard](#).

Creating a Report

1. From the main menu, choose **Report > Report Builder > Reports**.
2. Click **+New Report**.
3. Click a report template to select it. To find a report template, you can search or filter the list.
 - **Search:** Type terms in the search bar.
 - **Filter:** Click **Security**, **Network**, or **System** to review report templates grouped by type.
4. **No Parameters:** If the report does not include parameters, click **Run Report**. Go to step 6.

Choose Parameters: Select the date or time range, appliance/device, subject, connection, and peer parameters.

If the following parameters are available in the template you selected, note the following:

- **Time Range:** If you choose **Custom**, select a short time range.
- **Host Group:** Click **Select host groups**, and check a host group check box. Click **Apply**.
- **Host IP Addresses or Range:** You can enter the following (for more information about IPv4, IPv6, and entering ranges, refer to [Building Queries](#)):

Multiple IP Addresses: For each address, type the IP address and press Enter.

IP Range: Type <IP address > - <range limit>. For example, **10.52.1.0-10** includes any IP address from 10.52.1.0 to 10.52.1.10.

- **Include/Exclude:** Select **Include** or **Exclude**. If you select Include, add at least 1 parameter.

 The available parameters are determined by the [report template](#) you select.

5. Click **Run Report**.

6. To refine your search or review details:

- **Edit:** To edit your parameters and refine your search, click the **⋮ Ellipsis** icon. Choose **Edit**.
- **Investigate:** To review details, change the view, and/or filter results, refer to [Investigating the Details](#).
- **Endpoint Traffic (NVM) Report:** The report shows the most recent 300 records based on the end time. The number of records updates as you scroll through the report. Refer to [Endpoint Traffic \(NVM\) Report](#) and [Filtering the Endpoint Traffic \(NVM\) Report](#) for more information.

7. Save or close the report:

- **Save:** To save the report and parameters, click **Save Report**. Follow the on-screen prompts to enter a report title and description. Your changes to the views and filters are not saved.
- **Close:** To close a report without saving it, click the **x** icon or **Discard**.

Related Tasks

[Creating a TrustSec Report](#)

[Editing a Report](#)

[Investigating the Details](#)

[Saving a Report](#)

[Exporting a Report](#)

[Related Topics](#)

[Report Builder](#)

[Report Templates](#)

Building Queries

Each report template has a set of parameters you can use to define the search criteria. We use Boolean logic to build your search criteria, and the available parameters are determined by the template you select (for details, refer to [Report Templates](#)).

You can select some of the parameters directly, such as Date, Flow Collector, and Interface. Use the Subject, Connection, and Peer sections to build your query further. For example, within each section, you can include or exclude an IP address or multiple IP addresses.

- [Subject, Connection, Peer](#)
- [Parameters](#)
- [IPv4 IP Addresses](#)
- [IPv6 IP Addresses](#)
- [Example](#)

Subject, Connection, Peer

If the report you selected includes the Subject, Connection, and/or Peer sections, note the following:

- **Subject, Connection, and Peer** are ANDed together, so the report returns results that match the search criteria in all of them. The Connection section returns results matching the server OR client with the given connection search criteria.
- The **Orientation** field specifies if the Subject sections are for client, server, or either. For example, if you select server, the search returns results that match the Subject section as the server and the Peer section as the client. If you select either, the search returns results matching the Subject section as the client and Peer as the server UNIONed with Subject section as the server and Peer as the client.

Parameters

- **Multiple Sets:** Each section includes parameters, such as Host Group, Applications, and Services. Each of these parameters functions as a set. If you enter search criteria in multiple sets, the sets are ANDed together within the section.
- **Include:** You can enter multiple sets of search criteria in certain parameter fields (such as Applications or Services). By default, these are designated as Include and are ORed together. The search returns results that match one, some, or all of these sets. For example, if you want to search for three IP addresses, add three IP

addresses to the parameter field. The search returns results matching any of the three IP addresses.

- **Exclude:** You can also specify that one or more sets of search criteria be excluded. These individual Excludes are ORed together (just as the individual "includes" are ORed together). However, each of these Exclude sets functions as an AND NOT with the other Include sets within that one attribute. Note that If a user is excluded from the Subject in a query, it is automatically excluded from the Peer, and vice versa.

Example

To run a search on one IP address in the Subject section and two IP addresses in the Peer section:

- **Subject Section:** In the Subject section, in the Host IP Address or Range fields, enter the IP address you want to include.
- **Peer Section:** In the Peer section, in the Host IP Address or Range fields, enter the two IP addresses you want to include.

Orientation

- **Server:** If you select server, the results match the Subject IP address that is sending flows to the first OR the second Peer IP address.
- **Client:** If you select client, the search returns results that match the first Peer IP address OR the second Peer IP address that is sending flows to the Subject IP address.
- **Either:** If the orientation is either, the search returns results that match the following:
 - The Subject IP address is sending flows to the first OR the second Peer IP address.

UNION

- The first Peer IP address OR the second Peer IP address is sending flows to the Subject IP address.

IPv4 IP Addresses

IP Address Type	Formatting Requirements
Single IP address	10.52.1.55 Includes only the host at 10.52.1.55.

IP Address Type	Formatting Requirements
Trailing dot subnet	10.52. Includes any IP address from 10.52.0.0 to 10.52.255.255. Note: Make sure you include the trailing period (.) at the end.
Classless Inter-Domain Routing (CIDR) notation (i.e., using a “/” to indicate the masked bits)	10.52.1.0/24 Includes any IP address from 10.52.1.0 to 10.52.1.255. Note: Use this notation if you prefer to enter the IP address using a mask.
Range of networks	10-11. Includes any IP address from 10.0.0.0 to 11.255.255.255 10.52-53. Includes any IP address from 10.52.0.0 to 10.52.255.255, and from 10.53.0.0 to 10.53.255.255 10.52-55.3. Includes any IP address from 10.52.3.0 to 10.52.3.255; from 10.53.3.0 to 10.53.3.255; from 10.54.3.0 to 10.54.3.255; and from 10.55.3.0 to 10.55.3.255 Note: Make sure you include the trailing period (.) at the end.
Range of hosts	10.52.1.0-10 Includes any IP address from 10.52.1.0 to 10.52.1.10 10.52.0.0-10.52.255.255 Includes any IP address from 10.52.0.0 to 10.52.255.255
Comma-separated list	10.52.1.10,10.52.1.50,10.100.1.20 Includes only these three hosts.

IPv6 IP Addresses

IP Address Type	Formatting Requirements
Single IP address	2001:0DB8:0000:0056:0000:ABCD:EF12:3456
	2001:DB8:0:56:0:ABCD:EF12:3456
	2001:DB8::56:0:ABCD:EF12:3456
	2001:DB80:0:56::ABCD:EF12:3456
	2001:DB80:0:56::ABCD:239.18.52.86

Related Tasks

[Creating a Report](#)

[Editing a Report](#)

Related Topics

[Report Builder](#)

[Report Templates](#)

Editing a Report

After you run a report or save a report, you can edit the parameters and run it again to refine your search.

Before You Start

Note the following:

- **Report Details:** For information about each report template, parameters, and results, refer to [Report Templates](#).
- **Building Queries:** For more information about subject, connection, and peer parameters, refer to [Building Queries](#).
- **TrustSec Reports:** To create a [TrustSec Analytics Report](#) or [TrustSec Policy Analytics Report](#), follow the instructions in [Creating a TrustSec Report](#).



If your report does not include parameters, go to [Investigating the Details](#) for more information.

Editing a Report

Use the following instructions to edit saved report or a report saved to a dashboard.

1. From the main menu, choose **Report > Report Builder > Reports** or **Dashboards**.
2. Open a saved report or dashboard.
3. Click the **⋮ Ellipsis** icon.
4. Choose **Edit** or **Edit Report**.
5. **Choose Parameters:** Choose the parameters to update your search.

If the following parameters are available in the template you selected, note the following:

- **Time Range:** If you choose **Custom**, select a short time range.
- **Egress Policy Matrix:** To include policy analysis in your [TrustSec Analytics Report](#) or [TrustSec Policy Analytics Report](#), select an egress policy matrix.
- **ISE Security Group Tag:** If you are creating a [TrustSec Policy Analytics Report](#), select at least 1 security group tag.

- **Host Group:** Click **Select host groups**, and check a host group check box. Click **Apply**.
- **Host IP Addresses or Range:** You can enter the following (for more information about IPv4, IPv6, and entering ranges, refer to [Building Queries](#)):

Multiple IP Addresses: For each address, type the IP address and press Enter.

IP Range: Type <IP address > - <range>. For example, **10.52.1.0-10** includes any IP address from 10.52.1.0 to 10.52.1.10.

- **Include/Exclude:** Select **Include** or **Exclude**. If you select Include, add at least 1 parameter.

 The available parameters are determined by the [report template](#) you select.

6. Click **Run Report**.

7. To refine your search or review details:

- **Edit:** To edit your parameters and refine your search, click the **⋮ Ellipsis** icon. Choose **Edit**.
- **Investigate:** To review details, change the view, and/or filter results, refer to [Investigating the Details](#) or [Investigating TrustSec Report Details](#).
- **Endpoint Traffic (NVM) Report:** The report shows the most recent 300 records based on the end time. The number of records updates as you scroll through the report. Refer to [Endpoint Traffic \(NVM\) Report](#) and [Filtering the Endpoint Traffic \(NVM\) Report](#) for more information.

8. Save or close the report:

- **Save:** To save the report and parameters, click **Save Report**. Follow the on-screen prompts to enter a report title and description. Your changes to the views and filters are not saved.
- **Close:** To close a report without saving it, click the **x** icon or **Discard**.

Related Tasks

[Investigating the Details](#)

[Investigating TrustSec Report Details](#)

[Saving a Report](#)

[Exporting a Report](#)

[Related Topics](#)

[Report Builder](#)

[Report Templates](#)

Deleting a Report

Use the following instructions to delete a saved report.

1. From the main menu, choose **Report > Report Builder > Reports**.
2. Find your report in the list.
3. Click the **⋮ Ellipsis** icon.
4. Choose **Delete** or **Delete Report**.

The report is deleted immediately.

Related Tasks

[Saving a Report](#)

Related Topics

[Report Builder](#)

[Report Templates](#)

Saving a Report

Use the following instructions to save your report and parameters. Your user name is also saved. Your changes to the views and filters are not saved.

1. [Create a new report](#) or [edit a saved report](#).
2. Click **Save Report**.
3. Enter a title and description.
4. Click **Save Report**.

To open a saved report, refer to [Opening a Report](#).

Related Tasks

[Investigating the Details](#)

[Exporting a Report](#)

Related Topics

[Report Builder](#)

[Report Templates](#)

Opening a Report

Use the following instructions to open a saved report.

1. From the main menu, choose **Report > Report Builder > Reports**.
2. Click the report name.

To edit the parameters and run the report, refer to [Editing a Report](#).

Related Tasks

[Investigating the Details](#)

[Investigating TrustSec Report Details](#)

[Saving a Report](#)

[Exporting a Report](#)

Related Topics

[Report Builder](#)

[Report Templates](#)

Creating a TrustSec Report

Use the following instructions to run and save a **TrustSec Analytics Report** or **TrustSec Policy Analytics Report**. The available parameters and results are determined by the report template you select. Refer to **TrustSec Analytics Report** and **TrustSec Policy Analytics Report** for details.

 To run all other report types, refer to **Creating a Report**.

1. From the main menu, choose **Report > Report Builder > Reports**.
2. Click **+New Report**.
3. Click a TrustSec report template to select it.
4. **Choose Parameters:** Select the day range, ISE cluster, and/or egress policy matrix.

If the following parameters are available in the template you selected, note the following:

- **Egress Policy Matrix:** To include policy analysis in your report, select an egress policy matrix.
- **ISE Security Group Tag:** If you are creating a **TrustSec Policy Analytics Report**, select at least 1 security group tag.

5. Click **Run Report**.

First Run: If this is the first time you are running a TrustSec report, click **Save Report**. Follow the on-screen prompts to enter a report name and description. The report is available within approximately 24 hours.

 If this is the first time you are running this report, make sure you save it to start the data analysis.

6. To refine your search or review details:
 - **Edit:** To edit your parameters and refine your search, click the **⋮ Ellipsis** icon. Select **Edit**.
 - **Investigate:** To review details, change the view, and/or filter results, refer to **Investigating TrustSec Report Details**.

7. Save or close the report:

- **Save:** To save the report and parameters, click **Save Report**. Follow the on-screen prompts to enter a report title and description. Your changes to the views and filters are not saved.
- **Close:** To close a report without saving it, click the **x** icon.

Related Tasks

[Editing a Report](#)

[Investigating TrustSec Report Details](#)

[Saving a Report](#)

[Exporting a Report](#)

Related Topics

[Report Builder](#)

[Report Templates](#)

Investigating TrustSec Report Details

After you run a report, you can [edit the parameters](#) to refine your search. The **TrustSec Analytics Report** and **TrustSec Policy Analytics Report** provide options for reviewing traffic and investigating security group activity.

- **Reviewing Colors and Icons**
- **Hiding/Showing Security Group Tags**
- **Reviewing Security Group Activity**
- **Investigating Possible Policy Violations**
- **Reviewing the Default Policy**
- **Running a Flow Search**
- **Running Top Reports**

Related Tasks

Editing a Report

Exporting a Report

Related Topics

Report Builder

Report Templates

TrustSec Analytics Report

TrustSec Policy Analytics Report

Hiding/Showing Security Group Tags

If you've run the [TrustSec Analytics Report](#), you can hide or show security group tags on your report.



Make sure you select 1 tag minimum to display on your TrustSec reports.

1. Click **Manage Columns**.

This filter is unavailable in the [TrustSec Policy Analytics Report](#). You will select the security group tags in the parameters before you run the report.

2. On the Show/Hide Columns panel, you can choose which tags to hide or show on your report as follows:

- **Hide:** Uncheck the check box.
- **Show:** Check the check box.
- **Auto Select:** Click Auto Select to automatically select the 30 security group tags with the highest volume of observed traffic.
- **Deselect:** Click Deselect to uncheck/clear all check boxes. Make sure you select at least 1 tag to show on your report.
- **Search:** To search for a tag name, click the Search field and enter search text.

Related Tasks

[Exporting a Report](#)

[Reviewing Security Group Activity](#)

Related Topics

[Investigating TrustSec Report Details](#)

[Report Builder](#)

[Report Templates](#)

[TrustSec Analytics Report](#)

[TrustSec Policy Analytics Report](#)

Investigating Possible Policy Violations

If you've run the **TrustSec Policy Analytics Report**, the current policies are reviewed for compliance on the observed traffic.

1. Run the **TrustSec Policy Analytics Report** and select specific security group tags (SGTs) to investigate.
2. Review the report and look for cells with any of the following:
 - **Red Cells:** Denied Traffic.
 - **(⚠) Warning Icon:** Offending Traffic



The instructions in this section show how to identify possible policy violations based on a warning icon and/or denied traffic. You may also want to investigate unknown, custom policies, etc. For more information, refer to **Reviewing Colors and Icons** and **Reviewing Security Group Activity**.

3. Click the cell.

DESTINATION ►	Department_E	Developers	Development_Servers	DomainComputer	Employees	Guests	Network_Services
SOURCE ▼							
Department_E							
Developers			✓ ⚠			✓ ⚠	
Development_Serv...							
DomainComputer							
Employees							
Guests		✓					✓ ⚠
Network_Services				✓ ⚠			

4. Review the security group activity details.

Section	Description
Traffic Information	Direction: The arrows indicate the direction of traffic between the source and destination group. If there is only one arrow, traffic was observed in only one direction.

	Volume: The total amount of traffic (in bytes) is shown for each traffic direction. Flow Search: You can run a Flow Search on the selected traffic and add criteria. Refer to Running a Flow Search for details. Top Reports: You can run a Top Report by selecting from the menu. Refer to Running Top Reports for details.
ISE Policy	The ISE policy name and whether or not it is enabled. To configure the policy, refer to the instructions in your ISE Help.
Security Group ACLS	The policy details for the security group access control lists (ACLs), including the name, IP version, and access control entries (ACEs).

5. Scroll to the **Policy Violations** section and review each entry.

Refer to the table for a description of each policy violation status.

- To show more information, click **View Details**.
- Click **Flow Search** or select a Flow Search with additional search criteria from the drop-down menu. For details, refer to [Running a Flow Search](#).

Policy Violation Status	Description
Analysis Pending	Because the security group pair was just added to the report, it was not analyzed. It will be analyzed in the next update.
No Violations	We did not find any flows that violate the policy.
Policy Not Analyzed	The policy was not analyzed because the security group was not included in the report at the time of analysis.
Suspected Violations	We found flows that are not permitted in the policy. Click View Details for more information. To run a Flow Search, refer to Running a Flow Search .

Unknown	We could not confirm if there were any flows that violate the policy. We may have encountered a processing error, or the policy configuration is not supported on this report. Refer to TrustSec Policy Analytics Report for policy analysis information.

Related Tasks

[Exporting a Report](#)

[Reviewing Security Group Activity](#)

Related Topics

[Investigating TrustSec Report Details](#)

[Report Builder](#)

[Report Templates](#)

[TrustSec Analytics Report](#)

[TrustSec Policy Analytics Report](#)

Reviewing Colors and Icons

If you've run a [TrustSec Analytics Report](#) or [TrustSec Policy Analytics Report](#), the colors and icons indicate the traffic status and policy compliance analytics for each security group pair.

Color or Icon	Legend Information	Description
 Blue	Traffic with Custom Policy	There was traffic between security groups, and there is a custom policy enabled for the security group tag pair.
 Green	Traffic	There was traffic in one direction or both directions between the security groups.
 Gray	No Traffic	There was no traffic between the security groups.
 Red	Denied Traffic	There was traffic, but the ISE policy is configured to deny that traffic.
	Policy Analysis Pending	Because the security group pair was just added to the report, it was not analyzed. It will be analyzed in the next update. This information is only available on the TrustSec Policy Analytics Report .
	Policy Disabled	There is an ISE policy defined for the security group tag pair, but it is disabled.
	Policy Enabled	There is an ISE policy defined and enabled for the security group tag pair.
	Policy Monitor Mode	There is an ISE policy defined with Monitor Mode enabled for the

		security group tag pair. If Monitor Mode is configured for all policies, Monitor Mode is shown in the report heading.
	Possible Policy Violation	There was traffic in both directions between security groups that does not comply with the defined policy. This information is only available on the TrustSec Policy Analytics Report.
	Unknown	It is unknown whether the policy was honored or violated in observed traffic, the ISE policy configuration is not supported on this report, or we encountered a processing error. This information is only available on the TrustSec Policy Analytics Report.

Related Tasks

[Exporting a Report](#)

[Investigating Possible Policy Violations](#)

[Reviewing Security Group Activity](#)

Related Topics

[Investigating TrustSec Report Details](#)

[Report Builder](#)

[Report Templates](#)

[TrustSec Analytics Report](#)

[TrustSec Policy Analytics Report](#)

Reviewing the Default Policy

To review the default policy information for the ISE egress policy matrix, click the **Default Policy** check box. Refer to [TrustSec Analytics Report](#) or [TrustSec Policy Analytics Report](#) for more information.

Related Tasks

[Exporting a Report](#)

[Investigating Possible Policy Violations](#)

[Reviewing Security Group Activity](#)

Related Topics

[Investigating TrustSec Report Details](#)

[Report Builder](#)

[Report Templates](#)

[TrustSec Analytics Report](#)


[TrustSec Policy Analytics Report](#)

Reviewing Security Group Activity

If you've run the [TrustSec Analytics Report](#) or [TrustSec Policy Analytics Report](#), you can review the activity details between two security groups and the related policy information.

 Information about the icons and colors are shown in the legend. For more information, refer to [Reviewing Colors and Icons](#).

1. Click the cell where the two security group tags intersect.



2. Review the security group activity details.



Section	Description
Traffic Information	Direction: The arrows indicate the direction of traffic between the source and destination group. If there is only one arrow, traffic was observed in only one direction. Volume: The total amount of traffic (in bytes) is shown for each traffic direction. Flow Search: You can run a Flow Search on the selected traffic and add criteria. Refer to Running a Flow Search for details. Top Reports: You can run a Top Report by selecting from the menu. Refer to Running Top Reports for details.
ISE Policy	The ISE policy name and whether or not it is enabled.

	To configure the policy, refer to the instructions in your ISE Help.
Security Group ACLS	The policy details for the security group access control lists (ACLs), including the name, IP version, and access control entries (ACEs).

Related Tasks

[Exporting a Report](#)

[Investigating Possible Policy Violations](#)

Related Topics

[Investigating TrustSec Report Details](#)

[Report Builder](#)

[Report Templates](#)


[TrustSec Analytics Report](#)

[TrustSec Policy Analytics Report](#)

Running a Flow Search

Use the following instructions to run a Flow Search from a [TrustSec Analytics Report](#) or [TrustSec Policy Analytics Report](#). If you're investigating a possible policy violation, there may be additional search criteria available for your Flow Search.

1. Click the cell where the two security group tags intersect.
2. On the Security Group Activity panel, click **Flow Search**.



Policy Violations: Scroll to the list of policy violations. Click **Flow Search for Offending Traffic**, and select a specific search from the menu (if available).

- Peer Protocols/Ports
- Connection Protocols



For more information about your Flow Search Results, refer to [Flow Search Results: Overview](#).

Related Tasks

[Exporting a Report](#)

[Investigating Possible Policy Violations](#)

[Reviewing Security Group Activity](#)

Related Topics

[Investigating TrustSec Report Details](#)

[Report Builder](#)

[Report Templates](#)

[TrustSec Analytics Report](#)


[TrustSec Policy Analytics Report](#)

Running Top Reports

Use the following instructions to run a top report from a [TrustSec Analytics Report](#) or [TrustSec Policy Analytics Report](#).

1. Click the cell where the two security group tags (SGTs) intersect.
2. On the Security Group Activity panel, click **Top Reports**.

For more information, refer to [Flow Search Results: View Flow Search Results Overview](#).



Related Tasks

Exporting a Report

Investigating Possible Policy Violations

Reviewing Security Group Activity

Related Topics

Investigating TrustSec Report Details

[Report Builder](#)

[Report Templates](#)

[TrustSec Analytics Report](#)

[TrustSec Policy Analytics Report](#)

Investigating the Details

After you run a report, you can [edit the parameters](#) to refine your search. You can also review the details and investigate issues, as follows:

Charts

- [Reviewing Chart Details](#)
- [Running a Flow Search](#)
- [Running Top Reports](#)
- [Zooming In/Out](#)
- [Filtering the Chart](#)
- [Adding a New Chart Type](#)

Tables

- [Filtering a Table](#)
- [Filtering the Endpoint Traffic \(NVM\) Report](#)
- [Adjusting the Column Width](#)
- [Sorting Table Columns](#)
- [Hiding/Showing Columns](#)
- [Moving a Column](#)



To review TrustSec Report data, refer to [Investigating TrustSec Report Details](#).

Related Tasks

[Editing a Report](#)

[Exporting a Report](#)


Related Topics

[Report Templates](#)

[Report Builder](#)

Adding a New Chart Type

1. Click **+Add Chart**.
2. Click a chart type to select it.



3. Click **Select**.

i Some chart types may not be relevant for the report template you are using.

Related Tasks

[Exporting a Report](#)

[Investigating the Details](#)


Related Topics

[Report Builder](#)


[Report Templates](#)

Filtering the Chart

- **X Axis:** To show data exclusively from one source, click the name in the legend.



- **Y Axis:** To change the view criteria, click the y-axis menu and select a description.



Related Tasks

[Exporting a Report](#)

[Investigating the Details](#)


Related Topics

[Report Builder](#)

[Report Templates](#)

Reviewing Chart Details

- **Show Details:** Move your cursor over the data and pause.



Related Tasks

[Exporting a Report](#)

[Investigating the Details](#)

Related Topics


[Report Builder](#)

[Report Templates](#)

Running a Flow Search

1. Click a data point in the chart.
2. Select **View Flows**.

For more information, refer to [Flow Search Results: View Flow Search Results Overview](#).



Related Tasks

[Exporting a Report](#)

[Investigating the Details](#)

Related Topics

[Report Builder](#)


[Report Templates](#)

Running Top Reports

1. Click a data point in the chart.
2. Select **Top Reports**.

Select a report from the menu.

For more information, refer to [Flow Search Results: View Flow Search Results Overview](#).



Related Tasks

[Exporting a Report](#)

[Investigating the Details](#)

Related Topics

[Report Builder](#)

[Report Templates](#)

Zooming In/Out

Use zoom to focus on specific data in a chart. If a table is included in the report, the table also updates to show your selected data.

- **Zoom In:** Click and drag your cursor over an area on the chart.
- **Show Full View:** Click **Reset Zoom**.

Related Tasks

[Exporting a Report](#)

[Investigating the Details](#)

Related Topics

[Report Builder](#)

[Report Templates](#)

Adjusting the Column Width

1. Move your cursor in between the column until it changes to an **Arrow** icon.
2. Click and drag the column width left or right.

Host Group Application Traffic (0)

Showing: 0 - 0 of 0

Time 1	Application ↑ 2	:	Traffic Within (bps)	↔	Traffic Inbound (bps)
month/day/year					

Related Tasks

[Exporting a Report](#)

[Investigating the Details](#)

Related Topics

[Report Builder](#)

[Report Templates](#)

Hiding/Showing Columns

1. Click **Manage Columns**.
2. To hide a column, uncheck the check box.
To show a column, check the check box.

Related Tasks

[Exporting a Report](#)

[Investigating the Details](#)


Related Topics

[Report Builder](#)

[Report Templates](#)

Moving a Column

Click the column name and drag it to the new insertion point.



Related Tasks

[Exporting a Report](#)

[Investigating the Details](#)

Related Topics

[Report Builder](#)

[Report Templates](#)

Sorting Table Columns

When you first open a report, the gray highlighted column shows the default sort. You can sort additional columns as follows:

- **Sort:** Click the column heading. Each time you click it, the column switches from ascending to descending or descending to ascending. The (↑) **Arrow** icon shows the sort direction.
- **Remove Sort:** Click the column heading until the number in the heading (1, 2, 3) is removed.

Time 1	Application ↑ 2	Traffic Within (bps)
5/14/2021 4:50:00 PM	Undefined	105
5/14/2021 4:45:00 PM	Undefined	106
5/14/2021 4:40:00 PM	Undefined	85

Related Tasks

[Exporting a Report](#)

[Investigating the Details](#)

Related Topics

[Report Builder](#)

[Report Templates](#)

Filtering the Endpoint Traffic (NVM) Report

Any time you run a new query in the **Endpoint Traffic (NVM) Report**, it shows the most recent 300 records based on the end time. You can [scroll to see more records](#), filter the current records, or filter and query all data in the database.

There are two ways to filter your data. The following instructions include how to filter the current records or filter and query all data in the database.


- [Entering a Search Query](#)
- [Filtering by Parameter](#)

i To filter all other reports, refer to [Filtering a Table](#).

Entering a Search Query

1. Click a column field. Enter your search criteria.

For example, to search the column for a word or number, type it into the field.



IP Addresses

- **IP Addresses:** To search for Source IP Addresses or Destination IP Addresses, type the individual IP address or multiple IP addresses into the fields.
- **IP Address Ranges:** Filtering by IP address range is not available here. To search an IP address range, [edit the report](#), and enter the range in the Source IP Address or Destination IP Address fields. Click **Run**.

2. To add a query, click the  (**Filter**) icon. Select a query from the menu.

3. The table updates to show your search results based on the most recent 300 records (based on the end time).
 - **Search All Data:** To search all records in the database, click **Apply Filters to Total**.
 - **See More Data:** To show more results, scroll down. The number of records updates as you scroll through the report.
 - **Clear:** To clear the filter, click the  (**Clear Filter**) icon.

Filtering by Parameter

1. Click the column field.
2. Select a parameter from the drop-down list.
3. Repeat steps 1 and 2 to add another parameter.
4. The table updates to show your search results based on the most recent 300 records (based on the end time).
 - **Search All Data:** To search all records in the database, click **Apply Filters to Total**.
 - **See More Data:** To show more results, scroll down. The number of records updates as you scroll through the report.
 - **Clear:** To clear the filter, click the **x** icon.

Related Tasks

[Exporting a Report](#)

[Exporting an Endpoint Traffic \(NVM\) Report](#)

[Investigating the Details](#)

Related Topics

[Report Builder](#)

[Report Templates](#)

Filtering a Table

When you apply filters to the table, your changes are also shown in your charts.

- [Entering a Search Query](#)
- [Filtering by Parameter](#)




To filter the Endpoint Traffic (NVM) report, refer to [Filtering the Endpoint Traffic \(NVM\) Report](#).

Entering a Search Query

1. Click a column field. Enter your search criteria.

For example, to search the column for a word or number, type it into the field. The table updates to show your search results.

2. To add a query, click the **Filter** icon.



3. Select a query from the menu.

To clear the filter, click the (**Clear Filter**) icon.

Filtering by Parameter

1. Click the column field.

Host Group Applications (576)

Showing: 1 - 13 of 576

Time	Application
month/day/year	
9/14/2020, 3:30:00 PM	HTTP
9/14/2020, 3:30:00 PM	HTTP (unclassified)
9/14/2020, 3:30:00 PM	HTTPS
9/14/2020, 3:30:00 PM	HTTPS (unclassified)
9/14/2020, 3:30:00 PM	HTTPS (unclassified)

2. Select a parameter from the drop-down list.
3. Repeat steps 1 and 2 to add another parameter.

To clear the filter, click the x icon.

Host Group Applications (576)

Showing: 1 - 13 of 144

Time	Application
month/day/year	HTTPS
9/14/2020, 3:30:00 PM	HTTPS
9/14/2020, 3:25:00 PM	HTTPS
9/14/2020, 3:20:00 PM	HTTPS
9/14/2020, 3:15:00 PM	HTTPS

Related Tasks

Exporting a Report

Exporting an Endpoint Traffic (NVM) Report

Investigating the Details

Related Topics

Report Builder

Report Templates

Sorting Table Columns

When you first open a report, the gray highlighted column shows the default sort. You can sort additional columns as follows:

- **Sort:** Click the column heading. Each time you click it, the column switches from ascending to descending or descending to ascending. The (↑) **Arrow** icon shows the sort direction.
- **Remove Sort:** Click the column heading until the number in the heading (1, 2, 3) is removed.

Time 1	Application ↑ 2	Traffic Within (bps)
5/14/2021 4:50:00 PM	Undefined	105
5/14/2021 4:45:00 PM	Undefined	106
5/14/2021 4:40:00 PM	Undefined	85

Related Tasks

[Exporting a Report](#)

[Investigating the Details](#)

[Investigating TrustSec Report Details](#)

Related Topics

[Report Builder](#)

[Report Templates](#)

Dashboards

All dashboards (saved by you or other users) are listed on the **Report > Report Builder > Dashboards** page. For details, refer to [Report Builder](#).

Use dashboards to group your critical reports or create your own Network Insights Dashboard. You can also quickly refresh the report data on your dashboards.

Click the dashboard name to open a saved dashboard, or add a new dashboard as follows:

- [Creating a Report Dashboard](#)
- [Creating a Network Insights Dashboard](#)

Related Tasks

[Creating a Report](#)

[Creating a TrustSec Report](#)

[Editing a Report](#)

[Investigating the Details](#)

[Investigating TrustSec Report Details](#)

Related Topics

[Report Builder](#)

[Report Templates](#)

Creating a Report Dashboard

Use the following instructions to create a dashboard with one or more customized reports. After you have saved a dashboard, you can add, edit, rearrange, and re-size the reports.



To create a dashboard using network reports, refer to [Creating a Network Insights Dashboard](#).

Before You Start

Note the following:

- **Report Details:** For information about each report template, parameters, and results, refer to [Report Templates](#).
- **Domains:** Some of the reports are exclusive to a Data Store domain or Non-Data Store domain. To run a report with Data Store data, make sure you have a Data Store domain selected. For more information, refer to [Report Templates](#).
- **Building Queries:** For more information about subject, connection, and peer parameters, refer to [Building Queries](#).
- **TrustSec Reports:** To create a [TrustSec Analytics Report](#) or [TrustSec Policy Analytics Report](#), follow the instructions in [Creating a TrustSec Report](#).

Creating a Report Dashboard

1. From the main menu, choose **Report > Report Builder > Dashboards**.
2. Click **+New Dashboard**.
3. Click the **Dashboard** template.



To create a dashboard using network reports, refer to [Creating a Network Insights Dashboard](#).

4. Click **+Add Report**.
5. Scroll through the **Select Report Type** menu. Click a report to add it to your dashboard. To find a report, you can search or filter the list.

- **Search:** Type terms in the search bar.
 - **Filter:** Click **Security**, **Network**, or **System** to review report templates grouped by type.
6. **No Parameters:** If the report does not include parameters, click **Run Report**. Go to [step 7](#).

Choose Parameters: Select the date or time range, appliance/device, subject, connection, and peer parameters.

If the following parameters are available in the template you selected, note the following:

- **Time Range:** If you choose **Custom**, select a short time range.
- **Host Group:** Click **Select host groups**, and check a host group check box. Click **Apply**.
- **Host IP Addresses or Range:** You can enter the following (for more information about IPv4, IPv6, and entering ranges, refer to [Building Queries](#)):

Multiple IP Addresses: For each address, type the IP address and press Enter.

IP Range: Type <IP address > - <range limit>. For example, **10.52.1.0-10** includes any IP address from 10.52.1.0 to 10.52.1.10.

- **Include/Exclude:** Select **Include** or **Exclude**. If you select Include, add at least 1 parameter.

 The available parameters are determined by the [report template](#) you select.

7. Click **Save Dashboard**.
8. To refine your search or review details:
- **Edit:** To edit your parameters and refine your search, choose **Actions > Edit**.
 - **Investigate:** To review details, change the view, and/or filter results, refer to [Investigating the Details](#).
 - **Endpoint Traffic (NVM) Report:** The report shows the most recent 300 records based on the end time. The number of records updates as you scroll through the report. Refer to [Endpoint Traffic \(NVM\) Report](#) and [Filtering the Endpoint Traffic \(NVM\) Report](#) for more information.

9. You can refresh all reports on your dashboard or individual reports.

- **Refresh All Reports:** To refresh all reports on the dashboard, choose **Actions > Refresh**.
- **Refresh All Reports Automatically:** To refresh all reports on the dashboard at a specific interval (for example, every 4 hours), choose **Actions > Refresh Automatically**. Then set the auto refresh interval in the timer box and choose **Save**.
- **Refresh an Individual Report:** To refresh an individual report, choose **Actions > Refresh** on any individual report.

10. Save or close the report:

- **Save:** To save the dashboard and parameters, click **Save Dashboard**. Follow the on-screen prompts to enter a dashboard title and description.
- **Close:** To close a dashboard without saving it, click **Discard**.

11. Schedule or delete a dashboard:

- **Schedule:** To schedule a dashboard, click the **Actions** icon, and choose **Schedule** on the dashboard and complete the information. For more details, refer to [Scheduling a Report](#).
- **Delete a schedule:** To delete a schedule, click the **Actions** icon, and choose **Delete Schedule**.
- **Delete a dashboard:** To delete a dashboard, click the **Actions** icon, and choose **Delete Dashboard**.

Related Tasks

[Creating a TrustSec Report](#)

[Creating a Network Insights Dashboard](#)

[Editing a Dashboard Layout](#)

[Investigating the Details](#)

[Saving a Report](#)

[Scheduling a Report](#)

[Exporting a Report](#)

Related Topics

[Report Builder](#)

Report Templates

Creating a Network Insights Dashboard

The Network Insights dashboard is a customizable dashboard template that contains the following reports by default:

- [Host Group Application Traffic Report](#)
- [Host Group Flow Traffic Report](#)
- [Network and Server Performance Report](#)
- [Flow Collection Trend by Flow Collector Report](#)
- [Flow Collection Trend by Exporter Report](#)
- [Firewall Log Collection Trend Report](#)
- [NVM Collection Trend Report](#)

After you have saved a dashboard, you can add, edit, rearrange, and re-size reports.

 To add reports to a blank dashboard, refer to [Creating a Report Dashboard](#).

Before You Start

Note the following:

- **Report Details:** For information about each report template, parameters, and results, refer to [Report Templates](#).
- **Domains:** Some of the reports are exclusive to a Data Store domain or Non-Data Store domain. To run a report with Data Store data, make sure you have a Data Store domain selected. For more information, refer to [Report Templates](#).
- **Building Queries:** For more information about subject, connection, and peer parameters, refer to [Building Queries](#).
- **TrustSec Reports:** To create a [TrustSec Analytics Report](#) or [TrustSec Policy Analytics Report](#), follow the instructions in [Creating a TrustSec Report](#).

Creating a Network Insights Dashboard

1. From the main menu, choose **Report > Report Builder > Dashboards**.
2. Click **+New Dashboard**.
3. Click the **Network Insights** template.
4. Each report has parameters that are pre-selected by default. You can change your dashboard or individual reports by clicking the **Actions** icon.

- **Change Parameters:** To change parameters on a report, click the **Actions** icon, and choose **Edit**. Change the parameters, and click **Run Report**.
 - **Remove Reports:** To remove a report from your dashboard, click the **Actions** icon, and choose **Remove**.
 - **Unavailable Data:** Some reports may not apply to your environment. For example, if you don't have firewall or NVM telemetry, you won't need the trend reports.
5. You can refresh all reports on your dashboard or individual reports.
- **Refresh All Reports:** To refresh all reports on the dashboard, choose **Actions > Refresh**.
 - **Refresh All Reports Automatically:** To refresh all reports on the dashboard at a specific interval (for example, every 4 hours), choose **Actions > Refresh Automatically**. Then set the auto refresh interval in the timer box and choose **Save**.
 - **Refresh an Individual Report:** To refresh an individual report, choose **Actions > Refresh** on any individual report.
6. Save or close the dashboard:
- **Save:** To save the dashboard and parameters, click **Save Dashboard**. Follow the on-screen prompts to enter a dashboard title and description.
 - **Close:** To close a dashboard without saving it, click **Discard**.
7. Schedule or delete a dashboard:
- **Schedule:** To schedule a dashboard, click the **Actions** icon, and choose **Schedule** on the dashboard and complete the information. For more details, refer to [Scheduling a Report](#).
 - **Delete a schedule:** To delete a schedule, click the **Actions** icon, and choose **Delete Schedule**.
 - **Delete a dashboard:** To delete a dashboard, click the **Actions** icon, and choose **Delete Dashboard**.

Related Tasks

[Creating a TrustSec Report](#)

[Creating a Report Dashboard](#)

[Editing a Dashboard Layout](#)

[Investigating the Details](#)

[Saving a Report](#)

[Scheduling a Report](#)

[Exporting a Report](#)

[Related Topics](#)

[Report Builder](#)

[Report Templates](#)

Editing a Dashboard Layout

Use the following instructions to edit the report layout in a saved dashboard.

 To edit a report in a dashboard, refer to [Editing a Report](#).

1. From the main menu, choose **Report > Report Builder > Dashboards**.
2. Click the saved dashboard name to open it.
3. **Move:** To move a report tile to a different section of the dashboard, click the report header and drag it to a new location.
4. **Resize:** To re-size a report tile, click the bottom corner of the tile and drag it.
5. Click **Save Dashboard**.

Related Tasks

[Creating a Report Dashboard](#)

[Creating a Network Insights Dashboard](#)

[Saving a Report](#)

[Exporting a Report](#)

Related Topics

[Report Builder](#)

[Report Templates](#)

Schedule History

All scheduled reports (saved by you or other users) are listed on the **Report > Report Builder > Schedule History** page. For details, refer to [Report Builder](#).

If your report supports scheduling, you can create a custom schedule and email delivery list, as follows:

- [Scheduling a Report](#)
- [Deleting a Scheduled Report](#)

Related Tasks

[Editing a Report](#)

[Exporting a Report](#)

[Investigating the Details](#)

[Investigating TrustSec Report Details](#)

Related Topics

[Report Builder](#)

[Reports](#)

Scheduling a Report

After you have saved a report, you can create a schedule and email delivery list so that it is sent to recipients at a specific frequency (for example, on a daily basis) and time. You can also choose the report format (PDF, CSV).

Requirements

Before you schedule a report, review the following:

- Confirm that SMTP is configured. Refer to [SMTP Configuration](#) for details.
- Confirm the report is saved in your Reports list. Refer to [Creating a Report](#) for details.

Supported Reports

Scheduling is supported in the following reports:

- [Alarms Report](#)
- [DSCP Status Report](#)
- [Firewall Log Database Ingest Trend Report](#)
- [Flow Collection Trend by Exporter Report](#)
- [Flow Database Ingest Trend Report](#)
- [Host Group Application Traffic Report](#)
- [Interface Application Traffic Report](#)
- [Interface Service Traffic Report](#)
- [NVM Collection Trend Report](#)
- [Security Events Report](#)
- [Zeek Database Ingest Trend Report](#)
- [Alerts Report](#)
- [Firewall Log Collection Trend Report](#)
- [Flow Collection Status Report](#)
- [Flow Collection Trend by Flow Collector Report](#)
- [Flow Rate Usage Report](#)
- [Host Group Flow Traffic Report](#)
- [Interface Network Traffic Report](#)
- [Network and Server Performance Report](#)
- [NVM Database Ingest Trend Report](#)
- [System Alarms Report](#)
- [Zeek Log Collection Trend Report](#)

Schedule a Report

1. From the main menu, choose **Report > Report Builder**.
2. On the Reports list, find the saved report and click the **⋮ (Ellipsis)** icon.
3. Choose **Schedule**.

4. In the **Recurrence** drop-down menu, choose a frequency.
5. In the **Time of Day** field, click the timer box to set the time.
6. In the **Email Recipients** field, type the email address of the first recipient and press **Enter**. Repeat this step to enter additional email addresses.
7. In the **Format** section, check the checkboxes (PDF, CSV).
8. Click **Schedule** to save your changes.

Related Tasks

[Creating a Report](#)

[Deleting a Scheduled Report](#)

[Deleting a Scheduled Report](#)

Related Topics

[Report Builder](#)

[Report Templates](#)

Deleting a Scheduled Report

Follow the steps below to delete a scheduled report.

1. From the main menu, select **Report > Report Builder**.
2. On the Reports list, find the saved report and click the **⋮ (Ellipsis)** icon.
3. Select **Delete Schedule**.
4. Follow the on-screen prompts to delete the schedule.

Related Tasks

[Scheduling a Report](#)

[Creating a Report](#)

Related Topics

[Report Builder](#)

[Report Templates](#)

Exporting a Report

Use the following instructions to export a chart (PDF) or table (CSV).



To export an Endpoint Traffic (NVM) report, refer to [Exporting an Endpoint Traffic \(NVM\) Report](#).

1. [Create a new report](#) or [open a saved report](#).
2. On a chart or table, click **Export**.
3. Follow the on-screen prompts to download and save the file.

Related Tasks

[Exporting an Endpoint Traffic \(NVM\) Report](#)

[Saving a Report](#)

Related Topics

[Report Builder](#)

[Report Templates](#)

[Cisco Support](#)


Exporting an Endpoint Traffic (NVM) Report

Use the following instructions to export your **Endpoint Traffic (NVM) Report** to a CSV file. You can choose to export the current view or all data in the database.

i To export all other reports, refer to [Exporting a Report](#).

1. [Create a new report](#) or [open a saved report](#).
2. Click **Export**. Choose **Current View** or **All Data**.
 - **Current View:** Exports the current 300 records shown from your query and filters.
 - **All Data:** Exports all records in the database (from your query and filters). Check the **Export** menu periodically. When **Ready for Download** is shown, click the file name to download it.

The report may take a long time to prepare depending on the Time Range parameter you've selected and the number of records in the database. If the data set is too large, you will receive an error message within 10 minutes. If this happens, filter your report or edit your report to choose a shorter time range.



3. Follow the on-screen prompts to download and save the file.

Related Tasks

[Editing a Report](#)

[Filtering the Endpoint Traffic \(NVM\) Report](#)

[Saving a Report](#)

Related Topics

[Endpoint Traffic \(NVM\) Report](#)

[Report Builder](#)

[Report Templates](#)

[Cisco Support](#)

Visibility Assessment: Overview



You must be a Primary Admin or Power Analyst to access Visibility Assessment.

Modern organizations need internal visibility to understand the state of their network and its traffic. Secure Network Analytics continuously monitors and protects internal assets, observes data transferred between servers and the Internet, processes traffic flows, and then displays this data in Visibility Assessment. These metrics and others help security and network personnel quantify the hosts, systems, and resources on their network, making sure there is nothing present that they don't know about. They also help identify critical assets, validate policies, audit and demonstrate compliance, and make better decisions based on data. Visibility Assessment provides network visibility that requires minimal server classification and no policy tuning.

To view additional information on this page or on any of the tables, hover your cursor over any blue graph bar.

Open Visibility Assessment

From the main menu, choose **Report > Visibility Assessment**.

Visibility Assessment: Table of Contents

Visibility Assessment: Critical Summary Data Overview

To view key information about the hosts in a particular category, choose the desired category at the top of the page. The charts in these categories show trends for the associated category.

- The large number represents the total number of unique incidents for the last 14 days, unless you have [changed the time range](#).
- If you hover your cursor over any graph bar within a category, you can see the date and the number of unique incidents that occurred on that day. The chart displays data for the last 14 days, unless you have changed the time range.

View the summary for the associated category in the lower half of the page.

This summary contains data for the last 14 days that is related to the category you chose, including the number of associated security events and hit counts. The date at the top of the page indicates when the assessment was last run. Visibility Assessment runs hourly.

The security events for the top 10 hosts for the last 14 days are displayed. The top hosts can differ from day to day, so it is most probable that there will be more than 10 hosts that fall into this category.

If you want to view the security events for any other hosts, click an IP address to access the associated [Host Report](#), from which you can access the Security Events Details page. Keep in mind that the Security Events Details page lists only the security events for the current day. If you want to view the Host Report for a country, click the applicable country on the map.

To learn more about why you might want to access the Host Report, see [Visibility Assessment Host Report: View Host Details](#) in this topic.

To return to the Home page after clicking a category, click the  (**Home**) icon.

You can view summaries for the following five categories:

- [Internal Network Scanners](#)
- [Remote Access Breach](#)
- [Possible Malware](#)
- [Vulnerable Protocol Servers](#)
- [DNS Risk](#)

Related Tasks

[Visibility Assessment: Change the Time Range](#)

[Visibility Assessment: Filter Columns](#)

[Visibility Assessment: Sort Columns](#)

Related Topics

[Security Events Versus Hits](#)

[Visibility Assessment: Overview](#)

Security Events Versus Hits

Refer to the following examples to learn the difference between a security event and a hit.

- A security event is an event which contributes to a possible incident, based on many factors. One of these factors is the number of hits.
- A hit is recorded when a host attempts to communicate to another host on a particular port or in a particular subnet. The port or the subnet is used to determine the hit count, based on the situation.

Scenario		Number of security events that are recorded	Number of hits that are recorded
A host has connected to a group (10) of hosts		1	10
• in subnet X	• on port 80		
A host has connected to a group (10) of hosts		3 (one for each port)	10
• in subnet X	• on ports 23, 80, and 445		
A host has connected		2 (one for each subnet)	20
to a group (10) of hosts	and to another group (10) of hosts		
• in subnet X	• in subnet Z		
• on port 80	• on port 80.		
A host has connected		4 (one for each subnet/port pair)	40
to a group (10) of hosts	and to another group (10) of hosts		
• in subnet X	• in subnet Z		
• on port 80	• on ports 23, 80, and 445.		

Relates Topics

Visibility Assessment: Critical Summary Data Overview

Visibility Assessment: Overview

Visibility Assessment: Change the Time Range

Use the Time Range drop-down list to specify that you want the data to display data for either today or the last 7 days instead of for the last 14 days (which is the default setting) for all data displayed on this page.

To change the default data collection period of 14 days, click the  (**Settings**) icon . The Data Collection Period dialog opens. To change the collection period, click **30 days**; to turn off the collection period, click **Off**.

Relates Topics

[Visibility Assessment: Critical Summary Data Overview](#)

[Visibility Assessment: Overview](#)

Visibility Assessment Reset All Data

After you perform configuration changes (for example, when you add DNS servers to the appropriate host group), it's useful to reset all data. To do this, click the  (**Settings**) icon and click **Reset All Data**. This deletes all data (**except** [Internal Monitored Network data](#)) for all registered domains and begins anew with the following rules:

- Recalculates internal network scanners and host counts for the current day
- Recalculates all other data for the last hour



When you reset all data, this applies only to the primary Manager, not the secondary Manager.

If the reset hour settings have been changed but data has not yet been reset, you will receive a message requesting you to reset all data.

Relates Topics

[Visibility Assessment: Critical Summary Data Overview](#)

[Visibility Assessment: Overview](#)

Visibility Assessment: Filter Columns

You can filter some (but not all) of the columns to narrow your search. To do this, complete the following steps.

1. Click the  (**Vertical Ellipsis**) icon to open the Filter menu.

From this menu you can view a list of available filter options (for example, Is equal to; Is not equal to; Contains) from which you can choose to help you filter the list according to a specific rule. The "Contains" rule is the default filter mode.

2. After you have entered your filter criteria, click **Filter**.

Clearing entries

To clear your filter entries, click the **See More** link and then click **Clear** in the Filter menu.

Relates Topics

[Visibility Assessment: Critical Summary Data Overview](#)

[Visibility Assessment: Overview](#)

Visibility Assessment: Sort Columns

To reverse the sort order for a column, click the column name. Use the column name to toggle between ascending and descending order. You can also click the arrow that is displayed to the right of the column name (after you click the column name). An up arrow indicates that the column is sorted in ascending order, and a down arrow indicates descending order.

Relates Topics

[Visibility Assessment: Critical Summary Data Overview](#)

[Visibility Assessment: Overview](#)

Visibility Assessment: View Internal Network Scanners Data

Internal scanning can be the result of malware installed on internal machines, malicious users searching for additional resources inside the network, or advanced attacks looking for additional systems to connect to and steal data from. Secure Network Analytics can uncover any internal systems performing reconnaissance through network scanning to help find misbehaving systems. These systems might include any number of legitimate servers, which can then be placed on the allowed list.

The large number indicates the number of affected hosts.

When you hover over any bar in the 14-day trend, the number of incidents and the associated date are displayed.

What data does the Summary contain?

This summary contains Inside Hosts that are **not** in the Network Scanners host group, have at least one addr_scan/tcp event, and have accumulated over 300,000 concern index (CI) points.

When you hover over the 14-day trend, the total scans and the date for the associated day is displayed.

How the summary is sorted

The summary is sorted by Concern Index.

Additional details you can view

If you click **View** in the Details column for a host, the Internal Network Scanners Details table opens. It includes the following details for that host for the last 14 days:

- Subnets and ports that received a hit
- First time that scanning was observed
- Last time that scanning was observed
- Number of [hit counts](#) related to the scanning that occurred

To return to the Home page after clicking a category, click the  (**Home**) icon.

You can view summaries for the following five categories:

- [Internal Network Scanners](#)
- [Remote Access Breach](#)

- [Possible Malware](#)
- [Vulnerable Protocol Servers](#)
- [DNS Risk](#)

Related Tasks

[Visibility Assessment: Change the Time Range](#)

[Visibility Assessment: Filter Columns](#)

[Visibility Assessment: Sort Columns](#)

Related Topics

[Security Events Versus Hits](#)

[Visibility Assessment: Critical Summary Data Overview](#)

[Visibility Assessment: Overview](#)

Visibility Assessment: View Remote Access Breach Data

Remote network access services such as VNC and RDP are common in today's organizations, but they are also popular among attackers. If compromised, these services give threat actors the same privileges of legitimate users. By classifying authorized remote access servers exposed to the Internet, Secure Network Analytics can quickly identify illegitimate access.

The large number indicates the number of affected servers.

When you hover over any bar in the 14-day trend, the number of incidents and the associated date are displayed.

What data does the Summary contain?

This summary contains bidirectional communication from an Outside Host (client) to an Inside Host (server) for which the packet count is a minimum of 10 (5 transmitted in each direction), contains servers that are associated with a minimum of 10 clients, and that uses any of the following services:

- pcAnywhere (5631/tcp)
- pcAnywhere (5631/udp)
- RDP (3389/tcp)
- Telnet (23/tcp)
- VNC (5800-5999/tcp)
- XWindows (6000-6063/tcp)
- XWindows (6000-6063/udp)

When you hover over the 14-day trend, the total bytes and the date for the associated day is displayed.

How the summary is sorted

The summary is sorted by total bytes.

Additional details you can view

If you click **View** in the Details column for a host, the Remote Access Breach Details table opens. It includes the following details for activity that has occurred within the last 14 days:

- Client IP address (and name) that has communicated with an internal server
- Host groups to which the client belongs
- Total bytes
- Flow count
- First time that the communication was observed
- Last time that communication was observed

To return to the Home page after clicking a category, click the  (**Home**) icon.

You can view summaries for the following five categories:

- [Internal Network Scanners](#)
- [Remote Access Breach](#)
- [Possible Malware](#)
- [Vulnerable Protocol Servers](#)
- [DNS Risk](#)

Related Tasks

[Visibility Assessment: Change the Time Range](#)

[Visibility Assessment: Filter Columns](#)

[Visibility Assessment: Sort Columns](#)

Related Topics

[Security Events Versus Hits](#)

[Visibility Assessment: Critical Summary Data Overview](#)

[Visibility Assessment: Overview](#)

Visibility Assessment: View Possible Malware Data

Threat actors frequently target organizations by exploiting the Server Message Block (SMB) protocol to gain control of hosts. SMB is commonly used in many organizations, and attackers use it to mask their activities on the network. Targeted destructive malware (e.g., Concker) exploit vulnerabilities in SMB to deploy mechanisms such as proxy tools and backdoors. Secure Network Analytics monitors for hosts involved in many SMB sessions with hosts outside the network, which is consistent with worm propagation.

Malicious entities use targeted destructive malware to exploit vulnerabilities in SMB to deploy the following:

- Backdoors
- Proxy tools
- Hard drive destructive tools
- Target cleaning tools

The large number indicates the number of affected hosts.

When you hover over any bar in the 14-day trend, the number of incidents and the associated date are displayed.

What data does the Summary contain?

This summary contains attempts by an Inside Host (client) to communicate with an Outside Host for which the packet count is greater than 0 for the client and in which the client is communicating with a minimum of 10 servers, using 445/tcp (SMB).

When you hover over the 14-day trend, the server count and the date for the associated day is displayed.

How the summary is sorted

The summary is sorted by server count.

Additional details you can view

If you click **View** in the Details column for a host, the SMB Risk Details table opens. It includes the following details for activity that has occurred within the last 14 days:

- Server IP address (and name) that the internal host has communicated with
- Host groups to which the server belongs
- Total bytes

- Flow count
- First time that communication was observed
- Last time that communication was observed

To return to the Home page after clicking a category, click the  (**Home**) icon.

You can view summaries for the following five categories:

- [Internal Network Scanners](#)
- [Remote Access Breach](#)
- [Possible Malware](#)
- [Vulnerable Protocol Servers](#)
- [DNS Risk](#)

Related Tasks

[Visibility Assessment: Change the Time Range](#)

[Visibility Assessment: Filter Columns](#)

[Visibility Assessment: Sort Columns](#)

Related Topics

[Security Events Versus Hits](#)

[Visibility Assessment: Critical Summary Data Overview](#)

[Visibility Assessment: Overview](#)

Visibility Assessment: View Vulnerable Protocol Servers Data

Telnet is a vulnerable protocol that puts you at risk for credential compromise and data loss. Mainframes and other systems that store sensitive data, financial systems, and customer information often run Telnet, leaving them open to attacks. Since 1994, the CERT Division of Software Engineering Institute has recommended using something other than plain text authentication, such as with Telnet, due to potential network monitoring attacks.

The large number indicates the number of affected servers.

When you hover over any bar in the 14-day trend, the number of incidents and the associated date are displayed.

What data does the Summary contain?

This summary contains bidirectional communication from any host (client) to an Inside Host (server) for which the packet count is greater than 10 (5 transmitted in each direction) and in which the server is communicating with a minimum of 10 clients, using Telnet (23/tcp).

When you hover over the 14-day trend, the total servers and the date for the associated day is displayed. The security events for only the top 10 server hosts are displayed. If a security event is not part of the top 10, they won't be displayed.

How the summary is sorted

The summary is sorted by total bytes.

Additional details you can view

If you click **View** in the Details column for a host, the Vulnerable Protocol Servers Details table opens. It includes the following details for activity that has occurred within the last 14 days:

- Client IP address (and name) that has communicated with an internal server
- Host groups to which the client belongs
- Total bytes
- Flow count
- First time that communication was observed
- Last time that communication was observed

To return to the Home page after clicking a category, click the  (**Home**) icon.

You can view summaries for the following five categories:

- [Internal Network Scanners](#)
- [Remote Access Breach](#)
- [Possible Malware](#)
- [Vulnerable Protocol Servers](#)
- [DNS Risk](#)

Related Tasks

[Visibility Assessment: Change the Time Range](#)

[Visibility Assessment: Filter Columns](#)

[Visibility Assessment: Sort Columns](#)

Related Topics

[Security Events Versus Hits](#)

[Visibility Assessment: Critical Summary Data Overview](#)

[Visibility Assessment: Overview](#)

Visibility Assessment: View DNS Risk Data

DNS servers are critical to normal network function because they translate host names to IP addresses. Many organizations use specific DNS servers to safeguard their network and enforce policies. Hosts using an unauthorized DNS server could indicate malicious activity or a policy violation. Malware can make a host's DNS server forward requests to sites used for phishing or exploit delivery. Likewise, network users may use unauthorized DNS servers to access web resources forbidden by internal policies.

Unauthorized DNS servers can do the following:

- Direct hosts to bad websites to download malware or exploitation tools
- Prevent monitoring of DNS traffic for data loss, command and control activity, and exploitation
- Control or block access to software updates from vendors

The large number indicates the number of affected servers.

When you hover over any bar in the 14-day trend, the number of incidents and the associated date are displayed.

What data does the Summary contain?

This summary contains communication from an Inside Host (client) to any host (server) that is not in the internal DNS host group or authorized external DNS host group, and in which the server is communicating with a minimum of 10 clients, using DNS (53/udp).

When you hover over the 14-day trend, the client count and the date for the associated day is displayed. The security events for the top 10 hosts for the last 14 days are displayed. The top hosts can differ from day to day, so it is most probable that there will be more than 10 hosts that fall into this category. If a security event is not part of the top 10, they won't be displayed.

How the summary is sorted

The summary is sorted by client count.

Additional details you can view

If you click **View** in the Details column for a host, the DNS Risk Details table opens. It includes the following details for activity that has occurred within the last 14 days:

- Client IP address (and name) that has communicated with an unauthorized DNS server
- Host groups to which the client belongs

- Total bytes
- Flow count
- First time that communication was observed
- Last time that communication was observed

To return to the Home page after clicking a category, click the  (**Home**) icon.

You can view summaries for the following five categories:

- [Internal Network Scanners](#)
- [Remote Access Breach](#)
- [Possible Malware](#)
- [Vulnerable Protocol Servers](#)
- [DNS Risk](#)

Related Tasks

[Visibility Assessment: Change the Time Range](#)

[Visibility Assessment: Filter Columns](#)

[Visibility Assessment: Sort Columns](#)

Related Topics

[Security Events Versus Hits](#)

[Visibility Assessment: Critical Summary Data Overview](#)

[Visibility Assessment: Overview](#)

Visibility Assessment: Internal Monitored Network Data Overview

General summary data is displayed on the left side of the page in the Internal Monitored Network section. All data represents activity that has occurred within the last 14 days. The different information rotates in groups every three seconds. To freeze a view, hover your mouse pointer over the text. The graph represents the number of hosts communicating within your network today.

What to do Next

[Visibility Assessment: View Internal Monitored Network Data](#)

Related Topics

[Visibility Assessment: Overview](#)

Visibility Assessment: View Internal Monitored Network Data

The following general summary data is displayed on the left side of the page in the Internal Monitored Network section.

Category	Description
Hosts communicating within your network	This number represents all hosts that have been active on your network today. The trend line beside this number represents activity that has occurred within the last 14 days.
Internal traffic occurring on your network	This number represents the cumulative number of bytes observed within your network for the last 14 days. If more than one Flow Collector exists in your network, the sum of all Flow Collectors is included.
Traffic exchanged between your network and the Internet	This number represents the sum of the bytes sent and bytes received for the last 14 days.
Encrypted traffic exchanged between your network and the Internet	This number represents the sum of the encrypted bytes sent and bytes received for the last 14 days.
Max 95th percentile flows per second	This number represents the max 95th percentile of flows per second (FPS) for the last 14 days. This number can give you the average rate of traffic flowing into your network.
Flow records analyzed	This number represents the cumulative number of flows for the last 14 days. If more than one Flow Collector exists in your network, the sum of all Flow Collectors is included.
Flow history for incident response (days per TB storage)	How many days of flows can you store in 1TB.

Related Tasks

[Visibility Assessment: Change the Time Range](#)

[Visibility Assessment: Filter Columns](#)

[Visibility Assessment: Sort Columns](#)

Related Topics

[Security Events Versus Hits](#)

[Visibility Assessment: Critical Summary Data Overview](#)

[Visibility Assessment: Overview](#)

Visibility Assessment Report: Overview

You can save Visibility Assessment Report results to a PDF file.

What to do Next

[View the Visibility Assessment Report](#)

Related Topics

[Visibility Assessment: Overview](#)

View the Visibility Assessment Report

The Visibility Assessment Report contains the following information:

- Key information selected from both the internal network monitored data and the category summary.
- The top 10 of whatever item that category sorts on (e.g., the SMB Report would include the top 10 servers; the DNS Report would include the top 10 clients, etc.).
- The top 10 high-risk countries, sorted by total bytes.

To generate a report, complete the following steps.

1. Click **Generate Report**.

The results are generated and displayed on a new tab.

2. (Optional) You can do one of the following using the associated browser's print feature:

- Print the report.
- Save the report as a PDF file.



This report includes information for the last 14 days, regardless of the option you choose in the **Time Range** drop-down list.

Related Topics

[Visibility Assessment Report: Overview](#)

[Visibility Assessment: Overview](#)

Visibility Assessment: Host Report Overview

You might want to view the Host Report to gather more information about a host. To access the Host Report, click a host in the Subject column.

Access the Host Report to do any of the following:

- [Classify hosts](#)
- [View security events for non-top 10 hosts](#)
- [View host details](#)

Related Tasks

[Visibility Assessment Host Report: Classify a Host](#)

[Visibility Assessment Host Report: View Security Events for non-Top 10 Hosts](#)

[Visibility Assessment Host Report: View Host Details](#)

Related Topics

[Visibility Assessment: Overview](#)

Visibility Assessment Host Report: Classify a Host

You may find that you want to classify a host (e.g., you have determined that a host is a legitimate network scanner, so you want to classify that host as a network scanner).

To classify a host, click **Classify** in the Host Summary section. The host is removed from the Host Report after 14 days.

Related Tasks

[Visibility Assessment Host Report: View Security Events for non-Top 10 Hosts](#)

[Visibility Assessment Host Report: View Host Details](#)

Related Topics

[Visibility Assessment: Host Report Overview](#)

[Visibility Assessment: Overview](#)

Visibility Assessment Host Report: View Security Events for non-Top 10 Hosts

1. Do one of the following:

- Click the number in the applicable alarm category at the top of the page to view the security events that are contributing to that alarm category for the host.

OR

- In the Alarms by Type component, click on a bar for the alarm and date for which you want to investigate.
3. On the Alarm dashboard, for the host you want to examine, click the link in the Details column.

The Security Events Details page opens filtered by the type of alarm you chose.



Keep in mind that the Security Events Details page lists only the security events for the current day.

Related Tasks

[Visibility Assessment Host Report: Classify a Host](#)

[Visibility Assessment Host Report: View Host Details](#)

Related Topics

[Visibility Assessment: Host Report Overview](#)

[Visibility Assessment: Overview](#)

Visibility Assessment Host Report: View Host Details

Use the Host Report to find answers to the following questions:

- Based on the Traffic by Peer Host Group, are there any host groups that raise concerns or questions? Is the volume of traffic reasonable?
- Are any of the alarm categories at the top of the page showing a volume of alarms that causes concern? If so, click the number of active alarms for the category you want to investigate. The Alarm dashboard displays security events for that category filtered by this host.
- Does the Alarms by Type bar chart display any spikes or trends that you need to investigate?
- In the Users & Sessions component, is the user valid for that device? Are the start and end times and session duration reasonable for this host?
- In the Application Traffic component, which applications is this host using the most? Which ones are transferring the most data? How many are being used to contact inside hosts, and how many are being used to contact outside hosts? Do the results represent normal trends for your system and this host?
- Do the 7-day and 24-hour trend lines for application traffic between this host and inside or outside hosts show spikes or dangerous trends that you need to investigate? Do you notice absence of activity that seems unusual for the type of host and certain types of applications?

Related Tasks

[Visibility Assessment Host Report: Classify a Host](#)

[Visibility Assessment Host Report: View Security Events for non-Top 10 Hosts](#)

Related Topics

[Visibility Assessment: Host Report Overview](#)

[Visibility Assessment: Overview](#)

Visibility Assessment: High Risk Countries Overview

Attacks can come from distant regions of the world, and a common way to identify compromises is to detect traffic coming from those hosts located in high-risk countries as well as from internal hosts. When organizations do not conduct business with a certain region, there should be minimal traffic with hosts located there. If you observe a large amount of data being sent from countries you have designated as high risk, this could indicate that data exfiltration is occurring.

Traffic involving high-risk countries could be a sign of any of the following:

- Data exfiltration
- Advanced persistent threats (APTs)
- Command and control (C&C) activity

Related Tasks

[Visibility Assessment: View High Risk Countries Traffic](#)

[Visibility Assessment: Configure the High Risk Countries Map](#)

Related Topics

[Visibility Assessment: Overview](#)

Visibility Assessment: View High Risk Countries Traffic

On the world map you can quickly see the total traffic that has occurred between other countries and your network within the last 14 days, and whether or not these countries are high risk. The map key indicates what each color represents.

 This map is neither drawn to scale nor intended to represent geopolitical borders.

Color	Indication
dark red	high-risk country; has traffic
light red	high-risk country; has <i>no</i> traffic
dark gray	not a high-risk country; has traffic
light gray	not a high-risk country; has <i>no</i> traffic

Related Tasks

[Visibility Assessment: Configure the High Risk Countries Map](#)

Related Topics

[Visibility Assessment: High Risk Countries Overview](#)

[Visibility Assessment: Overview](#)

Visibility Assessment: Configure the High Risk Countries Map

You can specify which countries you want to designate as suspect countries. To do this, click the  (**Settings**) icon in the map component. The Data Collection Period dialog opens. Click **Select High-Risk Countries**. The Select High-Risk Countries dialog opens, from which you can choose the applicable countries or continents. The map incorporates these changes when Visibility Assessment begins its next hourly run. If you select more than 10, the map shows only the top 10 high-risk countries.

You can type the country name in the search field in the list. As you begin to type the country name, the countries that match what you have begun to type remain in the dialog. You can then choose the appropriate entry.

Related Tasks

[Visibility Assessment: View High Risk Countries Traffic](#)

Related Topics

[Visibility Assessment: High Risk Countries Overview](#)

[Visibility Assessment: Overview](#)

Network Diagrams Overview

One of the biggest challenges facing an organization is determining if a performance issue is related to an application, the network, or a security concern. Network diagrams enable you to more easily pinpoint the cause of a network problem, such as slow traffic, by visually showing relationship policies between two host groups.

Using network diagrams, you can graphically monitor the status of your entire system in near real-time and easily view active alarms and your organization's network traffic based on various criteria such as location, function, time, applications, services, or virtual environment.

When creating network diagrams, you can add images, shapes, and text boxes to aid in the visual representation of your environment.

The following alarms are displayed in a network diagram:

- all active alarms associated with the relationship policies that exist in a network diagram
- all active alarms associated with the host groups that exist in a network diagram



The Network Diagrams page refreshes every minute to show the latest traffic data.

Who can access network diagrams?



If you are a Secure Network Analytics user, you can use the Network Diagrams feature. However, you will see network diagrams which contain only the host groups for which you have user permissions.

- To determine if a particular data role has permissions to a specific host group, do the following:
 1. Access the User Management: Data Roles tab.
 2. In the row of the applicable data role, click **Edit** from the menu in the Actions column.
 3. On the User Management: Data Roles page, find the host group in the Data Role configuration tree and confirm if the Read Only check box, Read & Write check box, or both, are selected.

For more information, see the "Configuring Data Roles" topic in the Web App Help.

- To determine if a particular web role can currently view network diagrams, do the following:
 1. Access the User Management: Users tab.
 2. Click the  (**Information**) icon next to the applicable Web Role heading.

For more information, see the "Managing Users" topic in the Web App Help.

Users most likely to use network diagrams

Users most likely to use network diagrams are the following:

- System engineers and architects who configure and deploy security assets, identify current and emerging security threats, and keep the network running by monitoring it and troubleshooting issues.
- Network administrators who are responsible for keeping the network (e.g., servers, key infrastructure) up and running, ensuring that traffic reaches its appropriate destination, and ensuring that the network maintains adherence to compliance rules.

Open Network Diagrams

From the main menu, choose **Report > Network Diagrams**.



If you want to enlarge the canvas so that you can view more of a network diagram, you can collapse the Network Diagrams pane that is located on the left side of the canvas. Simply click the arrow located on the right border of the Network Diagrams pane to close the pane. To reopen the Network Diagrams pane, click the arrow that now rests on the left border of the canvas.

Network Diagrams: Table of Contents

Use Cases

Following are some possible scenarios for which you might want to create network diagrams.

Case 1

Your organization has many branch offices, and you want to quickly determine if any traffic is occurring among these satellite offices.

Case 2

You want to monitor all traffic that is occurring over specific ports and protocols as well as traffic that is using specific applications.

Case 3

You want to monitor traffic between specific segments of your internal network and various locations in the Internet.

Related Tasks

[Viewing Video Tutorials](#)

Related Topics

[Network Diagrams Overview](#)

Viewing Video Tutorials

Use the following video tutorials as needed to help you learn how to use some of the network diagrams functionality. You can also find each of these links in its associated section in this topic.

- [Create a basic network diagram](#)
- [Bend edges](#)
- [Create a relationship policy that points back to itself](#)

Related Topics

Use Cases

Interpreting Network Diagrams Overview

Network Diagrams Overview

Default Network Diagrams

The following network diagrams exist by default in Secure Network Analytics.

This default network diagram...	Shows...
Functional Diagram	traffic occurring between your server host groups and inside and outside hosts.
Internet Usage	Internet traffic occurring on your network. The number in the circle on the right represents your internal network traffic.
Internet World Map	traffic occurring between your network and the rest of the world.
Servers (simple)	traffic occurring between your public web servers and specified locations throughout the world. It also shows internal traffic between the server layers.

Related Topics

[Creating Network Diagrams Overview](#)

[Network Diagrams Overview](#)

Creating Network Diagrams Overview

Creating network diagrams consists of the following steps:

1. [Beginning steps](#)
2. [Adding host groups](#)
3. [Creating relationship policies](#)

Related Topics

[Network Diagrams Overview](#)

1 Beginning Steps

Watch this [video tutorial](#) to see how to create a basic network diagram.

1. At the bottom of the Network Diagrams pane on the left side of the page, click **Create New Diagram**.
2. At the top of the canvas, complete the following two fields. Listed below is the placeholder text displayed on the canvas.
 - Untitled Diagram** You can duplicate network diagram names.
 - Enter diagram description** This is an optional field.
3. (Optional) Name the diagram. You can allow Secure Network Analytics to assign a default name, but in order to more quickly identify a specific diagram, we advise that you name each one.
 - By default, the first diagram you save without naming it is assigned the name of *Untitled Diagram*, and, with the exception of the first diagram saved, a sequential number is appended at the end of the name.
4. Click **Save Diagram**.
5. Go to the next step, [2 Adding Host Groups](#).

What to do Next

[2 Adding Host Groups](#)

Related Topics

[Creating Network Diagrams Overview](#)

[Network Diagrams Overview](#)

2 Adding Host Groups

i You must be in edit mode to add host groups.

1. In the Network Diagrams pane on the left side of the page, if the Host Groups tab is not already selected, click it.
2. Click and drag all desired host groups to the canvas. You can resize and drag host groups to any location.
 - To change the host group shape (which by default is circular) to a rectangle, right-click the host group and choose **Change Shape to Rectangle** from the menu.
 - To make edits to the host group name, double-click the name.
 - To enlarge the host group, click the dotted outline around the host group so that the small navigation squares appear. Click a square and drag.
 - To delete a host group (and all associated edges) from the canvas, right-click the host group and choose **Delete Host Group** from the menu.
3. Go to the next step, [3 Creating Relationship Policies Overview](#). Within this section is described how to create relationship policies between different host groups and how to create a relationship policy that points back to itself.

If you delete a host group in Secure Network Analytics

If you delete a host group using Host Group Management or delete a relationship policy using Policy Management, the following occurs:

- The host group is removed from any network diagram in which it exists.
- All edges that were connected to the host group in a network diagram are deleted.
- All relationship policies that are associated with a map and to which the host group belongs are deleted from Policy Management.

What to do Next

[Creating Relationship Policies Between Different Host Groups](#)

[Creating Relationship Policies Within a Single Host Group](#)

Related Topics

[3 Creating Relationship Policies Overview](#)

[Creating Network Diagrams Overview](#)

[Network Diagrams Overview](#)

3 Creating Relationship Policies Overview

After you have added host groups to your network diagram, you need to create the desired relationship policies. You can create relationship policies in either of the two following ways:

- Between different host groups.
- Within a single host group (you can view activity among all hosts contained in this host group).

- After you create a relationship policy, by default the Relationship High Total Traffic event is automatically configured (but not enabled) for that relationship policy. For information about how to add one or more additional events, see [\(Optional\) Adding One or More Additional Events to a Relationship Policy](#).
- For information about how to enable an event associated with a relationship policy (so that you can see any related active alarms), see [\(Optional\) Enabling Events for a Relationship Policy to See Event Alarms for the Policy](#).



Related Tasks

[Creating Relationship Policies Between Different Host Groups](#)

[Creating Relationship Policies Within a Single Host Group](#)

Related Topics

[Creating Network Diagrams Overview](#)

[Network Diagrams Overview](#)

Creating Relationship Policies Between Different Host Groups

After you have added your host groups to a network diagram, you can establish relationship policies between any host groups of your choosing.

1. Click one of the ports (small blue circles) displayed on the host group that represents the first host group of the relationship and drag an edge (line) to one of the ports on the host group for which you want to establish a relationship policy. It does not matter which host group you start with.

You can create as many edges (relationship policies) as you want between the same two host groups. You would do this if you want to assign a different set of services or application to each policy.

- If you want to create more than one edge between two points (e.g., you want to define a different service or application for the relationship policy), right-click the edge and choose **Add Point** from the menu or double-click the edge. Click and hold the new point with your cursor, using it as "handle," and drag the point in the desired direction. Now you have room to create another edge between the two existing ports.

Watch this [video tutorial](#) to see how to bend edges.

- To add a label, right-click the edge and choose **Add Custom Label** from the menu.

Example: You filter a specific relationship policy to show only Citrix. You would probably want to add a label named "Citrix" so that you know when viewing that edge that it represents traffic associated with the Citrix application.

- To delete an edge, right-click the edge and choose **Delete Edge** from the menu.
- To delete a point, right-click the element and choose **Delete Item** from the menu.
- To edit or delete a custom label, right-click the custom label and choose the appropriate option from the menu.

2. Continue to drag edges between applicable host groups to create all desired relationship policies.
3. To save the network diagram, click **Save Diagram**.

As soon as you save a network diagram, the Manager sends the relationship information to the Flow Collectors, and the network diagram begins to display data in a short while.

For information about how to determine what the reset hour is for your Manager, see the "Domain Properties" topic in the Desktop Client Help.

What to do Next

(Optional) Adding One or More Additional Events to a Relationship Policy

(Optional) Enabling Events for a Relationship Policy to See Event Alarms for the Policy

(Optional) Adding Elements to Your Diagram

(Optional) Adding Existing Icons or Pictures to a Diagram

(Optional) Uploading New Icons or Pictures

Related Tasks

Creating Relationship Policies Within a Single Host Group

Related Topics

(Optional) Adding Images to Your Diagram Overview

Creating Network Diagrams Overview

Network Diagrams Overview

Creating Relationship Policies Within a Single Host Group

If you want to view internal traffic within a host group, you can create a relationship policy that points back to itself (you can then view activity among all hosts contained in this host group).

To create this type of relationship policy, complete the following steps.

1. Double-click one of the ports (small blue circles) on the host group.
A circle (self-pointing edge) is created that serves to represent a relationship to itself.
2. (Optional) If you want to configure another set of services and applications, create another self-pointing edge (which creates another relationship policy).
3. To save the network diagram, click **Save Diagram**.

Watch this [video tutorial](#) to see how to create a relationship policy that points back to itself.

What to do Next

(Optional) Adding One or More Additional Events to a Relationship Policy

(Optional) Enabling Events for a Relationship Policy to See Event Alarms for the Policy

(Optional) Adding Elements to Your Diagram

(Optional) Adding Existing Icons or Pictures to a Diagram

(Optional) Uploading New Icons or Pictures

Related Tasks

Creating Relationship Policies Between Different Host Groups

Related Topics

(Optional) Adding Images to Your Diagram Overview

Creating Network Diagrams Overview

Network Diagrams Overview

(Optional) Adding One or More Additional Events to a Relationship Policy

1. In the network diagram, right-click the applicable edge.
2. From the menu that opens, choose **Policy Management**.

The Relationship Policy page in Policy Management opens with the following fields pre-populated with the values configured for the corresponding network diagram:

- Name Secure Network Analytics creates this policy name based on the two associated host groups. Therefore, there can be more than one policy name with the same name.)
- Host Group - Side 1 (read-only)
- Host Group - Side 2 (read-only)
- Map or Diagram Name (read-only)

3. Click **Select Events** in the Relationship table.

The Events panel opens on the left side of the page from which you can select the desired events.

4. Add the desired event or events.
5. (Optional) Display an active alarm for an event. To do this, you must enable that event. Go to the next step, [\(Optional\) Enabling Events for a Relationship Policy to See Event Alarms for the Policy](#)

For more information about configuring relationship policies, see "Configuring Relationship Policies" topic in the Web App Help.

Related Topics

[Creating Network Diagrams Overview](#)

[Network Diagrams Overview](#)

(Optional) Enabling Events for a Relationship Policy to See Event Alarms for the Policy

If you want Secure Network Analytics to display an active alarm for a particular event that is associated with a specific relationship policy, you must enable that event within Policy Management for that specific policy.

On the Policy Management: Relationship Policy page for the applicable relationship policy, do the following:

If you want to...	Click the toggle switch in the Status column to display...
Enable an event	On
Disable an event	Off

You must click **Save** for your change to take effect.

Related Topics

[Creating Network Diagrams Overview](#)

[Network Diagrams Overview](#)

(Optional) Adding Elements to Your Diagram

 You must be in edit mode to add elements.

You can add graphical elements to your network diagram. These elements consist of rectangles, ellipses, diamonds, triangles, and any text you may want to add.

1. In the Network Diagrams pane on the left side of the page, click the Elements tab.
2. Click the element you want to create and drag it to the desired location on the canvas. You can do the following:
 - Resize elements.
 - Add text to an element.
3. To add text, do one of the following:
 - To add text to the interior of an element, double-click inside the element and begin typing.
 - To add text to an area on the canvas, click **Text** in the Network Diagrams pane on the left side of the page and drag to the desired location. A square with a dotted-line border appears. Double-click inside the square and begin typing.
 - To create another line of text, press **Shift-Enter**.
4. To delete text or an element, right-click it and choose **Delete Item** from the menu.
5. To save the network diagram, click **Save Diagram**.

Related Topics

[Creating Network Diagrams Overview](#)

[Network Diagrams Overview](#)

(Optional) Adding Images to Your Diagram Overview

You can use two different types of images in a network diagram: icons and pictures. The icons are preloaded in the Network Diagrams app and you cannot delete them or add any additional ones. You can add as many pictures as you want and can also delete any of them at any time.

Related Topics

[Creating Network Diagrams Overview](#)

[Network Diagrams Overview](#)

(Optional) Adding Existing Icons or Pictures to a Diagram

 You must be in edit mode to add images to a diagram.

1. In the Network Diagrams pane on the left side of the page, click the Images tab.
2. Click either the Icons tab or the Pictures tab, depending on which image type you want to add.
3. Drag and drop the icon or picture to the desired location on the canvas. You can do the following:
 - Resize icons and pictures.
 - Drag and drop an icon or picture on top of a host group. As you move the icon or picture over the host group, the host group turns blue to indicate you can continue with this action.
 - To delete an icon or picture, right-click it and choose one of the following options from the menu:
 - If the picture or icon is assigned to a host group, choose **Remove Image** or **Remove Icon**.
 - If the icon or picture is free-standing on the canvas, choose **Delete Item**.
4. To save the network diagram, click **Save Diagram**.

Related Topics

[Creating Network Diagrams Overview](#)

[Network Diagrams Overview](#)

(Optional) Uploading New Icons or Pictures

 You must be in edit mode to upload a new icon or picture.

1. In the Network Diagrams pane on the left side of the page, click the Images tab.
2. Click the Upload tab.
3. Do one of the following:

- Drag and drop the desired file onto the Upload tab.
- Click **Select files** (displayed on the Upload tab) and choose the image you want to add.

If the file successfully uploaded, then on the Upload tab this is indicated by a check mark and the word Done. The file that you uploaded is displayed at the bottom of the Upload tab.

4. Click the Pictures tab.

You see the image you just uploaded displayed in this list.

Related Topics

[Creating Network Diagrams Overview](#)

[Network Diagrams Overview](#)

Interpreting Network Diagrams Overview

Refer to the following topics for information about viewing network diagrams:

[Filtering Your View](#)

[Viewing Relationship Policy Traffic and Relationship Event Alarms](#)

[Viewing Core and Custom Security Event Alarms](#)

[Viewing Preferences](#)

[Related Topics](#)

[Network Diagrams Overview](#)

Filtering Your View

From the Displayed Metric drop-down list located next to the Start Editing button at the top of the canvas, choose the data type you want to view in the associated diagram. You can view the following types of data:

- Traffic:** Rate of traffic measured in bytes per second. Secure Network Analytics displays the total rate that has occurred within the last minute.
- Packet Rate:** Rate of traffic measured in packets per second. Secure Network Analytics displays the total rate that has occurred within the last minute.
- Round Trip Time:** The length of time it takes for the TCP handshake to complete. Secure Network Analytics displays the average rate that has occurred within the last minute.
- Server Response Time:** The length of time it takes for the server to respond. Secure displays the average rate that has occurred within the last minute.

Related Tasks

[Viewing Relationship Policy Traffic and Relationship Event Alarms](#)

[Viewing Core and Custom Security Event Alarms](#)

[Viewing Preferences](#)

Related Topics

[Interpreting Network Diagrams Overview](#)

[Network Diagrams Overview](#)

Viewing Relationship Policy Traffic and Relationship Event Alarms

Traffic Volume

You create a relationship policy as soon as you connect two host groups in a network diagram with an edge.

- For every edge for which traffic is occurring, a label shows the value of the currently selected metric type. For more information, see [Filtering Your View](#).
- If no traffic is occurring, then no label is displayed. (If it is a custom label, however, the label text remains.)

Refer to the following table to determine the meaning of the color associated with an edge.

Edge Color	Description
Gray	No active alarms exist for the relationship policy.
Red	One or more alarms exist for the relationship policy.

Alarms

If any relationship event alarms exist for a relationship policy, a number and an arrow are displayed in the label located on the edge. The number displayed on the label indicates the total number of active alarm types (in red) associated with the relationship policy (see [\(Optional\) Adding One or More Additional Events to a Relationship Policy](#)). Click the drop-down arrow next to this number to view a list of the alarming relationship events. The number next to each event indicates the total number of times this event has triggered an alarm.

Related Tasks

[Filtering Your View](#)

[Viewing Core and Custom Security Event Alarms](#)

[Viewing Preferences](#)

Related Topics

[Interpreting Network Diagrams Overview](#)

Network Diagrams Overview

Viewing Core and Custom Security Event Alarms

If any core or custom security event alarms exist for a host group, a number and an arrow are displayed in the label located beside the host group. The number displayed on the label indicates the total number of active alarm types (in red) associated with the host group. Click the drop-down arrow next to this number to view a list of the alarming security events. The number next to each event indicates the total number of times this event has triggered an alarm.

Refer to the following table to determine the meaning of the color associated with a host group.

Host Group Border Color	Description
Green	No active alarms exist for the host group.
Red	One or more alarms exist for the host group.

Related Tasks

Filtering Your View

Viewing Relationship Policy Traffic and Relationship Event Alarms

Viewing Preferences

Related Topics

Interpreting Network Diagrams Overview

Network Diagrams Overview

Viewing Preferences

Located in the bottom right corner of this page are five icons you can click to change how much of the network diagram you can see as well as the size of the text. Refer to the following table.

Icon	Function	Description
	Fullscreen	Enables you to see the entire diagram using your monitor's entire surface area (expands beyond the canvas).
	Minimap	It displays the entire diagram in a scaled view within a small window at the bottom right of the page. Within the small window you can click on a spot in the diagram, and the view on the canvas takes you there. You can also drag the diagram around in the small window to control the current view on the canvas. If you are already zoomed in on a diagram, then this is a helpful view to use since it enables you to see the entire diagram while still zoomed in on a specific area.
	Zoom to Fit (default view)	Enables you to see the entire diagram in its most magnified view (within the canvas) without any part of the network diagram disappearing around the perimeter of the viewing pane.
	Zoom Out	Zoom out to see more of the diagram on the page, which makes the content smaller.
	Zoom In	Zoom in to view the diagram in a magnified view.
	Current Zoom Level	Displays the current zoom level. Note: You cannot click this icon.

Related Tasks

[Filtering Your View](#)

[Pivoting to Relationship Policies Associated With an Edge](#)

[Viewing Core and Custom Security Event Alarms](#)

Related Topics

[Interpreting Network Diagrams Overview](#)

[Network Diagrams Overview](#)

What You Can Pivot to From a Network Diagram Overview

You can pivot to the following from a network diagram:

- [Policies to which a host belongs](#)
- [Relationship policies associated with an edge](#)
- [Relate flow traffic](#)
- [Alarm List](#)
- [Top reports](#)
- [Host Group Report](#)

Related Topics

Network Diagrams Overview

Pivoting to Policies to Which a Host Belongs

1. Right-click the host group.
2. Choose **Policy Management** from the menu that opens.

The Policy Management: Core Events tab opens, which lists all the core policies to which the host group belongs as well as the events related to these policies.

3. To see a list of all the custom policies to which the host group belongs, as well as the events related to these policies, click the Custom Events tab.
4. To see a list of all the relationship policies to which the host group belongs, as well as the events related to these policies, click the Relationship Events tab.

Related Tasks

[Pivoting to Relationship Policies Associated With an Edge](#)

[Pivoting to Related Flow Traffic](#)

[Pivoting to the Alarm List](#)

[Pivoting to Top Reports](#)

[Pivoting to the Host Group Report](#)

Related Topics

[What You Can Pivot to From a Network Diagram Overview](#)

[Network Diagrams Overview](#)

Pivoting to Relationship Policies Associated With an Edge

1. Right-click the edge.
2. From the menu that opens, choose **Policy Management**.

The Relationship Policy page associated with this policy opens in Policy Management, listing all the events that have been configured for this policy.

Related Tasks

[Pivoting to Policies to Which a Host Belongs](#)

[Pivoting to Related Flow Traffic](#)

[Pivoting to the Alarm List](#)

[Pivoting to Top Reports](#)

[Pivoting to the Host Group Report](#)

Related Topics

[What You Can Pivot to From a Network Diagram Overview](#)

[Network Diagrams Overview](#)

Pivoting to Related Flow Traffic

From a network diagram, you can run a flow search for two host groups in the same relationship policy or for a single host group.

1. Right-click the applicable edge or host group, depending on which entity you want to run the flow search for.

2. From the menu that opens, do one of the following:

- To run a flow search with the default criteria, click **View Flows**.

The Flow Search page opens. If you are pivoting from a host group, the Subject Host Groups field is pre-populated with that host group. If you are pivoting from an edge, the Subject Host IP Address and Peer Host IP Address fields are pre-populated with the host groups associated with that edge. The flow search begins to run automatically.

- To run a flow search with edited criteria, click **Edit** (located next to the View Flows option). After the Flow Search page opens, make the necessary edits.

3. To run the flow search, click **Search**.

If desired, you can change the name of the flow search (you can duplicate flow search names).

For more information about running a flow search, see the "Building a Flow Search" topic in the Web App Help.

Related Tasks

[Pivoting to Policies to Which a Host Belongs](#)

[Pivoting to Relationship Policies Associated With an Edge](#)

[Pivoting to the Alarm List](#)

[Pivoting to Top Reports](#)

[Pivoting to the Host Group Report](#)

Related Topics

[What You Can Pivot to From a Network Diagram Overview](#)

[Network Diagrams Overview](#)

Pivoting to the Alarm List

Use the Alarm List to investigate active alarms that have occurred since the last reset hour and are associated with either of the following:

- a specific alarm type associated with a particular host group
- a specific alarm type associated with a particular relationship policy

For more information about the Alarm List, see the "Investigating Today's Active Alarms" topic in the Web App Help.

View alarms associated with a relationship policy

Remember that each edge in a network diagram represents a relationship policy between the two host groups on either end of the edge.

1. From inside the label located on the edge for the associated host group pair, choose the drop-down arrow.
2. From the drop-down list, click the desired alarm name or number.

The Alarm List opens in a new tab and displays all active alarms of the alarm type you chose that are associated with the relationship policy (edge) you clicked.

The Alarm List is sorted by the most recent alarm in descending order.

View alarms associated with a host group

1. From inside the label located beside the applicable host group node, choose the drop-down arrow.
2. From the drop-down list, click the desired alarm name or number.

The Alarm List opens in a new tab and displays all active alarms of the alarm type you chose that are associated with the host group you clicked.

The Alarm List is sorted by the most recent alarm in descending order.

Related Tasks

[Pivoting to Policies to Which a Host Belongs](#)

[Pivoting to Relationship Policies Associated With an Edge](#)

[Pivoting to Related Flow Traffic](#)

[Pivoting to Top Reports](#)

[Pivoting to the Host Group Report](#)

Related Topics

[What You Can Pivot to From a Network Diagram Overview](#)

Network Diagrams Overview

Pivoting to Top Reports

From a network diagram, you can run a top report for two host groups in the same relationship policy or for a single host group.

1. Right-click the applicable edge or host group, depending on which entity you want to run the top report for.
2. From the menu that opens, do one of the following:

- To run a top report search with the default criteria, click **Top Reports > [Report name]**.

The Top Report Search page opens. If you are pivoting from a host group, the Subject Host Groups field is pre-populated with that host group. If you are pivoting from an edge, the Subject Host IP Address and Peer Host IP Address fields are pre-populated with the host groups associated with that edge. The top report search begins to run automatically.

- To run a top report search with edited criteria, click **[Top Reports > [Report name] > Edit]** (located next to the relevant top report name).

*After the Top Report Search page opens, make the necessary edits, and click **Search**.*

If desired, you can change the name of the top report search (you can duplicate top report search names).

For more information about running a top report search, see the "Building a Top Report Search" topic in the Web App Help.

Related Tasks

[Pivoting to Policies to Which a Host Belongs](#)

[Pivoting to Relationship Policies Associated With an Edge](#)

[Pivoting to Related Flow Traffic](#)

[Pivoting to the Alarm List](#)

[Pivoting to the Host Group Report](#)

Related Topics

[What You Can Pivot to From a Network Diagram Overview](#)

[Network Diagrams Overview](#)

Pivoting to the Host Group Report

Use the Host Group Report to view detailed information about a host group.

1. Right-click the applicable host group.
2. Choose **Host Group Report** from the menu that opens.

The Host Group Report opens.

Related Tasks

[Pivoting to Policies to Which a Host Belongs](#)

[Pivoting to Relationship Policies Associated With an Edge](#)

[Pivoting to Related Flow Traffic](#)

[Pivoting to the Alarm List](#)

[Pivoting to Top Reports](#)

Related Topics

[What You Can Pivot to From a Network Diagram Overview](#)

[Network Diagrams Overview](#)

Generating Network Diagrams From Flow Search Results Overview

This feature is available only if you have installed the Network Diagrams app.



You can generate a network diagram based on a flow search and then monitor the network diagram to see current traffic and any alarms that are raised. The network diagram consists of all host groups that appeared in your flow search result and any associated relationship policies.

Related Tasks

[Generating Network Diagrams From Flow Search Results](#)

Related Topics

[Network Diagrams Overview](#)

Generating Network Diagrams From Flow Search Results

1. After you have run a flow search in the Web App, click **More > View Diagram** the table on the Flow Search Results page.
 1. Manager begins to generate the network diagram. A progress indicator is displayed during this process
 2. After the diagram is finished generating, it opens in Edit mode in Network Diagrams so that you can make any further changes if necessary. This diagram contains the deepest child host group of any parent host group that is involved in the flow search.
2. (Optional) If you decide that you want to cancel the generation of the network diagram for any reason (e.g., you want to edit the flow configuration settings), simply click **Cancel** (located inside the progress indicator).
3. In order for Secure Network Analytics to create the associated policies and begin to monitor the traffic, you must click **Save Diagram** after the network diagram is generated.

Related Topics

[Generating Network Diagrams From Flow Search Results Overview](#)

[Network Diagrams Overview](#)

Editing Relationship Policies Linked to a Network Diagram Overview

You must be in **View** mode to access a relationship policy in Policy Management.

When a relationship policy is associated with a network diagram or map, you cannot do the following on the Policy Management: Relationship Policy page:

- change the relationship policy host groups
- delete the relationship policy

Related Tasks

[Accessing Relationship Policies](#)

[Enabling and Disabling Alarms](#)

[Adding and Deleting Events](#)

[Editing Service and Application Filters](#)

[Editing Event Details](#)

Related Topics

[Network Diagrams Overview](#)

Accessing Relationship Policies

1. Access the Relationship Policy page (on which you can make edits) using one of the following methods:

- Open a network diagram.

- a. Right-click the edge.
- b. From the menu that opens, choose **Policy Management**.

The Relationship Policy page associated with this policy opens in Policy Management, listing all the events that have been configured for this policy.

- Open Policy Management from the dashboard menu.

- a. From the **Configure** menu, choose **Policy Management**.
- b. Click the Relationship Events tab.
- c. Perform a search in either the Policy Name field for the applicable policy or in the Map or Diagram Name field for the applicable network diagram.

The table is filtered according to your search criteria.

- d. Click the policy name link in the Policy Name column.

The Relationship Policy page associated with this policy opens in Policy Management, listing all the events that have been configured for this policy.

2. Make your desired edits.

Related Tasks

[Enabling and Disabling Alarms](#)

[Adding and Deleting Events](#)

[Editing Service and Application Filters](#)

[Editing Event Details](#)

Related Topics

[Accessing Relationship Policies](#)

[Network Diagrams Overview](#)

Enabling and Disabling Alarms

- To enable the event, click the switch so that the status displays *On*.
- To disable an event that is enabled, click the toggle switch in the Status column so that the status displays *Off*.

You must click **Save** after enabling or disabling for your change to take effect.

Related Tasks

[Accessing Relationship Policies](#)

[Adding and Deleting Events](#)

[Editing Service and Application Filters](#)

[Editing Event Details](#)

Related Topics

[Enabling and Disabling Alarms](#)

[Network Diagrams Overview](#)

Adding and Deleting Events

Add an event

1. Click **Select Events** in the Relationship table.

The Events panel opens on the left side of the page from which you can select the desired events.

If you type the beginning letters or the complete name of the event name, it moves to the top of the list.

2. After you have added the desired events to the relationship policy, click **Save**.

The Policy Manager: Relationship tab opens again. A row is displayed for every event you added to the new relationship policy you just created. Also, the number (in the parentheses) on the Policy Manager: Relationship tab now reflects the new number of existing events in your system.

Delete an event

Click **Delete** in the Actions column of the row for the applicable event.

*After you have deleted the desired events from the relationship policy and clicked **Save**, the Policy Manager: Relationship page opens again, and the events you deleted are no longer displayed. Also, the number (in the parentheses) on the Policy Manager: Relationship tab now reflects the new number of relationship events in your system.*

Related Tasks

[Accessing Relationship Policies](#)

[Enabling and Disabling Alarms](#)

[Editing Service and Application Filters](#)

[Editing Event Details](#)

Related Topics

[Adding and Deleting Events](#)

[Network Diagrams Overview](#)

Editing Service and Application Filters

1. Click the plus sign (+). From the drop-down list, select one of the following options:

- Include All Except
- Exclude All Except

Whichever option you choose applies to both services and applications. So, for instance, you cannot choose to use Include All Except for services and Exclude All Except for applications. If you choose the Include All Except option, for instance, the selections you make from both the Services tab and the Application tab use that option to determine which selections are used in your configuration.

After you have chosen one of these options, a panel opens for you to make your selections for both services and applications. The Services tab is selected by default. The subtitle at the top of the panel reflects the option you selected from the drop-down list.

2. Make the appropriate selections to ensure the service traffic and the application traffic that you want observed for this relationship policy are returned. If you type the beginning letters or the complete application name or service name, it moves to the top of the list.

Secure Network Analytics displays up to three services or applications defined for the relationship policy in this field. The number of any remaining number of services or applications is displayed within parentheses, and a link entitled "More" is displayed beside this number. If you click this link, you see a list of all the services or applications that comprise the traffic between the specified host groups.

The total number of defined services and applications is displayed in parentheses on the associated tab. If a tab does not contain any defined entries, (0) is displayed on the tab.

Related Tasks

[Accessing Relationship Policies](#)

[Enabling and Disabling Alarms](#)

[Adding and Deleting Events](#)

[Editing Event Details](#)

Related Topics

[Editing Service and Application Filters](#)

[Network Diagrams Overview](#)

Editing Event Details

1. Click the arrow at the beginning of the row for the event that you want to edit.

The row expands to display a description of the relationship event as well as the settings you have configured for that event.

2. Make your edits.

When you edit any settings or add any events to an existing policy, a green dog ear is displayed in the top right corner of the row. A red dog ear indicates an input error.

Related Tasks

[Accessing Relationship Policies](#)

[Enabling and Disabling Alarms](#)

[Adding and Deleting Events](#)

[Editing Service and Application Filters](#)

Related Topics

[Editing Event Details](#)

[Network Diagrams Overview](#)

Working with Existing Network Diagrams Overview

You can perform the following actions with existing network diagrams.

Related Tasks

[Opening a Network Diagram](#)

[Editing Components in a Network Diagram](#)

[Duplicating Network Diagrams](#)

[Deleting Network Diagrams](#)

[Importing Diagrams](#)

[Exporting Network Diagrams](#)

Related Topics

[Network Diagrams Overview](#)

Opening a Network Diagram

i When you open an existing network diagram, by default it opens in View mode.

1. To open a diagram, click the diagram name displayed on the Current tab in the List of Diagrams slide-out pane.
2. If you need to search for a diagram, click the  (**Search**) icon to display the Search field. This icon acts as a toggle and switch. If you click it again, the Search field disappears.
3. Type any portion of the diagram name in the Search field. If you type more letters, it narrows down the number of results.

As you type an entry, any diagram names that contain your entry anywhere in the name are displayed and sorted alphabetically, numbers before letters.

Search criteria persists when you toggle between non-legacy and legacy diagrams.

4. To clear your entry, click **X** located at the end of the Search field.

Related Tasks

[Editing Components in a Network Diagram](#)

[Deleting Network Diagrams](#)

[Deleting Network Diagrams](#)

[Importing Diagrams](#)

[Exporting Network Diagrams](#)

Related Topics

[Opening a Network Diagram](#)

[Network Diagrams Overview](#)

Editing Components in a Network Diagram

To edit a diagram component, such as add or delete host groups, edges, or events, do the following.

1. From the Network Diagrams pane on the left side of the page, click the diagram name.

The network diagram opens on the canvas.

2. At the top of the canvas, click **Start Editing**.

Related Tasks

[Opening a Network Diagram](#)

[Duplicating Network Diagrams](#)

[Deleting Network Diagrams](#)

[Importing Diagrams](#)

[Exporting Network Diagrams](#)

Related Topics

[Editing Components in a Network Diagram](#)

[Network Diagrams Overview](#)

Deleting Network Diagrams

1. From the List of Diagrams slide-out pane on the left side of the page, click the **...** (**Ellipsis**) icon to the right of the applicable network diagram.
2. From the menu that opens, click **Delete**.

When you delete an imported network diagram that exists in the Desktop Client, you can reimport it again.

Related Tasks

[Opening a Network Diagram](#)

[Editing Components in a Network Diagram](#)

[Duplicating Network Diagrams](#)

[Importing Diagrams](#)

[Exporting Network Diagrams](#)

Related Topics

[Deleting Network Diagrams](#)

Network Diagrams Overview

Duplicating Network Diagrams

1. From the List of Diagrams slide-out pane on the left side of the page, click the **...** (**Ellipsis**) icon to the right of the diagram you want to duplicate.
2. From the menu that opens, click **Duplicate**.
A copy of the diagram opens in Edit mode with "Duplicate" appended to the end of the name.
3. Rename the diagram.
4. After making your edits, click **Save Diagram** to save the network diagram.

Related Tasks

[Opening a Network Diagram](#)

[Editing Components in a Network Diagram](#)

[Deleting Network Diagrams](#)

[Importing Diagrams](#)

[Exporting Network Diagrams](#)

Related Topics

[Duplicating Network Diagrams](#)

[Network Diagrams Overview](#)

Importing Diagrams

In the Network Diagrams slide-out pane on the left side of the page, complete the following steps.

1. Click the Import tab.

The list of all the network diagrams (referred to as maps in the Desktop Client (former known as Stealthwatch Desktop Client) that were created in the Desktop Client is displayed.

2. To import a diagram that you have previously exported, click **Select files...** at the top of the Import tab to find and import the desired file.

If you import the diagram at any time in the future, and any host groups included in the version of that diagram that currently exists in the Network Diagrams app contains different names than when the diagram was exported, you receive a warning stating that this may indicate you are importing the diagram into a different Secure Network Analytics cluster than from which the diagram was exported.

3. To import one or more network diagrams from the Desktop Client, click **Import** beside each applicable network diagram name.

The text on the Import button changes to say "Open." Any diagram that you imported are now listed only on the Current tab.

4. To import all diagrams from the Desktop Client, click **Import All**.

When importing a diagram, if a diagram contains any host group that no longer exists. you receive an error.

Refer to the following guidelines regarding the importing of network diagrams:

- When you import a network diagram, it is not deleted from the Desktop Client.
- When you delete an imported network diagram in the Web App, you can reimport it again.
- If you make a change to a diagram in the Web App that you have imported from the Desktop Client, the change is not duplicated in the Desktop Client.
- If you make a change to a diagram in the Desktop Client (where a diagram is referred to as a "map"), that change is not duplicated in the Network Diagrams app. You need to delete the diagram in the Network Diagrams app and reimport it.

Related Tasks

Opening a Network Diagram

Editing Components in a Network Diagram

[Duplicating Network Diagrams](#)

[Deleting Network Diagrams](#)

[Exporting Network Diagrams](#)

[Related Topics](#)

[Importing Diagrams](#)

[Network Diagrams Overview](#)

Exporting Network Diagrams

You can export a diagram and save it to a file. The file contains all images, policies, and alarm settings related to the diagram. You can import the diagram either into the same Secure Network Analytics cluster or to a different cluster. For example, you may have a beta lab environment into which you want to import a particular diagram.

1. From the List of Diagrams slide-out pane on the left side of the page, click the **Ellipsis** icon to the right of the diagram you want to export.
2. From the menu that opens, click **Export**.

A zipped file containing the diagram and related components is created and placed in your Downloads folder on your computer.

If you import the diagram at any time in the future, and any host groups included in the version of that diagram that currently exists in the Network Diagrams app contains different names than when the diagram was exported, you receive a warning stating that this may indicate you are importing the diagram into a different Secure Network Analytics cluster than from which the diagram was exported.

Related Tasks

[Opening a Network Diagram](#)

[Editing Components in a Network Diagram](#)

[Duplicating Network Diagrams](#)

[Deleting Network Diagrams](#)

[Importing Diagrams](#)

Related Topics

[Exporting Network Diagrams](#)

[Network Diagrams Overview](#)

Viewing Release Notes and the Secure Network Analytics Apps Compatibility Matrix

Before being moved to the core Secure Network Analytics in v7.5.0, Network Diagrams was a separate app, and for each of its releases, release notes were published. To access these Network Diagrams Release Notes, click [here](#).

To view the Secure Network Analytics Apps Version Compatibility Matrix that was also published for each released version of Network Diagrams, click [here](#).

Related Topics

[Network Diagrams Overview](#)

Managing and Configuring Host Groups

Use this page to logically organize hosts into host groups so that you can monitor different areas of your network and manage host behavior more efficiently. For example, when you add a host to a host group, it automatically inherits the policy associated with that host group. When you need to modify a host's policy, it is easier to only have to re-associate the host to a more relevant host group rather than have to re-associate the host group and then modify any number of individual policies that may have been configured with the particular host included.

All host groups and related details that have been configured in...	Are reflected in the...
The Desktop Client	Manager.
The Manager	Desktop Client with the exception of the Description field, which appears only in the Manager.

How to find this page

Do one of the following to access this page

Do either of the following:

- From the main menu, choose **Configure > Detection > Host Group Management**.
- On the Host Group Report page, from the Actions drop-down menu, select **Host Group Management**.



This page is not available on the secondary Manager.

Which web roles can access this page, and what permissions do they have

For a complete list of permissions assigned to each web role, on the User Management:

Users tab, click the  (**Information**) icon beside the web role whose assigned

permissions you want to view.

If a user has this web role...	Then the user...
Primary Administrator	Can view and make edits within Host Group Management.
Power Analyst	Can view and make edits within Host Group Management.
Analyst	Cannot access the Host Group Management page.
Configuration Manager	Cannot access the Host Group Management page.
Multiple web roles	Possesses permissions that reflect the union of those web roles (i.e., whichever permissions are common among all the web roles assigned to the user).

Concepts you need to understand to use this page

About host groups

Host groups offer flexibility in the way you can organize hosts. In general, hosts can belong to multiple groups. In addition, you can define policies per host group and/or per host. A host group is essentially a virtual container of multiple host IP addresses or IP address ranges that have similar attributes, such as location, function, or topology. By

grouping hosts into host groups, you can control how the Flow Collectors monitor and respond to the behavior of those hosts as a group, rather than individually.

Users can organize hosts in any way that makes sense for their organization. With this freedom, reporting and traffic management gain unlimited flexibility. In addition, policy management is much easier, allowing administrators to set policies for hosts based on the roles those hosts play in the network.

The Manager displays the host group structure in the Host Group tree. By default, each domain contains the following top-level host groups to which you can add subordinate host groups:

- Inside Hosts – Contains all host groups whose hosts have been specifically defined as being a part of your network
- Outside Hosts – Contains all host groups whose hosts have not been specifically defined as being a part of your network

Your data role determine whether or not you see all of the host groups in the Host Group tree.

You can add as many top-level host groups as desired, each containing as many sub-host groups as desired, which can also contain sub-host groups, etc. You can use duplicate host group names, but not at the same host group level (i.e., below the same parent host group).

By default, Secure Network Analytics does not create policies for hosts that are outside your network. However, if you need to track outside hosts that cause traffic on your network on a regular basis, you can place them in an Outside Host group and establish a policy for the group. Then, you can adjust settings just as you would for inside hosts. Additionally, you can also create a top-level host group (i.e., at the same level as Inside Hosts or Outside Hosts) if you have a need to do so for special reporting purposes.

Following are some situations for which you might want to track the behavior of specific outside hosts:

- When you are using outside DNS servers.
- When you have third-party consultants or vendors who regularly access your network.
- When you have partner firms that regularly access your network.

About the Catch All host group

The Inside Hosts host group contains the default Catch All sub-host group. Users can use the Catch All host group to help refine the host group structure.

We recommend that you initially place all large IP ranges that correspond to your network in the Catch All host group. Then, as you create other host groups with more narrowly defined IP ranges or specific IP addresses, those ranges/addresses move out of the Catch All host group automatically. New Manager installations for Secure Network Analytics v6 place the following IP ranges (RFC 1918 and RFC 4193) into the Catch All host group by default:

- 10.0.0.0/8
- 172.16.0.0/12
- 192.168.0.0/16
- fc00::/7

View the IP ranges and IP addresses in your Catch All host group

If you have registered any public IP addresses, we recommend that you also manually place these ranges in the Catch All host group. Depending on your login privileges, you can see which IP ranges and IP address have been defined in your Catch All host group. To do this, do the following:

1. Log in to your Manager.
2. In the Host Group tree, click the Catch All host group.
3. In the IP Addresses and Ranges field, view the IP ranges and IP addresses that comprise the Catch All host group.

Identify rogue host IP addresses

Ideally, when you finish defining your host group structure, there should be no more active IP addresses in the Catch All host group. To identify any rogue host IP addresses, you can look at the Active Hosts document for the Catch All host group. To do this, do the following:

1. Log in to the Desktop Client.
2. In the Enterprise tree, click the context menu next to the Catch All host group and select **Hosts > Active Hosts**.
 - Every network is different, of course, but here are some things to consider when assigning hosts to the Catch All Host Group:
 - Which areas within the network are more sensitive than others?
 - Which areas of the network undergo change frequently vs. those areas that are fairly stable?
 - Where are your critical assets?
 - Which hosts perform similar functions?

- What are the different functions that your hosts perform?
- Do you have any hosts that are simply quirky and routinely behave “strangely”?

About the host groups involved in the setup of custom security events

Secure Network Analytics new deployments (beginning with v7.1) come with a number of pre-defined custom security events. For more information about these pre-defined custom security events, see the "Default custom security events" section in [Configuring Custom Policies \(Custom Security Events\)](#).

Each of the following host groups is used in conjunction with one of these custom security events. If you are upgrading your Secure Network Analytics system, depending on which custom security events you want to add and populate, you may need to add one or more of the following host groups to your environment.

For information about how to define Internet IP ranges for these host groups, see [View host group details](#)

Host Group Name	Description
Added prior to v7.1	
DHCP Servers	Use this host group to define which DHCP servers your organization should be using. Use this host group in conjunction with the CSE: Unauthorized DHCP Traffic custom security event to generate an alarm when any host within Inside Hosts (except 255.255.255.255) and those within DHCP Servers (through 67/UDP) communicates with any host within Inside Hosts (except 255.255.255.255) and those within DHCP Servers (through 68/UDP).
Added as of v7.1	
Authorized External DNS Servers	Malware can cause a host's DNS server to forward requests to sites used for phishing or exploit delivery. Likewise, network users may use unauthorized DNS servers to access web resources forbidden by internal

Host Group Name	Description
	policies. Use the Authorized External DNS Servers host group to define which public DNS resolvers your organization should be using. Use this host group in conjunction with the CSE: Unauthorized DNS Traffic custom security event to generate an alarm when an internal host is using an unauthorized public DNS server. This event helps detect DNS changer type of malware.
Custom Reputation List	Many organization maintain their own Internet watch list of suspicious Internet hosts. Use the Custom Reputation List host group in conjunction with the “CSE: Custom Reputation List” custom security event to detect whenever an Inside host sends or receives traffic from hosts within the Custom Reputation List host group.
Internet Services	Use this host group and child host groups (DNS Servers, NAT Gateways, Proxies) to define internet services within your organization. Use the Internet Services host group in conjunction with the CSE: Unauthorized DNS Traffic custom security event to generate an alarm when an internal host other than a known internal internet service is using a public DNS server.

About Threat Feed

Threat Feed contains IP addresses, port numbers, protocols, host names, and URLs known to be used for malicious activity. The following host groups are included in Threat Feed:

- Bogon – A bogon is an IP address that has not been officially assigned on the public Internet.
- Command & Control Servers – A C&C server is the centralized computer that issues commands to a botnet and receives reports back from the hijacked computers.
- Tor – Tor is an Internet anonymization service.

If you want to investigate a host that has communicated with a malicious host in one of the previously mentioned host groups, but the malicious host no longer appears in the

relevant host group, go to the Alarm Table in the Desktop Client and filter on the following components:

- Types – Select the applicable bogon, Command & Control, or Tor alarm(s), depending on which type of malicious host(s) you are wanting to filter on.
- Date/Time – Filter according to the time period for which you want to investigate.



- To detect URLs in Threat Feed that may be contacting your hosts, you must have a Flow Sensor or router installed that is configured to export IPFIX (vs. NetFlow). (By default, the Flow Sensor is configured to export IPFIX.).
- The Threat Feed host group branches cannot be renamed, changed, moved, or deleted.

Strategies for creating host groups

All Flow Collectors ship with a default host group structure. Depending on your login privileges, you can modify the default host groups to suit the needs of your network. You can create additional host groups as well as delete any of the default host groups except for these: Inside Hosts, Catch All, Outside Hosts, Countries, and Command & Control Servers.

We recommend that hosts that behave similarly to each other be placed in a host group together. However, you could create different host groups for each department within your network, geographical regions, IP segments, or any other categories that makes sense for your organization.

IPv4 IP address

If you enter this type of IP address...	You must use this formatting...
Single IP address	10.52.1.55 Includes only the host at 10.52.1.55.
Trailing dot subnet	10.52. Includes any IP address from 10.52.0.0 to 10.52.255.255.

If you enter this type of IP address...	You must use this formatting...
	Note: You must include the trailing period (.) at the end.
Classless Inter-Domain Routing (CIDR) notation (i.e., using a “/” to indicate the masked bits)	10.52.1.0/24 Includes any IP address from 10.52.1.0 to 10.52.1.255. Note: Use this notation if you prefer to enter the IP address using a “mask.”
Range of networks	10-11. Includes any IP address from 10.0.0.0 to 11.255.255.255 10.52-53. Includes any IP address from 10.52.0.0 to 10.52.255.255, and from 10.53.0.0 to 10.53.255.255 10.52-55.3. Includes any IP address from 10.52.3.0 to 10.52.3.255; from 10.53.3.0 to 10.53.3.255; from 10.54.3.0 to 10.54.3.255; and from 10.55.3.0 to 10.55.3.255 Note: Be sure to include the trailing period (.) at the end.
Range of hosts	10.52.1.0-10 Includes any IP address from 10.52.1.0 to 10.52.1.10 10.52.0.0-10.52.255.255 Includes any IP address from 10.52.0.0 to 10.52.255.255
Comma-separated list	10.52.1.10,10.52.1.50,10.100.1.20 Includes only these three hosts.

IPv6 IP address

If you enter this type of IP address...	You must use this formatting...
Single IP address	2001:0DB8:0000:0056:0000:ABCD:EF12:3456 2001:DB8:0:56:0:ABCD:EF12:3456 2001:DB8::56:0:ABCD:EF12:3456 2001:DB80:0:56::ABCD:EF12:3456 2001:DB80:0:56::ABCD:239.18.52.86
Global Routing Prefix subnet	2001:DB8:0:56::/64
Range of hosts	2001:DB8:0:56:ABCD:EF12:3456:1-10 2001:DB8:0:56:ABCD:EF12:3456:1- 2001:DB8:0:56:ABCD:EF12:3456:10

About IPv4 and IPv6 addresses

You can use either IPv4 or IPv6 addresses for hosts or host groups. The notation forms that must be used for each type are listed in the applicable table. You must separate the IP addresses by a comma, comma space, space, or return.

The Manager displays IPv6 addresses and IPv4 addresses simultaneously. Of the advantages provided by IPv6, the fact that it supports a larger address space is probably the most well known. IPv6 supports 128-bit addresses, whereas IPv4 supports only 32-bit addresses.

The Manager continues to support IPv4. The Flow Collectors can monitor IPv6 traffic, but they are not IPv6 addressable at this time.

Examples

- An IPv6 address consists of eight blocks of hexadecimal digits separated by colons, as shown in this example (where the x's represent hexadecimal digits):

Example: xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx

- You can replace blocks of 4-digit zeros with double colons (e.g., ::), but only once per IPv6 address. (Leading zeros can also be omitted.) Consider the following address (where the x's represent hexadecimal digits):

Example: xxxx:xxxx:0000:0000:0000:0000:0000:xxxx

This address can be rewritten like this:

Example: xxxx:xxxx::xxxx.

- Like an IPv4 address, an IPv6 address may include a prefix length that specifies how many left-most bits of the address make up the prefix. In IPv4 addresses, the prefix is used to determine the subnet mask (a.k.a., CIDR notation). In IPv6 addresses, it identifies the subnet that an interface belongs to, and routers use it for forwarding.

Example: xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx/40

- Take care when compressing an IPv6 address that contains a prefix, because if not done correctly, it can result in an incorrect interpretation when the address is uncompressed. You can write an IPv4 address as an IPv6 address, like this:

Example: xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:192.168.0.2

- If the first 6 blocks are zeros, you can compress the address like this:

Example: ::192.168.0.2

CIDR notation

Use the following table to understand classless inter-domain routing (CIDR) notation:

C I A S S	C I D R L e n g t h	M a s k	N u m b e r o f A d d r e s s e s
A	/1	128.0.0.0	2,147,483,648
	/2	192.0.0.0	1,073,741,824

Cl as s	CI DR Le ngt h	Mask	Numbe r of Addres ses
		0	741,824
	/3	224.0.0.0	536,870,912
	/4	240.0.0.0	268,435,456
	/5	248.0.0.0	134,217,728
	/6	252.0.0.0	67,108,864
	/7	254.0.0.0	33,554,432
	/8	255.0.0.0	16,777,216
	/9	255.128.0.0	8,388,608
	/10	255.192.0.0	4,194,304
	/11	255.224.0.0	2,097,152
	/12	255.240.0.0	1,048,576
	/13	255.248	524,288

Cl as s	CI DR Le ngt h	Mask	Numbe r of Addres ses
		.0.0	8
	/14	255.252 .0.0	262,14 4
	/15	255.254 .0.0	131,07 2
B	/16	255.255 .0.0	65,536
	/17	255.255 .128.0	32,768
	/18	255.255 .192.0	16,384
	/19	255.255 .224.0	8,192
	/20	255.255 .240.0	4,096
	/21	255.255 .248.0	2,048
	/22	255.255 .252.0	1,024
	/23	255.255 .254.0	512
C	/24	255.255 .255.0	256

Class	CIDR Length	Mask	Number of Addresses
	/25	255.255.255.128	128
	/26	255.255.255.192	64
	/27	255.255.255.224	32
	/28	255.255.255.240	16
	/29	255.255.255.248	8
	/30	255.255.255.252	4
	/31	255.255.255.254	2

Baselining

Baselining is critical to network monitoring because it builds a profile of what is considered normal behavior for your network. This allows Secure Network Analytics to trigger alarms when it observes abnormal behavior.

As soon as Secure Network Analytics is installed on a network, it begins identifying every host on the network. During the first seven days, Secure Network Analytics establishes a baseline of what appears to be normal behavior, based on approximately 90 attributes, such as the following:

- Regular bandwidth consumption.
- Communication with other hosts.

- Number of concurrent flows.
- Packets per second.


This baseline represents the expected behavior for the day. Baseline is used in conjunction with the tolerance to calculate the threshold values to use for that day.

These 90 attributes become part of the host profile mentioned earlier. By default, Secure Network Analytics baselines every host within the Inside Hosts host group. However, for the Outside Hosts host group, Secure Network Analytics baselines only aggregate host behavior at the host group level. You can change the baselining method at any time on this page.

A unique host level baseline is established for each host in a host group only if you select the "Enable baselining for Hosts in this Group" check box. Otherwise, Secure Network Analytics baselines aggregate host behavior at the host group level. As previously stated, by default, Secure Network Analytics baselines every host within the Inside Hosts group; therefore, by default this option is enabled for Inside Hosts.

If you don't want baselining switched on for a certain subgroup, we highly recommend that you create the applicable subgroup immediately under the Inside Hosts group. You must then uncheck the "Enabling baselining for hosts in the group" checkbox for that subgroup. If you don't create the subgroup immediately under the Inside Hosts group, you will have to uncheck this checkbox for every parent subgroup for that subgroup (with the exception of the Inside Hosts group). For detailed information, see [Disable baselining](#).

You can disable this feature for very dynamic environments, such as DHCP scopes, where IPs change frequently. If you do, this causes a baseline for expected behavior of any DHCP host to behave like all DHCP hosts.



Secure Network Analytics uses the following guidelines when baselining:

- For Host Baselining, Secure Network Analytics stores the host daily maximum value seen for each alarm that has been enabled (e.g., the High Total Traffic alarm).
- For Host Group Baselining, Secure Network Analytics stores the average of the maximum values for all hosts in the group (e.g., the max high total traffic for all hosts, averaged out).
- If a host has no daily values, it uses the baseline from the group with the least amount of hosts that it belongs to. For example, if a host belongs to two groups (Group A defined as 10.201.0.0/16, and Group B defined as 10.201.3.0/24), the baseline inherits Group B, which has fewer hosts.
- If a host group baseline is zero (0), the maximum is used (e.g., the High Total Traffic maximum bytes in 24 hours).

- For a new installation, all hosts use the configuration policy maximum for the first day until the baselines are established. Hosts do not alarm unless they exceed the maximum values.

Going forward, Secure Network Analytics looks for and highlights changes in behavior, such as the following:

- One host contacting large numbers of other hosts in a short period of time (e.g., peer-to-peer applications, worms).
- Long flow durations (e.g., covert channels, VPNs).
- Use of unauthorized ports (e.g., rogue servers/applications).
- Bandwidth anomalies (e.g., Warezserver, denial of service).
- Unauthorized communications (e.g., a VPN host communicating with an accounting server).

Whenever a host exceeds its threshold of what Secure Network Analytics has baselined as “normal” behavior, Secure Network Analytics triggers an alarm. By observing a host’s behavior as the behavior occurs and by using several proprietary algorithms, Secure Network Analytics avoids generating false positive alarms like those often generated by signature-based solutions.

What you can do on this page

View host group details

Refer to the following information about the fields on this page. All information entered on this page is reflected in the Desktop Client as well as in the Manager, with the exception of the Description field, which appears only in the Manager.

Host Group Tree

To view host group details

In the left section of the page, within the Host Group tree, select the host group whose details you want to view.

If you type any letters that are anywhere within a host group name into the Filter field and press **Enter**, the Host Group tree expands to display (in italicized red font) any host group name that contains that entry.

Example: If you entered "South" into the Search field, some of the host group names that would then be displayed in italicized red font would be Southern Africa, South Africa, South-Eastern Asia, French Southern Territories.

To return an exact match

To return an exact match, enter the entire host group name into the Search field and press **Enter**. If you want to return to the previously selected host group, delete the entry in the Search field and press **Enter**.

To change the domain name

You can change the domain to which the host group tree you want to view belongs by clicking the drop-down arrow and selecting the desired domain. The domain you are viewing is also displayed at the top of the Host Group tree. You can choose only one host group at a time.

Host Group Name

You can duplicate host group names as long as they are not located under the same parent host group.

Parent Host Group

The immediate parent of the host group selected in the Host Group tree. A double dash (-) is displayed in this field when viewing a top-level parent host group (Inside, Outside, Bogn, Command & Control Servers, Tor).

This field displays "- -" if a host group at the highest level (e.g., Inside Hosts) has been selected.

Description

You can enter whatever content is pertinent to your need. There is a limit of 512 characters.

IP Addresses and Ranges

You can use either IPv4 or IPv6 addresses. (Refer to the "IPv4 and IPv6 address formats" section in this topic for the correct notation forms to use.)

Refer to the following guidelines:

- You must separate the IP addresses by a comma, comma space, space, or return.
- You can import files that use only a .csv or .txt file extension.
- You can paste entries from an Excel, text, or csv file.

Advanced Options

Refer to the following table for descriptions of the advanced options. You can use some of these options to prevent certain alarms from being triggered for a host group when

specific criteria are met.

Option	Description
Enable baselining for hosts in this group.	If you select this check box, a unique host level baseline is established only for each host in a host group. Otherwise, Secure Network Analytics baselines aggregate host behavior at the host group level. As previously stated, by default, Secure Network Analytics baselines every host within the Inside Hosts group; therefore, by default this option is enabled for Inside Hosts.
Disable security events using excluded services.	When this setting is enabled, security events for the host group(s) are disabled for any security event using a service that has the "Exclude" setting enabled on the TCP/UDP Service dialog.
Disable flood alarms and security events.	When this setting is enabled, flood alarms and flood security events are disabled when any host in the host group(s) is the target of a flood attack.
Trap hosts that scan unused addresses in this group.	Select this check box to enable the "trapping" of hosts that appear to be performing "low and slow" scans of the host group(s). ("Low and slow" scanning is the process of scanning the hosts/ports on a network over a period of hours to months. These types of scans may be performed in order to prevent discovery of the attacker, the ports being scanned, or the hosts being scanned, or even to pre-scan before release of a worm.) Any host that attempts to communicate with an unused address in the host group(s) (on any port) is counted. The scanning host receives a Trapped Host alert if the daily count threshold is exceeded. Once the host exceeds the number of days in that month that it can receive a Trapped Host alert, a Trapped Host alarm is triggered. After a host triggers a Trapped Host alarm, its Concern Index is raised to its threshold, making it a High Concern Index host. As such, the host generates

Option	Description
	Touched alarms that help you track any further activity. This is an advanced feature that should only be enabled for select small, well-controlled host groups with fixed IP addresses and a stable number of hosts, such as critical servers. This allows you to detect hosts intruding into the most important part of your network. Hosts that scan unused addresses in the Outside Hosts host group are not included in trapped host calculations and therefore do not generate trapped host alerts and alarms, regardless of whether this setting is enabled in the Outside Hosts host group.

Add a new host group

1. Do one of the following:

- If you are adding a top-level host group, click the context menu next to the domain in the Host Group tree and select **Add Host Group**.
- If you are adding a host group at any other level, click the context menu next to the host group that is to be the parent of the one you are adding and select **Add Host Group**.

The parent is displayed in the Parent Host Group field, and the Advanced Options that have been configured for the highest-level host group of the newly created host group are selected accordingly. All other fields are empty.

- ### 2. Next, configure the necessary settings on this page. The Host Group Name field is required; all other fields are optional. You can use duplicate host group names, but not at the same host group level (i.e., below the same parent host group).



The Advanced Options are not available to configure for Inside Hosts, Outside Hosts, Threat Feed, and Countries.

- ### 3. After you click **Save**, you see the newly created host group in the Host Group tree, listed alphabetically with all the other host groups that share that same parent. A check mark is displayed beside the new host group.

Disable baselining

If you don't want baselining switched on for a certain subgroup, we highly recommend that you create the applicable subgroup immediately under the Inside Hosts group. You must then uncheck the "Enabling baselining for hosts in the group" checkbox for that subgroup.

1. Click the host group for which you want to disable baselining.
2. Click **Edit** in the the top right corner of the page.

Its associated information is displayed in the Host Group Details section on the right side of the page.

3. Uncheck the "Enabling baselining for hosts in this group" checkbox.



If you don't create the subgroup for which you want to disable baselining immediately under the Inside Hosts group, you must uncheck this checkbox for every parent subgroup for that subgroup (with the exception of the Inside Hosts group); otherwise, baselining will not be disabled.

Edit a host group

In the Host Group tree, do the following:

1. Click the host group you want to edit.
2. Click **Edit** in the the top right corner of the page.

Its associated information is displayed in the Host Group Details section on the right side of the page.

3. Make your edits and click **Save**.

Refer to the following table for a list of which fields are editable for each host group type.

For this host group...	You can edit...
Inside Outside Threat Feed Countries	Not editable.

For this host group...	You can edit...
Catch All	Description, IP Address and Ranges, Advanced Options.
Any custom host group	All fields and Advanced Options.

Delete a host group

In the Host Group tree, click the context menu next to the host group you want to delete and select **Delete Host Group**.

After confirming that you want to delete the host group, it is deleted from the Host Group tree.

You cannot delete the following host groups (the Delete Host Group option is disabled for them):

- Inside Hosts
- Outside Hosts
- Threat Feed
- Countries
- Catch All

When you delete a custom host group, all of its child host groups are deleted as well.



Deletion of a host group can invalidate custom security events and cause missing or false alarms. For more information, see the "When you delete a host group, service, or application" section in [Configuring Custom Policies \(Custom Security Events\)](#).

Move host group

In the Host Group tree, drag and drop the host group over the host group which is to be its parent. A green check mark indicates that you can move the host group that is highlighted below your cursor. If it is not allowed, a red X is displayed. You cannot move the following host groups:

- Inside Hosts
- Outside Hosts

- Threat Feed
- Countries
- Catch All

Add IP Addresses and Ranges

To enable you to more easily set up your host groups, you can add IP addresses and ranges. You can perform an import or enter them manually.

To import IP addresses and ranges, do the following:

1. Click **Import IP Addresses and Ranges**.
2. Select the file you want to import. You can import files that use only a .csv or .txt file extension.

If you prefer, you can type, or cut and paste, IP addresses and ranges into the IP Addresses and Ranges field. (Refer to the "IPv4 and IPv6 address formats" section in this topic for the correct notation forms to use when importing IP address.)

Import entire host group tree

At times you may want to replace your current host group tree. Perhaps you want to replace it with the host group tree that is currently being used in a different environment. When you perform an import, all host groups are imported, including Inside Hosts, Outside Hosts, Countries, and Catch All.

To import the entire host group tree for the current domain, click **Import All** (located under the host group tree).



- Secure Network Analytics overwrites any existing host groups after the import. If you receive any validation errors while importing, the entire import fails.
- Host groups under different parents can have the same name, but two host groups under the same parent cannot share the same name.

Refer to the following table to learn how the Threat Feed is handled during an import.

If you are importing a host group tree from an environment where Threat Feedis...	And in your environment Threat Feed is...	Then...
On	Off	Threat Feed host groups are not imported.
Off	On	Threat Feed host groups remain after the import.

Import a single host group, including any child host group(s)

To import a specific set of host groups into the Host Group tree, do the following:

1. Click the context menu next to the host group into which you are importing.
2. Click **Import Host Group Sub Tree**.
3. Select the applicable file to import.

 If you receive any validation errors while importing, the entire import fails.

Export entire host group tree

At times you may want to export the host group tree. Perhaps you want to use it to replace the host group tree that is currently being used in a different environment. When you perform an export, all host groups are exported, including Inside Hosts, Outside Hosts, Countries, and Catch All.

To export the entire host group tree for the current domain, do the following:

1. Click **Export All** (located under the host group tree).

The XML file is generated.

2. Save the file to the desired location.

Export a single host group, including any child host group(s)

To export the selected host group, do the following:

1. Click the context menu next to the host group to be exported in the Host Group tree.
2. Click **Export Host Group Sub Tree**.

You are asked if you want to include its child host groups. After you respond to this question, the XML file is generated.

3. Save the file to the desired location.

What should you do next

To view alarm activity associated with a particular host group, go to the [Host Group Report](#).

Related topics

[Investigating Host Groups](#)

[User Management: Users Tab Overview](#)

[User Management: User Configuration Overview](#)

[Data Role Management: Data Roles Tab Overview](#)

[Data Role Management: Data Role Configuration Overview](#)

[Domain Configuration: Overview](#)

Configuring Host Group Automation

Requirement: Make sure you've configured Cisco Identity Services Engine (ISE).
 For more information about Cisco ISE, refer to the [ISE and ISE-PIC Configuration Guide](#).

Use this page to configure host group automation. Organizing hosts into host groups allows you to monitor different areas of your network and manage host behavior more efficiently. For example, when you add a host to a host group, it automatically inherits the policy associated with that host group. When you need to change a host's policy, it is easier to re-associate the host to a more relevant host group rather than to re-associate the host group along with modifying any number of individual policies that may have been configured with that particular host included.

Note the following:

- New host groups can only be created when active ISE endpoints exist.
- Host group automation is not supported on a secondary Manager.
- Manager appliance backup files include host group automation services.
- If the archive hour is changed, the **Next Run** time stamp won't be updated until a manual or automatic run occurs.

For more details about host groups, refer to [Managing and Configuring Host Groups](#).

Do any of the following to access this page

- From the main menu:
 1. Choose **Configure > Detection > Host Group Management**.
 2. Click **Automation**.
- From the [Host Group Report](#) Actions drop-down menu:
 1. Choose **Host Group Management**.
 2. Click **Automation**.

Which web roles can access this page, and what permissions do they have

For a complete list of permissions assigned to each web role, on the User Management:

Users tab, click the  (**Information**) icon beside the web role whose assigned permissions you want to view.

If a user has this web role...	Then the user...
Primary Admin	Can view and make edits within Host Group Management.
Power Analyst	Can view and make edits within Host Group Management.
Analyst	Cannot access the Host Group Management page.
Config Manager	Cannot access the Host Group Management page.
Multiple web roles	Possesses permissions that reflect the union of those web roles (i.e., whichever permissions are common among all the web roles assigned to the user).

Configure host group automation

To configure host group automation, do the following:

1. Click **Create > ISE**.
2. Complete the following fields for the **ISE Service**:
 - Name:** Type a unique name for the ISE service. (There is a limit of 255 characters.)
 - Description:** Provide a description for the ISE service. (There is a limit of 255 characters.)
 - Parent Host Group Name:** Choose **Inside Hosts** or **Outside Hosts** for the immediate parent of the host group.
 - Host Group Name:** Type a name for the host group. (There is a limit of 255 characters.) You can duplicate host group names as long as they are not located under the same parent host group.
3. Choose an **ISE Cluster Name** from the drop-down menu. These are the ISE clusters you configured in Secure Network Analytics on the Cisco ISE configuration page.
4. Check or uncheck the **Advanced Options** check boxes. Refer to the following table for descriptions of the advanced options. You can select some of these options to prevent certain alarms from being triggered for a host group when specific criteria are met.

Option	Description
Enable baselining for hosts in this group.	If you select this check box, a unique host level baseline is established only for each host in a host group. Otherwise, Secure Network Analytics baselines aggregate host behavior at the host group level. As previously stated, by default, Secure Network Analytics baselines every host within the Inside Hosts group. This option is in enabled by default for Inside Hosts.
Disable security events using excluded services.	When this setting is enabled, security events for the host group(s) are disabled for any security event using a service that has the "Exclude" setting enabled on the TCP/UDP Service dialog. This option is enabled by default for both Inside Hosts and Outside Hosts.
Disable flood alarms and security events.	When this setting is enabled, flood alarms and flood security events are disabled when any host in the host group(s) is the target of a flood attack.

Option	Description
Trap hosts that scan unused addresses in this group.	Select this check box to enable the "trapping" of hosts that appear to be performing "low and slow" scans of the host group(s). ("Low and slow" scanning is the process of scanning the hosts/ports on a network over a period of hours to months. These types of scans may be performed in order to prevent discovery of the attacker, the ports being scanned, or the hosts being scanned, or even to pre-scan before release of a worm.) Any host that attempts to communicate with an unused address in the host group(s) (on any port) is counted. The scanning host receives a Trapped Host alert if the daily count threshold is exceeded. Once the host exceeds the number of days in that month that it can receive a Trapped Host alert, a Trapped Host alarm is triggered. After a host triggers a Trapped Host alarm, its Concern Index is raised to its threshold, making it a High Concern Index host. As such, the host generates Touched alarms that help you track any further activity. This is an advanced feature that should only be enabled for select small, well-controlled host groups with fixed IP addresses and a stable number of hosts, such as critical servers. This allows you to detect hosts intruding into the most important part of your network. Hosts that scan unused addresses in the Outside Hosts host group are not included in trapped host calculations and therefore do not generate trapped host alerts and alarms, regardless of whether this setting is enabled in the Outside Hosts host group.

- Click **Test Connection** to confirm connection to the ISE cluster.



If the connection fails, make sure the ISE cluster status shows as connected in Secure Network Analytics on the Cisco ISE configuration page.

- Click **Create**.

 If you click **Cancel**, your host group automation service will be discarded.

7. Make sure to **Enable** the host group automation when ready for it to run automatically.

Enable

 When you enable a host group automation service, it schedules a recurring job that runs once a day (one hour after archive hour).

From the Host Group Automation table, go to the row that contains the one you want to enable and choose **Enable** option in the Enabled column.

Run Now

From the Host Group Automation table, go to the row that contains the one you want to run:

1. Click the **⋮ (Ellipsis)** icon in last column, then choose **Run Now** from the drop-down menu.
2. Click **Continue** when the dialog box displays, or click **Cancel** if you don't want to run the host group automation.

 While the **Status** column displays Queued or Running, the only action you can perform is Delete to remove the host group automation service. Click the  **(Refresh)** icon above the table to view the updated status. When the job has completed, the **Status** column displays Success or Failed.

Edit

From the Host Group Automation table, go to the row that contains the one you want to edit:

1. Click the **⋮ (Ellipsis)** icon in the last column, then choose **Edit** from the drop-down menu.
2. Make the necessary edits when the Host Group Automation Setup slide-in pane displays.
3. Click **Save** to save your changes.

Disable



If you disable a host group automation service, all future jobs scheduled to run will be canceled.

From the Host Group Automation table, go to the row that contains the one you want to disable and choose **Disable** option in the Enabled column.

Delete

From the Host Group Automation table, go to the row that contains the one you want to delete:

1. Click the **⋮ (Ellipsis)** icon in the last column, then choose **Delete** from the drop-down menu.
2. Click **Delete** when the dialog box displays.
3. Click **Cancel** if you don't want to delete the host group automation service.

What should you do next

Go to the Host Group Tree to view the automation results. Refer to [Managing and Configuring Host Groups](#) for details about the Host Group Tree.

Related topics

[Managing and Configuring Host Groups](#)

About Alarm Severities

Use this page to configure an alarm severity for individual alarms. The alarms listed here are based on the domain that is currently selected, so keep in mind that when you assign a severity to an alarm on this page, you are doing so for the selected domain.

Manager System alarms are the only type of alarm that are configured for your entire system instead of per domain, since they affect your entire system. Therefore, when you configure an alarm severity for this alarm type, you need to do so for *all your domains*, because a Manager System alarm will display the highest severity assigned to it in any given domain.



So, if you have various severities configured for a Manager System alarm across different domains, the severity displayed when the alarm is triggered will not always be the one you assigned to it for a particular domain. For example, if you configure the SMC Failover Channel Down alarm for Domain 1 with a Major severity, and you configure this alarm for Domain 2 with a Trivial severity, then when the alarm is triggered, it will display a Major severity in both Domain 1 and Domain 2, since this is the highest severity it has been assigned among all existing domains.

The following alarm categories exist in Secure Network Analytics. For more information about the various alarm categories, see [Alarm List: Overview](#).

- Exporter/Interface
- Flow Collector System
- Host (also referred to as security events)
- Host Group Relationship
- Manager System
- UDP Director

Access the Alarm Severity page

From the main menu, choose **Configure > Detection > Alarm Severity**.

The Alarm Severity page opens.

Assign an alarm severity to an alarm type

For each alarm type that you are classifying, choose the desired severity from its drop-down menu.

At the archive hour, all security events associated with the alarm are cleared.

Most security events contribute index points to an alarm. An alarm is a notification of unusual network activity that meets or exceeds a defined set of criteria that can indicate unacceptable behavior on your network. Any alarm number or associated alarm trend in the Manager that reflects alarm severity depends on the following criteria:

- The number of index points it receives from security events
- Alarm severity
- An "always bad" security event is associated with the alarm

Refer to the following table to learn the definition for each alarm severity.

Alarm Severity	Alarm Definition
Informational	No index value.  When indicated by a color, it is light blue. This severity is not used by all features in the Manager.
Trivial	No index value.  When indicated by a color, it is dark blue. This severity is not used by all features in the Manager.
Minor	Index value is less than the configured threshold and greater than 0. * These are your catch-all alarms that do not meet the requirements of the higher-priority categories. Minor alarms may or may not be tuned or be of interest. They may be useful for a general correlation of network events. For example, if you have had relatively few Minor High Traffic alarms, and then one day there are suddenly dozens or hundreds of them occurring, this may be indicative of a security threat on your network. As a Minor alarm is identified to be of interest, it can be elevated to Major or even directly to Critical.

Alarm Severity	Alarm Definition
	 When indicated by a color, it is yellow.
Major	Index value has crossed the threshold and is less than twice the difference between the threshold and baseline. * A Major alarm should be of interest to you. When you have tuned a Major alarm to the point that you believe it is a valuable source of intelligence, you can re-assign it to Critical.  When indicated by a color, it is orange.
Critical	Index value is above the baseline plus twice the difference between the threshold and baseline. * A Critical alarm is well-tuned, well-understood, and typically a low-volume alarm. The chance of a false positive is generally quite low.  When indicated by a color, it is red.
Critical	An "always bad" security event is active. Note: This alarm severity type is not shown in the Manager and you cannot assign this category of alarm severity to an alarm. For more information about this type of alarm, see the "Always Bad" security events section in this topic.  When indicated by a color, it is black. This severity is not used by all features in the Manager.

* For more information about thresholds, see [Threshold categories](#).

Filter columns if you need to narrow your search

At the top of each of the columns there is a filter field in which you can enter the criteria by which you want to filter that column. As you start to type your entry, the column dynamically filters to display a list of entries that contains any of the characters you have entered.

You can click the  (**Filter**) icon to see a list of available filter options (for example, Is equal to; Is not equal to; Contains) from which you can choose to help you filter the list according to a specific rule. The "Contains" rule is the default filter mode. When you select an option, the  (**Clear Filter**) icon is displayed.

Clearing entries

To clear the entry, click the Clear Filter icon.

Sort columns

To reverse the sort order for a column, click the column name. Use the column name to toggle between ascending and descending order. You can also click the arrow that is displayed to the right of the column name (after you click the column name). An up arrow indicates that the column is sorted in ascending order, and a down arrow indicates descending order.

"Always Bad" security events

Some security events are so serious ("always bad") that they are always considered a high severity security event, and thus the alarms they trigger are automatically assigned the highest severity. These security events do not have index points associated with them or contribute to the percentage of threshold reached for an alarm. However, this type of security event immediately increases the alarm count for a host and assigns it a severity of "always bad." The following security events are considered "always bad":

- All custom security events
- Any security event contributing to the Policy Violation alarm



These security events may contribute to other alarms as well, but they cause only the Policy Violation alarm to turn black in the Desktop Client.

Threshold categories

There are two categories of threshold: minimum and maximum

- The minimum threshold is a static value that indicates the lowest number of index points to allow before an alarm is no longer triggered. An alarm is not triggered when the number of index points falls below this minimum threshold value.
- The maximum threshold is a static value that indicates the highest number of index points to allow without triggering an alarm. An alarm is always triggered when the number of index points exceeds this maximum threshold value.

For more information about threshold values, see [About Alarm Categories](#)

Related topics

[About Alarm Categories](#)

Alarm List: Overview

Both alarm categories and security events can trigger alarms. Security events are also referred to as alarms.

An alarm is simply a security event for which an alarm has been issued once index points have reached a pre-configured level. Using the Host Policy Manager, you can configure the Manager (formerly Stelathwatch Management Console) to issue an alarm for individual security events, alarm categories (to which particular security events contribute), or both. When an alarm is received, it remains active until the original alarm condition no longer exists.

A security event is the mechanism that assigns index points to an alarm category. If you disable a security event, index points will not accumulate against it or any alarm category that the security event is associated with.

Security events can be disabled for specific services and disabled at the host group level. For more information, see the "How do I know which options to select for source and target" section in [Managing Core Policies](#).

You can also assign a severity level to individual security event alarms. For more information, see [About Alarm Severities](#).

The following supporting files are available:



- The [Internal Alarm IDs document](#) lists each alarm with its respective internal alarm ID and name.
- The [MIBs file](#) provides the MIB files that are included in this release. MIBs allow SNMP-based network alarm management systems to interpret the messages that the Manager sends using SNMP traps. If you want your Manager to send the respective SNMP traps (listed in the table below) when an alarm occurs, you must enable that feature using the Response Management feature in Secure Network Analytics (formerly Stealthwatch).

Following are the system alarms that exist in Secure Network Analytics. Any of them may be referenced by a system notification.

[Exporter/Interface Alarms](#)

[Flow Collector System Alarms](#)

[Host Alarms](#)

[Host Group Relationship Alarms](#)

[Manager System Alarms](#)

[UDP Director Alarms](#)

Alarm List: Exporter/Interface Alarms

These alarms occur when problems affect exporters or interfaces.

Alarm	Alarm Category	SNMP Trap (sent by Manager)
Flow Collector Flow Data Lost	Exporter or Interface Alarm	exporterAlarmCondition
Indicates that no exports have been received by the corresponding Flow Collector (NetFlow) or Flow Collector (sFlow) from a previously active exporter within the last 30 minutes. You can enable/disable this alarm on the Flow Collector Configuration page.		
Flow Collector Longest Export Exceeded	Flow Collector System Alarm	systemAlarmCondition
Triggered when the flow duration from an exporter has exceeded the threshold setting. If not remedied, inaccurate flow and interface stats will be generated. You can enable/disable this alarm on the Flow Collector Configuration page.		
Flow Sensor Management Channel Down	Exporter or Interface Alarm	exporterAlarmCondition
Indicates that communication is lost between the Manager and a Flow Sensor Virtual Edition or Flow Sensor. The Details column in the Alarm Table provides one of the previous reasons for the alarm. • No route to host • Connection error • Socket error • Http protocol error • Http authorization error • Http client error • Http server error • Http other error		

Alarm	Alarm Category	SNMP Trap (sent by Manager)
• Response handler error (something wrong with the response handler running on the device) • Other error (in this case, a detailed java exception description is also included) • The certificate for the FlowSensor is not trusted, has expired, or has been revoked You can enable/disable this alarm on the Flow Collector Configuration page.		
Flow Sensor RAID Failure	Exporter or Interface Alarm	systemAlarmCondition
Indicates that a failure has been detected in the RAID configuration of a Flow Sensor within the last hour.		
Flow Sensor RAID Rebuilding	Exporter or Interface Alarm	systemAlarmCondition
Indicates that a RAID drive has started the process of rebuilding after a RAID failure occurred on a Flow Sensor.		
Flow Sensor Time Mismatch	Exporter or Interface Alarm	exporterAlarmCondition
Indicates that there is more than 30 seconds difference between the time on the Flow Sensor Virtual Edition or Flow Sensor and the time that the Flow Collector (NetFlow) expects for the FlowSensor. If this alarm occurs, the Network Time Protocol (NTP) server may not be enabled on the device that has drifted (i.e., on either the Flow Collector (NetFlow) or the Flow Sensor). You can enable/disable this alarm on the Flow Collector Configuration page.		
Flow Sensor Traffic Lost	Exporter or Interface Alarm	exporterAlarmCondition
Indicates that the Flow Sensor has not seen packets in a 30-minute period from the indicated interface. You can enable/disable this alarm on the Flow Collector Configuration page.		

Alarm	Alarm Category	SNMP Trap (sent by Manager)
Interface Utilization Exceeded Inbound	Exporter or Interface Alarm	exporterAlarmCondition
Indicates that the average bits per second rate for inbound traffic has exceeded the specified utilization threshold setting for the exporter interface (i.e., the interface on the exporter that sends data to the Flow Collector) for the specified minimum amount of time. The alarm becomes inactive if the inbound utilization falls below the threshold. You can enable/disable this alarm on the Flow Collector Configuration page.		
Interface Utilization Exceeded Outbound	Exporter or Interface Alarm	exporterAlarmCondition
Indicates that the average bits per second rate for outbound traffic has exceeded the specified utilization threshold setting for the exporter interface (i.e., the interface on the exporter that sends data to the Flow Collector) for the specified minimum amount of time. The alarm becomes inactive if the outbound utilization falls below the threshold. You can enable/disable this alarm on the Flow Collector Configuration page.		

Related Topics

[Alarm List: Overview](#)

Alarm List: Flow Collector System Alarms

These alarms occur when a problem affects a Flow Collector.

Alarm	Alarm Category	SNMP Trap (sent by Manager)
Flow Collector Data Deleted	Flow Collector System Alarm	systemAlarmCondition
Indicates that data has been deleted from the database on the Flow Collector to keep the hard drive from running out of disk space. This means that less than the specified number of days of database data is available. The number of days threshold is set on the Flow Collector Appliance Administration Web UI. To access it, right-click the Flow Collector in the Enterprise Tree in the Manager and select Configuration > Appliance Administration. When the Flow Collector Web UI opens, click Support > Advanced Settings. Scroll to the summary_retention_days row and change the value as necessary. You can enable/disable this alarm on the Flow Collector Configuration page.		
Flow Collector Database Channel Down	Flow Collector System Alarm	systemAlarmCondition
Indicates that the network connection between a Flow Collector and its database node is not working. Check to see if a cable is unplugged or plugged into the wrong location. You can enable/disable this alarm on the Flow Collector Configuration page.		
Flow Collector Database Unavailable	Flow Collector System Alarm	systemAlarmCondition
Indicates that the Flow Collector's database is not running. You can enable/disable this alarm on the Flow Collector Configuration page.		
Flow Collector Database Updates Dropped	Flow Collector System Alarm	systemAlarmCondition
Database updates for the following telemetry types (if enabled) are currently being dropped:		

Alarm	Alarm Category	SNMP Trap (sent by Manager)
• Firewall log event updates • NVM flow updates • NetFlow flow updates This condition typically occurs when your Flow Collector either cannot reach the Data Store database or your Data Store database has remained unreachable for an extended period of time. This could be due to either a network connectivity problem or the Data Store database being down. When this occurs, the Flow Collector has already buffered one hour's worth of database updates and can no longer buffer additional updates. When the Flow Collector Database Updates Dropped alarm is triggered, the Flow Collector Performance Degraded alarm is subsequently triggered.		
Flow Collector Exporter Count Exceeded	Flow Collector System Alarm	systemAlarmCondition
Indicates that the number of active NetFlow or sFlow exporters has exceeded the maximum allowed for your model of Flow Collector (NetFlow) or Flow Collector (sFlow). This alarm prevents the Flow Collector from adding exporters, but it does not prevent the Flow Collector from adding exporters to another Flow Collector that is not full. When this condition exists, there are no more exporter slots left in the Flow Collector, so no more exporters can be added. This alarm usually occurs when the Flow Collector has been configured to allow any exporter.		
Flow Collector Interfaces Count Exceeded	Flow Collector System Alarm	systemAlarmCondition
Indicates that the specified exporter interface cannot be added to the Flow Collector because the limit permitted for this Flow Collector has been reached. The alarm information provided in the Alarm Table includes the exporter IP address for the interface. To resolve this issue, you can move this exporter interface or others to another Flow Collector that has not reached its capacity.		
Flow Collector Performance Degraded	Flow Collector System Alarm	systemAlarmCondition
Indicates that the Flow Collector is experiencing conditions that are preventing it from		

Alarm	Alarm Category	SNMP Trap (sent by Manager)
performing at its best. Further problems are likely to occur if the conditions are not corrected. If you receive this alarm, provide the sw.log file to Cisco Support. This alarm can occur under any of the following conditions as indicated in the Details column in the Alarm Table: - Database Connect – Secure Network Analytics was unable to connect to the database. Flows will not be logged. - Database Inserts Delayed – Secure Network Analytics took too long to insert the data into the database. - Database Inserts Disabled For Interface Statistics – Secure Network Analytics attempted 10 times in a row unsuccessfully to insert the data into the database. Database inserts for interface statistics have been disabled. - Error – A serious error was detected and logged. - Flow Analysis Time – Secure Network Analytics took a long time to process flows. - Low Memory – Secure Network Analytics is dangerously low on memory. - Manager Connect – Secure Network Analytics was unable to connect to the Manager. - Oversubscribed – Sustained input data rate is too high for system configuration to process all input data. - Program Error – An internal error has been detected in the engine code. - Some flows cannot be stored because the flow cache is full – For 5 minutes or more, the Flow Collector has not been able to store some of the flows it has processed because there is no more room in the flow cache. All flows are processed, but some data may not be visible in the Manager client user interface. If this is a Flow Collector VE, you can try to remedy the situation by allocating more memory resources to the Flow Collector VE. Otherwise, contact Cisco Support. - Telemetry Generation on Flow Collector Engine Disabled – Low disk space on the Flow Collector. - Unknown – The cause of the alarm cannot be determined.		
Flow Collector RAID Failure	Flow Collector System Alarm	systemAlarmCondition

Alarm	Alarm Category	SNMP Trap (sent by Manager)
Indicates that a failure has been detected in the RAID configuration of the Flow Collector within the last hour.		
Flow Collector RAID Rebuilding	Flow Collector System Alarm	systemAlarmCondition
Indicates that a RAID drive has started the process of rebuilding after a RAID failure occurred on a Flow Collector.		
Flow Collector Stopped	Flow Collector System Alarm	systemAlarmCondition
Indicates that the Flow Collector is not running. If you receive this alarm, attempt to reboot the Flow Collector and provide the sw.log file to Cisco Support. The Details column in the Alarm Table indicates the specific problem, as described below: • Program – Secure Network Analytics encountered an unrecoverable problem that required an immediate restart. • Time Change – Secure Network Analytics either detected a significant change in the system time or was prevented by the system from running scheduled processes in time. • Error – Secure Network Analytics encountered an unrecoverable problem. • Capture – Secure Network Analytics encountered an unrecoverable problem that required an immediate restart. • Processing – Secure Network Analytics encountered an unrecoverable problem that required an immediate restart. • Scheduler – Secure Network Analytics encountered an unrecoverable problem that required an immediate restart. • Messaging – Secure Network Analytics encountered an unrecoverable problem that required an immediate restart. • Unknown – The cause of the alarm cannot be determined.		

Related Topics

[Alarm List: Overview](#)

Alarm List: Host Alarms

For information about the security events that can trigger host alarms, see [Security Event List](#).

Related Topics

[Alarm List: Overview](#)

Alarm List: Host Group Relationship Alarms

These alarms alert you to possible problems affecting the data transmitted between two host groups.

Alarm	Alarm Category	SNMP Trap (sent by Manager)
Relationship High Total Traffic*	Host Group Relationship Alarm	N/A
For a host group relationship, this alarm indicates that the total traffic between the two host groups in the relationship exceeds the threshold. The alarm is raised if the alarm condition exists for longer than a user-specified duration. You can enable/disable this alarm, or modify the minimum and maximum thresholds, on the Host Group Relationship Properties dialog. The accumulated total traffic for all host group relationships is cleared at the archive hour.		
Relationship High Traffic*	Host Group Relationship Alarm	N/A
For a host group relationship, this alarm indicates that the traffic rate between the two host groups in the relationship exceeds the threshold. The alarm is raised if the alarm condition exists for longer than a user-specified duration. You can enable/disable this alarm, or modify the minimum and maximum thresholds, on the Host Group Relationship Properties dialog.		
Relationship ICMP Flood*	Host Group Relationship Alarm	N/A
For a host group relationship, this alarm indicates that an excessive number of ICMP has been sent, in either direction, between the two host groups in the relationship. The alarm is raised if the alarm condition exists for longer than a user-specified duration. You can enable/disable this alarm, or modify the minimum and maximum thresholds, on the Host Group Relationship Properties dialog.		
Relationship Low Traffic*	Host Group Relationship Alarm	N/A

Alarm	Alarm Category	SNMP Trap (sent by Manager)
For a host group relationship, this alarm indicates that the traffic rate between the two host groups in the relationship has fallen below the minimum acceptable threshold. The alarm is raised if the alarm condition exists for longer than a user-specified duration. You can enable/disable this alarm, or modify the minimum and maximum thresholds, on the Host Group Relationship Properties dialog.		
Relationship Max Flows*	Host Group Relationship Alarm	N/A
For a host group relationship, this alarm indicates that the maximum number of flows between the two host groups in the relationship has exceeded the threshold. The alarm is raised if the alarm condition exists for longer than a user-specified duration. You can enable/disable this alarm, or modify the minimum and maximum thresholds, on the Host Policy Manager dialog.		
Relationship New Flows*	Host Group Relationship Alarm	N/A
For a host group relationship, this alarm indicates that the total number of new flows between the two host groups in the relationship has exceeded the threshold. The alarm is raised if the alarm condition exists for longer than a user-specified duration. You can enable/disable this alarm, or modify the minimum and maximum thresholds, on the Host Policy Manager dialog.		
Relationship Round Trip Time*	Host Group Relationship Alarm	N/A
For a host group relationship, this alarm indicates that the average amount of time (in milliseconds) required for all the TCP connections to occur in the flows between the two host groups in the relationship exceeds the threshold. The alarm is raised if the alarm condition exists for longer than a user-specified duration. You can enable/disable this alarm, or modify the minimum and maximum thresholds, on the Host Group Relationship Properties dialog. This alarm is not supported by Flow Collector (NetFlow) unless it is used in conjunction with a Flow Sensor or Flow Sensor Virtual Edition.		

Alarm	Alarm Category	SNMP Trap (sent by Manager)
Relationship Server Response Time*	Host Group Relationship Alarm	N/A
For a host group relationship, this alarm indicates that the average amount of time (in milliseconds) between the first client request and the first server response, for all the TCP connections in the flows between the two host groups in the relationship, exceeds the threshold. The alarm is raised if the alarm condition exists for longer than a user-specified duration. You can enable/disable this alarm, or modify the minimum and maximum thresholds, on the Host Group Relationship Properties dialog. This alarm is not supported by Flow Collector (NetFlow) unless it is used in conjunction with a Flow Sensor or Flow Sensor Virtual Edition.		
Relationship SYN Flood*	Host Group Relationship Alarm	N/A
For a host group relationship, this alarm indicates that the difference between the number of TCP connection requests (SYN packets) and the number of TCP acknowledgments (SYN-ACK packets), in either direction between the two host groups in the relationship, exceeds the threshold. The alarm is raised if the alarm condition exists for longer than a user-specified duration. You can enable/disable this alarm, or modify the minimum and maximum thresholds, on the Host Group Relationship Properties dialog.		
Relationship TCP Retransmission Ratio*	Host Group Relationship Alarm	N/A
For a host group relationship, this alarm indicates that the percentage of TCP packets that were retransmitted during the flows between the two host groups in the relationship exceeds the threshold. Retransmission usually occurs because packets are either damaged or lost. The alarm is raised if the alarm condition exists for longer than a user-specified duration. You can enable/disable this alarm, or modify the minimum and maximum thresholds, on the Host Group Relationship Properties dialog.		

Alarm	Alarm Category	SNMP Trap (sent by Manager)
Relationship UDP Flood*	Host Group Relationship Alarm	N/A

For a host group relationship, this alarm indicates that an excessive number of short UDP packets has been sent, in either direction, between the two host groups in the relationship. The alarm is raised if the alarm condition exists for longer than a user-specified duration.

You can enable/disable this alarm, or modify the minimum and maximum thresholds, on the Host Group Relationship Properties dialog.

* These alarms are variance-based. You can adjust the tolerance for these alarms:

- For a host – on the Host Policy Manager dialog.
- For a host group relationship – on the Host Group Relationship Properties dialog.

Related Topics

[Alarm List: Overview](#)

Alarm List: Manager System Alarms

These alarms occur when there are issues with the Manager or communication channels. You cannot enable or disable these alarms.

Alarm	Alarm Category	SNMP Trap (sent by Manager)
Analytics does not support more than 1 Data Store domain	Manager System Alarm	ManagerAlarm
Indicates that your system detected more than 1 Data Store domain.		
Analytics Performance has degraded	Manager System Alarm	ManagerAlarm
This alarm can indicate one of three issues: 1. When there is no traffic available to be sent from the Flow Collector to Analytics. In this case there is no problem with your system, and you can ignore the alarm. 2. When you change the role of the primary Manager to that of the secondary Manager, and more than 5 hours has passed before the original primary Manager has been recovered and re-assigned to the primary role, this alarm is triggered. Analytics will recover and run the jobs that occurred during the last 6 hours, while the original primary Manager was down. Job performance will continue to lag until your system has processed all jobs from the last 6 hours and begins to process jobs in real time. 3. If your system has been experiencing degradation for more than 5 hours (usually due to insufficient resources such as CPU or memory), jobs will begin to lag. If this lag exceeds 5 hours, then this alarm is triggered. At this point, results will be incomplete and unreliable. A possible cause for this failure is that you have increased the flows per second beyond what is supported in your setup. To resolve this, either reduce the flows per second or increase the resources on the Manager, the Data Store, or both. If you cannot resolve the issue, contact Cisco Support.		
Analytics results are incomplete	Manager System Alarm	ManagerAlarm

Alarm	Alarm Category	SNMP Trap (sent by Manager)
Indicates that Analytics job results are incomplete. Check your system resources. If the problem persists, contact Cisco Support .		
Appliance Certificate Expiration less than 3 days	Manager System Alarm	ManagerAlarm
Indicates that you have fewer than 3 days until one or more of your appliance identity certificates will be expiring. Refer to the SSL/TLS Certificates Guide for Managed Appliances for details about certificates.		
Appliance Certificate Expiration less than 14 days	Manager System Alarm	ManagerAlarm
Indicates that you have fewer than 14 days until one or more of your appliance identity certificates will be expiring. Refer to the SSL/TLS Certificates Guide for Managed Appliances for details about certificates.		
Appliance Certificate Expiration less than 30 days	Manager System Alarm	ManagerAlarm
Indicates that you have fewer than 30 days until one or more of your appliance identity certificates will be expiring. Refer to the SSL/TLS Certificates Guide for Managed Appliances for details about certificates.		
Appliance Certificate Expiration less than 60 days	Manager System Alarm	ManagerAlarm
Indicates that you have fewer than 60 days until one or more of your appliance identity certificates will be expiring. Refer to the SSL/TLS Certificates Guide for Managed Appliances for details about certificates.		
Appliance Certificate Expiration less than 90 days	Manager System Alarm	ManagerAlarm
Indicates that you have fewer than 90 days until one or more of your appliance identity certificates will be expiring. Refer to the SSL/TLS Certificates Guide for Managed Appliances for details about certificates.		

Alarm	Alarm Category	SNMP Trap (sent by Manager)
Appliance Certificate has expired	Manager System Alarm	ManagerAlarm
Indicates that one or more of your appliance identity certificates has expired. Refer to the SSL/TLS Certificates Guide for Managed Appliances for details about certificates.		
Cisco® ISE Management Channel Down	Manager System Alarm	ManagerAlarm
Indicates that the Manager has not received a pxGrid message from Cisco Identity Services Engine (ISE) within the last five minutes. From the main menu, choose Configuration > INTEGRATIONS Cisco® ISE to review connection errors.		
Data Node down	Manager System Alarm	ManagerAlarm
Indicates that one or more of the Data Node's status is Down. Verify the Data Nodes' status in the Central Management Inventory. Log into that Data Node's console and restart the appliance. Contact Cisco Support if this persists.		
Data Node down; remaining Data Node count critical	Manager System Alarm	ManagerAlarm
Indicates that at least one Data Node is down, the Data Node count in the Data Store is now the minimum required to retain fault tolerance, and if any additional Data Nodes go down, the Data Store database automatically shuts down. Check your Central Management inventory to determine which Data Nodes are Down, and bring them back Up as quickly as possible. Contact Cisco Support for assistance.		
Data Node recovering	Manager System Alarm	ManagerAlarm
Indicates that one or more of the Data Nodes restarted after an unexpected shutdown, and is recovering data from other Data Nodes. The data recovery process may take a while to complete. Review that Data Node's log files to determine why it shut down, and verify the power settings for this Data Node to prevent or minimize future unplanned power outages. Contact Cisco Support if this persists.		
Data Store excessive	Manager System Alarm	ManagerAlarm

Alarm	Alarm Category	SNMP Trap (sent by Manager)
timestamp skew		
Indicates that there are several seconds or more difference among the Data Nodes and your Manager. Check that your Data Nodes and Manager are all configured with the same NTP server, and that they can all reach that NTP server. Contact Cisco Support if this persists.		
Data Store reaching limit for ROS container operational files	Manager System Alarm	ManagerAlarm
Indicates that the Data Store is reaching a limit for operational files, after which it cannot update a portion of the database if it reaches the limit. Contact Cisco Support for assistance.		
Data Store recovery error	Manager System Alarm	ManagerAlarm
Indicates that a Data Node unexpectedly went down, or the Data Store database shut down and restarted, but the Data Store database encountered an error during restart. Contact Cisco Support for assistance.		
Data Store recovery failure	Manager System Alarm	ManagerAlarm
Indicates that a Data Node unexpectedly went down, or the Data Store database shut down and restarted, but the Data Store database encountered multiple errors and failed to properly restart and recover data for all Data Nodes. Contact Cisco Support for assistance.		
Data Store recovery lock error	Manager System Alarm	Manager Alarm
Indicates that a Data Node could not properly recover data after it went down and restarted, due to a lock on certain data preventing the recovery process from copying that data. Restart the Data Node appliance during a time when the Data Store database is not ingesting or not busy. Contact Cisco Support for assistance.		
Data Store refresh failure	Manager System Alarm	ManagerAlarm

Alarm	Alarm Category	SNMP Trap (sent by Manager)
Indicates that a Data Node attempted to refresh data in its database, but encountered an error. Restart the Data Node appliance. Contact Cisco Support for assistance if the issue persists.		
Data Store shut down due to too many Data Nodes down	Manager System Alarm	ManagerAlarm
Indicates that either the number of Data Nodes with a status of Down is more than half of your total Data Nodes, or two adjacent configured Data Nodes are down, meaning that the Data Store database shuts down due to loss of fault tolerance. Check your Central Management inventory to determine which Data Nodes are Down, and bring them back Up as quickly as possible. Contact Cisco Support for more information on restarting the Data Store database.		
Flow Collector Management Channel Down	Manager System Alarm	ManagerAlarm
When communication is reestablished, changes made in the interim are synchronized between the Manager and Flow Collectors, and this alarm becomes inactive. An admin user can manually force a synchronization. Indicates that communication is lost between a Flow Collector and the Manager. The Details column in the Alarm Table provides one of the previous reasons for the alarm: • No route to host • Connection error • Socket error • Http protocol error • Http authorization error • Http client error • Http server error • Http other error • Response handler error (something wrong with the response handler running on the device, such as in the Agent)		

Alarm	Alarm Category	SNMP Trap (sent by Manager)
- Other error (in this case, a detailed java exception description is also included) - The certificate for the Flow Collector is not trusted, has expired, or has been revoked		
Flow Collector Time Mismatch	Manager System Alarm	ManagerAlarm
Indicates that there is more than 30 seconds difference between the Flow Collector time and the time that the Manager expects for the Flow Collector. If this alarm occurs, the Network Time Protocol (NTP) server may not be enabled on the device that has drifted (i.e., on either the Manager or the Flow Collector).		
SLIC Channel Down	Manager System Alarm	ManagerAlarm
Indicates that the Manager server is not retrieving data from the Threat Feed. This alarm clears when either of the following two conditions is true: - The Manager server begins retrieving data from the Threat Feed again. - You disable the Threat Feed.		
SMC database ingest and maintenance connection with Data Store lost	Manager System Alarm	ManagerAlarm
Indicates that the Manager could not properly establish a connection with the Data Store for data ingestion and maintenance. Verify that you properly configured the <code>dbadmin</code> account password and certificate for secure Data Store communications in SystemConfig, and that your Data Nodes are all Up in the Central Management inventory. Contact Cisco Support for assistance.		
SMC Disk Space Low	Manager System Alarm	ManagerAlarm
Indicates that the free disk space percentage on the Manager is at 20% or less. This alarm becomes inactive when the free disk space rises above 20%.		

Alarm	Alarm Category	SNMP Trap (sent by Manager)
If the free disk space percentage on the Manager drops below 15%, the Manager begins automatically freeing up disk space by deleting and condensing data beyond your retention settings. An alarm note is appended to the alarm to list these changes to the data retention settings. You can configure these settings on the Data Retention page in the Manager. You can back up/restore the Manager database as described in the Manager Appliance Administration Online Help.		
SMC Duplicate Primary	Manager System Alarm	ManagerAlarm
Indicates that an Manager that is in a failover pair and is in the primary role detects that the other Manager in the pair is also in the primary role. The Manager that raises this alarm will immediately change its role to secondary. This alarm, will therefore go inactive immediately.		
SMC Failover Channel Down	Manager System Alarm	ManagerAlarm
Indicates that communication is lost between the primary Manager and the secondary Manager.		
SMC query connection with Data Store lost	Manager System Alarm	ManagerAlarm
Indicates that the Manager could not properly establish a connection with the Data Store to query the database. Verify that you properly configured the <code>readonlyuser</code> account password and certificate for secure Data Store communications in SystemConfig, and that your Data Nodes are all Up in the Central Management inventory. Contact Cisco Support for assistance.		
SMC RAID Failure	Manager System Alarm	ManagerAlarm
Indicates that a failure has been detected in the RAID configuration of the Manager within the last hour.		
SMC RAID Rebuilding	Manager System Alarm	ManagerAlarm
Indicates that a RAID drive has started the process of rebuilding after a RAID failure		

Alarm	Alarm Category	SNMP Trap (sent by Manager)
occurred on the Manager.		

Related Topics

[Alarm List: Overview](#)

Alarm List: UDP Director Alarms

Alarm	Alarm Category	SNMP Trap (sent by Manager)
UDP Director Degradation	UDP Director System Alarm	ManagerAlarm
Indicates the UDP Director is experiencing degradation for one of the following reasons: - Overloaded: triggered when the UDP Director is continuously using 90%+ of CPU over a 5 min period - Oversubscribed: triggered when the UDP Director's input rates over a 5 minute period exceed published numbers - Remedy: Adjust inputs to the recommended numbers - Engine error: triggered when the UDP Director experiences any other issues not specifically addressed by one of the alarms - Remedy: Restart immediately - Low memory: triggered when the UDP Director is using 90% or more of its total memory - Remedy: Update memory - Packet loss: triggered when the appliance is low on resources and is unable to process all packets it is receiving - Remedy: Update resources		
UDP Director HA Service Down	UDP Director System Alarm	ManagerAlarm
Indicates the High Availability service between a primary and secondary UDP Director that are managed by the Manager is currently down. The alarm clears when either of the following two conditions is true: - The HA service becomes active. - The HA service is turned off on both the primary and secondary UDP Directors.		
UDP Director Management Channel Down	UDP Director System Alarm	ManagerAlarm

Alarm	Alarm Category	SNMP Trap (sent by Manager)
Indicates the management channel between the UDP Director and the Manager is currently down. This alarms clears when either of the following two conditions is true: - The management channel becomes active. - The UDP Director is removed as a managed appliance from the Manager.		
UDP Director RAID Failure	UDP Director System Alarm	ManagerAlarm
Indicates that a failure has been detected in the RAID configuration of the UDP Director within the last hour.		
UDP Director RAID Rebuilding	UDP Director System Alarm	ManagerAlarm
Indicates that a RAID drive has started the process of rebuilding after a RAID failure occurred on a UDP Director.		
UDP Director Stopped	UDP Director System Alarm	ManagerAlarm
Indicates when 1 or more UDP Director processes has stopped working. Remedy: Restart immediately		

Related Topics

[Alarm List: Overview](#)

Alarm Suppression: Overview

You can configure rules, based on alarm attributes, for known communication between devices which is expected and known to never be malicious. When the communication matches the rule criteria (such as ports, protocols, IP addresses), Analytics suppresses any alarms that would normally be generated, resulting in a less noisy and more efficacious system.

Each time Analytics processes a configuration file, it overwrites the previous file. No Flow Collector syncing occurs since this action is limited to the Manager (formerly Stealthwatch Management Console).

You can suppress [23 of the available alarms](#) for any of your domains.



- Analytics supports alarm suppression for Cisco Secure Network Analytics (formerly Stealthwatch) version 7.4.0. and later.
- Syslogs are not sent out for Response Management.
- A log of suppressions is preserved but is not recorded in the audit log.
- When you perform a configuration backup, the alarm suppression file should be included (but only if you have successfully previously processed this file). However, we recommend that you create a configuration backup of this file on the Manager.

Alarm Suppression: Table of Contents

Alarm Suppression: Alarms You Can Suppress

You can suppress the following alarms:

Alarm Type (Alarm ID)	Alarm Attributes you can Suppress
Addr (Address) Scan/tcp (276) The source host is attempting to contact multiple hosts (using TCP) within a natural class C network (/24) on the same port and most connection attempts are either being rejected (TCP Reset) or the target hosts are not responding at all. This is used to trigger the Worm Activity and Worm Propagation alarms. These are commonly seen during network scanning or enumeration.	• source_ip • source_groups • protocol • destination_ip_range • destination_port
Addr (Address) Scan/udp (286) The source host is attempting to contact multiple hosts (using UDP) within a natural class C network (/24) on the same port and most connection attempts are either being rejected (ICMP port unreachable) or the hosts are not responding at all. These are commonly seen during network scanning or enumeration. This is used to trigger the Worm Activity and Worm Propagation alarms.	• source_ip • source_groups • protocol • destination_ip_range • destination_port
Beaconing Host (39) An IP communication between an inside and outside host (with traffic in only one direction) exceeds the "Seconds required to qualify a flow as	• source_ip • source_groups • source_port • destination_ip

Alarm Type (Alarm ID)	Alarm Attributes you can Suppress
long duration" setting.	• destination_groups • destination_port • protocol
Bot Infected Host - Attempted C&C Activity (41) The source host is attempting to contact a command and control (C&C) server using a port identified in the C&C Server list. The communication is one-way only, indicating the C&C server has not responded. The inside host, as the initiator, accumulates Concern Index (CI) points. If the C&C server it attempts to contact is also an inside host, then that C&C server accumulates Target Index (TI) points.	• source_ip • source_groups • source_port • destination_ip • destination_groups • destination_port • protocol
Bot Infected Host - Successful C&C Activity (42) The source host has successfully contacted a C&C server using a port identified in the C&C Server list. The communication is two-way, indicating the C&C server has responded. The inside host, as the initiator, accumulates Concern Index (CI) points. If the C&C server it contacts is also an inside host, then that C&C server accumulates Target Index (TI) points.	• source_ip • source_groups • source_port • destination_ip • destination_groups • destination_port • protocol
Brute Force Login (58) The host detects a series of short TCP connections consistent with an	• source_ip • source_groups • source_port

Alarm Type (Alarm ID)	Alarm Attributes you can Suppress
attempt at brute force password cracking through repeated logins.	• destination_ip • destination_groups • destination_port • protocol
High SMB Peers (60) This security event indicates that host has many Server Message Block (SMB) sessions to the outside, which is consistent with worm propagation.	• source_ip • source_groups • source_port • protocol • flows_count
High Traffic (30) The host traffic rate averaged over a 5-minute period has exceeded the limit of acceptable traffic values.	• source_ip • source_groups • bytes_per_second
ICMP Flood (7) The source host has sent an excessive number of ICMP packets in a 5-minute period.	• source_ip • source_groups • packets
Packet Flood (8) The source host has sent an excessive number of short packets to the target host. This security event is seen as a result of brute force attacks, DoS attacks, and malfunctioning network applications.	• source_ip • source_groups • source_port • destination_ip • destination_groups • destination_port • protocol
Ping Oversized Packet (278) The source host has sent an ICMP	• source_ip • source_groups

Alarm Type (Alarm ID)	Alarm Attributes you can Suppress
echo request or reply that has more than 90 data bytes. These events may be harmless network health checks or may contain a covert data channel.	• destination_ip • destination_groups • protocol
Ping Scan (277) The source host is sending Echo Request packets to many hosts with a natural class C network (/24) range of addresses. This is often done to identify the active hosts on a network.	• source_ip • source_groups • protocol • destination_ip_range
Port Scan (55) The source IP has attempted to connect to an excessive number of ports on the target IP.	• source_ip • source_groups • source_port • destination_ip • destination_groups • destination_port • protocol
Scanner Talking (63) This security event indicates that a host that has been scanning your network now has a two-way conversation with one of the target hosts that it scanned. It is enabled by default in the Inside Hosts policy and the Outside Hosts policy. It is, however, disabled by default in the Network Management Scanners role policy.	• source_ip • source_groups • source_ports • destination_ip • destination_group • destination_port • protocol
SSH Reverse Shell (61) Detects an SSH session that appears	• source_ip • source_groups

Alarm Type (Alarm ID)	Alarm Attributes you can Suppress
to be a reverse shell. More data is being sent to the outside host than is being received.	• source_port • destination_ip • destination_groups • destination_port • protocol
Stealth Scan/tcp (272) The source host has used the same source port to connect to different ports on the target host at the same time. This behavior indicates applications that have used raw sockets to create TCP or UDP connections. The security event shows the last target port accessed before the security event was recognized.	• source_ip • source_groups • source_port • destination_ip • destination_groups • destination_port • protocol
Stealth Scan/udp (271) A crafted UDP scan with re-use of port numbers by the source host. This often indicates "packet crafting" is being done by the scanning host which may be looking for hosts to attack.	• source_ip • source_groups • source_port • destination_ip • destination_groups • destination_port • protocol
Suspect Data Hoarding (315) The source host has downloaded an unusual amount of data from one or more hosts.	• source IP Address • source Group List • bytes
Suspect Data Loss (40)	• source_ip

Alarm Type (Alarm ID)	Alarm Attributes you can Suppress
Indicates that an inside host has uploaded an abnormal amount of data to outside hosts.	• source_groups • bytes
Suspect UDP Activity (24) The source host has been identified scanning multiple hosts on a UDP port and has successfully sent a large UDP packet to another previously scanned host. This type of behavior is consistent with many single-packet UDP-based worms such as "SQL Slammer" and "Witty." Investigate this security event immediately.	• source_ip • source_groups • source_port • destination_ip • destination_groups • destination_port • protocol
SYN Flood (5) The source host has sent an excessive number of TCP connection requests (SYN packets) in a 5-minute period. This may indicate a DoS attack or non-stealthy scanning activity.	• source_ip • source_groups • packets
UDP Flood (6) The source host has sent an excessive number of UDP packets in a 5-minute period. This may indicate a DoS attack or non-stealthy scanning activity.	• source_ip • source_groups • packets
Worm Propagation (36) The host has scanned and connected on a particular port across more than one subnet, and the host was previously scanned and connected to by a host for which the Worm Activity alarm has been raised.	• source_ip • source_groups • source_port • destination_ip • destination_groups

Alarm Type (Alarm ID)	Alarm Attributes you can Suppress
	• destination_port • protocol • initial_infected_host

For more information about these alarms, see [Security Event List](#).

Related Tasks

[Alarm Suppression: Rule Configuration](#)

Related Topics

[Alarm Suppression: Overview](#)

Alarm Suppression: Rule Configuration

You must use APIs to configure alarm suppression rules. These APIs contain the following endpoints:

- **Get** Displays the current suppression configuration.
- **Put** Uploads and configures the rules to suppress alarms.
- **Delete** Deletes all suppression configuration.

 Both Admin and Primary Admin users can configure alarm suppression rules.

Example: Following is an example of an alarm type rule. Its definition is "For an Address Scan TCP alarm with a device ID corresponding to 400 or 401, AND a domain ID other than 300 or 301, AND source groups corresponding to 22,23, or 24, AND a source IP corresponding to 1.1.1.1 or 1.1.1.2."

```
{
  "rules": [
    {
      "type": 276,
      "filters": [
        { "field": "device_id", "operator": "In", "value": [400, 401] }
      ],
      { "field": "domain_id", "operator": "NotIn", "value": [300, 301] }
    },
    { "field": "source_groups", "operator": "In", "value": [22,23,24] }
  ],
  { "field": "source_ip", "operator": "In", "value": ["1.1.1.1", "1.1.1.2"] }
]
}
```

Rule Properties

Refer to the following guidelines when configuring rules:

- If an alarm type does have more than one rule, then either one of those rules (OR operation across rules) may suppress the alarm.
- The raised alarm should match all the attributes defined in the rule of a specific alarm type that you want to be suppressed.

- If you use the operator "IN" within a value, then any of the values within that expression can be true (OR operations across values within an expression).
- The same alarm type can have more than one rule.
- You can configure a maximum of 500 rules irrespective of alarm type.

Following are two rules for Address Scan TCP. The JSON for the full set of rules is shown in these two rules. Its definition is as follows:

" (For an Address Scan TCP alarm in a domain with ID 123 AND with a device ID of 12)

OR

(For an Address Scan TCP alarm in a domain with ID 234 AND with a device ID of 432)."

```
{
  "rules": [
    {
      "type": 276,
      "filters": [
        { "field": "domain_id", "operator": "Equal", "value": 123 }
      ],
      { "field": "device_id", "operator": "Equal", "value": 12 }
    },
    {
      "type": 276,
      "filters": [
        { "field": "domain_id", "operator": "Equal", "value": 234 }
      ],
      { "field": "device_id", "operator": "Equal", "value": 432 }
    }
  ]
}
```

When a host performs an address scan (TCP/UDP) on the same subnet with multiple ports, even though many security events are generated, only one condition is generated and one alarm is logged in host_alarm database on the Manager.

To suppress alarms of this kind, use the following rule:

Option 1: Give one port per expression

```
{
  "rules": [ {
    "type": 276,
    "filters": [
      { "field": "destination_port", "operator": "Equal", "value": 2049 }
    ]
  }, {
    "type": 276,
    "filters": [
      { "field": "destination_port", "operator": "Equal", "value": 2055 }
    ]
  },
  ,
```

Option 2: Give additional criteria (be more granular). See example below for available attributes

```
{
  "type": 276,
  "filters": [
    { "field": "device_id", "operator": "Equal", "value": 301 }
  ],
  ,
  { "field": "domain_id", "operator": "Equal", "value": 301 }
  ,
  { "field": "source_groups", "operator": "In", "value": [65534] }
  ,
  { "field": "source_ip", "operator": "In", "value": ["10.1.0.9", "10.1.0.8", "10.1.0.7", "10.1.0.6", "10.1.0.5"] }
  ,
  { "field": "destination_ip_range", "operator": "In", "value": ["10.10.0.0"] }
  ,
  { "field": "protocol", "operator": "In", "value": ["tcp"] }
  ,
  { "field": "destination_port", "operator": "In", "value": [50] }
  ]
}
```

Related Tasks

Alarm Suppression: Rule Configuration

Related Topics

Alarm Suppression: Overview

Alarm Suppression: Attribute Details

The attributes you can suppress on an alarm fall into one of the following attribute types. You can use any of these operators that match a particular attribute type when building a rule. The following attributes may or may not be common across alarms.



Entries are not case sensitive. As an example, for the protocol attribute you can use TCP, tcp, TcP, etc. These are all valid entries.

Attribute	Type	Operators	Description
bytes	Long	Equal, NotEqual, GreaterThan, GreaterThanEqual, LessThan, LessThanEqual	Positive whole number between 0 and 4294967296.
bytes_per_second	Long	Equal, NotEqual, GreaterThan, GreaterThanEqual, LessThan, LessThanEqual	Positive whole number between 0 and 4294967296.
device_id (This is the Flow Collector ID.)	integer or Set<Integer>	If Integer: Equal, NotEqual, GreaterThan, GreaterThanEqual, LessThan, LessThanEqual	Positive whole number.
		If Set<Integer>: In, NotIn	Set of whole numbers enclosed within [] up to 100 values.
destination_ip	Set<IPAddress>	In, NotIn	Set of IP addresses or subnet address range enclosed within [] up to 100

Attribute	Type	Operators	Description
			values.
destination_ ip_range	Set<IPAddress>	In, NotIn	Set of subnet addresses enclosed within [] up to 100 values.
destination_ groups	Set<Integer>	In, NotIn	Set of whole numbers representing host group ID enclosed within [].
destination_ port	Integer or Set<Integer>	If Integer: Equal, NotEqual, GreaterThan, GreaterThanEqual, LessThan, LessThanEqual	Positive whole number.
		If Set<Integer>: In, NotIn	Set of whole numbers enclosed within [] up to 100 values.
domain_id	integer or Set<Integer>	If Integer: Equal, NotEqual, GreaterThan, GreaterThanEqual, LessThan, LessThanEqual	Positive whole number.
		If Set<Integer>: In, NotIn	Set of whole numbers enclosed within [] up to 100 values.

Attribute	Type	Operators	Description
flows_count	Long	Equal, NotEqual, GreaterThan, GreaterThanEqual, LessThan, LessThanEqual	Positive whole number between 0 and 4294967296.
initial_infected_host	Set<IPAddress>	In, NotIn	Set of IP addresses enclosed within [] up to 100 values.
packets	Long	Equal, NotEqual, GreaterThan, GreaterThanEqual, LessThan, LessThanEqual	Positive whole number between 0 and 4294967296.
protocol	Protocol or Set<Protocol>	If Protocol: Equal, NotEqual	Single valid protocol name enclosed in double quotes Example: ["TCP"].
		If Set<Protocol> In, NotIn	Set of valid protocol names enclosed within [] up to 100 values Example: ["TCP", "UDP"].
source_ip	Set<IPAddress>	In, NotIn	Set of IP addresses or subnet address range enclosed within [] up to 100 values.
source_groups	Set<Integer>	In, NotIn	Set of whole numbers

Attribute	Type	Operators	Description
			representing host group ID enclosed within [] up to 100 values.
source_port	Integer or Set<Integer>	If Integer: Equal, NotEqual, GreaterThan, GreaterThanEqual, LessThan, LessThanEqual	Positive whole number.
		If Set<Integer>: In, NotIn	Set of whole numbers enclosed within [] up to 100 values.

Related Tasks

[Alarm Suppression: Rule Configuration](#)

Related Topics

[Alarm Suppression: Overview](#)

Alarm Suppression: API Examples

To read an overview of our APIs that contain endpoints (along with their request and response schema), do one of the following:

Option 1

1. Enter this URL: **`https://[manager_ip_address]/api-docs`**
2. Scroll to the link entitled "Alarm Suppression API" in the APIs section and click this link.

The corresponding page opens.

Option 2

Enter this URL: **`https://[manager_ip_address]/legacy-detections/v1/docs`**

The corresponding API documentation is displayed.

Related Tasks

[Alarm Suppression: Rule Configuration](#)

Related Topics

[Alarm Suppression: Overview](#)

Security Event List

A security event is an algorithm that looks for specific behavior and can alert on that behavior on your network, depending on settings applied in policies. It does this by directly generating a host alarm (if set to do so), or by assigning index points to alarm categories, which can in turn trigger a host alarm. (A host alarm is a security event that has triggered an alarm.)

Security events contribute index points for specific alarm categories. If you disable a security event, it will not accumulate index points against the alarm categories associated with it. Both alarm categories and security events can trigger host alarms, depending on settings applied in policies.

Security events can be disabled for specific services and disabled at the host group level.




- We are currently in the process of updating and expanding security event information. The new format includes information (presented in several tables) such as event description, what it means when an event is triggered, the steps you should take to investigate further, and much more. During this transition, some security events in this topic will contain more completed categories or tables than others.
- The Flow Collector (NetFlow) supports the security events marked with a double asterisk (**) only when it is used in conjunction with a Flow Sensor.

Cisco Secure Network Analytics (formerly Stealthwatch) contains the following default security events.

Addr (Address) Scan/tcp

Mitigation is not available for the alarm associated with this security event.



What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	If Secure Network Analytics detects activity that potentially indicates someone doing a network scan using TCP, the system records this as a scanning event. If enough scanning events affect a large enough number of hosts, the Manager (formerly Stealthwatch Management Console) raises a security event. Examples of "scanning events" include various combinations of bad TCP flags, or sending a SYN but not responding to the following SYN-ACK, as well as other indicators.
What does it mean when it triggers?	A host is attempting to discover which hosts are running particular services to potentially make use of.
What are next steps?	Determine what the host was scanning. Start this investigation with a broad net and then narrow it down. Begin by running a Top Ports (outbound) report. Set the time period to the day of the event and the client host to be the source IP of the event. Regardless of the target IP range that is listed, you can scan any type of

Questions about this event	Response
	host, so determine if you want to search on Inside Hosts, Outside Hosts, or both. Then, in the Filter dialog, set the Server filter as appropriate on the Hosts tab and set the "Order the records returned by" to <i>Flows</i> on the Advanced tab. Once the results come back, it may be helpful to sort by Peers. Regardless, start at the top of the list as sorted by Flows or Peers to find the standout ports that either don't belong or appear abnormally high. Right-click an IP address to pivot to the flows to view the various IP addresses and determine if particular host groups are those being primarily targeted. You can also right-click and pivot to the Top Peers report to determine whether or not the traffic is spread generally equally across the targets. It is also worth noting the percentage of hosts that responded to the scans. At this point you should be able to determine who the host was scanning and what ports were being scanned.
Non-standard flow data required? (FlowSensor, proxy, firewall, etc.)	Some scans are only recorded based on flags sent by the FlowSensor, and some by flags sent by firewalls. (These are noted in the event details). Scanning events that are not noted as such do not require any specific data.
Notes	N/A

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	Event must be enabled on the source host
Which policies have this event	Inside Hosts, Outside Hosts, Client IP Policy

Questions about event	Response
on by default?	
Which policies have this event off by default?	Antivirus & SMS Servers; Firewalls, Proxies, & NAT Devices; Network Management & Scanners
Which policies have this alarm on by default?	None
Which policies have this alarm off by default?	Inside Hosts, Outside Hosts, Client IP Policy; Antivirus & SMS Servers; Firewalls, Proxies, & NAT Devices; Network Management & Scanners
Can you tune this event?	No
Is the event variance-based, threshold-based, or other?	N/A
What are default values and units?	
Tolerance	N/A
Minimum threshold value	N/A
Maximum threshold value	N/A
Is event tunable with non-variance-based parameters?	N/A
If so, what are its tunable attributes and units?	N/A
Is event tunable outside of the normal policy editor	Yes
If so, what are the alternate locations of tunable values?	You can use the <code>disable_stealth_probe lc_threshold.txt</code> value to disable certain cases of stealth scan detection. This is only about 15% of the set of all cases that can

Questions about event	Response
	call out to register a scan event.
Does event have a default mitigation?	No
If so, what is it?	N/A

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The host performing the activity that may be an indication of scanning
What is the target?	The host being scanned by the source host.
Which policy causes the event to trigger?	The source host
What alarm categories does this security event contribute to for the source?	High Recon Index, High Concern Index
In what quantity?	• High Recon Index: True • CI: True
What alarm categories does this security event contribute to for the target?	
In what quantity?	

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: View details. Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source IP, the target's natural class C network (/24), the start time (reset hour) of the event's active day until the event's last active time, and TCP. Desktop Client: Flows are filtered by time period of last five minutes.


What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_TCP_ADDR_SCAN (276)

Addr (Address) Scan/udp

Mitigation is not available for the alarm associated with this security event.

Addr_Scan/udp



Resulting potential security event entry:

Source Host Groups	Source Host	Target Host Groups	Target Host	Concern Index	Security Events
Singapore	203.125.234.84	Lancpe Corporate	50.56.217.21	<CI value>*	Addr_Scan/udp-80(6)

* <CI value> represents a numerical point value that Secure Network Analytics assigns to an event based on various conditions.

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	If Secure Network Analytics detects activity that potentially indicates someone doing a network scan using UDP, the system records this as a scanning event. If enough scanning events affect a large enough number of hosts, the Manager raises a security event. Examples of "scanning events" include sending UDP packets that are responded to by various type of ICMP rejections or firewall flow- denied messages, among other indicators.
What does it mean when it triggers?	A host attempting to discover which hosts are running particular services to potentially make use of.
What are next steps?	Determine what the host was scanning. Start this investigation with a broad net and then narrow it down. Begin by running a Top Ports (outbound) report. Set the time period to the day of the event and the client host to be the source IP of the event. Regardless of the target IP range that is listed, you can scan any type of host, so determine if you want to search on Inside Hosts, Outside Hosts, or both. Then, in the Filter dialog,

Questions about this event	Response
	set the Server filter as appropriate on the Hosts tab and set the "Order the records returned by" to <i>Flows</i> on the Advanced tab. Once the results come back, it may be helpful to sort by Peers. Regardless, start at the top of the list as sorted by Flows or Peers to find the standout ports that either don't belong or appear abnormally high. Right-click an IP address to pivot to the flows to view the various IP addresses and determine if particular host groups are those being primarily targeted. You can also right-click and pivot to the Top Peers report to determine whether or not the traffic is spread generally equally across the targets. It is also worth noting the percentage of hosts that responded to the scans. At this point you should be able to determine who the host was scanning and what ports were being scanned.
Non-standard flow data required? (FlowSensor, proxy, firewall, etc.)	Some scans are only recorded based on flags sent by the FlowSensor, and some by flags sent by firewalls. (These are noted in the event details). Scanning events that are not noted as such do not require any specific data.
Notes	N/A

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	Event must be enabled on the source host
Which policies have this event on by default?	Inside Hosts, Outside Hosts, Client IP Policy

Questions about event	Response
Which policies have this event off by default?	Antivirus & SMS Servers; DHCP Server; Firewalls, Proxies, & NAT Devices; Network Management & Scanners
Which policies have this alarm on by default?	None
Which policies have this alarm off by default?	Inside Hosts, Outside Hosts, Client IP Policy; Antivirus & SMS Servers; DHCP Server; Firewalls, Proxies, & NAT Devices; Network Management & Scanners
Can you tune this event?	No
Is the event variance-based, threshold-based, or other?	N/A
What are default values and units?	N/A
Tolerance	N/A
Minimum threshold value	N/A
Maximum threshold value	N/A
Is event tunable with non-variance-based parameters?	N/A
If so, what are its tunable attributes and units?	N/A
Is event tunable outside of the normal policy editor	Yes
If so, what are the alternate locations of tunable values?	You can use the <code>disable_stealth_probe lc_threshold.txt</code> value to disable certain cases of stealth scan detection. This is only about 15% of the set of all cases that can

Questions about event	Response
	call out to register a scan event.
Does event have a default mitigation?	N/A
If so, what is it?	N/A

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The host performing the activity that may be an indication of scanning
What is the target?	The host being scanned by the source host.
Which policy causes the event to trigger?	The source host
What alarm categories does this security event contribute to for the source?	High Recon Index, High Concern Index
In what quantity?	• High Recon Index: True • CI: True
What alarm categories does this security event contribute to for the target?	
In what quantity?	

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: View details. Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source IP, the target's natural class C network (/24), the start time (reset hour) of the event's active day until the event's last active time, and UDP. Desktop Client: Flows are filtered by time period of last five minutes.


What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_UDP_ADDR_SCAN (286)

Bad Flag ACK (Acknowledge) **

Mitigation is not available for the alarm associated with this security event.

Bad_Flag_ACK



Resulting potential security event entry:

Source Host Groups	Source Host	Target Host Groups	Target Host	Concern Index	Security Events
Singapore	203.125.234.84	Lancope Corporate	50.56.217.21	<CI value>*	Bad_Flag_ACK-445(1)

* <CI value> represents a numerical point value that Secure Network Analytics assigns to an event based on various conditions.

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	If a packet is seen not containing the TCP acknowledge flag and containing any flag besides reset or synchronize, the security event is triggered on the originating host.
What does it mean when it triggers?	A TCP packet with invalid flags set should not normally happen. When it does happen, it has likely been done intentionally in an attempt to gather information about the machine to which the packet is sent since different system configurations may respond to different anomalous flag combinations differently.
What are next steps?	If the source of this event is an inside host, it is worth looking for other signs of recon activity. For example, is the host sending bad flag combinations to multiple hosts? Is it scanning many hosts on the same port? Is it scanning many ports on one host? An inside host performing reconnaissance can be a good sign of unwanted activity or compromise.
Non-standard flow data required?	For Flow Collector NetFlow edition, a FlowSensor is required. For Flow Collector sFlow, no special extras are required.

Questions about this event	Response
(FlowSensor, proxy, firewall, etc.)	
Notes	N/A

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	Enable Security event Bad_Flag_ACK on the originating host.
Which policies have this event on by default?	Inside Hosts, Outside Hosts, Client IP Policy
Which policies have this event off by default?	Antivirus & SMS Servers; Firewalls, Proxies, & NAT Devices; Network Management & Scanners
Which policies have this alarm on by default?	None
Which policies have this alarm off by default?	Inside Hosts; Outside Hosts; Client IP Policy; Antivirus & SMS Servers; Firewalls, Proxies, & NAT Devices; Network Management & Scanners
Can you tune this event?	No
Is the event variance-based, threshold-based, or other?	N/A
What are default values and units?	
Tolerance	N/A
Minimum threshold value	N/A

Questions about event	Response
Maximum threshold value	N/A
Is event tunable with non-variance-based parameters?	N/A
If so, what are its tunable attributes and units?	N/A
Is event tunable outside of the normal policy editor	N/A
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	N/A
If so, what is it?	N/A

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The host that sent the bad packet
What is the target?	The host listed as the bad packet's destination
Which policy causes the event to trigger?	The source's host policy
What alarm categories does this security event contribute to for the source?	High Recon Index, High Concern Index
In what quantity?	• High Recon Index: True • CI: True

Questions about event	Response
What alarm categories does this security event contribute to for the target?	High Target Index
In what quantity?	• TI: True

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: View details. Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source & target IP, a time range five minutes before the event's first active time until the event's last active time, the associated port as <port>/TCP and <port>/UDP, and TCP. Desktop Client: Flows are filtered by time period of last five minutes.


What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_BAD_FLAG_NO_ACK (267)

Bad Flag All **

Mitigation is not available for the alarm associated with this security event.

Bad_Flag_All



Resulting potential security event entry:

Source Host Groups	Source Host	Target Host Groups	Target Host	Concern Index	Security Events
Singapore	203.125.234.84	Lancop Corporate	50.56.217.21	<CI value>*	Bad_Flag_All-445(1)

* <CI value> represents a numerical point value that Secure Network Analytics assigns to an event based on various conditions.

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	If a packet is observed containing all of the TCP flags (synchronize, acknowledge, reset, push, urgent, and finish), the security event is triggered on the originating host.
What does it mean when it triggers?	A TCP packet with all flags set should not normally happen. When it does happen, it has likely been done intentionally in an attempt to gather information about the machine to which the packet is sent since different system configurations may respond to different anomalous flag combinations differently.
What are next steps?	If the source of this event is an inside host, it is worth looking for other signs of recon activity. For example, is the host sending bad flag combinations to multiple hosts? Is it scanning many hosts on the same port? Is it scanning many ports on one host? An inside host performing reconnaissance can be a good sign of unwanted activity or compromise.
Non-standard flow data required?	For Flow Collector NetFlow edition, a FlowSensor is required. For Flow Collector sFlow, no special extras are required.

Questions about this event	Response
(FlowSensor, proxy, firewall, etc.)	
Notes	N/A

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	Enable Security event Bad_Flag_All on the originating host.
Which policies have this event on by default?	Inside Hosts, Outside Hosts
Which policies have this event off by default?	N/A
Which policies have this alarm on by default?	Inside Hosts, Outside Hosts
Which policies have this alarm off by default?	N/A
Can you tune this event?	No
Is the event variance-based, threshold-based, or other?	N/A
What are default values and units?	
Tolerance	N/A
Minimum threshold value	N/A

Questions about event	Response
Maximum threshold value	N/A
Is event tunable with non-variance-based parameters?	N/A
If so, what are its tunable attributes and units?	N/A
Is event tunable outside of the normal policy editor	N/A
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	N/A
If so, what is it?	N/A

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The host that sent the bad packet
What is the target?	The host listed as the bad packet's destination
Which policy causes the event to trigger?	The source's host policy
What alarm categories does this security event contribute to for the source?	High Recon Index, High Concern Index
In what quantity?	• High Recon Index: True • CI: True

Questions about event	Response
What alarm categories does this security event contribute to for the target?	High Target Index
In what quantity?	• TI: True

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: View details. Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source & target IP, a time range five minutes before the event's first active time until the event's last active time, the associated port as <port>/TCP and <port>/UDP, and TCP. Desktop Client: Flows are filtered by time period of last five minutes.


What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_BAD_FLAG_XMAS (263)

Bad Flag NoFlg **

Mitigation is not available for the alarm associated with this security event.

Bad_Flag_NoFlg



Resulting potential security event entry:

Source Host Groups	Source Host	Target Host Groups	Target Host	Concern Index	Security Events
Singapore	203.125.234.84	Lancope Corporate	50.56.217.21	<CI value>*	Bad_Flag_NoFlg-445(1)

* <CI value> represents a numerical point value that Secure Network Analytics assigns to an event based on various conditions.

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	If a packet is seen containing no TCP flags, the security event is triggered on the originating host.
What does it mean when it triggers?	A TCP packet with invalid flags set should not normally happen. When it does happen, it has likely been done intentionally in an attempt to gather information about the machine to which the packet is sent since different system configurations may respond to different anomalous flag combinations differently.
What are next steps?	If the source of this event is an inside host, it is worth looking for other signs of recon activity. For example, is the host sending bad flag combinations to multiple hosts? Is it scanning many hosts on the same port? Is it scanning many ports on one host? An inside host performing reconnaissance can be a good sign of unwanted activity or compromise.
Non-standard flow data required? (FlowSensor, proxy, firewall, etc.)	For Flow Collector NetFlow edition, a FlowSensor is required. For Flow Collector sFlow, no special extras are required.

Questions about this event	Response
Notes	N/A

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	Enable Security event Bad_Flag_NoFlg on the originating host.
Which policies have this event on by default?	Inside Hosts, Outside Hosts, Client IP Policy
Which policies have this event off by default?	Antivirus & SMS Servers; Firewalls, Proxies, & NAT Devices; Network Management & Scanners
Which policies have this alarm on by default?	None
Which policies have this alarm off by default?	Inside Hosts; Outside Hosts; Client IP Policy; Antivirus & SMS Servers; Firewalls, Proxies, & NAT Devices; Network Management & Scanners
Can you tune this event?	No
Is the event variance-based, threshold-based, or other?	N/A
What are default values and units?	
Tolerance	N/A
Minimum threshold value	N/A
Maximum threshold value	N/A

Questions about event	Response
Is event tunable with non-variance-based parameters?	N/A
If so, what are its tunable attributes and units?	N/A
Is event tunable outside of the normal policy editor	N/A
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	N/A
If so, what is it?	N/A

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The host that sent the bad packet
What is the target?	The host listed as the bad packet's destination
Which policy causes the event to trigger?	The source's host policy
What alarm categories does this security event contribute to for the source?	High Recon Index, High Concern Index
In what quantity?	• High Recon Index: True • CI: True
What alarm categories does	High Target Index

Questions about event	Response
this security event contribute to for the target?	
In what quantity?	• TI: True

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: View details. Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source & target IP, a time range five minutes before the event's first active time until the event's last active time, the associated port as <port>/TCP and <port>/UDP, and TCP. Desktop Client: Flows are filtered by time period of last five minutes.


What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_BAD_FLAG_NOFLAG (269)

Bad Flag Rsrvd (Reserved)

Mitigation is not available for the alarm associated with this security event.

Bad_Flag_Rsrvd



Resulting potential security event entry:

Source Host Groups ^{▲1}	Source Host [⬇]	Target Host Groups [⬇]	Target Host ^{▲1}	Concern Index ^{▼2}	Security Events [⬇]
Singapore	203.125.234.84	Lancop Corporate	50.56.217.21	<CI value>*	Bad_Flag_Rsrvd-445(1)

* <CI value> represents a numerical point value that Secure Network Analytics assigns to an event based on various conditions.

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	If a packet is seen containing TCP flags that were reserved in the initial TCP standard, the security event is triggered on the originating host.
What does it mean when it triggers?	A TCP packet with invalid flags set should not normally happen. When it does happen, it has likely been done intentionally in an attempt to gather information about the machine to which the packet is sent since different system configurations may respond to different anomalous flag combinations differently.
What are next steps?	If the source of this event is an inside host, it is worth looking for other signs of recon activity. For example, is the host sending bad flag combinations to multiple hosts? Is it scanning many hosts on the same port? Is it scanning many ports on one host? An inside host performing reconnaissance can be a good sign of unwanted activity or compromise.
Non-standard flow data required?	Flow Collector NetFlow edition will never trigger this event. For Flow Collector sFlow, no special extras are required.

Questions about this event	Response
(FlowSensor, proxy, firewall, etc.)	
Notes	N/A

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	Enable Security event Bad_Flag_Rsrvd on the originating host.
Which policies have this event on by default?	Inside Hosts, Outside Hosts, Client IP Policy
Which policies have this event off by default?	Antivirus & SMS Servers; Firewalls, Proxies, & NAT Devices; Network Management & Scanners
Which policies have this alarm on by default?	None
Which policies have this alarm off by default?	Inside Hosts; Outside Hosts; Client IP Policy; Antivirus & SMS Servers; Firewalls, Proxies, & NAT Devices; Network Management & Scanners
Can you tune this event?	No
Is the event variance-based, threshold-based, or other?	N/A
What are default values and units?	
Tolerance	N/A
Minimum threshold value	N/A

Questions about event	Response
Maximum threshold value	N/A
Is event tunable with non-variance-based parameters?	N/A
If so, what are its tunable attributes and units?	N/A
Is event tunable outside of the normal policy editor	N/A
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	N/A
If so, what is it?	N/A

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The host that sent the bad packet
What is the target?	The host listed as the bad packet's destination
Which policy causes the event to trigger?	The source's host policy
What alarm categories does this security event contribute to for the source?	High Recon Index, High Concern Index
In what quantity?	• High Recon Index: True • CI: True

Questions about event	Response
What alarm categories does this security event contribute to for the target?	High Target Index
In what quantity?	• TI: True

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: View details. Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source & target IP, a time range five minutes before the event's first active time until the event's last active time, the associated port as <port>/TCP and <port>/UDP, and TCP. Desktop Client: Flows are filtered by time period of last five minutes.


What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_BAD_FLAG_RESERVED (265)

Bad Flag RST (Reset) **

Mitigation is not available for the alarm associated with this security event.

Bad_Flag_RST



Resulting potential security event entry:

Source Host Groups	Source Host	Target Host Groups	Target Host	Concern Index	Security Events
Singapore	203.125.234.84	Lancop Corporate	50.56.217.21	<CI value>*	Bad_Flag_RST-445(1)

* <CI value> represents a numerical point value that Secure Network Analytics assigns to an event based on various conditions.

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	If a packet is seen containing the TCP reset and any other flags besides push or acknowledge, the security event is triggered on the originating host.
What does it mean when it triggers?	A TCP packet with invalid flags set should not normally happen. When it does happen, it has likely been done intentionally in an attempt to gather information about the machine to which the packet is sent since different system configurations may respond to different anomalous flag combinations differently.
What are next steps?	If the source of this event is an inside host, it is worth looking for other signs of recon activity. For example, is the host sending bad flag combinations to multiple hosts? Is it scanning many hosts on the same port? Is it scanning many ports on one host? An inside host performing reconnaissance can be a good sign of unwanted activity or compromise.
Non-standard flow data required? (FlowSensor, proxy, firewall,	For Flow Collector NetFlow edition, a FlowSensor is required. For Flow Collector sFlow, no special extras are required.

Questions about this event	Response
etc.)	
Notes	N/A

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	Enable Security event Bad_Flag_RST on the originating host.
Which policies have this event on by default?	Inside Hosts, Outside Hosts, Client IP Policy
Which policies have this event off by default?	Antivirus & SMS Servers; Firewalls, Proxies, & NAT Devices; Network Management & Scanners
Which policies have this alarm on by default?	None
Which policies have this alarm off by default?	Inside Hosts; Outside Hosts; Client IP Policy; Antivirus & SMS Servers; Firewalls, Proxies, & NAT Devices; Network Management & Scanners
Can you tune this event?	No
Is the event variance-based, threshold-based, or other?	N/A
What are default values and units?	
Tolerance	N/A
Minimum threshold value	N/A

Questions about event	Response
Maximum threshold value	N/A
Is event tunable with non-variance-based parameters?	N/A
If so, what are its tunable attributes and units?	N/A
Is event tunable outside of the normal policy editor	N/A
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	N/A
If so, what is it?	N/A

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The host that sent the bad packet
What is the target?	The host listed as the bad packet's destination
Which policy causes the event to trigger?	The source's host policy
What alarm categories does this security event contribute to for the source?	High Recon Index, High Concern Index
In what quantity?	• High Recon Index: True • CI: True

Questions about event	Response
What alarm categories does this security event contribute to for the target?	High Target Index
In what quantity?	• TI: True

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: View details. Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source & target IP, a time range five minutes before the event's first active time until the event's last active time, the associated port as <port>/TCP and <port>/UDP, and TCP. Desktop Client: Flows are filtered by time period of last five minutes.


What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_BAD_FLAG_BAD_RST (266)

Bad Flag SYN (Synchronize) FIN (Finish) **

Mitigation is not available for the alarm associated with this security event.

Bad_Flag_SYN_FIN



Resulting potential security event entry:

Source Host Groups	Source Host	Target Host Groups	Target Host	Concern Index	Security Events
Singapore	203.125.234.84	Lancop Corporate	50.56.217.21	<CI value>*	Bad_Flag_SYN_FIN-445(1)

* <CI value> represents a numerical point value that Secure Network Analytics assigns to an event based on various conditions.

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	If a packet is seen containing the TCP synchronize and finish flags, the security event is triggered on the originating host.
What does it mean when it triggers?	A TCP packet with invalid flags set should not normally happen. When it does happen, it has likely been done intentionally in an attempt to gather information about the machine to which the packet is sent since different system configurations may respond to different anomalous flag combinations differently.
What are next steps?	If the source of this event is an inside host, it is worth looking for other signs of recon activity. For example, is the host sending bad flag combinations to multiple hosts? Is it scanning many hosts on the same port? Is it scanning many ports on one host? An inside host performing reconnaissance can be a good sign of unwanted activity or compromise.
Non-standard flow data required? (FlowSensor, proxy, firewall,	For Flow Collector NetFlow edition, a FlowSensor is required. For Flow Collector sFlow, no special extras are required.

Questions about this event	Response
etc.)	
Notes	N/A

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	Enable Security event Bad_Flag_SYN_FIN on the originating host.
Which policies have this event on by default?	Inside Hosts, Outside Hosts, Client IP Policy
Which policies have this event off by default?	Antivirus & SMS Servers; Firewalls, Proxies, & NAT Devices; Network Management & Scanners
Which policies have this alarm on by default?	None
Which policies have this alarm off by default?	Inside Hosts; Outside Hosts; Client IP Policy; Antivirus & SMS Servers; Firewalls, Proxies, & NAT Devices; Network Management & Scanners
Can you tune this event?	No
Is the event variance-based, threshold-based, or other?	N/A
What are default values and units?	
Tolerance	N/A
Minimum threshold value	N/A

Questions about event	Response
Maximum threshold value	N/A
Is event tunable with non-variance-based parameters?	N/A
If so, what are its tunable attributes and units?	N/A
Is event tunable outside of the normal policy editor	N/A
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	N/A
If so, what is it?	N/A

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The host that sent the bad packet
What is the target?	The host listed as the bad packet's destination
Which policy causes the event to trigger?	The source's host policy
What alarm categories does this security event contribute to for the source?	High Recon Index, High Concern Index
In what quantity?	• High Recon Index: True • CI: True

Questions about event	Response
What alarm categories does this security event contribute to for the target?	High Target Index
In what quantity?	• TI: True

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: View details. Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source & target IP, a time range five minutes before the event's first active time until the event's last active time, the associated port as <port>/TCP and <port>/UDP, and TCP. Desktop Client: Flows are filtered by time period of last five minutes.


What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_BAD_FLAG_SYN_FIN (264)

Bad Flag URG (Urgent) **

Mitigation is not available for the alarm associated with this security event.

Bad_Flag_URG



Resulting potential security event entry:

Source Host Groups ^{*1}	Source Host	Target Host Groups	Target Host	Concern Index ^{*2}	Security Events
Singapore	203.125.234.84	Lancop Corporate	50.56.217.21	<CI value>*	Bad_Flag_URG-445(1)

* <CI value> represents a numerical point value that Secure Network Analytics assigns to an event based on various conditions.

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	If a packet is seen containing the TCP urgent flag and any other flags besides acknowledge, the security event is triggered on the originating host.
What does it mean when it triggers?	A TCP packet with invalid flags set should not normally happen. When it does happen, it has likely been done intentionally in an attempt to gather information about the machine to which the packet is sent since different system configurations may respond to different anomalous flag combinations differently.
What are next steps?	If the source of this event is an inside host, it is worth looking for other signs of recon activity. For example, is the host sending bad flag combinations to multiple hosts? Is it scanning many hosts on the same port? Is it scanning many ports on one host? An inside host performing reconnaissance can be a good sign of unwanted activity or compromise.
Non-standard flow data required?	For Flow Collector NetFlow edition, a FlowSensor is required. For Flow Collector sFlow, no special extras are required.

Questions about this event	Response
(FlowSensor, proxy, firewall, etc.)	
Notes	N/A

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	Enable Security event Bad_Flag_URG on the originating host.
Which policies have this event on by default?	Inside Hosts, Outside Hosts, Client IP Policy
Which policies have this event off by default?	Antivirus & SMS Servers; Firewalls, Proxies, & NAT Devices; Network Management & Scanners
Which policies have this alarm on by default?	None
Which policies have this alarm off by default?	Inside Hosts; Outside Hosts; Client IP Policy; Antivirus & SMS Servers; Firewalls, Proxies, & NAT Devices; Network Management & Scanners
Can you tune this event?	No
Is the event variance-based, threshold-based, or other?	N/A
What are default values and units?	
Tolerance	N/A
Minimum threshold value	N/A

Questions about event	Response
Maximum threshold value	N/A
Is event tunable with non-variance-based parameters?	N/A
If so, what are its tunable attributes and units?	N/A
Is event tunable outside of the normal policy editor	N/A
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	N/A
If so, what is it?	N/A

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The host that sent the bad packet
What is the target?	The host listed as the bad packet's destination
Which policy causes the event to trigger?	The source's host policy
What alarm categories does this security event contribute to for the source?	High Recon Index, High Concern Index
In what quantity?	• High Recon Index: True • CI: True

Questions about event	Response
What alarm categories does this security event contribute to for the target?	High Target Index
In what quantity?	• TI: True

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: View details. Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source & target IP, a time range five minutes before the event's first active time until the event's last active time, the associated port as <port>/TCP and <port>/UDP, and TCP. Desktop Client: Flows are filtered by time period of last five minutes.

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_BAD_FLAG_URG (268)

Beaconing Host

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	An IP communication between an Inside Host and Outside Host (with traffic in only one direction) exceeds the "Seconds required to qualify a flow as long duration" setting.
What does it mean when it triggers?	A beaconing host is host looking for updates or commands from another. This traffic can be used for a variety of reasons, such as keep-alive (heartbeat), to obtain new orders from a Command and Control (C&C) server, or to download updates. It's important to realize that this behavior can be caused by malware, but that is not always the case.
What are next steps?	When investigating a Beaconing Host event, the goal is to determine if the Outside Host is indeed a C&C server. It can often be helpful to look this host up both inside of Secure Network Analytics and outside of it. A good first step within Secure Network Analytics is to open a Top Peers (Outbound) report on the target host. This will provide a list of inside peers sorted by the amount of data that have communicated with the Outside Host. This will tell you whether it is a very common peer within your environment or not. A second step would be to run a Flow Traffic report to visualize traffic patterns between the source and the Target. Communication with a C&C server will exhibit a periodic pattern with a fixed amount of outbound traffic.
Non-standard flow data	None

Questions about this event	Response
required? (FlowSensor, proxy, firewall, etc.)	
Notes	This is the same as a unidirectional suspect long flow, but this event will trigger instead of a suspect long flow.

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	Enable security event Beaconing Host on the originating host.
Which policies have this event on by default?	Inside Hosts, Outside Hosts
Which policies have this event off by default?	none
Which policies have this alarm on by default?	none
Which policies have this alarm off by default?	Inside Hosts, Outside Hosts
Can you tune this event?	no
Is the event variance-based, threshold-based, or other?	N/A
What are default values and units?	N/A

Questions about event	Response
Tolerance	N/A
Minimum threshold value	N/A
Maximum threshold value	N/A
Is event tunable with non-variance-based parameters?	N/A
If so, what are its tunable attributes and units?	N/A
Is event tunable outside of the normal policy editor	N/A
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	
If so, what is it?	

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The host that is beaconing
What is the target?	The host receiving the traffic from the beaconing Host
Which policy causes the event to trigger?	The source host policy
What alarm categories does this security event contribute	

Questions about event	Response
to for the source?	
In what quantity?	• High Command And Control Index: True • CI: True
What alarm categories does this security event contribute to for the target?	
In what quantity?	• TI: True

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: Source host W is using X Service Name (Y Protocol) as Peer to target Z. Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source & target IP, the start time (reset hour) of the event's active day until the event's last active time, and client bytes & server bytes (both of which are greater than or equal to 0). Desktop Client: Flows are filtered by time period of last five minutes.

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_BEACONING_HOST (39)

Bot Command & Control Server

What is the security event?

Questions about this event	Response
What does it mean when it triggers?	Indicates that a host in your environment is being used to assist in the compromise of other hosts beyond your environment by acting as a command and control (C&C) server. This host is also very likely compromised itself.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	High Command And Control Index, High Concern Index
In what quantity?	• High Command And Control Index: True • CI: True
What alarm categories does this security event contribute to for the target?	
In what quantity?	

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_COMMAND_AND_CONTROL_HOST (43)

Bot Infected Host - Attempted C&C Activity

What is the security event?

Questions about this event	Response
What does it mean when it triggers?	A host on your network has attempted to talk to a known command and control (C&C) server. Although the host was not successful in doing so, this is still concerning since something caused that host to make the attempt. Note that malware or a malicious redirect can also cause this behavior.
What are next steps?	This security event generally doesn't require too much validation within the product, but you should perform a flow query for flows between the target of the event and any other host to see if other hosts are interacting with the suspected C&C server. Set the time period of the query to the day of the event or longer. Depending on the duration of the communication, you can also use this query to determine when the source host began reaching out to the suspected C&C server. If you have many hosts communicating with the target host or you see historical communication, it is possible that the C&C server has been falsely identified.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	'Full' Policy: High Command And Control Index, High Concern Index 'Partial' Policy: High Command And Control Index, High Concern Index
In what quantity?	'Full' Policy: • High Command And Control Index: True • CI: True 'Partial' Policy: • High Command And Control Index: True • CI: True
What alarm categories does this security event contribute to for the target?	'Full' Policy: High Target Index 'Partial' Policy: High Target Index
In what quantity?	'Full' Policy: • TI: True 'Partial' Policy: • TI: True

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_BOT_INFECTED_HOST (41)

Bot Infected Host - Successful C&C Activity

What is the security event?

Questions about this event	Response
What does it mean when it triggers?	A host on your network has communicated with a known command and control (C&C) server. It is possible that this behavior has been caused by malware or a malicious redirect, or that this communication was an attempt at infection by, for example, an exploit kit. This is almost certainly sign of a compromised host.
What are next steps?	This security event generally doesn't require too much validation within the product, but you should perform a flow query for flows between the target of the event and any other host to see if other hosts are interacting with the suspected C&C server. Set the time period of the query to the day of the event or longer. Depending on the duration of the communication, you can also use this query to determine when the source host began reaching out to the suspected C&C server. If you have many hosts communicating with the target host or you see historical communication, it is possible that the C&C server has been falsely identified.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	'Full' Policy: High Command And Control Index, High Concern Index 'Partial' Policy: High Command And Control Index, High Concern Index
In what quantity?	'Full' Policy: • High Command And Control Index: True • CI: True 'Partial' Policy: • High Command And Control Index: True • CI: True
What alarm categories does this security event contribute to for the target?	'Full' Policy: High Target Index 'Partial' Policy: High Target Index
In what quantity?	'Full' Policy: • TI: True 'Partial' Policy: • TI: True

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_BOT_INFECTED_HOST_CONTROLLED (42)

Brute Force Login

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	The host detects a series of short TCP connections consistent with an attempt at brute force password cracking through repeated logins.
What does it mean when it triggers?	This may be indicative of a host trying to guess login credentials against another in an attempt to gain access to the system. It may also indicate a misconfigured application on the client which is attempting numerous repeat connections to the server and failing authentication.
What are next steps?	The goal is to determine if the connection attempt to the server is legitimate. If the client is an outside host, is the client address part of a known trusted or business partner network that may be expected to initiate these connections? If so, it may be a misconfigured application. You should also run an external lookup against DShield (for the target host) to verify the owner of the target IP in the security event alarm. If the client is an inside host, would this host be expected to attempt these types of connections to the server in question? Running a Top Peers report for the source host may tell you if this host is connecting to other hosts on the network with similar byte counts or high number of flows. You can also run a Top Ports report to find other hosts that the source IP is connecting to over the same port for which the security event alarm triggered (probably SSH).
Non-standard flow data required? (FlowSensor, proxy, firewall, etc.)	none

Questions about this event	Response
Notes	None

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	
Which policies have this event on by default?	Inside Hosts, Outside Hosts
Which policies have this event off by default?	
Which policies have this alarm on by default?	
Which policies have this alarm off by default?	Inside Hosts, Outside Hosts
Can you tune this event?	Yes
Is the event variance-based, threshold-based, or other?	N/A
What are default values and units?	N/A
Tolerance	N/A
Minimum threshold value	1 (number of connections) and 1 (average bytes per connection)
Maximum threshold value	3,000 (number of connections) and 50,000 (average bytes per connection)

Questions about event	Response
Is event tunable with non-variance-based parameters?	Number of connections and average bytes per connection.
If so, what are its tunable attributes and units?	N/A
Is event tunable outside of the normal policy editor	N/A
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	
If so, what is it?	

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The attacking host
What is the target?	The victim host
Which policy causes the event to trigger?	
What alarm categories does this security event contribute to for the source?	
In what quantity?	• Attack Index: True • CI: True
What alarm categories does	

Questions about event	Response
this security event contribute to for the target?	
In what quantity?	• TI: True

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: View details. Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source & target IP, the start time (reset hour) of the event's active day until the event's last active time, and the associated port as <port>/TCP. Desktop Client: Flows are filtered by time period of last five minutes.

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_BRUTE_FORCE_LOGIN (58)

Connection from Bogon Address Attempted

What is the security event?

Questions about this event	Response
What does it mean when it triggers?	A bogon address is an IP address which has not been assigned and should not be used. Unidirectional communication from one is generally not very concerning, however, it should not happen during regular use. In fact, this sort of traffic is often blocked at a network's perimeter. A specific type of behavior to look for when seen in bulk is a denial-of-service (DoS) attack in which an attacker is spoofing source IP addresses.
What are next steps?	There are two things to initially look for when investigating unidirectional traffic from a bogon address: 1) Is the bogon address (or other bogon addresses) attempting to communicate with other hosts in your network? 2) Is the host to which traffic is being sent from a bogon address receiving an abnormally large amount of traffic? To answer the first question, you can run either a Top Hosts query or a flow query against the bogon host group. You can run it specifically against the bogon IP from the security event, but in the event of spoofed traffic it is trivial for an attacker to use many different bogon IPs. Assess the results returned from your query to determine whether or not this is part of a DoS attack or is potentially a misconfiguration / unclassified host in your network. Generally, a large volume of data or packets is a sign of a DoS attack. To answer the second question, focus your query on the target host of the security event, since the goal is to find out if there is an abnormally large amount of traffic, whether from that bogon address or not. A quick way to get some helpful information is to view the target

Questions about this event	Response
	host's other security events and look for DoS events like SYNs Received or New Flows Served, or look for an abundance of communication from bogons. To investigate a bit more thoroughly, run a Flow Traffic report. Set the Date/Time filter to include the time of the bogon communication as well as two hours before this. A longer time window will result in a longer query but will provide you with better context. Here you can also filter the hosts to include traffic only from outside hosts if you are not concerned about an internally launched DoS attack. A large spike or even a gradual but large volume of traffic may be indicative of a DoS attack.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	High Concern Index
In what quantity?	• CI: True
What alarm categories does this security event contribute to for the target?	Anomaly Index, High Ddos Target Index, High Target Index

Questions about event	Response
In what quantity?	• Anomaly Index: True • High Ddos Target Index: True • TI: True

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_CONN_FROM_BOGON_ATTEMPTED (519)

Connection from Bogon Address Successful

What is the security event?

Questions about this event	Response
What does it mean when it triggers?	A bogon address is an IP address which has not been assigned and should not be used. Bidirectional communication with a bogon address should not occur. This behavior is most likely a sign of a misconfiguration within either your network or your Secure Network Analytics deployment.
What are next steps?	Determine if your environment actually uses a particular bogon IP address range internally or whether or not you use Carrier Grade NAT. If your environment uses a bogon range internally and that is expected, simply add the range to your Inside Hosts group. A way to check this within Secure Network Analytics is to look for active flows in the /24 of the bogon. If a large part of the range is active, it is likely a range being purposefully used within your environment. If the bogon IP is within 100.64.0.0/10, it is using IP space reserved for Carrier Grade NAT and is not an issue if your environment is behind one. If neither of these is the case, the Host Snapshot (accessed within the Desktop Client) will contain information to help you identify the host. For example, the Exporter Interface tab displays which exporters and interfaces are seeing the host's flows to help you determine where the device is hosted. In addition, check out the Security Events tab, Alarms tab, and Identification tabs. These will show you any other behavior with which the bogon host has been involved. Finally, view the security events of the bogon host to see if there are any other hosts it has interacted with.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	High Concern Index
In what quantity?	• CI: True
What alarm categories does this security event contribute to for the target?	Anomaly Index, High Ddos Target Index, High Target Index
In what quantity?	• Anomaly Index: True • High Ddos Target Index: True • TI: True

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_CONN_FROM_BOGON_SUCCEEDED (517)

Connection from Tor Attempted

Detects attempted connections to host(s) inside your network from Tor exit nodes.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	High Policy Violation Index, High Concern Index
In what quantity?	• High Policy Violation Index: True • CI: True
What alarm categories does this security event contribute to for the target?	High Target Index
In what quantity?	• TI: True

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_TOR_EXIT_ATTEMPTED (317)

Connection from Tor Successful

Detects successful connections to host(s) inside your network from Tor exit nodes.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	High Policy Violation Index, High Concern Index
In what quantity?	• High Policy Violation Index: True • CI: True
What alarm categories does this security event contribute to for the target?	High Target Index
In what quantity?	• TI: True

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_TOR_EXIT_SUCCEEDED (318)

Connection to Bogon Address Attempted

What is the security event?

Questions about this event	Response
What does it mean when it triggers?	A bogon address is an IP address which has not been assigned and should not be used. Unidirectional communication to one is a bit peculiar, since something must have prompted a host in your network to do so. It is possible that this is backscatter from a denial-of-service (DoS) attack. This occurs when an attacker sends you packets with randomized source address as part of a distributed denial-of-service (DDoS) and you respond to those randomized addresses. If an attacker is randomizing the entire IP space, some of these will be bogon addresses. Unidirectional communication to a bogon IP may also be a sign of reconnaissance or a misconfigured host.
What are next steps?	Determine if your environment actually uses a particular bogon IP address range internally or whether or not you use Carrier Grade NAT. If your environment uses a bogon range internally and that is expected, simply add the range to your Inside Hosts group. A way to check this within Secure Network Analytics is to look for active flows in the /24 of the bogon. If a large part of the range is active, it is likely a range being purposefully used within your environment. If the bogon IP is within 100.64.0.0/10, it is using IP space reserved for Carrier Grade NAT and is not an issue if your environment is behind one. If neither of these are the case, review the other security events on the source of the event. Look for either an abundance of communication to bogons or denial-of-service (DoS) events like SYNs Received or New Flows Served. Also, look for recon-related events such as Ping Scan or Addr_Scan. To investigate a

Questions about this event	Response
	potential DDoS bit more thoroughly, run a Flow Traffic report. Set the Date/Time filter to include the time of the Bogon communication as well as two hours before this. A longer time window will result in a longer query but will provide you with better context. Here you can also filter the hosts to include traffic only from outside hosts if you are not concerned about an internally launched DoS attack. A large spike or even a gradual but large volume of traffic may be indicative of a DoS attack. If neither a DDoS or a large number of reconnaissance activity seem likely, view the security events of the bogon IP to see if there are any other hosts that have attempted to interact with the bogon host. A large number of hosts attempting to interact with a particular bogon suggest either an application had a bad configuration pushed out or an application hosted on a bogon range disappeared. You can investigate this further by viewing the bogon's historical traffic.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	Anomaly Index, High Concern Index

Questions about event	Response
In what quantity?	- Anomaly Index: True - CI: True
What alarm categories does this security event contribute to for the target?	High Target Index
In what quantity?	TI: True

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_CONN_TO_BOGON_ATTEMPTED (518)

Connection to Bogon Address Successful

What is the security event?

Questions about this event	Response
What does it mean when it triggers?	A bogon address is an IP address which has not been assigned and should not be used. Bidirectional communication with a bogon address should not occur. This behavior is most likely a sign of a misconfiguration within either your network or your Secure Network Analytics deployment.
What are next steps?	Determine if your environment actually uses a particular bogon IP address range internally or whether or not you use Carrier Grade NAT. If your environment uses a bogon range internally and that is expected, simply add the range to your Inside Hosts group. A way to check this within Secure Network Analytics is to look for active flows in the /24 of the bogon. If a large part of the range is active, it is likely a range being purposefully used within your environment. If the bogon IP is within 100.64.0.0/10, it is using IP space reserved for Carrier Grade NAT and is not an issue if your environment is behind one. If neither of these is the case, the Host Snapshot (accessed within the Desktop Client) will contain information to help you identify the host. For example, the Exporter Interface tab displays which exporters and interfaces are seeing the host's flows to help you determine where the device is hosted. In addition, check out the Security Events tab, Alarms tab, and Identification tabs. These will show you any other behavior with which the bogon host has been involved. Finally, view the security events of the bogon host to see if there are any other hosts it has interacted with.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	Anomaly Index, High Concern Index
In what quantity?	- Anomaly Index: True - CI: True
What alarm categories does this security event contribute to for the target?	High Target Index
In what quantity?	TI: True

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_CONN_TO_BOGON_SUCCEEDED (516)

Connection to Tor Attempted

What is the security event?

Questions about this event	Response
What does it mean when it triggers?	Tor, formerly The Onion Router, is a network used for anonymizing Internet connections which works by sending a connection through multiple relays before exiting the Tor network. A Tor entry node is the first server a Tor connection transits through before navigating and exiting the Tor network. This security event involves a host that is being monitored by Secure Network Analytics attempting to communicate with a Tor entry node but that was not observed establishing a successful connection. It is possible that this was either user-driven or malware attempting to find command and control traffic. If it was user-driven, this was likely an attempt at obfuscating the destination of the user's traffic or obfuscating the location from where the user was browsing. It is important to note that some Tor entry nodes run on servers that also run other services. For example, DuckDuckGo runs Tor entry nodes on the same server that is accessed to perform searches.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	High Policy Violation Index, High Concern Index
In what quantity?	• High Policy Violation Index: True • CI: True
What alarm categories does this security event contribute to for the target?	
In what quantity?	

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_TOR_ENTRY_ATTEMPTED (513)

Connection to Tor Successful

What is the security event?

Questions about this event	Response
What does it mean when it triggers?	Tor, formerly The Onion Router, is a network used for anonymizing Internet connections which works by sending a connection through multiple relays before exiting the Tor network. A Tor entry node is the first server a Tor connection transits through before navigating and exiting the Tor network. This security event involves a host that is being monitored by Secure Network Analytics communicating with a Tor entry node. It is possible that this is either user-driven or malware attempting to find command and control traffic. If it was user-driven, this was likely an attempt at obfuscating the destination of the user's traffic or obfuscating the location from where the user was browsing. It is important to note that some Tor entry nodes run on servers that also run other services. For example, DuckDuckGo runs Tor entry nodes on the same server that is accessed to perform searches.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	High Policy Violation Index, High Concern Index
In what quantity?	• High Policy Violation Index: True • CI: True
What alarm categories does this security event contribute to for the target?	
In what quantity?	

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_TOR_ENTRY_SUCCEEDED (514)

Fake Application Detected

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	When processing updates to the application list for a host, a method is called to determine if the application data being processed is a fake application. The fake app logic takes the data from the current flow being analyzed along with the information for the hosts involved. Using this data the algorithm determines if the protocol (TCP or UDP) and ports being used match the expected application identified. The application identification comes from one of the following devices depending on how the network is configured (FlowSensor, Palo Alto Appliance, Packetshaper Appliance, or NBar Appliance). Currently monitored applications include telnet, SSH FTP, mail services, NTP, and DNS. The alarm will trigger if a standard application uses a non-standard port (For example, DNS over 123/udp), or if a particular port has a non-standard application communicating over it (For example, SSH over 80/tcp).
What does it mean when it triggers?	This behavior is often a sign of a person or application attempting to circumvent egress filtering by sending traffic out of a commonly permitted port using a service other than the one intended. It will also look for standard applications over ports other than their standard which can cause it to occasionally fire on applications using secondary ports, like SSH over TCP 8022.
What are next steps?	The goal of your investigation is to find the relevant flows, and then to attempt to determine whether or not they appear to be signs of exfiltration or command and control.

Questions about this event	Response
	A first step in investigating this event is to consider the port attached to the event and the hosts involved. For example, if the target host is an offsite backup server and the event triggered on port 8022 that may not be of concern. On the other hand, if the target host is unfamiliar and you see non-DNS going out over 53 UDP it's likely worth a deeper look. There are a few options to investigate this- the event will list an associated port, and you can use it to build a very specific filter. For example, if you filter for flows between the two involved hosts over the involved port but also exclude flows using the 'proper' application (e.g. searching for flows over port 53 UDP that are not DNS) you'll find the flows that caused the event to fire. It may also be useful to run a more general query instead, filtering on just the source and target host of the event on the date of the event. This may turn up flows beyond just any that have a port / application mismatch, but the existence of any such flows may actual be valuable context in an investigation. Once you have a selection of flows to work with, examine the destination host, history, and data volume for signs of misuse.
Non-standard flow data required? (FlowSensor, proxy, firewall, etc.)	Requires additional information from a FlowSensor, Palo Alto Appliance, Packetshaper Appliance, or NBar Appliance.
Notes	None.

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	
Which policies have this event on by default?	Inside Hosts, Outside Hosts
Which policies have this event off by default?	
Which policies have this alarm on by default?	
Which policies have this alarm off by default?	Inside Hosts, Outside Hosts
Can you tune this event?	no
Is the event variance-based, threshold-based, or other?	N/A
What are default values and units?	N/A
Tolerance	N/A
Minimum threshold value	N/A
Maximum threshold value	N/A
Is event tunable with non-variance-based parameters?	N/A
If so, what are its tunable attributes and units?	N/A

Questions about event	Response
Is event tunable outside of the normal policy editor	N/A
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	
If so, what is it?	

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	N/A
What is the target?	N/A
Which policy causes the event to trigger?	N/A
What alarm categories does this security event contribute to for the source?	High Command And Control Index, High Concern Index
In what quantity?	• High Command And Control Index: True • CI: True
What alarm categories does this security event contribute to for the target?	High Target Index
In what quantity?	• TI: True

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: A host has used X over Port/Protocol (ServiceName). Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source & target IP, a time range 35 minutes before the event's first active time until the event's last active time, and the associated port as <port>/TCP and <port>/UDP. Desktop Client: Flows are filtered by time period of last five minutes.

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_FAKE_APP (62)

Flow Denied

Mitigation is not available for the alarm associated with this security event.

What is the security event?

Questions about this event	Response
What does it mean when it triggers?	This is largely dependent on context. For example, Flow Denied from a host on the Internet is likely not that interesting, especially as a single event. Flow Denied from an inside host to an inside host is different from inside to outside, and trying the same host/port pairing over and over again is much different from getting blocked to a variety of destinations. If an inside host has many denied flows talking to another inside host over one port, it is likely a misconfiguration. Many denied flows to a variety of ports or inside hosts is likely reconnaissance. Apart from that, it is a sign of a host doing things that you have determined it should not do, which makes it inherently interesting. Flows blocked from an inside host to the Internet require the context of what rules it is attempting to violate, but regardless, the host is doing something that your security policy has determined it should not be able to do.
What are next steps?	A good way to investigate Flows Denied is to perform a security event query on the source host of the event. Run the query for the day of the event, set the source host to be the source of the initial Flow Denied event, and then set the type to include only other Flow Denied events. This will show you all of the Flow Denied events for the host. From the returned results you can determine whether it was a single denied flow, many denied flows over the same port to the same host, many denied flows to different hosts over the same port, many denied flows to the same host group, etc. This will enable you to figure out what behavior the host is actually engaging in.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	High Recon Index, High Concern Index
In what quantity?	• High Recon Index: True • CI: True
What alarm categories does this security event contribute to for the target?	High Target Index
In what quantity?	• TI: True


What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_FLOW_DENIED (310)

Frag Packet Too Long **

Mitigation is not available for the alarm associated with this security event.

Frag:Packet_Too_Long



What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	If a packet is seen with fragmentation values that would cause the packet to exceed the maximum packet length, the security event triggers on the originating host.
What does it mean when it triggers?	An IP packet with an invalid fragmentation value is not something that should normally happen. When it happens it has likely been done intentionally in an attempt to gather information about the machine the packet is sent to as different operating systems and versions may respond differently.
What are next steps?	If the source of this event is an inside host, it is worth looking for other signs of recon activity. For example, is the host sending invalid fragments to multiple hosts? Is it scanning many hosts on the same port? Is it scanning many ports on one host? An inside host performing reconnaissance can be a good sign of unwanted activity or compromise.
Non-standard flow data required?	For Flow Collector NetFlow edition, a FlowSensor is required. For Flow Collector sFlow, no special extras are required.

Questions about this event	Response
(FlowSensor, proxy, firewall, etc.)	
Notes	N/A

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	Security event Frag:Packet_Too_Long must be enabled on the originating host
Which policies have this event on by default?	Inside Hosts, Outside Hosts, Client IP Policy
Which policies have this event off by default?	Antivirus & SMS Servers; Firewalls, Proxies, & NAT Devices; Network Management & Scanners
Which policies have this alarm on by default?	None
Which policies have this alarm off by default?	Inside Hosts; Outside Hosts; Client IP Policy; Antivirus & SMS Servers; Firewalls, Proxies, & NAT Devices; Network Management & Scanners
Can you tune this event?	No
Is the event variance-based, threshold-based, or other?	N/A
What are default values and units?	
Tolerance	N/A
Minimum threshold value	N/A

Questions about event	Response
Maximum threshold value	N/A
Is event tunable with non-variance-based parameters?	N/A
If so, what are its tunable attributes and units?	N/A
Is event tunable outside of the normal policy editor	N/A
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	N/A
If so, what is it?	N/A

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The host that sent the bad packet
What is the target?	The host listed as the bad packet's destination
Which policy causes the event to trigger?	Source host policy
What alarm categories does this security event contribute to for the source?	Attack Index, High Concern Index
In what quantity?	• Attack Index: True • CI: True

Questions about event	Response
What alarm categories does this security event contribute to for the target?	High Target Index
In what quantity?	• TI: True

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: View details. Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source & target IP, a time range one minute before the event's first active time until the event's last active time, and TCP. Desktop Client: Flows are filtered by time period of last five minutes.

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_FRAG_PKT_TOO_LONG (282)

Frag First Packet Too Short **

Mitigation is not available for the alarm associated with this security event.

Frag:First_Too_Short

1500-byte payload

Fragment

50 bytes Offset 0

725 bytes Offset 1

725 bytes Offset 2

203.125.234.84

50.56.217.21

Resulting potential security event entry:

Source Host Groups	Source Host	Target Host Groups	Target Host	Concern Index**1	Security Events
Singapore	203.125.234.84	Lancpe	50.56.217.21	<CI value>*	Frag:First_Too_Short-445(1)

* <CI value> represents a numerical point value that Secure Network Analytics assigns to an event based on various conditions.

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	If a packet is seen with a fragmentation length so short that the protocol headers are truncated, the security event triggers on the originating host.
What does it mean when it triggers?	An IP packet with an invalid fragmentation value is not something that should normally happen. When it happens it has likely been done intentionally in an attempt to gather information about the machine the packet is sent to as different operating systems and versions may respond differently.
What are next steps?	If the source of this event is an inside host, it is worth looking for other signs of recon activity. For example, is the host sending invalid fragments to multiple hosts? Is it scanning many hosts on the same port? Is it scanning many ports on one host? An inside host performing reconnaissance can be a good sign of unwanted activity or compromise.
Non-standard flow data required? (FlowSensor, proxy, firewall,	For Flow Collector NetFlow edition, a FlowSensor is required. For Flow Collector sFlow, no special extras are required.

Questions about this event	Response
etc.)	
Notes	N/A

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	Security event Frag:First_Too_Short must be enabled on the originating host
Which policies have this event on by default?	Inside Hosts, Outside Hosts, Client IP Policy
Which policies have this event off by default?	Antivirus & SMS Servers; Firewalls, Proxies, & NAT Devices; Network Management & Scanners
Which policies have this alarm on by default?	None
Which policies have this alarm off by default?	Inside Hosts; Outside Hosts; Client IP Policy; Antivirus & SMS Servers; Firewalls, Proxies, & NAT Devices; Network Management & Scanners
Can you tune this event?	No
Is the event variance-based, threshold-based, or other?	N/A
What are default values and units?	
Tolerance	N/A
Minimum threshold value	N/A

Questions about event	Response
Maximum threshold value	N/A
Is event tunable with non-variance-based parameters?	N/A
If so, what are its tunable attributes and units?	N/A
Is event tunable outside of the normal policy editor	N/A
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	N/A
If so, what is it?	N/A

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The host that sent the bad packet
What is the target?	The host listed as the bad packet's destination
Which policy causes the event to trigger?	Source host policy
What alarm categories does this security event contribute to for the source?	Attack Index, High Concern Index
In what quantity?	• Attack Index: True • CI: True

Questions about event	Response
What alarm categories does this security event contribute to for the target?	High Target Index
In what quantity?	• TI: True

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: View details. Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source & target IP, a time range one minute before the event's first active time until the event's last active time, and TCP. Desktop Client: Flows are filtered by time period of last five minutes.


What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_FRAG_PKT_TOO_SHORT (281)

Frag Sizes Differ **

Mitigation is not available for the alarm associated with this security event.

Frag:Sizes_Differ



Resulting potential security event entry:

Source Host Groups	Source Host	Target Host Groups	Target Host	Concern Index*	Security Events
Singapore	203.125.234.84	Lancopé	50.56.217.21	<CI value>*	Frag:Sizes_Differ-848(2)

* <CI value> represents a numerical point value that Secure Network Analytics assigns to an event based on various conditions.

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	If a packet is seen with fragmentation sizes that differ between segments, the security event triggers on the originating host.
What does it mean when it triggers?	An IP packet with different fragmentation sizes for each segment is not something that should normally happen. When it happens it has likely been done intentionally in an attempt to evade packet inspection tools along its route.
What are next steps?	If the source of this event is an inside host, it is worth looking for other signs of recon activity. For example, is the host sending invalid fragments to multiple hosts? Is it scanning many hosts on the same port? Is it scanning many ports on one host? An inside host performing reconnaissance can be a good sign of unwanted activity or compromise.
Non-standard flow data required? (FlowSensor, proxy, firewall, etc.)	For Flow Collector NetFlow edition, a FlowSensor is required. For Flow Collector sFlow, no special extras are required.

Questions about this event	Response
Notes	N/A

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	Security event Frag:Sizes_Differ must be enabled on the originating host.
Which policies have this event on by default?	Inside Hosts, Outside Hosts, Client IP Policy
Which policies have this event off by default?	Antivirus & SMS Servers; Firewalls, Proxies, & NAT Devices; Network Management & Scanners
Which policies have this alarm on by default?	None
Which policies have this alarm off by default?	Inside Hosts; Outside Hosts; Client IP Policy; Antivirus & SMS Servers; Firewalls, Proxies, & NAT Devices; Network Management & Scanners
Can you tune this event?	No
Is the event variance-based, threshold-based, or other?	N/A
What are default values and units?	
Tolerance	N/A
Minimum threshold value	N/A
Maximum threshold value	N/A

Questions about event	Response
Is event tunable with non-variance-based parameters?	N/A
If so, what are its tunable attributes and units?	N/A
Is event tunable outside of the normal policy editor	N/A
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	N/A
If so, what is it?	N/A

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The host that sent the bad packet
What is the target?	The host listed as the bad packet's destination
Which policy causes the event to trigger?	Source host policy
What alarm categories does this security event contribute to for the source?	Attack Index, High Concern Index
In what quantity?	• Attack Index: True • CI: True
What alarm categories does	High Target Index

Questions about event	Response
this security event contribute to for the target?	
In what quantity?	• TI: True

What information is available in Secure Network Analytics?


Questions about event	Response
What information is displayed in the alarm details?	Manager: View details. Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source & target IP, a time range one minute before the event's first active time until the event's last active time, and TCP. Desktop Client: Flows are filtered by time period of last five minutes.

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_FRAG_DIFFERENT_SIZES (283)

Half Open Attack

Half_Open_Attack



Resulting potential security event entry:

Source Host Groups	Source Host	Target Host Groups	Target Host	Concern Index	Security Events
Singapore	203.125.234.84	Lancopé	50.56.217.21	<CI value>*	Half_Open_Attack

* <CI value> represents a numerical point value that Secure Network Analytics assigns to an event based on various conditions.

What is the security event?

Questions about this event	Response
What does it mean when it triggers?	This is a denial-of-service (DoS) attack. It is possible that a half open attack can be an attempt to exhaust either bandwidth or connection handlers since it opens connections but never carries them out, forcing the victim host to wait long enough to cause the malicious connections to time out.
What are next steps?	The goal is to determine if the connections to the target are legitimate and if the target experiences any performance degradation. While it is possible that this can be malicious behavior between two inside hosts, the event is more likely to be interesting between an inside and outside host. First, run a flow query between the source and target of the event for the day of the event. Take note of whether the flow is ongoing. If it is not, and there was no service outage, it was likely not an issue. You can also check the security event history of both the source and target hosts. Does the source host have this event against multiple targets or historically? If so, this is something to research further. Does the target host have many hosts triggering this event against it at once but not historically? This may

Questions about this event	Response
	be sign of a distributed denial-of-service (DDoS) attack. If the event is also firing historically, it may just be normal behavior for whatever application the service is hosting. If this is the case, consider turning the event off for the particular host. If the relevant traffic was observed by a Flow Sensor, run a performance report to see if the SRT (Server Response Time) has been affected due to the nature of the traffic. This could indicate whether or not this is legitimate traffic.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	'A' Policy: High Ddos Source Index, High Concern Index 'B' Policy: High Ddos Source Index, High Concern Index
In what quantity?	'A' Policy: • High Ddos Source Index: True • CI: True 'B' Policy: • High Ddos Source Index: True • CI: True

Questions about event	Response
What alarm categories does this security event contribute to for the target?	'A' Policy: High Ddos Target Index, High Target Index 'B' Policy: High Ddos Target Index, High Target Index
In what quantity?	'A' Policy: • High Ddos Target Index: True • TI: True 'B' Policy: • High Ddos Target Index: True • TI: True

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_HALF_OPEN (26)

High File Sharing Index

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	It searches for large data transfers that do not originate from a file server. The source (the host downloading the large amount of data) may be using peer-to-peer techniques or simple file sharing methods and protocols to download large files from the target.
What does it mean when it triggers?	The source is likely downloading an unusual amount of data from the target. This is a potential precursor to Data Hoarding and Data Exfiltration.
What are next steps?	Determine how much data was downloaded. As this may be a precursor to Suspect Data Hoarding, the steps are the same, except the initial target is known in this case. An ideal way to investigate this is to run a Top Peers (inbound) report for the source. Set the time period of the query to be the day of the event. Filter the report so that <i>Client</i> or <i>Server Host</i> is the source IP of the security event and <i>Other Host</i> is the Inside Hosts' host group. The goal is to find whether or not the source has done this across many targets. If the target is the only host sending data to the source, run a Top Peers report for the target to see if it has exhibited the same activity with other hosts. If so, it is worth investigating the target as the file sharing server to see if the activity is expected.
Non-standard flow data required? (FlowSensor, proxy, firewall, etc.)	none
Notes	N/A

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	
Which policies have this event on by default?	Inside Host
Which policies have this event off by default?	Outside Host
Which policies have this alarm on by default?	Inside Host on by default
Which policies have this alarm off by default?	Outside Host
Can you tune this event?	yes
Is the event variance-based, threshold-based, or other?	variance
What are default values and units?	N/A
Tolerance	75
Minimum threshold value	100,000 FSI points in 24 hours
Maximum threshold value	1,000,000,000 FSI points in 24 hours
Is event tunable with non-variance-based parameters?	N/A
If so, what are its tunable attributes and units?	N/A

Questions about event	Response
Is event tunable outside of the normal policy editor	N/A
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	
If so, what is it?	

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The host that is receiving the files.
What is the target?	The hosts that are sending the files.
Which policy causes the event to trigger?	
What alarm categories does this security event contribute to for the source?	
In what quantity?	• High Policy Violation Index: True • CI: True
What alarm categories does this security event contribute to for the target?	
In what quantity?	

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: Observed W points. Expected X points, tolerance of Y allows up to Z points. Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source IP and Outside Hosts and the start time (reset hour) of the event's active day until the event's last active time. Desktop Client: Flows are filtered by time period of last five minutes.

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_FILE_SHARING (20)

High SMB Peers

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	The host has many Server Message Block (SMB) sessions to the outside, which is consistent with worm propagation.
What does it mean when it triggers?	Unless you know that this should be happening in a particular scenario, or you have set the threshold very low, this event is almost certainly a sign of a compromised host. It is often a sign of malware looking to spread automatically via SMB.
What are next steps?	This security event is a probable sign of compromise, but you can investigate it to see if the outside hosts belong to a particular range or to see how many of them have been contacted. To do this, start a flow query. Set the start time of the query to the day of the event. On the Filter dialog, on the Hosts tab, set Client or Server Host as the source of the event and Other Host is the Inside Hosts host group. On the Services & Applications tab, filter by Services and include SMB. The query returns all relevant flows.
Non-standard flow data required? (FlowSensor, proxy, firewall, etc.)	None
Notes	None

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	
Which policies have this event on by default?	Inside Host, Outside Host
Which policies have this event off by default?	
Which policies have this alarm on by default?	
Which policies have this alarm off by default?	Inside Host, Outside Host
Can you tune this event?	Yes
Is the event variance-based, threshold-based, or other?	Threshold
What are default values and units?	
Tolerance	N/A
Minimum threshold value	100 SMB flows from an Inside Host to an Outside Host.
Maximum threshold value	N/A
Is event tunable with non-variance-based parameters?	N/A
If so, what are its tunable attributes and units?	N/A

Questions about event	Response
Is event tunable outside of the normal policy editor	N/A
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	
If so, what is it?	

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The Inside Host with many Manager peers
What is the target?	The hosts that are acting as SMB peers.
Which policy causes the event to trigger?	
What alarm categories does this security event contribute to for the source?	
In what quantity?	• High Recon Index: True • Attack Index: True • CI: True
What alarm categories does this security event contribute to for the target?	
In what quantity?	• TI: True

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: View details. Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source IP & Outside Hosts, the start time (reset hour) of the event's active day until the event's last active time, 445/TCP, and 445/UDP. Desktop Client: Flows are filtered by time period of last five minutes.

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_HIGH_SMB_PEERS (60)

High Total Traffic

What is the security event?

Questions about this event	Response
What does it mean when it triggers?	Due to counting both inbound and outbound traffic, High Total Traffic is a general sign of anomalous behavior. It may be predominantly inbound, predominantly outbound, or a mixture of the two. The IP is simply involved in an abnormally higher amount of traffic.
What are next steps?	Determine how much data is moving and where that data is going to and coming from. An ideal way to investigate this is to perform a flow query for the source host. Set the time period to the day of the associated security event, and set the source host to be the client host or server host. Once the results have returned, find where the highest volume of data was going to or coming from; sort by Total Bytes to do this. The goal is to find the standout flow(s). After finding the standout flow(s), determine whether or not this is expected behavior.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	Anomaly Index, High Concern Index
In what quantity?	• Anomaly Index: True • CI: True
What alarm categories does this security event contribute to for the target?	
In what quantity?	

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_TOTAL_TRAFFIC (16)

High Traffic

What is the security event?

Questions about this event	Response
What does it mean when it triggers?	Due to counting both inbound and outbound traffic, High Traffic is a general sign of anomalous behavior. It may be predominantly inbound, predominantly outbound, or a mixture of the two. The IP is simply involved in an abnormally higher amount of traffic. It is a possible indicator of incoming or outgoing DoS.
What are next steps?	Determine how much data is moving and where that data is going to and coming from. An ideal way to investigate this is to perform a flow query for the source host. Set the start time of the query to five minutes before the start time of the associated security event, and set the end time to the end time of the alarm (if there is an alarm). Set the source host to be the client host or server host. Once the results have returned, find where the highest volume of data was going to or coming from; sort by Total Bytes to do this. The goal is to find the standout flow(s). After finding the standout flow(s), determine whether or not this is expected behavior.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	Anomaly Index, High Concern Index
In what quantity?	- Anomaly Index: True - CI: True
What alarm categories does this security event contribute to for the target?	
In what quantity?	

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_HI_TRAFFIC (30)

High Volume Email

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	It searches for hosts that have not been specified as Email servers but are displaying very large outgoing email bursts in relation to a small or non-existent incoming email ratio.
What does it mean when it triggers?	The host alerted on is potentially infected with a piece of email worm malware, or the host has been re-purposed for spamming.
What are next steps?	Determine the amount of data per flow going out and over which ports it is sending them out. Run two Top Peer reports—one filtered by protocol and one filtered by flows—to see whether the host is propagating a worm across the network or is being used as an SMTP relay for spam purposes. Determine whether or not data being sent from these peers in the observed quantity is expected behavior.
Non-standard flow data required? (FlowSensor, proxy, firewall, etc.)	FlowSensor is required for acquiring and working with NetFlow
Notes	N/A

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event	

Questions about event	Response
to trigger?	
Which policies have this event on by default?	Inside Hosts, Outside Hosts
Which policies have this event off by default?	
Which policies have this alarm on by default?	Inside Hosts, Outside Hosts
Which policies have this alarm off by default?	
Can you tune this event?	yes
Is the event variance-based, threshold-based, or other?	Threshold
What are default values and units?	N/A
Tolerance	N/A
Minimum threshold value	N/A
Maximum threshold value	N/A
Is event tunable with non-variance-based parameters?	N/A
If so, what are its tunable attributes and units?	Five-minute periods with High Volume Email Alerts to trigger the alarm; defaults to 1.
Is event tunable outside of the normal policy editor	N/A
If so, what are the alternate	N/A

Questions about event	Response
locations of tunable values?	
Does event have a default mitigation?	
If so, what is it?	

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The host that is being scanned
What is the target?	No target host.
Which policy causes the event to trigger?	
What alarm categories does this security event contribute to for the source?	
In what quantity?	• High Policy Violation Index: True • CI: True
What alarm categories does this security event contribute to for the target?	
In what quantity?	

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: View details Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source IP, the start time (reset hour) of the event's active day until the event's last active time, and 25/TCP. Desktop Client: Flows are filtered by time period of last five minutes.

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_HIGH_VOLUME_EMAIL (9)

ICMP Comm (Communication) Admin **

The source host has received an ICMP message that "communication is administratively prohibited." This is generated if a router cannot forward a packet due to administrative filtering.

Mitigation is not available for the alarm associated with this security event.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	High Recon Index, High Concern Index
In what quantity?	• High Recon Index: True • CI: True
What alarm categories does this security event contribute to for the target?	High Target Index
In what quantity?	• TI: True

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_COMM_ADMIN_PROHIBITED (302)

ICMP Dest (Destination) Host Admin **

The source host has received an ICMP message stating that the "destination host is administratively prohibited."

Mitigation is not available for the alarm associated with this security event.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	High Recon Index, High Concern Index
In what quantity?	• High Recon Index: True • CI: True
What alarm categories does this security event contribute to for the target?	High Target Index
In what quantity?	• TI: True

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_DEST_HOST_ADMIN_PROHIBITED (299)

ICMP Dest (Destination) Host Unk (Unknown)**

The source host has received an ICMP message that a "destination host unknown" error has occurred.

Mitigation is not available for the alarm associated with this security event.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	High Recon Index, High Concern Index
In what quantity?	• High Recon Index: True • CI: True
What alarm categories does this security event contribute to for the target?	High Target Index
In what quantity?	• TI: True

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_DEST_HOST_UNKNOWN (296)

ICMP Dest (Destination) Net Admin **

The source host has received an ICMP message that the "destination network is administratively prohibited."

Mitigation is not available for the alarm associated with this security event.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	High Recon Index, High Concern Index
In what quantity?	• High Recon Index: True • CI: True
What alarm categories does this security event contribute to for the target?	High Target Index
In what quantity?	• TI: True

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_DEST_NETWK_ADMIN_PROHIBITED (298)

ICMP Dest (Destination) Net Unk (Unknown) **

The source host has received an ICMP message that a "destination network unknown" error has occurred.

Mitigation is not available for the alarm associated with this security event.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	High Recon Index, High Concern Index
In what quantity?	• High Recon Index: True • CI: True
What alarm categories does this security event contribute to for the target?	High Target Index
In what quantity?	• TI: True

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_DEST_NETWORK_UNKNOWN (295)

ICMP Flood

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	The source host has sent an excessive number of ICMP packets in a 5-minute period.
What does it mean when it triggers?	This indicates that the source host is engaged in a DoS attack (potentially caused by either a compromised host or driven by a user), a misconfigured network application, or a reconnaissance (during which the packets are sent across a very wide range of IPs rather than across a small group).
What are next steps?	Determine how many ICMP packets are being sent at what rate, when they were sent, and where they are going. An ideal way to investigate this is to perform a flow query for the source host. Set the start time of the query to five minutes before the start time of the associated security event, and set the end time to the end time of the alarm (if there is an alarm). Set the protocol filter to ICMP only. Once the results have returned, find where the most ICMP packets went during the searched time period. If the subject of the event is the client host in its ICMP-heavy flows, you can sort on Client Packets or Client Packet Rate (pps). The goal is to find the standout flow(s). After finding the standout flow(s), determine whether or not this is expected behavior.
Non-standard flow data required? (FlowSensor, proxy, firewall, etc.)	None
Notes	N/A

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	
Which policies have this event on by default?	Inside Hosts, Outside Hosts
Which policies have this event off by default?	
Which policies have this alarm on by default?	Inside Hosts, Outside Hosts
Which policies have this alarm off by default?	
Can you tune this event?	Yes
Is the event variance-based, threshold-based, or other?	Behavioral and Threshold
What are default values and units?	
Tolerance	75
Minimum threshold value	1,800 ICMP Packets in 5 minutes
Maximum threshold value	10 Million ICMP Packets in 5 minutes
Is event tunable with non-variance-based parameters?	N/A
If so, what are its tunable attributes and units?	N/A

Questions about event	Response
Is event tunable outside of the normal policy editor	N/A
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	
If so, what is it?	

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The host receiving the ICMP Packets
What is the target?	N/A
Which policy causes the event to trigger?	
What alarm categories does this security event contribute to for the source?	
In what quantity?	• High Ddos Source Index: True • CI: True
What alarm categories does this security event contribute to for the target?	
In what quantity?	

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: Observed X pp5m. Expected Y pp5m, tolerance of Z allows up to Y pp5m. Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source IP, a time range five minutes before the event's first active time until the event's last active time, and ICMP. Desktop Client: Flows are filtered by time period of last five minutes.

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_ICMP_FLOOD (7)

ICMP Frag (Fragmentation) Needed **

The source host has received an ICMP message stating that "the IP datagram is too big". Packet fragmentation is required, but the DF bit in the IP header is set.

Mitigation is not available for the alarm associated with this security event.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	Anomaly Index, High Concern Index
In what quantity?	- Anomaly Index: True - CI: True
What alarm categories does this security event contribute to for the target?	High Target Index
In what quantity?	TI: True

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_FRAG_NEEDED_DF_SET (293)

ICMP Host Precedence **

The source host has received an ICMP message that "host precedence violation" has occurred. Sent by the first hop router to a host to indicate that a requested precedence is not permitted for the particular combination of source/destination host or network, upper layer protocol, and source/destination port.

Mitigation is not available for the alarm associated with this security event.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	Anomaly Index, High Concern Index
In what quantity?	- Anomaly Index: True - CI: True
What alarm categories does this security event contribute to for the target?	High Target Index
In what quantity?	TI: True

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_HOST_PRECEDENCE_VIOLATION (303)

ICMP Host Unreach **

The source host has received an ICMP message that the "target host is unreachable."

Mitigation is not available for the alarm associated with this security event.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	High Recon Index, High Concern Index
In what quantity?	• High Recon Index: True • CI: True
What alarm categories does this security event contribute to for the target?	High Target Index
In what quantity?	• TI: True

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_HOST_UNREACHABLE (290)

ICMP Host Unreach TOS (Type of Service) **

The source host has received an ICMP message stating that "target host is unreachable due to the specified Type of Service."

Mitigation is not available for the alarm associated with this security event.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	Anomaly Index, High Concern Index
In what quantity?	- Anomaly Index: True - CI: True
What alarm categories does this security event contribute to for the target?	High Target Index
In what quantity?	TI: True

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_HOST_UNREACHABLE_FOR_SVC (301)

ICMP Net Unreach **

The source host has received an ICMP message that "the target network is unreachable."

Mitigation is not available for the alarm associated with this security event.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	High Recon Index, High Concern Index
In what quantity?	• High Recon Index: True • CI: True
What alarm categories does this security event contribute to for the target?	High Target Index
In what quantity?	• TI: True

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_NETWORK_UNREACHABLE (289)

ICMP Net Unreach TOS **

The source host has received an ICMP message that the "network is unreachable for the specified Type Of Service."

Mitigation is not available for the alarm associated with this security event.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	Anomaly Index, High Concern Index
In what quantity?	- Anomaly Index: True - CI: True
What alarm categories does this security event contribute to for the target?	High Target Index
In what quantity?	TI: True

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_NETWORK_UNREACHABLE_FOR_SVC (300)

ICMP Port Unreach **

The source host has received an ICMP message stating that "the destination port is unreachable." The designated transport protocol (e.g., UDP) is unable to demultiplex the datagram, but has no protocol mechanism to inform the sender.

Mitigation is not available for the alarm associated with this security event.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	High Recon Index, High Concern Index
In what quantity?	• High Recon Index: True • CI: True
What alarm categories does this security event contribute to for the target?	High Target Index
In what quantity?	• TI: True

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_PORT_UNREACHABLE (292)

ICMP Precedence Cutoff **

The source host has received an ICMP message that "precedence cutoff" is in effect. The network operators have imposed a minimum level of precedence required for operation; the datagram was sent with a precedence below this level.

Mitigation is not available for the alarm associated with this security event.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	Anomaly Index, High Concern Index
In what quantity?	- Anomaly Index: True - CI: True
What alarm categories does this security event contribute to for the target?	High Target Index
In what quantity?	TI: True

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_PRECEDENCE_CUTOFF (304)

ICMP Proto (Protocol) Unreach **

The source host has received an ICMP message stating that "the destination protocol is unreachable."

Mitigation is not available for the alarm associated with this security event.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	Anomaly Index, High Concern Index
In what quantity?	- Anomaly Index: True - CI: True
What alarm categories does this security event contribute to for the target?	High Target Index
In what quantity?	TI: True

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_PROTOCOL_UNREACHABLE (291)

ICMP Received

What is the security event?

Questions about this event	Response
What does it mean when it triggers?	This indicates that the target host is being targeted in a DoS attack or a misconfigured network application.
What are next steps?	Determine how many ICMP packets are being received at what rate, when they were received, and from where they are coming. An ideal way to investigate this is to perform a flow query for the target host. Set the start time of the query to five minutes before the start time of the associated security event, and set the end time to the end time of the alarm (if there is an alarm). Set the protocol filter to <i>ICMP only</i>. Once the results have returned, find where the most ICMP packets came from during the searched time period. If the subject of the event is the client host in its ICMP-heavy flows, you can sort on Client Packets or Client Packet Rate (pps). The goal is to find the standout flow(s). After finding the standout flow(s), determine whether or not this is expected behavior.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	
In what quantity?	
What alarm categories does this security event contribute to for the target?	High Ddos Target Index, High Target Index
In what quantity?	• High Ddos Target Index: True • TI: True

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_ICMP_RECEIVED (50)

ICMP Src (Source) Host Isolated **

The source host has received an ICMP message that a "source host isolated error" has occurred.

Mitigation is not available for the alarm associated with this security event.

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	High Recon Index, High Concern Index
In what quantity?	• High Recon Index: True • CI: True
What alarm categories does this security event contribute to for the target?	High Target Index
In what quantity?	• TI: True

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_SOURCE_HOST_ISOLATED (297)

ICMP Src (Source) Route Failed **

The source host has received an ICMP message stating that a "source route failed" error has occurred.

Mitigation is not available for the alarm associated with this security event.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	Anomaly Index, High Concern Index
In what quantity?	- Anomaly Index: True - CI: True
What alarm categories does this security event contribute to for the target?	High Target Index
In what quantity?	TI: True


What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_SOURCE_ROUTE_FAIL (294)

ICMP Timeout

Mitigation is not available for the alarm associated with this security event.

ICMP_Timeout



Resulting potential security event entry:

Source Host Groups	Source Host	Target Host Groups	Target Host	Concern Index	Security Events
Singapore	203.125.234.84	Lancop	50.56.217.21	<CI value>*	ICMP_Timeout(10)

* <CI value> represents a numerical point value that Secure Network Analytics assigns to an event based on various conditions.

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	It searches for the ICMP_TIME_EXCEEDED message generated as a result of the Traceroute tool, malfunctioning network devices (routing loops, etc), and firewalking techniques (using ICMP timeouts to determine ACLs allowed layer 4 protocols).
What does it mean when it triggers?	The host is potentially scanning / mapping the network to determine which layer 4 protocols are allowed through gateways/firewalls, etc.
What are next steps?	Determine if the host has been scanning ports along with the ICMP invocation. Run reports that show peers by port and protocol to discern a pattern of ICMP as well as various ports, usually repeating across peers. Determine whether or not data patterns being sent from the host is expected behavior (security tool/device, etc).
Non-standard flow data required? (FlowSensor, proxy, firewall,	For analyzing Netflow, a FlowSensor is required.

Questions about this event	Response
etc.)	
Notes	None

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	
Which policies have this event on by default?	Inside Hosts, Outside Hosts
Which policies have this event off by default?	
Which policies have this alarm on by default?	
Which policies have this alarm off by default?	Inside Hosts, Outside Hosts
Can you tune this event?	no
Is the event variance-based, threshold-based, or other?	N/A
What are default values and units?	N/A
Tolerance	N/A
Minimum threshold value	N/A
Maximum threshold value	N/A

Questions about event	Response
Is event tunable with non-variance-based parameters?	N/A
If so, what are its tunable attributes and units?	N/A
Is event tunable outside of the normal policy editor	N/A
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	
If so, what is it?	

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The host that is the packet's destination
What is the target?	The host that is the packet's source
Which policy causes the event to trigger?	
What alarm categories does this security event contribute to for the source?	
In what quantity?	• High Recon Index: True • CI: True
What alarm categories does	

Questions about event	Response
this security event contribute to for the target?	
In what quantity?	• TI: True

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: View details Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source & target IP, a time range five minutes before the event's first active time until the event's last active time, and ICMP. Desktop Client: Flows are filtered by time period of last five minutes.

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_ICMP_TO (258)

Inside Tor Entry Detected

What is the security event?

Questions about this event	Response
What does it mean when it triggers?	Tor, formerly The Onion Router, is a network used for anonymizing Internet connections which works by sending a connection through multiple relays before exiting the Tor network. A Tor entry node is the first server a Tor connection transits through before navigating and exiting the Tor network. This security event involves a host that is being monitored by Secure Network Analytics and is being advertised as a Tor entry node. Hosting a Tor entry node is not necessarily an issue, however it's almost certainly not something that should be occurring without a network's administrators being aware. Unlike a Tor exit node, Tor entry nodes only forward traffic within the Tor network, so the main concern is usually bandwidth.
What are next steps?	The main goal of investigating this event is largely non-technical. Determine whether or not anyone is aware of the Tor entry node and if it is permitted. If so, disable the event for the specific permitted host. You can also view the host's security event history to find out when it began to communicate while being advertised as a Tor entry node.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	High Policy Violation Index, High Concern Index
In what quantity?	• High Policy Violation Index: True • CI: True
What alarm categories does this security event contribute to for the target?	High Target Index
In what quantity?	• TI: True

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_TOR_ENTRY_INSIDE_HOST (515)

Inside Tor Exit Detected

What is the security event?

Questions about this event	Response
What does it mean when it triggers?	Tor, formerly The Onion Router, is a network used for anonymizing Internet connections which works by sending a connection through multiple relays before exiting the Tor network. A Tor exit node is the last relay through which a Tor connection exits the Tor network, and it is what the destination of a network connection ultimately see as as the source of the connection. This security event involves a host that is being monitored by Secure Network Analytics and is being advertised as a Tor exit node. Hosting a Tor exit node is not necessarily an issue, however it's almost certainly not something that should be occurring without a network's administrators being aware. Concerns about hosting a Tor exit node often includes bandwidth usage and potential legal concerns around the activity being proxied through the node.
What are next steps?	The main goal of investigating this event is largely non-technical. Determine whether or not anyone is aware of the Tor exit node and if it is permitted. If so, disable the event for the specific permitted host. You can also view the host's security event history to find out when it began to communicate while being advertised as a Tor exit node. In addition, you can run a Top Peers report to view the amount of bandwidth the server has consumed as well as where that bandwidth has been going.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	High Policy Violation Index, High Concern Index
In what quantity?	• High Policy Violation Index: True • CI: True
What alarm categories does this security event contribute to for the target?	High Target Index
In what quantity?	• TI: True

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_TOR_EXIT_INSIDE_HOST (319)

Low Traffic

The host's traffic average over the last 5 minutes has fallen below the minimum acceptable values.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	Anomaly Index, High Concern Index
In what quantity?	- Anomaly Index: True - CI: True
What alarm categories does this security event contribute to for the target?	
In what quantity?	

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_LOW_TRAFFIC (29)

MAC Address Violation

The host has changed its MAC address more than an acceptable number of times since the last archive hour.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	High Policy Violation Index, High Concern Index
In what quantity?	• High Policy Violation Index: True • CI: True
What alarm categories does this security event contribute to for the target?	High Target Index
In what quantity?	• TI: True

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_MAC_ADDRESS (25)

Mail Rejects

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	It searches for excess Mail Rejected messages from hosts that are not SMTP (mail) servers.
What does it mean when it triggers?	A host has received a number of Mail Rejected messages that exceed the threshold and the host is not a mail server. These messages may occur for a number of reasons. For example, the sending server/IP is registered as a spam server and the intended mail recipient is no longer available.
What are next steps?	Determine whether the host is intended to be sending/receiving mail. Run a Top Peers report by Port / Protocol (25/SMTP) to see who the server has been communicating with. If the host is intended to be sending/receiving mail, check the SMTP configurations to ensure they are correct. Check the larger spam and mail blacklists to see if the host IP is listed. If the server is not intended to send and receive email, investigate the host to see when the services were installed and if it was compromised in order to accomplish this.
Non-standard flow data required? (FlowSensor, proxy, firewall, etc.)	A Flow Sensor is required to analyze NetFlow.
Notes	N/A

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	
Which policies have this event on by default?	Inside Hosts
Which policies have this event off by default?	Outside Hosts
Which policies have this alarm on by default?	Inside Hosts
Which policies have this alarm off by default?	Outside Hosts
Can you tune this event?	Yes
Is the event variance-based, threshold-based, or other?	Five-minute periods with Email Reject alerts to trigger alarm: Default to 5. Rejected email deliveries in 5 minutes to trigger alarm: Default to 5.
What are default values and units?	
Tolerance	N/A
Minimum threshold value	N/A
Maximum threshold value	N/A
Is event tunable with non-variance-based parameters?	N/A
If so, what are its tunable attributes and units?	

Questions about event	Response
Is event tunable outside of the normal policy editor	
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	
If so, what is it?	

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The host that is being scanned
What is the target?	No target host
Which policy causes the event to trigger?	
What alarm categories does this security event contribute to for the source?	
In what quantity?	• High Policy Violation Index: True • CI: True
What alarm categories does this security event contribute to for the target?	
In what quantity?	

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: View details Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source IP, the start time (reset hour) of the event's active day until the event's last active time and 25/TCP. Desktop Client: Flows are filtered by time period of last five minutes.

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_MAIL_REJECTS (12)

Mail Relay

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	It searches for hosts that have accumulated a high volume of email traffic, both inbound and outbound, and are sending larger volumes of email than they are receiving.
What does it mean when it triggers?	The host is potentially being used as a mail relay. If a mail server is not configured to use authentication, it is left open to the outside world with the SMTP port (25) open and/or does not have configurations to limit the networks that can relay mail. It may be taken advantage of by spammers.
What are next steps?	Determine the volume of email that is normal and whether the host is intended to send and receive email. If the server is intended to send and receive email, verify the SMTP server configurations. Enable authentication and network relay limits if they are not already enabled. If the server is not intended to send and receive email, investigate the host to see when the services were installed and whether the host was compromised to accomplish this.
Non-standard flow data required? (FlowSensor, proxy, firewall, etc.)	A Flow Sensor is required to analyze NetFlow.
Notes	N/A

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	
Which policies have this event on by default?	Inside Hosts
Which policies have this event off by default?	Outside Hosts
Which policies have this alarm on by default?	Inside Hosts
Which policies have this alarm off by default?	Outside Hosts
Can you tune this event?	Yes
Is the event variance-based, threshold-based, or other?	Five-minute periods with Email Alerts to trigger alarm: Default to 5.
What are default values and units?	N/A
Tolerance	N/A
Minimum threshold value	N/A
Maximum threshold value	N/A
Is event tunable with non-variance-based parameters?	N/A
If so, what are its tunable attributes and units?	N/A

Questions about event	Response
Is event tunable outside of the normal policy editor	N/A
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	
If so, what is it?	

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The host being scanned
What is the target?	No target host
Which policy causes the event to trigger?	
What alarm categories does this security event contribute to for the source?	
In what quantity?	• High Policy Violation Index: True • CI: True
What alarm categories does this security event contribute to for the target?	
In what quantity?	

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: View details Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source IP, the start time (reset hour) of the event's active day until the event's last active time and 25/TCP. Desktop Client: Flows are filtered by time period of last five minutes.

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_MAIL_RELAY (10)

Max Flows Initiated

What is the security event?

Questions about this event	Response
What does it mean when it triggers?	This can indicate several different issues, but you should first determine if this is a sign of reconnaissance or a host engaging in a DoS attack.
What are next steps?	Determine what is causing these types flows by finding a common link between them. For example, are most of these flows going to the same host? Are they going to different hosts using the same port? An ideal way to determine this is to run one of the top reports. You can start with either the Top Peers report or Top Ports report. Access the report and then, on the Hosts tab in the Filter dialog, set the Direction field to <i>Total</i>, where <i>Client</i> is the source of the event and there are no exclusions for <i>Server Host</i>. On the Date/Time tab set the time period to five minutes before the start active time of the event up to the start active time of the event. On the Advanced tab, sort by <i>Flows</i>. Determine if the majority of the flows applies to a certain set of rows. (If this is not the case, run the other of the two suggested Top reports.) After finding the standout ports or peers, determine whether or not this is expected behavior.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	Anomaly Index, High Concern Index
In what quantity?	• Anomaly Index: True • CI: True
What alarm categories does this security event contribute to for the target?	
In what quantity?	

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_MAX_FLOWS_INIT (17)

Max Flows Served

What is the security event?

Questions about this event	Response
What does it mean when it triggers?	This can indicate several different issues, but you should first determine if this is a sign of reconnaissance or a host being targeted in DoS attacks. For hosts that normally serve large volumes of traffic, it may just be an increase in normal use depending on threshold and tolerance settings.
What are next steps?	Determine what is causing these types flows by finding a common link between them. For example, are most of these flows coming from the same host? Are they coming from different hosts targeting the same port? An ideal way to determine this is to run one of the top reports. You can start with either the Top Peers report or Top Ports report. Access the report and then, on the Hosts tab in the Filter dialog, set the Direction field to <i>Total</i>, where <i>Server Host</i> is the source of the event and there are no exclusions for <i>Client</i>. On the Date/Time tab set the time period to five minutes before the start active time of the event up to the start active time of the event. On the Advanced tab, sort by <i>Flows</i>. Determine if the majority of the flows applies to a certain set of rows. (If this is not the case, run the other of the two suggested Top reports.) After finding the standout ports or peers, determine whether or not this is expected behavior.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	
In what quantity?	
What alarm categories does this security event contribute to for the target?	Anomaly Index, High Target Index
In what quantity?	- Anomaly Index: True - TI: True

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_MAX_FLOWS_SERVED (37)

New Flows Initiated

What is the security event?

Questions about this event	Response
What does it mean when it triggers?	This can indicate several different issues, but you should first determine if this is a sign of reconnaissance or a host engaging in a DoS attack.
What are next steps?	Determine what is causing these types flows by finding a common link between them. For example, are most of these flows going to the same host? Are they going to different hosts using the same port? An ideal way to determine this is to run one of the top reports. You can start with either the Top Peers report or Top Ports report. Access the report and then, on the Hosts tab in the Filter dialog, set the Direction field to <i>Total</i>, where <i>Client</i> is the source of the event and there are no exclusions for <i>Server Host</i>. On the Date/Time tab set the time period to five minutes before the start active time of the event up to the start active time of the event. On the Advanced tab, sort by <i>Flows</i>. Determine if the majority of the flows applies to a certain set of rows. (If this is not the case, run the other of the two suggested Top reports.) After finding the standout ports or peers, determine whether or not this is expected behavior.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	Anomaly Index, High Concern Index
In what quantity?	• Anomaly Index: True • CI: True
What alarm categories does this security event contribute to for the target?	
In what quantity?	

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_NEW_FLOWS_INIT (18)

New Flows Served

What is the security event?

Questions about this event	Response
What does it mean when it triggers?	This can indicate several different issues, but you should first determine if this is a sign of reconnaissance or a host being targeted in DoS attacks. For hosts that normally serve large volumes of traffic, it may just be an increase in normal use depending on threshold and tolerance settings.
What are next steps?	Determine what is causing these types flows by finding a common link between them. For example, are most of these flows going to the same host? Are they going to different hosts using the same port? An ideal way to determine this is to run one of the top reports. You can start with either the Top Peers report or Top Ports report. Access the report and then, on the Hosts tab in the Filter dialog, set the Direction field to <i>Total</i>, where <i>Server Host</i> is the source of the event and there are no exclusions for <i>Client</i>. On the Date/Time tab set the time period to five minutes before the start active time of the event up to the start active time of the event. On the Advanced tab, sort by <i>Flows</i>. Determine if the majority of the flows applies to a certain set of rows. (If this is not the case, run the other of the two suggested Top reports.) After finding the standout ports or peers, determine whether or not this is expected behavior.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	
In what quantity?	
What alarm categories does this security event contribute to for the target?	High Ddos Target Index, High Target Index
In what quantity?	• High Ddos Target Index: True • TI: True

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_NEW_FLOWS_SERVED (38)

New Host Active

A new host has been detected.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?


Questions about event	Response
What alarm categories does this security event contribute to for the source?	High Policy Violation Index, High Concern Index
In what quantity?	• High Policy Violation Index: True • CI: True
What alarm categories does this security event contribute to for the target?	
In what quantity?	

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_NEW_HOST (14)

Packet Flood

Packet_Flood



What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	It looks for a large volume of small-sized packets being sent over a long period of time to a host or hosts that are not SYN Flood attacks. (A SYN Flood attack occurs when a three-way TCP handshake is not completed, which then leaves the TCP connection open).
What does it mean when it triggers?	The source is sending out a large volume of packets, where each packet is smaller in size than a traditional TCP packet. This is occurring over a long period of time, indicating that the source is attempting a DoS (Denial of Service) or Brute Force attack against the target.
What are next steps?	Run a Top Peers report on the source to see if it has executed any similar flows, perhaps slightly shorter in length and close in size to other hosts. This can be a precursor to the attack, as the attacker is trying to gauge resources on various machines and identify weak targets. Run a Top Peers report on the target to see if any other machines have been compromised or re-purposed to join in the attack but have not yet reached the timing threshold for the alarm to fire.

Questions about this event	Response
	Investigate the source to see if the flow activity is behaving as expected. Configure DoS mitigation on the targets, if possible.
Non-standard flow data required? (FlowSensor, proxy, firewall, etc.)	A Flow Sensor is required to analyze NetFlow.
Notes	N/A

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	
Which policies have this event on by default?	Inside Hosts, Outside Hosts
Which policies have this event off by default?	
Which policies have this alarm on by default?	Inside Hosts, Outside Hosts
Which policies have this alarm off by default?	
Can you tune this event?	Yes
Is the event variance-based, threshold-based, or other?	This alarm is triggered when the PPS is greater than 75,000 for Inside Hosts. or when the PPS is greater than 50,000 for Outside Hosts.

Questions about event	Response
What are default values and units?	
Tolerance	N/A
Minimum threshold value	N/A
Maximum threshold value	N/A
Is event tunable with non-variance-based parameters?	N/A
If so, what are its tunable attributes and units?	N/A
Is event tunable outside of the normal policy editor	N/A
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	
If so, what is it?	

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The attacker sending the short packets.
What is the target?	The receiver of the packets.
Which policy causes the event to trigger?	

Questions about event	Response
What alarm categories does this security event contribute to for the source?	High Ddos Source Index, High Concern Index
In what quantity?	• High Ddos Source Index: True • CI: True
What alarm categories does this security event contribute to for the target?	
In what quantity?	• High Ddos Target Index: True • TI: True

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: Target Host is X. Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source & target IP and a time range five minutes before the event's first active time until the event's last active time. Desktop Client: Flows are filtered by time period of last five minutes.

What information is available for responses for this security event?


Questions about event	Response
What is the event ID?	SEC_ID_PACKET_FLOOD (8)

Ping

The source host has sent an ICMP echo reply and has received an ICMP echo response from the target.

Mitigation is not available for the alarm associated with this security event.

Ping



What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	High Recon Index, High Concern Index
In what quantity?	- High Recon Index: True - CI: True
What alarm categories does this security event contribute to for the target?	High Target Index
In what quantity?	TI: True


What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_PING_PROBE (257)

Ping Oversized Packet

Mitigation is not available for the alarm associated with this security event.

Ping_Oversized_Packet



Ping_Oversized_Packet events also apply to ICMP Echo Replies

Resulting potential security event entry:

Source Host Groups	Source Host	Target Host Groups	Target Host	Concern Index	Security Events
Singapore	203.125.234.84	Lancpe Corporate	50.56.217.21	<CI value>*	Ping_Oversized_Packet(58)

* <CI value> represents a numerical point value that Secure Network Analytics assigns to an event based on various conditions.

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	It searches for ICMP packets that are larger than the standard size of 90 bytes, either as an ICMP echo request (if the host is the destination of the packet) or as an ICMP echo reply (if the host is the source of the packet).
What does it mean when it triggers?	This may indicate that a covert channel for data exfiltration or secretive communication is being used. ICMP is used for standard control or error response purposes, but it can be tunneled through to send and receive other data.
What are next steps?	Run a Top Services report to distinguish the excessive ICMP traffic from the all the other traffic sent to and from the host. Investigate the source and target hosts to see if unauthorized services are running. For example, a program or service must be running on the host, target, or both to capture or send the ICMP traffic. Run a Top Peers report filtered by the ICMP application (see note below) to see if the source or target have been communicating with other peers over excessive

Questions about this event	Response
	ICMP or to unfavorable peers outside the network (data exfiltration). Filter the report by ICMP for the Applications Note: When configuring the Top Peers report, in the Connection section, click Select under Applications.
Non-standard flow data required? (FlowSensor, proxy, firewall, etc.)	None
Notes	None

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	
Which policies have this event on by default?	Inside Hosts, Outside Hosts
Which policies have this event off by default?	
Which policies have this alarm on by default?	
Which policies have this alarm off by default?	Inside Hosts, Outside Hosts
Can you tune this event?	No
Is the event variance-based,	N/A

Questions about event	Response
threshold-based, or other?	
What are default values and units?	N/A
Tolerance	N/A
Minimum threshold value	N/A
Maximum threshold value	N/A
Is event tunable with non-variance-based parameters?	N/A
If so, what are its tunable attributes and units?	N/A
Is event tunable outside of the normal policy editor	N/A
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	
If so, what is it?	

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The host that sent the packet
What is the target?	The host that received the packet

Questions about event	Response
Which policy causes the event to trigger?	
What alarm categories does this security event contribute to for the source?	
In what quantity?	• High Command And Control Index: True • CI: True
What alarm categories does this security event contribute to for the target?	
In what quantity?	• TI: True

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: View details Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source & target IP, the start time (reset hour) of the event's active day until the event's last active time, and ICMP. Desktop Client: Flows are filtered by time period of last five minutes.


What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_LONG_PING (278)

Ping Scan

Mitigation is not available for the alarm associated with this security event.

Ping_Scan



Resulting potential security event entry:

Source Host Groups	Source Host	Target Host Groups	Target Host	Concern Index	Security Events
Singapore	203.125.234.84	Lancope Corporate	50.56.217.21	<CI value>*	Ping_Scan(72)

* <CI value> represents a numerical point value that Secure Network Analytics assigns to an event based on various conditions.

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	The scanning event is recorded and the alarm is triggered if the following is true: 1) the timeout is exceeded for an ICMP datagram; 2) it is not a service request; 3) the source host is not the destination host; 4) the source/destination port is not 38293 (used for Norton AV Host Discovery).
What does it mean when it triggers?	An ICMP scan is generally an early type of reconnaissance since, unlike scans of a particular port across hosts or many ports on one host, it is simply searching for any responsive host rather than a particular service. A host ping scanning is generally just trying to find other active hosts on a network, possibly for further investigation.
What are next steps?	Determine what the host was scanning. Since this is an ICMP-based scan, your main goal is to figure out which hosts are being scanned rather than what particular services are being targeted, as you would with a UDP or TCP scan. You can investigate this by performing a flow query where the client host or server host is the source of the event, the time period is the day of the event, and the protocol is ICMP. Sort the results by IP

Questions about this event	Response
	or host group of the peer to try to determine whether or not there is a pattern to the hosts being scanned.
Non-standard flow data required? (FlowSensor, proxy, firewall, etc.)	For Flow Collector NetFlow edition, a Flow Sensor is required. For Flow Collector SFlow edition, nothing extra is required.
Notes	none

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	
Which policies have this event on by default?	Inside Hosts, Outside Hosts
Which policies have this event off by default?	
Which policies have this alarm on by default?	
Which policies have this alarm off by default?	Inside Hosts, Outside Hosts
Can you tune this event?	No
Is the event variance-based, threshold-based, or other?	N/A
What are default values and units?	N/A

Questions about event	Response
Tolerance	N/A
Minimum threshold value	N/A
Maximum threshold value	N/A
Is event tunable with non-variance-based parameters?	N/A
If so, what are its tunable attributes and units?	N/A
Is event tunable outside of the normal policy editor	N/A
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	
If so, what is it?	

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The host that performs the scan
What is the target?	The host that is being scanned
Which policy causes the event to trigger?	
What alarm categories does this security event contribute	

Questions about event	Response
to for the source?	
In what quantity?	• High Recon Index: True • CI: True
What alarm categories does this security event contribute to for the target?	
In what quantity?	

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: View details. Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source IP, the target's natural class C network (/24), the start time (reset hour) of the event's active day until the event's last active time, and ICMP. Desktop Client: Flows are filtered by time period of last five minutes.

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_PING_ADDR_SCAN (277)

Port Scan

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	Port Scan counts how many target ports under or equal to 1024 one host talks on to another host. When the configurable threshold of target ports is crossed, the event is triggered.
What does it mean when it triggers?	It means that a host has been observed "port scanning" another host. This means that the scanning host is attempting to identify which services the scanned host can provide. When this is done by an outside host, it can be the precursor to attempts at exploitation; scanning from the internet, however, is to be expected. When the scanning host is an inside host, it can be a sign of either an attacker or a compromised host trying to move laterally within an organization's network.
What are next steps?	If the source host of this security event is an outside host, check to see if it has had continued communication with the host that it scanned. To do this, open a flow query between the source and target host. It may also be wise to simply view all traffic between the source host and your network. If the source of this host is an inside host, view the flows between the source and target using a Flow Query to determine whether or not this is expected behavior.
Non-standard flow data required? (FlowSensor, proxy, firewall, etc.)	None
Notes	The source port will likely be over or equal to port 1024.

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	Port scan must be enabled on host performing the scan
Which policies have this event on by default?	Inside Hosts, Outside Hosts
Which policies have this event off by default?	None
Which policies have this alarm on by default?	None
Which policies have this alarm off by default?	Inside Hosts, Outside Hosts
Can you tune this event?	Yes
Is the event variance-based, threshold-based, or other?	Threshold based
What are default values and units?	N/A
Tolerance	N/A
Minimum threshold value	N/A
Maximum threshold value	N/A
Is event tunable with non-variance-based parameters?	Yes
If so, what are its tunable attributes and units?	The tunable parameter is the number of scanned ports required to fire, and the default is 10.

Questions about event	Response
Is event tunable outside of the normal policy editor	N/A
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	N/A
If so, what is it?	N/A

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The host performing the scan
What is the target?	The host being scanned
Which policy causes the event to trigger?	The source host
What alarm categories does this security event contribute to for the source?	High Recon Index, High Concern Index
In what quantity?	• High Recon Index: True • CI: True
What alarm categories does this security event contribute to for the target?	High Target Index
In what quantity?	• TI: True

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: View details. Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source & target IP and the start time (reset hour) of the event's active day until the event's last active time. Desktop Client: Flows are filtered by time period of last five minutes.


What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_PORT_SCAN (55)

Reset/tcp

Mitigation is not available for the alarm associated with this security event.

Reset/tcp



What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	It searches for TCP Reset packets coming from a target port that 1) is not used for HTTP; 2) has not had TCP activity on it; 3) has an ICMP type of "Destination Unreachable"; 4) is not SNMP; and 5) is not using unprivileged ports (1024+).
What does it mean when it triggers?	TCP Reset packets are sent when a TCP conversation is torn down or when there is a problem with the conversation. These can be sent by malfunctioning devices in the middle of a conversation (such as load balancers), or they can be sent by the target if a port is not open or available. In the case of the latter, if multiple ports are responding this way and are in groups or are congruent, it usually indicates a network scan is being performed to determine which services, if any, are available on a host.
What are next steps?	Run a Top Peers report for the source to determine if there is a pattern to the ports that are being used in the flows. Determine whether the source is a security host, device, or appliance that has been authorized to scan.

Questions about this event	Response
	If not, investigate the host for scanning software or services, as well as evidence of compromise.
Non-standard flow data required? (FlowSensor, proxy, firewall, etc.)	A Flow Sensor is required to analyze NetFlow.
Notes	None

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	
Which policies have this event on by default?	Inside Hosts, Outside Hosts
Which policies have this event off by default?	
Which policies have this alarm on by default?	
Which policies have this alarm off by default?	Inside Hosts, Outside Hosts
Can you tune this event?	No
Is the event variance-based, threshold-based, or other?	N/A
What are default values and units?	N/A

Questions about event	Response
Tolerance	N/A
Minimum threshold value	N/A
Maximum threshold value	N/A
Is event tunable with non-variance-based parameters?	N/A
If so, what are its tunable attributes and units?	N/A
Is event tunable outside of the normal policy editor	N/A
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	
If so, what is it?	

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The host that is the packet's destination
What is the target?	The host that is the packet's source
Which policy causes the event to trigger?	
What alarm categories does this security event contribute	

Questions about event	Response
to for the source?	
In what quantity?	• High Recon Index: True • CI: True
What alarm categories does this security event contribute to for the target?	
In what quantity?	• TI: True

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: View details Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source & target IP, a time range five minutes before the event's first active time until the event's last active time, the associated port as <port>/TCP, and TCP. Desktop Client: Flows are filtered by time period of last five minutes.


What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_TCP_PROBE (262)

Reset/udp

Mitigation is not available for the alarm associated with this security event.

Reset/udp



Resulting potential security event entry:

Source Host Groups	Source Host	Target Host Groups	Target Host	Concern Index	Security Events
Singapore	203.125.234.84	Lancop	50.56.217.21	<CI value>*	Reset/udp-110(3) Reset/udp-143(6)

* <CI value> represents a numerical point value that Secure Network Analytics assigns to an event based on various conditions.

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	It searches for UDP Reset packets coming from a target port that 1) is not used for HTTP; 2) has not had UDP activity on it; 3) has an ICMP type of "Destination Unreachable"; 4) is not SNMP; and 5) is not using unprivileged ports (1024+).
What does it mean when it triggers?	UDP Reset packets are sent when communication to a UDP port that is closed or not available has been attempted. These packets can be sent by malfunctioning devices in the middle of a conversation or they can be sent by the target if a port is not open or available. In the case of the latter, if multiple ports are responding this way and are in groups or are congruent, it usually indicates a network scan is being performed to determine which services, if any, are available on a host.
What are next steps?	Run a Top Peers report for the source to determine if there is a pattern to the ports that are being used in the flows. Determine whether the source is a security host, device, or appliance that has been authorized to scan.

Questions about this event	Response
	If not, investigate the host for scanning software or services, as well as evidence of compromise.
Non-standard flow data required? (FlowSensor, proxy, firewall, etc.)	A Flow Sensor is required to analyze NetFlow.
Notes	None

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	
Which policies have this event on by default?	Inside Hosts, Outside Hosts
Which policies have this event off by default?	
Which policies have this alarm on by default?	
Which policies have this alarm off by default?	Inside Hosts, Outside Hosts
Can you tune this event?	No
Is the event variance-based, threshold-based, or other?	N/A
What are default values and units?	N/A

Questions about event	Response
Tolerance	N/A
Minimum threshold value	N/A
Maximum threshold value	N/A
Is event tunable with non-variance-based parameters?	N/A
If so, what are its tunable attributes and units?	N/A
Is event tunable outside of the normal policy editor	N/A
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	
If so, what is it?	

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The host that is the packet's destination
What is the target?	The host that is the packet's source
Which policy causes the event to trigger?	
What alarm categories does this security event contribute	

Questions about event	Response
to for the source?	
In what quantity?	• High Recon Index: True • CI: True
What alarm categories does this security event contribute to for the target?	
In what quantity?	• TI: True

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: View details Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source & target IP, a time range one minute before the event's first active time until the event's last active time, the associated port as <port>/UDP, and UDP. Desktop Client: Flows are filtered by time period of last five minutes.

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_UDP_PROBE (261)

Scanner Talking

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	A host that has been scanning your network now has a two-way conversation with one of the target hosts that it scanned. It is enabled by default in both the Inside Hosts policy and the Outside Hosts policy. It is, however, disabled by default in the Network Management Scanners role policy.
What does it mean when it triggers?	A machine which has been flagged as an address scanner is communicating with one of the machines that has been scanned. If the source is not an authorized scanner, it can indicate the scanner has found a target matching its attack vector and has successfully compromised it.
What are next steps?	Determine whether the source is an authorized scanner. If the source is not an authorized scanner, run a Peer report for the source to see other hosts communicated to, making sure to choose peers that have increased flows or flow sizes. Investigate the source and target for signs of compromise.
Non-standard flow data required? (FlowSensor, proxy, firewall, etc.)	
Notes	

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	
Which policies have this event on by default?	Inside Hosts, Outside Hosts
Which policies have this event off by default?	none
Which policies have this alarm on by default?	Inside Hosts, Outside Hosts
Which policies have this alarm off by default?	none
Can you tune this event?	no
Is the event variance-based, threshold-based, or other?	N/A
What are default values and units?	N/A
Tolerance	N/A
Minimum threshold value	N/A
Maximum threshold value	N/A
Is event tunable with non-variance-based parameters?	N/A
If so, what are its tunable attributes and units?	N/A

Questions about event	Response
Is event tunable outside of the normal policy editor	lc_threshold value addr_scan_talking_stale_timeout controls the amount of time before the global scan list is reset.
If so, what are the alternate locations of tunable values?	none
Does event have a default mitigation?	
If so, what is it?	

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The scanner that is communicating
What is the target?	The scanned host that is being communicated with
Which policy causes the event to trigger?	The source host policy
What alarm categories does this security event contribute to for the source?	
In what quantity?	• Attack Index: True • CI: True
What alarm categories does this security event contribute to for the target?	
In what quantity?	• TI: True

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: View details Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source & target IP, a time range five minutes before the event's first active time until the event's last active time, and client bytes & server bytes (both of which are greater than or equal to 1). Desktop Client: Flows between the client and the server, starting at the alarm time.

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_ADDR_SCAN_TALKING (63)

Slow Connection Flood

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	The source host has sent an excessive number of simultaneous connections to a single destination. Each connection has a very low packet-per-second rate.
What does it mean when it triggers?	Indicates that a host is initiating multiple connections to a target with a very low packet rate. The goal of connections like these are to maintain open connections without requiring much bandwidth. This is a type of application denial-of-service (DoS) attack which can render a service unavailable without requiring a large amount of bandwidth, which makes the attack easier to carry out against a vulnerable host and often more difficult to detect.
What are next steps?	The goal is to determine if the connections to the target are legitimate and if the target experiences any performance degradation. While it is possible that this can be malicious behavior between two inside hosts, the event is more likely to be interesting between an inside and an outside host. First, run a flow query between the source and target of the event for the day of the event. Take note of whether the flow is ongoing. If it is not, and there was no service outage, it was likely not an issue. You can also check the security event history of both the source and target hosts. Does the source host have this event against multiple targets or historically? If so, this is something to research further. Does the target host have many hosts triggering this event against it at once but not historically? This may be sign of a distributed denial-of-service (DDoS) attack. If the event is also firing historically, it may just be normal behavior for whatever application the service is hosting. If this is the case, consider turning t

Questions about this event	Response
	he event off for the particular host. If the relevant traffic was observed by a Flow Sensor, run a performance report to see if the SRT (Server Response Time) has been affected due to the nature of the traffic. This could indicate whether or not this is legitimate traffic.
Non-standard flow data required? (FlowSensor, proxy, firewall, etc.)	None
Notes	None

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	
Which policies have this event on by default?	Inside Hosts, Outside Hosts
Which policies have this event off by default?	
Which policies have this alarm on by default?	Outside Hosts
Which policies have this alarm off by default?	
Can you tune this event?	Yes
Is the event variance-based,	Threshold

Questions about event	Response
threshold-based, or other?	
What are default values and units?	N/A
Tolerance	N/A
Minimum threshold value	30
Maximum threshold value	N/A
Is event tunable with non-variance-based parameters?	Yes
If so, what are its tunable attributes and units?	The number of slow connections needed to trigger the alarm
Is event tunable outside of the normal policy editor	N/A
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	
If so, what is it?	

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The host sending the excessive number of slow connections
What is the target?	The host receiving the excessive number of slow connections

Questions about event	Response
Which policy causes the event to trigger?	
What alarm categories does this security event contribute to for the source?	
In what quantity?	• High Ddos Source Index: True • CI: True
What alarm categories does this security event contribute to for the target?	
In what quantity?	• High Ddos Target Index: True • TI: True

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: The target host X has observed Y low-pps connections using port Z. Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source & target IP, a time range two minutes before the event's first active time until the event's last active time, the ports 80/TCP, 443/TCP, and 8080/TCP. Desktop Client: Flows are filtered by time period of last five minutes.

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_SLOW_CONNECTION_FLOOD (44)

Spam Source

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	It is searching for sources with an excessive volume of outgoing email or for an excessive address-to-message ratio (i.e., many recipients per single message).
What does it mean when it triggers?	It can indicate that a server is being used to send spam email. This can be accomplished through compromise or misconfiguration. This can cause public IPs to be blocked by blacklists and spam lists.
What are next steps?	Determine if the source is intended to send and receive email. If the source is not meant to do so, determine when the increase in volume of outgoing email began. Search for signs of compromise or malicious software installation around that time. If the source is intended to send and receive email, determine when the source first appeared on the network and search for signs of compromise.
Non-standard flow data required? (FlowSensor, proxy, firewall, etc.)	A Flow Sensor is required to analyze NetFlow.
Notes	N/A

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event	

Questions about event	Response
to trigger?	
Which policies have this event on by default?	Inside Hosts
Which policies have this event off by default?	Outside Hosts
Which policies have this alarm on by default?	Inside Hosts
Which policies have this alarm off by default?	Outside Hosts
Can you tune this event?	Yes
Is the event variance-based, threshold-based, or other?	Five-minute periods with Spam Source alerts to trigger alarm: Default 5. Addresses per email for Spam Email alert: Default 10.
What are default values and units?	
Tolerance	N/A
Minimum threshold value	N/A
Maximum threshold value	N/A
Is event tunable with non-variance-based parameters?	N/A
If so, what are its tunable attributes and units?	N/A
Is event tunable outside of the normal policy editor	N/A

Questions about event	Response
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	
If so, what is it?	

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The host that is being scanned
What is the target?	No target host
Which policy causes the event to trigger?	
What alarm categories does this security event contribute to for the source?	
In what quantity?	• High Policy Violation Index: True • CI: True
What alarm categories does this security event contribute to for the target?	
In what quantity?	

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: View details Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source IP, the start time (reset hour) of the event's active day until the event's last active time, and 25/TCP. Desktop Client: Flows are filtered by time period of last five minutes.

What information is available for responses for this security event?


Questions about event	Response
What is the event ID?	SEC_ID_SPAM_SOURCE (11)

Src=Des (Source=Destination)

The source and target hosts of an IP datagram are the same. This security event usually results from a crafted packet intended to disrupt routing.

Mitigation is not available for the alarm associated with this security event.

Src=Des



What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	Anomaly Index
In what quantity?	Anomaly Index: True
What alarm categories does this security event contribute to for the target?	High Target Index
In what quantity?	TI: True

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_BOOMERANG (273)

SSH Reverse Shell

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	Detects an SSH session that appears to be a reverse shell. More data is being sent to the outside host than is being received.
What does it mean when it triggers?	The intent of this security event alarm is to find a reverse SSH shell. An attacker may use an SSH reverse shell as a method of establishing on-demand access to the compromised network through what looks like an outbound SSH connection. This allows an attacker to maintain access to a system even when incoming connections are blocked off by something like a firewall.
What are next steps?	When investigating this security event, it is important to keep the role of the source host, presumably on the inside network, in perspective. If the host in question is an SSH or SFTP server in an official capacity, this could be a false alarm, and further host policy tuning may be required to limit these alarms from such devices. However, in a situation where inside hosts do not have any obvious capacity as such a device, the identity of the outside host in question must be established. Is it owned by a well-known entity? If so, it could be a false alarm. If it is not a well-known entity, run a Top Peers (outbound) report on the target host. This report provides a list of inside peers sorted by the amount of data that have communicated with the outside host. This indicates whether or not it a very common peer within your environment and can potentially assist you in determining if this behavior is expected.
Non-standard flow data required?	None

Questions about this event	Response
(FlowSensor, proxy, firewall, etc.)	
Notes	None

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	
Which policies have this event on by default?	Inside Hosts, Outside Hosts
Which policies have this event off by default?	
Which policies have this alarm on by default?	
Which policies have this alarm off by default?	Inside Hosts, Outside Hosts
Can you tune this event?	Yes
Is the event variance-based, threshold-based, or other?	Threshold
What are default values and units?	N/A
Tolerance	N/A
Minimum threshold value	Traffic exceeds 10,000 bytes AND client data percentage exceeds 60.

Questions about event	Response
Maximum threshold value	Client data percentage exceeds 80.
Is event tunable with non-variance-based parameters?	N/A
If so, what are its tunable attributes and units?	N/A
Is event tunable outside of the normal policy editor	N/A
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	
If so, what is it?	

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The inside host with traffic to an SSH port of an outside host
What is the target?	The outside host receiving SSH traffic
Which policy causes the event to trigger?	
What alarm categories does this security event contribute to for the source?	
In what quantity?	• High Command And Control Index: True • CI: True

Questions about event	Response
What alarm categories does this security event contribute to for the target?	
In what quantity?	• TI: True

What information is available in Secure Network Analytics?

Questions about event	Response
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source & target IP, a time range five minutes before the event's first active time until the event's last active time, the port 22/TCP. Desktop Client: Flows are filtered by time period of last five minutes.


What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_SSH_REV_SHELL (61)

Stealth Scan/tcp

Mitigation is not available for the alarm associated with this security event.

Stealth_Scan/tcp



Resulting potential security event entry:

Source Host Groups	Source Host	Target Host Groups	Target Host	Concern Index*	Security Events
Singapore	203.125.234.84	Lancop	50.56.217.21	<CI value>*	Stealth_Scan/tcp-848(2) Stealth_Scan/tcp-231(1)

* <CI value> represents a numerical point value that Secure Network Analytics assigns to an event based on various conditions.

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	The source host has used the same source port to connect to different ports on the target host at the same time. This behavior indicates applications that have used raw sockets to create TCP or UDP connections. The security event shows the last target port accessed before the security event was recognized.
What does it mean when it triggers?	N/A
What are next steps?	N/A
Non-standard flow data required? (FlowSensor, proxy, firewall, etc.)	For Flow Collector NetFlow edition, a Flow Sensor is required. For Flow Collector sFlow edition, nothing extra is required.
Notes	N/A

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	
Which policies have this event on by default?	Inside Hosts, Outside Hosts
Which policies have this event off by default?	
Which policies have this alarm on by default?	
Which policies have this alarm off by default?	Inside Hosts, Outside Hosts
Can you tune this event?	no
Is the event variance-based, threshold-based, or other?	N/A
What are default values and units?	N/A
Tolerance	N/A
Minimum threshold value	N/A
Maximum threshold value	N/A
Is event tunable with non-variance-based parameters?	N/A
If so, what are its tunable attributes and units?	N/A

Questions about event	Response
Is event tunable outside of the normal policy editor	N/A
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	
If so, what is it?	

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The host that performs the scan
What is the target?	The host that is being scanned
Which policy causes the event to trigger?	
What alarm categories does this security event contribute to for the source?	
In what quantity?	'1' Policy: • CI: True '8000' Policy: • High Recon Index: True • CI: True
What alarm categories does this security event contribute to for the target?	

Questions about event	Response
In what quantity?	'1' Policy: • TI: True '8000' Policy: • TI: True

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: View details. Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source & target IP, a time range five minutes before the event's first active time until the event's last active time, the associated port as <port>/TCP, and TCP. Desktop Client: Flows are filtered by time period of last five minutes.

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_TCP_STEALTH (272)

Stealth Scan/udp

Mitigation is not available for the alarm associated with this security event.

Stealth_Scan/udp



Resulting potential security event entry:

Source Host Groups	Source Host	Target Host Groups	Target Host	Concern Index*	Security Events
Singapore	203.125.234.84	Lancop	50.56.217.21	<CI value>*	Stealth_Scan/udp-5000(2) Stealth_Scan/udp-5001(2)

* <CI value> represents a numerical point value that Secure Network Analytics assigns to an event based on various conditions.

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	The event fires when Secure Network Analytics detects traffic which would not be generated naturally by an operating system, suggesting that someone is deliberately crafting packets in an attempt to evade detection.
What does it mean when it triggers?	N/A
What are next steps?	N/A
Non-standard flow data required? (FlowSensor, proxy, firewall, etc.)	For Flow Collector NetFlow edition, a Flow Sensor is required. For Flow Collector sFlow edition, nothing extra is required.
Notes	N/A

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	
Which policies have this event on by default?	Inside Hosts, Outside Hosts
Which policies have this event off by default?	
Which policies have this alarm on by default?	
Which policies have this alarm off by default?	Inside Hosts, Outside Hosts
Can you tune this event?	no
Is the event variance-based, threshold-based, or other?	N/A
What are default values and units?	N/A
Tolerance	N/A
Minimum threshold value	N/A
Maximum threshold value	N/A
Is event tunable with non-variance-based parameters?	N/A
If so, what are its tunable attributes and units?	N/A

Questions about event	Response
Is event tunable outside of the normal policy editor	N/A
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	
If so, what is it?	

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The host that performs the scan.
What is the target?	The host that is being scanned
Which policy causes the event to trigger?	
What alarm categories does this security event contribute to for the source?	
In what quantity?	'1' Policy: • CI: True '8000' Policy: • High Recon Index: True • CI: True
What alarm categories does this security event contribute to for the target?	

Questions about event	Response
In what quantity?	'1' Policy: • TI: True '8000' Policy: • TI: True

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: View details. Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source & target IP, a time range one minute before the event's first active time until the event's last active time, the associated port as <port>/UDP, and UDP. Desktop Client: Flows are filtered by time period of last five minutes.

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_UDP_STEALTH (271)

Suspect Data Hoarding

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	Suspect Data Hoarding monitors how much TCP/UDP data an inside host, while acting as a client, downloads from other inside hosts in a reset period. The event fires when the amount of data surpasses the threshold for a given host. This threshold can either be built automatically by baselining or it can be set manually.
What does it mean when it triggers?	This event is potentially an indication of a particular host gathering data to prepare for exfiltration or other large-than-normal downloads of internal data.
What are next steps?	Determine how much data was downloaded and where that data was downloaded from. An ideal way to investigate this is to run a Top Peers (inbound) report. Set the time period of the query to be the day of the event. Filter the report so that <i>Client</i> or <i>Server Host</i> is the source IP of the security event and <i>Other Host</i> is the Inside Hosts host group. The goal is to find the standout peers who sent the majority of the data. After finding the standout peers, determine whether or not data being sent from these peers in the observed quantity is expected behavior.
Non-standard flow data required? (FlowSensor, proxy, firewall, etc.)	N/A
Notes	N/A

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	The event must be enabled for Inside Host host group.
Which policies have this event on by default?	Inside Hosts, Outside Hosts
Which policies have this event off by default?	Firewalls, Proxies, & NAT Devices
Which policies have this alarm on by default?	None
Which policies have this alarm off by default?	Inside Hosts; Outside Hosts; Firewalls, Proxies, & NAT Devices
Can you tune this event?	Yes
Is the event variance-based, threshold-based, or other?	Variance or threshold-based.
What are default values and units?	
Tolerance	92
Minimum threshold value	500M client payload bytes in 24 hours.
Maximum threshold value	1T downloaded payload bytes in 24 hours.
Is event tunable with non-variance-based parameters?	Yes
If so, what are its tunable attributes and units?	The event can be set to trigger off a particular number of bytes that does not take baselines or tolerance into account.

Questions about event	Response
Is event tunable outside of the normal policy editor	No
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	No
If so, what is it?	N/A

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The inside host that is receiving the payload data.
What is the target?	The inside host or hosts that are sending the payload data.
Which policy causes the event to trigger?	The source's policy.
What alarm categories does this security event contribute to for the source?	High Data Hoarding Index, High Concern Index
In what quantity?	• High Data Hoarding Index: True • CI: True
What alarm categories does this security event contribute to for the target?	
In what quantity?	

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: Observed W bytes. Expected X bytes, tolerance of Y allows up to Z bytes. Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source vs. Inside Hosts, the start time (reset hour) of the event's active day until the event's last active time, 53/UDP, 67/UDP, 68/UDP, 161/TCP, 161/UDP, 162/TCP, 162/UDP, client bytes & server bytes (both of which are greater than or equal to 1), and total bytes greater or equal to 1000. Desktop Client: Flows are filtered by time period of last five minutes.

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_SUSPECT_DATA_HOARD (315)

Suspect Data Loss

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	When this event triggers, an inside host acting as a client has uploaded a cumulative amount of TCP or UDP payload data to an outside host, and the cumulative amount exceeds the threshold value set in the policy applied to the inside host. This can be used as a variance-based alarm.
What does it mean when it triggers?	A host is being used to upload more information to the Internet than is acceptable. This can be anything from someone using external backup services to maliciously exfiltrating corporate data.
What are next steps?	Determine how much data was uploaded and where the bulk of that data went. An ideal way to determine this is to run a Top Peers (outbound) report on the host that was the source of the security event. Filter the report so that the client is the source of the event and the server host group is set to <i>Outside Hosts</i>. The goal is to find the standout peers who received the majority of the data. After finding the standout peers, determine whether or not data being received by these peers in the observed quantity is expected behavior.
Non-standard flow data required? (FlowSensor, proxy, firewall, etc.)	No
Notes	N/A

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	This event must be enabled in policies applied to both the source and target hosts.
Which policies have this event on by default?	Inside Hosts, Outside Hosts, Client IP Policy
Which policies have this event off by default?	Firewalls, Proxies, & NAT Devices; Guest Wireless; Mail Server Policy; Trusted Internet Hosts
Which policies have this alarm on by default?	Inside Hosts, Outside Hosts, Client IP Policy
Which policies have this alarm off by default?	Firewalls, Proxies, & NAT Devices; Guest Wireless; Mail Server Policy; Trusted Internet Hosts
Can you tune this event?	Yes
Is the event variance-based, threshold-based, or other?	Variance or threshold-based
What are default values and units?	
Tolerance	50
Minimum threshold value	1G client payload bytes in 24 hours
Maximum threshold value	100G client payload bytes in 24 hours
Is event tunable with non-variance-based parameters?	Yes
If so, what are its tunable attributes and units?	The event can be set to trigger off a particular number of bytes that does not take baselines or tolerance into account.

Questions about event	Response
Is event tunable outside of the normal policy editor	N/A
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	No
If so, what is it?	N/A

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The inside host, acting as a client, that is sending the payload data.
What is the target?	The outside host or hosts that are receiving the payload data from your inside host.
Which policy causes the event to trigger?	The source's policy.
What alarm categories does this security event contribute to for the source?	High Exfiltration Index, High Concern Index
In what quantity?	• High Exfiltration Index: True • CI: True
What alarm categories does this security event contribute to for the target?	
In what quantity?	

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: Observed W bytes. Expected X bytes, tolerance of Y allows up to Z bytes. Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source IP and Outside Hosts and the start time (reset hour) of the event's active day until the event's last active time. Desktop Client: Flows are filtered by time period of last five minutes.

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_DATA_LOSS (40)

Suspect Long Flow

What is the security event?

Questions about this event	Response
What does it mean when it triggers?	This security event detects long-established connections such as the ones used by remote desktop technologies and VPN, but it can also detect communication channels such as spyware, IRC botnets, and other covert means of communication. This event is not necessarily indicative of malicious behavior or compromise by itself, but it contributes to the identification of interesting hosts. When you perform an upgrade to your system, this security event should start firing again after 6 days.
What are next steps?	The goal is to determine if this traffic represents normal activity. Oftentimes you can do this by identifying the destination of the traffic. A good starting point is to pivot into the associated flow table for the security event to get more details of the flow such as source and destination IP addresses, ports, services, starting time of the flow, geo-location information, and username (if available). If the target is an outside host, run an external lookup to verify the owner of the IP. If the target is a known business partner or trusted network, you should classify the target into its corresponding group. If an external lookup of the host is not enough to identify it, run a Top Peer reports to identify other hosts communicating with the target to identify similar behaviors.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	High Command And Control Index, High Concern Index
In what quantity?	• High Command And Control Index: True • CI: True
What alarm categories does this security event contribute to for the target?	High Target Index
In what quantity?	• TI: True

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_SUSPECT_LONG_FLOW (33)

Suspect Quiet Long Flow

What is the security event?

Questions about this event	Response
What does it mean when it triggers?	This security event identifies heartbeat connections that are used in some types of command and control (C&C) activity as well as other suspicious communication channels such as spyware, IRC botnets, and other covert means of communication. This event is similar to a beaconing host except that bidirectional communications are included, provided a low volume of data is transferred. Note that this event can also identify background website heartbeats.
What are next steps?	The goal is to determine if this traffic represents normal activity. Oftentimes you can do this by identifying the destination of the traffic. A good starting point is to pivot into the associated flow table for the security event to get more details of the flow such as source and destination IP addresses, ports, services, starting time of the flow, geo-location information, and username (if available). If the target is an outside host, run an external lookup to verify the owner of the IP. If the target is a known business partner or trusted network, you should classify the target into its corresponding group. If an external lookup of the host is not enough to identify it, run a Top Peer reports to identify other hosts communicating with the target to identify similar behaviors. You can also run a Flow Traffic report to visualize suspicious traffic patterns. Heartbeats to a C&C server will exhibit a periodic pattern with a low amount of bytes transferred.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	High Command And Control Index, High Concern Index
In what quantity?	• High Command And Control Index: True • CI: True
What alarm categories does this security event contribute to for the target?	High Target Index
In what quantity?	• TI: True

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_QUIET_LONG_DURATION_FLOW (48)

Suspect UDP Activity

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	The source host has been identified scanning multiple hosts on a UDP port and has successfully sent a large UDP packet to another previously scanned host. This type of behavior is consistent with many single-packet UDP-based worms such as "SQL Slammer" and "Witty." Investigate this security event immediately.
What does it mean when it triggers?	The source host has successfully sent a single, large UDP packet to one of the previously scanned hosts.
What are next steps?	Investigate this activity immediately, since it is most commonly identified with many UDP-based worms.
Non-standard flow data required? (FlowSensor, proxy, firewall, etc.)	A Flow Sensor is required for acquiring and working with NetFlow.
Notes	N/A

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	
Which policies have this event on by default?	Inside Hosts, Outside Hosts
Which policies have this event	

Questions about event	Response
off by default?	
Which policies have this alarm on by default?	Inside Hosts, Outside Hosts
Which policies have this alarm off by default?	
Can you tune this event?	No
Is the event variance-based, threshold-based, or other?	N/A
What are default values and units?	N/A
Tolerance	N/A
Minimum threshold value	N/A
Maximum threshold value	N/A
Is event tunable with non-variance-based parameters?	N/A
If so, what are its tunable attributes and units?	N/A
Is event tunable outside of the normal policy editor	N/A
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	
If so, what is it?	

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The host that sent the packet
What is the target?	The host being scanned and receiving the packet
Which policy causes the event to trigger?	
What alarm categories does this security event contribute to for the source?	
In what quantity?	• Attack Index: True • CI: True
What alarm categories does this security event contribute to for the target?	
In what quantity?	• TI: True

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: Source host is using X service (Y Protocol) as Peer to Z. Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source & target IP, a

Questions about event	Response
	time range 35 minutes before the event's first active time until the event's last active time, and the associated port as <port>/UDP. Desktop Client: Flows are filtered by time period of last five minutes.

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_SUSPECT_UDP_ACTIVITY (24)

SYN Flood

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	The source host has sent an excessive number of TCP connection requests (SYN packets) in a 5-minute period. This may indicate a Denial of Service (DoS) attack or non-stealthy scanning activity.
What does it mean when it triggers?	This indicates that the source host is engaged in a DoS attack (potentially caused by either a compromised host or driven by a user) or a misconfigured network application. Many SYN packets are often sent to consume a target's bandwidth as part of larger, distributed denial-of- service (DDoS) attacks. However, it is possible that an application is failing to establish a TCP connection and is trying to re-establish one at a very high rate. Relatively low packet counts relative to a large number of IPs is also likely to be a sign of reconnaissance.
What are next steps?	Determine how many SYNs are being sent at what rate, when they were sent, and where they are going. An ideal way to investigate this is to perform a flow query for the source host. Set the start time of the query to five minutes before the start time of the associated security event, and set the end time to the end time of the alarm (if there is an alarm). Once the results have returned, find where the most SYN packets went during the searched time period. If the subject of the event is the client host in its SYN- heavy flows, you can sort on Client SYN Packets, Client Packets, or Client Packet Rate (pps). The goal is to find the standout flow (s). After finding the standout flow(s), determine whether or not this is a misconfiguration. An Inside to Inside SYN Flood is often a misconfiguration; an Inside to Outside SYN Flood occasionally is.

Questions about this event	Response
Non-standard flow data required? (FlowSensor, proxy, firewall, etc.)	None
Notes	N/A

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	
Which policies have this event on by default?	Inside Hosts, Outside Hosts
Which policies have this event off by default?	
Which policies have this alarm on by default?	Inside Hosts, Outside Hosts
Which policies have this alarm off by default?	
Can you tune this event?	Yes
Is the event variance-based, threshold-based, or other?	behavioral and threshold
What are default values and units?	N/A
Tolerance	75

Questions about event	Response
Minimum threshold value	10 SYN packets per period
Maximum threshold value	4 Million SYN packets per period
Is event tunable with non-variance-based parameters?	N/A
If so, what are its tunable attributes and units?	N/A
Is event tunable outside of the normal policy editor	N/A
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	
If so, what is it?	

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The host sending SYN Packets
What is the target?	None
Which policy causes the event to trigger?	
What alarm categories does this security event contribute to for the source?	High Ddos Source Index, High Concern Index
In what quantity?	High Ddos Source

Questions about event	Response
	Index: True • Cl: True
What alarm categories does this security event contribute to for the target?	
In what quantity?	

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: Observed W pp5m. Expected X pp5m, tolerance of Y pp5m allows up to Z pp5m. Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source IP, a time range five minutes before the event's first active time until the event's last active time, and TCP. Desktop Client: Flows are filtered by time period of last five minutes.

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_SYN_FLOOD (5)

SYNs Received

What is the security event?

Questions about this event	Response
What does it mean when it triggers?	This indicates that the target host is being targeted in a DoS attack or misconfigured network application. Many SYN packets are often sent to consume a target's bandwidth as part of larger DDoS or DoS attacks. However, it is possible that an application is failing to establish a TCP connection and trying to re-establish one at a very high rate.
What are next steps?	Determine how many SYNs are being sent at what rate, when they were sent, and from where they are coming. An ideal way to investigate this is to perform a flow query for the target host. Set the start time of the query to five minutes before the start time of the associated security event, and set the end time to the end time of the alarm (if there is an alarm). Once the results have returned, find where the most SYN packets came from during the searched time period. If the target of the event is the server host in its SYN-heavy flows, you can sort on Client SYN Packets, Client Packets, or Client Packet Rate (pps). The goal is to find the standout flow(s). After you find the standout flow(s), determine whether or not this is a misconfiguration. An Inside to Inside SYN Flood is often a misconfiguration.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	
In what quantity?	
What alarm categories does this security event contribute to for the target?	High Ddos Target Index, High Target Index
In what quantity?	• High Ddos Target Index: True • TI: True

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_SYNS_RECEIVED (19)

Talks to Phantoms

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	Any time a flow exists where one host is sending non-trivial traffic to a host which has never communicated with anything, add the silent host to the "phantom list" of the host that is communicating. If the phantom list of a host grows past the configured allowable maximum, then trigger the event from the talking host to the most recently seen phantom host.
What does it mean when it triggers?	Since computers rarely attempt to send unsolicited traffic to hosts that do not exist, a host involved in such activity is noteworthy. It is possible that this is a sign of someone attempting to perform reconnaissance, however, it is also possible that a host has been misconfigured to talk to nonexistent hosts.
What are next steps?	This security event can be a little tricky to investigate since it is not possible today to specify "phantoms" as part of a flow query, but you can get close. The goal of the investigation is to find the flows that were sent to unresponsive hosts and draw meaning from the common thread between the flows. Find the unidirectional flows. Build a flow query where the time range is the day of the security event. Make the client host the source IP of the security event, and allow the server host to be any IP. Find the common thread between these flows. For example, were they all going to the same server port? Did they all occur at approximately the same time? Were they going to hosts that you know once existed? How many hosts were there? In the Desktop Client: In the Traffic tab, filter so that the number of server packets are less than or equal to

Questions about this event	Response
	0 (leave the "Greater than or equal to" part of the setting empty). To make sure everything is captured, you should repeat this query where the source of the event is the server of the flow query, and client packets are less than or equal to 0. Note that this is unlikely to produce additional results too frequently. In the Web App: On the Flow Search page, in the Advanced Subject Options section, set the Packets field to < 1 and the Orientation field to Server. To make sure everything is captured, you should repeat this query where the source of the event is the server of the flow query, and client packets are < 1. Note that this is unlikely to produce additional results too frequently.
Non-standard flow data required? (FlowSensor, proxy, firewall, etc.)	None
Notes	N/A

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	
Which policies have this event on by default?	Inside Host, Outside Hosts
Which policies have this event off by default?	
Which policies have this alarm	None

Questions about event	Response
on by default?	
Which policies have this alarm off by default?	Inside Host, Outside Hosts
Can you tune this event?	Yes
Is the event variance-based, threshold-based, or other?	Threshold
What are default values and units?	N/A
Tolerance	N/A
Minimum threshold value	Three phantom hosts that a particular host has attempted to reach.
Maximum threshold value	N/A
Is event tunable with non-variance-based parameters?	N/A
If so, what are its tunable attributes and units?	N/A
Is event tunable outside of the normal policy editor	N/A
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	
If so, what is it?	

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The host that attempts to talk to the phantom host(s).
What is the target?	The phantom host
Which policy causes the event to trigger?	
What alarm categories does this security event contribute to for the source?	
In what quantity?	• High Recon Index: True • CI: True
What alarm categories does this security event contribute to for the target?	
In what quantity?	• TI: True

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: View details. Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source IP, the start time (reset hour) of the event's active day until the event's

Questions about event	Response
	last active time, and bytes & packets (both of which are less than 1). Desktop Client: Flows are filtered by time period of last five minutes.

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_TALKS_TO_PHANTOMS (59)

Target Data Hoarding

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	Target Data Hoarding monitors how much TCP/UDP data an inside host, while acting as a server, serves to other inside hosts in a reset period. The event fires when the amount of data surpasses the threshold for a given host. This threshold can either be built automatically by baselining or it can be set manually.
What does it mean when it triggers?	This event is potentially an indication of one or many Inside Hosts gathering more data than normal from a particular Inside Host, potentially in preparation for exfiltration or misuse.
What are next steps?	Determine how much data was transferred and where the bulk of that data went. An ideal way to determine this is to run a Top Peers (outbound) report on the host that was the source of the security event. Filter the report so that <i>Client</i> or <i>Server Host</i> is the target IP of the security event and <i>Other Host</i> is the Inside Hosts host group. The goal is to find the standout peers who received the majority of the uploaded data. After finding the standout peers, determine whether or not data being received by these peers in the observed quantity is expected behavior.
Non-standard flow data required? (FlowSensor, proxy, firewall, etc.)	N/A
Notes	N/A

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	The event must be enabled for Inside Host host group.
Which policies have this event on by default?	Inside Hosts, Outside Hosts
Which policies have this event off by default?	Firewalls, Proxies, & NAT Devices
Which policies have this alarm on by default?	None
Which policies have this alarm off by default?	Inside Hosts; Outside Hosts; Firewalls, Proxies, & NAT Devices
Can you tune this event?	Yes
Is the event variance-based, threshold-based, or other?	Variance or threshold-based.
What are default values and units?	
Tolerance	92
Minimum threshold value	500M client payload bytes in 24 hours.
Maximum threshold value	1T downloaded payload bytes in 24 hours.
Is event tunable with non-variance-based parameters?	Yes
If so, what are its tunable attributes and units?	The event can be set to trigger off a particular number of bytes that does not take baselines or tolerance into account.

Questions about event	Response
Is event tunable outside of the normal policy editor	No
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	No
If so, what is it?	N/A

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The inside host that is sending the payload data.
What is the target?	The inside host or hosts that are receiving the payload data.
Which policy causes the event to trigger?	The source's policy.
What alarm categories does this security event contribute to for the source?	
In what quantity?	
What alarm categories does this security event contribute to for the target?	High Data Hoarding Index, High Target Index
In what quantity?	• High Data Hoarding Index: True • TI: True

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: Observed W bytes. Expected X bytes, tolerance of Y allows up to Z bytes. Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the target. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by target IP vs. Inside Hosts, the start time (reset hour) of the event's active day until the event's last active time, server bytes greater than or equal to 1, and total bytes greater than or equal to 1k. Desktop Client: Flows are filtered by time period of last five minutes.


What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_TARGET_DATA_HOARD (316)

Timeout/tcp

Mitigation is not available for the alarm associated with this security event.

Timeout/tcp



Resulting potential security event entry:

Source Host Groups ^{A1}	Source Host ^B	Target Host Groups ^B	Target Host ^{A3}	Concern Index ^{A2}	Security Events ^B
Singapore	203.125.234.84	Lancpe Corporate	50.56.217.21	<CI value>* ^{A2}	Timeout/tcp-520(10)

* <CI value> represents a numerical point value that Secure Network Analytics assigns to an event based on various conditions.

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	The source host has sent a TCP packet in which the TTL was less than 2. Because tracerouting occurs over UDP, TCP packets with short TTLs usually indicate malicious activity (i.e., firewalking) or a broken network (i.e., routing loops).
What does it mean when it triggers?	N/A
What are next steps?	N/A
Non-standard flow data required? (FlowSensor, proxy, firewall, etc.)	For analyzing NetFlow, a FlowSensor is required.
Notes	none

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	
Which policies have this event on by default?	Inside Hosts, Outside Hosts
Which policies have this event off by default?	
Which policies have this alarm on by default?	
Which policies have this alarm off by default?	Inside Hosts, Outside Hosts
Can you tune this event?	No
Is the event variance-based, threshold-based, or other?	N/A
What are default values and units?	N/A
Tolerance	N/A
Minimum threshold value	N/A
Maximum threshold value	N/A
Is event tunable with non-variance-based parameters?	N/A
If so, what are its tunable attributes and units?	N/A

Questions about event	Response
Is event tunable outside of the normal policy editor	N/A
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	
If so, what is it?	

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The host that is the packet's source
What is the target?	The host that is the packet's destination
Which policy causes the event to trigger?	
What alarm categories does this security event contribute to for the source?	
In what quantity?	• High Recon Index: True • CI: True
What alarm categories does this security event contribute to for the target?	
In what quantity?	• TI: True

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: View details. Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source & target IP, a time range five minutes before the event's first active time until the event's last active time, the associated port as <port>/TCP, and TCP. Desktop Client: Flows are filtered by time period of last five minutes.


What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_TCP_TO (260)

Timeout/udp

Mitigation is not available for the alarm associated with this security event.

Timeout/udp



Resulting potential security event entry:

Source Host Groups ^{^1}	Source Host ^{^2}	Target Host Groups ^{^3}	Target Host ^{^4}	Concern Index ^{^5}	Security Events ^{^6}
Singapore	203.125.234.84	Lancope Corporate	50.56.217.21	<CI value>*	Timeout/udp-520(10)

* <CI value> represents a numerical point value that Secure Network Analytics assigns to an event based on various conditions.

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	The source host has sent a UDP packet in which the TTL was less than 2. Because tracerouting occurs over UDP, this security event can be fairly common; however, high numbers tend to imply a broken network (i.e., routing loops). Port 38293 is excluded from the UDP timeout to accommodate Norton Antivirus servers.
What does it mean when it triggers?	N/A
What are next steps?	N/A
Non-standard flow data required? (FlowSensor, proxy, firewall, etc.)	To analyze NetFlow, you need a Flow Sensor.
Notes	None

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	
Which policies have this event on by default?	Inside Host, Outside Hosts
Which policies have this event off by default?	
Which policies have this alarm on by default?	
Which policies have this alarm off by default?	Inside Host, Outside Hosts
Can you tune this event?	No
Is the event variance-based, threshold-based, or other?	N/A
What are default values and units?	N/A
Tolerance	N/A
Minimum threshold value	N/A
Maximum threshold value	N/A
Is event tunable with non-variance-based parameters?	N/A
If so, what are its tunable attributes and units?	N/A

Questions about event	Response
Is event tunable outside of the normal policy editor	N/A
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	
If so, what is it?	

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The host that is the packet's source
What is the target?	The host that is the packet's destination
Which policy causes the event to trigger?	
What alarm categories does this security event contribute to for the source?	
In what quantity?	• High Recon Index: True • CI: True
What alarm categories does this security event contribute to for the target?	
In what quantity?	• TI: True

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: View details. Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source & target IP, a time range one minute before the event's first active time until the event's last active time, the associated port as <port>/UDP, and UDP. Desktop Client: Flows are filtered by time period of last five minutes.

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_UDP_TO (259)

Touched

The alarmed host (a host with a [High Concern Index](#) alarm or a [Trapped Host](#) security event) has initiated a connection with the target host and exchanged data.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	
In what quantity?	
What alarm categories does this security event contribute to for the target?	Attack Index, High Target Index
In what quantity?	• Attack Index: True • TI: True

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_TOUCHED (28)

Trapped Host

The host exceeded the accepted daily threshold of attempts to communicate with a host in a Trap host group on a port that is not in the configured services, which indicates potential “low and slow” scanning activity. The system performs these types of scans in order to prevent discovery of the attacker, the ports being scanned, or the hosts being scanned, or even to run a preliminary scan before the release of a worm.

Select the “Trap hosts that scan unused addresses in this group” check box on the Host Group Management page for a host group to enable the tracking of this type of activity for that particular host group. This is an advanced feature that you should enable only for specific, small, well-controlled host groups with fixed IP addresses and a stable number of hosts, such as critical servers. Doing so allows you to detect hosts intruding into the most important areas of your network.

Hosts that scan unused addresses in the Outside Hosts host group are not included in trapped host calculations and therefore do not generate trapped host alerts and alarms, regardless of whether this setting is enabled in the Outside Hosts host group.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	High Recon Index, High Concern Index
In what quantity?	• High Recon Index: True • CI: True
What alarm categories does this security event contribute to for the target?	High Target Index

Questions about event	Response
In what quantity?	• TI: True

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_TRAPPED_HOST (34)

UDP Flood

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	The source host has sent an excessive number of UDP packets in a 5-minute period. This may indicate a Denial of Service (DoS) attack or non-stealthy scanning activity.
What does it mean when it triggers?	This indicates that the source host is engaged in a DoS attack (potentially caused by either a compromised host or driven by a user), a misconfigured network application, or a connection with abnormally high UDP packets per second from the source. Many packets are often sent to consume a target's bandwidth as part of larger DDoS attacks. However, high UDP packet rates are more typical than high SYN packet rates, so a low threshold or tolerance setting may cause a high volume of false positives. Relatively low packet counts relative to a large amount of IPs is also likely to be a sign of reconnaissance.
What are next steps?	Determine how many UDP packets are being sent at what rate, when they were sent, and where they are going. An ideal way to investigate this is to perform a flow query for the source host. Set the start time of the query to five minutes before the start time of the associated security event, and set the end time to the end time of the alarm (if there is an alarm). Set the protocol filter to UDP only. Once the results have returned, find where the most UDP packets went during the searched time period. If the subject of the event is the client host in its UDP-heavy flows, you can sort on Client Packets or Client

Questions about this event	Response
	Packet Rate (pps). The goal is to find the standout flow (s). After finding the standout flow(s), determine whether or not this is expected behavior. Examples of expected behavior is a UDP-based VPN connection or uploads to a web server using UDP (like Google often does).
Non-standard flow data required? (FlowSensor, proxy, firewall, etc.)	None
Notes	N/A

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	
Which policies have this event on by default?	Inside Hosts, Outside Hosts
Which policies have this event off by default?	
Which policies have this alarm on by default?	Inside Hosts, Outside Hosts
Which policies have this alarm off by default?	

Questions about event	Response
Can you tune this event?	Yes
Is the event variance-based, threshold-based, or other?	Both
What are default values and units?	N/A
Tolerance	75
Minimum threshold value	3.6k UDP Packets in 5 minutes
Maximum threshold value	10M UDP Packets in 5 minutes
Is event tunable with non-variance-based parameters?	N/A
If so, what are its tunable attributes and units?	N/A
Is event tunable outside of the normal policy editor	N/A
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	
If so, what is it?	

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The host sending UDP packets

Questions about event	Response
What is the target?	None
Which policy causes the event to trigger?	
What alarm categories does this security event contribute to for the source?	
In what quantity?	• High Ddos Source Index: True • CI: True
What alarm categories does this security event contribute to for the target?	
In what quantity?	

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: Observed W pp5m. Expected X pp5m, tolerance of Y allows up to Z pp5m. Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source IP, a time range five minutes before the event's first active time until the event's last active time, and UDP. Desktop Client: Flows are filtered by time period of last five minutes.

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_UDP_FLOOD (6)

UDP Received

What is the security event?

Questions about this event	Response
What does it mean when it triggers?	This indicates that the target hosted is being targeted in a DoS attack, a misconfigured network application, or a connection with abnormally high UDP packets per second from the source(s). Many packets are often sent to consume a target's bandwidth as part of larger DDoS attacks. However, high UDP packet rates are more typical than high SYN packet rates, so a low threshold or tolerance setting may cause a high volume of false positives.
What are next steps?	The goal is to determine how many UDP packets are being received at what rate, when they were received, and where they are coming from. An ideal way to investigate this is to perform a flow query for the target host. Set the start time of the query to five minutes before the start time of the associated security event, and set the end time to the end time of the alarm (if there is an alarm). Set the protocol filter to <i>UDP only</i>. Once the results have returned, find where the most UDP packets came from during the searched time period. If the target of the event is the server host in its UDP-heavy flows, you can sort on either Client Packets or Client Packet Rate (pps). The goal is to find the standout flow(s). After finding the standout flow(s), determine whether or not this is expected behavior. An example of expected behavior is a UDP-based VPN connection.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	
In what quantity?	
What alarm categories does this security event contribute to for the target?	High Ddos Target Index, High Target Index
In what quantity?	• High Ddos Target Index: True • TI: True

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_UDP_RECEIVED (49)

Watch Host Active

What is the security event?

Questions about this event	Response
What does it mean when it triggers?	This is a user configured event that notifies you when a particular host is observed communicating. The meaning of the event can vary depending on why a user decided to put a particular host on the watch list, but it is generally a sign that any communication to or from the monitored host is inappropriate.
What are next steps?	The investigatory steps for this security event are entirely dependent on the context in which the host was added to the watch list.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	High Policy Violation Index, High Concern Index
In what quantity?	• High Policy Violation Index: True • CI: True
What alarm categories does this security event contribute	

Questions about event	Response
to for the target?	
In what quantity?	

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_WATCH_LIST (31)

Watch Port Active

What is the security event?

Questions about this event	Response
What does it mean when it triggers?	This is a user configured event that notifies you when a particular host is observed communicating over a particular port. The meaning of the event can vary depending on why a user decided to put a particular host on the watch list, but it is generally a sign that any communication to or from the monitored host is inappropriate.
What are next steps?	The investigatory steps for this security event are entirely dependent on the context in which the port was added to the watch list.

What policy settings are available for this security event?

Questions about event	Response
Is the event variance-based, threshold-based, or other?	

How does this security event contribute to categories?

Questions about event	Response
What alarm categories does this security event contribute to for the source?	High Policy Violation Index, High Concern Index
In what quantity?	• High Policy Violation Index: True • CI: True
What alarm categories does	

Questions about event	Response
this security event contribute to for the target?	
In what quantity?	

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_WATCH_PORT (13)

Worm Activity

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	The host has scanned and connected on a particular port across more than one subnet.
What does it mean when it triggers?	This security event indicates that a host appears to be performing an excess amount of reconnaissance across various internal networks, which may indicate that the host is compromised and attempting to propagate an infection throughout a network.
What are next steps?	This event can be noisy and its success largely depends on how well network scanners have been identified and placed into the network scanners host group. By doing this, these devices will inherit a policy that disables worm activity. Worm activity largely hinges upon a host scanning for similar ports across different logical /24 ranges. Begin by running a Top Ports (outbound) report. Set the time period to the day of the event and the client host to be the source IP of the event. Regardless of the target IP range that is listed, you can scan any type of host, so determine if you want to search on inside hosts, outside hosts, or both. On the Hosts tab in the Filter dialog, set the Server filter as appropriate, and on the Advanced tab, set the "Order the records returned by" to Flows. Sort the results by peers. Regardless, start at the top of the list as sorted by flows or peers to find the standout ports that either do not belong or appear abnormally high. Right-click an IP address to pivot to the flows to view the various IP addresses and determine if particular host groups are those being primarily targeted. You can also right-click and pivot to the Top Peers report to determine whether or not the traffic is spread generally equally across the targets. It is also

Questions about this event	Response
	worth noting the percentage of hosts that responded to the scans. At this point you should be able to determine who the host was scanning and what ports were being scanned.
Non-standard flow data required? (FlowSensor, proxy, firewall, etc.)	
Notes	None

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	
Which policies have this event on by default?	Outside Hosts
Which policies have this event off by default?	
Which policies have this alarm on by default?	Outside Hosts
Which policies have this alarm off by default?	
Can you tune this event?	No
Is the event variance-based, threshold-based, or other?	N/A

Questions about event	Response
What are default values and units?	N/A
Tolerance	N/A
Minimum threshold value	N/A
Maximum threshold value	N/A
Is event tunable with non-variance-based parameters?	N/A
If so, what are its tunable attributes and units?	N/A
Is event tunable outside of the normal policy editor	N/A
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	
If so, what is it?	

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The host that is the worm
What is the target?	The host that is the victim of the worm
Which policy causes the event to trigger?	

Questions about event	Response
What alarm categories does this security event contribute to for the source?	
In what quantity?	• Attack Index: True • CI: True
What alarm categories does this security event contribute to for the target?	
In what quantity?	

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: Worm activity on port X ServiceName (Y Protocol). Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source & target IP, the start time (reset hour) of the event's active day until the event's last active time, and the associated port as <port>/TCP & <port>/UDP. Desktop Client: Flows are filtered by time period of last five minutes.

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_WORM_ACTIVITY (35)

Worm Propagation

What is the security event?

Questions about this event	Response
What type of behavior is this event looking for?	The host has scanned and connected on a particular port across more than one subnet, and the host was previously scanned and connected to by a host for which the Worm Activity alarm has been triggered.
What does it mean when it triggers?	This event seeks to find hosts that have been compromised by another compromised host, which results in the host attempting to compromise other hosts. This involves the following sequence of events: 1) A host is scanned by another host that is in the process of scanning multiple hosts; 2) The host that was just scanned now starts to scan other hosts.
What are next steps?	Run a Top Peers report on the host in question. Examine what other peer hosts contacted the subject on the specified port, filtering the report by the port in question. Verify that these hosts are not known scanners and should not have been participating in the observed scanning activity. If they are known scanners, the risk may be reduced. If they are not known scanners, the risk may increase since the scan was not legitimate or sanctioned. In this case, the event may indicate that the host is compromised.
Non-standard flow data required? (FlowSensor, proxy, firewall, etc.)	None
Notes	None

What policy settings are available for this security event?

Questions about event	Response
What policy settings are required for this security event to trigger?	
Which policies have this event on by default?	Inside Hosts, Outside Hosts
Which policies have this event off by default?	
Which policies have this alarm on by default?	Inside Hosts, Outside Hosts
Which policies have this alarm off by default?	
Can you tune this event?	No
Is the event variance-based, threshold-based, or other?	N/A
What are default values and units?	N/A
Tolerance	N/A
Minimum threshold value	N/A
Maximum threshold value	N/A
Is event tunable with non-variance-based parameters?	N/A
If so, what are its tunable attributes and units?	N/A

Questions about event	Response
Is event tunable outside of the normal policy editor	N/A
If so, what are the alternate locations of tunable values?	N/A
Does event have a default mitigation?	
If so, what is it?	

How does this security event contribute to categories?

Questions about event	Response
What is the source of the event?	The host that is the worm
What is the target?	The host that is the victim of the worm
Which policy causes the event to trigger?	
What alarm categories does this security event contribute to for the source?	
In what quantity?	• Attack Index: True • CI: True
What alarm categories does this security event contribute to for the target?	
In what quantity?	• TI: True

What information is available in Secure Network Analytics?

Questions about event	Response
What information is displayed in the alarm details?	Manager: Worm propagated X direction Source Host using Y ServiceName (Z Protocol). Desktop Client: N/A
What happens when you click the alarm details?	Manager: Displays the security events for the source host. Desktop Client: N/A
What information is shown for associated flows?	Manager: Flows are filtered by source & target IP, the start time (reset hour) of the event's active day until the event's last active time, and client bytes & server bytes (both of which are greater than or equal to 1). Desktop Client: Flows are filtered by time period of last five minutes.

What information is available for responses for this security event?

Questions about event	Response
What is the event ID?	SEC_ID_WORM_PROPAGATION (36)

Investigating Today's Active Alarms

Use the Alarm List to investigate alarms that have been active on your network today (i.e., for the period since the last archive hour, which by default occurs at 12:00 a.m.). By default, information is sorted based on the most recent alarm, whether for security events or alarm categories. If a column has arrows in the heading, you can sort on that column. The columns, however, cannot be changed.

The number of hosts represents each host. If a specific host is routed through two Flow Collectors, for example, it will appear twice in the alarm list on this page, but only once in the alarm category's host count on the Dashboard.

This page provides information similar to the information available in the Alarm Table on the Desktop Client. It also shows whether an alarm has been acknowledged. If you need to acknowledge an alarm, you must do so in the Desktop Client.



You must install an identity service in order to see data in the Source User column.

How to find this page

Do the following to access this page

On the Security Insight Dashboard, in either the Alarms by Type or Today's Alarms component, click the section for the specific type of alarm you want to investigate.

The Alarm List opens filtered by the type of alarm you selected.

The information on this page is read-only. You cannot acknowledge or close an alarm. The Source and Target columns provide links to the Host Report page, and the Source User column provides links to the User Entity View page.

Concepts you need to understand to use this page

Refer to the following parameters when viewing information on this page

- The severity column only appears if the alarms are for a category, such as Recon or Anomaly. It will not appear if the alarms are associated with a security event.
- The severity column is sorted in a manner similar to that of the Host List or the Alarm List. "Always bad" alarms are at the top. Remaining hosts are sorted by severity of alarms, then percent of threshold, and finally by IP address.

- If alarms are occurring for a host group, the column for each affected alarm category displays one of the following colors:
 - Black – Critical, or “always bad.”
 - Red – Major
 - Orange – Minor
 - Yellow – Trivial
- You should always investigate alarms with the severity level of Critical.
- If a host does not have active alarms, it will not have a severity color assigned to it in the alarm category columns.
- The color of the column indicates the presence of active alarms for a specific host, but does not indicate the type of security event generating the alarm or the volume of alarms. For additional information, click the alarm category column for the host you want to examine to view the Alarm List page in the Desktop Client.
- The numerical value in each cell shows the percent of threshold for a specific alarm category and host.



For more details about alarm severities, see [About Alarm Severities](#).

What you can do on this page

View the following information on this page

- The severity of alarms.
- What type of alarm has occurred?
- Based on the First Active and Last Active columns, how long ago did the alarm start? How long did it last? Do either of these convey a sense of concern or urgency?
- Which host or host group did the alarm affect? Is that normal for your network, or is it a source of concern?
- Does the alarm for the type of security event cause concern? If so, in the Details column, click the link to access the [Security Events Details](#) page for more information.
- Is a user associated with the alarm? If so, does that raise questions or concerns?

What you can do next

View associated policies

To view or tune the policies that have caused an alarm to trigger, do the following:

1. Identity the type of policy and event for which the alarm has triggered.
2. Click the **⋮ (Ellipsis)** icon > **Tune Alarm** for any of the following:
 - An IP address in the Source or Target column
 - The applicable Actions column

If you pivot from the ...	Then your query filters on the ...	And this page opens...	Filtered by the...
Source or Target	IP address	Policy Management: Core Events tab	All effective policies for the IP address
Actions column, and the security event or category (displayed at the top of the page) is a <i>core</i>	security event or category	Policy Management: Core Events tab	Event, and then Policy Name

If you pivot from the ...	Then your query filters on the ...	And this page opens...	Filtered by the...
event			
Actions column, and the security event or category (displayed at the top of the page) is a <i>custom</i> security event	security event or category	Policy Management: Custom Events tab	Event

Investigate a host that is the source of an alarm

1. Click the IP address in the Source column.
2. Evaluate the information on the [Host Report](#).
3. If necessary, in the Actions column for the applicable host, click the ... (Ellipsis) icon > [View Flows](#). Click Edit if you first want to make edits to the Flow Search page.

Investigate a user who is the source of an alarm

1. Click the user name in the Source User column.
2. Evaluate the information on the [User Report](#).

Investigate a host that is the target of an alarm

1. Click the IP address in the Target column.
2. Evaluate the information on the [Host Report](#).
3. If necessary, in the Actions column for the applicable host, click the **Ellipsis** icon > [View Flows](#). Click Edit if you first want to make edits to the Flow Search page.

Investigate a security event for a particular host

1. Click the link in the Details column.
2. Evaluate the information on the [Security Event Details](#) page.
3. If necessary, use the links to further investigate a source or target host.

 To close or acknowledge an alarm, use the Desktop Client. Refer to the Desktop Client online help for more information.

Investigate flows related to a particular host or alarm

To run a flow search associated with a row of data, in the Actions column, click the **Ellipsis** icon to access the relevant options.

- To run a flow search associated with a single IP address, click the context menu next to the applicable IP address to access the relevant options. The flow filters on the host, start date, and start time.
- To run a flow search associated with a row of data, in the Actions column, click the **Ellipsis** icon to access the relevant options. The flow filters on the alarm, start date, and start time.

How the start date and time on which the flow filters are calculated:

Start date and time

If the time listed in the First Active column is...	Then the start date filtered on is the...	And the start time filtered on is ...
1. Before the last reset hour	day before the date listed in the First Active column	Last reset hour

If the time listed in the First Active column is...	Then the start date filtered on is the...	And the start time filtered on is ...
2. After the last reset hour	date listed in the First Active column	Last reset hour

Examples: For the following examples, assume today's date is April 26, 2018 and the reset hour is 12:00 a.m.

- The First Active date and time are 4/24/18, 9:00 p.m. | The Last Active date and time are 4/26/18, 2:00 p.m.
 - The flow will **filter on** 4/23/18 as the *start date* and 12:00 a.m. as the *start time*.
- The First Active date and time are 4/26/18, 9:00 p.m. | The Last Active date and time are 4/26/18, 2:00 p.m.
 - Then the flow will **filter on** 4/26/18 as the *start date* and 12:00 a.m. as the *start time*.

End date and time

If ...	And the Last Active time is...	Then the end date filtered on is the...	And the end time filtered on is ...
3. The day after the date listed in the First Active column is <i>before</i> the Last Active date	Before the current time	Day after the First Active date	24 hours after the First Active time.
4. The day listed in the First Active column is today or yesterday	Current	Last Active date	Current

Examples: For the following examples, assume today's date is April 26, 2018 and the current time is 4:00 p.m.

3. The First Active date and time are 4/25/18, 9:00 p.m. | The Last Active date and time are 4/26/18, 2:00 p.m.
 - The flow will **filter on** 4/26/18 as the *end date* and 9:00 p.m. as the *end time*.
4. The First Active date and time are 4/26/18, 9:00 p.m. | The Last Active date and time are 4/26/18, 2:00 p.m.
 - Then the flow will **filter on** 4/26/18 as the *end date* and 4:00 p.m. as the *end time*.

View reports that will give you more in-depth information about specific areas of interest

- To [run a top report](#) search associated with a single IP address, click the **Ellipsis** icon next to the applicable IP address to access the relevant options.
- To run a top report search associated with a row of data, in the Actions column, click the Ellipsis icon to access the relevant options.

Related topics

[About Alarm Categories](#)

[Security Event List](#)

[Investigating a Single Host](#)

[Investigating a Single User](#)

Using Alarm Dashboard to Assess Potential Threats

Use the information on this page to investigate alarms that are currently contributing to a specific alarm category for a specific host. From here you can quickly assess the possibility of an attack or other threats against your network and devices.

If a single host's flows are detected by multiple Flow Collectors, you could see duplicate alarms. The number of alarms listed on this page may not match the number of hosts shown in the alarm category on the dashboard.

The number of hosts represents each host that is currently receiving alarms; if a specific host is routed through two Flow Collectors, for example, it appears twice in the alarm list on this page, but only once in the alarm category's host count on the dashboard.



- This page provides information similar to the information available in the Alarm Table on the [Desktop Client](#).
- You must have an identity service installed if you want the Source User column to display data.

How to find this page

Do one of the following to access this page

On the [Host Report](#), do one of the following:

- Click the number in the corresponding alarm category at the top of the page.

The Alarm dashboard opens, displaying all the alarms for that category that are currently active for that host.

- In the Alarms by Type component, click on a bar for the alarm and date for which you want to investigate.

The Alarm dashboard opens, filtered by the alarm you selected.

On the [Host List](#), click the alarm category column on the row for the applicable host.

The Alarm dashboard for that host opens.

What you can view on this page

The information displayed on the Alarm dashboard is based on the alarm category you clicked within on the Security Insight dashboard. The information that is shown may differ among various Alarm dashboards.

Use the Alarm dashboard to answer the following questions for each alarm category.

Anomaly

Use the Anomaly Alarm dashboard to evaluate the following:

- Based on the First Active column, how long ago did the alarm start? Does this convey a sense of concern or urgency?
- Which host or host group did the alarm affect? Is that normal for your network, or is it a source of concern?
- Does the alarm for the type of security event cause concern?
- Is a policy assigned to the host group, and if so, is it one that is appropriate for the host group?
- Are there any alarms for security events? If so, does the type of security event seem suspicious for the associated host?
- Is a user associated with the alarm? If so, does that raise questions or concerns?
- Information in the Details column compares Anomaly (AN) points associated with the active alarm to the points allowed in the policy assigned to the host producing the alarm. Is the volume of points actually occurring a source of concern? Should the level of points allowed in the policy be changed?

Command & Control (C&C)

Use the C&C Alarm dashboard to evaluate the following:

- Based on the column, how long ago did the alarm start? Does this convey a sense of concern or urgency?
- Which host or host group did the alarm affect? Is that normal for your network, or is it a source of concern?
- Does the alarm for the type of security event cause concern?
- Is a policy assigned to the host group, and if so, is it one that is appropriate for the host group?
- Are there any alarms for security events? If so, does the type of security event seem suspicious for the associated host?
- Is a user associated with the alarm? If so, does that raise questions or concerns?
- Information in the Details column compares C&C points associated with the active alarm to the points allowed in the policy assigned to the host producing the alarm. Is

the volume of points actually occurring a source of concern? Should the level of points allowed in the policy be changed?



If either the Bot Infected Host - Attempted C&C Activity or Bot Infected Host - Successful C&C Activity alarm occurs, click the event name in the Event Alarms column to access information about a confidence rating. The rating indicates Cisco's confidence that *[name of entry]* is classified appropriately.

Concern Index

Use the Concern Index Alarm dashboard to evaluate the following:

- Based on the column, how long ago did the alarm start? Does this convey a sense of concern or urgency?
- Which host or host group did the alarm affect? Is that normal for your network, or is it a source of concern?
- Does the alarm for the type of security event cause concern?
- Is a policy assigned to the host group, and if so, is it one that is appropriate for the host group?
- Are there any alarms for security events? If so, does the type of security event seem suspicious for the associated host?
- Is a user associated with the alarm? If so, does that raise questions or concerns?
- Information in the Details column compares Concern Index (CI) points associated with the active alarm to the points allowed in the policy assigned to the host producing the alarm. Is the volume of points actually occurring a source of concern? Should the level of points allowed in the policy be changed?

Data Hoarding

Use the Data Hoarding Alarm dashboard to evaluate the following:

- Based on the column, how long ago did the alarm start? Does this convey a sense of concern or urgency?
- Which host or host group did the alarm affect? Is that normal for your network, or is it a source of concern?
- Is a policy assigned to the host group, and if so, is it one that is appropriate for the host group?
- Are there any alarms for security events? If so, does the type of security event seem suspicious for the associated host?

- Is a user associated with the alarm? If so, does that raise questions or concerns?
- Information in the Details column compares Data Hoarding (DH) points associated with the active alarm to the points allowed in the policy assigned to the host producing the alarm. Is the volume of points actually occurring a source of concern? Should the level of points allowed in the policy be changed?

DDoS Source

Use the DDoS Source Alarm dashboard to evaluate the following:

Application Traffic

- Which top 10 applications have incurred the most inbound and outbound traffic within the last two hours, in five-minute increments?
- Are there any spikes or anomalies in your traffic?

Refer to the following guidelines:

- Each application type is represented by a different color.
- When you hover over the graph, a window appears listing each application represented in the graph.
- Information shows the volume of both inbound and outbound traffic carried by each application.
- You can drag the "handles" on the graph to change the time frame of the data. Click **Reset zoom** to restore the default view of the graph.
- This graph updates every five minutes.
- Traffic traveling from the Internet to Inside Hosts appears above the zero line.
- Traffic traveling from Inside Hosts to the Internet appears below the zero line.
- When you place your cursor over an area in the graph, a pop-up appears. It displays, in bps, the inbound and outbound traffic for all top ten applications for that time frame.

Summary of Sources

- How many hosts have triggered DDoS Source alarms that are still active and have occurred within the last hour?
- Which hosts have triggered multiple DDoS Source alarms?
- What is the IP address for a particular source host?

This chart updates every five minutes.

Alarms Table

When you access the Alarm dashboard by clicking ...	This table displays ...
Within the DDoS Source alarm category on the Security Insight dashboard	All active DDoS Source alarms.
In any other section on the Security Insight dashboard	All active and inactive alarms that have occurred only since the last reset hour.

- Based on the First Active column, how long ago did the alarm start? Does this convey a sense of concern or urgency?
- Which host or host group did the alarm affect? Is that normal for your network, or is it a source of concern?
- Is a policy assigned to the host group, and if so, is it one that is appropriate for the host group?
- Are there any alarms for security events? If so, does the type of security event seem suspicious for the associated host?
- Is a user associated with the alarm? If so, does that raise questions or concerns?
- Information in the Details column compares DDOS Source (DS) points associated with the active alarm to the points allowed in the policy assigned to the host producing the alarm. Is the volume of points actually occurring a source of concern? Should the level of points allowed in the policy be changed?

DDoS Target

Use the DDoS Target Alarm dashboard to evaluate the following:

Application Traffic

- Which top 10 applications have incurred the most inbound and outbound traffic within the last two hours, in five-minute increments?
- Are there any spikes or anomalies in your traffic?

Refer to the following guidelines:

- Each application type is represented by a different color.
- When you hover over the graph, a window opens listing each application represented in the graph.
- Information shows the volume of both inbound and outbound traffic carried by each application.
- You can drag the "handles" on the graph to change the time frame of the data. Click **Reset zoom** to restore the default view of the graph.
- This graph updates every five minutes.
- Traffic traveling from the Internet to Inside Hosts appears above the zero line.
- Traffic traveling from Inside Hosts to the Internet appears below the zero line.
- When you place your cursor over an area in the graph, a pop-up appears. It displays, in bps, the inbound and outbound traffic for all top ten applications for that time frame.

Summary of Targets

- How many hosts have triggered DDoS Target alarms that are still active and have occurred within the last hour?
- Which hosts have triggered multiple DDoS Target alarms?
- What is the IP address for a particular target host?

This chart updates every five minutes.

Alarms Table

When you access the Alarm dashboard by clicking ...	This table displays ...
Within the DDoS Target alarm category on the Security Insight dashboard	All active DDoS Target alarms.
In any other section on the Security Insight dashboard	All active and inactive alarms that have occurred only since the last reset hour.

- Based on the First Active column, how long ago did the alarm start? Does this convey a sense of concern or urgency?

- Which host or host group did the alarm affect? Is that normal for your network, or is it a source of concern?
- Is a policy assigned to the host group, and if so, is it one that is appropriate for the host group?
- Are there any alarms for security events? If so, does the type of security event seem suspicious for the associated host?
- Is a user associated with the alarm? If so, does that raise questions or concerns?
- Information in the Details column compares DDOS Target (DT) points associated with the active alarm to the points allowed in the policy assigned to the host producing the alarm. Is the volume of points actually occurring a source of concern? Should the level of points allowed in the policy be changed?

Exfiltration

Use the Exfiltration Alarm dashboard to evaluate the following:

- Based on the column, how long ago did the alarm start? Does this convey a sense of concern or urgency?
- Which host or host group did the alarm affect? Is that normal for your network, or is it a source of concern?
- Is a policy assigned to the host group, and if so, is it one that is appropriate for the host group?
- Are there any alarms for security events? If so, does the type of security event seem suspicious for the associated host?
- Is a user associated with the alarm? If so, does that raise questions or concerns?
- Information in the Details column compares Exfiltration (EX) points associated with the active alarm to the points allowed in the policy assigned to the host producing the alarm. Is the volume of points actually occurring a source of concern? Should the level of points allowed in the policy be changed?

Exploitation

Use the Exploitation Alarm dashboard to evaluate the following:

- Based on the column, how long ago did the alarm start? Does this convey a sense of concern or urgency?
- Which host or host group did the alarm affect? Is that normal for your network, or is it a source of concern?

- Is a policy assigned to the host group, and if so, is it one that is appropriate for the host group?
- Are there any alarms for security events? If so, does the type of security event seem suspicious for the associated host?
- Is a user associated with the alarm? If so, does that raise questions or concerns?
- Information in the Details column compares Exploitation (EP) points associated with the active alarm to the points allowed in the policy assigned to the host producing the alarm. Is the volume of points actually occurring a source of concern? Should the level of points allowed in the policy be changed?

Policy Violation

Use the Policy Violation Alarm dashboard to evaluate the following:

- Based on the column, how long ago did the alarm start? Does this convey a sense of concern or urgency?
- Which host or host group did the alarm affect? Is that normal for your network, or is it a source of concern?
- Is a policy assigned to the host group, and if so, is it one that is appropriate for the host group?
- Are there any alarms for security events? If so, does the type of security event seem suspicious for the associated host?
- Is a user associated with the alarm? If so, does that raise questions or concerns?
- Information in the Details column shows the Policy Violation (PV) points associated with the active alarm. Is the volume of points actually occurring a source of concern? Should the level of points allowed in the policy be changed?

Recon

Use the Recon Alarm dashboard to evaluate the following:

- Based on the column, how long ago did the alarm start? Does this convey a sense of concern or urgency?
- Which host or host group did the alarm affect? Is that normal for your network, or is it a source of concern?
- Is a policy assigned to the host group, and if so, is it one that is appropriate for the host group?
- Are there any alarms for security events? If so, does the type of security event seem suspicious for the associated host?

- Is a user associated with the alarm? If so, does that raise questions or concerns?
- Information in the Details column compares Recon (RC) points associated with the active alarm to the points allowed in the policy assigned to the host producing the alarm. Is the volume of points actually occurring a source of concern? Should the level of points allowed in the policy be changed?

Target Index

Use the Target Index Alarm dashboard to evaluate the following:

- Based on the column, how long ago did the alarm start? Does this convey a sense of concern or urgency?
- Which host or host group did the alarm affect? Is that normal for your network, or is it a source of concern?
- Is a policy assigned to the host group, and if so, is it one that is appropriate for the host group?
- Are there any alarms for security events? If so, does the type of security event seem suspicious for the associated host?
- Is a user associated with the alarm? If so, does that raise questions or concerns?
- Information in the Details column compares Target Index (TI) points associated with the active alarm to the points allowed in the policy assigned to the host producing the alarm. Is the volume of points actually occurring a source of concern? Should the level of points allowed in the policy be changed?

What you can do next

View associated policies

To view or tune the policies that have caused an alarm to trigger, do the following:

1. Identity the type of policy and event for which the alarm has triggered.
2. Click **Tune Alarm** from the menu for any of the following:
 - An IP address in the Source or Target column
 - The applicable Actions column

If you pivot from the ...	Then your query filters on the ...	And this page opens...	Filtered by the...
Source or Target	IP address	Policy Management: Core Events tab	All effective policies for the IP address
Actions column, and the security event or category (displayed at the top of the page) is a <i>core</i> event	security event or category	Policy Management: Core Events tab	Event, and then Policy Name
Actions column, and the security event or category (display	security event or category	Policy Management: Custom Events tab	Event

If you pivot from the ...	Then your query filters on the ...	And this page opens...	Filtered by the...
ed at the top of the page) is a <i>custom</i> security event			

Investigate a host that is the source of an alarm

1. Click the IP address in the Source column.
2. Evaluate the information on the [Host Report](#).
3. If necessary, in the Actions column for the applicable host, click the **...** (**Ellipsis**) icon > [View Flows](#). Click **Edit** if you first want to make edits to the Flow Search page.

Investigate a user who is the source of an alarm

1. Click the user name in the Source User column.
2. Evaluate the information on the [User Report](#).

Investigate a host that is the target of an alarm

1. Click the IP address in the Target column.
2. Evaluate the information on the [Host Report](#).
3. If necessary, in the Actions column for the applicable host, click the **Ellipsis** icon > [View Flows](#). Click Edit if you first want to make edits to the Flow Search page.

Investigate a security event for a particular host

1. Click the link in the Details column.
2. Evaluate the information on the [Security Event Details](#) page.
3. If necessary, use the links to further investigate a source or target host.

Investigate the security event alarms for the day

1. Click the link in the Event Alarm column, if one is available.
2. Evaluate the information on the [Event Alarms](#) page.
3. If necessary, use the links to investigate further.

To close or acknowledge an alarm, use the Desktop Client. Refer to the Desktop Client Help for more information.

Run a flow search

- To run a flow search associated with a single IP address, click the menu next to the applicable IP address to access the relevant options.
- To run a flow search associated with a row of data, click the menu in the Actions column to access the relevant options.

Run a top report search

- To [run a top report](#) search associated with a single IP address, click the **Ellipsis** icon next to the applicable IP address to access the relevant options.
- To run a top report search associated with a row of data, in the Actions column, click the Ellipsis icon to access the relevant options.

Related Topics

[About Alarm Categories](#)

[Security Event List](#)

[Examining Details of Security Events](#)

[Investigating a Single Host](#)

[Investigating a Single User](#)

Investigating Event Alarms

Use the Event Alarms page to investigate a directly alarming security event (as opposed to a non-alarming security event) related to a Category Alarm which has alarmed. This page will help to provide additional context for your category alarm by displaying details of the analogous security event that has also alarmed.

The Event Alarms page provides information similar to the information available on the Alarm List. Note that although this page shows whether an alarm has been acknowledged, to acknowledge the alarm you must use the Desktop Client.



- If you want a security event to trigger its associated alarm, you must perform this configuration in the Desktop Client. Refer to the topic entitled "Effective Host Policy" in the Desktop Client Help.
- You must have an identity service installed if you want the Source User column to display data.

How to find this page

Do one of the following to access this page

From either of the following pages, click the applicable security event alarm entry in the Event Alarms column:

- Alarms List View that is associated with a category alarm
- Alarm dashboard that is associated with a category alarm

The Event Alarms page opens.

What you can view on this page

Use this page to evaluate the following

- What type of directly alarming event has occurred?
- Based on the First Active and Last Active columns, how long ago did the alarm start? How long did it last? Do either of these convey a sense of concern or urgency?
- Which host or host group did the alarm affect? Is that normal for your network, or is it a source of concern?

- Does the alarm for the type of security event cause concern? If so, view the Details column for more information about the alarm, such as protocol or number of bytes observed compared to the number allowed by the policy.
 - If either the Bot Infected Host – Attempted C&C Activity or Bot Infected Host – Successful C&C Activity alarm occurs, the Details column provides information about the event, including a confidence rating. The rating indicates Cisco’s confidence that [name of entry] is classified appropriately. The scale is 1 to 10, with 1 representing the lowest level of confidence and 10 representing the highest level.
- Is a user associated with the alarm? If so, does that raise questions or concerns?

What you can do next

View associated policies

To view or tune the policies that have caused an alarm to trigger, do the following:

1. Identity the type of policy and event for which the alarm has triggered.
2. Click **Tune Alarm** from the menu for any of the following:
 - An IP address in the Source or Target column
 - The applicable Actions column

If you pivot from the ...	Then your query filters on the ...	And this page opens...	Filtered by the...
Source or Target	IP address	Policy Management: Core Events tab	All effective policies for the IP address
Actions	security	Policy	Event,

If you pivot from the ...	Then your query filters on the ...	And this page opens...	Filtered by the...
column, and the security event or category (displayed at the top of the page) is a <i>core</i> event	event or category	Management: Core Events tab	and then Policy Name
Actions column, and the security event or category (displayed at the top of the page) is a <i>custom</i> security event	security event or category	Policy Management: Custom Events tab	Event

Investigate a host that is the source of an alarm

1. Click the IP address in the Source column.
2. Evaluate the information on the [Host Report](#).
3. If necessary, run a [flow search](#) for that host by clicking the menu next to the applicable IP address and clicking [View Flows](#). Click Edit if you first want to make edits to the Flow Search page.
4. To run a flow search associated with a row of data, click the menu in the Actions column to access the relevant options.

Investigate a user who is the source of an alarm

1. Click the user name in the Source User column.
2. Evaluate the information on the [User Report](#).

Investigate a host that is the target of an alarm

1. Click the IP address in the Target column.
2. Evaluate the information on the [Host Report](#).
3. If necessary, run a flow search for that host by clicking the menu next to the applicable IP address and clicking [View Flows](#). Click Edit if you first want to make edits to the Flow Search page.

Investigate a security event for a particular host

1. Click the link in the Details column.
2. Evaluate the information on the [Security Event Details](#) page.
3. If necessary, use the links to further investigate a source or target host.



To close or acknowledge an alarm, use the Desktop Client.

Run a top report search

- To run a top report search associated with a *single IP address*, click the menu next to the applicable IP address to access the relevant options.
- To run a top report search associated with a *row of data*, click the menu in the Actions column to access the relevant options.

Related topics

[About Alarm Categories](#)

[Security Event List](#)

[Investigating a Single Host](#)

[Investigating a Single User](#)

Examining Details of Security Events

Use the Security Events Details page to view all the [security events](#) affecting a host that have been active on your network since the last reset hour, whether the host is the source or target. Information on this page provides insight into what is causing the security events to occur. Use the data to diagnose and triage specific security events. Secure Network Analytics collects and displays security event information regardless of whether or not the events have caused an alarm to fire.

The security events listed are the ones associated with the [alarm category](#) you selected on the [Security Insight dashboard](#). The data displayed is sorted by the source or target host, the alarm category selected on the Security Insight dashboard, and the security events associated with that alarm category.



This page provides information similar to the information available in the Security Events page on the [Desktop Client](#).

How to find this page

Do one of the following to access this page:

On the Security Insight Dashboard

1. In either the Alarms by Type or Today's Alarms component, click the section for the specific type of alarm you want to investigate.

The Alarm List opens with the results filtered by the type of alarm you selected.

2. On the Alarm List page, for the host you want to examine, click the link in the Details column.

The Security Events Details page opens filtered by the host you selected.

On the Host List

1. Click the alarm category column in the row for the applicable host.

The Alarm dashboard opens filtered by the type of alarm you selected and the associated host.

2. On the Alarm dashboard, for the host you want to examine, click the link in the Details column.

The Security Events Details page opens filtered by the type of alarm you selected.

On the Host Report

Option 1

1. Do one of the following:

- Click the number in the applicable alarm category at the top of the page to view the security events that are contributing to that alarm category for the host.

OR

- In the Alarms by Type component, click on a bar for the alarm and date for which you want to investigate.

3. On the Alarm dashboard, for the host you want to examine, click the link in the Details column.

The Security Events Details page opens filtered by the type of alarm you selected.

Option 2

Click **View More Security Events** at the bottom of the Top Security Events component.

The Security Events Details page opens filtered by the associated host.

What you can view on this page

View the following information on this page

Please note that if Secure Network Analytics returns more than 2,000 records, the message (*Reached 2,000 Max Records*) is displayed at the top of the page beside the IP address. In this case, even if more than 2,000 records are returned, then the 2000 records that contain the largest number of concern index points are displayed.

- View the source and target hosts associated with the security events.
 - Are there any that might be misclassified?
- View the date and time the security events became active. Is there a pattern that might indicate scanning activity?
- View the type of security event and the number of times the event occurred (displayed in the Count column).
- View the volume of Concern Index (CI) points for this host that have been active on your network since the last reset hour.
- Click the arrow beside the applicable security event to expand it and view the following information:
- **Details** Information indicating why the security event was triggered. Depending on the event, some details that may be provided are the number of concern index point

this event has accumulated since the last reset hour as well as the tolerance level for which it has been set.

Detailed information is not provided for every security event type, since some do not require further details.

- **Description.** A description of the security event.

What you can do next

View associated policies

To view or tune the policies that have caused a security event to trigger, do the following:

1. Identify the type of policy and event for which the alarm has triggered.
2. Click **Tune Event** from the menu (represented by the **Ellipsis** icon) for any of the following:
 - An IP address in the Source or Target column
 - The applicable Actions column

If you pivot from a...	Then your query filters on the ...	And this page opens...	Filtered by ...
Source or Target	IP address	Policy Management: Core Events tab	All effective policies for the IP address
The Actions column, and the associated	security event	Policy Management: Core Events tab	IP address and Event

If you pivot from a...	Then your query filters on the ...	And this page opens...	Filtered by ...
ed security event is a <i>core</i> event			
The Actions column, and the associated security event is a <i>custom</i> security event	security event	Policy Management: Custom Events tab	Event

View associated flows

To view the flows associated with a specific security event, do one of the following:

- In the applicable Actions column, click the **Ellipsis** icon > **Associated Flows**.
- In the Action column for an IP address in the Source or Target column, click the **Ellipsis** icon > **View Flow**.

A Flow Search runs with the criteria for that security event or IP address.

Related topics

[About Alarm Categories](#)

[Security Event List](#)

Investigating Relationship Policy Alarms

This page does the following:

- Provides information similar to the information available on the Alarm List in the Manager or in the Alarm Table on the Desktop Client. The columns on this page cannot be changed.
- Shows whether an alarm has been acknowledged, but to acknowledge or close the alarm, you need to use the Desktop Client.

Use the Relationship Policy Alarm List to investigate relationship policy alarms that have been active on your network today (i.e., for the period since the last archive hour, which by default occurs at 12:00 a.m.).

Relationship policies provide a way for you to monitor the current state of traffic between host groups with the ability to filter the traffic by applications and services. This policy type offers more granular relationships than core policies, because not only can you configure the policy to trigger an alarm when a specified amount of traffic has occurred between host groups, you can also direct Secure Network Analytics to send an alarm when the traffic exceeds a threshold that has been specified for the designated applications or services.

By default, information is sorted based on the most recent alarm. If a column has arrows in the heading, you can sort on that column.

How to find this page

Do the following to access this page

On the Security Insight Dashboard, in either the Alarms by Type or Today's Alarms component, click the section for the specific type of alarm you want to investigate.

The Relationship Policy Alarm List opens filtered by the type of alarm you selected.

What you can view on this page

Use the Relationship Policy Alarm List to evaluate the following

- What type of alarm has occurred?
- Based on the First Active and Last Active columns, how long ago did the alarm start? How long did it last? Do either of these convey a sense of concern or

urgency?

- Which host group did the alarm affect? Is that normal for your network, or is it a source of concern?
- Is the amount of traffic that has occurred between any two host groups abnormal behavior?
- Has the amount of traffic between any two host groups exceeded the threshold that has been specified for the designated applications or services?

What you can do next

View associated policies

To view or tune the policies that have caused a relationship alarm to trigger, do the following:

1. Identity the type of policy and event for which the alarm has triggered.
2. Click **Tune Alarm** from the context menu for any of the following:
 - The applicable Actions column
 - An IP address in the Source or Target column

If you pivot from a...	Then your query filters on the ...	And this page opens...	Filtered by the...
Source or Target	IP address	Policy Management: Core Events tab	IP address
The Actions column for a relationship event	relationship event	Policy Management: Relationship Events tab	Event name, and then Policy Name

Investigate flows related to a particular host group or alarm

To run a flow search associated with a row of data, click the context menu in the Actions column to access the relevant options.

The Subject Host Group listed in the pop-up menu represents the Host Group – Side 1 on the [Policy Management: Relationship Events tab](#), and the Peer Host Group represents Host Group – Side 2. Flows filter on the following criteria:

- Host Group – Side 1 (Subject Host Group)
- Host Group – Side 2 (Peer Host Group)
- Start date and time
- End date and time

The following subsections explain how the start date, start time, end date, and end time on which the flow filters are calculated.

Start date and time

If the First Active time is...	Then the start date filtered on is the...	And the start time filtered on is the...
1. Before the last reset hour	day before the date listed in the First Active column	Last reset hour
2. After the last reset hour	date listed in the First Active column	Last reset hour

Examples: For the following examples, assume today's date is April 26, 2018 and the reset hour is 4:00 a.m.

- The First Active date and time are 4/26/18, 3:00 a.m. | The Last Active date and time are 4/26/18, 2:00 p.m.
 - Then the flow will **filter on** 4/25/18 as the *start date* and 4:00 a.m. as the *start time*.
- The First Active date and time are 4/24/18, 9:00 p.m. | The Last Active date and time are 4/26/18, 2:00 p.m.
 - The flow will **filter on** 4/24/18 as the *start date* and 4:00 a.m. as the *start time*.

End date and time

If the First Active day...	And the Last Active time is...	Then the end date filtered on is the...	And the end time filtered on is ...
3. After the date listed in the First Active column is <i>before</i> the Last Active date	Before the current time	Day after the First Active date	24 hours after the First Active time.
4. Listed in the First Active column is today or yesterday	Current	Last Active date	Current

Examples: For the following examples, assume today's date is April 26, 2018 and the current time is 4:00 p.m.

3. The First Active date and time are 4/25/18, 9:00 p.m. | The Last Active date and time are 4/26/18, 2:00 p.m.
 - The flow will **filter on** 4/26/18 as the *end date* and 9:00 p.m. as the *end time*.
4. The First Active date and time are 4/26/18, 9:00 p.m. | The Last Active date and time are 4/26/18, 2:00 p.m.
 - Then the flow will **filter on** 4/26/18 as the *end date* and 2:00 p.m. as the *end time*.

View reports that will give you more in-depth information about specific areas of interest

To run a top report search associated with a row of data, click the context menu in the Actions column to access the relevant options. The Subject Host Group listed in the pop-up menu represents the Host Group - Side 1 on the [Policy Management: Relationship Events tab](#), and the Peer Host Group represents Host Group - Side 2. Flows filter on the following criteria:

- Host Group – Side 1 as Subject Host Group
- Host Group – Side 2 as Peer Host Group
- Time frame: Refer to the Start date and Time and the End date and time tables shown in the previous section, [Investigate flows related to a particular host group or alarm](#).

Related topics

[Relationship Policy Management: Overview](#)

[Relationship Policy Configuration: Overview](#)

System Notifications: Overview

System notifications inform users immediately of any current general or system alarms.

These notifications are indicated by a numeral displayed beside the  (**Alert**) icon . The numeral indicates how many messages you have. When an alarm becomes inactive, it is no longer reflected in the number displayed in the Alert icon.



- You cannot access system notifications on a failover Cisco Secure Network Analytics Manager (formerly Stealthwatch Management Console).
- To view the online help topic for the System Notifications home page (the Notifications Panel), see [System Notifications: Notifications Panel Overview](#).

System Notifications: Table of Contents

System Notifications: Notification Categories

The two main notification categories displayed in the Notification panel are General and System Alarms.

General

General notification types fall into the following two subtypes:

Smart Licensing

- **Evaluation Period Expired:** Your evaluation period has expired. Flow collection has stopped and your UDP Director has stopped forwarding flow. To start flow collection again, register your product instance.
- **Authorization Expired:** If Secure Network Analytics loses communication with your Cisco Smart Account, your authorization may expire. Authorization Expired indicates the communication status. It does not indicate license status.
- **Out of Compliance:** An appliance or feature has a license shortage and is using more licenses than are allocated in your Cisco Smart Account.

For more information about these system notifications and Smart Licensing, please refer to the [Cisco Secure Network Analytics Smart Software Licensing Guide](#).

Password Expiry

Ten days remain until the **Password Expires After** value is reached.

System Alarms

To see the list of system alarms that users can view in the Notifications panel, see [Manager System Alarms](#) and [Flow Collector System Alarms](#).



The DataStore alarms are functional only in Secure Network Analytics systems deployed with a Data Store.

Related Topics

[System Alarm Severity Descriptions](#)

[System Notifications: Overview](#)

System Alarm Severity Descriptions

Refer to the table below for alarm severity definitions associated with system alarms.

Icon Color	Alarm Severity
Red	Critical
Orange	Major
Yellow	Minor
Dark Blue	Trivial
Light Blue	Informational

For more details about alarm severities, see [About Alarm Severities](#).

Related Topics

[System Notifications: Notification Categories](#)

System Notifications: Notifications Panel Overview

To view the list of current notifications, open the Notifications panel.

Open the Notifications Panel

To view the [Notifications Panel](#), click the  (**Alert**) icon .

The Notifications panel slides out on the right side of the page.

Related Tasks

[System Notifications: View the Notifications Panel](#)

Related Topics

[System Notifications: Overview](#)

System Notifications: View the Notifications Panel

SMC database ingest and maintenance connection with Data Store lost
Device: Stealthwatch Management Console

 **Critical**

13 hours ago

Alarm ID

0-LZ5E-FYA7

Date & time

08/09/2024 12:36 AM

Duration

14:34:07

Description

sw-datastore unreachable

[View report](#)

SLIC channel down
Device: StealthWatch Management Console

 **Major**

15 hours ago

Flow Collector management channel down
Device: FlowCollector (10.0.88.143)

 **Minor**

16 hours ago

On this panel you can view the following information about an active notification:

- System notification name.
- [Alarm severity](#) (if the system notification is for a system alarm).
 - A  (**Status**) icon is displayed beside the notification name, indicating the severity level.

- A down arrow is displayed beside the **Status** icon. Click the down arrow to view the Alarm details. The details are displayed below the name of the System notification.
- If it is a system alarm, the time or day the system alarm last occurred (password expiry notifications show only for the current day).
- The device on which your Secure Network Analytics system observed it.

What to do Next

[System Notifications: View the Alarm Detail Panel](#)

Related Tasks

[System Notifications: Investigate System Notifications](#)

Related Topics

[System Notifications: Notifications Panel Overview](#)

[System Notifications: Overview](#)

System Notifications: Alarm Detail Panel Overview

To view the details for a particular system alarm, open the Alarm Detail panel for the associated system alarm.

Open the Alert Details Panel

To view the [system alarm details](#) for a system alarm, click **View Details** located under the associated system alarm for which you want to view details.

The Alarm Detail panel opens beside the Notifications panel.

Related Tasks

[System Notifications: View the Alarm Detail Panel](#)

Related Topics

[System Notifications: Overview](#)

System Notifications: View the Alarm Detail Panel

SMC database ingest and maintenance connection with Data Store lost

 **Critical** ^

13 hours ago

Device: Stealthwatch Management Console

Alarm ID 0-LZ5E-FYA7

Date & time 08/09/2024 12:36 AM

Duration 14:34:07

Description sw-datastore unreachable

[View report](#)

SLIC channel down

 **Major** v

15 hours ago

Device: StealthWatch Management Console

Flow Collector management channel down

 **Minor** v

16 hours ago

Device: FlowCollector (10.0.88.143)

On this panel you can view the following details about the system alarm:

- [Alarm severity](#) (if the system notification is for a system alarm)
- The device on which your Secure Network Analytics system observed it
- Alarm ID
- Date and time the alarm last occurred

- Duration of the alarm
- [Alarm description](#)



You can click **View Report** at the bottom of each Alarm Detail section to access the **System Alarms Report** in Report Builder. This report contains a list of all active alarms.

What to do Next

[System Notifications: Investigate System Notifications](#)

Related Tasks

[System Notifications: View the Notifications Panel](#)

Related Topics

[System Notifications: Alarm Detail Panel Overview](#)

[System Notifications: Overview](#)

System Notifications: Investigate System Notifications

Consider the following questions when investigating system notifications:

- Does this alarm indicate unusual behavior for the associated device on your network?
- Is this an exceptional deviation from normal behavior for this device?
- How severe is the impact to your network if this behavior is allowed to continue?

Related Tasks

[System Notifications: View the Notifications Panel](#)

[System Notifications: View the Alarm Detail Panel](#)

Related Topics

[System Notifications: Overview](#)

Secure Network Analytics Policies

A policy is a set of rules that you can apply to entities within your network that monitors and responds to host behavior that it observes. The Manager uses three types of event-related policies: custom (security event), relationship, and core. Refer to the table below to learn more about each policy type.

The Primary Admin and Power Analyst web roles have unlimited privileges within policy management.

Policy Type	Description
Custom (security event)	You can create custom (security event) policies to address specific concerns of your organization. They are the most granular in nature since they build off a single flow's attributes. The following list includes some examples (this list is not exhaustive) of the criteria you can configure: • total bytes • total packets • time of day • duration of flow
Relationship	Relationship policies provide a way for you to monitor the current state of traffic between host groups with the ability to filter the traffic by applications and services. This policy type offers more granular relationships than core policies, because not only can you configure the policy to trigger an alarm when a specified amount of traffic has occurred between host groups, you can also direct Secure Network Analytics to send an alarm when the traffic exceeds a threshold that you have specified for the designated applications or services. These policies trigger only relationship alarms. For more information, see the "Host Group Relationship Alarms" section in the Alarm List

Policy Type	Description						
	topic in the Desktop Client Help.						
Core	Core policies control how Secure Network Analytics monitors and responds to host behavior that it observes. Core policies can reduce unwanted alarms on your system as well as ensure that alarms are triggered in certain instances (when configured criteria is observed). The following table lists the three different types of core policies that Secure Network Analytics uses. These policies are listed in order of precedence; the one that takes the highest precedence is listed first. <table><tr><th>This type of policy...</th><th>Pertains to...</th></tr><tr><td>Host</td><td>Specific IP addresses. You cannot assign a host to more than one host policy.</td></tr><tr><td>Role</td><td>Collections of hosts (IP addresses) that serve a common purpose (e.g., Web servers, firewalls, trusted Internet hosts, etc.). It is always preferable to create a role policy for a host group rather than for an individual host. When you add a host to a host group, it automatically inherits the policy associated with that host group. When you need to modify a host's policy, it is easier to only have to re-</td></tr></table>	This type of policy...	Pertains to...	Host	Specific IP addresses. You cannot assign a host to more than one host policy.	Role	Collections of hosts (IP addresses) that serve a common purpose (e.g., Web servers, firewalls, trusted Internet hosts, etc.). It is always preferable to create a role policy for a host group rather than for an individual host. When you add a host to a host group, it automatically inherits the policy associated with that host group. When you need to modify a host's policy, it is easier to only have to re-
This type of policy...	Pertains to...						
Host	Specific IP addresses. You cannot assign a host to more than one host policy.						
Role	Collections of hosts (IP addresses) that serve a common purpose (e.g., Web servers, firewalls, trusted Internet hosts, etc.). It is always preferable to create a role policy for a host group rather than for an individual host. When you add a host to a host group, it automatically inherits the policy associated with that host group. When you need to modify a host's policy, it is easier to only have to re-						

Policy Type	Description	
		associate the host to a more relevant host group rather than both re-associating the host group and then modifying any number of individual policies that may have been configured with the particular host included.
	Default	All Inside and Outside hosts. These settings will apply when no role policies or host policies have been created.

Related Topics

[Managing Custom Policies \(Custom Security Events\)](#)

[Configuring Custom Policies \(Custom Security Events\)](#)

[Relationship Policy Management: Overview](#)

[Relationship Policy Configuration: Overview](#)

[Managing Core Policies](#)

[Editing Default Policies](#)

[Configuring Role Policies](#)

[Configuring Host Policies](#)

Managing Custom Policies (Custom Security Events)

Use this page to view and maintain the custom policies (custom security events) you have created for your network as well as create custom security events that address specific concerns of your organization.

A custom security event is a unique policy created for your system that is configured based on rule types and rule values. These rules evaluate the related flow attributes. If these attributes match the rule parameters you have set for a policy, then Secure Network Analytics can potentially trigger the alarm for that custom security event as well as for the Policy Violation alarm category. Unlike standard events, this evaluation is limited to a single flow and is not based on baselining or trends.

Custom security events begin as a blank slate and allow you to designate the who, what, and how for the flow that will cause an event or alarm to trigger. You make these configurations based on key flow data points, and they can be defined based on only a subject or can be configured to evaluate the subject in relation to a peer. By default, this evaluation disregards the flow orientation, although you can define the subject orientation.

For information about configuring custom events, see [Configuring Custom Policies \(Custom Security Events\)](#).

How to find this page

Do any of the following to access this page

- From the main menu on any page, click **Configure > Detection > Policy Management**.

The Policy Management: Custom Events tab opens.

- From the context menu for any IP address in the Web App, choose **Show Effective Policy**.

The Policy Management: Core Events tab opens. The IP address from which you pivoted is displayed in the search box at the top of the page, and all core events related to that IP address are listed in the Core Events table. To see all custom events related to the IP address, click the Custom Events tab. (The relationship events on the Policy Management: Relationship tab cannot filter on IP addresses.)

- Next, click the Core Events tab.

The IP address from which you pivoted is displayed in the search box at the top of the page, and all core events related to that IP address are listed in the Core Events table. (The relationship events on the Policy Management: Relationship tab cannot filter on IP addresses.)

- On the Alarm List, the Security Events Details page, or the Top Security Event component on the Host Report, from the Actions drop-down menu for the applicable custom security event, choose **Tune Alarm**.

Depending on which policy type triggered the event, you will be redirected to either the Policy Management: Custom Events tab, Relationship tab, or Core Events tab. If it was a custom event that triggered the alarm, then that custom event type is displayed in the Event column's filter box, and the custom policy for that event is displayed in the Custom Events table.

- On the Security Events Details page or the Top Security Event component on the Host Report, from the Actions drop-down menu for the applicable host, role, or default event, choose **Tune Event**.

Depending on which policy type triggered the event, you will be redirected to either the Policy Management: Custom Events tab, Relationship tab, or Core Events tab. If it was a custom event that triggered the alarm, then that custom event type is displayed in the Event column's filter box, the IP address is displayed in the Search box, and the custom policy for that event is displayed in the Custom Events table.

- On the Host Group Report page, from the Actions drop-down menu, choose **Policy Management**.

The Policy Management: Core Events tab opens by default. The host group from which you pivoted is displayed in the search box at the top of the page, and all custom, relationship, and core policies that are applied to this host group are displayed on the associated tab.

Click the Custom Events tab.



- By default, custom policies are sorted by date modified.
- Only users assigned the role of Primary Admin or Power Analyst can access this page.
- The number displayed on the tab indicates the number of defined custom security events.

Using the menu accessed from the ellipsis icon

When you click an ellipsis icon in the Web App, a menu containing submenu options opens. If you want to open a page in a new tab when choosing one of these options, do

one of the following:

- On a Windows-based PC, perform a Ctrl+Click.
- On a Mac, perform a Cmd+Click.

 This functionality does not work if you are connected to a virtual machine.

What you can do on this page

View your current custom security events

Refer to the following when viewing information on this page.

Event

The event's unique name. Click the arrow to expand the row and view the summary statement. The summary statement summarizes the rules you have configured for the custom security event.

 A small number of custom security events exist by default on your system.

Description

Indicates the purpose of this custom security event.

Date Modified

Indicates the time when the custom security event was last modified. You can toggle this column to use ascending or descending order.

Subject

Lists the host group(s) and/or IP address(es) that are acting as the subject for this custom security event. These subjects comprise those implicitly included and excluded by the user. If there are more than three defined subjects, the number of additional subjects is displayed within parentheses, and a link entitled "More" is displayed beside this number. If you click this link, the Subject panel (non-editable) opens from which you can view the complete list of subject host groups and/or hosts that are configured for the applicable custom security event. The total number of defined host groups or hosts will be displayed in parentheses on the associated tab. If a tab does not contain any defined entries, (0) will be displayed on the tab.

Peer

Lists the host group(s) and/or IP address(es) that are acting as the peer for this custom security event. These subjects comprise those implicitly included and excluded by the user. If there are more than three defined peers, the total number of additional peers is

displayed within parentheses, and a link entitled "More" is displayed beside this number. If you click this link, the Peer panel (non-editable) opens, from which you can view this complete list of peer host groups and/or hosts that are configured for the applicable custom security event. The total number of defined host groups or hosts will be displayed in parentheses on the associated tab. If a tab does not contain any defined entries, (0) will be displayed on the tab.

Status

Indicates if the custom event is being evaluated.

If you want to...	Click the toggle so that the following status is displayed...
Enable a custom security event	<i>On.</i>
Disable a custom security event	<i>Off.</i>

Find all custom security events associated with a host or host group

To find all the custom policies associated with a host or host group, do one of the following:

- For a host, type the IP address into the search box at the top of the page and click **Search**.

This will return all custom events which contain the following:

- The IP configured as an include or exclude value.
 - An IP range configured as include or exclude, which contains the IP.
 - A host group configured as include or exclude, which contains the IP address, OR parents of that host group.
- For a host group, click the  (**List Menu**) icon in the search box, select the host group from the host group tree that is displayed, and click **Search**.

This will return all custom events which contain the following:

- The specific host group configured as an include or exclude value.

Once filtered, all the custom policies that are associated with that host or host group are displayed. The number displayed on the tab changes to indicate the number of core policies configured for the host or host group. If you also enter filter criteria (see [Filter columns if you need to narrow your search](#) in this topic), keep in mind that the system will return results that match both the host or host group you selected in the search box and all the criteria you entered in the filter fields at the top of each column.

The page will return to the default view if you do either of the following:

- Clear the host group in the search box and press **Enter**.
- Make a change and click **Save**.

Filter columns if you need to narrow your search

At the top of each of the filterable columns, there is a filter field in which you can enter the criteria by which you want to filter that column. As you start to type your entry, the column dynamically filters to display a list of entries that match what you have begun to type. This match can be for any word in the entry (beginning, middle, or last word). Type the complete entry into the field to return results that are an exact match.

To return events that contain only the value you entered into a field for that field, you must enclose the value in quotation marks. Otherwise, the returned results will also include events which contain your filtered value as only ONE of several entries in that field.

Example:

Let's say you want to filter only on 10.3.4.5 as a value for the Subject. In other words, you want the results to include events that contain only 10.3.4.5 as a value for the Subject field.

- You must type "10.3.4.5" in the field. The results will include events with only 10.3.4.5 as a value for the Subject.

Now you want the results to include every event that has 10.3.4.5 as a value for the Subject even if there are other values in the Subject field along with 10.3.4.5.

- You must type 10.3.4.5 in the field. The results will include events with 10.3.4.5 as either the only value or as one of two or more other values for the Subject.



Filtering policies with multiple entries

You can type multiple entries within one column as well as enter values across multiple columns. When you type multiple values within one column, these values are ORed together; therefore, all returned results need only match one of the entries within the column. Remember to leave a space between each entry. When you filter more than one column, filters are ANDed together; therefore, all returned results must match the search criteria in all of the columns.

Clearing entries

Click the X in the field to clear the filter criteria and return to the previous view for that column. To quickly clear multiple filter criteria, refresh the page.

Sort columns

To reverse the sort order for a column, click the column name. Use the column name to toggle between ascending and descending order. You can also click the arrow that is displayed to the right of the column name (after you click the column name). An up arrow indicates that the column is sorted in ascending order, and a down arrow indicates descending order.

Enable or disable an event

To disable an event that is enabled, click the  (**Toggle**) icon in the Status column so that the status displays *Off*. To enable the event, click the switch so that the status displays *On*.

Delete a custom security event policy

In the applicable row, click the **...** (**Ellipsis**) icon in the Actions column and click **Delete** from the context menu. The custom security event policy is deleted from the Custom Security Event table.

Duplicate a custom security event

Duplicate a custom security event if you want to add a custom security event that is similar to an existing one so you can use it as a template and therefore create the new custom security event more easily and quickly.

In the applicable row, click the **Ellipsis** icon in the Actions column and click **Duplicate** from the context menu. The [Custom Security Event page](#) opens, showing the duplicated event's details populated in the applicable fields. Change the policy name (each custom security event name must be unique) and confirm that the event's status is correct. After you finish making edits, click **Save**.

What you can do next

Create a custom security event

From the main menu, click **Configure > DETECTION Policy Management** and then click **Create New Policy > Custom Security Event**. The Custom Security Event page opens.

For information about how to create a custom security event, see [Configuring Custom Policies \(Custom Security Events\)](#) .

Edit a custom security event

In the applicable row, click the **Ellipsis** icon in the Actions column and click **Edit** from the context menu. The Custom Security Event page opens, showing the custom security event's details populated in the applicable field. Make your edits and click **Save**.

View a triggered custom security event

Do any of the following:

- On the [Security Insight Dashboard](#), in the Alarms by Type or Today Alarms components, click on a bar or section that represents a Policy Violation alarm or a custom security event alarm. The [Alarm List](#) opens.

- On the [Host Report](#), do any of the following:
 - Click a number in the Policy Violation category (top row of the page). The [Alarm Dashboard](#) opens.
 - In the Alarms by Type component, click on a bar that represents a policy violation or custom security event alarm. The Alarm Dashboard opens.
 - Look in the Top Security Events table for the custom security events.

Related topics

[Secure Network Analytics Policies](#)

[Configuring Custom Policies \(Custom Security Events\)](#)

[Security Event List](#)

[About Alarm Categories](#)

[About Alarm Severities](#)

Configuring Custom Policies (Custom Security Events)

Use this page to create and manage your custom policies (also known as custom security events). Within Secure Network Analytics, custom policies are also referred to as custom security events.

A custom security event is a unique policy created for your system that you configure based on rule types and rule values. These rules evaluate the related flow attributes. If these attributes match the rule parameters you have set for a policy, then Secure Network Analytics triggers the alarm for that custom security event and potentially a corresponding Policy Violation alarm. Unlike standard events, this evaluation is limited to a single flow and is not based on baselining or trends.

Custom security events begin as a blank slate and allow you to designate the who, what, and how for the flow that causes an event or alarm to trigger. You make these configurations based on key flow data points, and you can define them based on only a subject or you can configure them to evaluate the subject in relation to a peer. By default, this evaluation disregards the flow orientation, although you can define the subject orientation.

For information about viewing and maintaining custom events, see [Managing Custom Policies \(Custom Security Events\)](#).



- Only users assigned the role of Primary Admin or Power Analyst can access this page.
- You can define up to a lifetime maximum of 2750 custom security events. If you create a custom security event and then delete it, this still counts toward the 2750 cap. The maximum number of custom security events that can be enabled at one time are 182.

How to find this page

Do any of the following to access this page

- From the main menu:
 1. Choose **Configure > Detection > Policy Management**.
The Policy Management: Custom Events tab opens.

2. From the **Create New Policy** drop-down menu, choose **Custom Security Event**.

The Custom Security Event page opens.

- On the Policy Management: Custom tab:
 1. In the row of the applicable custom security event, click the **⋮ (Ellipsis)** icon in the Actions column.
 2. From the context menu, click **Edit**.

The Custom Security Event page opens.

Concepts you need to understand to use this page

When do custom security events trigger

- Alarms are triggered within the first few minutes after a flow begins.
- A positively evaluated custom security event triggers a direct alarm and potentially a corresponding Policy Violation alarm (if the host alarming for the custom security event belongs to any policy for which the Policy Violation alarm has been turned off, the Policy Violation alarm will not trigger).
- Custom security events do not contribute individual index points to any Alarm Category.
- If a host is associated with an active custom security event alarm, that same custom security event does not trigger another alarm for a flow matching the same criteria for that host.
- After a custom security event becomes inactive, if the same conditions matching that event's rules are seen again, Secure Network Analytics triggers a new custom security event alarm. The Policy Violation alarm, however, is simply be reactivated.
- If a host matches a different custom security event rule from any it has matched on previously, Secure Network Analytics triggers both a custom security event alarm and a Policy Violation alarm that match on this rule.

Refer to the following guidelines when configuring a custom security event

- **[IMPORTANT]** Run a flow search for the last 24 hours matching the criteria you want to configure for the custom event. If any matching criteria exists, the results show the number of events that would have been generated if the custom event had

already existed. Running this report before you configure the custom event prevents a voluminous amount of unnecessary events being generated if you determine that this is a common flow occurrence.

- If possible, add additional rules and values to your custom security event to ensure your results are relevant to your purpose and that you do not overload your Flow Collector.
- If flows matching the desired rules exist regularly, consider building a scheduled report using a filtered Flow Table or Top Report (these are built in the Desktop Client. You can then receive the reports on an hourly, daily, or weekly basis instead of having the Alarm Table fill with events that you would not immediately address.
- If you need to filter on a list of IP addresses, it is best to build host groups and then apply rules to the host groups.
- The custom security event rule builder follows familiar interactive patterns used on the Flow Search page.
- Build custom security events with Subject rules as the main subject of your configuration and Peer rules as the secondary side of the communication.
- Though Subject Orientation is available to designate orientation (subject as client, server, or either), you do not have to specify a direction in your configuration. Leaving a custom security event generic for client and server allows the evaluation to happen in either direction. Eliminating the need to duplicate and flip your configuration. The peer is always opposite of the subject.
- When creating a custom security event, you have the ability to configure the following rules:

Subject only

- Orientation

Subject or peer

• Host	• TrustSec Name
• Host Group	• Application
• User	• File Hash
• Devices	• Process Name

• Port/Protocol	• Bytes
• TrustSec ID	• Packets
• Byte Ratio	

Flow connection

• Flow Action	• Time of Day
• Protocol	• Duration
• Service	• Server Response Time
• Total Bytes	• Round Trip Time
• Total Packets	• Encryption TLS/SSL Version

See [here](#) for detail on the availability of specific rules (e.g., TrustSec, TLS/SSL Version).

What query logic does the system use between and within sections

Logic Terms

The rule builder relies on Boolean logic. Refer to the following for the definitions of terms used when discussing the logic that is used on this page.

Term	Definition
Rule	A flow attribute that can be used to define a custom security event configuration, specific to the Subject, Peer, or Connection of the conversation. Subject Host Group, Peer Process Name, Protocol, Service, Duration
Value	An included or excluded item configured for a given rule. Protocol: TCP, UDP, ICMP; Total Bytes: >5k

Term	Definition
Inclusion	A value designated within a rule for which you would like the custom security event to alarm (an individual bad attribute you want to track). I want to alarm on all communications to the Americas.
Exclusion	A specific exception value designated within a rule for which you would NOT like the custom security event to alarm. These are designated with an exclamation point (!). I want to alarm on all traffic coming from the Outside Host Group except the traffic coming from Trusted Networks. Notes: • Items that you do not specifically exclude are not considered as an include or exclude by the system, so they may still be included in the results. • Once you include a given value for a rule, all other possible values for that rule are automatically excluded.
Operator	Designates the context of a values range on which you want to evaluate. Options include <, <=, =, >=, >. Only certain fields accept operators. Total Bytes: >5k; Duration: <1 Second

Rules vs. Values

The custom security event rule builder contains rule types for the Subject of your Query, a Peer to your Subject, and Connection type rules (e.g., Subject Host Group, Peer Port/Protocol, Duration). The rule builder relies on Boolean logic while evaluating rules. All configured rules must pass for the custom security event to alarm. Refer to the following for an explanation.

Term	Definition
Rules	Rules are mostly ANDed together, so a query must return results that match the search criteria in all of them. For more information, see the "Query logic used for rules" section in Create or edit a custom security event .
Rule Values	Within each rule, you can add one and possible more values, depending on the rule. These values are evaluated according to the following:
	INCLUSIONS Within each rule, included values are treated as "OR" conditions. There must be a match on at least one of the included values to satisfy the rule (though matching on more than one inclusion also satisfies the rule).
	EXCLUSIONS Within each rule, excluded values are treated as "AND" conditions. To satisfy a rule, a flow cannot contain any of the exclusion values set in the rule.
	COMBINATIONS (Of inclusions and exclusions) In the case of parent-child relationships (Subject Host, Subject Host Group, Peer Host, And Peer Host Group fields), rules can contain both include and exclude values simultaneously. To satisfy the rule, the following must be true: • There must be a match on at least one inclusion (though there can be a match on more than one inclusion) • The flow cannot contain any of the exclusion values set in the rule.

Policy Management | Custom Security Event Cancel Save

NAME * DESCRIPTION STATUS ☐ OFF

When any *subject host* or host within a *subject host group*; through *80/tcp* communicates with any *peer host*; and the entire flow contains *>50K total bytes*, an alarm is raised.

FIND +

SUBJECT HOSTS + OR

SUBJECT HOST GROUPS + OR

EXCEPT AND

AND

TOTAL BYTES + AND

SUBJECT PORT/PROTOCOLS +

ACTIONS + ☐ Alarm when a single flow matches this event.

Rules are ANDed together.

Exception: When a host and a host group are added to a policy, they are ORed together.

Show me some examples of creating a custom event that will alarm on certain desired criteria

Example 1

You want to create a rule to trigger on two IP address ranges, four user names, and a port/protocol; you need, however, to exclude two specific IP addresses. You also want to only trigger when the subject communicates with a specific IP address.

1. Add the main subject of your custom security event by adding your subject IPs.
 - a. In the Find section, click the + (**Add**) icon.
 - b. From the drop-down list, click **Subject Host** to add it to your configuration.
 - b. Type or paste in the required IP address ranges.
 - Valid formats include IPv4, IPv6, or CIDR notation.
 - You can enter an IP address or a range. Some examples are 192.168.10.10, 192.168.10/16, 192.168.10-100.2.
 - d. Precede an IP address with an exclamation point (!) to indicate the exclusion from the included ranges.
2. Add the users related to the Subject Hosts.

- a. Click the **Add** icon.
 - b. From the drop-down list, click **Subject User** to add it to your configuration.
 - c. Enter the user names you want to use in your query.
 - Separate user names with a space, comma, or return.
3. Add the port/protocol related to the Subject Hosts.
 - a. Click the **Add** icon.
 - b. From the drop-down list, click **Subject Port/Protocol** to add it to your configuration.
 - c. Enter the port/protocol combination you want to use in your query.
 - Separate entries with a space, comma, or return.
4. Add the IP address to the opposite side of the conversation.
 - a. Click the **Add** icon.
 - b. From the drop-down list, click **Peer Host** to add it to your configuration.
 - c. In the Peer Host rule, enter the IP Address you want to use in your query.

The custom security event triggers alarms for any flow that matches ALL of the following criteria:

- Any IP addresses that fall within one OR both of the Subject Host ranges are those other than the two IP addresses that you excluded.

AND

- Any of the four user names entered are on the same side of the conversation as that host.

AND

- The host communicated through the defined port/protocol.

AND

- The specific Peer Host is always on the opposite side of the conversation.

Example 2

You want to trigger a custom security event when an internal server is connected to on an unauthorized port, excluding an authorized group. For example, a server should only be running ntp, dns, snmp, ssh, and icmp from all groups except from within the Server host group itself.

1. Add the host groups which are the subject of your configuration.
 - a. In the Find section, click the **Add** icon.
 - b. From the drop-down list, click **Subject Host Group** to add it to your configuration.
 - c. From the slide-out panel that opens, click **Inside Hosts**.
 - b. Exclude the Host Group named *Servers*.
 - If the selection is already green, click it once to exclude it.
 - If the selection is not yet green, click it twice to exclude it.
 - e. Click **Apply**.
2. Define the subject orientation.
 - a. Click the **Add** icon.
 - b. From the drop-down list, click **Subject Orientation** to add it to your configuration.
 - c. From the drop-down list, select **Client**.
3. Configure the host group on the other side of the communication.
 - a. Click the **Add** icon.
 - b. From the drop-down list, click **Peer Host Group** to add it to your configuration.
 - c. Select your server host group.
 - Click once to include the host group.
 - d. Click **Apply**.
4. Add the port/protocol by which the server is communicating.
 - a. Click the **Add** icon.
 - b. From the drop-down list, click **Peer Port/Protocol** to add it to your configuration.
 - c. Enter the port/protocol combinations you want to use in your query.
 - Precede a value with an exclamation point (!) to exclude it.
 - Separate entries with a space, comma, or return.
5. Add the application by which the server is communicating.

- a. Click the **Add** icon.
- b. From the drop-down list, click **Peer Application** to add it to your configuration.
- c. Click **Exclude** from the rule's popup menu.
- d. From the slide-out panel that opens, click **ICMP**.
- e. Click **Apply**.

The custom security event triggers alarms for any flow that matches ALL of the following criteria:

- The host is in any host group within the Inside Host Group except for the Servers host group.

AND

- The host is acting as the client.

AND

- The server is within the Server host group.

AND

- The port/protocol being used by the server is something other than those excluded (e.g., 22/TCP, 53/UDP, 123/UDP, or 161/UDP)

AND



- The application being used by the server is any other than ICMP.


Default custom security events

Secure Network Analytics new deployments (beginning with v7.1) come with the following pre-defined custom security events. If you want to add these to your upgraded system, make sure the host groups used with these custom security events exist in your environment. See the "About the host groups involved in the setup of custom security events" section in [Managing and Configuring Host Groups](#).

You can manually add each of the CSEs defined below. For more information, see [Create or edit a custom security event](#).

To enable or disable any of these CSEs, use the  (Toggle) icon on the Custom Events tab. For more information, see [Enable or disable a custom security event](#).

Custom Security Event Name	Description
Added as of v7.1	
	
CSE: Custom Reputation List	Triggers any time an inside host sends or receives traffic from hosts within the Outside > Custom Reputation List host group. When does this custom security event trigger an alarm? When any host within Inside Hosts communicates with any host within Custom Reputation List, an alarm is raised. What must I do for this custom security event to be able to alarm? You must enable this custom security event and populate the Custom Reputation List host group.
	

Custom Security Event Name	Description
CSE: P2P to Internet	Triggers any time a Peer to Peer (P2P) file sharing application is detected between Inside hosts and the Internet. P2P file sharing applications can be vulnerable to remote exploits, or usage could indicate a policy violation. When does this custom security event trigger an alarm? When any host within Inside Hosts communicates with any host within Outside Hosts; using P2P file, an alarm is raised. What must I do for this custom security event to be able to alarm? You must enable this custom security event and have a Flow Sensor.
CSE: Possible Remote Access Breach	 The screenshot shows the configuration page for the custom security event 'CSE: Possible Remote Access Breach'. The description is 'Internet Hosts connecting to internal servers on Remote Desktop applications, which is a sign of a possible breach.' The configuration rules are as follows: Find: - Subject Host Groups: Outside Hosts - Subject Packets: >10 - Subject Orientation: Client - Peer Host Groups: Inside Hosts - Peer Port/Protocols: 23/TCP, 3389/TCP, 6000-6063/TCP, 6000-6063/UDP, 5800-5803/TCP, 5900-5903/UDP - Peer Packets: >10 The configuration is set to trigger an alarm when any host within Outside Hosts, acting as a client, with >10 packet communicates with any host within Inside Hosts; through any of the monitored port/protocols; with >10 packet, an alarm is raised.
CSE: Possible Remote Access Breach	Triggers any time an Internet Host connects to internal servers on Remote Desktop applications, which is a sign of a possible breach. When does this custom security event trigger an alarm? When any host within Outside Hosts, acting as a client communicates with any host within Inside Hosts; through

Custom Security Event Name	Description
	any of the monitored port/protocols, an alarm is raised. What must I do for this custom security event to be able to alarm? You must enable this custom security event.
Name * CSE: Unauthorized DHCP Server Description Detects possible rogue DHCP servers. Classify authorized DHCP Servers with the DHCP Servers host group to tune out When any host within <i>Inside Hosts</i> except 255.255.255.255 and those within <i>DHCP Servers</i>; through 67/UDP communicates with any host within <i>Inside Hosts</i> except 255.255.255.255 and those within <i>DHCP Servers</i>; through 68/UDP, an alarm is raised. Find Subject Host Groups Inside Hosts X EXCEPT DHCP Servers X AND Subject Hosts 1255.255.255.255 X AND Subject Port/Protocols 67/UDP X AND Peer Host Groups Inside Hosts X EXCEPT DHCP Servers X AND Peer Hosts 1255.255.255.255 X AND Peer Port/Protocols 68/UDP X	Detects possible rogue DHCP servers and indicates possible malicious activity or policy violations. When does this custom security event trigger an alarm? When any host within Inside Hosts (except 255.255.255.255 and those within DHCP Servers); through 67/UDP communicates with any host within Inside Hosts (except 255.255.255.255 and those within DHCP Servers); through 68/UDP, an alarm is raised. What must I do for this custom security event to be able to alarm? You must enable this custom security event and populate the DHCP host group with any authorized DHCP servers (to tune out authorized DHCP server traffic).

Custom Security Event Name	Description
CSE: Unauthorized DNS Traffic	Name * CSE: Unauthorized DNS Traffic Description Generate an alarm when an internal host is using an unauthorized public DNS server. This event will help detect DNS cha When any host within <i>Inside Hosts</i> except those within <i>Internet Services</i>, acting as a <i>client</i> communicates with any host within <i>Outside Hosts</i> except those within <i>Authorized External DNS Servers</i>; through <i>53/TCP</i> or <i>53/UDP</i>, an alarm is raised. Find Subject Host Groups Inside Hosts X EXCEPT Internet Services X AND Subject Orientation Client AND Peer Host Groups Outside Hosts X EXCEPT Authorized External DNS Servers X AND Peer Port/Protocols 53/TCP X 53/UDP X AND
CSE: Unauthorized DNS Traffic	Generates an alarm when an internal host is using an unauthorized public DNS server. This event helps detect DNS changer type of malware or possible policy violations. When does this custom security event trigger an alarm? When any host within Inside Hosts except those within Internet Services, acting as a client; through 53/UDP or 53/TCP communicates with any host within Outside Hosts except those within Authorized External DNS Servers, an alarm is raised. What must I do for this custom security event to be able to alarm? You must enable this custom security event and populate the Authorized External DNS Servers host group or the child host groups (DNS Servers, NAT Gateways, Proxies).

IPv4 and IPv6 address formats

You can use either IPv4 or IPv6 addresses for hosts or host groups. The notation forms that must be used for each type are listed in the applicable table. You must separate the IP addresses by a comma, comma space, space, or return. Configuring custom security events with rules containing large ranges of IP addresses (e.g., X.X.X.X/1) could impact the performance of your Flow Collector.

IPv4 IP address

If you enter this type of IP address...	You must use this formatting...
Single IP address	10.52.1.55 Includes only the host at 10.52.1.55.
Trailing dot subnet	10.52. Includes any IP address from 10.52.0.0 to 10.52.255.255. Note: You must include the trailing period (.) at the end.
Classless Inter-Domain Routing (CIDR) notation (i.e., using a “/” to indicate the masked bits)	10.52.1.0/24 Includes any IP address from 10.52.1.0 to 10.52.1.255. Note: Use this notation if you prefer to enter the IP address using a “mask.”
Range of networks	10-11. Includes any IP address from 10.0.0.0 to 11.255.255.255 10.52-53. Includes any IP address from 10.52.0.0 to 10.52.255.255, and from 10.53.0.0 to 10.53.255.255 10.52-55.3. Includes any IP address from 10.52.3.0 to 10.52.3.255; from 10.53.3.0 to 10.53.3.255; from 10.54.3.0 to 10.54.3.255; and from 10.55.3.0 to 10.55.3.255

If you enter this type of IP address...	You must use this formatting...
	Note: Be sure to include the trailing period (.) at the end.
Range of hosts	10.52.1.0-10 Includes any IP address from 10.52.1.0 to 10.52.1.10 10.52.0.0-10.52.255.255 Includes any IP address from 10.52.0.0 to 10.52.255.255
Comma-separated list	10.52.1.10,10.52.1.50,10.100.1.20 Includes only these three hosts.

IPv6 IP address

If you enter this type of IP address...	You must use this formatting...
Single IP address	2001:0DB8:0000:0056:0000:ABCD:EF12:3456 2001:DB8:0:56:0:ABCD:EF12:3456 2001:DB8::56:0:ABCD:EF12:3456 2001:DB80:0:56::ABCD:EF12:3456 2001:DB80:0:56::ABCD:239.18.52.86
Global Routing Prefix subnet	2001:DB8:0:56::/64
Range of hosts	2001:DB8:0:56:ABCD:EF12:3456:1-10 2001:DB8:0:56:ABCD:EF12:3456:1- 2001:DB8:0:56:ABCD:EF12:3456:10

About IPv4 and IPv6 addresses

You can use either IPv4 or IPv6 addresses for hosts or host groups. The notation forms that must be used for each type are listed in the applicable table. You must separate the IP addresses by a comma, comma space, space, or return.

The Manager displays IPv6 addresses and IPv4 addresses simultaneously. Of the advantages provided by IPv6, the fact that it supports a larger address space is probably the most well known. IPv6 supports 128-bit addresses, whereas IPv4 supports only 32-bit addresses.

The Manager continues to support IPv4. The Flow Collectors can monitor IPv6 traffic, but they are not IPv6 addressable at this time.

Examples

- An IPv6 address consists of eight blocks of hexadecimal digits separated by colons, as shown in this example (where the x's represent hexadecimal digits):

Example: xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx

- You can replace blocks of 4-digit zeros with double colons (e.g., ::), but only once per IPv6 address. (Leading zeros can also be omitted.) Consider the following address (where the x's represent hexadecimal digits):

Example: xxxx:xxxx:0000:0000:0000:0000:0000:xxxx

This address can be rewritten like this:

Example: xxxx:xxxx::xxxx.

- Like an IPv4 address, an IPv6 address may include a prefix length that specifies how many left-most bits of the address make up the prefix. In IPv4 addresses, the prefix is used to determine the subnet mask (a.k.a., CIDR notation). In IPv6 addresses, it identifies the subnet that an interface belongs to, and routers use it for forwarding.

Example: xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx/40

- Take care when compressing an IPv6 address that contains a prefix, because if not done correctly, it can result in an incorrect interpretation when the address is uncompressed. You can write an IPv4 address as an IPv6 address, like this:

Example: xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:192.168.0.2

- If the first 6 blocks are zeros, you can compress the address like this:

Example: ::192.168.0.2

CIDR notation

Use the following table to understand classless inter-domain routing (CIDR) notation:

Class	CIDR Length	Mask	Number of Addresses
A	/1	128.0.0.0	2,147,483,648
	/2	192.0.0.0	1,073,741,824
	/3	224.0.0.0	536,870,912
	/4	240.0.0.0	268,435,456
	/5	248.0.0.0	134,217,728
	/6	252.0.0.0	67,108,864
	/7	254.0.0.0	33,554,432
	/8	255.0.0.0	16,777,216

Cl as s	CI DR Le ngt h	Mask	Numbe r of Addres ses
	/9	255.128 .0.0	8,388, 608
	/10	255.192 .0.0	4,194, 304
	/11	255.224 .0.0	2,097, 152
	/12	255.240 .0.0	1,048, 576
	/13	255.248 .0.0	524,28 8
	/14	255.252 .0.0	262,14 4
	/15	255.254 .0.0	131,07 2
B	/16	255.255 .0.0	65,536
	/17	255.255 .128.0	32,768
	/18	255.255 .192.0	16,384
	/19	255.255 .224.0	8,192
	/20	255.255	4,096

CI as s	CI DR Le ngt h	Mask	Numbe r of Addres ses
		.240.0	
	/21	255.255 .248.0	2,048
	/22	255.255 .252.0	1,024
	/23	255.255 .254.0	512
C	/24	255.255 .255.0	256
	/25	255.255 .255.128	128
	/26	255.255 .255.192	64
	/27	255.255 .255.224	32
	/28	255.255 .255.240	16
	/29	255.255 .255.248	8
	/30	255.255 .255.252	4
	/31	255.255 .255.254	2

What you can do on this page

Create or edit a custom security event



- You must assign a name and add a minimum of one event to create a custom security event. (Rules with only exclusions do not satisfy this requirement.)
- A newly created custom security event is disabled by default. You must enable it to make it active (see the Status option in this section).

To create or edit a custom security event, do the following:

- Complete or edit the fields at the top of the page.
- Add (or delete) the desired rules in the Find section.

As you add rules, they are displayed in this section. As you add a rule, a sentence is displayed at the top of the page that summarizes the rules you have configured for the custom security event (see the following image).

After you have saved the information on this page, the [Policy Management: Custom tab](#) opens and the policy you just added is displayed in the Custom Policy table. The events are sorted by date modified, with the most recently modified event at the top.

Refer to the following for guidelines on entering data on this page.

Name (required)

The name you enter must be unique not only among all custom security events but also among all non-custom security events. Cisco suggests that you use a standard prefix for every rule/event name (such as *CSE:*) to enable easy grouping for response management and when filtering reports.

Example: CSE: Unauthorized Connection to Servers

Description

Indicates the purpose of this custom security event. This description is displayed on the [Policy Management: Custom tab](#) as well as on the Security Events Details page for a triggered custom security event.

Status

The **Toggle** icon indicates if this custom security event is enabled (On, icon bar is blue ) or disabled (Off, icon bar is gray ). If you click the **Toggle** icon when the value displayed is *Off*, it displays *On*. Click it again to display *Off*. You must click **Save** after enabling or disabling for your change to take effect.

A newly created custom security event is disabled by default. You must enable it to make it active.

Find

Click the  (**Plus**) icon for each new rule you want to add. A drop-down search box is displayed, listing the first six rules that have not yet been used. To choose a rule, either you can begin to type any part of a rule name into the drop-down list (upon which any rule names that match your entry move to the top of the drop-down list), or you can select a rule. To view all of the rules, click **View All**. A slide-out panel opens that lists all of the rules by category from which to select: subject, peer, and connection (this mirrors the Flow Search page).

- You need only add the rule once. After you have added a rule, it is no longer displayed in either the search list or the slide-out panel.
- If you select a rule that does not require input using the slide-out panel, a new, empty field is displayed in the Find section in which you can enter the applicable information.
- If you need to configure the rule using a slide-out panel, then after you have done so, simply click **Apply** to accept your configuration and close the panel.

When entering a value directly into a field, the following applies:

- When you apply more than ten values from the slide-out panel, the "More" link is displayed in that option's field to indicate that one or more additional configurations exist that are not displayed in the field.
- When entering a value directly into a field, use a space, comma, or press Enter to separate multiple entries within that field.

When deleting a rule or value, refer to the following:

- To delete a rule you have added, click the *X* that is displayed to the right of the rule's text box.
- To delete a value you have added within a rule, click the *X* that is displayed within the specific pill.

Query logic used for rules

What logic is used for values within fields?	What logic is used for rules?
1. Included values within a single field are ORed together; therefore, a result need only match one of the included entries within the rule. (Remember to leave a space between each entry.) 2. Excluded values within a single field are ANDed together; therefore, a result must match all of the excluded entries within the rule. (Remember to leave a space between each entry.) 3. When a field contains both included and excluded values, a result must match at least one included value and all excluded values.	4. Rules are ANDed together; therefore, a result must match the search criteria configured for all of the rules. 5. There are two scenarios where there is an exception to this rule and where the rules are ORed together: Exceptions: • When the Subject Host AND Subject Host Group contain at least one include value each, and they are added to a policy. • When the Peer Host AND Peer Host Group contain at least one include value each, and they are added to a policy.

Example: In the screen image below, the results must adhere to the following logic:

- Any result must include 10.0.0.0 or 70.70.70.70 OR a Command & Control host or an Inside Host (#1 and #5).
 - All Catch All hosts within Inside Hosts must be excluded in the results, AND, all Outside Hosts must be excluded in the results (#2 and #3).
- The total bytes must be over 50,000. (#4)

- The subject port/protocol used must be 80/tcp. (#4)

Policy Management | Custom Security Event

NAME * Store # 2 DESCRIPTION ex. Event will trigger based on this rule STATUS OFF

When any *subject host* or host within a *subject host group*; through 80/tcp communicates with any *peer host*; and the entire flow contains >50K total bytes, an alarm is raised.

FIND

SUBJECT HOSTS 10.0.0.128 X 70.70.70.70 X

SUBJECT HOST GROUPS Servers X VoIP X Other X

OR Client IP Ranges (DHCP Range) X EXCEPT Guest Wireless Networks X Remote VPN IP Pool X

AND Countries X

TOTAL BYTES >50K X

SUBJECT PORT/PROTOCOLS 80/tcp X

ACTIONS Alarm when a single flow matches this event.

Rules are ANDed together.

Exception: When a host and a host group are added to a policy, they are ORed together.

When you delete a host group, service, or application

When you delete a host group, service, or application, it is not deleted automatically from related custom security events, which can invalidate the custom security event configuration and cause missing or false alarms. Similarly, when you disable Threat Feed, this removes the host groups that Threat Feed added, so you will need to update any related custom security events.

We recommend the following when you are either considering deleting a host group, service, or application (planning), or for when you are checking to see if any customer security events have been impacted and need to be updated (reviewing):

- **Planning:** Before you delete a host group, service, or application, or disable Threat Feed, review your custom security events to determine if you need to update them.

If the rule value (such as a host group or service) you are planning to delete or disable is included in the custom security event, delete the event or edit it to use a valid rule value (click the **⋯ (Ellipsis)** icon next to the applicable custom security event and choose **Edit**).

- **Reviewing:** For each custom security event, click the **⋯ (Ellipsis)** icon and choose **Edit**. If the custom security event is blank or missing rule values, delete the event or edit it to use valid rule values.

Refer to the following parameters when configuring a custom security event

See the placeholder text in the Web App for examples of valid entries for each rule.

- Once you add a rule, you can enter one or more values to it, and it is no longer displayed in the Add Rule menu.
- In some fields you can enter more than one value, and in some fields you can only enter one; a message displays if you are attempting to enter more than one value into a field that accepts only a single value.
- To exclude a value, precede your entry with an exclamation point (!). You can only do this in fields that do not accept operator symbols (<, <=, =, >=, >).
- The only fields for which you can have an included value and an excluded value (for the same rule) are the Subject Host, Peer Host, Subject Host Group, and Peer Host Group fields.
- You can use some or all of these special characters as operators, depending on which field you are configuring: >, <, =, >=, <=. If you do not include an operator with your entry (null), Secure Network Analytics treats this entry as an equal (=) operator.
- Using "<" in a rule requires the flow to end before being evaluated, so therefore the flow may not immediately trigger an event.
- For fields requiring a numeric value, you can enter either the number of units (e.g., 1) or the number of units followed by the applicable abbreviation (e.g., 1k). Units of measure allowed are K, M, and G.
 - 1k = 1,000 seconds.
- You can use either one or more IP addresses, IP ranges, and/or host groups to define the Subject Host, Peer Host, Subject Host Group, or Peer Host Group. See the [IPv4_and_IPv6_address_formats](#) section for guidelines to use when entering IP addresses. The notation forms that must be used for each type are listed in the "IPv4 and IPv6 address formats section."
- To delete a selection within a rule, click the corresponding X.
- To delete an entire rule, click the X that is displayed at the end of the text box.

Information about options available in the Find section

Subject Host Groups / Peer Host Groups

Remember, if you add to a policy both a Subject Host and a Subject Host Group that each contain at least one include value, or you add to a policy both a Peer Host and a Peer Host Group that each contain at least one include value, the two rules are ORed together (all other rules are ANDed together). For more information, see the description for the Find field in the [Create or edit a custom security event](#) section.

To return an exact match when you are configuring host groups in the slide-out panel, enter the entire host group name into the Search field and press **Enter**. If you want to return to the previously selected host group, delete the entry in the Search field and press **Enter**. If you type any letters that are anywhere within a host group name into the Search field and press **Enter**, the Host Group tree expands to display (in italicized red font) any host group name that contains that entry.

Example: If you entered South into the Search field, some of the host group names that would then be displayed in italicized red font would be Southern Africa, South Africa, South-Eastern Asia, French Southern Territories.

When making a selection in the Host Group panel, the following applies:

- A check mark or a plain green button means that the associated host group is included in your custom security event. (When you mark a parent with a check mark, by default all of its children are marked with a green button, even if previously you had marked any of these children with an x.)
- An x or a gray button means that host group is excluded from your custom security event. (When you mark a parent with an x, by default all of its children are marked with a gray button, even if previously you had marked any of these children with a check mark.)

When you click a host group's check box, its subsequent appearance depends on the following factors:

- Its status before you clicked it
- Its location in the tree
- If its parent or any of its children are selected

After you click **Apply**, then your selections is displayed in the following order in the Subject Host Group field:

1. Host groups without any excluded children.
2. Host groups with one or more excluded children.
3. Excluded host groups without an included parent.

Policy Management | Custom Security Event Cancel Save

NAME * DESCRIPTION STATUS ☐ OFF

When any *subject host* or host within a *subject host group*, through *80/tcp* communicates with any *peer host*, and the entire flow contains *>50K total bytes*, an alarm is raised.

FIND ⓘ

SUBJECT HOSTS ⓘ OR

SUBJECT HOST GROUPS ⓘ

1 Servers x VoIP x Other x

OR

2 Client IP Ranges (DHCP Range) x EXCEPT ! Guest Wireless Networks x ! Remote VPN IP Pool x

AND

3 ! Countries x

TOTAL BYTES ⓘ AND

SUBJECT PORT/PROTOCOLS ⓘ ⓘ

ACTIONS ⓘ

Alarm when a single flow matches this event.

Subject Hosts / Peer Hosts

Remember, if you add to a policy both a Subject Host and a Subject Host Group that each contain at least one include value, or you add to a policy both a Peer Host and a Peer Host Group that each contain at least one include value, the two rules are ORed together (all other rules are ANDed together). For more information, see the description for the Find field in the "Create or edit a custom security event" section.

You can use either IPv4 or IPv6 addresses. (Refer to the "IPv4 and IPv6 address formats" section in this topic for the correct notation forms to use.)

Refer to the following guidelines:

- You must separate the IP addresses by a comma, comma space, space, or return.
- Type or paste the entries from Excel, CSV, or text files.
- You can include a range of hosts and then exclude a portion of it. To exclude a value, precede your entry with an exclamation point (!).

Subject Users / Peer Users

If you want to find specific user information or create a custom security event that contains specific user information, you must have implemented ISE or integrated Active Directory with Secure Network Analytics.

Subject TrustSec IDs / Peer TrustSec IDs

In order to search for TrustSec information or to create a custom event using it, you must have already implemented Cisco TrustSec and have created Security Tag Groups (SGT) on your network.

Subject TrustSec Names / Peer TrustSec Names

In order to search for TrustSec information or to create a custom event using it, you must have already implemented Cisco TrustSec and have created Security Tag Groups (SGT) on your network.

Subject Applications / Peer Applications

These two rules enable you to configure applications for either the subject or the peer. You should not configure an application for both the subject and peer. Click inside the field and select either **Include** or **Exclude**. A panel opens for you to make your selections. If you type the beginning letters or the complete name of the application name, it moves to the top of the list.

Secure Network Analytics must have a Flow Sensor in order for you to add application details to your NetFlow record.

Subject File Hashes / Peer File Hashes

Secure Network Analytics must have NVM telemetry in order to create a Custom (Security Event) Policy with a file hash to your NetFlow record. For more information about how to configure Secure Network Analytics and NVM, refer to the [Network Visibility Module \(NVM\) Configuration Guide](#).

Subject Process Names / Peer Process Names

Secure Network Analytics must have NVM telemetry in order to create a Custom (Security Event) Policy with a process name details to your NetFlow record. For more information on how to configure Secure Network Analytics and NVM, refer to the [Network Visibility Module \(NVM\) Configuration Guide](#).

Subject Orientation

Specify whether you want to return results for instances where the subject is the client, server, or either one (the client is considered to be the host that initiated the session). If you do not specify an orientation, then *Either* is selected by default. This allows you to alert in both scenarios off the same custom security event.

Flow Action

If you select *Permitted and Denied*, the event queries for any individual flow that is tagged as both *Permitted* and *Denied*. In the case of one flow traversing multiple routers, your network configuration may permit the flow to cross one router and deny it from crossing another.

Secure Network Analytics must have firewall support in order for you to add flow action details to your NetFlow record.

Time of Day

The time is converted to UTC time. You can add multiple time ranges.

Encryption TLS/SSL Version

The entry must be a supported ETA (encrypted traffic analytics) TLS or SSL version. You can use operators in your entries.

Examples:

- SSL 1.0
 - <TLS 1.2
-

Using this rule on its own may affect the performance of your system. It is best to configure this rule with specific host groups that are in an ETA-supported area of your network or paired with protocols, services, or applications which support encryption. The Encryption TLS/SSL Version field contains a double dash (--) on the Flow Search Results page (we always evaluate a double dash as the lowest value) if you do not have an ETA-enabled switch and router. Certain encryption fields may also contain a double dash if the flow record does not have information for that particular field. For more information about Secure Network Analytics and ETA, refer to the [Encrypted Traffic Analytics white paper](#) and the [Encrypted Traffic Analytics deployment guides](#).



Secure Network Analytics must have ETA in order for you to add encryption details to your NetFlow record.

Enable or disable a custom security event

If you want to...	Slide the Toggle icon to display...
Enable a custom security event	On, icon bar is blue  .
Disable a custom security event	Off, icon bar is gray  .

Delete a custom security event

From the Actions drop-down menu, click **Delete**.

The custom security event is deleted from your system and removed from the Custom Events table.

Duplicate a custom security event

Duplicate a custom security event if you want to add a custom security event that is similar to an existing one so you can use it as a template and therefore create the new custom security event more easily and quickly.

To duplicate a custom security event, do the following:

1. From the Actions drop-down menu, click **Duplicate**.
The Custom Security Event page opens, showing the duplicated event's details populated in the applicable fields.
2. Change the policy name (each custom security event name must be unique) and confirm that the event's status is correct.
3. Make any other edits, if necessary,
4. Click **Save**.

What you can do next

View triggered custom security events

- On the [Security Insight Dashboard](#), in the Alarms by Type or Today Alarms components, click on a bar or section that represents a Policy Violation alarm or a custom security event alarm. The [Alarm List](#) opens.
- On the Host Report, do any of the following:
 - Click a number in the Policy Violation category (top row of the page). The [Alarm Dashboard](#) opens.
 - In the Alarms by Type component, click on a bar that represents a policy violation or custom security event alarm. The [Alarm Dashboard](#) opens.
 - Look in the Top Security Events table for the custom security events.

If your policies are not behaving correctly, synchronize your Flow Collector with the Manager

At times you may need to synchronize the Flow Collectors on your network with your Manager. Some reasons may be as follows:

- You observe that your Flow Collectors are not working properly.
- You have added a new configuration (e.g., role policy, host group), but while you were doing so your Manager stopped communicating with any or all of your Flow Collectors. In this case, the policy will have been saved to your Manager, but it may not have propagated to any or all of your Flow Collectors.

To synchronize one or more Flow Collectors to your Manager, log in to your Manager client interface and complete the following steps:

1. From the navigation menu, click **Configure > Flow Collectors**.
2. Click the **Advanced** tab.
3. In the top left corner, from the Flow Collector drop-down list, choose the applicable Flow Collector.
4. In the Synchronize section, click **Synchronize**.
5. Repeat Step 3 through Step 4 for the remaining Flow Collectors that need to be synchronized with the Manager.

Related topics

[Managing Custom Policies \(Custom Security Events\)](#)

[Security Event List](#)

[About Alarm Categories](#)

[About Alarm Severities](#)

Relationship Policy Management: Overview



For information about configuring relationship policies, see [Relationship Policy Configuration: Overview](#).

Relationship policies provide a way for you to monitor the current state of traffic between host groups with the ability to filter the traffic by applications and services. You can also create network diagrams, which are visual representations that show active alarms and traffic bandwidth within host groups. You must download the Network Diagrams app in order to create network diagrams. This policy type offers more granular relationships than core policies, because not only can you configure the policy to trigger an alarm when a specified amount of traffic has occurred between host groups, you can also direct Secure Network Analytics to send an alarm when the traffic exceeds a threshold that you have specified for the designated applications or services.

All edits you make to a relationship policy in either the Manager or Desktop Client are displayed in both interfaces.

Open Relationship Policy Management

Do any of the following to access this page

- From the main menu on any page,
 1. Click **Configure > Detection > Policy Management**.
 2. Click the Relationship Events tab.

The Policy Management: Relationship Events tab opens.
- On the Alarm List, from the Actions drop-down menu for the applicable relationship event, click **Tune Alarm**.

Depending on which policy type triggered the event, you are redirected to either the Policy Management: Custom Events tab, Relationship tab, or Core Events tab. If it was a relationship event that triggered the alarm, then that relationship event type is displayed in the Event column's filter box, the policy name is displayed in the Policy Name column's filter box, and the relationship policy for that event is displayed in the Relationship Events table.
- On the Host Group Report page, from the Actions drop-down menu, select **Policy Management**.

The Policy Management: CoreEvents tab opens by default. The host group from which you pivoted is displayed in the search box at the top of the page, and all custom, relationship, and core policies that are applied to this host group are displayed on the associated tab.

Click the Relationship Events tab.



- The number displayed on the tab indicates the number of defined relationship events.
- By default, relationship policies are sorted by event. You can toggle this column to use ascending or descending order.

Relationship Policy Management: Table of Contents

[Relationship Policy Configuration: Overview](#)

[Security Event List](#)

[Alarm List: Overview](#)

[About Alarm Categories](#)

[About Alarm Severities](#)

Relationship Policy Management: Use Context Menus

If you want to open a page in a new tab when pivoting from the context menu, do one of the following:

- On a Windows-based PC, perform a Ctrl+Click.
- On a Mac, perform a Cmd+Click.



This functionality does not work if you are connected to a virtual machine.

Related Topics

[Relationship Policy Management: Overview](#)

Relationship Policy Management: Filter Columns

Filter columns if you need to narrow your search.

At the top of each of the filterable columns, there is a filter field in which you can enter the criteria by which you want to filter that column. Some display a drop-down list as you begin to type and some do not.

For the Event, Policy Name, and Map or Diagram Name fields, as you start to type your entry in any of these fields, a drop-down list appears displaying a list of entries that match what you have begun to type. You can then choose the appropriate entry from the list so that you don't have to type the complete entry. The remaining columns dynamically filter to display only the entries that match your entry.

For the Host Groups, Traffic By Services, and Traffic By Applications fields, as you start to type your entry, the column dynamically filters to display a list of entries that match what you have begun to type. This match can be for any word in the entry (beginning, middle, or last word). Type the complete entry into the field to return results that are an exact match.

Filtering with multiple entries

You can type multiple entries within one column as well as enter values across multiple columns. When you type multiple values within one column, these values are ORed together; therefore, all returned results need only match one of the entries within the column. Remember to leave a space between each entry. When you filter more than one column, filters are ANDed together; therefore, all returned results must match the search criteria in all of the columns. Enclose an entry within quotation marks to return results that exactly match your entry.

For the Host Groups column, the table will first list all events that contain your entry as the first host group, and then will display all events that contain your entry as the second host group. Also, the Host Groups column will filter all host groups, not just those displayed in the table.

Editing a policy with filtered criteria

If you edit a relationship policy after having filtered criteria, the entries in the table does not change until you save your edits. Click the X in the field to clear the filter criteria and return to the previous view for that column. To quickly clear multiple filter criteria, refresh the page.

Clearing entries

Click the *X* in the field to clear the filter criteria and return to the previous view for that column. To quickly clear multiple filter criteria, refresh the page.

Related Topics

[Relationship Policy Management: Overview](#)

Relationship Policy Management: Sort Columns

To reverse the sort order for a column, click the column name.

Use the column name to toggle between ascending and descending order. You can also click the arrow that is displayed to the right of the column name (after you click the column name). An up arrow indicates that the column is sorted in ascending order, and a down arrow indicates descending order.

Related Topics

[Relationship Policy Management: Overview](#)

Relationship Policy Management: What You Can Do On This Page Overview

You can perform the following actions on this page:

- [View current policies.](#)
- [View or edit event details of a relationship policy.](#)
- [Find all relationship policies associated with a host group.](#)
- [Create a policy.](#)
- [Edit a policy.](#)

Related Topics

[Relationship Policy Management: Overview](#)

Relationship Policy Management: View Current Policies

When viewing a relationship policy, consider the following questions:

- Which hosts comprise the relationship pair?
- Which network diagrams, services, and applications have been configured for this policy?



To view a relationship policy's details and all of the events related to the policy, click the policy name in the Policy Name column.

- Which two host groups are involved in this event?
- What services and/or applications are involved in the traffic between the two specified host groups?

If the total number of services or applications defined for the relationship policy exceeds three entries, the remaining number of services or applications is displayed within parentheses, and a link entitled "More" is displayed beside this number. If you click this link, a list of all the services or applications that comprise the traffic between the specified host groups is displayed. The total number of defined services and applications is displayed in parentheses on the associated tab in the Traffic by Services and Applications slide-out panel. If a tab does not contain any defined entries, (0) is displayed on the tab.

What is a network diagram?

Network diagrams enable you to more easily pinpoint the cause of a network problem, such as slow traffic, by visually showing relationship policies between two host groups. Using network diagrams, you can graphically monitor the status of your entire system in near real-time and easily view active alarms and your organization's network traffic based on various criteria, such as location, function, time, applications, services, or virtual environment.

When creating network diagrams, you can add images, edges (lines), and text boxes to aid in the visual representation of your environment. You can also change the domain or network diagram for which you want to view data.



To use the network diagrams feature, you must purchase the Network Diagrams app. For more information, call your local Cisco Partner.

Related Tasks

[Relationship Policy Management: View or Edit Event Details of a Relationship Policy](#)

[Relationship Policy Management: Find All Relationship Policies Associated With a Host Group](#)

[Relationship Policy Management: Create a Policy](#)

[Relationship Policy Management: Edit a Policy](#)

Related Topics

[Relationship Policy Management: What You Can Do On This Page Overview](#)

[Relationship Policy Management: Overview](#)

Relationship Policy Management: View or Edit Event Details of a Relationship Policy

To view or edit the event details associated with a particular entry in the Relationship Policy table, do the following:

1. Click the arrow at the beginning of the row.

The row expands to display a description of the relationship event as well as the settings you have configured for that event.

2. Make any edits, if necessary.

When you edit any settings or add any events to an existing policy, a green dog ear is displayed in the top right corner of the row. A red dog ear indicates an input error. Hover your cursor over a field name to see the minimum and maximum allowable value

The following table contains descriptions of the various event setting details.

Option Name	Description
Behavioral and Threshold	When this option is selected, the system uses the tolerance setting, minimum threshold, and maximum threshold. Note: When you change your selection from "Behavioral and Threshold" to "Threshold Only," the following fields become disabled: • Tolerance • Never trigger alarm when less than
Threshold only	The alarm triggers only when a minimum or maximum threshold is reached.
Tolerance	Tolerance is the relative number between 0 and 100 that indicates how much to allow actual behavior to exceed expected behavior before alarming. This allows the user to define what is "significantly different." It can also be defined as "the number of standard deviations from the norm." Note: Standard deviation shows how much variation there is from the average (i.e., mean, or expected value). A low standard

Option Name	Description
	deviation indicates that the data points tend to be very close to the mean, whereas high standard deviation indicates that the data points are spread out over a large range of values. • A tolerance of 0 causes an alarm to be triggered for any values over the expected value; it is very sensitive and results in a lot of alarms. • A tolerance of 100 allows the behavior to exceed expected behavior to a high degree before it causes an alarm to be triggered, thus greatly reducing the number of times an alarm is triggered. If you never want the alarm to trigger, you must disable the alarm.
Never trigger alarm when less than	Also known as the minimum threshold, this is a static value that indicates the lowest value to allow without triggering an alarm. The alarm does not trigger when the observed value falls below this setting. In other words, even if a host is greatly over its expected value, if it is not more than the minimum indicated in this dialog, then do not trigger an alarm.
Always trigger alarm when greater than	Also known as the maximum threshold, this is a static value that indicates the highest value to allow without triggering an alarm. The alarm triggers when the observed value exceeds this setting. In other words, if a host's value exceeds the maximum indicated in this dialog, even if it is expected for that host, then trigger an alarm.
Trigger alarm when duration greater than	Duration is the amount of time that the associated value can exceed the calculated threshold before the Manager triggers an alarm.

Related Tasks

[Relationship Policy Management: View Current Policies](#)

[Relationship Policy Management: Find All Relationship Policies Associated With a Host Group](#)

[Relationship Policy Management: Create a Policy](#)

[Relationship Policy Management: Edit a Policy](#)

Related Topics

[Relationship Policy Management: What You Can Do On This Page Overview](#)

[Relationship Policy Management: Overview](#)

Relationship Policy Management: Find All Relationship Policies Associated With a Host Group

To find all the relationship policies associated with a host group, do the following:

1. Click the  (**List Menu**) icon in the search box.
2. Select the host group from the host group tree that is displayed.
3. Click **Search**.

All the relationship policies that are associated with that host group are displayed, sorted by event.

The host group can be on either side of the relationship pair in the Host Groups column to be included in the results. The number displayed on the tab changes to indicate the number of relationship policies configured for the host group.

If you also enter filter criteria (see the next section, [Filter columns if you need to narrow your search](#)), keep in mind that the system returns results that match both the host group you selected in the search box and all the criteria you entered in the filter fields at the top of each column.

The page returns to the default view if you do either of the following:

- Clear the host group in the search box and press **Enter**.
- Make a change and click **Save**.

 You cannot search for a relationship policy associated with a single host.

Related Tasks

[Relationship Policy Management: View Current Policies](#)

[Relationship Policy Management: View or Edit Event Details of a Relationship Policy](#)

[Relationship Policy Management: Create a Policy](#)

[Relationship Policy Management: Edit a Policy](#)

Related Topics

[Relationship Policy Management: What You Can Do On This Page Overview](#)

[Relationship Policy Management: Overview](#)

Relationship Policy Management: Create a Policy

1. From the main menu, click **Configure > Policy Management**.
2. Click **Create New Policy > Relationship Policy**.

The Relationship Policy page opens.

For information about how to create a relationship policy, see [Relationship Policy Configuration: Overview](#).

Related Tasks

[Relationship Policy Management: View Current Policies](#)

[Relationship Policy Management: View or Edit Event Details of a Relationship Policy](#)

[Relationship Policy Management: Find All Relationship Policies Associated With a Host Group](#)

[Relationship Policy Management: Edit a Policy](#)

Related Topics

[Relationship Policy Management: What You Can Do On This Page Overview](#)

[Relationship Policy Management: Overview](#)

Relationship Policy Management: Edit a Policy



You cannot change the host groups assigned to a relationship policy if that policy is linked to a map or network diagram.

1. In the applicable row, click the policy name in the Policy Name column.

The [Relationship Policy page](#) opens, showing the relationship policy's details populated in the applicable fields and all of the events related to the selected relationship policy listed alphabetically below that.

2. Expand the applicable row to view more details about an event and decide if you need to make edits to these settings as well.
3. Make your edits.
4. Click **Save**. (You can make edits on different pages before you click **Save**.)

Related Tasks

[Relationship Policy Management: View Current Policies](#)

[Relationship Policy Management: View or Edit Event Details of a Relationship Policy](#)

[Relationship Policy Management: Find All Relationship Policies Associated With a Host Group](#)

[Relationship Policy Management: Create a Policy](#)

Related Topics

[Relationship Policy Management: What You Can Do On This Page Overview](#)

[Relationship Policy Management: Overview](#)

Relationship Policy Configuration: Overview

Use this page to create and manage your relationship policies and individual events related to a policy.



For information about viewing and maintaining relationship policies, see [Relationship Policy Management: Overview](#).

Open Relationship Policy Configuration

Do any of the following to access this page

- From the main menu:
 1. Click **Configure > Detections > Policy Management**.
The Policy Management: Custom Events tab opens.
 2. Click the Relationship Events tab.
The Policy Management: Relationship Events tab opens.
 3. Click **Create New Policy** from the drop-down menu and select **Relationship Policy**.
The Relationship Policy page opens.
- On the Policy Management: Relationship tab, click the policy name.
The Relationship Policy page opens.

Relationship Configuration: Table of Contents

[Relationship Policy Management: Overview](#)

[Security Event List](#)

[Alarm List: Overview](#)

[About Alarm Categories](#)

[About Alarm Severities](#)

Relationship Policy Configuration: Filter Columns

Filter columns if you need to narrow your search.

At the top of each of the filterable columns, there is a filter field in which you can enter the criteria by which you want to filter that column. Some display a drop-down list as you begin to type and some do not.

For the Event, Policy Name, and Map or Diagram Name fields, as you start to type your entry in any of these fields, a drop-down list appears displaying a list of entries that match what you have begun to type. You can then choose the appropriate entry from the list so that you don't have to type the complete entry. The remaining columns dynamically filter to display only the entries that match your entry.

For the Host Groups, Traffic By Services, and Traffic By Applications fields, as you start to type your entry, the column dynamically filters to display a list of entries that match what you have begun to type. This match can be for any word in the entry (beginning, middle, or last word). Type the complete entry into the field to return results that are an exact match.

Filtering with multiple entries

You can type multiple entries within one column as well as enter values across multiple columns. When you type multiple values within one column, these values are ORed together; therefore, all returned results need only match one of the entries within the column. Remember to leave a space between each entry. When you filter more than one column, filters are ANDed together; therefore, all returned results must match the search criteria in all of the columns. Enclose an entry within quotation marks to return results that exactly match your entry.

For the Host Groups column, the table will first list all events that contain your entry as the first host group, and then will display all events that contain your entry as the second host group. Also, the Host Groups column will filter all host groups, not just those displayed in the table.

Editing a policy with filtered criteria

If you edit a relationship policy after having filtered criteria, the entries in the table does not change until you save your edits. Click the X in the field to clear the filter criteria and return to the previous view for that column. To quickly clear multiple filter criteria, refresh the page.

Clearing entries

Click the *X* in the field to clear the filter criteria and return to the previous view for that column. To quickly clear multiple filter criteria, refresh the page.

Related Topics

[Relationship Policy Configuration: Overview](#)

Relationship Policy Configuration: Sort Columns

To reverse the sort order for a column, click the column name.

Use the column name to toggle between ascending and descending order. You can also click the arrow that is displayed to the right of the column name (after you click the column name). An up arrow indicates that the column is sorted in ascending order, and a down arrow indicates descending order.

Related Topics

[Relationship Policy Configuration: Overview](#)

Relationship Policy Configuration: Concepts You Need to Understand

Refer to the following guidelines when configuring a relationship event:

- You can use the following units of measure: K, M, G, T, P, and E.
- You can enter the number of units (e.g., 1) with the unit of measurement immediately after it (no space).

Example: 10M, not 10 M

- You can enter decimal numbers only if you use a unit with it (e.g., 1.2 M, not 1.2).
- In all fields, 1K = 1000.
- When you enter a number over 999, Secure Network Analytics changes that number to the nearest unit after you have clicked away from that field.

*Example: 1000 changes to 1K; 1024 changes to 1.02K; 2,000,000,000 = 2 G;
2,000,000,000,000,000,000 = 2 E*

Related Topics

[Relationship Policy Configuration: What You Can Do On This Page Overview](#)

[Relationship Policy Configuration: Overview](#)

Relationship Policy Configuration: What You Can Do On This Page Overview

You can perform the following actions on this page:

- [Create or edit a relationship policy.](#)
- [Add an event to a relationship policy.](#)
- [View or edit the event setting details of a relationship policy.](#)
- [Enable or disable an event .](#)
- [Delete a relationship policy.](#)
- [Delete an event from a relationship policy.](#)
- [Duplicate a relationship policy.](#)

Related Topics

[Relationship Policy Configuration: Overview](#)

Relationship Policy Configuration: Create or Edit a Policy

To create or edit a relationship policy, do the following:

1. Complete the fields at the top of the page.
2. Add the desired events in the lower half of the page. See the next section, [Add an event to a relationship policy](#).

As you add events, they are displayed in the table, sorted by event.

After you have saved the information on this page, the Policy Manager: Relationship tab opens and each event you have configured for this relationship policy are displayed in the Core Policies table.



- To create a relationship policy, you must assign a name and two host groups and add a minimum of one event.
- To see the default values for the relationship policies, access this table: [Relationship Policy Default Values](#).
- For a description of the relationship events, see the Host Group Relationship Alarms section in the Alarm List topic in the Desktop Client Help.

Refer to the following information about the fields on this page:

Name (required)

You cannot duplicate a relationship policy name that already exists if that policy is associated with a map (in the Desktop Client) or network diagram (in the Manager).

Host Group - Side 1 (required), **Host Group - Side 2** (required)

This is the first host group and second host group, respectively, in the relationship pair.

Click the  (**Plus**) icon to open the Host Group panel to make your selection. You can select only one host group for side 1 and one host group for side 2.

If you type any letters that are anywhere within a host group name into the Search field and press **Enter**, any host group name that contains that entry expand and display in italicized red font.

Example: If you entered South into the Search field, some of the host group names that would then be displayed in italicized red font would be Southern Africa, South Africa, South-Eastern Asia, French Southern Territories.

To return an exact match, enter the entire host group name into the Search field and press **Enter**. If you want to return to the previously selected host group, delete the entry in the Search field and press **Enter**.

Map or Diagram Name

Indicates the name of the configured network diagram (in the Desktop Client they are called maps).

For more information about network diagrams, see the "What is a network diagram" topic in.

Description

Indicates the purpose of this relationship policy.

Traffic by Services and Applications

This setting enables you to create a relationship policy with events that monitor the traffic based on services and/or applications.

1. Click the **Plus** icon. From the drop-down list, select one of the following options:

- Include All Except
- Exclude All Except

Whichever option you choose applies to both services and applications. So, for instance, you cannot choose to use Include All Except for services and Exclude All Except for applications. If you choose the Include All Except option, for instance, the selections you make from both the Services tab and the Application tab use that option to determine which selections are used in your configuration.

After you have chosen one of these options, a panel opens for you to make your selections for both services and applications. The Services tab is selected by default. The subtitle at the top of the panel reflects the option you selected from the drop-down list.

2. Make the appropriate selections to ensure that the service traffic and the application traffic that you want observed for this relationship policy are returned. If you type the beginning letters or the complete application name or service name, it moves to the top of the list.

Secure Network Analytics displays up to three services or applications defined for the relationship policy in this field. The number of any remaining number of services or applications is displayed within parentheses, and a link entitled "More" is displayed beside this number. If you click this link, you see a list of all the services or applications that comprise the traffic between the specified host groups.

The total number of defined services and applications are displayed in parentheses on the associated tab. If a tab does not contain any defined entries, (0) is displayed on the tab.

Related Topics

[Relationship Policy Configuration: Add an Event to a Relationship Policy](#)

[Relationship Policy Configuration: View or Edit Event Setting Details of a Relationship Policy](#)

[Relationship Policy Configuration: Enable or Disable an Event](#)

[Relationship Policy Configuration: Delete a Relationship Policy](#)

[Relationship Policy Configuration: Delete an Event from a Relationship Policy](#)

[Relationship Policy Configuration: Duplicate a Relationship Policy](#)

[Relationship Policy Configuration: Overview](#)

Relationship Policy Configuration: Add an Event to a Relationship Policy

View or edit the event setting details of a relationship policy

To view or edit the event setting details for a relationship policy, click the arrow at the beginning of the row.

The row expands to display the settings you have configured for the event.

When you edit any settings or add any events to an existing policy, a green dog ear is displayed in the top right corner of the row. A red dog ear indicates an input error. Hover your cursor over a field name to see the minimum and maximum allowable values.

Refer to the following table for information about the event setting detail values.

Option Name	Description
Behavioral and Threshold	When this option is selected, the system uses the tolerance setting, minimum threshold, and maximum threshold. Note: When you change your selection from "Behavioral and Threshold" to "Threshold Only," the following fields become disabled: • Tolerance • Never trigger alarm when less than
Threshold only	The alarm triggers only when a minimum or maximum threshold is reached.
Tolerance	Tolerance is the relative number between 0 and 100 that indicates how much to allow actual behavior to exceed expected behavior before alarming. This allows the user to define what is "significantly different." It can also be defined as "the number of standard deviations from the norm." Note: Standard deviation shows how much variation there is from the average (i.e., mean, or expected value). A low standard deviation indicates that the data points tend to be very close to the mean, whereas high standard deviation indicates that the data points are spread out over a large range of values.

Option Name	Description
	• A tolerance of 0 causes an alarm to be triggered for any values over the expected value; it is very sensitive and results in a lot of alarms. • A tolerance of 100 allows the behavior to exceed expected behavior to a high degree before it causes an alarm to be triggered, thus greatly reducing the number of times an alarm is triggered. If you never want the alarm to trigger, you must disable the alarm.
Never trigger alarm when less than	Also known as the minimum threshold, this is a static value that indicates the lowest value to allow without triggering an alarm. The alarm does not trigger when the observed value falls below this setting. In other words, even if a host is greatly over its expected value, if it is not more than the minimum indicated in this dialog, then do not trigger an alarm.
Always trigger alarm when greater than	Also known as the maximum threshold, this is a static value that indicates the highest value to allow without triggering an alarm. The alarm triggers when the observed value exceeds this setting. In other words, if a host's value exceeds the maximum indicated in this dialog, even if it is expected for that host, then trigger an alarm.
Trigger alarm when duration greater than	Duration is the amount of time that the associated value can exceed the calculated threshold before the Manager triggers an alarm.

Related Topics

[Relationship Policy Configuration: Add an Event to a Relationship Policy](#)

[Relationship Policy Configuration: View or Edit Event Setting Details of a Relationship Policy](#)

[Relationship Policy Configuration: Enable or Disable an Event](#)

[Relationship Policy Configuration: Delete a Relationship Policy](#)

[Relationship Policy Configuration: Delete an Event from a Relationship Policy](#)

Relationship Policy Configuration: Duplicate a Relationship Policy

Relationship Policy Configuration: Overview

Relationship Policy Configuration: View or Edit Event Setting Details of a Relationship Policy

To view or edit the event setting details for a relationship policy, click the arrow at the beginning of the row.

The row expands to display the settings you have configured for the event.

When you edit any settings or add any events to an existing policy, a green dog ear is displayed in the top right corner of the row. A red dog ear indicates an input error. Hover your cursor over a field name to see the minimum and maximum allowable values.

Refer to the following table for information about the event setting detail values.

Option Name	Description
Behavioral and Threshold	When this option is selected, the system uses the tolerance setting, minimum threshold, and maximum threshold. Note: When you change your selection from "Behavioral and Threshold" to "Threshold Only," the following fields become disabled: • Tolerance • Never trigger alarm when less than
Threshold only	The alarm triggers only when a minimum or maximum threshold is reached.
Tolerance	Tolerance is the relative number between 0 and 100 that indicates how much to allow actual behavior to exceed expected behavior before alarming. This allows the user to define what is "significantly different." It can also be defined as "the number of standard deviations from the norm." Note: Standard deviation shows how much variation there is from the average (i.e., mean, or expected value). A low standard deviation indicates that the data points tend to be very close to the mean, whereas high standard deviation indicates that the data points are spread out over a large range of values. • A tolerance of 0 causes an alarm to be triggered for any

Option Name	Description
	values over the expected value; it is very sensitive and results in a lot of alarms. • A tolerance of 100 allows the behavior to exceed expected behavior to a high degree before it causes an alarm to be triggered, thus greatly reducing the number of times an alarm is triggered. If you never want the alarm to trigger, you must disable the alarm.
Never trigger alarm when less than	Also known as the minimum threshold, this is a static value that indicates the lowest value to allow without triggering an alarm. The alarm does not trigger when the observed value falls below this setting. In other words, even if a host is greatly over its expected value, if it is not more than the minimum indicated in this dialog, then do not trigger an alarm.
Always trigger alarm when greater than	Also known as the maximum threshold, this is a static value that indicates the highest value to allow without triggering an alarm. The alarm triggers when the observed value exceeds this setting. In other words, if a host's value exceeds the maximum indicated in this dialog, even if it is expected for that host, then trigger an alarm.
Trigger alarm when duration greater than	Duration is the amount of time that the associated value can exceed the calculated threshold before the Manager triggers an alarm.

Related Topics

[Relationship Policy Configuration: View or Edit Event Setting Details of a Relationship Policy](#)

[Relationship Policy Configuration: Add an Event to a Relationship Policy](#)

[Relationship Policy Configuration: Enable or Disable an Event](#)

[Relationship Policy Configuration: Delete a Relationship Policy](#)

[Relationship Policy Configuration: Delete an Event from a Relationship Policy](#)

[Relationship Policy Configuration: Duplicate a Relationship Policy](#)

Relationship Policy Configuration: Overview

Relationship Policy Configuration: Enable or Disable an Event

To disable an event that is enabled, click the  (**Toggle**) icon in the Status column so that the status displays *Off*. To enable the event, click the switch so that the status displays *On*. You must click **Save** after enabling or disabling for your change to take effect.

Related Topics

[Relationship Policy Configuration: Enable or Disable an Event](#)

[Relationship Policy Configuration: Add an Event to a Relationship Policy](#)

[Relationship Policy Configuration: View or Edit Event Setting Details of a Relationship Policy](#)

[Relationship Policy Configuration: Delete a Relationship Policy](#)

[Relationship Policy Configuration: Delete an Event from a Relationship Policy](#)

[Relationship Policy Configuration: Duplicate a Relationship Policy](#)

[Relationship Policy Configuration: Overview](#)

Relationship Policy Configuration: Delete a Relationship Policy



You cannot delete a relationship policy that is linked to a map or network diagram.

From the Actions drop-down menu, click **Delete**. The relationship policy is deleted from the Relationship Policy table.

Related Topics

[Relationship Policy Configuration: Delete a Relationship Policy](#)

[Relationship Policy Configuration: Add an Event to a Relationship Policy](#)

[Relationship Policy Configuration: View or Edit Event Setting Details of a Relationship Policy](#)

[Relationship Policy Configuration: Enable or Disable an Event](#)

[Relationship Policy Configuration: Delete an Event from a Relationship Policy](#)

[Relationship Policy Configuration: Duplicate a Relationship Policy](#)

[Relationship Policy Configuration: Overview](#)

Relationship Policy Configuration: Delete an Event from a Relationship Policy

To delete one or more events from a relationship policy, click **Delete** in the Actions column of the row for the applicable event.

*After you have deleted the desired events from the relationship policy and clicked **Save**, the Policy Manager: Relationship page opens again, and the events you deleted are no longer displayed. Also, the number (in the parentheses) on the Policy Manager: Relationship tab now reflect the new number of relationship events in your system.*

Related Topics

[Relationship Policy Configuration: Delete an Event from a Relationship Policy](#)

[Relationship Policy Configuration: Add an Event to a Relationship Policy](#)

[Relationship Policy Configuration: View or Edit Event Setting Details of a Relationship Policy](#)

[Relationship Policy Configuration: Enable or Disable an Event](#)

[Relationship Policy Configuration: Delete a Relationship Policy](#)

[Relationship Policy Configuration: Duplicate a Relationship Policy](#)

[Relationship Policy Configuration: Overview](#)

Relationship Policy Configuration: Duplicate a Relationship Policy

Duplicate a relationship policy if you want to add a relationship policy that is similar to an existing policy so you can use it as a template and therefore create the new policy more easily and quickly.

To duplicate a relationship policy, do the following:

1. From the Actions drop-down menu, click **Duplicate**.

The Relationship Policy page opens, showing the policy's details populated in the applicable fields and all of the events related to the selected relationship policy listed below that.

2. Change the policy name (each policy name must be unique).
3. Assign host groups 1 and 2.
4. Make any other edits, if necessary.
5. Click **Save**.

Related Topics

[Relationship Policy Configuration: Duplicate a Relationship Policy](#)

[Relationship Policy Configuration: Add an Event to a Relationship Policy](#)

[Relationship Policy Configuration: View or Edit Event Setting Details of a Relationship Policy](#)

[Relationship Policy Configuration: Enable or Disable an Event](#)

[Relationship Policy Configuration: Delete a Relationship Policy](#)

[Relationship Policy Configuration: Delete an Event from a Relationship Policy](#)

[Relationship Policy Configuration: Overview](#)

Relationship Policy Configuration: Synchronize Your Flow Collector with the Manager

If your policies are not behaving correctly, synchronize your Flow Collector with the Manager

At times you may need to synchronize the Flow Collectors on your network with your Manager. Some reasons may be as follows:

- You observe that your Flow Collectors are not working properly.
- You have added a new configuration (e.g., role policy, host group), but while you were doing so your Manager stopped communicating with any or all of your Flow Collectors. In this case, the policy will have been saved to your Manager, but it may not have propagated to any or all of your Flow Collectors.

To synchronize one or more Flow Collectors to your Manager, log in to your Manager client interface and complete the following steps:

1. From the navigation menu, click **Configure > Flow Collectors**.
2. Click the **Advanced** tab.
3. In the top left corner, from the Flow Collector drop-down list, choose the applicable Flow Collector.
4. In the Synchronize section, click **Synchronize**.
5. Repeat Step 3 through Step 4 for the remaining Flow Collectors that need to be synchronized with the Manager.

Related Topics

[Relationship Policy Configuration: What You Can Do On This Page Overview](#)

[Relationship Policy Configuration: Overview](#)

Managing Core Policies



- For information about configuring role policies, see [Configuring Role Policies](#).
- For information about configuring host policies, see [Configuring Host Policies](#).

Use this page to view and make edits to the core policies you have created for your network. You can use core policies to control how the system monitors and responds to host behavior that it observes. Core policies can reduce unwanted alarms on your system as well as ensure that alarms are triggered in certain instances (when configured criteria are observed). If you see an alarm triggered for a particular host that should not be triggered or should be triggered at a different threshold, you should consider modifying the effective policy for that particular host.

All role or host policies created before v7.0 retain their mitigation settings, if any have been configured. Note that you can view and update these mitigation settings only in the Desktop Client and not in the Web App.

Any core policies you add to either the Web App or Desktop Client are displayed in both interfaces.

How to find this page

Do any of the following to access this page

- From the main menu on any page,
 1. Choose **Configure > Detection > Policy Management**.
 2. Click the Core Events tab.

The Policy Management: Core Events tab opens.
- From the context menu for any IP address in the Web App, choose **Show Effective Policy**.

The Policy Management: Core Events tab opens. The IP address from which you pivoted is displayed in the search box at the top of the page, and all core events related to that IP address are listed in the Core Events table. (The relationship events on the Policy Management: Relationship tab cannot filter on IP addresses.)
- On the Alarm List, from the Actions drop-down menu for the applicable host, role, or default event, choose **Tune Alarm**.

Depending on which policy type triggered the event, you are redirected to either the Policy Management: Custom Events tab, Relationship tab, or Core Events tab. If it was a core event that triggered the alarm, then that core event type is displayed in the Event column's filter box, the policy name is displayed in the Policy Name column's filter box, and the core policy for that event is displayed in the Core Events table.

- On the Security Events Details page or the Top Security Event component on the Host Report, from the Actions drop-down menu for the applicable host, role, or default event, choose **Tune Event**.

Depending on which policy type triggered the event, you are redirected to either the Policy Management: Custom Events tab, Relationship tab, or Core Events tab. If it was a core event that triggered the alarm, then that core event type is displayed in the Event column's filter box, the IP address is displayed in the Search box, and the core policy for that event is displayed in the Core Events table.

- On the Host Group Report page, from the Actions drop-down menu, choose **Policy Management**.

The Policy Management: Core Events tab opens by default. The host group from which you pivoted is displayed in the search box at the top of the page, and all custom, relationship, and core policies that are applied to this host group are displayed on the associated tab.



- The number displayed on the tab indicates the number of defined core events.
- To view a description of a core event, hover your cursor over the associated row. Ensure you are not hovering over any text.

By default, the table on this page displays every policy, along with every event that has been configured for each policy, configured for your system. The results are sorted in the following order:

1. Event, alphabetically
2. Policy Type (first host, then role, and lastly default)
3. Most restrictive to least restrictive policy, which is determined by the combination of settings between the "When Host is Source" column and the "When Host Is Target" column. For more information, see [Determining the effective policy when two or more role policies exist for a host](#) and [Find the effective policy for a host](#) in this topic.

Using the menu accessed from the ellipsis icon

When you click an ellipsis icon in the Web App, a menu containing submenu options opens. If you want to open a page in a new tab when choosing one of these options, do one of the following:

- On a Windows-based PC, perform a Ctrl+Click.
- On a Mac, perform a Cmd+Click.



This functionality does not work if you are connected to a virtual machine.

Concepts you need to understand to use this page

What are the three core policy types

Secure Network Analytics uses the following three types of core policies, which you can modify whenever needed. These policies are listed in order of precedence; the one that takes the highest precedence is listed first.

This type of policy...	Pertains to...
Host	Specific IP addresses. A host cannot be assigned to more than one host policy. The host policy always overrides the default policy or any role policy that may apply to the host of this host policy, even if you disable, enable, or change any of the security event settings in any of these policies.
Role	Groups of hosts (IP addresses) that serve a common purpose (for example, Web servers, firewalls, trusted Internet hosts, etc.). The role policy always overrides the default policy, even if you disable, enable, or change any of the security event settings in either policy.
Default	All Inside and Outside hosts. The default policy applies when no role policies or host policies have been created. For descriptions of all the default security events for Secure Network Analytics, see Security Event List.

What are the three host alarm types

Host alarms can be divided into the following sub-categories based on configuration type.

Behavior-based security events

Some events are triggered when host activity experiences a significant change from past behavior. You can change the tolerance (i.e., sensitivity) of these types of events using Policy Management. The main advantage of this approach is that you can adjust the system so that the number of alarms that occur matches your organizational needs. That is, if you prefer many alarms (i.e., you can tolerate only a slight change from the expected behavior), you can lower the tolerance setting in the associated policy. Conversely, if you prefer fewer alarms (i.e., you can tolerate significant change from the expected behavior), you can raise the tolerance setting. Essentially, each of the settings for the behavior-based event is attached to a numerical value, and you can adjust this value up or down.

Behavior-based events use a baseline based on recent activity and a configured tolerance. This gives a host the ability to change its behavior over time without losing the capability to alarm if its behavior changes too radically. Tolerance provides a way to control how much change is acceptable. In essence, you have the ability to adjust the sensitivity of the event's threshold level (i.e., to "turn down the noise" to whatever level you want).

With behavior-based events, a host must reach a certain level of deviation from its baseline before an alarm triggers. As you increase the tolerance for an event, you see less alarms, and as you decrease the tolerance, you see more alarms. For more information about baselining, see "Baselining" in [Managing and Configuring Host Groups](#).

For information about the various settings used in behavior-based events, refer to the following table.

Option Name	Description
Behavioral and Threshold	When this option is selected, the system uses the tolerance setting, minimum threshold, and maximum threshold. Note: When you change your selection from "Behavioral and Threshold" to "Threshold Only," the following fields become disabled:

Option Name	Description
	• Tolerance • Never trigger alarm when less than
Threshold only	The alarm triggers only when a minimum or maximum threshold is reached.
Tolerance	Tolerance is the relative number between 0 and 100 that indicates how much to allow actual behavior to exceed expected behavior before alarming. This allows the user to define what is "significantly different." It can also be defined as "the number of standard deviations from the norm." Note: Standard deviation shows how much variation there is from the average (i.e., mean, or expected value). A low standard deviation indicates that the data points tend to be very close to the mean, whereas high standard deviation indicates that the data points are spread out over a large range of values. • A tolerance of 0 causes an alarm to be triggered for any values over the expected value; it is very sensitive and results in a lot of alarms. • A tolerance of 100 allows the behavior to exceed expected behavior to a high degree before it causes an alarm to be triggered, thus greatly reducing the number of times an alarm is triggered. If you never want the alarm to trigger, you must disable the alarm.
Never trigger alarm when less than	Also known as the minimum threshold, this is a static value that indicates the lowest value to allow without triggering an alarm. The alarm does not trigger when the observed value falls below this setting. In other words, even if a host is greatly over its expected value, if it is not more than the minimum indicated in this dialog, then do not trigger an alarm.
Always	Also known as the maximum threshold, this is a static value that

Option Name	Description
trigger alarm when greater than	indicates the highest value to allow without triggering an alarm. The alarm triggers when the observed value exceeds this setting. In other words, if a host's value exceeds the maximum indicated in this dialog, even if it is expected for that host, then trigger an alarm.
Trigger alarm when duration greater than	Duration is the amount of time that the associated value can exceed the calculated threshold before the Manager triggers an alarm.

Below are two examples of the Edit Settings dialog for a behavior-based event. In the first image the Behavioral and Threshold option is selected, and in the second image the Threshold Only option is selected.

The following images are examples of the ICMP Received security event.

The first screenshot shows the 'Edit Settings' dialog for the 'ICMP Received' event. The 'POLICY TYPE' is set to 'Host'. The 'WHEN HOST IS SOURCE' and 'WHEN HOST IS TARGET' are both set to 'On'. The 'Behavioral and Threshold' option is selected. The 'Tolerance' is 75 / 100. The 'Never trigger alarm when less than' field is 1000, and the 'Always trigger alarm when greater than' field is 1000. The description states: 'The host has received ICMP packets above the acceptable number in a 5-minute period.'

The second screenshot shows the same dialog, but the 'Threshold Only' option is selected. The 'Tolerance' is 75 / 100. The 'Never trigger alarm when less than' field is 1000, and the 'Always trigger alarm when greater than' field is 1000. The description states: 'The host has received ICMP packets above the acceptable number in a 5-minute period.'

When you edit any settings or add any events to an existing policy, a green dog ear is displayed in the top right corner of the row (refer to the first image). Also, when you change your selection from "Behavioral and Threshold" to "Threshold Only," the following fields are disabled and populated with its respective default value (refer to the second image).

- Tolerance
- Never trigger alarm when less than

Condition-based security events

The second type of event is the event which you can either turn on or off. There are no settings associated with this type of event. The criteria for triggering an alarm for this type of alarm differs from that of the behavior-based event. When an event simply has an on/off setting, the host's behavior must match certain conditions that must all coincide with each other before this type of event triggers an alarm. If all of these conditions do not exist, then an alarm does not trigger. For instance, in order for a worm activity event to trigger an alarm, all of the following behavior must have occurred:

- A source host has scanned multiple sub-nets.
- At least one of the target hosts has connected to the source host.
- This target host has transferred information to the source host.

If even one of these conditions does not exist, the alarm does not trigger.

You can determine which events are behavior-based and which events are condition-based by expanding an entry in the System table. If an event is behavior-based, then the tolerance values that have been indicated for that event are displayed. If an event is condition-based, then the entry "NA" is displayed.

Below is an example of the verbiage displayed in the expanded section for a condition-based security event. These type of events do not have a Edit Settings dialog.

All other alarms

The events that do not fall into the previous two categories are generated based on any of the following parameters:

- A minimum threshold – a static value that indicates the lowest value to allow without triggering an alarm.
- maximum threshold – a static value that indicates the highest value to allow without triggering an alarm.
- Duration – the amount of time that the associated value can exceed the calculated threshold before triggering an alarm.

Below are examples of what the Edit Settings dialogs for this type of security event can look like.

EVENT	EVENT TYPE	POLICY NAME	POLICY TYPE	HOSTS	WHEN HOST IS SOURCE	WHEN HOST IS TARGET
SSH Reverse Shell	Ex. Category	Ex. Outside Hosts	Ex. Role	Ex. Network Scanners	Ex. On + Alarm	Ex. On + Alarm
SSH Reverse Shell	Security	0.40	Host	0.40	On	On
Description Detects an SSH session that appears to be a reverse shell. More data is being sent to the outside host than is being received. Trigger alarm when traffic exceeds this value: 1000 AND, the percentage of client data sent exceeds: 60 Do not trigger alarm when percentage of client data sent exceeds: 80						

EVENT	EVENT TYPE	POLICY NAME	POLICY TYPE	HOSTS	WHEN HOST IS SOURCE	WHEN HOST IS TARGET
Talks to Phantoms	Ex. Category	Ex. Outside Hosts	Ex. Role	Ex. Network Scanners	Ex. On + Alarm	Ex. On + Alarm
Talks to Phantoms	Security	0.40	Host	0.40	On	On
Description A host that is talking to other hosts who have not been seen before and do not respond. Trigger alarm when the number of phantom hosts that a particular host has attempted to reach exceeds: 3 Notes: - To reduce false positives, increase the number of attempted connections required to trigger the alarm.						

What are the three security event types

One-to-one

This security event type involves one source host initiating a suspicious conversation with one other host.

Common one-to-one security events include the following:

- Port Scan
- Beaconing Host
- Brute Force Login
- Suspect Long Flow
- Ping

Example:

Hosts running scripts that frequently log in and out of another system are prone to trigger a Brute Force Login alarm. People do not want to be bothered when these systems serve their purpose, but still want to know if someone tries to brute force into the box.

1. Alarm at source and at target.
2. Contribute points at source (but do not alarm) and contribute points to target (but do not alarm).
3. Turn off at source and at target.

	WHEN HOST IS SOURCE	WHEN HOST IS TARGET
1	Ex. On + Alarm On + Alarm	Ex. On + Alarm On + Alarm
2	Ex. On + Alarm On	Ex. On + Alarm On
3	Ex. On + Alarm Off	Ex. On + Alarm Off

Many-to-one

This security event type involves many sources initiating suspicious conversations with the same host in a short time. You cannot alarm at the source with one-to-many events.

Common many-to-one security events include the following:


- SYNs Received
- Max Flows Served
- UDP Received
- ICMP Received

Example:

You may want to do the following:

1. Track large quantities of SYN packets sent from external hosts to internal hosts.
2. Not track large quantities of SYN packets sent from approved internal network scanners to internal hosts.

To achieve this, you can set up different policies for the SYNs Received security event for each host as demonstrated below. Any of these policies can be a host, role, or default policy; it does not matter. What is important is that each policy is the effective policy for the associated host.



One-to-many

This security event type involves one source host initiating a suspicious conversation with many other hosts. You cannot alarm at the target with one-to-many events.

Common one-to-many security events include the following:

- Addr_Scan/tcp
- Addr_Scan/udp


- Suspect Data Hoarding
- High SMB Peers
- Max Flows Initiated

Example:

You may want to track the following:

1. When an external host is scanning an internal host.
2. When an external host is scanning a network scanner.
3. When an internal host is scanning a network scanner.

To achieve this, you can set up different policies to detect scanning for each host as demonstrated below. Any of these policies can be a host, role, or default policy; it does not matter. What is important is that each policy is the effective policy for the associated host.



Effective policy for the external host

Alarm when this host is the source (we want to know) and don't alarm when this host is the target (we don't care).

Effective policy for the network scanner

Don't alarm when this host is the source (since this is expected behavior for the network scanner) and enable index points (On) to contribute to the target (we want to collect index points but we don't want to receive an alarm).

Effective policy for the internal host

Alarm when this host is the source (we want to know) and enable index

points (*On*) to contribute to the target (we want to collect index points but we don't want to receive an alarm).

When an internal host (in this case, the source) scans an *external host*, we don't want to receive an alarm. However, when an internal host (again, the source) scans a *network scanner*, we do want to receive an alarm.

Therefore, we need to set the "When Host Is Source" option to *On + Alarm* to cover this type of scenario.

What is Directional Policy

Directional Policy is not used with custom security events.

Directional Policy is a solution to either of the following scenarios when you are investigating a security event:

- In general, the behavior you discovered is of concern to you, but not for the current situation.
- The behavior you discovered is of concern to you in only a few specific situations.

Secure Network Analytics security events possess two sides:

- Source - this is generally the client(s).
- Target - this is generally the server(s).

Prior to v6.9, the basis upon which Secure Network Analytics triggered security events was not uniform across types of events. Some events considered the policy associated with the source(s), some considered the policy associated with the target(s), and yet others considered the policy of both the source and target. As of v6.9, nearly all the security events now check both the source host's and target host's policies independently in order to determine how to proceed.

Directional Policy has also changed how Secure Network Analytics makes policy decisions for "multiple hosts" security events (events that are either "One-to-Many" or "Many-to-One"). "One-to-Many" security events involve one source host but more than one target host. "Many-to-One" security events involve more than one source host and one target host. For more information about "One-to-Many" and "Many-to-One" alarms, see [What are the three security event types](#) in this topic.

For example, consider the SYNs Received security event, which tracks how many SYN packets a single server receives in a five-minute period from any source. Prior to v6.9, you could not use a policy to direct the system to ignore SYN packets that a particular source was sending. Beginning with v6.9, Secure Network Analytics considers both the

source host and target host policies (as described in the previous scenario). If, for example, you want to track SYNs Received using a web server, but do not want SYN packets from your network scanners to count in the tracking process, you can now make that distinction without turning off the event.

How do I know which options to select for source and target

Following are the options you can select for both the source host and the target host on this page. (You can also edit these options on the [Host Policy](#), [Role Policy](#), and [Default Policy](#) pages.) The following list shows the most restrictive setting (least likely to trigger an alarm) at the top and ends with the least restrictive setting (most likely to trigger an alarm). In addition to using this hierarchical list of options, Secure Network Analytics system also takes into consideration whether the policy is a host, role, or default policy, when determining the effective policy. For more information about policy type hierarchy, see [What is an effective policy](#) in this topic.

- Off
- Ignore
- On
- On + Alarm

Refer to the following table for an explanation of the options.

If this option...	Is selected for "When Host Is Source," then this security event...	Is selected for "When Host is Target," then this security event...
Off	Does not collect index points nor alarm.	Does not contribute index points to alarm category events nor alarm.
Ignore	Is bypassed for the associated policy. Secure Network Analytics ignores the settings configured for this event in this policy (to which the associated host belongs). It	Is bypassed for the associated policy. Secure Network Analytics ignores the settings configured for this event in this policy (to which the associated host belongs). It

	instead uses the settings configured for this event in the next policy (in order of precedence) to which the associated host belongs. For more information about order of precedence used for host policies, refer to the "What are the three core policy types" section in this topic. For example, if the ignored policy is a host policy, Secure Network Analytics uses the settings configured for this event in the role policy to which the associated host belongs. If no such role policy exists, then Secure Network Analytics uses the settings configured for this event in the default policy.  The Ignore option does not apply to alarm category events nor is not available for default policies.	instead uses the settings configured for this event in the next policy (in order of precedence) to which the associated host belongs. For more information about order of precedence used for host policies, refer to the "What are the three core policy types" section in this topic. For example, if the ignored policy is a host policy, Secure Network Analytics uses the settings configured for this event in the role policy to which the associated host belongs. If no such role policy exists, then Secure Network Analytics uses the settings configured for this event in the default policy.
On	Collects index points but does not alarm.	Contributes index points to alarm category events but does not generate alarms.
On + Alarm	Collects index points and generates alarms.	Contributes index points to alarm category events and generates alarms.

Default Policy, Inside Hosts

When you set the Target to OFF for a default policy, it disables the alarm for all inside hosts, even if the policy is set to alarm on the Source, because the Target is also a member of the Inside Hosts group.

In this case, you would create a role policy which would contain the applicable IP addresses, and then you would set the Source to ON and the Target to OFF.

Example:

Say you want to turn off the Brute Force Login security event when your monitoring server is the source (since it is frequently logging in to devices), but you still want to know if anyone tries to brute force their way into it. Provided you had your monitoring servers assigned to a host group, you could [create a role policy](#) for that host group and add the Brute Force Login security event to that policy. You would then configure the settings as shown in the following table.

The settings for Brute Force Login security event for this role policy would be...	
When Host Is Source	When Host Is Target
Off (This would prevent Bruce Force Login from firing when your monitoring server was the source.)	On + Alarm (This would cause Bruce Force Login to trigger an alarm when your monitoring server was the target.)

What is an effective policy

The effective policy is derived from all policies defined for a host. The effective policy represents the policy settings for a host which have been determined by the system to take precedence over all other defined settings. The effective policy determines if and when any alarms are triggered by the associated host. The effective policy is associated with only a host, not a host group.

There may be situations where the following is true:

- A host is not assigned to a host policy but is assigned to two or more differently configured role policies.
- A host has one policy for when it is the source and one policy for when it is the target.

In the first instance, since more than one role policy can be applied to a specific host and the threshold settings for each policy can be different, multiple alarms could be triggered if the host behavior exceeds values defined within each role policy. In the second instance, depending on whether the host is the source or the target, a different policy would apply.

Secure Network Analytics takes into consideration both the policy type and the settings configured for the source host and target host when determining the effective policy. These are discussed below. For more information about determining the effective policy, see [Determining the effective policy when two or more role policies exist for a host](#) and [Find the effective policy for a host](#) in this topic.

Policy type

A [host policy](#) takes precedence over all other policies, and a role policy takes precedence over a default policy. Therefore, a host policy is more specific than a [role policy](#), and a role policy is more specific than a [default policy](#). Only the most specific policy settings for a host trigger an alarm.

Order of precedence for policy type:

- host
- role
- default

For instance, if you add a [security event](#) to a role policy, this role policy always overrides the default policy, even if you disable, enable, or change any of the security event settings in either policy. Likewise, if you add a security event to a host policy, the host policy always overrides the default policy or any role policy that may apply to the host of this host policy, even if you disable, enable, or change any of the security event settings in any of these policies.

So, if a host is associated with one or more host policies, Secure Network Analytics would look first at those host policies before moving on to any role policies associated with the host to determine the effective policy or policies. And, if Secure Network Analytics determines that no host policy or role policy could be used as an effective source or target policy, it would at that point use the default policy, which is the last type to be considered.

"When Source is Host" and "When Source is Target" settings

If a host is not assigned to a host policy but is assigned to two or more role policies, Secure Network Analytics determines for each alarm which policy's settings are used in the host's effective policy by using a hierarchical list of options, shown below. This list

shows the most restrictive setting (least likely to trigger an alarm) at the top and ends with the least restrictive setting (most likely to trigger an alarm). For more information about these options, see [How do I know which options to select for source and target](#) this topic.

- Off
- Ignore
- On
- On + Alarm

How do I determine the effective policy when two or more role policies exist for a host

If necessary, review the [What is an effective policy](#) section in this topic to ensure your understanding of the following examples. You can also find additional information in [Find the effective policy for a host](#).

Assume that all policies apply to the same event (e.g., Port Scan) in each scenario.

Scenario 1

In this scenario, assume the host belongs to a host policy as well as two role policies.

If the settings for the Host Policy are...	
When Host Is Source	When Host Is Target
Ignore	Off
And the settings for Role Policy 1 are...	
When Host Is Source	When Host Is Target
Off	On + Alarm
And the settings for Role Policy 2 are...	
When Host Is Source	When Host Is Target
On + Alarm	Ignore

Then the effective policy for the source host is...	Then the effective policy for the target host is...
Role Policy 1 Since a host policy exists, the system looks at that policy first. Since the source in the host policy says <i>Ignore</i>, the system moves on to the role policies. <i>Off</i> takes precedence over <i>On + Alarm</i>, so Role Policy 1 is the effective source policy.	the Host Policy Since a host policy exists, the system looks at that policy first. Since the source setting in the host policy says <i>Off</i>, the system uses the host policy as the effective target policy and does not move on to look at the role policies.

In the remaining scenarios, assume the host belongs to two role policies.

Scenario 2

If the settings for Role Policy 1 are...	
When Host Is Source	When Host Is Target
Off	On + Alarm
And the settings for Role Policy 2 are...	
When Host Is Source	When Host Is Target
On + Alarm	Ignore
Then the effective policy for the source host is...	Then the effective policy for the target host is...
Role Policy 1	Role Policy 1

Scenario 3

If the settings for Role Policy 1 are...	
When Host Is Source	When Host Is Target

If the settings for Role Policy 1 are...	
On	Off
And the settings for Role Policy 2 are...	
When Host Is Source	When Host Is Target
On + Alarm	Ignore
Then the effective policy for the source host is...	Then the effective policy for the target host is...
Role Policy 1	Role Policy 1

Scenario 4

If the settings for Role Policy 1 are...	
When Host Is Source	When Host Is Target
On	Ignore
And the settings for Role Policy 2 are...	
When Host Is Source	When Host Is Target
Ignore	Ignore
Then the effective policy for the source host is...	Then the effective policy for the target host is...
Role Policy 1	The Default Policy Since <i>Ignore</i> tells the system to not use either of these role policies for the effective target policy, the system moves on to the default policy for the effective target policy.

Scenario 5

If the settings for Role Policy 1 are...	
When Host Is Source	When Host Is Target
Off	NA
And the settings for Role Policy 2 are...	
When Host Is Source	When Host Is Target
On	NA
Then the effective policy for the source host is...	Then the effective policy for the target host is...
Role Policy 1	There is no effective target policy. Whenever the source and target settings say <i>NA</i> in any host's policy for a particular security event, this indicates that <i>NA</i> is also the setting in all the policies for this security event. Therefore, there is no effective policy in this case.

Scenario 6

If the settings for Role Policy 1 are...	
When Host Is Source	When Host Is Target
On	NA
And the settings for Role Policy 2 are...	
When Host Is Source	When Host Is Target
On	NA

If the settings for Role Policy 1 are...	
Then the effective policy for the source host is...	Then the effective policy for the target host is...
Secure Network Analytics considers both policies to be equivalent, so it arbitrarily chooses one to use as the effective policy.	There is no effective target policy. Whenever the source and target settings say <i>NA</i> in any host's policy for a particular security event, this indicates that <i>NA</i> is also the setting in all the policies for this security event. Therefore, there is no effective policy in this case.

What you can do on this page

View or edit the current core policies

By default, the table on this page displays every policy, along with every event that has been configured for each policy, configured for your system. The results are sorted in the following order:

1. Event, alphabetically, so that you can easily see all the policies related to a specific security event type.
2. Policy Type (first host, then role, and lastly default)
3. Most restrictive to least restrictive policy, which is determined by the combination of settings between the "When Host is Source" column and the "When Host Is Target" column. For more information, see [Determining the effective policy when two or more role policies exist for a host](#) and [Find the effective policy for a host](#) in this topic.

The following list shows the most restrictive setting at the top and ends with the least restrictive setting.

- Off
- Ignore
- On
- On + Alarm

Refer to the following information about the fields on this page.

Event

Security event name.

Event Type

Indicates the event type: alarm category (displayed as *Category*) or security event (displayed as *Security*).

Policy Name

This must be a unique name.

- To see view the description for the policy, hover your cursor over the policy name. (You define the description for a policy when you create the policy.)
- To view the details for this policy and all of the events related to the policy, click the policy name.

Policy Type

Indicates whether this is a host, role, or default policy. To see an image that illustrates the hierarchical relationship between the three policy types, click the  (**Information**) icon next to this field.

Hosts

IP address(es) or host group(s) to which the core policy applies. If the core policy applies to more than three hosts, the remaining number of hosts is displayed within parentheses, and a link entitled "More" is displayed beside this number. If you click this link, the Hosts panel (non-editable) opens on which you can see all the host groups and/or hosts to which the core policy applies. For the applicable tab, the total number of host groups or hosts to which the core policy applies are displayed in parentheses. If no entries exist on one of the tabs, (0) is displayed on it.

When Host is Source

Secure Network Analytics uses this field to determine a host's effective source policy. The following table shows which options are available for alarm categories and security events when the host is the source. They are listed from the most restrictive setting (least likely to trigger an alarm) to the least restrictive setting (most likely to trigger an alarm).

The settings you configure for this field are duplicated in the Desktop Client. For more information about these options, see [How do I know which options to select for source and target](#) in this topic.

These options are available for an alarm category when the host is the source	These options are available for a security event when the host is the source
Off (This option is selected by default.)	Off
The Ignore option is not available in this case.	Ignore (This option is selected by default.)
On	On
On + Alarm	On + Alarm " Many-to-One security events "* (e.g., SYNs Received) cannot alarm at the source, so none of these options are available. In these cases, NA is displayed. * "Many-to-One" security events involve one target host and more than one source host

When Host is Target

Secure Network Analytics uses this field to determine a host's effective target policy. The following table shows which options are available for alarm categories and security events when the host is the target. They are listed from the most restrictive setting (least likely to trigger an alarm) to the least restrictive setting (most likely to trigger an alarm).

The settings you configure for this field are duplicated in the the Desktop Client. For more information about these options, see the "How do I know which options to select" section in this topic.

These options are available for an alarm category when the host is the target	These options are available for a security event when the host is the target
NA is displayed since none of the settings apply to an alarm category when the host is the target.	Off
	Ignore (This option is selected by default.)
	On
	On + Alarm "One-to-Many" security events * (e.g., Max Flows Initiated) cannot alarm at the target, so this option is not available. In these cases, NA is displayed. * "One-to-Many" security events involve one source host but more than one target host.

View or edit the event setting details of a core policy

When you edit any settings or add any events to an existing policy, a green dog ear is displayed in the top right corner of the row. A red dog ear indicates an input error. Hover your cursor over a field name to see the minimum and maximum allowable values.

To edit the settings in this expanded section, the value in at least one of the following columns must be set to *On* or *On+Alarm*:

- When Host is Source
- When Host is Target

To view or edit the details associated with a particular entry in the Core Policy table, click the arrow at the beginning of the row.

The row expands to display the settings you have configured for the event as well as the following:

- If it is a security event, a description of the core policy is listed here.
- If it is an alarm category, all the security events that can potentially contribute to this category are listed here.

For more information about a particular security event, click the **Information** icon next to the Description header and then click the link that appears.

The following table contains descriptions of the various event setting details.

Option Name	Description
Behavioral and Threshold	When this option is selected, the system uses the tolerance setting, minimum threshold, and maximum threshold. Note: When you change your selection from "Behavioral and Threshold" to "Threshold Only," the following fields become disabled: • Tolerance • Never trigger alarm when less than
Threshold only	The alarm triggers only when a minimum or maximum threshold is reached.
Tolerance	Tolerance is the relative number between 0 and 100 that indicates how much to allow actual behavior to exceed expected behavior before alarming. This allows the user to define what is "significantly different." It can also be defined as "the number of standard deviations from the norm." Note: Standard deviation shows how much variation there is from the average (i.e., mean, or expected value). A low standard deviation indicates that the data points tend to be very close to the mean, whereas high standard deviation indicates that the data points are spread out over a large range of values. • A tolerance of 0 causes an alarm to be triggered for any values over the expected value; it is very sensitive and results in a lot of alarms. • A tolerance of 100 allows the behavior to exceed expected behavior to a high degree before it causes an alarm to be triggered, thus greatly reducing the number of times an alarm is triggered. If you never want the alarm to trigger, you must disable the alarm.
Never trigger	Also known as the minimum threshold, this is a static value that indicates the lowest value to allow without triggering an alarm.


Option Name	Description
alarm when less than	The alarm does not trigger when the observed value falls below this setting. In other words, even if a host is greatly over its expected value, if it is not more than the minimum indicated in this dialog, then do not trigger an alarm.
Always trigger alarm when greater than	Also known as the maximum threshold, this is a static value that indicates the highest value to allow without triggering an alarm. The alarm triggers when the observed value exceeds this setting. In other words, if a host's value exceeds the maximum indicated in this dialog, even if it is expected for that host, then trigger an alarm.
Trigger alarm when duration greater than	Duration is the amount of time that the associated value can exceed the calculated threshold before the Manager triggers an alarm.

Find the core policies (including the Effective Policies) associated with a host

To find the effective policies for a host, do the following:

1. Type the IP address into the search box at the top of the page.
2. Click **Search**.

You can enter either an IPv4 or IPv6 address. An orange rectangle with the word *Effective* in it is displayed in the Event column for each effective policy.



- At least one effective policy is displayed per security event, since a host always has at least one effective policy for each security event.

- Two effective policies are displayed if the effective target policy is a different policy from that of the source host effective policy. In these cases, the source policy is always listed first and the target policy second.

The results are sorted in the following order:

1. Effective policy, by name, from most restrictive policy to least restrictive (there is always one effective policy, and sometimes two)
2. Event, in alphabetical order
3. Policy Type (first host, then role, and lastly default)
4. Most restrictive to least restrictive policy, which is determined by the combination of settings between the "When Host is Source" column and the "When Host Is Target" column. For more information, see [Determining the effective policy when two or more role policies exist for a host](#) and [Find the effective policy for a host](#) in this topic.

The number displayed on the tab changes to indicate the number of both effective and noneffective policies configured for the host. Listed immediately after the effective policy or policies are all the noneffective policies for that host. The events are displayed in order from most effective to least effective.

If you also enter filter criteria (see the "Filter columns if you need to narrow your search" section), keep in mind that the system returns results that match both the IP address you entered in the search box and all the criteria you entered in the filter fields at the top of each column.

The page returns to the default view if you do either of the following:

- Clear the IP address in the search box and press **Enter**.
- Make a change and click **Save**.

When looking at the examples below that illustrate the logic used to determine the effective policy or policies, remember that the order of precedence for policy type is as follows:

- [host policy](#)
- [role policy](#)
- [default policy](#)

The hierarchy of options for source and target is as follows:

- Off
- Ignore

- On
- On + Alarm

Example 1:

1. Since there is a host policy, Secure Network Analytics looks at that policy first. The target setting says *On*, so Secure Network Analytics uses the host policy as the effective target policy. Since the source setting says *Ignore* and because there are no more host policies, the system next looks at any role policies to find the effective source policy.
2. No role policy contains the source setting of *Off* which is first in the hierarchy of options, so the system moves on. One of the role policies contains the source setting of *Ignore*, which is next in the hierarchy. Secure Network Analytics does that and ignores this policy, and then moves on to the second role policy. The source setting for this role policy says *On + Alarm*. Therefore, the Antivirus & SMS Servers role policy is the effective source policy.

Secure Network Analytics does not have to use the default policy for an effective source or target policy in this scenario.

i Remember that the default policy is the policy used if there are no host or role policies that can be used as an effective source or target policy.

- Effective SOURCE policy is the Antivirus & SMS Servers role policy.
- Effective TARGET policy is the 60.60 host policy.

Policy Management | 60.60

Custom (1) Relationship (353) **System (5)** [Create New Policy](#)

EVENT	EVENT TYPE	POLICY NAME	POLICY TYPE	HOSTS	WHEN HOST IS SOURCE	WHEN HOST IS TARGET	ACTIONS
Talks to Phantoms	Security	Ex. Default Outside Policy	Ex. Role	Ex. Network Scanners	Ex. On + Alarm	Ex. On + Alarm	
Talks to Phantoms	Security	60.60	Host	60.60	Ignore	On	
Talks to Phantoms	Security	Trusted Internet Hosts	Role	Trusted Internet Hosts, 60.60	On + Alarm	On	
Talks to Phantoms	Security	Role A	Role	60s Host Group, 60.60	Ignore	Off	
Talks to Phantoms	Security	Network Management & Scanners	Role	Network Scanners, 60.60	Ignore	Ignore	
Talks to Phantoms	Security	Default Outside Policy	Default	Outside Hosts	On	On	

50 items per page 1 - 5 of 5 items 1 / 1

Example 2:

1. No host policies exist for this host, so Secure Network Analytics looks next to any role policies that exist, of which two exist for this host.
2. The *On* source setting in the Antivirus & SMS Servers role policy takes precedence over the *On + Alarm* source setting in the Network Management & Scanners role policy, so Secure Network Analytics uses the Antivirus & SMS Servers role policy as the effective source policy.
3. The *Off* target setting in the Network Management & Scanners role policy takes precedence over the *On* target setting in the Antivirus & SMS Servers role policy, so Secure Network Analytics uses the Network Management & Scanners role policy as the effective target policy.

Secure Network Analytics does not have to use the default policy for an effective source or target policy in this scenario.

- i** Remember that the default policy is the policy used if there are no host or role policies that can be used as an effective source or target policy.

- Effective SOURCE policy is the Antivirus & SMS Servers role policy.
- Effective TARGET policy is the Network Management & Scanners role policy.

Policy Management | 60.60

Custom (1) Relationship (353) System (3) [Create New Policy](#)

EVENT	EVENT TYPE	POLICY NAME	POLICY TYPE	HOSTS	WHEN HOST IS SOURCE	WHEN HOST IS TARGET	ACTIONS
Addr_Scan/vuln	Security	Antivirus & SMS Servers	Role	SMS Servers, Antivirus Servers, 60.60	On	On	
Addr_Scan/vuln	Security	Network Management & Scanners	Role	Network Scanners, 60.60	On + Alarm	Off	
Addr_Scan/vuln	Security	Default Outside Policy	Default	Outside Hosts	On	On	

50 Items per page 1 - 3 of 3 items

Find all core policies associated with a host group

You must assign a host group to a role policy in order for you to be able to search for that host group on any of the Policy Management pages.

To find all the core policies associated with a host group, do the following:

1. Click the  (**List Menu**) icon in the search box.

The Host Groups slide-out panel opens.

2. Select the desired host group.
3. Click **Search**.

All the core policies that are associated with that host group are displayed, sorted by event. The number displayed on the tab changes to indicate the number of core policies configured for the host group.

If you also enter filter criteria, keep in mind that the system returns results that match both the host group you selected in the search box and all the criteria you entered in the filter fields at the top of each column. For more information, see the next section, [Filter columns if you need to narrow your search](#).

The page returns to the default view if you do either of the following:

- Clear the host group in the search box and press **Enter**.
- Make a change and click **Save**.

Filter columns if you need to narrow your search

At the top of each of the filterable columns, there is a filter field in which you can enter the criteria by which you want to filter that column. Some display a drop-down list as you begin to type and some do not.

For all fields except for the Hosts field, as you start to type your entry in any of these fields, a drop-down list appears displaying a list of entries that match what you have begun to type. You can then choose the appropriate entry from the list so that you don't have to type the complete entry. The remaining columns dynamically filter to display only the entries that match your entry.

For the Hosts field, as you start to type your entry, the column dynamically filters to display a list of entries that match what you have begun to type. This match can be for any word in the entry (beginning, middle, or last word). Type the complete entry into the field to return results that are an exact match.

Filtering with multiple entries

You can type multiple entries within one column as well as enter values across multiple columns. When you type multiple values within one column, these values are ORed together; therefore, all returned results need only match one of the entries within the column. Remember to leave a space between each entry. When you filter more than one column, filters are ANDed together; therefore, all returned results must match the search

criteria in all of the columns. Enclose an entry within quotation marks to return results that exactly match your entry.

For the Host Groups column, the table first lists all events that contain your entry as the first host group, and then displays all events that contain your entry as the second host group. Also, the Host Groups column filters all host groups, not just those displayed in the table.

Editing a policy with filtered criteria

If you edit a core policy after having filtered criteria, the entries in the table do not change until you save your edits. Click the *X* in the field to clear the filter criteria and return to the previous view for that column. To quickly clear multiple filter criteria, refresh the page.

Clearing entries

Click the *X* in the field to clear the filter criteria and return to the previous view for that column. To quickly clear multiple filter criteria, refresh the page.

What you can do next

Create a host or role policy

To create a role policy or host policy, click **Create New Policy > [Host or Role] Policy**.

The associated Policy page opens, showing the policy's details populated in the applicable fields and all of the events related to the selected policy listed alphabetically below that.



- For information about how to create a role policy, see [Configuring Role Policies](#).
- For information about how to create a host policy, see [Configuring Host Policies](#).

Edit a host or role policy

To edit a [host policy](#), [role policy](#), or [default policy](#), do the following:

1. Click the policy name in the Policy Name column.

The associated Policy page opens, showing the policy's details populated in the applicable fields and all of the events related to the selected policy listed alphabetically below that.

2. Expand the applicable row to view more details about an event and decide if you need to make edits to these settings as well.
3. Make any necessary edits.
4. Click **Save**. (You can make edits on different pages before you click **Save**.)



If you click **Edit** from a row that contains an effective policy that is displayed as a result of using the Search box, the Policy page that opens lists all of the events related to that policy, not related to the effective policy.

Related topics

[Secure Network Analytics Policies](#)

[Configuring Host Policies](#)

[Configuring Role Policies](#)

[Editing Default Policies](#)

[Security Event List](#)

[About Alarm Categories](#)

[About Alarm Severities](#)

Configuring Host Policies

Use this page to create and manage your host policies.

If you find that you are about to assign a host policy to a range of IP addresses or to multiple ranges of IP addresses, then instead create a host group for these IP addresses, and then assign to this host group a host policy.

It is always preferable to create a role policy for a host group (where hosts share a common function) rather than a host policy for an individual host. When you add a host to a host group, it automatically inherits the host policy associated with that host group. When you need to modify a host's host policy, it is easier to only have to re-associate the host to a more relevant host group rather than both re-associating the host group and then modifying any number of individual policies that may have been configured with the particular host included.

For information about viewing and maintaining core policies, see [Managing Core Policies](#)

How to find this page

Do any of the following to access this page

- From the main menu:
 1. Click **Configure > Detection > Policy Management**.
The Policy Management: Custom Events tab opens.
 2. Click **Create New Policy** from the drop-down menu and select **Single Host Policy**.
The Host Policy page opens. The host policy's events are listed alphabetically in the table at the bottom of the page, first by Category events and then by Security events.
- On the Policy Management: Core tab:
 1. In the row of the applicable event, click the policy name.
 2. From the menu, click **Edit**.
The Host Policy page opens. The host policy's events are listed alphabetically in the table at the bottom of the page, first by Category events and then by Security events.

Concepts you need to understand to use this page

Refer to the following guidelines when configuring a host policy

- You can use the following units of measure: K, M, G, T, P, and E.
- You can enter decimal numbers only if you use a unit with it (e.g., 1.2 M, not 1.2).
- In all fields, 1k = 1000.
- When you enter a number over 999, Secure Network Analytics changes that number to the nearest unit after you have clicked away from that field.

*Example: 1000 changes to 1K; 1024 changes to 1.02K; 2,000,000 = 2 M;
2,000,000,000 = 2 G*

IPv4 IP address

If you enter this type of IP address...	You must use this formatting...
Single IP address	10.52.1.55 Includes only the host at 10.52.1.55.
Trailing dot subnet	10.52. Includes any IP address from 10.52.0.0 to 10.52.255.255. Note: You must include the trailing period (.) at the end.
Classless Inter-Domain Routing (CIDR) notation (i.e., using a “/” to indicate the masked bits)	10.52.1.0/24 Includes any IP address from 10.52.1.0 to 10.52.1.255. Note: Use this notation if you prefer to enter the IP address using a “mask.”
Range of networks	10-11. Includes any IP address from 10.0.0.0 to 11.255.255.255

If you enter this type of IP address...	You must use this formatting...
	10.52-53. Includes any IP address from 10.52.0.0 to 10.52.255.255, and from 10.53.0.0 to 10.53.255.255 10.52-55.3. Includes any IP address from 10.52.3.0 to 10.52.3.255; from 10.53.3.0 to 10.53.3.255; from 10.54.3.0 to 10.54.3.255; and from 10.55.3.0 to 10.55.3.255 Note: Be sure to include the trailing period (.) at the end.
Range of hosts	10.52.1.0-10 Includes any IP address from 10.52.1.0 to 10.52.1.10 10.52.0.0-10.52.255.255 Includes any IP address from 10.52.0.0 to 10.52.255.255
Comma-separated list	10.52.1.10,10.52.1.50,10.100.1.20 Includes only these three hosts.

IPv6 IP address

If you enter this type of IP address...	You must use this formatting...
Single IP address	2001:0DB8:0000:0056:0000:ABCD:EF12:3456 2001:DB8:0:56:0:ABCD:EF12:3456 2001:DB8::56:0:ABCD:EF12:3456 2001:DB80:0:56::ABCD:EF12:3456 2001:DB80:0:56::ABCD:239.18.52.86
Global Routing	2001:DB8:0:56::/64

If you enter this type of IP address...	You must use this formatting...
Prefix subnet	
Range of hosts	2001:DB8:0:56:ABCD:EF12:3456:1-10 2001:DB8:0:56:ABCD:EF12:3456:1- 2001:DB8:0:56:ABCD:EF12:3456:10

About IPv4 and IPv6 addresses

You can use either IPv4 or IPv6 addresses for hosts or host groups. The notation forms that must be used for each type are listed in the applicable table. You must separate the IP addresses by a comma, comma space, space, or return.

The Manager displays IPv6 addresses and IPv4 addresses simultaneously. Of the advantages provided by IPv6, the fact that it supports a larger address space is probably the most well known. IPv6 supports 128-bit addresses, whereas IPv4 supports only 32-bit addresses.

The Manager continues to support IPv4. The Flow Collectors can monitor IPv6 traffic, but they are not IPv6 addressable at this time.

Examples

- An IPv6 address consists of eight blocks of hexadecimal digits separated by colons, as shown in this example (where the x's represent hexadecimal digits):

Example: xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx

- You can replace blocks of 4-digit zeros with double colons (e.g., ::), but only once per IPv6 address. (Leading zeros can also be omitted.) Consider the following address (where the x's represent hexadecimal digits):

Example: xxxx:xxxx:0000:0000:0000:0000:0000:xxxx

This address can be rewritten like this:

Example: xxxx:xxxx::xxxx.

- Like an IPv4 address, an IPv6 address may include a prefix length that specifies how many left-most bits of the address make up the prefix. In IPv4 addresses, the prefix is used to determine the subnet mask (a.k.a., CIDR notation). In IPv6 addresses, it identifies the subnet that an interface belongs to, and routers use it for forwarding.

Example: xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx/40

- Take care when compressing an IPv6 address that contains a prefix, because if not done correctly, it can result in an incorrect interpretation when the address is uncompressed. You can write an IPv4 address as an IPv6 address, like this:

Example: xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:192.168.0.2

- If the first 6 blocks are zeros, you can compress the address like this:

Example: ::192.168.0.2

CIDR notation

Use the following table to understand classless inter-domain routing (CIDR) notation:

Class	CIDR Length	Mask	Number of Addresses
A	/1	128.0.0.0	2,147,483,648
	/2	192.0.0.0	1,073,741,824
	/3	224.0.0.0	536,870,912
	/4	240.0.0.0	268,435,456

Cl as s	CI DR Le ngt h	Mask	Numbe r of Addres ses
	/5	248.0.0. 0	134,21 7,728
	/6	252.0.0. 0	67,108 ,864
	/7	254.0.0. 0	33,554 ,432
	/8	255.0.0. 0	16,777 ,216
	/9	255.128 .0.0	8,388, 608
	/10	255.192 .0.0	4,194, 304
	/11	255.224 .0.0	2,097, 152
	/12	255.240 .0.0	1,048, 576
	/13	255.248 .0.0	524,28 8
	/14	255.252 .0.0	262,14 4
	/15	255.254 .0.0	131,07 2
B	/16	255.255	65,536

Cl as s	CI DR Le ngt h	Mask	Numbe r of Addres ses
		.0.0	
	/17	255.255 .128.0	32,768
	/18	255.255 .192.0	16,384
	/19	255.255 .224.0	8,192
	/20	255.255 .240.0	4,096
	/21	255.255 .248.0	2,048
	/22	255.255 .252.0	1,024
	/23	255.255 .254.0	512
C	/24	255.255 .255.0	256
	/25	255.255 .255.128	128
	/26	255.255 .255.192	64
	/27	255.255 .255.224	32

Class	CIDR Length	Mask	Number of Addresses
	/28	255.255. .255.240	16
	/29	255.255. .255.248	8
	/30	255.255. .255.252	4
	/31	255.255. .255.254	2

What you can do on this page

Create or edit a host policy

Refer to the following guidelines when using this page:

- You must configure the following to create a host policy:
 - A host
 - A minimum of one event
- You cannot create a host policy for the same IP address within the same domain.
- All mitigation configuration must be completed in the Desktop Client. Any edits you make to a host policy in the Web App does not alter any mitigation settings associated with that host policy.



To see the default values for the core policies, access this table: [Core Policy Default Values](#).

To create or edit a host policy, do the following:

1. Complete the fields at the top of the page.
2. Add the desired events. See the next section, [Add an event to a host policy](#).

As you add events, they are displayed in the associated table.

After you have saved the information on this page, the [Policy Manager: Core tab](#) opens and each event you have configured for this host policy are displayed in the Core Policies table.

Refer to the following information about the fields on this page:

Name

This must be a unique name.

IP Address

You can use either an IPv4 or IPv6 address. (Refer to the "IPv4 and IPv6 address formats" section in this topic for the correct notation forms to use.)

Description

Indicates the purpose of this host policy.

Add an event to a host policy

To add one or more events to a host policy, do the following:

1. Click **Select Events** in the applicable table.

The Events panel opens on the left side of the page which contains a list of all the alarm categories on one tab and a list of all the default security events on another tab. Any events already defined for the policy are pre-selected.

There is no maximum number of items you can select. If you type the beginning letters or the complete name of the alarm category or security event name into the text box, all other items are removed from the list and the one you are filtering on remains for you to select.

2. After you have added the desired events to the host policy and clicked **Save**, the Host Policy page opens again.

A row is displayed for every event you added to the host policy. Also, the number (in the parentheses) on the Policy Manager: System tab now reflects the new number of existing events in your system.

View or edit the event setting details of a host policy

To view or edit the event setting details for a host policy, click the arrow at the beginning of the row.

The row expands to display the settings you have configured for the event as well as the following:

- If it is a security event, a description of the core policy is listed here.
- If it is an alarm category, all the security events that can potentially contribute to this category are listed here.

For more information about a particular security event, click the  (**Information**) icon next to the Description header and then click the link that appears.

To edit the event setting details in this expanded section, the value in at least one of the following columns must be set to *On* or *On+Alarm*:

- When Host is Source
- When Host is Target

When you edit any settings or add any events to an existing policy, a green dog ear is displayed in the top right corner of the row. A red dog ear indicates an input error. Hover your cursor over a field name to see the minimum and maximum allowable values.

Refer to the following table for information about the event setting detail values.

Option Name	Description
Behavioral and Threshold	When this option is selected, the system uses the tolerance setting, minimum threshold, and maximum threshold. Note: When you change your selection from "Behavioral and Threshold" to "Threshold Only," the following fields become disabled: • Tolerance • Never trigger alarm when less than
Threshold only	The alarm triggers only when a minimum or maximum threshold is reached.
Tolerance	Tolerance is the relative number between 0 and 100 that indicates how much to allow actual behavior to exceed expected behavior before alarming. This allows the user to define what is "significantly different." It can also be defined as "the number of standard deviations from the norm." Note: Standard deviation shows how much variation there is

Option Name	Description
	from the average (i.e., mean, or expected value). A low standard deviation indicates that the data points tend to be very close to the mean, whereas high standard deviation indicates that the data points are spread out over a large range of values. • A tolerance of 0 causes an alarm to be triggered for any values over the expected value; it is very sensitive and results in a lot of alarms. • A tolerance of 100 allows the behavior to exceed expected behavior to a high degree before it causes an alarm to be triggered, thus greatly reducing the number of times an alarm is triggered. If you never want the alarm to trigger, you must disable the alarm.
Never trigger alarm when less than	Also known as the minimum threshold, this is a static value that indicates the lowest value to allow without triggering an alarm. The alarm does not trigger when the observed value falls below this setting. In other words, even if a host is greatly over its expected value, if it is not more than the minimum indicated in this dialog, then do not trigger an alarm.
Always trigger alarm when greater than	Also known as the maximum threshold, this is a static value that indicates the highest value to allow without triggering an alarm. The alarm triggers when the observed value exceeds this setting. In other words, if a host's value exceeds the maximum indicated in this dialog, even if it is expected for that host, then trigger an alarm.
Trigger alarm when duration greater than	Duration is the amount of time that the associated value can exceed the calculated threshold before the Manager triggers an alarm.

How do I know which options to select for source and target

Following are the options you can select for both the source host and the target host. The following list shows the most restrictive setting (least likely to trigger an alarm) at the top and ends with the least restrictive setting (most likely to trigger an alarm).

- Off
- Ignore
- On
- On + Alarm

Refer to the following table for an explanation of the options.

If this option...	Is selected for "When Host Is Source," then this security event...	Is selected for "When Host is Target," then this security event...
Off	Does not collect index points nor alarm.	Does not contribute index points to alarm category events nor alarm.
Ignore	Is bypassed for the associated policy. Secure Network Analytics ignores the settings configured for this event in this policy (to which the associated host belongs). It instead uses the settings configured for this event in the next policy (in order of precedence) to which the associated host belongs. For more information about order of precedence used for host policies, refer to the "What are the three core	Is bypassed for the associated policy. Secure Network Analytics ignores the settings configured for this event in this policy (to which the associated host belongs). It instead uses the settings configured for this event in the next policy (in order of precedence) to which the associated host belongs. For more information about order of precedence used for host policies, refer to the "What are the three core

	policy types" section in this topic. For example, if the ignored policy is a host policy, Secure Network Analytics uses the settings configured for this event in the role policy to which the associated host belongs. If no such role policy exists, then Secure Network Analytics uses the settings configured for this event in the default policy.  The Ignore option does not apply to alarm category events nor is not available for default policies.	policy types" section in this topic. For example, if the ignored policy is a host policy, Secure Network Analytics uses the settings configured for this event in the role policy to which the associated host belongs. If no such role policy exists, then Secure Network Analytics uses the settings configured for this event in the default policy.
On	Collects index points but does not alarm.	Contributes index points to alarm category events but does not generate alarms.
On + Alarm	Collects index points and generates alarms.	Contributes index points to alarm category events and generates alarms.

Default Policy, Inside Hosts

When you set the Target to OFF for a default policy, it disables the alarm for all inside hosts, even if the policy is set to alarm on the Source, because the Target is also a member of the Inside Hosts group.

In this case, you would create a role policy which would contain the applicable IP addresses, and then you would set the Source to ON and the Target to OFF.

Example:

Say you want to turn off the Brute Force Login security event when your monitoring server is the source (since it is frequently logging in to devices), but you still want to know if anyone tries to brute force their way into it. Provided you had your monitoring servers assigned to a host group, you could [create a role policy](#) for that host group and add the Brute Force Login security event to that policy. You would then configure the settings as shown in the following table.

The settings for Brute Force Login security event for this role policy would be...	
When Host Is Source	When Host Is Target
Off (This would prevent Bruce Force Login from firing when your monitoring server was the source.)	On + Alarm (This would cause Bruce Force Login to trigger an alarm when your monitoring server was the target.)

Filter columns if you need to narrow your search

At the top of each of the filterable columns, there is a filter field in which you can enter the criteria by which you want to filter that column. As you start to type your entry in any of these fields, a drop-down list appears displaying a list of entries that match what you have begun to type. You can then chose the appropriate entry from the list so that you don't have to type the complete entry. The remaining columns dynamically filter to display only the entries that match your entry.

Filtering with multiple entries

You can type multiple entries within one column as well as enter values across multiple columns. When you type multiple values within one column, these values are ORed together; therefore, all returned results need only match one of the entries within the column. Remember to leave a space between each entry. When you filter more than one column, filters are ANDed together; therefore, all returned results must match the search criteria in all of the columns. Enclose an entry within quotation marks to return results that exactly match your entry.

Editing a policy with filtered criteria

If you edit a host policy after having filtered criteria, the entries in the table do not change until you save your edits. Click the X in the field to clear the filter criteria and return to the previous view for that column. To quickly clear multiple filter criteria, refresh the page.

Clearing entries

Click the X in the field to clear the filter criteria and return to the previous view for that column. To quickly clear multiple filter criteria, refresh the page.

Delete a host policy

From the Actions drop-down menu, click **Delete**.

The host policy is deleted from the Host Policy table.

Delete an event from a host policy

To delete one or more events from a host policy, click **Delete** in the Actions column of the row for the applicable event.

*After you have deleted the desired events from the host policy and clicked **Save**, the [Policy Manager: Core Events tab](#) opens again, and the events you deleted are no longer displayed. Also, the number (in the parentheses) on the Core Events tab now reflects the new number of core events in your system.*

Duplicate a host policy

Duplicate a host policy if you want to add a host policy that is the same or similar to an existing host policy so you can use it as a template and therefore create the new one more easily and quickly.

To duplicate a host policy, do the following:

1. From the Actions drop-down menu, click **Duplicate**.

The Host Policy page opens, showing the policy's details populated in the applicable fields and all of the events related to the selected host policy listed below that.

2. Enter the applicable IP address into the empty IP Address field.
3. Make any other edits, if necessary.
4. Click **Save**.

What you can do next

If your policies are not behaving correctly, synchronize your Flow Collector with the Manager

At times you may need to synchronize the Flow Collectors on your network with your Manager. Some reasons may be as follows:

- You observe that your Flow Collectors are not working properly.
- You have added a new configuration (e.g., role policy, host group), but while you were doing so your Manager stopped communicating with any or all of your Flow Collectors. In this case, the policy will have been saved to your Manager, but it may not have propagated to any or all of your Flow Collectors.

To synchronize one or more Flow Collectors to your Manager, log in to your Manager client interface and complete the following steps:

1. From the navigation menu, click **Configure > Flow Collectors**.
2. Click the **Advanced** tab.
3. In the top left corner, from the Flow Collector drop-down list, choose the applicable Flow Collector.
4. In the Synchronize section, click **Synchronize**.
5. Repeat Step 3 through Step 4 for the remaining Flow Collectors that need to be synchronized with the Manager.

Related topics

[Managing Core Policies](#)

[Configuring Role Policies](#)

[Editing Default Policies](#)

[Security Event List](#)

[About Alarm Categories](#)

[About Alarm Severities](#)

Configuring Role Policies

Use this page to create and manage your role policies.

You create role policy when you want a group of hosts that share a common function or similar attributes to be assigned the same alarm thresholds. It is always preferable to create a role policy for a host group rather than for an individual host. When you add a host to a host group, it automatically inherits the policy associated with that host group. When you need to modify a host's policy, it is easier to only have to re-associate the host to a more relevant host group rather than both re-associating the host group and then modifying any number of individual policies that may have been configured with the particular host included.

To prevent confusion, it is best if you do not use multiple role policies with the same alarm unless there is a need to trigger alarms based on different values (e.g., for different teams).

For information about viewing and maintaining core policies, see [Managing Core Policies](#).

How to find this page

Do any of the following to access this page

- From the main menu:
 1. Click **Configure > Detection > Policy Management**.
The Policy Management: Custom Events tab opens.
 2. Click **Create New Policy** from the drop-down menu and select **Role Policy**.
The Role Policy page opens. The role policy's events are listed alphabetically in the table at the bottom of the page, first by Category events and then by Security events.
- On the Policy Management: Core tab:
 1. In the row of the applicable event, click the policy name.
 2. From the menu, click **Edit**.
The Role Policy page opens. The role policy's events are listed alphabetically in the table at the bottom of the page, first by Category events and then by Security events.

Concepts you need to understand to use this page

Refer to the following guidelines when configuring a role policy

- You can use the following units of measure: K, M, G, T, P, and E.
- You can enter decimal numbers only if you use a unit with it (e.g., 1.2 M, not 1.2).
- In all fields, 1k = 1000.
- When you enter a number over 999, Secure Network Analytics changes that number to the nearest unit after you have clicked away from that field.

*Example: 1000 changes to 1K; 1024 changes to 1.02K; 2,000,000 = 2 M;
2,000,000,000 = 2 G*

IPv4 IP address

If you enter this type of IP address...	You must use this formatting...
Single IP address	10.52.1.55 Includes only the host at 10.52.1.55.
Trailing dot subnet	10.52. Includes any IP address from 10.52.0.0 to 10.52.255.255. Note: You must include the trailing period (.) at the end.
Classless Inter-Domain Routing (CIDR) notation (i.e., using a “/” to indicate the masked bits)	10.52.1.0/24 Includes any IP address from 10.52.1.0 to 10.52.1.255. Note: Use this notation if you prefer to enter the IP address using a “mask.”
Range of networks	10-11. Includes any IP address from 10.0.0.0 to 11.255.255.255

If you enter this type of IP address...	You must use this formatting...
	10.52-53. Includes any IP address from 10.52.0.0 to 10.52.255.255, and from 10.53.0.0 to 10.53.255.255 10.52-55.3. Includes any IP address from 10.52.3.0 to 10.52.3.255; from 10.53.3.0 to 10.53.3.255; from 10.54.3.0 to 10.54.3.255; and from 10.55.3.0 to 10.55.3.255 Note: Be sure to include the trailing period (.) at the end.
Range of hosts	10.52.1.0-10 Includes any IP address from 10.52.1.0 to 10.52.1.10 10.52.0.0-10.52.255.255 Includes any IP address from 10.52.0.0 to 10.52.255.255
Comma-separated list	10.52.1.10,10.52.1.50,10.100.1.20 Includes only these three hosts.

IPv6 IP address

If you enter this type of IP address...	You must use this formatting...
Single IP address	2001:0DB8:0000:0056:0000:ABCD:EF12:3456 2001:DB8:0:56:0:ABCD:EF12:3456 2001:DB8::56:0:ABCD:EF12:3456 2001:DB80:0:56::ABCD:EF12:3456 2001:DB80:0:56::ABCD:239.18.52.86
Global Routing	2001:DB8:0:56::/64

If you enter this type of IP address...	You must use this formatting...
Prefix subnet	
Range of hosts	2001:DB8:0:56:ABCD:EF12:3456:1-10 2001:DB8:0:56:ABCD:EF12:3456:1- 2001:DB8:0:56:ABCD:EF12:3456:10

About IPv4 and IPv6 addresses

You can use either IPv4 or IPv6 addresses for hosts or host groups. The notation forms that must be used for each type are listed in the applicable table. You must separate the IP addresses by a comma, comma space, space, or return.

The Manager displays IPv6 addresses and IPv4 addresses simultaneously. Of the advantages provided by IPv6, the fact that it supports a larger address space is probably the most well known. IPv6 supports 128-bit addresses, whereas IPv4 supports only 32-bit addresses.

The Manager continues to support IPv4. The Flow Collectors can monitor IPv6 traffic, but they are not IPv6 addressable at this time.

Examples

- An IPv6 address consists of eight blocks of hexadecimal digits separated by colons, as shown in this example (where the x's represent hexadecimal digits):

Example: xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx

- You can replace blocks of 4-digit zeros with double colons (e.g., ::), but only once per IPv6 address. (Leading zeros can also be omitted.) Consider the following address (where the x's represent hexadecimal digits):

Example: xxxx:xxxx:0000:0000:0000:0000:0000:xxxx

This address can be rewritten like this:

Example: xxxx:xxxx::xxxx.

- Like an IPv4 address, an IPv6 address may include a prefix length that specifies how many left-most bits of the address make up the prefix. In IPv4 addresses, the prefix is used to determine the subnet mask (a.k.a., CIDR notation). In IPv6 addresses, it identifies the subnet that an interface belongs to, and routers use it for forwarding.

Example: xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx/40

- Take care when compressing an IPv6 address that contains a prefix, because if not done correctly, it can result in an incorrect interpretation when the address is uncompressed. You can write an IPv4 address as an IPv6 address, like this:

Example: xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:192.168.0.2

- If the first 6 blocks are zeros, you can compress the address like this:

Example: ::192.168.0.2

CIDR notation

Use the following table to understand classless inter-domain routing (CIDR) notation:

Class	CIDR Length	Mask	Number of Addresses
A	/1	128.0.0.0	2,147,483,648
	/2	192.0.0.0	1,073,741,824
	/3	224.0.0.0	536,870,912
	/4	240.0.0.0	268,435,456

Cl as s	CI DR Le ngt h	Mask	Numbe r of Addres ses
	/5	248.0.0. 0	134,21 7,728
	/6	252.0.0. 0	67,108 ,864
	/7	254.0.0. 0	33,554 ,432
	/8	255.0.0. 0	16,777 ,216
	/9	255.128 .0.0	8,388, 608
	/10	255.192 .0.0	4,194, 304
	/11	255.224 .0.0	2,097, 152
	/12	255.240 .0.0	1,048, 576
	/13	255.248 .0.0	524,28 8
	/14	255.252 .0.0	262,14 4
	/15	255.254 .0.0	131,07 2
B	/16	255.255	65,536

Cl as s	CI DR Le ngt h	Mask	Numbe r of Addres ses
		.0.0	
	/17	255.255 .128.0	32,768
	/18	255.255 .192.0	16,384
	/19	255.255 .224.0	8,192
	/20	255.255 .240.0	4,096
	/21	255.255 .248.0	2,048
	/22	255.255 .252.0	1,024
	/23	255.255 .254.0	512
C	/24	255.255 .255.0	256
	/25	255.255 .255.128	128
	/26	255.255 .255.192	64
	/27	255.255 .255.224	32

Class	CIDR Length	Mask	Number of Addresses
	/28	255.255. .255.240	16
	/29	255.255. .255.248	8
	/30	255.255. .255.252	4
	/31	255.255. .255.254	2

What you can do on this page

Create or edit a role policy

Refer to the following guidelines when using this page:

- You must configure the following to create a role policy:
 - Name
 - One host group, IP address, or IP range
 - A minimum of one event
- All mitigation configuration must be completed in the Desktop Client. Any edits you make to a role policy in the Manager does not alter any mitigation settings associated with that role policy.



To see the default values for the core policies, access this table: [Core Policy Default Values](#).

To create or edit a role policy, do the following:

- Complete the fields at the top of the page.
- Add the desired events. See the next section, [Add an event to a role policy](#)

As you add events, they are displayed in the applicable table.

After you have saved the information on this page, the [Policy Manager: Core Events tab](#) opens. The Core Policies table is filtered by your new role policy, and each event you have configured for the new policy is displayed in this table.

Refer to the following information about the fields on this page:

Name (required)

This must be a unique name.

Description

Indicates the purpose of this role policy.

Host Groups

The host groups to which this policy applies. Click the  **(Plus)** icon to open the Host Group panel and add one or more host groups. Any host groups already defined for the policy are pre-selected.

If you type any letters that are anywhere within a host group name into the Search field and press **Enter**, the Host Group tree expands to display (in italicized red font) any host group name that contains that entry.

Example: If you entered "South" into the Search field, some of the host group names that would then be displayed in italicized red font would be Southern Africa, South Africa, South-Eastern Asia, French Southern Territories.

To return an exact match, enter the entire host group name into the Search field and press **Enter**. If you want to return to the previously selected host group, delete the entry in the Search field and press **Enter**.

IP Address or Range

You can use either IPv4 or IPv6 addresses. (Refer to the "IPv4 and IPv6 address formats" section in this topic for the correct notation forms to use.)

- Paste the entries from another source or type them.
- Include a range of hosts.

Add an event to a role policy

To add one or more events to a role policy, do the following:

1. Click **Select Events** in the associated table.

The Events panel opens on the left side of the page which contains a list of all the alarm categories on one tab and a list of all the default security events on another tab. Any events already defined for the policy are pre-selected.

There is no maximum number of items you can select. If you type the beginning letters or the complete name of the alarm category or security event name into the text box, all other items are removed from the list and the one you are filtering on remain for you to select.

2. After you have added the desired events to the role policy and clicked **Save**, the Role Policy page opens again.

A row is displayed for every event you added to the role policy. Also, the number (in the parentheses) on the Policy Manager: System tab now reflects the new number of existing events in your system.

View or edit the event setting details of a role policy

To view or edit the event setting details for a role policy, click the arrow at the beginning of the row.

The row expands to display the settings you have configured for the event as well as the following:

- If it is a security event, a description of the core policy is listed here.
- If it is an alarm category, all the security events that can potentially contribute to this category are listed here.

For more information about a particular security event, click the  (**Information**) icon next to the Description header and then click the link that appears.

To edit the event setting details in this expanded section, the value in at least one of the following columns must be set to *On* or *On+Alarm*:

- When Host is Source
- When Host is Target

When you edit any settings or add any events to an existing policy, a green dog ear is displayed in the top right corner of the row. A red dog ear indicates an input error. Hover your cursor over a field name to see the minimum and maximum allowable values.

Refer to the following table for information about the event setting detail values.

Option Name	Description
Behavioral and Threshold	When this option is selected, the system uses the tolerance setting, minimum threshold, and maximum threshold. Note: When you change your selection from "Behavioral and Threshold" to "Threshold Only," the following fields become disabled: • Tolerance • Never trigger alarm when less than
Threshold only	The alarm triggers only when a minimum or maximum threshold is reached.
Tolerance	Tolerance is the relative number between 0 and 100 that indicates how much to allow actual behavior to exceed expected behavior before alarming. This allows the user to define what is "significantly different." It can also be defined as "the number of standard deviations from the norm." Note: Standard deviation shows how much variation there is from the average (i.e., mean, or expected value). A low standard deviation indicates that the data points tend to be very close to the mean, whereas high standard deviation indicates that the data points are spread out over a large range of values. • A tolerance of 0 causes an alarm to be triggered for any values over the expected value; it is very sensitive and results in a lot of alarms. • A tolerance of 100 allows the behavior to exceed expected behavior to a high degree before it causes an alarm to be triggered, thus greatly reducing the number of times an alarm is triggered. If you never want the alarm to trigger, you must disable the alarm.
Never trigger alarm when	Also known as the minimum threshold, this is a static value that indicates the lowest value to allow without triggering an alarm. The alarm does not trigger when the observed value falls below

Option Name	Description
less than	this setting. In other words, even if a host is greatly over its expected value, if it is not more than the minimum indicated in this dialog, then do not trigger an alarm.
Always trigger alarm when greater than	Also known as the maximum threshold, this is a static value that indicates the highest value to allow without triggering an alarm. The alarm triggers when the observed value exceeds this setting. In other words, if a host's value exceeds the maximum indicated in this dialog, even if it is expected for that host, then trigger an alarm.
Trigger alarm when duration greater than	Duration is the amount of time that the associated value can exceed the calculated threshold before the Manager triggers an alarm.

How do I know which options to select for source and target

Following are the options you can select for both the source host and the target host. The following list shows the most restrictive setting (least likely to trigger an alarm) at the top and ends with the least restrictive setting (most likely to trigger an alarm).

- Off
- Ignore
- On
- On + Alarm

Refer to the following table for an explanation of the options.

If this option...	Is selected for "When Host Is Source," then this security event...	Is selected for "When Host is Target," then this security event...
Off	Does not collect index points nor alarm.	Does not contribute index points to alarm category events

		nor alarm.
Ignore	Is bypassed for the associated policy. Secure Network Analytics ignores the settings configured for this event in this policy (to which the associated host belongs). It instead uses the settings configured for this event in the next policy (in order of precedence) to which the associated host belongs. For more information about order of precedence used for host policies, refer to the "What are the three core policy types" section in this topic. For example, if the ignored policy is a host policy, Secure Network Analytics uses the settings configured for this event in the role policy to which the associated host belongs. If no such role policy exists, then Secure Network Analytics uses the settings configured for this event in the default policy.  The Ignore option does not apply to alarm category events nor is not available for default policies.	Is bypassed for the associated policy. Secure Network Analytics ignores the settings configured for this event in this policy (to which the associated host belongs). It instead uses the settings configured for this event in the next policy (in order of precedence) to which the associated host belongs. For more information about order of precedence used for host policies, refer to the "What are the three core policy types" section in this topic. For example, if the ignored policy is a host policy, Secure Network Analytics uses the settings configured for this event in the role policy to which the associated host belongs. If no such role policy exists, then Secure Network Analytics uses the settings configured for this event in the default policy.

On	Collects index points but does not alarm.	Contributes index points to alarm category events but does not generate alarms.
On + Alarm	Collects index points and generates alarms.	Contributes index points to alarm category events and generates alarms.

Default Policy, Inside Hosts

When you set the Target to OFF for a default policy, it disables the alarm for all inside hosts, even if the policy is set to alarm on the Source, because the Target is also a member of the Inside Hosts group.

In this case, you would create a role policy which would contain the applicable IP addresses, and then you would set the Source to ON and the Target to OFF.

Example:

Say you want to turn off the Brute Force Login security event when your monitoring server is the source (since it is frequently logging in to devices), but you still want to know if anyone tries to brute force their way into it. Provided you had your monitoring servers assigned to a host group, you could [create a role policy](#) for that host group and add the Brute Force Login security event to that policy. You would then configure the settings as shown in the following table.

The settings for Brute Force Login security event for this role policy would be...	
When Host Is Source	When Host Is Target
Off (This would prevent Bruce Force Login from firing when your monitoring server was the source.)	On + Alarm (This would cause Bruce Force Login to trigger an alarm when your monitoring server was the target.)

Filter columns if you need to narrow your search

At the top of each of the columns, there is a filter field in which you can enter the criteria by which you want to filter that column. As you start to type your entry in any of these

fields, a drop-down list appears displaying a list of entries that match what you have begun to type. You can then choose the appropriate entry from the list so that you don't have to type the complete entry. The remaining columns dynamically filter to display only the entries that match your entry.

Filtering with multiple entries

You can type multiple entries within one column as well as enter values across multiple columns. When you type multiple values within one column, these values are ORed together; therefore, all returned results need only match one of the entries within the column. Remember to leave a space between each entry. When you filter more than one column, filters are ANDed together; therefore, all returned results must match the search criteria in all of the columns. Enclose an entry within quotation marks to return results that exactly match your entry.

Editing a policy with filtered criteria

If you edit a role policy after having filtered criteria, the entries in the table do not change until you save your edits. Click the *X* in the field to clear the filter criteria and return to the previous view for that column. To quickly clear multiple filter criteria, refresh the page.

Clearing entries

Click the *X* in the field to clear the filter criteria and return to the previous view for that column. To quickly clear multiple filter criteria, refresh the page.

Delete a role policy

From the Actions drop-down menu, click **Delete**.

The role policy is deleted from the Role Policy table.

Delete an event from a role policy

To delete one or more events from a role policy, click **Delete** in the Actions column of the row for the applicable event.

*After you have deleted the desired events from the role policy and clicked **Save**, the Policy Manager: Role page opens again, and the events you deleted are no longer displayed. Also, the number (in the parentheses) on the [Policy Manager: Core Events tab](#) now reflects the new number of core events in your system.*

Duplicate a role policy

Duplicate a role policy if you want to add a role policy that is the same or similar to an existing role policy so you can use it as a template and therefore create the new one more easily and quickly.

To duplicate a role policy, do the following:

1. From the Actions drop-down menu, click **Duplicate**.

The Role Policy page opens, showing the policy's details populated in the applicable fields and all of the events related to the selected role policy listed below that.

2. Change the policy name (each policy name must be unique).
3. Select one or more host groups from the Host Groups field or enter the applicable IP address(es) or IP range(s).
4. Make any other edits, if necessary.
5. Click **Save**.

What you can do next

If your policies are not behaving correctly, synchronize your Flow Collector with the Manager

At times you may need to synchronize the Flow Collectors on your network with your Manager. Some reasons may be as follows:

- You observe that your Flow Collectors are not working properly.
- You have added a new configuration (e.g., role policy, host group), but while you were doing so your Manager stopped communicating with any or all of your Flow Collectors. In this case, the policy will have been saved to your Manager, but it may not have propagated to any or all of your Flow Collectors.

To synchronize one or more Flow Collectors to your Manager, log in to your Manager client interface and complete the following steps:

1. From the navigation menu, click **Configure > Flow Collectors**.
2. Click the **Advanced** tab.
3. In the top left corner, from the Flow Collector drop-down list, choose the applicable Flow Collector.
4. In the Synchronize section, click **Synchronize**.

5. Repeat Step 3 through Step 4 for the remaining Flow Collectors that need to be synchronized with the Manager.

Related topics

[Managing Core Policies](#)

[Configuring Host Policies](#)

[Editing Default Policies](#)

[Security Event List](#)

[About Alarm Categories](#)

[About Alarm Severities](#)

Editing Default Policies

Use this page to edit the default policies. For information about viewing and maintaining core policies, see [Managing Core Policies](#).

How to find this page

Do the following to access this page

1. From the main menu, choose **Configure > Detection > Policy Management**.
2. Click the Core Events tab.

The Policy Management: Core Events tab opens. The role policy's events are listed alphabetically, first by Category events and then by Security events.

What you can do on this page

Edit a default policy

On this page you can edit the options for the When Host is Source and When Host is Target fields. When you edit any settings or add any events to an existing policy, a green dog ear is displayed in the top right corner of the row. A red dog ear indicates an input error. You cannot delete an event from a default policy.

The following list shows the most restrictive setting at the top and ends with the least restrictive setting.

- Off
- On
- On + Alarm

Refer to the following table for an explanation of the options.

If this option...	Is selected for "When Host Is Source," then this security event will...	Is selected for "When Host is Target," then this security event will...
Off	Not collect index points nor alarm.	Not contribute index points to alarm category events nor alarm.

On	Collect index points but not alarm.	Contribute index points to alarm category events but not generate alarms.
On + Alarm	Collect index points and generate alarms.	Contribute index points to alarm category events and generate alarms.

 The Ignore option is not available for default policies.

Restore a default policy to its default settings

If you have previously changed a default policy and want to return to the default settings that it was configured with when your system was first set up, then from the Actions menu click **Restore Defaults**.

What you can do next

Related topics

Managing Core Policies

If your policies are not behaving correctly, synchronize your Flow Collector with the Manager

At times you may need to synchronize the Flow Collectors on your network with your Manager. Some reasons may be as follows:

- You observe that your Flow Collectors are not working properly.
- You have added a new configuration (e.g., role policy, host group), but while you were doing so your Manager stopped communicating with any or all of your Flow Collectors. In this case, the policy will have been saved to your Manager, but it may not have propagated to any or all of your Flow Collectors.

To synchronize one or more Flow Collectors to your Manager, log in to your Manager client interface and complete the following steps:

1. From the navigation menu, click **Configure > Flow Collectors**.
2. Click the **Advanced** tab.

3. In the top left corner, from the Flow Collector drop-down list, choose the applicable Flow Collector.
4. In the Synchronize section, click **Synchronize**.
5. Repeat Step 3 through Step 4 for the remaining Flow Collectors that need to be synchronized with the Manager.

[Configuring Role Policies](#)

[Configuring Host Policies](#)

[Security Event List](#)

[About Alarm Categories](#)

[About Alarm Severities](#)

Response Management: Overview



- You must be an Admin or Configuration Manager to use Response Management.
- Response Management is not active on a secondary Cisco Secure Network Analytics Manager (formerly Stealthwatch Management Console). If your primary Manager fails or is shut down and rebooted, and you want Response Management to continue responding to alarms, you must manually promote your secondary Manager to primary Manager.

As of Cisco Secure Network Analytics (formerly Stealthwatch) v7.3.0, you can perform Response Management tasks only in the Manager. When you upgrade to v7.3.0, Secure Network Analytics migrates all Response Management configurations from the Desktop Client to the Manager. If any migrated rules contain conditions that are incompatible with v7.3.0, Secure Network Analytics disables them during the migration process. If this occurs, please correct the affected rules before you re-enable them.

Example: Previous Manager System Alarm default rules contained some UDPD alarm type conditions. These conditions are now associated with UDPD System Alarm rules, so you can no longer use them for an Manager System Alarm rule.

Use Response Management to receive alerts and prompt the Manager to perform pre-defined actions when abnormal behavior occurs on your network.

To achieve this, Secure Network Analytics first collects baseline measurements to determine what constitutes normal behavior on your network. Secure Network Analytics then observes activity that deviates from expected behavior and alerts you to this anomalous behavior using the rules and actions you configure in Response Management.

A rule is a network, security, or system event defined in the Manager that contains one or more conditions. Conditions can be stand-alone or organized into condition sets. Both conditions and condition sets can contain additional sub conditions and sub condition sets; you can nest conditions and condition sets to as many levels as needed. See [Response Management: Add Sub Conditions.](#))

Rule is triggered if:

Domain that originated the alarm is and:

ALL of the following is true: **Main parent row**

Processing time is between 12:00 AM and 11:59 AM **A condition directly under the main parent row**

ALL of the following is true: **Condition set # 1 (contains two conditions)**

Severity is Major or higher

Host Group of Source Host is or

NONE of the following is true: **A sub condition of set # 1**

IP Address of Source Host is 192.168.1.1

ANY of the following is true: **Condition set # 2 (contains two conditions)**

Type is Addr_Scan/tcp

Type is Addr_Scan/udp

The operator in the main parent row applies to all conditions and condition sets that follow. Since this operator is "ALL," all the conditions and condition sets need to be true in order for this rule's associated action(s) to be triggered.

You specify one or more actions that you want to associate with the rule. When network behavior matches a rule's conditions, the Manager generates the response (action) associated with the rule. For example, the Manager might send an email to a specific person or transmit a customizable syslog message to a server.

Open Response Management

From the main menu, choose **Configure > Response Management**.

The Rules tab is selected by default. This page displays information related to the rules you have configured.



To view the help topic for the Response Management home page (the Rules tab), see [Response Management Rules: Overview](#).

Response Management: Table of Contents

Response Management: Sort Columns

To reverse the sort order for a column, click the column name. Use the column name to switch between ascending and descending order. You can also click the arrow displayed to the right of the column name (after you click the column name). An up arrow indicates that the column is sorted in ascending order, and a down arrow indicates descending order.

Related Tasks

[Response Management: Add Rules](#)

[Response Management: Add Actions to a Rule](#)

[Response Management: Assign Actions to a Rule](#)

Related Topics

[Response Management: Overview](#)

Response Management: IPv4 and IPv6 Address Notation

You can use either IPv4 or IPv6 addresses for hosts or host groups. The notation forms that must be used for each type are listed in the applicable table. You must separate the IP addresses by a comma, comma space, space, or return.

The Manager displays IPv6 addresses and IPv4 addresses simultaneously. Of the advantages provided by IPv6, the fact that it supports a larger address space is probably the most well known. IPv6 supports 128-bit addresses, whereas IPv4 supports only 32-bit addresses.

The Manager continues to support IPv4. The Flow Collectors can monitor IPv6 traffic, but they are not IPv6 addressable at this time.

Examples

- An IPv6 address consists of eight blocks of hexadecimal digits separated by colons, as shown in this example (where the x's represent hexadecimal digits):

Example: xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx

- You can replace blocks of 4-digit zeros with double colons (e.g., ::), but only once per IPv6 address. (Leading zeros can also be omitted.) Consider the following address (where the x's represent hexadecimal digits):

Example: xxxx:xxxx:0000:0000:0000:0000:0000:xxxx

This address can be rewritten like this:

Example: xxxx:xxxx::xxxx.

- Like an IPv4 address, an IPv6 address may include a prefix length that specifies how many left-most bits of the address make up the prefix. In IPv4 addresses, the prefix is used to determine the subnet mask (a.k.a., CIDR notation). In IPv6 addresses, it identifies the subnet that an interface belongs to, and routers use it for forwarding.

Example: xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx/40

- Take care when compressing an IPv6 address that contains a prefix, because if not done correctly, it can result in an incorrect interpretation when the address is uncompressed. You can write an IPv4 address as an IPv6 address, like this:

Example: xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:192.168.0.2

- If the first 6 blocks are zeros, you can compress the address like this:

Example: ::192.168.0.2

Response Management: IPv4 Address Notation

Response Management: IPv6 Address Notation

Response Management: IPv4 Address Notation

If you enter this type of IP address...	You must use this formatting...
Single IP address	10.52.1.55 Includes only the host at 10.52.1.55.
Trailing dot subnet	10.52. Includes any IP address from 10.52.0.0 to 10.52.255.255. Note: You must include the trailing period (.) at the end.
Classless Inter-Domain Routing (CIDR) notation (i.e., using a “/” to indicate the masked bits)	10.52.1.0/24 Includes any IP address from 10.52.1.0 to 10.52.1.255. Note: Use this notation if you prefer to enter the IP address using a “mask.”
Range of networks	10-11. Includes any IP address from 10.0.0.0 to 11.255.255.255 10.52-53. Includes any IP address from 10.52.0.0 to 10.52.255.255, and from 10.53.0.0 to 10.53.255.255 10.52-55.3. Includes any IP address from 10.52.3.0 to 10.52.3.255; from 10.53.3.0 to 10.53.3.255; from 10.54.3.0 to 10.54.3.255; and from 10.55.3.0 to 10.55.3.255 Note: Be sure to include the trailing period (.) at the end.
Range of hosts	10.52.1.0-10 Includes any IP address from 10.52.1.0 to 10.52.1.10

If you enter this type of IP address...	You must use this formatting...
	10.52.0.0-10.52.255.255 Includes any IP address from 10.52.0.0 to 10.52.255.255
Comma-separated list	10.52.1.10,10.52.1.50,10.100.1.20 Includes only these three hosts.

Response Management: IPv4 and IPv6 Address Notation

Response Management: IPv6 Address Notation

Response Management: IPv6 Address Notation

If you enter this type of IP address...	You must use this formatting...
Single IP address	2001:0DB8:0000:0056:0000:ABCD:EF12:3456 2001:DB8:0:56:0:ABCD:EF12:3456 2001:DB8::56:0:ABCD:EF12:3456 2001:DB80:0:56::ABCD:EF12:3456 2001:DB80:0:56::ABCD:239.18.52.86
Global Routing Prefix subnet	2001:DB8:0:56::/64
Range of hosts	2001:DB8:0:56:ABCD:EF12:3456:1-10 2001:DB8:0:56:ABCD:EF12:3456:1- 2001:DB8:0:56:ABCD:EF12:3456:10

Response Management: IPv4 and IPv6 Address Notation

Response Management: IPv4 Address Notation

Response Management: Workflow

To add a rule, its conditions, and an associated action (and syslog format if applicable), complete the following work flow. You must complete all of the steps to create a functioning rule.

 You can reverse steps 1 and 2.

1. [Add the rule.](#)
2. [Add the associated conditions.](#)
3. [Add the associated action.](#)
 - [Add a syslog format to the rule IF you chose a syslog message as an action for the rule.](#)
4. [Assign the action to the rule.](#)

What to do Next

[Response Management: Add Rules](#)

Related Tasks

[Response Management: Add the Main Parent Operator and Conditions](#)

[Response Management: Add Condition Sets](#)

[Response Management: Add Sub Conditions](#)

[Response Management: Add Actions to a Rule](#)

[Response Management: Add a Syslog Format to a Rule](#)

[Response Management: Add Actions to a Rule](#)

Related Topics

[Response Management: Overview](#)

Alarm Variables and Rule Types

The following table shows which alarm variables can be used with the various rule types in Response Management. Use this information when you are formatting syslog formats or the body or subject in an email action.

Refer to the following key:

- An X indicates that the alarm variable applies to the rule type displayed in the applicable column.
- A – (dash) indicates that the alarm field does *not* apply to the rule type displayed in the applicable column or it is optionally configurable for your system.

This alarm variable...	Is associated with this rule type...				
	Manager System Alarm	Flow Collector System Alarm	Exporter or Interface Alarm	Host Alarm	Host Group Relationship Alarm
alarm_category_id	X	X	X	X	X
alarm_category_name	x	X	X	X	X
alarm_id	X	X	X	X	X
alarm_note	X	X	X	X	X
alarm_severity_id	X	X	X	X	X
alarm_severity_name	X	X	X	X	X
alarm_status	X	X	X	X	X

This alarm variable...	Is associated with this rule type...				
	Manager System Alarm	Flow Collector System Alarm	Exporter or Interface Alarm	Host Alarm	Host Group Relationship Alarm
alarm_type_id	X	X	X	X	X
alarm_type_description	X	X	X	X	X
alarm_type_name	X	X	X	X	X
details	X	X	X	X	X
device_id	X	X	X	X	X
device_name	X	X	X	X	X
device_ip	X	X	X	X	X
device_type_id	X	X	X	X	X
device_type_name	X	X	X	X	X
domain_id	X	X	X	X	X
domain_name	X	X	X	X	X
end_active_time	X	X	X	X	X
end_active_time_local	X	X	X	X	X

This alarm variable...	Is associated with this rule type...				
	Manager System Alarm	Flow Collector System Alarm	Exporter or Interface Alarm	Host Alarm	Host Group Relationship Alarm
exporter_hostname	X	X	X	—	—
exporter_ip	X	X	X	—	—
exporter_label	X	X	X	—	—
host_group_relationship_description	—	—	—	—	X
host_group_relationship_id	—	—	—	—	X
host_group_relationship_label	—	—	—	—	X
host_group_relationship_map_id	—	—	—	—	X
host_group_relationship_map_name	—	—	—	—	X
host_group_relationship_name	—	—	—	—	X
interface_	—	—	X	—	—

This alarm variable...	Is associated with this rule type...				
	Manager System Alarm	Flow Collector System Alarm	Exporter or Interface Alarm	Host Alarm	Host Group Relationship Alarm
index					
policy_description	—	—	—	X	—
policy_id	—	—	—	X	—
policy_name	—	—	—	X	—
port	—	—	—	X *	—
protocol	—	—	—	X *	—
smc_ip	—	—	—	X	—
source_country_code	—	—	—	X	—
source_device_type	—	—	—	X	—
source_hostname	—	—	—	X	—
source_host_group_ids	—	—	—	X	—
source_host_group_names	—	—	—	X	—
source_host_group_paths	—	—	—	X	—
source_ip	—	—	—	X	—

This alarm variable...	Is associated with this rule type...				
	Manager System Alarm	Flow Collector System Alarm	Exporter or Interface Alarm	Host Alarm	Host Group Relationship Alarm
source_label	—	—	—	X	—
source_mac_address	—	—	—	X	—
source_mac_address_vendor	—	—	—	X	—
source_multiple_ips	—	—	—	X	—
source_url	—	—	—	X	—
source_username	—	—	—	X	—
source_zone_id	—	—	—	X	—
source_zone_name	—	—	—	X	—
source_zone_path	—	—	—	X	—
spi	—	—	—	X *	—
start_active_time	X	X	X	X	X
start_active_time_local	X	X	X	X	X

This alarm variable...	Is associated with this rule type...				
	Manager System Alarm	Flow Collector System Alarm	Exporter or Interface Alarm	Host Alarm	Host Group Relationship Alarm
target_country_code	—	—	—	X	—
target_country_name	—	—	—	X	—
target_device_type	—	—	—	X	—
target_hostname	—	—	—	X	—
target_host_group_ids	—	—	—	X	—
target_host_group_names	—	—	—	X	—
target_host_group_paths	—	—	—	X	—
target_ip	—	—	—	X	—
target_label	—	—	—	X	—
target_mac_address	—	—	—	X	—
target_mac_address_vendor	—	—	—	X	—

This alarm variable...	Is associated with this rule type...				
	Manager System Alarm	Flow Collector System Alarm	Exporter or Interface Alarm	Host Alarm	Host Group Relationship Alarm
target_multiple_ips	—	—	—	X	—
target_url	—	—	—	X	—
target_username	—	—	—	X	—
target_zone_id	—	—	—	X	—
target_zone_name	—	—	—	X	—
target_zone_path	—	—	—	X	—
time	X	X	X	X	X
time_local	X	X	X	X	X

Those variables marked with an asterisk (*) in the Host Alarm column apply only to the following alarms:

- Beaconsing Host
- Bot Infected Host - Attempted C&C Activity
- Bot Infected Host - Successful C&C Activity
- Brute Force Login
- Connection From Bogon Address Attempted
- Connection From Bogon Address Successful
- Connection From Tor Attempted
- Connection From Tor Successful
- Connection To Bogon
- Connection To Bogon
- Connection To Tor
- Connection To Tor

Address Attempted	Address Successful	Attempted	Successful
• Fake Application Detected • Inside Tor Entry Detected • Port Scan	• Half Open Attack • Inside Tor Exit Detected • SSH Reverse Shell	• High SMB Peers • MAC Address Violation • Scanner Talking	• Host Lock Violation • New Host Active • Slow Connection Flood
• Suspect Long Flow • Trapped Host	• Suspect Quiet Long Flow • Watch Host Active	• Suspect UDP Activity • Watch Port Active	• Talks to Phantoms • Worm Activity
• Worm Propagation	• All custom security events		

Related topics

[Response Management: Overview](#)

Alarm Variables Table

The following table contains the list of alarm variables you can use when formatting syslog formats or the body or subject in an email action in Response Management.

Field	Description
{alarm_category_id}	Numeric ID of the alarm category (e.g., 0).
{alarm_category_name}	String name of the category (e.g., Anomaly).
{alarm_id}	Unique ID assigned to each alarm (e.g., 3Y-13Y1-QJJ2-YYA9-U).
{alarm_note}	Any note attached to this alarm.
{alarm_severity_id}	Numeric ID of the alarm severity (e.g., 4).
{alarm_severity_name}	String name of the alarm severity (e.g., Major).
{alarm_status}	Status of the alarm event. Options are ACTIVE or INACTIVE.
{alarm_type_description}	String description of the alarm type (e.g., The source IP has attempted to connect on an excessive number of ports on the target IP).
{alarm_type_id}	Numeric ID of the alarm type (e.g., 4).
{alarm_type_name}	String name of the alarm type (e.g., Port Scan).
{details}	Miscellaneous information about the alarm, such as how many concern index points had accumulated at the time of the

Field	Description
	alarm, or how much traffic was seen going through the host at that time.
{device_hostname}	Host name of the device related to this alarm.
{device_id}	Numeric ID of the device related to the alarm or alert.
{device_ip}	IP Address of the device related to this alarm (e.g., 192.168.100.99).
{device_name}	String name of the device related to this alarm (e.g., Core NetFlow Collector).
{device_type_id}	String ID of the type of device related to this alarm (e.g., swa-netflow).
{device_type_name}	String name of the type of device related to this alarm (e.g., Flow Collector (NetFlow)).
{domain_id}	ID of the domain.
{domain_name}	Name of the domain.
{end_active_time}	End active time of the alarm in ISO 8601 format, UTC time (e.g., 2020-04-21T15:48:58Z).
{end_active_time_local}	End active time of the alarm in ISO 8601 format including UTC time offset, in UTC time (e.g., 2020-04-21T15:48:58 (UTC)).
{exporter_hostname}	Name of the exporter host.
{exporter_ip}	IP address of the exporter.
{exporter_label}	Name of the exporter host, if available, followed by the exporter's IP address.

Field	Description
{host_group_relationship_description}	Description of the host group relationship.
{host_group_relationship_id}	ID for the host group relationship.
[host_group_relationship_label]	Label (i.e., that was configured on the map) for the host group relationship.
{host_group_relationship_map_id}	ID for the map in which the host group relationship exists.
{host_group_relationship_map_name}	Name of the map in which the host group relationship exists.
{host_group_relationship_name}	Name for the host group relationship.
{interface_index}	For exporter alarms only. The interface number on the exporter.
{policy_description}	Description of the policy applied to the host or host group.
{policy_id}	ID of the policy applied to the host or host group.
{policy_name}	Name of the policy applied to the host or host group.
{policy_type}	Policy type. You can choose from one of the following three types: • Default: Can apply to Inside Hosts and Outside Hosts.

Field	Description
	• Role: Can apply to groups of hosts that have similar behavior. • Host: Can apply to single hosts that have anomalous behavior and need their own policy settings.
{port}	TCP or UDP port number.
{protocol}	IP protocol number.
{sgt_id}	Group ID of subject as configured in ISE.
{sgt_name}	Group name of subject as configured in ISE.
{sgt_peer_id}	Group ID of peer as configured in ISE.
{sgt_peer_name}	Group name of peer as configured in ISE.
{smc_ip}	IP address of the Manager (e.g., 192.168.1.19).
{source_country_code}	Country code for the source host.
{source_country_name}	Country in which the host that caused the alarm is located.
{source_device_type}	Source host's device type as detected by Cisco® Identity Services Engine (ISE). To access this data, you must add a Cisco ISE as a device within the Desktop Client.
{source_host_group_ids}	IDs of the source host groups where the source host is located.
{source_host_group_names}	Names of the source host groups where the source host is located.

Field	Description
{source_host_group_paths}	Hierarchy of all the parent host groups where the source host is located.
{source_host_port}	Port used by the source host to communicate.
{source_hostname}	Name of the source host.
{source_ip}	IP address of the source host.
{source_label}	Name of the source host, if available, followed by the source's IP address.
{source_mac_address}	MAC address of the source host device (not available if the user is connected through a VPN). To access this data, you must add a Cisco ISE as a device within the Desktop Client.
{source_mac_address_vendor}	Company name associated with the source host MAC address. To access this data, you must add a Cisco ISE as a device within the Desktop Client.
{source_multiple_ips}	Boolean value that indicates the alarm has more than one IP address associated with the source.
{source_url}	URL of the web-version of the Host Snapshot for the source host.
{source_username}	User name associated with the source host in the monitored network. To access this data, you must add a Cisco ISE as a device within the Desktop Client.
{source_zone_id}	ID of the zone where the source host is located.

Field	Description
{source_zone_name}	Name of the zone where the source host is located.
{source_zone_path}	Hierarchy of all the parent zones for the zone in which the source host belongs.
{spi}	Service Profile Index (ranges from 1-163).
{start_active_time}	Date and time that the alarm condition became active (if available) in ISO 8601 format, UTC time (e.g., 2020-04-21T15:48:58Z).
{start_active_time_local}	Date and time that the alarm condition became active (if available) in ISO 8601 format including UTC time offset, in UTC time (e.g., 2020-04-21T15:48:58 (UTC)).
{target_country_code}	Country code for the target host.
{target_country_name}	Country in which the host is located that was the target of the condition that caused the alarm.
{target_device_type}	Target host's device type as detected by the Cisco ISE. To access this data, you must add a Cisco ISE as a device within the Desktop Client.
{target_host_group_ids}	IDs of the source host groups where the target host is located.
{target_host_group_names}	Names of the source host groups where the target host is located.
{target_host_group_paths}	Hierarchy of all the parent host groups to which the target host belongs.
{target_hostname}	Name, if available, assigned to the host that was the target of the condition that caused the alarm.

Field	Description
{target_ip}	IP address of the host that was the target of the condition that caused the alarm.
{target_label}	Name of the target host, if available, followed by the target's IP address.
{target_mac_address}	MAC address of the target host device (not available if the user is connected through a VPN). To access this data, you must add a Cisco ISE as a device within the Desktop Client.
{target_mac_address_vendor}	Company name associated with the target host MAC address. To access this data, you must add a Cisco ISE as a device within the Desktop Client.
{target_multiple_ips}	Boolean value that indicates the alarm has more than one IP address associated with the target.
{target_url}	URL of the web-version of the Host Snapshot for the target host.
{target_username}	User name associated with the target host in the monitored network. To access this data, you must add a Cisco ISE as a device within the Desktop Client.
{target_zone_id}	ID of the zone where the target host is located.
{target_zone_name}	Name of the zone where the target host is located.
{target_zone_path}	Hierarchy of all the parent zones for the zone in which the target host belongs.
{time}	Start or end active time of the alarm in ISO 8601 format, UTC time (e.g., 2020-04-21T15:48:58Z).

Field	Description
{time_local}	Start or end active time of the alarm in ISO 8601 format including UTC time offset, in UTC time (e.g., 2020-04-21T15:48:58 (UTC)).
{translated_destination_ip}	Translated IP for the subject.
{translated_destination_port}	Translated port for the subject.
{translated_source_ip}	Translated IP for the peer.
{translated_source_port}	Translated port for the peer.

Related topics

[Response Management: Overview](#)

Alert Variables Table

The following table contains the list of alert variables you can use when formatting syslog formats or the body or subject in an email action in Response Management.

Field	Description	Example
{alert_description}	Alert description.	Device tried to communicate with a host in a user-supplied, watchlisted country on an illegal protocol/port combination (e.g., UDP on port 22). This alert uses the Bad Protocol observation and may indicate the presence of a malicious covert communications channel.
{alert_id}	Unique ID assigned to each alert.	123
{alert_next_steps}	Next steps the customer needs to take for the alert.	Reference the supporting observations to determine why the entity used the unusual protocol/port combination to communicate with the entity in the watchlisted country. Determine what was transferred in the communication. If deemed malicious, update your firewall and blocklist rules to prevent further access with both this protocol/port combination and geolocation, unless there is a business reason for allowing it.

Field	Description	Example
{alert_priority}	Priority setting for the alert.	Low, Normal, or High
{alert_tags}	List of tags for the alert.	tag 1, tag 2
{alert_type}	Alert type.	Internal Port Scanner
{comments_count}	Number of comments for the alert.	2
{device_id}	Numeric ID of the device related to the alarm or alert.	792
{device_ips}	Source IP Addresses of the device related to the alert.	10.70.27.182, 10.70.27.183
{mitre_tactic_cti_id}	Industry-standard CTI ID of the MITRE tactic.	x-mitre-tactic--ffd5bcee-6e16-4dd2-8eca-7b3beedf33ca
{mitre_tactic_description}	Industry-standard description of the MITRE tactic.	The adversary is trying to get into your network.
{mitre_tactic_id}	Industry-standard ID of the MITRE tactic.	TA0001
{mitre_tactic_name}	Industry-standard name of the MITRE tactic.	Initial Access
{mitre_technique_cti_id}	Industry-standard CTI ID of the MITRE technique.	attack-pattern--b17a1a56-e99c-403c-8948-561df0cffe81
{mitre_technique_}	Industry-standard description of the	Adversaries may

Field	Description	Example
description}	MITRE technique.	communicate using application layer protocols to avoid detection/network filtering by blending in with existing traffic. Commands to the remote system, and often the results of those commands, will be embedded within the protocol traffic between the client and server. Adversaries may utilize many different protocols, including those used for web browsing, transferring files, electronic mail, or DNS. For connections that occur internally within an enclave (such as those between a proxy or pivot node and other nodes), commonly used protocols are SMB, SSH, or RDP.
{mitre_ technique_id}	Industry-standard ID of the MITRE technique.	T1078
{mitre_ technique_ name}	Industry-standard name of the MITRE technique	Application Layer Protocol
{observation_ names}	Observations associated with the alert.	Remote Access
{resolved}	Boolean value that indicates if the alert was resolved.	false

Field	Description	Example
{smc_ip}	IP address of the Manager.	192.168.100.9
{subnets}	Subnets associated with the alert.	10.0.0.0/8 (Default RFC1918)
{telemetry_types}	Telemetry types associated with the alert.	NetFlow, North-South
{time_first_observed}	Time the alert was first observed.	2023-07-21T11:51:15Z
{time_last_modified}	Time the Manager last updated the alert.	2023-07-21T11:51:15Z
{time_last_observed}	Time the alert was last observed.	2023-07-21T11:51:15Z
{time_natural}	Time in days and hours ago the alert was seen.	2 days, 23 hours ago
{time_object_created}	Time the initial alert object was created.	2023-07-21T11:51:15Z
{time_published}	Time the alert was published.	2023-07-21T11:51:15Z
{time_resolved}	Time the alert was resolved.	2023-07-21T11:51:15Z
{user_or_network}	Conveys whether the alert is tied to a user or a network entity.	net-link - not tied to a user user - tied to a user

Related topics

[Response Management: Overview](#)

Response Management Rules: Overview

Use the Rules tab to add Response Management rules.

What to do Next

[Response Management: Add Rules](#)

Related Topics

[Response Management: Alarm Types](#)

[Alarm Tiering](#)

[Response Management: Overview](#)

Response Management: Alarm Types

See [Alarm List: Manager System Alarms](#) to view the different alarm types you can specify when configuring rules using response management.

Related Tasks

[Response Management: Add Rules](#)

Related Topics

[Response Management Rules: Overview](#)

[Response Management: Overview](#)

Response Management: Add Rules

Create the rule for which you want to define one or more conditions. Remember, conditions can be stand-alone or organized into condition sets.

You can do any of the following, but note that before you can add a sub condition you must first add its main parent operator.

1. Click the Rules tab.
2. In the Rules table, choose **Add New Rule > [Alarm type]**.

The Rules page opens.



- You might want to use alarm tiering for some rules. See the [Alarm Tiering](#) section at the end of this procedure.
- To learn which alarm types each alarm category contains, view the [Response Management: Alarm Types](#) table.
- For information about the host alarms, see [Security Event List](#) in the Manager Help.

You can choose from the following alarm types:

- Manager System Alarm
- Exporter or Interface Alarm
- Flow Collector System Alarm
- Host Alarm
- Host Group Relationship Alarm
- UDP Director Alarm

You can also choose the following non-alarm type:

- Alerts
- Findings



Alerts and findings are not domain-specific.

3. Enter a name for the rule.
4. (Optional) Enter a description for the rule.
5. (Optional) Disable the rule by clicking the  (**Toggle**) icon (located beside the Name field) to change to Disabled Mode (). By default, Secure Network

Analytics enables the rule. You can also enable or disable a rule in the list on the Rules tab by clicking the **Toggle** icon in the Enabled column.

- When the rule is enabled, the **Toggle** icon bar is blue.
- When the rule is disabled, the **Toggle** icon bar is gray.

Secure Network Analytics disables this rule for all associated actions. Disabled rules are not triggered even when associated conditions are met.

What to do Next

[Response Management: Add the Main Parent Operator and Conditions](#)

Related Topics

[Response Management: Alarm Types](#)

[Alarm Tiering](#)

[Response Management Rules: Overview](#)

[Response Management: Overview](#)

Alarm Tiering

Alarm tiering in the system is a method of categorizing various alarms into groups based on severity. This approach translates into operational readiness. By exposing your team only to alarms that are higher priority, you can focus on well-tuned alarms that represent scenarios you should immediately investigate.

At the same time, you are not disabling lower-priority alarms. You can still investigate, correlate, and migrate them through the alarm life cycle until you have properly categorized them.

Related Tasks

[Response Management: Add Rules](#)

Related Topics


[Response Management Rules: Overview](#)

[Response Management: Overview](#)

Response Management: Add Condition Sets

Add one or more condition sets

- To add a condition set, click the  (**Add Condition Set**) icon located at the end of the main parent row. The parent row of the first condition set appears.
 - Change the selection in the Operator drop-down list in the row you just added, if applicable (options are *ANY*, *ALL*, or *NONE*).



When you add a condition, you must choose a condition type. For information about condition types, refer to the [Condition types](#) section.

- Add as many conditions as you want in this first condition set. To do this, click the **Add Condition** icon located at the end of this first condition set's parent row (refer to the second image below). You can add as many conditions as you want.

Rules | Host Alarm Cancel Save

Name: Host Alarm Rule 1 Description:

☒ Enabled (Disabled rules are not triggered even when associated conditions are met.)

Rule is triggered if:
Domain that originated the alarm is Autobots and:

ALL of the following is true: + →

Processing time is between 12:00 AM and 11:59 AM −

ALL of the following is true: − + →

Severity is Major or higher −

Host Group of Source Host is + DNS Servers x or Proxies x −

- If you choose Host Group as the condition type, click **+Select** to open the Host Groups slide-out panel from which to make your selection(s).

Select Host Groups

Search

[Collapse All](#)
[Uncheck All](#)

Inside Hosts

By Function

Client IP Ranges (DHCP Range)

DMZ

Internet Services

DNS Servers

NAT Gateway

Proxies

Load Balancer VIPs

Network Scanners

Other

Servers

VoIP

By Location

Catch All

Protected Asset Monitoring

Protected Trapped Hosts - Honeygot

Outside Hosts

Authorized External DNS Servers

Content Networks

Countries

Custom Reputation List

Trusted Internet Hosts

Cancel

Apply

- To add a second condition set, click the **Add Condition Set** icon located at the end of the main parent row. The parent row of the second condition set appears.
 - Change the selection in the Operator drop-down list in the row you just added, if applicable (options are *ANY*, *ALL*, or *NONE*).

© 2026 Cisco Systems, Inc. and/or its affiliates. All rights reserved.

- 1137 -

Rules | Host Alarm Cancel Save

Name: Host Alarm Rule 1 Description:

☒ Enabled (Disabled rules are not triggered even when associated conditions are met.)

Rule is triggered if:

Domain that originated the alarm is Autobots and:

ALL of the following is true: + ->

Processing time is between 12:00 AM and 11:59 AM -

ALL of the following is true: - + ->

Severity is Major or higher -

Host Group of Source Host is + DNS Servers x or Proxies x -

→ ANY of the following is true: - + ->

4. Add as many conditions as you want in this second condition set. To do this, click the **Add Condition** icon located at the end of the second condition set's parent row.

Rules | Host Alarm Cancel Save

Name: Host Alarm Rule 1 Description:

☒ Enabled (Disabled rules are not triggered even when associated conditions are met.)

Rule is triggered if:

Domain that originated the alarm is Autobots and:

ALL of the following is true: + ->

Processing time is between 12:00 AM and 11:59 AM -

ALL of the following is true: - + ->

Severity is Major or higher -

Host Group of Source Host is + DNS Servers x or Proxies x -

ANY of the following is true: - + ->

Type is Addr_Scan/tcp -

Type is Addr_Scan/udp -

5. To move a condition to another condition set, simply drag and drop it to where you want it to be.
6. To delete a rule or condition, click the **Minus** icon at the end of the associated row.
7. To save the configuration, click **Save**.

What to do Next

Response Management: Add Sub Conditions

Related Tasks

Response Management: Add the Main Parent Operator and Conditions

Related Topics

Response Management: Condition Types

Response Management Conditions: Overview

Response Management: Overview

Response Management Conditions: Overview

Use the Rules tab to add Response Management conditions.

What to do Next

[Response Management: Add the Main Parent Operator and Conditions](#)

Related Tasks

[Response Management: Add Condition Sets](#)

[Response Management: Add Sub Conditions](#)

Related Topics

[Response Management: Condition Types](#)

[Response Management: Overview](#)

Response Management: Condition Types

For each condition that you add, whether it's directly under the main parent operator or under another parent row, you need to choose a condition type. Not all condition types are available for each rule type. Therefore, depending on the rule type you choose, you may be able to choose the following condition types:

Processing Time

When you configure a processing time, you are telling Response Management that you're interested in the associated alarms only between a specific start and end time during each day. You can enter the times directly in the Start field and End field, or you can click the corresponding  (**Clock**) icon to do so.

If using the **Clock** icon, scroll through the list and click the desired hour and minute. Click **Set** to accept your entries.

Severity

Use this option to specify that you want to trigger the associated rule when either of the following occurs:

- An alarm of only one type of severity is observed.
- Alarms that are of a severity higher than a designated severity is observed.

Type

Use this option to specify that you want to trigger the associated rule when a certain type of security event is observed.

IP address or Range

Use this option to specify that you want to trigger the associated rule for one or both of the following reasons. You can enter either one IP address or one IP address range per condition (i.e., if you want to specify multiple IP addresses or ranges, you must add a condition for each entry).

Alarms

- When a specific source's or target's IP address is observed.
- When a specific IP address is in a particular range for a source host or target host.

Alerts

- When a specific source's IP address is observed.
- When a specific IP address is in a particular range for a source host.

See [IPv4 and IPv6 address notation](#) for additional information about IP addresses and IP address ranges.

Host Group

 Host groups don't apply to alerts.

Use this option to specify that you want to trigger the associated rule when a particular source host group or target host group is observed.

- When you select Host Group as a value in a condition, click the  (**Select**) icon, and from the slide-out panel that opens, select the applicable host group(s). Click **Apply** to accept your selections.
- If you select multiple host groups in a condition, the condition is met if any of the selected host groups match.
- If you type any letters that are anywhere within a host group name into the Search field and press **Enter**, the Host Group tree expands to display any host group names that contain the letters you entered.
- If you want to search for another host group, highlight or delete your current entry in the Search field, type your next entry and press **Enter**.
- If you want to clear the filter and view the entire host group tree to add another host group, delete your entry and press **Enter**.
- A parent host group is marked with a dash inside its corresponding check box to indicate that one or more of its child host groups have been chosen.
- If you simply want to search for a particular host group and not necessarily want to add it, delete the entry that is currently in the Search field and press **Enter**.

Host Group Relationship

 Host group relationships don't apply to alerts.

Use this option to specify that you want to trigger the associated rule when a specified amount of traffic has occurred for a specific host group relationship policy or when the

traffic exceeds a threshold that you have specified for the designated applications or services within the policy.

- When you select Host Group Relationship as a value in a condition, click the **Select** icon, and from the Select Policy slide-out panel that opens, select the applicable policy. Click **Apply** to accept your selection. You can choose from any of the following types. You can choose only one policy.
 - Maps (These are policies that were created using the Maps feature in the Desktop Client.)
 - Network diagrams (These are policies that were created using the Network Diagrams app in the Manager.)
 - Relationship policies
- Once you select and apply your selection, a label is displayed in the associated row of the applicable rule that displays the host group name and the related map, network diagram, or policy name.
- To view the policy in Policy Management, click the **⋮ (Ellipsis)** icon that is shown after the policy name and choose **Policy Management**.
- If you type any letters that are anywhere within a policy name into the Search field and press **Enter**, the Host Group tree expands to display any policy names that contain the letters you entered.
- If you want to highlight or delete your current entry in the Search field, type your next entry and press **Enter**.
- If you want to clear the filter and view the entire host group tree to add a different policy, delete your entry and press **Enter**.
- If you simply want to search for a particular policy and not necessarily want to add it, delete the entry that is currently in the Search field and press **Enter**.

Alert Type

Use this option to specify that you want to trigger the associated rule when a certain type of alert is observed.



- If Analytics is not enabled and you edit an Alert Rule, you receive an error message stating that the Alert type is not valid (since Analytics is not enabled).
- If you enable Analytics and immediately attempt to choose an Alert Type entry, the available options don't display immediately. There will be a delay until Analytics has completely loaded onto your system.

Priority

Use this option to specify that you want to trigger the associated rule when this priority status of the associated observed security event or alert is associated with it. You can choose from Low, Medium, and High.

Related Tasks

[Response Management: Add the Main Parent Operator and Conditions](#)

[Response Management: Add Condition Sets](#)

[Response Management: Add Sub Conditions](#)

Related Topics

[Response Management Conditions: Overview](#)

[Response Management: Overview](#)

Response Management: Add the Main Parent Operator and Conditions

Add the main parent operator and related conditions.

1. In the "Rule is triggered if" section, from the Domain drop-down list, choose the applicable originating domain of the alarm. Depending on the alarm type you chose in step 2, this option may not be available.
2. In the main parent row, from the Operator drop-down list, choose one of the following operators that indicates which conditions must be met before the rule's associated action(s) are triggered (refer to the second image below for an example):
 - ALL** The Manager performs the associated action(s) only if all the conditions under this option are met. When you choose *ALL* and don't add at least one condition, the rule is triggered for all alarms.
 - ANY** The Manager performs the associated action(s) if any of the conditions under this option are met. When you choose *ANY* and don't add at least one condition, the rule is never triggered.
 - NONE** The Manager performs the associated action(s) only if none of the conditions under this option are met. When you choose *NONE* and don't add at least one condition, the rule is triggered for all alarms.

Rules | Host Alarm

Cancel Save

Name

Description

☒ Enabled Disabled rules are not triggered even when associated conditions are met.

Rule is triggered if:

Domain in which the alarm originated is Autobots and:

ALL of the following is true:

You've chosen the ALL operator. If you don't add at least one condition, this rule will be triggered for all alarms. To add a condition, click the + button.

3. To add a condition, click the (Add Condition) icon located at the end of the main parent row. The row containing the first condition directly under the main parent row appears.

Since this new condition has been added directly under the main parent operator of *ALL*, the new condition uses this operator.

Rules | Host Alarm Cancel Save

Name: Host Alarm Rule 1

Description:

☒ Enabled (Disabled rules are not triggered even when associated conditions are met.)

Rule is triggered if:

Domain that originated the alarm is Autobots and:

ALL of the following is true: + ->

Processing time is between 12:00 AM and 11:59 AM -

When you add a condition, you must choose a condition type. For information about condition types, refer to the [Condition types](#) section.

4. To add another condition directly under the main parent row, click the same **Add Condition** icon again. You can add as many conditions as you want.

What to do Next

[Response Management: Add Condition Sets](#)

Related Tasks

[Response Management: Add Sub Conditions](#)

Related Topics

[Response Management: Condition Types](#)

[Response Management Conditions: Overview](#)

[Response Management: Overview](#)

Response Management: Add Sub Conditions

1. Click the **Add Condition Set** icon located at the end of the parent row of the condition set under which you want to add the sub condition. The parent row of the new sub condition appears.
 - Change the selection in the Operator drop-down list in the row you just added, if applicable (options are *ANY*, *ALL*, or *NONE*).

The screenshot shows the 'Rules | Host Alarm' configuration page. At the top, there are fields for 'Name' (Host Alarm Rule 1) and 'Description'. Below these is a toggle for 'Enabled'. The main section, 'Rule is triggered if:', contains a series of nested conditions. The first condition is 'Domain that originated the alarm is Autobots'. This is followed by a row with the operator 'ALL' and the condition 'Processing time is between 12:00 AM and 11:59 AM'. Below this is another row with the operator 'ALL' and a nested condition. This nested condition has the operator 'Severity is Major or higher'. Below that is a row with the operator 'Host Group of Source Host is' followed by 'DNS Servers' and 'Proxies'. The next row, which is highlighted with a red arrow, has the operator 'NONE' and is followed by a nested condition. The final row has the operator 'ANY' and is followed by a nested condition. Each row has control icons (minus, plus, and right arrow) at its end. The 'Add Condition Set' icon (plus and right arrow) at the end of the 'NONE' row is circled in red.

When you add a condition, you must choose a condition type. For information about condition types, refer to the [Condition types](#) section.

2. Add the sub condition. To do this, click the **Add Condition** icon located at the end of the newly added parent row for the sub condition.

In the following image, there is one sub condition nested one level in.

Rule is triggered if:

Domain that originated the alarm is Autobots and:

ALL of the following is true:

Processing time is between 12:00 AM and 11:59 AM

ALL of the following is true:

Severity is Major or higher

Host Group of Source Host is DNS Servers or Proxies

NONE of the following is true:

IP Address of Source Host is 192.168.1.1

ANY of the following is true:

Type is Addr_Scan/tcp

Type is Addr_Scan/udp

- Repeat step 2 for as many nested sub conditions you want to add under the operator *None*.
- To add a sub condition under a different existing condition, repeat steps 1 and 2.

What to do Next

Response Management: Add Actions to a Rule

Related Tasks

Response Management: Add the Main Parent Operator and Conditions

Response Management: Add Condition Sets

Related Topics

Response Management: Condition Types

Response Management Conditions: Overview

Response Management: Overview

Response Management Actions: Overview

Use the Actions tab to add Response Management actions.

What to do Next

[Response Management: Add Actions to a Rule](#)

Related Topics

[Response Management: Alarm Types](#)

[Response Management: Overview](#)

Response Management: Action Types

The action types to choose from are listed below. The fields displayed on the page depend on which type you choose.

The following actions apply to *both* alarms and alerts.

- Email
- ISE ANC Policy
- Syslog Message
- Cisco Webex Webhook
- Webhook

The SNMP Trap action applies *only* to alarms.

Email

Before an email action can take effect, make sure you configure the SMTP server that the Manager can use:



Open the **Central Management Inventory**. For the applicable Manager, choose **Edit Appliance Configuration** from the Actions column. On the General tab, scroll to the **SMTP Configuration** section to configure the settings.

Use this action type to send an email to one or more designated recipients for each qualified rule (all of the rule's conditions have been met) with which this action is associated.

1. From the Add New Action dropdown menu, choose either **Email (Alarm)** or **Email (Alert)**.
2. Enter a name for this email action. It does not have to be unique.
3. (Optional) Enter a description for this email action.
4. Enter one or more email addresses in the To field.
 - You can specify any email address; it does not need to be an email address of an existing Secure Network Analytics user.
 - You can either click in the field to see a list containing Secure Network Analytics username and email addresses, or you can begin to type either the user name that the user enters when logging in to the system, the user's full name (first and last name), or the email address. As you do this, Secure Network Analytics

suggests a list of users in a drop-down list. Select the desired entry and press **Enter** to accept the email address.

- If you later delete or modify Secure Network Analytics usernames or email addresses (that exists in the drop-down list), Secure Network Analytics does not update them in any email actions in which you have used them.
5. In the Subject field, enter a subject. When you enter a value in this field, it replaces the default value. For alarms, the default value includes an alarm state (active or inactive), the alarm type name, and other details depending on the alarm or type. For alerts, it contains the alert type.
 6. Use the Body field to add additional information to the email.
 - Type or paste the content into the Body field. You can enter a maximum of 2000 characters.
 - You can also include alarm or alert variables in the Body field.
 - When using variables (which are replaced with values), you can click **+ Alarm Variables** or **+ Alert Variables**, or type the { character (opening curly bracket) to view the Alarm or Alert Variables table and descriptions (you can also hover over a variable to see a description). To search for a particular alarm or alert variable, type { and the beginning letters of the alarm or alert variable until you see the name appear. Click a variable to add it to the message.

The types of variables that are available in the Body field depend on whether you chose an email action for an alarm or alert from the drop-down list on the Actions tab.

If configuring an email for an...	Use this variables list...	Examples...
Alarm	alarm variables	• alarm_id={alarm_id} For instance, 123 • https://hostname/path/?source_ip={source_ip}
Alert	alert	• alert_priority={alert_priority}

	variables	For instance, Medium • alert_type={alert_type} For instance, Internal Port Scanner

7. To see a preview of the completed email, click **Preview**. You will see the text you entered in the Body field as well as the alarm or alert content (a sample alarm or alert is used).
8. To send a test email message to the specified email addresses using a sample alarm or alert, click **Test Action**.
9. Click **Save**.

ISE ANC Policy

Use this action to direct Cisco® ISE (Identity Services Engine) to apply an ANC (Adaptive Network Control) policy to a source host or target host for which a host alarm has been triggered.



Before you can create an ISE ANC Policy action, you must first configure a Cisco ISE Cluster using the **Configure > INTEGRATIONS Cisco® ISE** option from the Manager main menu.

Complete the following steps to configure an ISE ANC Policy action.

1. From the Add New Action dropdown menu, choose either **ISE ANC Policy (Alarm)** or **ISE ANC Policy (Alert)**.
2. Enter a name for the ISE ANC Policy action. It does not have to be unique.
3. (Optional) Enter a description for the automated ISE ANC Policy action.
4. Choose the applicable ISE cluster from the ISE Cluster drop-down list. The domain is shown in parentheses after the ISE cluster name. The ISE ANC Policy action executes only for rules created for the domain name shown in this field. Secure Network Analytics will not implement this action for any rule it's assigned to that is configured for another domain.
5. Choose the applicable ANC policy from the ANC Policy drop-down list. These are policies that have been created in the ISE management interface.
6. Indicate whether you want this action to apply to the source host or target host.



For alarms, the selected policies apply to a source host IP address or a target host IP address. However, since alerts can have multiple observations associated, the selected policies can apply to multiple source IP addresses.

You can use the ISE ANC Policy Assignments Report to monitor ISE ANC policy assignments and to find out which hosts were assigned an ANC policy using Response Management (so that you can quickly determine for which hosts you need to change their policy). For more information about this report, see [ISE ANC Policy Assignments Report](#).

SNMP Trap

Use this action type to configure an SNMP notification. The Manager acts as an SNMP agent that sends SNMP notifications (unsolicited, asynchronous messages) to one or more SNMP Managers. The SNMP Trap action can only be applied to alarms.

1. From the Add New Action dropdown menu, choose **SNMP Trap**.
2. Enter a name for the SNMP trap action. It does not have to be unique.
3. (Optional) Enter a description for the SNMP trap action.

Refer to the following table when completing the remaining fields.

Field	Description
Name	Enter a name for your SNMP trap action.
Description	Enter a description for your SNMP trap action.
Enabled	Use the Enabled toggle to enable or disable your SNMP trap action.
IP Address	IP address of the SNMP Manager (the server that can receive SNMP traps). You can use IPv4 or IPv6 formatting.
UDP Port	The destination UDP port. The SNMP Manager must listen using this port. The default port is 162.

Version	The SNMP version used by response management.
If you are adding a configuration for version 3, you also need to configure the following values for the user-based security model RFC 3414).	
Username	The unique name you assign to the user (the receiver).
Engine ID	SNMPv3 entities use this to uniquely identify them. SNMP agents running on the Manager include this Engine ID in all SNMPv3 notifications. For some SNMP Managers (e.g. net-snmp), this value must be entered when configuring an SNMP Trap Action. Secure Network Analytics Manager SNMP Engine ID is based on the enterprise OID and the IP address of the Manager.
Security Type	You can choose from the following options: • Authentication: Secure Network Analytics determines that messages were sent from a valid source. • Authentication & Encryption: Secure Network Analytics determines that messages were sent from a valid source and scrambles the packet contents to prevent access by unauthorized sources. • None: Secure Network Analytics does not perform any of the actions described in the previous two options.
Authentication Password	This password enables you to authenticate a user (8 characters minimum).

Authentication Protocol	Use this protocol to require user authentication. This protocol DOES NOT encrypt messages.
Encryption Password	This password enables you to encrypt and decrypt SNMP.
Encryption Protocol	Use this protocol to require encryption. This protocol DOES encrypt messages.

3. To send a test SNMP notification to the specified SNMP Manager, click **Test Action**.

Syslog Message

Use this action type to send a syslog message to a specified syslog server.

1. From the Add New Action dropdown menu, choose either **Syslog Message (Alarm)** or **Syslog Message (Alert)**.
2. Enter a name for the syslog message action. It does not have to be unique.
3. (Optional) Enter a description for the syslog message action.
4. Enter either the host name of the syslog server in the Syslog Server Address field or the appropriate IP address. You can use either an [IPv4 or IPv6 address](#).

If you enter a host name, the Manager must be able to resolve the host name to the IP address either via DNS (see [DNS Server and DNS Cache](#)) or the static host name list in Central Manager (see [Local Resolution](#)). This occurs when you save the rule and the Manager rejects unknown host names.

5. Define the UDP port number to be used for receiving syslog messages. The default port is 514.
6. Choose the applicable format.
 - Custom:** This is the default option. When you select this option, you can choose from the Syslog Format field the syslog format you have configured on the Syslog Formats page (from the Syslog Formats drop-down list).
 - CEF (Common Event Format):** When you choose this format, Response Management changes the content of the syslog message to adhere to ArcSight Common Event Format (CEF) standards. CEF messages contain generic alert information such as the Cisco product name, the version, etc.

7. To send a test message to the specified syslog server, click **Test Action**.
8. Click **Save**.

Cisco Webex Webhook

Use this action type to integrate Secure Network Analytics with Cisco Webex.



Before configuring the Cisco Webex integration in the Manager, connect the [Incoming Webhooks](#) app to your Webex.

1. From the Add New Action dropdown menu, choose either **Cisco Webex Webhook (Alarm)** or **Cisco Webex Webhook (Alert)**.
2. Enter a name for the Webhook action. It does not have to be unique.
3. (Optional) Enter a description for the Webhook action.
4. Enter the URL of the receiving endpoint. It must use HTTPS protocol.
5. (Optional) If you want Response Management to send requests using the proxy server that is configured for Secure Network Analytics, check the **Use Internet Proxy** check box. For information about how to configure Internet Proxy settings, see [Internet Proxy](#).
 - You cannot use the proxy server if it has been configured using NTLM (Windows Challenge/Response) authentication.
6. Use the Message field to define the message that you want to pass through Webex.
 - Type or paste the content into the Message field. You can enter a maximum of 1000 characters.
 - You can also include alarm variables or alert variables in the Message field, depending on what kind of rule you are using.
 - When using variables (which are replaced with values), you can click **+ Alarm Variables** or **+ Alert Variables**, or type the { character (opening curly bracket) to view the Alarm or Alert Variables table and descriptions (you can also hover over a variable to see a description). To search for a particular alarm or alert variable, type { and the beginning letters of the alarm or alert variable until you see the name appear. Click a variable to add it to the message.

The types of variables in the Message field depend on whether you chose a Webex Webhook for an alarm or alert from the drop-down list on the Actions tab.

If configuring a Webex message for an...	Use this variables list...	Examples...
Alarm	alarm variables	- alarm_id={alarm_id} For instance, 123 - https://hostname/path/?source_ip={source_ip}
Alert	alert variables	- alert_priority={alert_priority} For instance, Medium - alert_type={alert_type} For instance, Internal Port Scanner

- To see a preview of the completed Webex message, click **Preview**. You will see the text you entered in the Message field as well as the alarm or alert content (a sample alarm or alert is used).
- To send a test Webex message using a sample alarm or alert, click **Test Action**.
- Click **Save**.

Webhook

Use this action to integrate Secure Network Analytics with external systems over Web service, typically REST API. When you execute this action, Secure Network Analytics sends an HTTP POST payload with alarm data to the configured Webhook URL. Secure Network Analytics uses your Manager appliance identity certificate, so client-side certificate authentication can be used. The receiving endpoint must be trusted by the Manager, so you need to add its certificate to Manager Trust Store in Central Management.



For more information about adding certificates to the trust stores, refer to the [SSL/TLS Certificates for Managed Appliances Guide](#).

You can direct Secure Network Analytics to access the receiving endpoint or use an Internet proxy that you configure in Central Management for the applicable Manager. If it

is taking a significant amount of time for the receiving endpoint to process incoming requests or the Internet Proxy is experiencing high latency or low bandwidth, then the processing time for these alarms may lengthen. In this case, depending on the rate of alarms that are sent via the webhook, Secure Network Analytics may drop some alarms before they are processed.

1. From the Add New Action dropdown menu, choose **Webhook**.
2. Enter a name for the Webhook action. It does not have to be unique.
3. (Optional) Enter a description for the Webhook action.
4. Enter the URL of the receiving endpoint. It must use HTTPS protocol.
5. (Optional) If you want Response Management to send requests using the proxy server that is configured for Secure Network Analytics, check the **Use Internet Proxy** check box. For information about how to configure Internet Proxy settings, see [Internet Proxy](#).
 - You cannot use the proxy server if it has been configured using NTLM (Windows Challenge/Response) authentication.
6. (Optional) To configure a basic authentication for all requests that the Man makes to the Webhook receiving endpoint, check the **Basic Authentication** check box.
7. (Conditional) When you check the **Basic Authentication** check box, you must complete the Username and Password fields.
8. To send a test alarm using Webhook to the configured Webhook URL, click **Test Action**.

The Webhook action sends HTTP POST payload to the specified URL with alarm data in the following format:

```
{
  "event": {
    "alarm_category_id": "57",
    "alarm_category_name": "Anomaly",
    "alarm_id": "31-1FT4-D05H-FPYO-F",
    ...the rest of alarm fields
  }
}
```

Related Tasks

[Response Management: Add Actions to a Rule](#)

Related Topics

[Response Management Actions: Overview](#)

[Response Management: Overview](#)

Response Management: Add Actions to a Rule

Add the action(s) that you want the Manager to perform when responding to each qualified rule (a qualified rule is one where all its conditions have been met). An action can be assigned to multiple rules.

1. Click the Actions tab.
2. If you want to enable or disable any existing actions listed in the Actions table, in the Enabled column for the associated action, click the **Toggle** icon to switch to the desired mode.

The number in the Used by Rules column indicates the number of rules to which the associated action is currently assigned.

- When the rule is enabled, the **Toggle** icon bar is blue (.
- When the rule is disabled, the **Toggle** icon bar is gray (.

Response Management disables this action for all associated rules. It does not trigger disabled actions even when a rule's associated conditions are met.



You can also enable or disable an action in the list on the Actions tab by using the **Toggle** icon in the Enabled column.

3. In the Actions table, choose **Add New Action > Action Type**.

The Actions page for the associated rule opens. The fields displayed on this page vary depending on the action type you chose.

4. Go to the next section, [Action Types](#), to continue.

System performance regarding actions

As you add more rules, the length of time required for Secure Network Analytics to perform an action increases.

To prevent denial of service, Response Management performs actions until a certain limit is reached; then no further actions execute. The action response time lengthens as the volume of alarms increases. If this situation continues for more than 30 minutes, Secure Network Analytics begins to drop some alarms before they can be processed.

For example, let's say you have two rules that each contain two email actions. During the first minute, 1050 alarms occur for these two rules, but only 1000 emails are sent during the first minute since a maximum of 1000 emails can be sent per minute. Now 50 emails remain that need to be sent.

These 50 emails are sent during the second minute. But let's say that during this second minute an additional 1150 alarms occur. So now during the second minute, 1200 emails need to be sent. But again, because of the 1000 email per minute maximum, only 1000 are sent during this second minute. Now 200 remain that need to be sent.

If this cycle continues for more than 30 minutes, Secure Network Analytics begins to drop some alarms before they can be processed.

Response Management processes this type of action...	At a rate of...
Email	1000 alarms per minute.
ISE ANC	100 alarms per minute.
Webhook	5000 alarms per minute.

What to do Next

[Response Management: Add a Syslog Format to a Rule](#)

Related Topics

[Response Management: Alarm Types](#)

[Response Management Actions: Overview](#)

[Response Management: Overview](#)

Response Management: Syslog Formats Overview

Use the Syslog Formats tab to add a syslog format to a rule in Response Management.

What to do Next

[Response Management: Add a Syslog Format to a Rule](#)

Related Topics

[Facility Codes for Event Messages](#)

[Severity Levels for Event Messages](#)

[Response Management: Overview](#)

Response Management: Add a Syslog Format to a Rule

To use a Custom format when creating a Syslog Message action, you must configure the format first. The format defines the structure and information of the messages that will be sent to your syslog server.

1. Click the Syslog Formats tab.
2. In the Syslog Format table, click **Add New**.
3. From the dropdown menu, choose either **Syslog Message (Alarm)** or **Syslog Message (Alert)**.
4. Enter a name for the syslog format.
5. If desired, enter a description for the syslog format.
6. Choose a facility code. The pre-selected facility code is 16 (Local Use). Each event message is categorized into a facility code. This code denotes the program type that logged the message.
7. Choose a severity level that you want associated with each event message. The pre-selected severity level is 3 (Error). The severity level indicates the level of importance that Secure Network Analytics assigns it. The lower the value, the more severe the event.
8. Define the message that you want to pass through Syslog. This content consists of alarm fields or alert fields that are replaced with values. You can add content using any one of the following methods:
 - Type or paste the content into the Message field. You can enter a maximum of 1000 characters.
 - You can also include alarm variables or alert variables in the Message field, depending on what kind of rule you are using.
 - When using variables (which are replaced with values), you can click **+ Alarm Variables** or **+ Alert Variables**, or type the { character (opening curly bracket) to view the Alarm or Alert Variables table and descriptions (you can also hover over a variable to see a description). To search for a particular alarm or alert variable, type { and the beginning letters of the alarm or alert variable until you see the name appear. Click a variable to add it to the message.

The types of variables in the Message field depend on whether you chose a syslog message for an alarm or alert from the drop-down list on the Syslog Formats tab.

If configuring a syslog message for an...	Use this variables list...	Examples...
Alarm	alarm variables	- alarm_id={alarm_id} For instance, 123 - https://hostname/path/?source_ip={source_ip}
Alert	alert variables	- alert_priority={alert_priority} For instance, Medium - alert_type={alert_type} For instance, Internal Port Scanner

- To see a preview of the message with sample alarm data or alert data, click **Preview**. You will see the text you entered in the Body field as well as the alarm or alert content (a sample alarm or alert is used).
- To save the configuration, click **Save**.

What to do Next

Response Management: Assign Actions to a Rule

Related Topics

Facility Codes for Event Messages

Severity Levels for Event Messages

Response Management: Syslog Formats Overview

Response Management: Overview

Facility Codes for Event Messages

This table contains the various facility codes into which you can categorize event messages.

Facility Code	Description
0	Kernel Messages
1	User-Level Messages
2	Mail System
3	System Daemons
4	Security/Authentication Messages
5	Internal syslogd Messages
6	Line Printer Subsystem
7	Network News Subsystem
8	UUCP Subsystem
9	Clock Daemon
10	Security/Authorization Messages
11	FTP Daemon
12	NTP Subsystem
13	Log Audit
14	Log Alert
15	Clock Daemon
16	Local Use 0 (local0)

17	Local Use 1 (local1)
18	Local Use 2 (local2)
19	Local Use 3 (local3)
20	Local Use 4 (local4)
21	Local Use 5 (local5)
22	Local Use 6 (local6)
23	Local Use 7 (local7)

Related Tasks

[Response Management: Add a Syslog Format to a Rule](#)

Related Topics

[Severity Levels for Event Messages](#)

[Response Management: Syslog Formats Overview](#)

[Response Management: Overview](#)

Severity Levels for Event Messages

This table contains the various severity levels that you can associate with an event message.

Severity Level	Description
0	Emergency: System is Unusable
1	Alert: Action Must Be Taken Immediately
2	Critical: Critical Conditions
3	Error: Error Conditions
4	Warning: Warning Conditions
5	Notice: Normal But Significant Condition
6	Informational: Informational Messages
7	Debug: Debug-Level Messages

Related Tasks

[Response Management: Add a Syslog Format to a Rule](#)

Related Topics

[Facility Codes for Event Messages](#)

[Response Management: Syslog Formats Overview](#)

[Response Management: Overview](#)

Response Management: Assign Actions Overview

Now that you have created and enabled the action, on the Rules tab you bind the action to the rule.

1. On the Rules tab, in the Actions column for the appropriate rule, choose **Edit** from the context menu.

At the top of the page, the configured settings for the associated rule are displayed. At the bottom of the page, the Associated Actions table is displayed. This table contains two sections: active and inactive. The number in the Used by Rules column indicates the number of rules to which the associated action is currently assigned.

2. In the desired section in the Associated Actions table, in the Assigned column for the applicable action, click the **Toggle** icon to switch to Enabled Mode (, icon bar is blue) for each action you want the Manager to execute for the associated rule (alarm) where applicable.
3. To save the configuration, click **Save**.

You have now assigned the assigned an action or actions to the associated rule.

What to do Next

Response Management: Assign Actions to a Rule

Related Topics

Response Management: Overview

Response Management: Assign Actions to a Rule

Now that you have created and enabled the action, on the Rules tab you bind the action to the rule.

1. On the Rules tab, in the Actions column for the appropriate rule, choose **Edit** from the context menu.

At the top of the page, the configured settings for the associated rule are displayed. At the bottom of the page, the Associated Actions table is displayed. This table contains two sections: active and inactive. The number in the Used by Rules column indicates the number of rules to which the associated action is currently assigned.

2. In the desired section in the Associated Actions table, in the Assigned column for the applicable action, click the **Toggle** icon to switch to Enabled Mode (, icon bar is blue) for each action you want the Manager to execute for the associated rule (alarm) where applicable.
3. To save the configuration, click **Save**.

You have now assigned the assigned an action or actions to the associated rule.

Related Topics

[Response Management: Assign Actions Overview](#)

[Response Management: Overview](#)

Response Management: Rules, Actions, and Syslog Formats Overview

You can perform the following actions with existing rules, actions, and syslog formats.

Related Tasks

[Response Management: Edit a Rule, Action, or Syslog Format](#)

[Response Management: Delete a Rule, Action, or Syslog Format](#)

[Response Management: Duplicate a Rule or Action](#)

Related Topics

[Response Management: Overview](#)

Response Management: Edit a Rule, Action, or Syslog Format

1. On the applicable tab, in the Actions column for the component you want to edit, click **Edit** from the context menu.
2. Make the desired edits on the Edit page, and click **Save**.

Related Tasks

[Response Management: Delete a Rule, Action, or Syslog Format](#)

[Response Management: Duplicate a Rule or Action](#)

Related Topics

[Response Management: Overview](#)

Response Management: Delete a Rule, Action, or Syslog Format

Delete a rule, action, or syslog format

On the applicable tab, in the Actions column for the component you want to delete, click **Delete** from the context menu.

Related Tasks

[Response Management: Edit a Rule, Action, or Syslog Format](#)

[Response Management: Duplicate a Rule or Action](#)

Related Topics

[Response Management: Overview](#)

Response Management: Duplicate a Rule or Action

Duplicate a rule or action if you want to add a new rule or action that is the same or similar to an existing one so that you can use it as a template and therefore create the new one more easily and quickly.

1. On the applicable tab, in the Actions column for the rule or action you want to duplicate, click **Duplicate** from the context menu.

The Rules or Actions page opens (depending on what you are duplicating) with the same configured settings as the rule or action you're duplicating. (Copy of) [Flow Collector name] is displayed in the Name field.

2. If desired, change the rule or action name.
3. Make any other edits, if necessary.
4. Click **Save**.

Related Tasks

[Response Management: Edit a Rule, Action, or Syslog Format](#)

[Response Management: Delete a Rule, Action, or Syslog Format](#)

Related Topics

[Response Management: Overview](#)

Classify Hosts Inside the Network: Overview

Use the Network Scanners page to review a list of hosts presumed to be network scanners. You can classify them and assign them to a host group. When they are assigned to the appropriate host group, you can use [policies](#) to monitor and control their behavior.

How to find this page

From the main menu, choose **Configure > Detection > Network Scanners**.

Classifying Hosts Inside the Network: Table of Contents

Classify Hosts Inside the Network: Concepts You Need to Understand

Secure Network Analytics identifies hosts as possible network scanners based on the volume of Concern Index (CI) points. Because of their role, network scanners generate a high volume of Addr_Scan/tcp and Addr_Scan/udp security events. Any host with 300,000 or more index points is presumed to be a network scanner and will be listed on the Network Scanners page until you classify it.

The list of hosts is updated every day at the Manager's archive hour, 12 AM (2400) UTC. If the query has not run or if there are no new potential network scanners, you will see a message on the Network Scanners page advising you of the status. You cannot change how often or when the query is run.



The Manager's archive hour is automatically set to 12 AM (2400) UTC and may not match the Flow Collector's configured archive hour.

After you have classified hosts, you will receive a message to confirm the hosts have been successfully classified. You will not see the changes reflected on this page until the next query has run. If you need to verify the hosts have been classified before the query runs, access the Host Group Editor or the Host Group Membership Report in the [Desktop Client](#). If the classified hosts continue to appear in this list, you may need to adjust settings in the role policy applied to the host group.

After the network classification query has run, it will be listed on the [Job Management](#) table in the Finished Jobs table as ANC Query. The entry in the table and the data from the query are deleted after 24 hours from the time the query completed. You cannot save the ANC query results. If the Job Management page displays a banner message stating that the system's database is nearly full, new ANC queries will not run. In this case, you need to consider deleting completed jobs to make space for additional queries.

Related Topics

[Classify Hosts Inside the Network: Classify Hosts Overview](#)

[Classify Hosts Inside the Network: Overview](#)

Classify Hosts Inside the Network: Viewing Information

Consider the following suggestions and questions when viewing the information on this page:

- Review the list of presumed hosts. Use filters on the left to narrow the list by host groups or by full or partial IP address.
- Are the host names or IP addresses ones you expect to see for your organization?
- Are those host names or IP addresses valid as network scanners?

Example: As you scan the Potential Network Scanners list, you notice an IP address that you know should be a user work station. You would want to investigate further to find out if the workstation has been compromised and is being used to scan other devices on your network.

- Is the Host Group one that you expect to see on a list of potential network scanners?

Example: If Printers is shown as a host group on a list of potential network scanners, that is anomalous, and you would want to investigate further.

- Are the applications the ones you would expect to see a network scanner use to contact and scan other hosts? Is there a host on the list that is not really a scanner but shows the full list of applications? If so, you would want to investigate further; only a scanner should have the full list.

In the Potential Network Scanners table, in the Applications column on the row for the host you want to examine, click the  (**Expand**) icon. An information box displays the full list of applications and the percentage of incoming and outgoing traffic handled by each application. (The Total Bytes column shows the total volume of traffic.) To close the information box, click the **Expand** icon again.

- Examine the ratio between Packets Sent and Packets Received. If the volume under Packets Sent is significantly higher than the volume under Packets Received, it may indicate that scanning activity is taking place. Verify that scanning activity by that host is expected and reasonable.
- Is the volume for Max CI reasonable for a network scanner? If not, you would want to investigate further.

Related Topics

[Classify Hosts Inside the Network: Classify Hosts Overview](#)

[Classify Hosts Inside the Network: Overview](#)

Classify Hosts Inside the Network: Classify Hosts Overview

You can classify multiple hosts at one time, or you can classify them individually.



If you navigate away from this page at any point before Step 3 (in which you click **Classify Hosts**), all of your selections and entries will be cleared.

Related Topics

[Classify Hosts Inside the Network: Overview](#)

Classify Hosts Inside the Network: Classify Multiple Hosts

To classify multiple hosts at one time, complete the following steps:

1. Do one of the following:

- In the Pending Classifications box at the top of the table, enter the range of [IP addresses](#) and press **Enter**. If you want to remove an IP address from the Pending Classifications box, click the X to its right or click **Clear** beside the Classify Hosts button.

[Valid formats](#) include IPv4, IPv6, or CIDR notation.

- In the Potential Network Scanners table, in the first column, select the check box for each host you want to add or select the check box at the top of the column to select all the hosts. Expand the For Selected drop-down list and click **Add Hosts**.

The hosts are shown in the Pending Classifications box.

2. In the Classify As field, click **Select** to expand the Host Groups list to select a host group. You can select only one host group.

The Host Groups Selector appears on the right and displays a tree structure with host groups. The tree is similar to the Enterprise tree in the Desktop Client. If the host group you want to use is not shown, you must add it through the Desktop Client.

The Network Scanners host group is selected by default (Host Groups > Inside Hosts > By Function > Network Scanners).

If you need to use a different host group, you can select it. If the host group you want to use is not shown, you must add it on the Enterprise tree in the Desktop Client.

3. Click **Classify Hosts**.

What To Do Next

[Classify Hosts Inside the Network: Apply a Policy](#)

Related Topics

[Classify Hosts Inside the Network: Classify a Single Host](#)

[Classify Hosts Inside the Network: Overview](#)

Classify Hosts Inside the Network: Classify a Single Host

To classify a single host, complete the following steps:

1. Do one of the following:

- In the Potential Network Scanners table, in the first column on the row for the host you want to add, select the check box, expand the For Selected drop-down list, and click **Add Hosts**.
- In the Potential Network Scanners table, in the Actions column on the right, click **Add**

To quickly locate a host, enter a string of alphanumeric characters in the search box at the top of the table. The search will use either a full or partial string, so you do not need to use a wild card character. The search function will check for a match in any of the fields in the table. Note that the Clear button does not remove a search string. To remove a search string, you need to delete it manually.

- In the Pending Classifications box at the top of the table, enter the [IP address](#). If you want to remove an IP address from the Pending Classifications box, click the X to its right or click **Clear** beside the Classify Hosts button.

[Valid formats](#) include IPv4, IPv6, or CIDR notation.

The host is shown under Pending Classifications.

2. In the Classify As field, click **Select** to expand the Host Groups list to select a host group. You can select only one host group.

The Host Groups Selector appears on the right and displays a tree structure with host groups. The tree is similar to the Enterprise tree in the Desktop Client. If the host group you want to use is not shown, you must add it through the Desktop Client.

The Network Scanners host group is selected by default (Host Groups > Inside Hosts > By Function > Network Scanners).

If you need to use a different host group, you can select it. If the host group you want to use is not shown, you must add it on the Enterprise tree in the Desktop Client.

3. Click **Classify Hosts**.

What To Do Next

Classify Hosts Inside the Network: Apply a Policy

Related Topics

[Classify Hosts Inside the Network: Classify Multiple Hosts](#)

[Classify Hosts Inside the Network: Overview](#)

Classify Hosts Inside the Network: Apply a Policy

After you have classified the scanners on your network and assigned them to a host group, apply a policy that will reduce the number of alarms to a reasonable level.

To configure and apply policy settings, complete the following steps:

1. Log in to the Desktop Client.
2. Right-click the applicable host group and choose **Configuration > Host Policy Manager**.

Related Topics

[Classify Hosts Inside the Network: Classify Multiple Hosts](#)

[Classify Hosts Inside the Network: Classify a Single Host](#)

[Classify Hosts Inside the Network: Overview](#)

Analytics: Overview

We have introduced additional advanced analytics capabilities and user interface workflows for Secure Network Analytics. These improvements include alerts, observations, and profiling capabilities that provide greater efficacy. Additionally, these new alerts require less tuning and pre-configuration.

Some of the expanded capabilities are as follows:

- Automatic role classification
- New alerts that are aligned to modern threats and mapped to MITRE ATT&CK Framework Enterprise matrix
- Proprietary "Entity Modeling" capabilities that currently exists in our Cisco XDR product, which greatly enhances baselining for normal behavior

These new analytics can be enabled independent of the existing alarms and security events within your Secure Network Analytics deployment. When you enable Analytics, related features are switched on. These additional capabilities function in parallel with your existing detections and user interface pages. Continue to monitor your alarms and security events as you do today, while also taking advantage of our new detection outcomes. Or, alternatively, enable Analytics and adjust your alarm policies to focus on these newer alerts and observations.

When you open a new alert, you can view the supporting observations that led the system to generate the alert. From these observations, you can also view additional context about the entities involved, including the traffic they transmitted and, if available, external threat intelligence.

The new alerts consume additional system resources. Review your resource consumption prior to enabling this optional feature. Feature sets for the new alerts, as well as their general availability, are limited to specific deployment types and may change with each future release. For more information, see [Analytics: System Requirements](#).



- The new alerts support only systems that contain just 1 Cisco Secure Network Analytics Data Store domain. (Your system can also contain 1 or more non-Data Store domains.) If after initially creating 1 Data Store domain on your system, you later create one or more Data Store domains, be aware that we do not provide support for this scenario.
- You cannot enable Analytics until you have deployed a minimum of 1 node and added a minimum of 1 Data Store Flow Collector (NetFlow) on your network.
- If your system is consuming sFlow, the accuracy of your detections and

the performance of your system may be impacted.

Enable Analytics

1. From the main menu, choose **Configure > Detection > Analytics**.

The Analytics Welcome page opens.

2. Click the switch so that the label displays *Analytics On*, like this: .

The Alerts Summary opens.




- Only users assigned the role of **Primary Admin** can enable and disable Analytics.
- Users assigned the web role of Analyst cannot access any configuration pages using the menu option (**Configure > Alerts**). However, all users, regardless of their assigned data roles or function roles, can access all pages and data related to the new alerts.

Disable Analytics

1. From the main menu, choose **Configure > Detection > Analytics**.

The Analytics Welcome page opens.

2. Click the switch so that the label displays *Analytics Off*, like this: .

The Analytics Welcome page remains open.

Analytics: Table of Contents

Analytics: System Requirements

Feature sets for the new alerts and their general availability are limited to specific deployment types and may change with each future release. Please verify that your system adheres to the required specifications. Systems that don't adhere to applicable specifications, or are under heavy load, may experience adverse impacts to their performance, reliability, and retention capabilities.

For more information about the system requirements required for your deployment type, see the Resource Requirements section in the [Virtual Edition Appliance Installation Guide](#).



- Alerts and Observations data is processed and stored only on the Manager that is currently in the primary role. When you promote the original primary Manager back to the primary role, you will not be able to view any alerts and observations data that was processed on the original secondary Manager while it served in the primary role.
- Only users assigned the role of **Primary Admin** can enable and disable Analytics.

In order to enable Analytics, your deployment must be configured

- on a Virtual or a Hardware Data Store deployment with any number of Flow Collectors.
- with only 1 Secure Network Analytics Data Store domain.

To install appliances, follow the instructions in the [Virtual Edition Appliance Installation Guide](#), the [x2xx Series Hardware Appliance Installation Guide](#), or the [x3xx Series Hardware Appliance Installation Guide](#).

Related Topics

[Analytics: Overview](#)

Analytics: APIs

When using APIs in Secure Network Analytics, refer to the [Secure Cloud Analytics API document](#). This document includes all the API endpoints and related documentation that can be leveraged with the new alerts' engine within Secure Network Analytics.

These endpoints are shared among Secure Network Analytics and Secure Cloud Analytics, and can be used for either deployment.



- To use these APIs in Secure Network Analytics, you still need to use *api_key*, which is shown in the following command:
cat /lancope/var/services/detections/config/api_key.
- To perform the "cat" command, please open a support case in Support Case Manager at <https://mycase.cloudapps.cisco.com/case>.

Related Topics

[Analytics: Overview](#)

Analytics: Alerts and Observations Overview

Secure Network Analytics uses dynamic entity modeling to track the state of your network. In the context of Secure Network Analytics, an entity is something that can be tracked over time, such as a host or endpoint on your network. Dynamic entity modeling gathers information about entities based on the traffic they transmit and activities they perform on your network.

From this information, Secure Network Analytics identifies:

- Roles for the entity, which are a descriptor of what the entity usually does. For example, if an entity sends traffic that is generally associated with email servers, Secure Network Analytics assigns the entity an Email Server role. The role/entity relationship can be many-to-one, as entities may perform multiple roles.
- Observations for the entity, which are facts about the entity's behavior on the network, such as a heartbeat connection with an external IP address, an interaction with an entity on a watchlist, or a remote access session established with another entity. Observations on their own do not carry meaning beyond the fact of what they represent. A typical customer may have many thousands of observations and a few alerts.

Based on the combination of roles, observations, and other threat intelligence, Secure Network Analytics generates alerts, which are actionable items that represent possible malicious behavior as identified by the system.

Analytics: Alerts and Observations from Detections Findings

Secure Network Analytics also incorporates new detection capabilities through Detections Packs. These packs contain deterministic detection rules designed to identify pattern-based threats and anomalies in your network traffic. When the conditions of these rules are met by incoming telemetry, a Detection Finding is generated.

Starting with 7.6.0 release, Detection Findings are incorporated into the operational cycle and can be seen as Analytics Alerts and Observations in the user interface.

For information about Detections Packs and how Detections Findings are converted to Analytics Alerts and Observations, refer to the [Detections Pack Guide](#).

Related Topics

[Analytics: Alert Descriptions](#)

[Analytics: Alert Prerequisite Chart](#)

[Analytics: Observation Descriptions](#)

[Analytics: Overview](#)

[Detections Pack Management](#)

Analytics: Alert Prerequisite Chart

The Alert Prerequisites Chart provides the list of the new alerts sorted by their baseline requirements as well as a brief description of the alert's meaning.

The following table provides an overview of how much history is required to generate a given alert type and possible reasons for further investigation. Note that the reasons for investigation are not guarantees that this alert indicates the listed behavior or rationale; these reasons should be considered as you further investigate the alert.

This table also lists, where applicable, any MITRE ATT&CK tactics or techniques associated with an alert type.



- After you enable Analytics, you need to wait 24 hours for all jobs to run. After this, the time periods that are listed in the History column start. In other words, these time periods do not begin counting down from the moment you enable Analytics; they begin immediately after the initial 24-hour period has completed.
- If you have enabled Analytics less than 24 hours ago and you are not receiving flows into Cisco Secure Network Analytics, the Alert Priorities page will not be populated with data.

Alert	History	Telemetry	MITRE ATT&CK Tactics	MITRE ATT&CK Techniques
Amplification Attack	0 days	NetFlow North-South	Impact	Reflection Amplification
Anomalous Domain Controller Activity <i>Formerly known as Exceptional Domain Controller</i>	7 days	NetFlow	Credential Access Lateral Movement Persistence Privilege Escalation	Abuse Elevation Control Mechanism Domain Controller Authentication
Anomalous Windows Workstation	14 days	NetFlow North-South	Execution	
Country Set Deviation	36 days	NetFlow North-South	Command and Control Exfiltration Initial Access	Valid Accounts

Alert	History	Telemetry	MITRE ATT&CK Tactics	MITRE ATT&CK Techniques
Country Watchlist: Protocol Violation <i>Formerly known as Protocol Violation (Geographic)</i>	0 days	NetFlow North-South ETA	Defense Evasion Command and Control	Non-Standard Port Non-Application Layer Protocol
Country Watchlist: Remote Access <i>Formerly known as Remote Access (Geographic)</i>	0 days	NetFlow North-South	Initial Access	Valid Accounts
DNS Traffic to Tor Proxy	0 days	Zeek Logs	Exfiltration	Exfiltration Over Alternative Protocol
Emergent Profile	14 days	NetFlow	Lateral Movement Command and Control Initial Access	Exfiltration Over Alternative Protocol External Remote Services
Empire Command and Control	1 day	NetFlow North-South	Command and Control	Non-Application Layer Protocol
Exceptional Domain Controller <i>This alert is now known as Anomalous Domain Controller Activity</i>	7 days	NetFlow	Credential Access Lateral Movement Persistence Privilege Escalation	Abuse Elevation Control Mechanism Domain Controller Authentication
Excessive Access Attempts (External)	0 days	NetFlow North-South	Credential Access	Brute Force
Excessive Connections to Network Printers	0 days	NetFlow East-West	Discovery Impact Credential Access	Brute Force Endpoint Denial of Service Peripheral Device Discovery
Geographically	30 days	NetFlow	Initial Access	External Remote

Alert	History	Telemetry	MITRE ATT&CK Tactics	MITRE ATT&CK Techniques
Unusual Remote Access		North-South	Reconnaissance	Services Active Scanning
Heartbeat Connection Count	8 days	NetFlow North-South	Command and Control	Non-Application Layer Protocol Application Layer Protocol
High Bandwidth Unidirectional Traffic	0 days	NetFlow North-South	Exfiltration Impact	Automated Exfiltration
Inbound Port Scanner	1 day	NetFlow North-South	Discovery Reconnaissance	Network Service Discovery Gather Victim Network Information
Internal Connection Spike	0 days	NetFlow North-South	Discovery Lateral Movement	Network Service Discovery Remote Services
Internal Port Scanner	7 days	NetFlow North-South East-West	Discovery	Network Service Discovery
LDAP Connection Spike	9 days	NetFlow North-South	Discovery Lateral Movement Credential Access	Network Service Discovery Domain Account
Meterpreter Command and Control Success	1 day	NetFlow North-South	Command and Control	Non-Application Layer Protocol
NetBIOS Connection Spike	9 days	NetFlow North-South	Discovery	Network Service Discovery
Network Population Spike	36 days	NetFlow East-West	Impact Defense Evasion	Network Denial of Service Impair Defenses Resource Hijacking

Alert	History	Telemetry	MITRE ATT&CK Tactics	MITRE ATT&CK Techniques
Network Printer with Excessive Connections	0 days	NetFlow North-South		
New Internal Device	21 days	NetFlow East-West	Resource Development Persistence Initial Access	Hardware Additions
New IP Scanner	9 days	NetFlow North-South	Discovery	Network Service Discovery
New Remote Access	36 days	NetFlow North-South	Initial Access	External Remote Services Valid Accounts
New SNMP Sweep	9 days	NetFlow North-South	Discovery	Network Service Discovery
New Unusual DNS Resolver	7 days	NetFlow North-South Passive DNS	Resource Development Command and Control	Compromise Infrastructure Application Layer Protocol
Non-Service Port Scanner	9 days	NetFlow North-South	Discovery	Network Service Discovery
Outbound LDAP Spike	0 days	NetFlow North-South	Reconnaissance Impact	Active Scanning Automated Exfiltration Network Denial of Service
Outbound SMB Connection Spike <i>This alert is now known as SMB Connection Spike: Outbound</i>	0 days	NetFlow North-South	Impact Reconnaissance	Active Scanning Automated Exfiltration Network Denial of Service
Persistent Remote	7 days	NetFlow	Initial Access	External Remote

Alert	History	Telemetry	MITRE ATT&CK Tactics	MITRE ATT&CK Techniques
Control Connections		North-South		Services
PetitPotam Attack Via EFS RPC Calls	0 days	Zeek Logs	Credential Access	Forced Authentication LLMNR/NBT-NS Poisoning and SMB Relay
Possible Impacket SecretDump Activity	0 days	Zeek Logs	Credential Access	Security Account Manager NTDS LSA Secrets
Potential Data Exfiltration	0 days	NetFlow East-West North-South	Exfiltration	Automated Exfiltration Exfiltration Over C2 Channel Exfiltration Over Alternative Protocol
Potential Database Exfiltration	7 days	NetFlow North-South	Exfiltration	Exfiltration Over C2 Channel
Potential Gamaredon C2 Callout	0 days	Cisco NVM	Command and Control	Command and Scripting Interpreter
Protocol Forgery	1 day	NetFlow North-South ETA	Command and Control Defense Evasion	Non-Standard Port
Protocol Violation (Geographic) <i>This alert is now known as Country Watchlist: Protocol Violation</i>	0 days	NetFlow North-South ETA	Defense Evasion Command and Control	Non-Standard Port Non-Application Layer Protocol
Remote Access (Geographic)	0 days	NetFlow	Initial Access	Valid Accounts

Alert	History	Telemetry	MITRE ATT&CK Tactics	MITRE ATT&CK Techniques
<i>This alert is now known as Country Watchlist: Remote Access</i>		North-South		
Remote Task Creation via ATSVN Named Pipe	0 days	Zeek Logs	Persistence Lateral Movement	At
Repeated Umbrella Sinkhole Communications	0 days	NetFlow North-South	Command and Control	Application Layer Protocol
Repeated Watchlist Communications	0 days	NetFlow North-South ETA Firewall Passive DNS		
Role Violation	0 days	NetFlow		
SMB Connection Spike <i>This alert is now known as SMB Connection Spike: Internal</i>	9 days	NetFlow North-South	Discovery	Active Scanning Network Denial of Service
SMB Connection Spike: Internal <i>Formerly known as SMB Connection Spike</i>	9 days	NetFlow North-South	Discovery	Active Scanning Network Denial of Service
SMB Connection Spike: Outbound <i>Formerly known as Outbound SMB Connection Spike</i>	0 days	NetFlow North-South	Impact Reconnaissance	Active Scanning Automated Exfiltration Network Denial of Service
Suspected Botnet Interaction	0 days	NetFlow ETA	Command and Control	Application Layer Protocol Non-Application

Alert	History	Telemetry	MITRE ATT&CK Tactics	MITRE ATT&CK Techniques
		Firewall North-South Passive DNS		Layer Protocol
Suspected Cryptocurrency Activity	0 days	NetFlow ETA Firewall North-South Passive DNS	Impact	Resource Hijacking
Suspected Port Abuse: External <i>Formerly known as Suspected Port Abuse (External)</i>	1 day	NetFlow North-South	Reconnaissance	Active Scanning
Suspected Remote Access Tool Heartbeat	0 days	NetFlow North-South Passive DNS	Command and Control	Application Layer Protocol
Suspected Zerologon RPC Exploit Attempt	0 days	NetFlow North-South Passive DNS	Privilege Escalation	Exploitation for Privilege Escalation
Suspicious Curl Behavior	0 days	Cisco NVM	Execution	Exploitation for Client Execution
Suspicious MSHTA Execution	0 days	Cisco NVM	Command and Control Defense Evasion	Ingress Tool Transfer Mshta
Suspicious SMB Activity	14 days	NetFlow North-South	Lateral Movement	Remote Services System Network Connections Discovery

Alert	History	Telemetry	MITRE ATT&CK Tactics	MITRE ATT&CK Techniques
				Network Service Discovery
Suspicious PsExec Execution	0 days	Zeek Logs	Lateral Movement	SMB/Windows Admin Shares
Suspicious Process Path	0 days	Cisco NVM	Execution Defense Evasion	Masquerading
Talos Intelligence Watchlist Hits	0 days	NetFlow ETA Firewall North-South Passive DNS		
Unusual DNS Connection	1 day	NetFlow North-South Passive DNS	Command and Control	DNS
Unusual External Server	14 days	NetFlow ETA Firewall North-South Passive DNS		
Worm Propagation	9 days	NetFlow East-West	Lateral Movement	Exploitation of Remote Services Valid Accounts

Related Topics

[Analytics: Alert Descriptions](#)

[Analytics: Alert Prerequisite Chart](#)

[Analytics: Observation Descriptions](#)

[Analytics: Overview](#)

Analytics: Alert Descriptions

Listed below are the alert types that Secure Network Analytics can generate. This list contains both published and unpublished alerts.

Unpublished alerts are newer alerts that are still being evaluated for their optimal level of efficacy. They are Off by default. Unpublished alerts work just as other alerts do (for instance, you can snooze and close them). Additionally, changing the priority of the alert does not affect whether or not an alert is published or unpublished. However, they will behave inaccurately in single node deployments.

The unpublished alerts are as follows (they are also denoted by an asterisk [*] in the alert list below).

- Amplification Attack
- Anomalous Domain Controller Activity
- Country Set Deviation
- Country Watchlist Protocol Violation
- Emergent Profile
- Excessive Connections to Network Printers
- Heartbeat Connection Count
- High Bandwidth Unidirectional Traffic
- NetBIOS Connection Spike
- Network Population Spike
- Network Printer with Excessive Connections
- New IP Scanner
- New SNMP Sweep
- New Unusual DNS Resolver
- Non-Service Port Scanner
- Persistent Remote Control Connections
- Potential Data Exfiltration
- Potential Database Exfiltration
- Repeated Umbrella Sinkhole Communications
- SMB Connection Spike: Outbound

- Suspected Remote Access Tool Heartbeat
- Unusual DNS Connection
- Worm Propagation

Amplification Attack *

Description: This entity sent traffic with a profile that suggests participation in an amplification attack. An amplification attack attempts to overwhelm a server with a massive amount of packets in response to a request, usually involving spoofed IP addresses to allow multiple entities to send traffic in response to a request. Participation in an amplification attack may indicate that an entity has been infected with botnet malware, and is sending these packets unintentionally.

Prerequisites: This alert requires 0 days of history.

Associated Observations: Traffic Amplification Observations

Next Steps: Reference the entity information in the alert and supporting observations, and determine whether or not an external entity is responsible for spreading malware. If so, update your firewall rules to block traffic from the external entity, and any other entities if it is a distributed denial of service (DDoS) attack.

If the entity sending the amplification attack is internal to your network, quarantine the entity from your network, and any other entities if it is a DDoS attack. Examine the entities for, and remove, malware.

Anomalous Domain Controller Activity *

 This alert was formerly known as Exceptional Domain Controller.

Description: A device identified as a domain controller deviated from its usual behavior. This may indicate the device is compromised.

Prerequisite: This alert requires 7 days of history to establish an entity's normal activity level.

Associated Observations: New External Server Observations

Next Steps: From the supporting evidence, view the entity's traffic profile and connections with other entities to determine what types of traffic it is sending, and if it is malicious in nature. Determine if data has been exfiltrated from your network, and if so, the types of data, and how best to remediate the situation.

Anomalous Windows Workstation

Description: A Windows workstation used a new anomalous behavioral profile (e.g., the host connected to many entities over BitTorrent). This alert uses the Anomalous Profile observation and may be an indication of malware or misuse.

Prerequisites: This alert requires 14 days of history to establish an entity's normal activity level .

Associated Observations: Anomalous Profile Observations

Next Steps: Reference the supporting observations to determine the entity's roles and determine whether or not there is a legitimate business reason for the anomalous behavior. For example, if an entity used BitTorrent to connect to other entities, it may be a test entity or some type of possible testing of firewall rules or other security tests. If there is not a legitimate reason for the anomalous behavior, examine the entity and determine whether or not the entity is functioning as intended, and if it is free of malware.

Country Set Deviation *

Description: This entity has significantly deviated from the set of countries it usually communicates with. This alert requires 36 days of history.

Prerequisites: Device tried to communicate with a host in a user-supplied watchlisted country on an illegal protocol / port combination (e.g., UDP on port 22). This may indicate the presence of a malicious covert communications channel.

Associated Observations: Country Set Deviation Observations

Next Steps: Reference the supporting observations to find the entities to which the entity has established connections, and their geolocation. Determine why it established these connections, and remediate the issue if it was due to malicious behavior. Update your Country Watchlist as necessary to include any countries involved with malicious behavior.

Country Watchlist: Protocol Violation *



This alert was formerly known as Protocol Violation (Geographic).

Description: Device tried to communicate with a host in a user-supplied watchlisted country on an illegal protocol / port combination (e.g., UDP on port 22). This may indicate the presence of a malicious covert communications channel. You must configure the Country Watchlist with at least one country.

Prerequisite: This alert requires 0 days of history.

Associated Observations: Bad Protocol Observations

Next Steps: Reference the supporting evidence to determine why the entity used the unusual protocol/port combination to communicate with the entity in the watchlisted country. Determine what was transferred in the communication. If deemed malicious, update your firewall and blocklist rules to prevent further access with this protocol/port combination, and with this geolocation, unless there is a business reason for allowing it.

To Configure a Watchlist via an API

Currently there is no ability to configure watchlists or blocklists in the Manager. Therefore, you need to make an API call to the backend (onsite) as explained below. For example, to configure a watchlist for a country ("CN"- China in this case), the following call can be made on the Manager, which would add this country to a watchlist.



For both of the following commands, the beginning character of each line after the first line immediately follows the last character in the preceding line.

To get the API key from here:

```
cat /lancope/var/services/detections/config/api_key
```

Make a note of the key to pass into next commands.

1. Enter this command:

```
curl -X POST -d '{"identifier":"US", "list_on": "watchlist"}' -v
-H 'Content-Type: application/json' -H 'Authorization: ApiKey _
customer_01-api[key]'
http://0.0.0.0:8086/api/v3/watchlist/listedcountry/
```

2. To confirm, enter this command:

```
curl -X GET -v -H 'Content-Type: application/json' -H
'Authorization: ApiKey _customer_01-api:[key]'
http://0.0.0.0:8086/api/v3/watchlist/listedcountry/
```

Country Watchlist: Remote Access



This alert was formerly known as Remote Access (Geographic).

Description: This device has been accessed from a remote host in a user-supplied watchlisted country. This alert uses the Remote Access observation and may indicate a device is compromised.

Prerequisite: This alert requires 0 days of history. This alert requires configuring the Country Watchlist with at least one country.

Associated Observations: Remote Access Observations

Next Steps: Device has been accessed from a remote host in a user-supplied watchlisted country. This may indicate a device is compromised. Reference the supporting observations to identify the external entity, and how the external entity interacted with your internal entity. Determine if the behavior was malicious, and if any data was exfiltrated, as well as what actions were taken on the internal entity. If needed, add additional firewall or security group rules to prevent future access.

To Configure a Watchlist via an API

Currently there is no ability to configure watchlists or blocklists in the Manager. Therefore, you need to make an API call to the backend (onsite) as explained below. For example, to configure a watchlist for a country ("CN"- China in this case), the following call can be made on the Manager, which would add this country to a watchlist.



For both of the following commands, the beginning character of each line after the first line immediately follows the last character in the preceding line.

To get the API key from here:

```
cat /lancope/var/services/detections/config/api_key
```

Make a note of the key to pass into next commands.

1. Enter this command:

```
curl -X POST -d '{"identifier":"US", "list_on": "watchlist"}' -v
-H 'Content-Type: application/json' -H 'Authorization: ApiKey _
customer_01-api[key]'
http://0.0.0.0:8086/api/v3/watchlist/listedcountry/
```

2. To confirm, enter this command:

```
curl -X GET -v -H 'Content-Type: application/json' -H
'Authorization: ApiKey _customer_01-api:[key]'
http://0.0.0.0:8086/api/v3/watchlist/listedcountry/
```

DNS Traffic to Tor Proxy

Description: A device sent DNS query traffic for a known Tor proxy. This may indicate that an application is preparing to establish a connection via a Tor proxy. It could be a botnet attempting to contact other devices for command-and-control. Adversaries are known to leverage it for command-and-control and defense evasion. Even if utilized by a legitimate user, it can circumvent some security controls.

Prerequisites: This alert requires 0 days of history .

Associated Observations: DNS Traffic to Tor Proxy Observations

Next Steps: Review the traffic and supporting observations for the alert. Identify the devices involved is a Tor proxy involved. The endpoint could be infected with a botnet and attempting malicious command-and-control activity; or it could be a user using a Tor proxy.

Emergent Profile *

Description: A highly sensitive entity has traffic that fits a new profile. For example, an entity that starts accepting FTP connections may be exposing sensitive data.

Prerequisite: This alert requires 14 days of history to establish entity models and determine expected traffic profiles.

Associated Observations: New Profile Observations

Next Steps: Reference the entity's new traffic profile in the supporting observations, and whether it is expected, especially in light of the previous profile or role. For example, if an entity has been repurposed from an FTP server to a mail server, this shift in behavior is expected. If it is not expected, investigate why the entity's traffic has changed, and if it is malicious.

Empire Command and Control

Description: An entity has established new periodic connections that appear to be part of an Empire PowerShell Command and Control channel. This alert uses the Heartbeat observation and may indicate the device is compromised. This alert requires 1 day of history.

Prerequisite: This alert requires 1 day of history to establish entity models and determine expected traffic profiles.

Associated Observations: Heartbeat Observations

Next Steps: Review the entity's traffic in the supporting observations, identify the entity to which it is establishing the heartbeat connections, and determine if the traffic is anticipated or malicious. If malicious, determine if other entities on your network are similarly affected. Quarantine the entities and remove any malware. Update your block list and firewall rules to disallow the command and control servers' access to your network.

Excessive Access Attempts (External)

Description: This entity has many failed access attempts from an external entity. For example, a remote entity trying repeatedly to access an internal server using SSH or Telnet would trigger this alert.

Prerequisite: This alert requires 0 days of history.

Associated Observations: Multiple Access Failures Observations, Remote Access Observations

Next Steps: Reference the supporting observations and ensure that the external entity is abnormal and unexpected. If it is normal and expected, determine why a user or machine login keeps failing to login, such as if credentials changed, but the user or machine was not given the updated credentials. If the external entity is unknown, update your firewall or security group rules to limit access for the remote control protocol. Update your block list and firewall rules to disallow this entity's access to your network if the entity is potentially malicious.

Excessive Connections to Network Printers *

Description: This entity initiates too many connections to network printers. This behavior may indicate a denial-of-service attack, or an attempt to exfiltrate data by printing documents.

Prerequisite: This alert requires 0 days of history.

Associated Observations: Excessive Connections to Network Printers Observations

Next Steps: Reference the supporting observations and determine how the entity is communicating with the network printers. Quarantine the entity and remove malware if the communications are malicious. Examine the printer job queues to determine what actions they are performing. Clear the queues if the printer is tasked to print confidential documents. Disconnect the printers' internet access if they are tasked to transmit confidential information to external entities. Remove any malware from the printers as necessary.

Geographically Unusual Remote Access

Description: This entity has been accessed from a remote host in a country that doesn't normally access the local network. For example, a local server accepting an SSH connection from a foreign source would trigger this alert. Remote access from an unusual geolocation could be an indication of malicious access.

Prerequisite: This alert requires 30 days of history to establish sufficient traffic history, and determine normal traffic based on geolocation.

Associated Observations: Remote Access Observations

Next Steps: Reference the supporting observations and determine what action the entity took, and why it took the action. If the entity is expected, but is accessing the internet from another country than expected, update your firewall settings to allow this traffic.

Remediate the action, and update your blocklist and firewall rules to disallow the entity from accessing your network if this is malicious behavior.

Heartbeat Connection Count *

Description: This entity has established new periodic connections with many remote entities, which might indicate unauthorized P2P traffic or botnet activity.

Prerequisite: This alert requires 1 day of history to establish traffic models.

Associated Observations: Heartbeat Observations

Next Steps: Reference the supporting observations and determine the entities to which the affected entity is establishing the heartbeat connections, and confirm that they are not expected. Understand the purpose for the periodic connections, and update your firewall and blocklist rules to prevent further access.

High Bandwidth Unidirectional Traffic *

Description: This entity started sending large amounts of data to new remote hosts. This can indicate misuse or misconfiguration. For example, malware might cause an infected host to attack a website by directing a host to send lots of data to a vulnerable service.

Prerequisite: This alert requires 0 days of history.

Associated Observations: New High Throughput Connection Observations

Next Steps: Reference the supporting observations for flow details, and determine why the entity is sending large amounts of traffic. If the traffic is not permissible, investigate what software on the host is responsible for the malicious traffic.

Inbound Port Scanner

Description: This entity was port scanned by an external entity. If an external entity is scanning entities internal to your network, it may be scanning for unpatched vulnerabilities or other ways to infiltrate entities on your network.

Prerequisite: This alert requires 1 day of history to establish entity models and determine normal behavior.

Associated Observations: External Port Scanner Observations

Next Steps: Reference the supporting observations to identify the external entity that port scanned your internal entity. Determine if it is the result of planned penetration testing or other intended behavior, or if it is malicious. Update your IP scanner and allow list rules to allow the traffic if it is intended. Block the traffic if it is not intended. Update your firewall rules as necessary, including port access.

Internal Connection Spike

Description: This entity had a sudden increase in internal connections, which is suggestive of scanning activity.

Prerequisite: This alert requires 0 days of history.

Associated Observations: Record Metric Outlier Observations

Next Steps: Reference the supporting observations to determine why the entity is establishing multiple connections. Determine if it is performing scanning activity because of penetration testing or another allowed purpose, or if it is malicious behavior. Remediate the behavior as necessary.

Internal Port Scanner

Description: This entity has started a port scan on an entity internal to your network. If an internal entity is scanning entities internal to your network, it may be a penetration test by your network security team, or it may be malicious behavior from an entity on your network.

Prerequisite: This alert requires 7 days of history to establish entity models and normal entity behavior.

Associated Observations: Internal Port Scanner Observations

Next Steps: Reference the supporting observations to understand the type of scanning activity. Scanning activity is often associated with a compromised host that is searching for data or other hosts to infect. To gain more context, search for observations related to the entity that the system logged around the same time (such as watchlist interactions). This may provide additional information about the behavior.

LDAP Connection Spike

Description: This entity attempted to contact an unusually large number of internal LDAP servers. This alert may be an indication of malware or abuse.

Prerequisite: This alert requires 9 days of history to establish normal behavior.

Associated Observations: IP Scanner Observations

Next Steps: Reference the supporting observations and determine why the entity is establishing connections with multiple LDAP servers, what types of actions the entity is taking, and if this is malicious behavior. If data was exfiltrated, follow your organization's guidelines for dealing with data exfiltration. Quarantine the entity as necessary to remove malware.

Meterpreter Command and Control Success

Description: This Device has established new periodic connections that appear to be part of a Meterpreter Command and Control channel. This alert uses the Heartbeat observation and may indicate that the device is compromised.

Prerequisites: This alert requires 1 day of history.

Associated Observations: Heartbeat Observations

Next Steps: Review the entity's traffic in the supporting observations, identify the entity for which it is establishing the heartbeat connections, and determine if the traffic is anticipated or malicious. If malicious, determine if other entities on your network are similarly affected. Quarantine the entities and remove any malware. Update your block list and firewall rules to disallow the command and control servers' access to your network.

NetBIOS Connection Spike *

Description: Source attempted to contact large number of hosts using NetBIOS. This can be an indication of malware or abuse.

Prerequisite: This alert requires 9 days of history to establish entity traffic models and determine normal traffic behavior.

Associated Observations: IP Scanner Observations

Next Steps: Reference the supporting observations to determine the host and analyze the traffic flow details. NetBIOS is not a commonly used protocol, so any connection spike events would likely be malicious. If detected, review what applications are using NetBIOS and if that traffic is legitimate. If so, snooze this alert for the host.

Network Population Spike *

Description: A record number of IP addresses were observed communicating on the network. This might indicate spoofing of source addresses or scanning activity.

Prerequisite: This alert requires 36 days of history to establish a sufficient amount of days to count the total number of entities communicating on the network.

Associated Observations: Population Spike Observations

Next Steps: Reference the supporting observations associated with the alert and determine if the IP addresses are legitimate entities. If they are not legitimate, locate the source of the spoofed addresses, and remediate as necessary.

Network Printer with Excessive Connections *

Description: This printer initiates too many connections. This may indicate malicious behavior, such as botnet malware infection.

Prerequisite: This alert requires 0 days of history.

Associated Observations: Network Printer with Excessive Connections Observations

Next Steps: Review the established connections, and the entities that established connections with the printer. Reference the supporting observations to see what type of connections were established by the printer. If the connections indicate the printer is compromised, quarantine the printer and consider removing and re-installing the operating systems.

New Internal Device

Description: A new entity has appeared on a restricted subnet range after not being seen in the lookback period.

Prerequisite: This alert requires 21 days of history to learn which entities are normally seen on the network. This alert also requires selecting **New Internal Device** on the Subnet Configuration page.

Associated Observations: New Internal Device Observations

Next Steps: Reference the supporting observations to determine if this entity is an expected entity, and is merely new to your network. If the entity is expected and not malicious, close the alert; future new entities will continue to generate alerts. If the entity is suspicious, determine the MAC address by accessing the local switch.

New IP Scanner *

Description: This entity started scanning the local IP network. This could indicate, for example, reconnaissance by an attacker.

Prerequisite: This alert requires 9 days of history to establish entity traffic models and determine normal traffic behavior.

Associated Observations: IP Scanner Observations

Next Steps: Reference the supporting observations and investigate why the external entity is scanning the network. Determine if it is the result of penetration testing or other intended behavior, or if it is malicious. Update your IP scanner and firewall rules to allow the traffic if it is intended. If potentially malicious, search for associated observations for

the entity or user who owns the machine to determine what software caused the scanning activity.

New Remote Access

Description: This entity has been accessed (e.g., via SSH) from a remote host for the first time in recent history. This remote access may indicate malicious behavior, especially if the entity is not expected to accept connections from external entities.

Prerequisite: This alert requires 36 days of history to establish sufficient traffic history and entity models.

Associated Observations: Remote Access Observations

Next Steps: Reference the supporting observations to determine why the entity is being accessed by the external entity, and if it is a legitimate form of access. Also determine (based on the observations) if there were multiple access attempts to the source entity prior to this access, whether from this external entity or another external entity. Update your firewall and blocklist rules based on this information.

New SNMP Sweep *

Description: This entity attempted to reach a large number of hosts using SNMP. This can be an indication of network reconnaissance cause by malicious software. An SNMP sweep, when performed by a malicious actor, could result in gathering information about your network, or malicious entity configuration updates.

Prerequisite: This alert requires 9 days of history to establish entity traffic models and determine normal traffic behavior.

Associated Observations: IP Scanner Observations

Next Steps: Reference the supporting observations to determine if the entity is intended to track network entities over SNMP, and if this behavior is malicious. If the activity is not part of planned penetration testing or otherwise intended behavior, quarantine the entity and remediate the issue. Determine if any of the entities have been affected, such as updated configuration or compromised security settings, and remediate any issues. If the entity is expected to perform SNMP sweeps, add the entity to the Scanner allowed list.

New Unusual DNS Resolver *

Description: This entity contacted a DNS resolver that it doesn't normally use. This can indicate misconfiguration or the presence of malware. For example, an attacker could cause a DNS resolver to redirect a popular website to a domain that serves additional

malware.

Prerequisite: This alert requires 7 days of history to establish entity roles and model normal traffic.

Associated Observations: Unusual DNS Resolver Observations

Next Steps: Verify the entity's configuration to ensure that it is configured with the proper DNS settings. If so, determine what software is making the DNS lookup. Block the external IP address if the traffic is deemed malicious.

Non-Service Port Scanner *

Description: This device started scanning the local network on a port not associated with a common service. This alert uses the IP Scanner observation and may indicate that an attacker is inside the network, scanning for vulnerabilities.

Prerequisite: This alert requires 9 days of history to establish entity roles and model normal traffic

Associated Observations: IP Scanner Observation

Next Steps: Reference the supporting observations and investigate why the external entity is scanning the network. Determine if it is the result of penetration testing or other intended behavior, or if it is malicious. Update your IP scanner and firewall rules to allow the traffic if it is intended. If potentially malicious, search for associated observations for the entity or user who owns the machine to determine what software caused the scanning activity.

Outbound LDAP Connection Spike

Description: This entity is communicating with a large number of external hosts using an LDAP port. This alert may indicate a possible infected host or an internally-initiated port scan.

Prerequisite: This alert requires 0 days of history.

Associated Observations: IP Scanner Observations

Next Steps: Reference the supporting observations and determine to which entities the source entity is sending traffic, what type of traffic, and if this is an update to the entity's roles or responsibilities, or if this is unintended. If this is unintended, remediate the issue. Update your firewall and blocklist rules to prevent this access.

Persistent Remote Control Connections *

Description: This entity is receiving persistent connections from a new host on a remote control protocol like Remote Desktop or SSH. This may indicate that a firewall rule or ACL is overly permissive.

Prerequisite: This alert requires 7 days of history to establish traffic models and determine normal traffic behavior.

Associated Observations: New External Server Observations, Persistent External Server Observations

Next Steps: Adjust firewall or security group rules to prevent malicious attempts to repeatedly access the entity. Confirm that the local entity has not been breached by checking Remote Access Observations or the entity.

PetitPotam Attack Via EFS RPC Calls

Description: PetitPotam Attack Via EFS RPC Calls

A device sent a Remote Procedure Call (RPC) using the Encrypting File System Remote Protocol (EFSRPC) Protocol library. The PetitPotam attack is known to be related to this type of RPC traffic. PetitPotam is a tool that can exploit this library. It is also known as an NTLM relay attack. Since most organizations don't use this library at all, or limit the usage of it, any use is uncommon enough to indicate a possible PetitPotam attack.

Prerequisite: This alert requires 7 days of history to establish traffic models and determine normal traffic behavior.

Associated Observations: PetitPotam Attack Via EFS RPC Calls Observations

Next Steps: Review the traffic and supporting observations for the alert. Identify the devices involved in the traffic. Confirm how and when this library is routinely used in your organization. If a domain controller is involved in the related traffic, it might indicate an NTLM relay attack is under way and attempting to take over a domain controller. Validate that the external source is not a known originator of the PetitPotam attack. Determine whether any internal devices are involved in the traffic as a source because such an attack can be propagated using an internal device.

Possible Impacket SecretDump Activity

Description: A device is attempting a secrets dump using an impact tool such as secretdump.py, which allows dumping credentials from an Active Directory (AD) server. This is also referred to as a secrets-dump HKTL.

Prerequisite: This alert requires 0 days of history.

Associated Observations: Possible Impacket SecretDump Activity Observations

Next Steps: Review the traffic and supporting observations. Identify the devices involved in the traffic. If any of the devices involved seem questionable for AD traffic, it could be an attempt to exfiltrate secrets or credentials. Determine whether the devices involved are allowed to perform an AD credentials dump since it could be a valid credentials transfer.

Potential Data Exfiltration *

Description: This entity downloaded a sizeable chunk of data from an internal entity that it doesn't communicate with regularly. Shortly after that, the entity uploaded a similar amount of data to an external entity. This may indicate an unauthorized transfer of information, or other malicious behavior.

Prerequisite: This alert requires 0 days of history.

Associated Observations: Potential Data Forwarding Observations

Next Steps: Reference the supporting observation to determine the volume of traffic and the client entity to determine if the behavior is expected in the normal course of business, such as a new scheduled backup. If it is malicious, determine what was transferred. Follow your organization's guidelines on data exfiltration.

Potential Database Exfiltration *

Description: A statistically unusual amount of data was transferred from a database server to a client. This may indicate an unauthorized transfer of information, or other malicious behavior.

Prerequisite: This alert requires 7 days of history to establish which entities normally serve as databases, and what their normal traffic profiles are.

Associated Observations: New High Throughput Connection Observations

Next Steps: Examine the client entity to determine if the behavior is expected in the normal course of business, such as a new scheduled backup. If it is malicious, determine what was transferred. Follow your organization's guidelines on data exfiltration.

Potential Gamaredon C2 Callout

Description: A command line utility was used to contact a URL associated with the command-and-control servers of a threat actor known as Gamaredon. Gamaredon (also known as Armageddon, Primitive Bear, and ACTINIUM) is an APT active since 2013 known to leverage spearphishing to infect victims with custom malware.

Prerequisite: This alert requires 0 days of history.

Associated Observations: Potential Gamaredon C2 Callout Observations

Next Steps: Determine if this activity was legitimate, and if not, investigate and potentially quarantine the device that made the connection.

Protocol Forgery

Description: This entity was observed running a potentially restricted service (such as SSH) on a non-standard port. This can indicate evasion of security controls.

Prerequisite: This alert requires 1 day of history to establish entity models and see which entities use potentially restricted services. This alert also requires Encrypted Traffic Analytics capabilities.

Associated Observations: Insecure Transport Protocol Observations

Next Steps: Reference the supporting observations to determine why the entity used the unusual protocol/port combination to communicate. Update your firewall and blocklist rules to prevent further access with this protocol/port combination, if deemed a security risk.

Remote Task Creation via ATSVK Named Pipe

Description: A device is attempting to create a remote task using ATSVK named pipes, which could be a malicious attempt to use at.exe for performing task scheduling for initial or recurring execution of malicious code. The at.exe utility has been deprecated in current versions of Windows in favor of schtasks.

Prerequisite: This alert requires 0 days of history.

Associated Observations: Remote Task Creation via ATSVK Named Pipe Observations

Next Steps: Review the traffic and supporting observations. Identify the devices involved in the traffic. Verify that you no longer allow at.exe via older Windows deployments. It is possible older Windows versions or test lab environments continue to allow use of at.exe. If there aren't any older Windows deployments, then this is likely a suspicious use of at.exe.

Repeated Umbrella Sinkhole Communications *

Description: This entity has established periodic connections with a Cisco Umbrella Sinkhole. This alerts uses the Umbrella Sinkhole Hit and Heartbeat observations and may indicate a device is compromised.

Prerequisite: This alert requires 0 days of history.

Associated Observations: Heartbeat Observations, Umbrella Sinkhole Hit Observations

Next Steps: Reference the supporting observations and examine the affected entity and log information. Determine why the entity is establishing periodic communications to the sinkhole and remediate the situation.

Repeated Watchlist Communications

Description: This entity has established periodic connections with a watchlisted IP. This may indicate the presence of malware, or a compromised entity on your network.

Prerequisite: This alert requires 0 days of history.

Associated Observations: Watchlist Interaction Observations, Heartbeat Observations

Next Steps: Reference the supporting observations and examine the affected entity and log information. Determine why the entity is establishing periodic communications, and remediate the situation. As necessary, contact the organization that maintains a given watchlist, either for advice to remediate the situation, or to verify that the entity is no longer infected with malware.

Role Violation

Description: This entity is identified with a particular role (e.g., User entity), but was observed acting in an atypical manner for that role (e.g., SSH server). If an entity changes roles, it may be an indication of malicious behavior, such as malware changing how an entity functions.

Prerequisite: This alert requires 0 days of history.

Associated Observations: Role Violation Observations

Next Steps: Reference the supporting observations and determine whether the new role behavior is intended and part of the normal course of business. If it is not, quarantine the entity. If it is intended, snooze the alert.

SMB Connection Spike: Internal



This alert was formerly known as SMB Connection Spike.

Description: Device attempted to contact an unusually large number of SMB servers. This can be an indication of malware or abuse.

Prerequisite: This alert requires 9 days of history.

Associated Observations: IP Scanner Observations

Next Steps: Reference the supporting evidence and determine why the entity is establishing connections with multiple SMB servers, what types of actions the entity is

taking, and if this is malicious behavior. If data was exfiltrated, follow your organization's guidelines for dealing with data exfiltration. Quarantine the entity as necessary to remove malware.

SMB Connection Spike: Outbound*

 This alert was formerly known as Outbound SMB Connection Spike.

Description: This entity attempted to contact an unusually large number of SMB servers. This can be an indication of malware or abuse. As SMB is used primarily for file sharing, but can also be used for accessing network printers or browsing other hosts on a network, this could indicate data exfiltration or network resource misuse.

Prerequisite: This alert requires 0 days of history to establish entity traffic models and determine normal traffic behavior.

Associated Observations: IP Scanner Observations

Next Steps: Reference the supporting observations and determine why the entity is establishing connections with multiple SMB servers, what types of actions the entity is taking, and if this is malicious behavior. If data was exfiltrated, follow your organization's guidelines for dealing with data exfiltration. Quarantine the entity as necessary to remove malware.

Suspected Botnet Interaction

Description: This entity exchanged traffic with IP addresses associated with botnets, or attempted to resolve domain names associated with botnets.

Prerequisite: This alert requires 1 day of history to establish entity models.

Associated Observations: Watchlist Interaction Observations

Next Steps: Quarantine the entity, and remove all malware. Update your block list and firewall rules to disallow the botnet entities from accessing your network. Reference the supporting observations and determine if any other entities on your network are also infected, based on communications that the entity may have established, and remediate as necessary.

Suspected Cryptocurrency Activity

Description: Source exchanged a significant amount of traffic with multiple addresses known to be operating cryptocurrency nodes, based on Talos intelligence, and other sources. This behavior may indicate that an entity is being used to mine cryptocurrency.

Prerequisite: This alert requires 0 days of history.

Associated Observations: Watchlist Interaction Observations

Next Steps: Quarantine the entity, and remove all cryptocurrency mining software, whether it is malware or installed by a user.

Suspected Port Abuse: External

 This alert was formerly known as Suspected Port Abuse (External).

Description: This entity is communicating with an external host on unusual range of ports. This can indicate externally-initiated abuse (e.g., a spoof attack) or an internally-initiated port scan.

Prerequisite: This alert requires 1 day of history to establish entity models.

Associated Observations: Port Scanner Observations, External Port Scanner Observations

Next Steps: Reference the supporting observations to review the entity's activity, and determine if it is consistent with planned penetration testing, or malicious behavior. Determine the origin of the malicious behavior, and remediate the issue. Update your firewall and blocklist rules as needed.

Suspected Remote Access Tool Heartbeat *

Description: Traffic with a signature matching Remote Access Tools (e.g., RevengeRAT) was seen on this device. This alert uses the Suspicious Network Activity observation and may indicate that the device is compromised.

Prerequisites: This alert requires 0 days of history.

Associated Observations: Suspicious Network Activity Observations

Next Steps: Ensure this device has the most recent security updates applied and investigate the device for signs of compromise.

Suspected Zerologon RPC Exploit Attempt

Description: Traffic with a signature matching the Zerologon RPC exploit was seen on this device. This alert uses the Suspicious Network Activity observation and may indicate the device is being targeted for exploitation.

Prerequisite: This alert requires 0 days of history.

Associated Observations: Suspicious Network Activity Observations

Next Steps: Ensure this device has the most recent security updates applied. Follow mitigation steps in reference to CVE-2020-1472.

Suspicious Curl Behavior

Description: The system utility curl exhibited suspicious behavior that may be indicative of exploitation of CVE-2023-38545.

Prerequisite: This alert requires 0 days of history.

Associated Observations: Suspicious Curl Behavior Observations

Next Steps: Isolate the endpoint, investigate recent utilization of the curl process and ensure curl is updated to version 8.4 or newer on all devices.

Suspicious MSHTA Activity

Description: The built-in Windows application MSHTA.exe was executed interactively by a non-system user and utilized to make a network connection. While typically legitimate when run automatically by the system, it is also known to be utilized by threat actors including Advanced Persistent Threats (APTs).

Prerequisite: This alert requires 0 days of history.

Associated Observations: Suspicious MSHTA Activity Observations

Next Steps: Isolate the endpoint and investigate if executables were downloaded in non-standard directories and executed.

Suspicious Process Path

Description A process was executed on an endpoint from a directory that should not have executables.

Prerequisite: This alert requires 0 days of history.

Associated Observations: Suspicious Process Path Observations

Next Steps: Isolate the endpoint and investigate if executables were downloaded in non-standard directories and executed.

Suspicious PsExec Execution

Description: A device other than a Windows Sysinternal device is using psexec with a renamed service name, which could indicate a threat actor attempting to perform a remote execution.

Prerequisite: This alert requires 0 days of history.

Associated Observations: Suspicious PsExec Execution Observations

Next Steps: Review the traffic and supporting observations. Identify the devices involved in the traffic. If the devices aren't Windows Sysinternal devices, an external or infected internal device is likely exploiting a service.

Suspicious SMB Activity

Description: Multiple new SMB servers have communicated with common SMB peers. This can be an indication of malware or abuse.

Prerequisite: This alert requires 14 days of history.

Associated Observations: Suspicious SMB Activity Observations

Next Steps: Reference the supporting observations to examine the entity's traffic profile to determine if there is further evidence of botnet activity or other malicious behavior. Check for other entities on your network which may be exhibiting similar behavior, and remediate it.

Talos Intelligence Watchlist Hits

Description: This entity exchanged a significant amount of traffic with multiple addresses on the Cisco Talos IP Blocklist.

Prerequisite: This alert requires 0 days of history.

Associated Observations: Watchlist Interaction Observations

Next Steps: Quarantine the entity and remove all malware. Investigate the external IP address by selecting **Talos Intelligence** from the menu to see what the traffic indicates and take appropriate remediation actions.

Unusual DNS Connection *

Description: This entity contacted an unusual DNS resolver and then established periodic connections with a remote entity. This behavior may indicate a malicious redirect of traffic, or a malware infection on an entity.

Prerequisite: This alert requires 1 day of history to establish entity models.

Associated Observations: Unusual DNS Resolver Observations, Heartbeat Observations

Next Steps: Reference the supporting observations and determine if this behavior is malicious, and remove malware if it is present. Update your block list and firewall rules to disallow access.

Unusual External Server

Description: This entity has repeatedly communicated with a new external server with suspicious traffic profiles. This could indicate, for example, a new piece of software that is acting as a server to an external entity, such as syslog or TeamViewer.

Prerequisite: This alert requires 14 days of history to establish normal traffic patterns, and determine expected external entity traffic.

Associated Observations: New External Server Observations, Persistent External Server Observations, Watchlist Lookup Observations, Watchlist Interaction Observations

Next Steps: Reference the supporting observations to examine the entity's traffic profile to determine the nature of the traffic and if it is permitted. Quarantine the entity and remove offending software. Determine if other entities on your network exhibit similar behavior, and remediate that behavior.

Worm Propagation *

Description: A previously scanned device started scanning the local IP network. This alert uses the Worm Propagation observation and may indicate that a worm is propagating itself inside the network.

Prerequisites: This alert requires 9 days of history.

Associated Observations: Worm Propagation Observations

Next Steps: Reference the supporting observations and investigate why the internal entities are scanning the network. Determine if it is the result of penetration testing, other intended behavior, or if it is malicious. Update your IP scanner and firewall rules to allow the traffic if it is intended. If potentially malicious, search for associated observations for the entity or user who owns the machine to determine which software caused the scanning activity.

Related Topics

[Analytics: Alert Prerequisite Chart](#)

[Analytics: Observation Descriptions](#)

[Analytics: Alerts and Observations Overview](#)

[Analytics: Overview](#)

Analytics: Observation Descriptions

Listed below are the observation types that Secure Network Analytics can generate.

Anomalous Profile Observation

Description: An entity or entities used a profile for the first time which differs from typical behaviors seen in the network (e.g., an abnormally high number of entities using the profile for the first time, sending anomalous traffic).

Prerequisite: None.

Associated Alerts: Anomalous Windows Workstation Alerts

Bad Protocol Observation

Description: An entity used a non-standard protocol on a standard port (e.g., UDP on port 22).

Prerequisite: None.

Associated Alerts: Country Watchlist: Protocol Violation

Country Set Deviation Observation

Description: An entity communicated with a set of countries different from its usual one.

Prerequisite: None.

Associated Alerts: Country Set Deviation Alerts

DNS Traffic to Tor Proxy Observation

Description: A device sent DNS query traffic for a known Tor proxy. This may indicate that an application is preparing to establish a connection via a Tor proxy. It could be a botnet attempting to contact other devices for command-and-control. Adversaries are known to leverage it for command-and-control and defense evasion. Even if utilized by a legitimate user, it can circumvent some security controls.

Prerequisite: None.

Associated Alerts: DNS Traffic to Tor Proxy Alerts

Exceptional Domain Controller Observation

Description: Domain Controller entity communicated with unusual external ports.

Prerequisite: None.

Associated Alerts: Anomalous Domain Controller Activity Alerts

Excessive Connections to Network Printers Observation

Description: An entity initiated excessive connections to network printers.

Prerequisite: None.

Associated Alerts: Excessive Connections to Network Printers Alerts

External Port Scanner Observation

Description: An entity on the local network scanned (or was scanned by) a remote IP address.

Prerequisite: None.

Associated Alerts: Inbound Port Scanner Alerts, Suspected Port Abuse: External Alerts

Heartbeat Observation

Description: An entity maintained a heartbeat with a remote host.

Prerequisite: None.

Associated Alerts: Empire Command and Control Alert, Heartbeat Connection Count Alerts, Unusual DNS Connection Alerts, Meterpreter Command and Control Success, Repeated Umbrella Sinkhole Communications

Internal Port Scanner Observation

Description: An entity scanned a large number of ports.

Prerequisite: None.

Associated Alerts: Internal Port Scanner Alerts

IP Scanner Observation

Description: An entity scanned a large number of entities.

Prerequisite: None.

Associated Alerts: LDAP Connection Spike, NetBIOS Connection Spike Alerts, New IP Scanner Alerts, New SNMP Sweep Alerts, Outbound LDAP Connection Spike, SMB Connection Spike: Outbound Alerts, SMB Connection Spike: Internal Alerts

Multiple Access Failures Observation

Description: An entity had multiple failed application (e.g., FTP, SSH, RDP) access attempts.

Prerequisite: None.

Associated Alerts: Excessive Access Attempts (External) Alerts

Network Printer with Excessive Connections Observation

Description: Network printer initiated excessive connections to other entities.

Prerequisite: None.

Associated Alerts: Network Printer with Excessive Connections Alerts

New External Server Observation

Description: An entity started communicating with an external server.

Prerequisite: None.

Associated Alerts: Anomalous Domain Controller Activity Alerts, Persistent Remote Control Connections Alerts, Unusual External Server Alerts

New High Throughput Connection Observation

Description: An entity has exchanged a large amount of traffic with a new host.

Prerequisite: None.

Associated Alerts: High Bandwidth Unidirectional Traffic Alerts, Potential Database Exfiltration Alerts

New Internal Device Observation

Description: After not being seen in the lookback period, a new entity emerges on the network.

Prerequisite: None.

Associated Alerts: New Internal Device Alerts

New Profile Observation

Description: An entity matches a profile tag (e.g., FTP server) that it hasn't matched recently.

Prerequisite: None.

Associated Alerts: Emergent Profile Alerts

Persistent External Server Observation

Description: This entity has regularly communicated with the same external server (FTP, SSH, etc.).

Prerequisite: None.

Associated Alerts: Persistent Remote Control Connections Alerts, Unusual External Server Alerts

PetitPotam Attack Via EFS RPC Calls Observation

Description: A device sent a Remote Procedure Call (RPC) using the Encrypting File System Remote Protocol (EFSRPC) Protocol library. The PetitPotam attack is known to be related to this type of RPC traffic. PetitPotam is a tool that can exploit this library. It is also known as an NTLM relay attack. Since most organizations don't use this library at all, or limit the usage of it, any use is uncommon enough to indicate a possible PetitPotam attack.

Prerequisite: None.

Associated Alerts: PetitPotam Attack Via EFS RPC Calls Alerts

Population Spike Observation

Description: A record number of IP addresses were observed communicating on the local network.

Prerequisite: None.

Associated Alerts: Network Population Spike Alerts

Port Scanner Observation

Description: An entity scanned a large number of ports.

Prerequisite: None.

Associated Alerts: Internal Port Scanner Alerts

Possible Impacket SecretDump Activity Observation

Description: A device is attempting a secrets dump using an impact tool such as secretdump.py, which allows dumping credentials from an Active Directory (AD) server. This is also referred to as a secrets-dump HKTL.

Prerequisite: None.

Associated Alerts: Possible Impacket SecretDump Activity Alerts

Potential Data Forwarding Observation

Description: A similarly sized, and closely timed, data transfer was detected between an internal data source to this entity (the "download"), and then from this entity to an external data sink (the "upload").

Prerequisite: None.

Associated Alerts: Potential Data Exfiltration Alerts

Potential Gamaredon C2 Callout Observation

Description: A command line utility was used to contact a URL associated with the command-and-control servers of a threat actor known as Gamaredon. Gamaredon (also known as Armageddon, Primitive Bear, and ACTINIUM) is an APT active since 2013 known to leverage spearphishing to infect victims with custom malware.

Prerequisite: None.

Associated Alerts: Potential Gamaredon C2 Callout Alerts

Record Metric Outlier Observation

Description: An entity sent or received a record amount of traffic.

Prerequisite: None.

Associated Alerts: Internal Connection Spike Alerts

Remote Access Observation

Description: An entity was accessed from a remote source.

Prerequisite: None.

Associated Alerts: Geographically Unusual Remote Access Alerts, New Remote Access Alerts, Country Watchlist: Remote Access Alerts

Remote Task Creation via ATSVCS Named Pipe Observation

Description: A device is attempting to create a remote task using ATSVCS named pipes, which could be a malicious attempt to use at.exe for performing task scheduling for initial or recurring execution of malicious code. The at.exe utility has been deprecated in current versions of Windows in favor of schtasks.

Prerequisite: None.

Associated Alerts: Remote Task Creation via ATSVCS Named Pipe Alerts

Role Violation Observation

Description: An entity has new traffic that doesn't fit its role (e.g., FTP server communicating on port 80).

Prerequisite: None.

Associated Alerts: Role Violation Alerts

Suspicious Curl Behavior Observation

Description: The system utility curl exhibited suspicious behavior that may be indicative of exploitation of CVE-2023-38545.

Prerequisite: None.

Associated Alerts: Suspicious Curl Behavior Alerts

Suspicious MSHTA Activity Observation

Description: The built-in Windows application MSHTA.exe was executed interactively by a non-system user and utilized to make a network connection. While typically legitimate when run automatically by the system, it is also known to be utilized by threat actors including Advanced Persistent Threats (APTs).

Prerequisite: None.

Associated Alerts: Suspicious MSHTA Activity Alerts

Suspicious Process Path Observation

Description: A process was executed on an endpoint from a directory that shouldn't have executables.

Prerequisite: None.

Associated Alerts: Suspicious Process Path Alerts

Suspicious PsExec Execution Observation

Description: A device other than a Windows Sysinternal device is using psexec with a renamed service name, which could indicate a threat actor attempting to perform a remote execution.

Prerequisite: None.

Associated Alerts: Suspicious PsExec Execution Alerts

Suspicious Network Activity Observation

Description: Suspicious activity was detected that matches a Talos signature.

Prerequisite: None

Associated Alerts: Suspected Remote Access Tool Heartbeat Alerts

Suspicious SMB Activity Observation

Description: Multiple entities have performed anomalous activity using the SMB protocol for the first time.

Prerequisite: None.

Associated Alerts: Suspicious SMB Activity Alerts

Traffic Amplification Observation

Description: An entity's outbound and inbound traffic did not match the typical ratio associated with the profile it was using. This could indicate participation in an amplification attack. An amplification attack attempts to overwhelm a server with a massive amount of packets in response to a request, involving spoofed IP addresses or other identifying information. Participation in an amplification attack may also indicate that an entity has been infected with botnet malware, and is sending these packets unintentionally.

Prerequisite: None.

Associated Alerts: Amplification Attack Alerts

Umbrella Sinkhole Hit Observation

Description: The entity communicated with a known Cisco Umbrella sinkhole.

Prerequisite: None.

Associated Alerts: Repeated Umbrella Sinkhole Communications Alerts

Unusual DNS Resolver Observation

Description: An entity communicated with an unusual DNS resolver.

Prerequisite: None.

Associated Alerts: New Unusual DNS Resolver Alerts, Unusual DNS Connection Alerts

Watchlist Interaction Observation

Description: An entity communicated with an IP address that is on a watchlist (either explicitly or implicitly via a domain name).

Prerequisite: None.

Associated Alerts: Repeated Watchlist Communications Alerts, Suspected Botnet Interaction Alerts, Suspected Cryptocurrency Activity Alerts, Talos Intelligence Watchlist Hits Alerts, Unusual External Server Alerts

Worm Propagation Observation

Description: A previously scanned device started scanning the local IP network.

Prerequisite: None

Associated Alerts: Worm Propagation Alerts

Related Topics

[Analytics: Alerts and Observations Overview](#)

[Analytics: Alert Descriptions](#)

[Analytics: Alert Prerequisite Chart](#)

[Analytics: Overview](#)

Analytics: Alerts Summary

The Alerts Summary displays the alerts generated by the system according to your filter settings. It generates these alerts, representing potential malicious activity, based on an analysis of various information about your network, including the following:

- Monitored entities roles, and the observations logged for those entities
- Alert type priority
- IP scanner rules



Use Host Group Management to add known valid IP scanners to the default Network Scanners group (Host Group ID: 48). These devices should be listed by single comma separated IP or CIDR notation. By adding known scanners, you will prevent alerts from occurring when these devices scan external or internal hosts.

Open Alerts Summary

From the main menu, choose **Monitor > Alerts**.

OR

In the Supporting Observations section on the Alert Details page, from the drop-down list for the desired device, choose **Alerts**.

The Alerts Summary opens with a list of all the alerts related to that device.



For more information on how to use this page, see [Analytics: Investigating Alerts](#)

Filter the Alerts Table

Refer to the following information to learn how to filter for various settings.

Search - In the Search field, type the entry by which you want to filter the table. You can filter by any content that would potentially be displayed in any of the table columns, such as alert type, source type, time, etc. When finished, click **Apply** to the right of the Time Range fields.

Time Range - To filter by date and time:

- Enter the start date and time in the **Start** text box or click the calendar icon at the corner of the text box to select the start date and time.
- Enter the end date and time in the **End** text box or click the calendar icon at the

corner of the text box to select the end date and time.

- Click **Apply** . To revert, click **Reset**.

Set alerts by status – By clicking the desired option in this field, you can choose to view All alerts or only Open, Snoozed, Unpublished, or Closed alerts. By default, the table displays alerts that are assigned the All status.



When you open, re-open, or close an alert that matches a Response Management rule, all associated actions occur (email, syslog, webhook), informing you that the alert has been opened, re-opened, or closed, depending on the action you took.

- Snoozed alerts: When you close an alert (see [Edit Alerts Table Entries](#)), the Close Alert dialog opens. In this dialog you can specify if you want to snooze the alert for a certain period of time.
- Unpublished alerts: Unpublished alerts are newer alerts that are still being evaluated for their optimal level of efficacy. They are Off by default. Unpublished alerts work just as other alerts do (for instance, you can snooze and close them). Additionally, changing the priority of the alert does not affect whether or not an alert is published or unpublished. However, they will behave inaccurately in single node deployments.

To know which alerts are unpublished, see [Analytics: Alert Descriptions](#).

View the Alerts Table

The following information is displayed in the Alerts table.

Field	Description
Alert	The alert type that was generated. • To access the Alert Details page for an alert, click the alert. For information about the Alert Details page, see Analytics: Alert Details. • To see the status of an alert, hover over the (Information) icon that is beside the desired alert. Use the Filter drop-down list to filter by a specific alert type.
Source	The source entity that caused the alert to be generated. To view details for an alert, hover over the (Alert) icon or

	the  (Eye) icon that is beside the IP address. (If the source is associated with a minimum of one open alert, the Alert icon is displayed. If the source is not associated with an open alert, the Eye icon is displayed.) Click the drop-down list to access the following options: - Choose Device (Network) to access the Device Report filtered by the source. For information about the Device page, see Analytics: Device Report. - Choose Alerts to access the Alerts Summary that contains a list of all the alerts related to that source. For information about the Alerts Summary, see Analytics: Alerts Summary. - Choose Observations to access the Observations by Device page that contains a list of all observations related to the source. For information about the Observations by Device page, see Analytics: Observations by Device. - Choose Flow Analysis to access the Flow Search Results page that contains flow information related to the source. For information about the Flow Search Results page, see Flow Search Results: Overview.
Time	The time the alert was last updated.
Tags	A tag is any type of label or text that you add to an alert so that it can be filtered by that label or text. Use the Tags drop-down list to filter by a specific tag.

Edit Alerts Table Entries

You can assign tags to specific alerts as well as change the status of alerts by using the drop-down lists in the Actions field (located immediately above the Alert table).

1. In the Alerts table, check the checkbox next to each alert for which you want to make one or more changes. Note that if you check multiple alerts, the changes you make will be uniformly applied to all of the alerts you selected.
2. In the Actions field, make your selection(s) from one of the following drop-down lists, one at a time:
 - Assign Tag
 - Change Status

You can change the status to Open, Close, or Unsnnooze. When you close an alert, the Close Alert dialog opens. In this dialog you can specify if you want to snooze the alert for a certain period of time. If you snooze an alert and then decide you want to reopen it, choose the Unsnnooze option for that alert.

To choose more than one option at a time from a particular drop-down list, press the **Shift key** while clicking each option. To de-select a choice, click it again.

3. To save your options, click outside the drop-down list on any white space on the page.

The applicable entries in the table change to reflect your selections.

View Related Configuration Pages

Click the  (**Related Config Links**) icon and then click the appropriate option to access either the Alert Priorities Configuration page or the Alerts Country Watchlist Configuration page.

Download a .CSV File

To download a .csv file containing either all available observations or just the current filtered view (if you have already filtered the table), click the  (**Download CSV**) icon located in the top right corner of the page and click the appropriate option.

 You can download a maximum of 65,000 alert records from the Alerts Summary.

Related Tasks

[Analytics: Investigating Alerts](#)

[Analytics: Alert Details](#)

Related Topics

[Analytics: Alerts Workflow](#)

[Analytics: Alerts FAQs](#)

[Analytics: Device Report](#)

[Analytics: Overview](#)

Analytics: Alerts Workflow

An alert's workflow is based around its status. When the system generates an alert, its default status is Open. On the Alert Details page, all open alerts are displayed by default since these are of immediate concern.

As you review the Alerts Summary, you can update their statuses as an initial triage. You can use the filters and search functionality to locate specific alerts, or you can display alerts of different statuses or those with different tags.



When you close an alert, you can set the alert's status to Snoozed, in which case it does not reappear in the list of open alerts until the snooze period elapses. You can also remove Snoozed status from an alert to display it as an open alert again.

When you click an alert on the Alerts page, you can view its details. From this page you can click an alert's associated device.

This information can help you pinpoint the actual issue, in order to further research the issue on your network, and potentially resolve malicious behavior. As you research within the Manager and on your network, you can leave comments that describe your findings on the alert. This helps create a record for your research that you can reference in the future.

If you complete your analysis, you can update the status to Closed so that it no longer appears by default as an open alert.

Related Tasks

[Analytics: Investigating Alerts](#)

Related Topics

[Analytics: Alerts FAQs](#)

[Analytics: Alert Details](#)

[Analytics: Device Report](#)

[Analytics: Alerts Summary](#)

[Analytics: Overview](#)

Analytics: Alerts FAQs

Why are certain alerts disabled?

Certain alerts are disabled by default because only alerts that are likely to affect the broadest cross-section of customers have been enabled. Users can enable alerts per their network needs, especially as part of fine-tuning their deployment.

What does an alert's status mean?

An alert's workflow is based around its status. When the system generates an alert, its default status is Open. On the Alerts page, all open alerts are displayed by default since these are of immediate concern.

As you review the Alerts Summary, you can assign and tag alerts as well as update their statuses as an initial triage. You can use the filters and search functionality to locate specific alerts, or you can display alerts of different statuses or those with different tags. When you close an alert, you can set the alert's status to Snoozed, in which case it does not reappear in the list of open alerts until the snooze period elapses. You can also remove the Snoozed status from an alert to display it as an open alert again.

If you complete your analysis, you can update the status to Closed so that it no longer appears by default as an open alert. You can also re-open a closed alert in the future if circumstances change.

How does a subnet's sensitivity affect alerts?

The only subnets in Secure Network Analytics that this feature currently applies to are the following:



- 10.0.0.0/8 (Default RFC1918)
- 172.16.0.0/12 (Default RFC1918)
- 192.168.0.0/16 (Default RFC1918)
- fc00::/7 (Default RFC4193)

By default, the sensitivity type for these subnets is Medium.

A subnet's sensitivity and an alert's priority type determine when a specific alert is generated, based on that particular subnet's traffic. The alert priority type influences the degree to which subnet traffic is monitored.

Alerts default to Low or Normal. The alert priority is determined by the alert type. From the Priority drop-down list on the Alert Priorities Configuration page (**Configure > Detection**

> **Alerts**), you can configure any alert type to be Low, Normal, or High. You can also filter the Alert Settings page by priority using the **Filter** icon in the Priority column header. If no activity occurs for that alert after 14 days, the alert automatically closes.

Refer to the following matrix to learn which combinations of subnet sensitivity and alert priority type generate alerts for the previously mentioned subnets.

Subnet Sensitivity Status	Alert Priority Type		
	Low	Medium	High
Low	No open alert.	No open alert.	Generates an open alert.
Medium	No open alert.	Generates an open alert.	Generates an open alert.
High	Generates an open alert.	Generates an open alert.	Generates an open alert.

Related Topics

[Analytics: Alerts Workflow](#)

[Analytics: Investigating Alerts](#)

[Analytics: Alert Details](#)

[Analytics: Device Report](#)

[Analytics: Alerts Summary](#)

[Analytics: Overview](#)

Analytics: Investigating Alerts

The following presents general guidelines and suggestions for how to investigate a given alert. Because Secure Network Analytics provides additional context when it logs an alert, you can use this context to help guide your investigation.



These steps are meant to be neither comprehensive, nor all-inclusive. They merely offer a general framework with which to start investigating an alert.

In general, you can take the following steps when you review an alert:

Triage Open Alerts

It's particularly useful to triage the open alerts if more than one have yet to be investigated.

On the Alerts Summary, in the **See alerts by status** field, click **Open** to filter by Open status.

Ask the following questions:

- Have you configured this alert type as high priority?
- Did you set a high sensitivity for the affected subnet?
- Is this unusual behavior from a new entity on your network?
- What is the entity's normal role, and how does the behavior in this alert fit that role?
- Is this an exceptional deviation from normal behavior for this entity?
- Is protected or sensitive data at risk of being compromised?
- How severe is the impact to your network if this behavior is allowed to continue?
- If there is communication with external entities, have these entities established connections with other entities on your network in the past?
- If this is a high priority alert, consider quarantining the entity from the internet, or otherwise closing its connections, before continuing your investigation.

Snooze an Alert

Snooze alerts for later analysis when they are of lesser priority, as compared to other alerts. For example, if your organization is repurposing an email server as an FTP server, and the system generates an Emergent Profile alert, you can snooze this alert as it is intended behavior, and revisit it at a later date. A snoozed alert does not show up with the open alerts; you must specifically filter to review these snoozed alerts.

1. On the Alerts Summary, In the Alerts table, check the checkbox next to the applicable alert.
2. From the Change Status drop-down list, choose **Close**.
3. In the Close Alert dialog, specify a snooze period from the drop-down list and click **Submit**.

Unsnnooze a Snoozed Alert

When you are ready to review a snoozed alert, you can unsnooze it. This sets the alert's status to Open.

1. On the Alerts Summary, In the Alerts table, check the checkbox next to the applicable alert.
2. In the Change Status drop-down list, choose **Unsnnoozed**.

Close an Alert

When you are finished investigating an alert, you can close it.

1. On the Alerts Summary, In the Alerts table, check the checkbox next to the applicable alert.
2. In the Change Status drop-down list, choose **Close**.
3. On the Alerts Summary, in the Alerts Table, confirm that the alert is indeed closed.

Reopen a Closed Alert

If you discover additional information related to a closed alert, or want to add more comments related to that alert, you can reopen it, changing the status to Open. You can then make changes as necessary to the alert, then close it again when your additional investigation is complete.

1. Filter the Alert list to display all closed alerts.
2. Search for and click the alert you need to reopen to view its details.
3. In the Change Status drop-down list, choose **Open**.

Update an Alert

Based on your initial triage, do one or more of the following:

- On the Alerts Summary, in the Actions field, add tags to the alert. This enables you to better categorize your alerts for future identification as well as helps to establish long-term patterns in your alerts.
- At the bottom of the Alert Details page in the Comments text box, enter a comment for the alert.

Review an Alert

If you are reviewing an assigned alert, review the alert detail to understand why Secure Network Analytics generated an alert.

To do this, do one or both of the following:

- On the Observations Dashboard, click the  (**Arrow Right**) icon next to an observation type to view all logged observations of that type.
- On the Alerts Details page, view all logged observations for this alert's source entity.

Following are some suggestions for what to review:

- Review the supporting observations to understand what these observations mean for the source. View all observations for the source to understand its general behavior and patterns, and see if this activity may be part of a longer trend.
- Review the supporting observations. Understand what these observations mean for the source entity.
- View all of the observations for the applicable source to understand its general behavior and patterns, and see if this activity may be part of a longer trend.
- From the observations, view additional context surrounding the source, including other alerts and observations it may be involved in, information about the device itself, and what type of flow traffic it is transmitting. Determine if this behavior is indicative of malicious behavior. If the source entity established connections with multiple external entities, determine if the external entities are somehow related, such as if they all have similar geolocation information, or their IP addresses are from the same subnet.
- From the observations, review the context for the entities with which the source established a connection. Examine the geolocation information, and determine if any of the geolocation data identifies a malicious entity. View the traffic generated by these entities. Check whether Talos, AbuseIPDB, or Google have any information on these entities. Find the IP address on multiple days and see what other types of connections the external entity established with entities on your network. If

necessary, locate those internal entities and determine if there is any evidence of compromise or unintended behavior.

Review Supporting Observations and Contextual Details

Source Entities

Review the supporting observations to understand what these observations mean for the source entity. Determine if the source entity behavior indicates malicious behavior. If the source entity established connections with multiple external entities, determine if the external entities are somehow related, such as if they all have similar geolocation information, or their IP addresses are from the same subnet. View additional context surrounding the source entity, including other alerts and observations it may be involved in, information about the device itself, and the type of flow traffic it is transmitting.

From an observation, you have the following options:

From the Device drop-down list,

- Choose **Alerts** to view all alerts related to the entity.
- Choose **Observations** to view all observations related to the entity.
- Choose **Device** to view information about the device.
- Choose **Flow Analysis** to view flow traffic related to this entity.

From the IP address or hostname drop-down list,

- Click the  (**Copy**) icon to copy the IP address or hostname.

External Entities

From the observations, examine information about other external entities. Examine the geolocation information, and determine if any of the geolocation data identifies a malicious entity. View the traffic generated by these entities. Check whether Talos, AbuseIPDB, or Google have any information on these entities. Find the IP address on multiple days and see what other types of connections the external entity established with entities on your network. If necessary, locate those internal entities and determine if there is any evidence of compromise or unintended behavior. Review the context for the entities with which the source entity established a connection:

From an observation, you have the following options:

From the IP address or hostname drop-down list,

- Choose the  (**Copy**) icon to copy the IP address or hostname.
- Choose **Find IP on multiple days** to see the amount of traffic sent to and from the corresponding entity and the number of connections it was involved in, for the previous day, today, and the next day. If you click a date in the Day column, you can view additional traffic-related information for the associated day for that entity.
- Choose **IP Traffic** to view recent traffic information for this entity.
- Choose **Flow Analysis** to view the flow search results for this entity.
- Choose **Add IP to watchlist** to add this entity to the watchlist.
- Choose **AbuseIPDB** to view information about this entity on AbuseIPDB's website.
- Choose **Cisco Umbrella** from the IP address or hostname drop-down to view information about this entity on Cisco Umbrella's website.
- Choose **Google Search** to search for this IP address on Google.
- Choose **Talos Intelligence** to view information about this information on Talos's website.

Remediate an Issue

If malicious behavior caused the alert, remediate the malicious behavior.

- If a malicious entity attempted to log in from outside your network, update your firewall rules to prevent the entity from accessing your network.
- If you identify a vulnerability or exploit, update or patch the affected entity to remove the vulnerability, or update your firewall settings to prevent unauthorized access. Determine if other entities on your network may similarly be affected, and apply the same update or patch to those entities. If the vulnerability or exploit currently does not have a fix, contact the appropriate vendor to let them know.
- If you identify malware, quarantine the entity and remove the malware. Determine if other entities on your network are at risk, and update the entities or your security solution to prevent this malware from spreading. Update your security intelligence with information about this malware, or the entities that caused this malware. Alert vendors as necessary.
- If malicious behavior resulted in data exfiltration, determine the nature of the data sent to an unauthorized source. Follow your organization's protocols for unauthorized data exfiltration.
- Enter comments about what you find. To do this, scroll to the bottom of the Alert Details page and enter it in the Comment text box. When finished, click **Comment**.

Fine-Tune Your Secure Network Analytics Settings

Based on the alert and remediation, update your Secure Network Analytics settings to help identify this behavior in the future.

- Add countries to the country watchlist if multiple entities from a country caused malicious behavior. For more information, see [Analytics: Configuring Country Watchlist](#).
- Update your subnet sensitivity if a particular subnet is targeted. For more information, see the "How does a subnet's sensitivity affect alerts?" section in [Analytics: Alerts FAQs](#).
- Update your alert type priority settings if a specific alert becomes a concern.
- Add any known good scanners to the default Network Scanners Host Group (ID 48).

Related Tasks

[Analytics: Alert Details](#)

[Analytics: Alerts Summary](#)

Related Topics

[Analytics: Alerts Workflow](#)

[Analytics: Alerts FAQs](#)

[Analytics: Device Report](#)

[Analytics: Overview](#)

Analytics: Alert Details



For more information on how to use this page, see [Analytics: Investigating Alerts](#)

The Alert details presents an overview of the alerts reported by the system. You can search for specific text, or filter by status, or tags. You can also see all observations generated for the affected entity as well as view related device information.

The following information is displayed in the Alert Type Details section.

Open Alert Details

On the Alerts Summary, click an alert.

The Alert Details page for that alert opens.

View Alert Type Details

Field	Description
Description	A description of the alert.
Next Steps	The next steps you should take in investigating the alert.
MITRE Tactics	The MITRE tactics associated with the alert. To access the MITRE page for this tactic, either click the MITRE Tactic entry or hover your cursor over the entry and, in the pop-up window that opens, click the "See Full Details at" link.
MITRE Techniques	The MITRE techniques associated with the alert. To access the MITRE page for this technique, either click the MITRE Technique entry or hover your cursor over the entry and, in the pop-up window that opens, click the "See Full Details at" link.
Alert Type Priority	The alert priority: Low, Normal, High.

View Alert Rule Details

In the Alert Rule Details section, you can view additional information related to the alert.

The following information is displayed in the Alert Rule Details section.

Field	Description
Status	The alert status: Open, Closed, Snoozed.
ID	The alert ID number.
Latest Observation	The most recent time the alert was observed.
First Observation	The time the alert was first observed.
Detected At	The time the alert was detected by the engine.
Tags	A tag is any type of label or text that you add to an alert so that it can be filtered by that label or text.
Close Alert	To close the alert, click Close Alert .

View descriptions for the Supporting Observations table

An alert's details display a list of observations that led to this alert being generated. You can review these for more information about the network behavior that led to this alert.



The fields displayed in this table vary depending on the alert the table is associated with.

Option Name	Description
Affected Resource	The affected resource.
Affected Resource Type	The affected resource type.
Anomaly	The type of anomaly detected.
Bytes In	Amount of traffic (in bytes) that has been received by the device for a specific point in time.

Option Name	Description
Bytes Out	Amount of traffic (in bytes) that has been sent from the device for a specific point in time.
CIDR Range	The approximate CIDR notation scanning range (the actual range may be smaller).
Connected Device	The device to which the endpoint or source established a connection.
Connected IP	The IP with which this source communicated.
Connected Ports	The port with which this port communicated.
Corresponding Ports	The ports that were used in the communication.
Data Sink IP	The IP address of the external device to which data was uploaded.
Data Sink Profile	The Local profile of the observation source (this device) when uploading to the data sink.
Data Source	The internal device from which the data was downloaded.
Data Source IP	The IP address of the internal device from which the data was downloaded.
Data Source Profile	The local profile of the observation source (this device) when downloading from the data source.
Device	The associated endpoint or source. Click the drop-down list to access the following options: - Choose Device (Network) to access the Device Report filtered by the source. For information about the Device page, see Analytics: Device Report. - Choose Alerts to access the Alerts Summary that

Option Name	Description
	contains a list of all the alerts related to that source. For information about the Alerts Summary, see Analytics: Alerts Summary. - Choose Observations to access the Observations by Device page that contains a list of all observations related to the source. For information about the Observations by Device page, see Analytics: Observations by Device. - Choose Flow Analysis to access the Flow Search Results page that contains flow information related to the source. For information about the Flow Search Results page, see Flow Search Results: Overview.
Domain/URL	The domain/URL based on NGFW connection events or passive DNS.
Download Sec	The length of time required for the download to complete.
Download Speed bps	The download speed in bps.
Download Start	The time the download began for the external data sink.
External IP	The external IP address.
Failed Attempts	The number of times an entity has attempted to establish a connection to a device.
Heartbeat period (Seconds)	The time between heartbeats.
History Length (Days)	The number of days of history that was used to calculate the normal set.
Internal Port Set	The type of port set. For example, "connected internal" is

Option Name	Description
	the set of connected ports used for internal connections.
Last Active	The time that an observation was last active.
Local Device	The local device involved in the communication.
Local Port	The port over which the endpoint or source connected to the device.
Lookback Days	The number of days of history that was used to calculate the normal set.
Lost Port Sets	Ports that are no longer used on this date.
Matching Watchlists	If the watchlist is domain-based, the matching domain names are listed here.
Metric	The metric for this outlier. For example, an outlier for internal "Bytes In" indicates that the internal network traffic (where there is no internet) to the device has spiked.
New Connections	New connections on this date that weren't in the lookback period.
New Port Set	Ports that were used on this date that weren't used in the lookback period.
New Profile	A new device profile that differs from previous behavior.
Normal Connection Set	The connections found in the lookback period.
Normal Ports Set	The ports that were used in the lookback period.
Number of Heartbeats	The number of times the server was connected during this observed event.
Packets In	The packets received by the source.

Option Name	Description
Packets Out	The packets sent from the source.
Port	The source port used in the observed event.
Port Ranges	The device-scanned ports included in this range, and possibly others. Common targets may include web server targets.
Probability	The probability that you would see this outlier.
Profile	The role(s) associated with the endpoint or source to which the device connected.
Public Facing IP	A public IP address that was discovered on a watchlist.
Remote Device	The device this source communicated with. Click the drop-down list to access the following options: - Choose Device (Network) to access the Device Report filtered by the source. For information about the Device page, see Analytics: Device Report. - Choose Alerts to access the Alerts Summary that contains a list of all the alerts related to that source. For information about the Alerts Summary, see Analytics: Alerts Summary. - Choose Observations to access the Observations by Device page that contains a list of all observations related to the source. For information about the Observations by Device page, see Analytics: Observations by Device. - Choose Flow Analysis to access the Flow Search Results page that contains flow information related to the source. For information about the Flow Search Results page,

Option Name	Description
	see Flow Search Results: Overview .
Remote IP	The IP address with which this source communicated.
Remote Port	The port with which this source communicated.
Resource	The affected resource.
Retrieve Sessions	Provides a link to execute a supporting query on the Flow Search page. Use it to explore telemetry data related to the components of the observation, such as the observed time range, source and peer communication, observed port or protocol, and so on. Use the Flow Search for detailed session analysis.
Sample Size	The number of historical samples used in this calculation.
Scanned Device	The IP address or host name of the device that was scanned.
Scanned Ports	The device-scanned ports included in this range, and possibly others.
Scanner Device	The IP address or host name of the device that performed the scan.
Severity	The severity of the reported event.
Source	The associated endpoint or source. Click the drop-down list to access the following options: - Choose Device (Network) to access the Device Report filtered by the source. For information about the Device page, see Analytics: Device Report. - Choose Alerts to access the Alerts Summary that contains a list of all the alerts related to that source.

Option Name	Description
	For information about the Alerts Summary, see Analytics: Alerts Summary. - Choose Observations to access the Observations by Device page that contains a list of all observations related to the source. For information about the Observations by Device page, see Analytics: Observations by Device. - Choose Flow Analysis to access the Flow Search Results page that contains flow information related to the source. For information about the Flow Search Results page, see Flow Search Results: Overview.
Source IP	The IP of the source at the time of the observed event.
Source Port	The port of the source at the time of the observed event.
Suspect Connections	New connections to unusual external sources or that are on unusual ports.
Time	The time the observation occurred.
Time Window	The length of time during which the event was shared.
Transfer Size	The number of bytes transferred.
Upload Start	The start time of the upload to the external data sink.
Upload Sec	The duration of the upload.
Upload Speed bps	The upload speed in bps.
Violating Port	The port used by the source.
Violating Protocol	The network protocol used by the source (for example, TCP).

Option Name	Description
Violating Type	The type of violation.

Other Pages You Can Access from the Alert Details Page

Alert Priorities Configuration

To access the Alert Priorities Configuration page, click the **"Go to Alert Priorities page"** link at the bottom of the Alert Type Details section.

The Alert Priorities Configuration page opens.

Flow Search Results

To access the Flow Search results page for a particular time displayed in the Supporting Observations table, click the applicable entry in the Time column.

OR

From the Supporting Observations table, click the drop-down list for the applicable device, and from the menu that is displayed, choose **Flow Analysis**.

The Flow Search Results page opens.

For information about the Flow Search Results page, see [Flow Search Results: Overview](#).

Device Report

To access the Device Report for a particular device displayed in the Supporting Observations table, click the drop-down list for the applicable device, and from the menu that is displayed, choose **Device**.

The Device Report opens.

For information about the Device page, see [Analytics: Device Report](#).

Alerts Summary

To access the Alerts Summary that contains a list of all the alerts related to a particular device displayed in the Supporting Observations table, click the drop-down list for the applicable device, and from the menu that is displayed, choose **Alerts**.

The Alerts Summary opens, filtered by that device.

For information about the Alerts Summary, see [Analytics: Alerts Summary](#).

Observations by Device

To access the Observations by Device page that contains a list of all observations related to a particular device displayed in the Supporting Observations table, click the drop-down list for the applicable device, and from the menu that is displayed, choose **Observations**.

The Observations by Device page opens, filtered by that device.

For information about the Observations by Device page, see [Analytics: Observations by Device](#).

Enter Comments for an Alert

1. In the Comments text box at the bottom of the page, enter any applicable comments.
2. Click **Comment**.

Related Tasks

[Analytics: Investigating Alerts](#)

[Analytics: Alerts Summary](#)

Related Topics

[Analytics: Alerts Workflow](#)

[Analytics: Alerts FAQs](#)

[Analytics: Device Report](#)

[Analytics: Alert Prerequisite Chart](#)

Analytics: Post Alerts to Cisco XDR

Analytics alarms and observations are sent to Cisco XDR if the Secure Network Analytics XDR integration option is enabled.

For more information, refer to the following topics:

- [Analytics: Alert Details](#)
- [Analytics: Alerts Summary](#)
- [Managing the Cisco XDR Integration](#)

 For more information about Cisco XDR, go to the [Cisco XDR site](#).

Analytics: Device Report

The Device Report provides behavioral information about an entity within your network.

- To view additional context about an alert, access the Device Report from the Alert Details page.
- To view additional context about an observation, access this report from one of the Observation pages listed in the next section.

Open Device Report

On any of the following pages, click the drop-down list for the desired device displayed in the applicable table, and from the menu that is displayed, choose **Device**.

- Alert Details
- Observation Highlights
- Observations by Device
- Selected Observations

The Device Report opens.

Device Report Overview

Refer to the following sections to learn what you can view on this page.

Device Outline

In this section on the left side of the page, you can view general information about the device, such as data for the previous day and Secure Network Analytics generated data.

Click the  (**Copy**) icon beside an entry to copy it to the clipboard.

Alerts

In the Alerts section, you can see at a glance the number of open alerts, closed alerts, observations, total connections, and total traffic.

- If you click the Open or Closed number, the [Alerts Summary](#) opens, displaying all open alerts or all closed alerts (depending on which number you clicked) associated with this device.
- If you click the Observations number, the [Observations by Device](#) page opens, listing the observations associated with this device.

History

The History line graph displays the amount of traffic sent to and from an entity and the number of connections it was involved in, per 1-day intervals. Hover your cursor over points in the graph to view details about the entity's traffic for that day.

Edit Time Range

By default, the Summary, Traffic, Traffic Connections Visualization, Profiling, and IPs tabs display information for the current day. However, you can also choose to view information for a day in the past. To do this, complete the following steps.

1. From the **Selected Date** dropdown list located immediately under the History graph, use the left and right arrows to navigate to the desired month.
2. Click the desired date.
3. Click **Select Date**.
4. Click **Apply**, which is located on the right side of the Selected Date section.

Summary

The Summary tab displays an overview of the entity's model that is stored by the system. Refer to the following table for descriptions of the fields displayed on the Summary tab.

Entity Summary Fields

Field	Description
Attendance	
Normally Active	The time period during which this entity is normally active.
IP Addresses	The IP addresses associated with the entity.
Connectivity	
Connections	The number of connections this entity was involved in.
Internal Connections	The number of connections with internal entities this entity was involved in.
External Connections	The number of connections with external entities this entity was involved in.

Traffic	
Bytes In	The amount of traffic that the entity received.
Bytes Out	The amount of traffic that the entity sent.
Bytes Total	The amount of traffic that the entity transmitted in total.
Traffic Internal	
Bytes In	The amount of traffic that the entity received from internal entities.
Bytes Out	The amount of traffic that the entity sent to internal entities.
Traffic External	
Bytes In	The amount of traffic that the entity received from external entities.
Bytes Out	The amount of traffic that the entity sent to external entities .
DNS Names	
DNS Names	DNS domain names associated with the entity.
Roles	
Roles	The roles associated with this entity.
Profiles	
Profiles	The percentage of time that the entity acted corresponding to the listed profile.

Traffic

The Traffic line graph displays the amount of traffic sent to and from an entity and the number of connections it was involved in, per 10-minute intervals. You can also view information about the connections with which the entity was involved.

- Hover your pointer over areas on the various lines in the Traffic graph to view details on the entity's traffic for that 10-minute interval.
- To filter the graph, you have the following options:

- Click **All** to view information about all entities with which this entity established a connection.
- Click **Internal** to view information about internal entities with which this entity established a connection.
- Click **External** to view information about external entities with which this entity established a connection.

To export the results on the Traffic tab to a CSV file, click the  (Download CSV) icon. You can download a maximum of 65,000 results.

Refer to the following table for descriptions of the fields displayed on the Traffic tab.

Entity Traffic Fields

Field	Description
Connected IP	The IP address that this entity established a connection with.
Hostname/PDNS Record	The hostname for this IP address, if available.
Bytes In	The bytes received by the entity from the connected entity.
Bytes Out	The bytes sent by the entity to the connected entity.
Bytes Total	The total bytes transmitted by the entity in this connection.
First Connection	The time of the first connection with the connected IP on this day.
Last Connection	The time of the last connection with the connected IP on this day.
See Conversation in Flow Analysis (Ellipsis icon)	In the last column, click the ... (Ellipsis) icon > See Conversation in Flow Analysis to access the associated flow search results.

IPs

The IPs tab displays a list of the IPs associated with the device for each associated date as well as a consolidated list of all IPs associated with the device for the last 30 days. It also provides the date that each of these IPs was last active in association with the device.

Field	Description
IPs	The IP associated with the device.
Last Active	The date that the IP was last active in association with the device.

Related Topics

[Analytics: Alerts Workflow](#)

[Analytics: Alerts FAQs](#)

[Analytics: Investigating Alerts](#)

[Analytics: Alert Details](#)

[Analytics: Alerts Summary](#)

[Analytics: Overview](#)

Analytics: Roles

Use this page to view a list of all roles (in the Active Roles table), along with at least one associated device for each role, that were active on your system for a specific period of time. The entries that are displayed in the Active Roles table depend on the telemetry types your system is ingesting. You can also view a list of all roles that are currently inactive on your system (the Inactive Roles table). The default time frame for both tables is the last 7 days, though you can set the time frame to a maximum of 90 days.

- You can see the Roles page only if you've enabled Analytics.
- The following roles will always be inactive in Cisco Secure Network Analytics (and therefore will always be listed in the Inactive Roles table):



- AWS
- Azure
- GCP
- Any other Cloud telemetry-related roles

Open the Roles page

From the main menu, choose **Investigate > Assets > Roles**.

The Roles page opens.



After you enable Analytics, you need to wait 24 hours for all jobs to run before the Roles page displays information.

Set the time frame

The time frame for the results defaults to the last 7 days. To configure a different time frame, complete the following steps:

1. Click either **Start** or **End** text boxes.
The calendar opens.
2. Choose your desired start and end date from the calendar.
3. Click **Apply** to confirm and set the selected date range.
Click **Reset** , to revert to the default date settings.

View the results

To view the matching devices for a role, check the check box of the role listed in the Active Roles table and the matching sources and role names for this role are added to a table that opens to the right.

Each time you check the check box of the role to view its matching sources and role names, the matching sources and role names for previous selections continue to be displayed. To remove matching sources and role names for a role, uncheck the check box of the role listed in the Active Roles table.

To view details about a matching source, hover your cursor over the  (**Alert**) icon or the  (**Eye**) icon that is beside the applicable source. (If the source is associated with a minimum of one open alert, the **Alert** icon is displayed. If the source is not associated with an open alert, the **Eye** icon is displayed.)

If you click the menu to the right of a matching source IP address, you can choose from several options to access the following pages:

Choose...	To access this page...
Device	Device Report. In the Roles section you will see the roles associated with the device, and in the Profiles section you will see the processes or protocols associated with the roles.
Alerts	Alerts Summary
Observations	Observations by Device
Flow Analysis	Flow Search Results

Export results

To export the results on this page to a CSV file, click the  (**Download CSV**) icon.

Related Topics

[Analytics: Alerts Workflow](#)

[Analytics: Alerts FAQs](#)

[Analytics: Investigating Alerts](#)

[Analytics: Alert Details](#)

[Analytics: Alerts Summary](#)

[Analytics: Overview](#)

Analytics: Observations Dashboard



- For information about the Observation Types page, see [Analytics: Observation Types](#).
- For information about the Observations by Device page, see [Analytics: Observations by Device](#).
- For information about the Selected Observations page, see [Analytics: Selected Observations](#).

Open Observations Dashboard

From the main menu, choose **Monitor > Observations**.

The Observations Dashboard opens.

Observations Highlights Overview

Observations are facts about an entity's behavior on the network, such as a heartbeat connection with an external IP address, an interaction with an entity on a watchlist, or a remote access session established with another entity. Observations on their own do not carry meaning beyond the fact of what they represent. A typical customer may have many thousands of observations and a few alerts.

Alerts are generated from combinations of observations; observations in isolation do not necessarily constitute malicious behavior. The recent observations, by themselves, do not necessarily mean there is malicious behavior on your network. Review your alerts for a better picture of the potential malicious behavior.

As the system inspects your traffic, it logs observations, or facts, about the entities on your network. Because observations are logged per entity, your network may generate more observations than can be reasonably reviewed. The system presents a subset of the most notable observations logged for your network. You can review and filter these to gain a better understanding of the types of behavior that may result in alerts being generated.



A high number of observations does not necessarily correspond with a high number of alerts, or even any alerts. For example, an entity with many observations may be passing a large amount of varied traffic, but Dynamic Entity Modeling may have identified this behavior as normal and expected for this entity. Similarly, a low number of observations does not necessarily correspond with a lack of alerts. For example, if the only activity detected for an entity is a continual attempt to log into a server with improper credentials, this may



generate a relatively small number of observations, and an alert noting the multiple failed login attempts.

View Observations Highlights

You can drill down and view all observations of that type. If you drill down, the system opens a new tab on this page that displays those observations. If you select a different observation type to drill down into, the system updates that new tab with these observations.

- To view all observations of a specific type, click the  (**Arrow Right**) icon next to the applicable observation type.

The Selected Observations page opens, which by default displays a table that contains all observations of the type you selected that have occurred within the last 24 hours.

For more information about the Selected Observations page, see [Analytics: Selected Observations](#).

- To download a .csv file containing all observations for a specific type, click the  (**Download CSV**) icon located at the top right corner of the associated table.

View Descriptions for the Observations Highlights tables



The fields displayed in these tables vary depending on the observation each table is associated with.

Option Name	Description
Affected Resource	The affected resource.
Affected Resource Type	The affected resource type.
Anomaly	The type of anomaly detected.
Bytes In	Amount of traffic (in bytes) that has been received by the device for a specific point in time.
Bytes Out	Amount of traffic (in bytes) that has been sent from the

Option Name	Description
	device for a specific point in time.
CIDR Range	The approximate CIDR notation scanning range (the actual range may be smaller).
Connected Device	The device to which the endpoint or source established a connection.
Connected IP	The IP with which this source communicated.
Connected Ports	The port with which this port communicated.
Corresponding Ports	The ports that were used in the communication.
Data Sink IP	The IP address of the external device to which data was uploaded.
Data Sink Profile	The Local profile of the observation source (this device) when uploading to the data sink.
Data Source	The internal device from which the data was downloaded.
Data Source IP	The IP address of the internal device from which the data was downloaded.
Data Source Profile	The local profile of the observation source (this device) when downloading from the data source.
Device	The associated endpoint or source. Click the drop-down list to access the following options: - Choose Device (Network) to access the Device Report filtered by the source. For information about the Device page, see Analytics: Device Report. - Choose Alerts to access the Alerts Summary that contains a list of all the alerts related to that source.

Option Name	Description
	For information about the Alerts Summary, see Analytics: Alerts Summary. - Choose Observations to access the Observations by Device page that contains a list of all observations related to the source. For information about the Observations by Device page, see Analytics: Observations by Device. - Choose Flow Analysis to access the Flow Search Results page that contains flow information related to the source. For information about the Flow Search Results page, see Flow Search Results: Overview.
Domain/URL	The domain/URL based on NGFW connection events or passive DNS.
Download Sec	The length of time required for the download to complete.
Download Speed bps	The download speed in bps.
Download Start	The time the download began for the external data sink.
External IP	The external IP address.
Failed Attempts	The number of times an entity has attempted to establish a connection to a device.
Heartbeat period (Seconds)	The time between heartbeats.
History Length (Days)	The number of days of history that was used to calculate the normal set.
Internal Port Set	The type of port set. For example, "connected internal" is the set of connected ports used for internal connections.

Option Name	Description
Last Active	The time that an observation was last active.
Local Device	The local device involved in the communication.
Local Port	The port over which the endpoint or source connected to the device.
Lookback Days	The number of days of history that was used to calculate the normal set.
Lost Port Sets	Ports that are no longer used on this date.
Matching Watchlists	If the watchlist is domain-based, the matching domain names are listed here.
Metric	The metric for this outlier. For example, an outlier for internal "Bytes In" indicates that the internal network traffic (where there is no internet) to the device has spiked.
New Connections	New connections on this date that weren't in the lookback period.
New Port Set	Ports that were used on this date that weren't used in the lookback period.
New Profile	A new device profile that differs from previous behavior.
Normal Connection Set	The connections found in the lookback period.
Normal Ports Set	The ports that were used in the lookback period.
Number of Heartbeats	The number of times the server was connected during this observed event.
Packets In	The packets received by the source.
Packets Out	The packets sent from the source.

Option Name	Description
Port	The source port used in the observed event.
Port Ranges	The device-scanned ports included in this range, and possibly others. Common targets may include web server targets.
Probability	The probability that you would see this outlier.
Profile	The role(s) associated with the endpoint or source to which the device connected.
Public Facing IP	A public IP address that was discovered on a watchlist.
Remote Device	The device this source communicated with. Click the drop-down list to access the following options: - Choose Device (Network) to access the Device Report filtered by the source. For information about the Device page, see Analytics: Device Report. - Choose Alerts to access the Alerts Summary that contains a list of all the alerts related to that source. For information about the Alerts Summary, see Analytics: Alerts Summary. - Choose Observations to access the Observations by Device page that contains a list of all observations related to the source. For information about the Observations by Device page, see Analytics: Observations by Device. - Choose Flow Analysis to access the Flow Search Results page that contains flow information related to the source. For information about the Flow Search Results page, see Flow Search Results: Overview.

Option Name	Description
Remote IP	The IP address with which this source communicated.
Remote Port	The port with which this source communicated.
Resource	The affected resource.
Retrieve Sessions	Provides a link to execute a supporting query on the Flow Search page. Use it to explore telemetry data related to the components of the observation, such as the observed time range, source and peer communication, observed port or protocol, and so on. Use the Flow Search for detailed session analysis.
Sample Size	The number of historical samples used in this calculation.
Scanned Device	The IP address or host name of the device that was scanned.
Scanned Ports	The device-scanned ports included in this range, and possibly others.
Scanner Device	The IP address or host name of the device that performed the scan.
Severity	The severity of the reported event.
Source	The associated endpoint or source. Click the drop-down list to access the following options: - Choose Device (Network) to access the Device Report filtered by the source. For information about the Device page, see Analytics: Device Report. - Choose Alerts to access the Alerts Summary that contains a list of all the alerts related to that source. For information about the Alerts Summary, see Analytics: Alerts Summary.

Option Name	Description
	- Choose Observations to access the Observations by Device page that contains a list of all observations related to the source. For information about the Observations by Device page, see Analytics: Observations by Device. - Choose Flow Analysis to access the Flow Search Results page that contains flow information related to the source. For information about the Flow Search Results page, see Flow Search Results: Overview.
Source IP	The IP of the source at the time of the observed event.
Source Port	The port of the source at the time of the observed event.
Suspect Connections	New connections to unusual external sources or that are on unusual ports.
Time	The time the observation occurred.
Time Window	The length of time during which the event was shared.
Transfer Size	The number of bytes transferred.
Upload Start	The start time of the upload to the external data sink.
Upload Sec	The duration of the upload.
Upload Speed bps	The upload speed in bps.
Violating Port	The port used by the source.
Violating Protocol	The network protocol used by the source (for example, TCP).
Violating Type	The type of violation.

Related Topics

[Analytics: Observation Types](#)

[Analytics: Observations by Device](#)

[Analytics: Selected Observations](#)

[Analytics: Overview](#)

Analytics: Observation Types

Open Observation Types

1. From the main menu, choose **Monitor > Observations**.
2. On the Observations dashboard, click the **Types** tab.

The Observation Types page opens.

View Observation Types

The Observation Types lists all of the types of observations that the system can log, along with a description below the Observation Name, a count of how many of that observation it has logged, and Telemetry details.

To read details of a specific observation type, do the following:

Click the Observation Name you want to view.

The Selected Observations page opens, which by default displays a table that contains all observations of the type you selected that have occurred within the last 24 hours. For more information about the Selected Observation page, see [Analytics: Selected Observations](#).

Related Topics

[Analytics: Observations by Device](#)

[Analytics: Selected Observations](#)

[Analytics: Observations Dashboard](#)

[Analytics: Overview](#)

Analytics: Observations by Device

Open Observations by Device

1. From the main menu, choose **Monitor > Observations**.
2. On the Observations dashboard, click the **By Device** tab.

OR

In the Supporting Observations table on the Alert Details page, from the drop-down list for the desired device, choose **Observations**.

The Observations by Device page opens with a list of all observations related to that device.

View Observations by Device

You can use this page to view the observations of all the devices.

To download a .csv file containing all observations by a specific device, click the  **CSV** (Download CSV) icon located at the top right corner of the page.

Field	Description
Device	The endpoint or source associated with the observation. Click the drop-down list to access the following options: • Choose Device (Network) to access the Device Report filtered by the source. For information about the Device page, see Analytics: Device Report. • Choose Alerts to access the Alerts Summary that contains a list of all the alerts related to that source. For information about the Alerts Summary, see Analytics: Alerts Summary. • Choose Observations to access the Observations by Device page that contains a list of all observations related to the source. For information about the Observations by Device page, see Analytics: Observations by Device.

	Choose Flow Analysis to access the Flow Search Results page that contains flow information related to the source. For information about the Flow Search Results page, see Flow Search Results: Overview.
Count	The number of observations that have occurred for the associated device.
Time	The time at which the system last logged an observation for the associated device.

Related Topics

[Analytics: Observation Types](#)

[Analytics: Selected Observations](#)

[Analytics: Observations Dashboard](#)

[Analytics: Overview](#)

Analytics: Selected Observations

Open Selected Observations

1. From the main menu, choose **Monitor > Observations**.
2. On the Observations dashboard, click the **Selected Observation** tab.

The Selected Observations page opens.

View Selected Observations

You can view all observations of a given type from the Selected Observation window. This allows you to review the observations that Secure Network Analytics is logging based on your network traffic. All observations that have occurred within the last 24 hours are displayed (though you can change this time range using the Time field).

To view selected observations, do the following:

1. From the Observation Type drop-down list, select an **Observation Type**.
2. Enter a filter value in the **Search** field.
3. To set the time range:
 - Enter the start date and time in the **Start** text box or click the calendar icon at the corner of the text box to select the start date and time.
 - Enter the end date and time in the **End** text box or click the calendar icon at the corner of the text box to select the end date and time.
 - Click **Apply** . To revert, click **Reset**.

The time is displayed in UTC.

To download a .csv file for the associated observation, click the  **(Download CSV)** icon located at the top right corner of the table.



The fields displayed in this table vary depending on the observation the table is associated with.

View descriptions for the Selected Observations table

Option Name	Description
Affected Resource	The affected resource.

Option Name	Description
Affected Resource Type	The affected resource type.
Anomaly	The type of anomaly detected.
Bytes In	Amount of traffic (in bytes) that has been received by the device for a specific point in time.
Bytes Out	Amount of traffic (in bytes) that has been sent from the device for a specific point in time.
CIDR Range	The approximate CIDR notation scanning range (the actual range may be smaller).
Connected Device	The device to which the endpoint or source established a connection.
Connected IP	The IP with which this source communicated.
Connected Ports	The port with which this port communicated.
Corresponding Ports	The ports that were used in the communication.
Data Sink IP	The IP address of the external device to which data was uploaded.
Data Sink Profile	The Local profile of the observation source (this device) when uploading to the data sink.
Data Source	The internal device from which the data was downloaded.
Data Source IP	The IP address of the internal device from which the data was downloaded.
Data Source Profile	The local profile of the observation source (this device) when downloading from the data source.
Device	The associated endpoint or source.

Option Name	Description
	Click the drop-down list to access the following options: - Choose Device (Network) to access the Device Report filtered by the source. For information about the Device page, see Analytics: Device Report. - Choose Alerts to access the Alerts Summary that contains a list of all the alerts related to that source. For information about the Alerts Summary, see Analytics: Alerts Summary. - Choose Observations to access the Observations by Device page that contains a list of all observations related to the source. For information about the Observations by Device page, see Analytics: Observations by Device. - Choose Flow Analysis to access the Flow Search Results page that contains flow information related to the source. For information about the Flow Search Results page, see Flow Search Results: Overview.
Domain/URL	The domain/URL based on NGFW connection events or passive DNS.
Download Sec	The length of time required for the download to complete.
Download Speed bps	The download speed in bps.
Download Start	The time the download began for the external data sink.
External IP	The external IP address.
Failed Attempts	The number of times an entity has attempted to establish a connection to a device.

Option Name	Description
Heartbeat period (Seconds)	The time between heartbeats.
History Length (Days)	The number of days of history that was used to calculate the normal set.
Internal Port Set	The type of port set. For example, "connected internal" is the set of connected ports used for internal connections.
Last Active	The time that an observation was last active.
Local Device	The local device involved in the communication.
Local Port	The port over which the endpoint or source connected to the device.
Lookback Days	The number of days of history that was used to calculate the normal set.
Lost Port Sets	Ports that are no longer used on this date.
Matching Watchlists	If the watchlist is domain-based, the matching domain names are listed here.
Metric	The metric for this outlier. For example, an outlier for internal "Bytes In" indicates that the internal network traffic (where there is no internet) to the device has spiked.
New Connections	New connections on this date that weren't in the lookback period.
New Port Set	Ports that were used on this date that weren't used in the lookback period.
New Profile	A new device profile that differs from previous behavior.
Normal	The connections found in the lookback period.

Option Name	Description
Connection Set	
Normal Ports Set	The ports that were used in the lookback period.
Number of Heartbeats	The number of times the server was connected during this observed event.
Packets In	The packets received by the source.
Packets Out	The packets sent from the source.
Port	The source port used in the observed event.
Port Ranges	The device-scanned ports included in this range, and possibly others. Common targets may include web server targets.
Probability	The probability that you would see this outlier.
Profile	The role(s) associated with the endpoint or source to which the device connected.
Public Facing IP	A public IP address that was discovered on a watchlist.
Remote Device	The device this source communicated with. Click the drop-down list to access the following options: - Choose Device (Network) to access the Device Report filtered by the source. For information about the Device page, see Analytics: Device Report. - Choose Alerts to access the Alerts Summary that contains a list of all the alerts related to that source. For information about the Alerts Summary, see Analytics: Alerts Summary. - Choose Observations to access the Observations by Device page that contains a list of all

Option Name	Description
	observations related to the source. For information about the Observations by Device page, see Analytics: Observations by Device. Choose Flow Analysis to access the Flow Search Results page that contains flow information related to the source. For information about the Flow Search Results page, see Flow Search Results: Overview.
Remote IP	The IP address with which this source communicated.
Remote Port	The port with which this source communicated.
Resource	The affected resource.
Retrieve Sessions	Provides a link to execute a supporting query on the Flow Search page. Use it to explore telemetry data related to the components of the observation, such as the observed time range, source and peer communication, observed port or protocol, and so on. Use the Flow Search for detailed session analysis.
Sample Size	The number of historical samples used in this calculation.
Scanned Device	The IP address or host name of the device that was scanned.
Scanned Ports	The device-scanned ports included in this range, and possibly others.
Scanner Device	The IP address or host name of the device that performed the scan.
Severity	The severity of the reported event.
Source	The associated endpoint or source. Click the drop-down list to access the following options:

Option Name	Description
	- Choose Device (Network) to access the Device Report filtered by the source. For information about the Device page, see Analytics: Device Report. - Choose Alerts to access the Alerts Summary that contains a list of all the alerts related to that source. For information about the Alerts Summary, see Analytics: Alerts Summary. - Choose Observations to access the Observations by Device page that contains a list of all observations related to the source. For information about the Observations by Device page, see Analytics: Observations by Device. - Choose Flow Analysis to access the Flow Search Results page that contains flow information related to the source. For information about the Flow Search Results page, see Flow Search Results: Overview.
Source IP	The IP of the source at the time of the observed event.
Source Port	The port of the source at the time of the observed event.
Suspect Connections	New connections to unusual external sources or that are on unusual ports.
Time	The time the observation occurred.
Time Window	The length of time during which the event was shared.
Transfer Size	The number of bytes transferred.
Upload Start	The start time of the upload to the external data sink.
Upload Sec	The duration of the upload.

Option Name	Description
Upload Speed bps	The upload speed in bps.
Violating Port	The port used by the source.
Violating Protocol	The network protocol used by the source (for example, TCP).
Violating Type	The type of violation.

Related Topics

[Analytics: Observation Types](#)

[Analytics: Observations by Device](#)

[Analytics: Observations Dashboard](#)

[Analytics: Overview](#)

Analytics: Configuring Priorities



- For information about the Alerts Expiration Configuration page, see [Analytics: Configuring Expiration](#).
- For information about the Alerts Country Watchlist Configuration page, see [Analytics: Configuring Country Watchlist](#).

For information about alert statuses, see [Analytics: Alerts FAQs](#).

Open Alert Priorities Configuration

From the main menu, choose **Configure > Detection > Alerts**.

OR

From the main menu, choose **Monitor > Alerts**. Click the **Related Config Links** drop-down list and choose **Alert Priorities**.

The Alert Priorities Configuration page opens.

Configure Alert Priorities

Refer to the following table for descriptions of the fields that appear on this page.



For more information about alerts, see [Analytics: Alerts FAQs](#).

Field	Description
Alert Type	Click the  (Filter) icon and choose one of these options by which you want to filter the alerts: Contains, Does not contain, Is equal to, Is not equal to, Starts with, Ends with, Is null, Is not null, Is empty, and Is not empty.
Observation Types	The Observation Types column shows the observations detected and logged by a monitoring system for the specific Alert Type.
History	The number of days of data collection required to generate the alert type (also known as "soak time").

Priority	Alerts default to Low or Normal. The alert priority is determined by the alert type. From the Priority drop-down list, you can configure any alert type to be Low, Normal, or High. You can also filter the Alert Settings page by priority using the Filter icon in the Priority column header. If no activity occurs for that alert after 14 days, the alert automatically closes. Changing the priority does not affect whether or not an alert is published or unpublished. For more information about unpublished alerts, see the "Filter the Alerts Table" section in Analytics: Alerts Summary.
Enabled	To enable an alert, click the  (Toggle) icon so that the bar displays in blue ().
Telemetry	The Telemetry column shows the telemetry sources for your alerts. Some may have one or more telemetry sources, which are a union of telemetry sources from the observations, roles, etc. The alert fires only if your environment has integrated and is consuming that telemetry source.
MITRE ATT&CK Tactics	The MITRE tactics associated with the alert. To access the MITRE page for this tactic, either click the MITRE Tactic entry or hover your cursor over the entry and, in the pop-up window that opens, click the "See Full Details at" link.
MITRE ATT&CK Techniques	The MITRE techniques associated with the alert. To access the MITRE page for this technique, either click the MITRE Technique entry or hover your cursor over the entry and, in the pop-up window that opens, click the "See Full Details at" link.

Related Topics

[Analytics: Configuring Expiration](#)

[Analytics: Configuring Country Watchlist](#)

[Analytics: Overview](#)

Analytics: Configuring Expiration

Open Alerts Expiration Configuration

1. From the main menu, choose **Configure > Detection > Alerts**.

The Alert Priorities Configuration page opens.

2. On the Alert Priorities Configuration page, click the **Alert Expiration** tab.

The Alerts Expiration Configuration page opens.

Configure Alert Expiration Days

Use the **Days before alerts expire** drop-down list to specify the number of days an alert remains in any given table before it is tagged as expired. If you configure an alert rule for Response Management with an action (email, syslog, webhook), all associated actions occur when the alert is opened or automatically closes.

Related Topics

[Analytics: Configuring Priorities](#)

[Analytics: Configuring Country Watchlist](#)

[Analytics: Overview](#)

Analytics: Configuring Country Watchlist

Open Alerts Country Watchlist Configuration

1. From the main menu, choose **Configure > Detection > Alerts**.

The Alert Priorities Configuration page opens.

2. On the Alert Priorities Configuration page, click the **Country Watchlist** tab.

The Alerts Country Watchlist Configuration page opens.

View Watched Countries

Use the Country Watchlist to trigger alerts for any countries that are listed on this page.

For each country that you want to add to the Country Watchlist, click its  (**Toggle**) icon so that the bar displays a blue bar (.

To add all the countries listed on this page to the Country Watchlist, click the **Select All** button. To remove all the countries listed on this page from the Country Watchlist, click the **Unselect All** button.

Related Topics

[Analytics: Configuring Priorities](#)

[Analytics: Configuring Expiration](#)

[Analytics: Overview](#)

Troubleshooting

Use the following troubleshooting information as needed.

Analytics jobs are lagging

In both of the following instances, the "Analytics performance has degraded" system alarm will be triggered.

The secondary Manager has been promoted to primary Manager

When you change the role of the primary Manager to that of the secondary Manager, and more than 5 hours has passed before the original primary Manager has been recovered and re-assigned to the primary role, the "Analytics performance has degraded" system alarm will be triggered. Analytics will recover and run the jobs that occurred during the last 6 hours, while the original primary Manager was down. Job performance will continue to lag until your system has processed all jobs from the last 6 hours and begins to process jobs in real time.

Your system is under-resourced for the amount of data being ingested

If your system is experiencing degradation (which is usually due to insufficient resources such as CPU or memory), jobs will begin to lag. If this lag exceeds 5 hours, then the "Analytics performance has degraded" system alarm will be triggered. At this point, results will be incomplete and unreliable.

A possible cause for this failure is that you have increased the flows per second beyond what is supported in your setup. To resolve this, either reduce the flows per second or increase the resources on the Manager, the Data Store, or both. If you cannot resolve the issue, contact [Customer Support](#).

Related Topics

[Analytics: Overview](#)

Detections Pack Management

An initial detections pack is included and activated when Secure Network Analytics is first installed or when you upgrade to a new release.



Make sure you've installed the latest detections pack. For more information, refer to the [Detections Pack Guide](#).

Use this page to review the detections in the current detections pack. You can also use this page to review and/or activate another detections pack.

Review the Detections in the Pack

This page displays with table showing detections for the activated pack. If you'd like to view detections for a different pack, do the following:

1. From the main menu, choose **Configure > Detection > Pack Management**.
2. In the **Select pack** field, choose a pack.



Review **What's new in this pack?** for information about the selected pack.

3. Review the detections in the pack. The table provides details for each detection within the pack.

Activate a Pack

To activate another pack, do the following:

1. From the main menu, choose **Configure > Detection > Pack Management**.
2. Choose a pack in the **Select pack** field.



Review **What's new in this pack?** for information about the selected pack.

3. Review the detections in the pack. The table provides details for each detection within the pack.
4. Click **Activate pack** to activate the selected pack.
5. Confirm the pack activates for each Flow Collector by reviewing the Flow Collector pack status table.



After you've activated a new pack, go to Response Management to enable findings rules for the detections in the pack. For more information, refer to the [Detections Pack Guide](#).

Operationalize Detections Pack

The detections from the Detections pack are now operationalized as specific alerts and observations within Secure Network Analytics. To review the alerts, go to **Configure > Alerts**.

For information about Detections Pack alerts and observations types and their descriptions, refer to the [Detections Pack Guide](#).

Related Tasks

[Response Management: Add Rules](#)

Related Topics

[Response Management Rules: Overview](#)

[Response Management: Overview](#)

Configuring Services



Only users assigned the role of Admin or Configuration Manager can configure services.

Use the Services page to configure settings for traffic that uses a set of TCP ports or UDP ports, or both. By doing so, you create definitions for the services that you expect your system to observe. Your Secure Network Analytics system uses these service definitions for look-up purposes only. These definitions do not restrict the Flow Collectors from monitoring or reporting on services that have not been configured here. For example, if your network observes an undefined TCP or UDP service, this service is tracked in a predefined "bucket" where services are grouped as "Undefined TCP" or "Undefined UDP."

How to find this page

From the main menu, click **Configure > System > Services**.

View your currently configured services

These fields in the Services table are described as follows.

Name (Required) The name of the service.

Port/Protocol (Required) When defining the port and protocol for the associated service, you can enter a single pair (e.g., 80/tcp) or a range (e.g., 80-300/tcp).

P2P In Edit mode, select the "Include this service in File Sharing Index calculations (P2P)" check box to include the associated service in calculations for the File Sharing Index.

Watch In Edit mode, select the "Trigger the Watch Port Active alarm when this service is seen (Watch)" check box to direct Secure Network Analytics to trigger the Watch Port Active alarm when the associated service occurs.

Web In Edit mode, select the "This service is used by a web server (Web) check box" to direct the Web server to use the associated service and to direct the Flow Collector to refrain from generating alarms for traffic that is considered normal for Web servers.

Exclude Worm In Edit mode, select the "Exclude this service from the worm detection algorithm (Exclude Worm)" check box to exclude the associated service from the Worm Detection algorithm.

Exclude Security Event In the Edit mode, select the "Allow host group policies to exclude this service from the security event detection algorithm (Exclude Security Event)" check

box to enable host group policies to exclude the associated service from the algorithm used to detect security events.

Exclude Long Flow In the Edit mode, select the "Exclude this service from the Suspect Long Flow detection algorithm (Exclude Long Flow)" check box to exclude the associated service from the Suspect Long Flow Detection algorithm.

Filter columns if you need to narrow your search

You can filter only the Name and Port/Protocol columns. At the top of each of these columns there is a filter field in which you can enter the criteria by which you want to filter that column. As you start to type your entry, the column dynamically filters to display a list of entries that contains any of the characters you have entered.

You can click the  (**Filter**) icon to see a list of available filter options (for example, Is equal to; Is not equal to; Contains) from which you can choose to help you filter the list according to a specific rule. The "Contains" rule is the default filter mode. When you select an option, the  (**Clear Filter**) icon is displayed.

Clearing entries

To clear the entry, click the **Clear Filter** icon.

Sort columns

You can sort the Name and Port/Protocol columns. To reverse the sort order for either of these columns, click the column name. Use the column name to toggle between ascending and descending order. You can also click the arrow that is displayed to the right of the column name (after you click the column name). An up arrow indicates that the column is sorted in ascending order, and a down arrow indicates descending order.

Add a service configuration

1. In the Services table, choose **Add New**.

The configuration fields for the service you are adding appear above the Services table.

2. (Required) Enter a name for the service.
3. (Required) Enter the port and protocol for the associated service. You can enter a single pair (e.g., 80/tcp) or a range (e.g., 80-300/tcp).
4. Select the desired check boxes in the Settings section. For descriptions of these check boxes, see [View your currently configured services](#).
5. To save your configuration, click **Save**.

Edit a service configuration

1. In the row that contains the service you want to edit, click the **⋮ (Ellipsis)** icon in the Action column and choose **Edit**.

The Services table expands to display the settings you have configured for the associated service.

2. Make the necessary edits, and when you have finished, click **Save**

Delete a service configuration

In the row that contains the service you want to delete, click the **Ellipsis** icon in the Action column and choose **Delete**.



Deletion of a service can invalidate custom security events and cause missing or false alarms. For more information, see the "When you delete a host group, service, or application" section in [Configuring Custom Policies \(Custom Security Events\)](#).

Custom Applications: Overview

Use this page to do the following:

- View a list of the applications that are not included in the default list provided by Secure Network Analytics.
- Add and define applications that are not included in the default list provided by Secure Network Analytics. You use these custom applications to monitor traffic that matches the criteria you define for each one. You can view, add, edit, delete, and re-order these applications.

Open Custom Applications

Do one of the following to access this page:

- From the main menu, choose **Configure > System > Applications**.
- You can also open this Web App page from the Desktop Client.
 1. Log in to the Desktop Client.
 2. From the main menu, choose **Configuration > Applications**.
 3. On the dashboards page that opens, choose **Tools > Custom Applications**.

Custom Applications: Table of Contents

Custom Applications: View Applications

You can view both user-defined (custom) and default applications to determine if this list contains the applications you need. If necessary, you can add applications to the Custom Applications section.

Before you navigate away from or close this page, you must click **Save** (located near the top of the page) to complete the submission of all changes made on this page. Click **Cancel** if you want to discard the changes you have made (but have not yet submitted).

- **Custom Applications:** This section contains a list of all user-defined custom applications. Each custom application is domain-specific and contains a group of criteria (a definition set) that identifies the application as unique. The system finds each application by searching through the list in order, from top to bottom, until a match occurs. If no match occurs, the system searches through the Secure Network Analytics applications. Hover over a custom application name to view the rules that apply to that application.
- **Secure Network Analytics Applications:** This section contains a list of the [built-in applications](#) that are in Secure Network Analytics. These applications are read-only and are listed in the order in which the system searches for these applications. The application at the top of the list contains the most specific criteria, and the application at the bottom contains the least specific criteria. You cannot reorder this list of applications.

You can reorder the custom applications, if desired. Simply click either the ▲ (**Move Up**) icon or the ▼ (**Move Down**) icon within the ▲ ▼ (**Move Up/Move Down**) icon in the row for the application you want to move (depending on if you want to move it up or down within the table).

Related Topics

[Custom Applications: Overview](#)

Custom Applications: Add an Application

Refer to the following information when adding and configuring applications that are not included in the default list provided by Secure Network Analytics:

- When defining the port and protocol, you can enter a single pair (e.g., 80/tcp) or a range (80-300/tcp).
- For the Server field, you can define an IP address or a host group.
 - When defining an IP address, valid formats include IPv4, IPv6, or CIDR notation.
 - You can enter an IP address or a range. (Some examples include 192.168.10.10, 192.168.10/16, 192.168.10-100.)
 - You can select only one host group.
 - You can use the Search box to quickly search for a host group.
- If you want to add one or more rules to further define the custom application, click **Add to Rules** (located under the DPI Classification field). Note the following:
 - Criteria within each rule, or block, are ANDed together, so all the criteria within a rule must be met for a rule to be a match.
 - Rules, or blocks, are ORed together; therefore, the search will return results that match one, some, or all of the rules.



Once you have configured a new application and you want to see if there is any traffic associated with it, access either the [Host Report](#) or the [User Report](#), and view the data in the Network Activity section.

Related Topics

[Custom Applications: Overview](#)

Custom Applications: Edit and Delete an Application

In the Actions column, click the **⋮ (Ellipsis)** icon to open the context menu, and then select the appropriate option.



Deletion of an application can invalidate custom security events and cause missing or false alarms. For more information, see the "When you delete a host group, service, or application" section in [Configuring Custom Policies \(Custom Security Events\)](#).

Related Topics

[Custom Applications: Overview](#)

Domain Configuration: Overview



Only users assigned the role of Admin or Configuration Manager can configure domains. Power Analysts can only view the domains.

Use the Domain Properties page to define settings that are associated with a domain. A domain is a grouping of hosts and other devices that you want to monitor and manage. Flow Collectors exist within domains, and you can have multiple domains within one Secure Network Analytics system. Domains are completely independent of other domains, and every domain contains the Host Group tree. For information about which host groups exist in the Host Group tree, see [Managing and Configuring Host Groups](#).

Opening Domain Properties

1. From the main menu, choose the applicable **[Current domain name]**.
2. From the main menu, choose **Configure > System > Domain Properties**.

Deleting a Desktop Client Domain

Use caution when deciding which Desktop Client domains you want to delete as you will lose access to all data which has been collected for the domain you are deleting.



Workaround: If you accidentally delete all of your domains in the Desktop Client and lock yourself out of the Manager Web App, create a new Non-Data Store domain in the Desktop Client. This will allow you to regain access into the Manager Web App. For information on creating a domain refer to the Add a Domain topic in the Desktop Client help.

Domain Configuration: Table of Contents

Domain Configuration: Configure Domain Settings

Complete the following settings for the domain you are adding.

Domain Name

Name for the domain you are currently in.

Archive Hour

Allows you to set the time at which each Flow Collector in the domain clears all counts. You can enter whole numbers between 0 and 23, where 0 is midnight in your local time zone. The local time zone is indicated to the right of the Archive hour field.

At the defined time, the Flow Collector resets all index counts to 0. In addition, the Flow Collector saves the log files and Web files that it has gathered during the preceding 24 hours and then begins a new day of data collection.

Internal autonomous system (AS) number

Click inside the Internal AS Numbers field and type your AS numbers. Separate multiple entries with commas or by pressing **Enter** after each entry to place each one on a separate line.

You can assign internal autonomous system (AS) numbers only to domains that contain [Flow Collector \(NetFlow\)](#) or [Flow Collector \(sFlow\)](#) appliances. When Secure Network Analytics encounters traffic containing these numbers in NetFlow or sFlow data, it categorizes the traffic as "origin" traffic on the Autonomous System Traffic document. Origin traffic signifies traffic from or within your network as opposed to traffic from an external network that is passing through your network (transit traffic).

For information about the Autonomous System Traffic document, see the "Autonomous System Traffic" topic in the Desktop Client help.

Export settings

The Export page on the Domain Properties dialog allows you to export specific domain content. You might want to do this so you can use the content as a template for any additional domains you add in the future.

1. Make your desired selections.

Refer to the following table for information about the available settings.

If you select this check box...	Secure Network Analytics exports this data...
Export All configurations*	All of the data listed in "Export the Domain configuration" below. In addition, a list of your flow collectors as well as your exporters and their interfaces are also exported.
Export the Host Group configuration*	The entire host group definition structure, including the host group names and IP address ranges. This output does not include policies.
Export the Domain configuration*	• Archive hour setting from the Domain Properties dialog. • All Service definitions. For information about services, see the "Services" topic in the Desktop Client help. • All Alarm Configuration settings. For information about configuring alarms, see the "About Alarm Severities" topic in the Desktop Client help. • The entire host group structure, including the host group names and IP address ranges. • All policies, including custom security events. Mitigation alarm actions are only exported when they have been manually changed from the defaults (set to <i>Not inherited</i>).

* You can use any of the XML files resulting from these commands to replace the host group configuration. For more information, see the "How to Replace the Host Group Configuration" topic in the Desktop Client help.

2. Click **Export**.

Secure Network Analytics saves the corresponding settings in an XML file that is downloaded to your Downloads folder.



Exporting is not the same thing as backing up. For more information, see the "Backing Up/Restoring Information" topic in the Desktop Client help.

Related Topics


[Domain Configuration: Overview](#)

[Managing and Configuring Host Groups](#)

[Investigating Host Groups](#)

Domain Configuration: Add a Domain

1. From the main menu, choose **[Current domain name] > Add Domain**.



2. Configure the following fields:

Domain Name: The name to be assigned to the domain. This name is shown on the Host Group tree.

Select Method: Select one of the methods described in the table below to designate which host group structure you want to use for the domain you are adding.

If you want to make this a Data Store domain, make sure you check the **Configure as a Data Store domain** check box.

If you select this method...	Then...
Default	Secure Network Analytics creates the domain with the default host group structure but without any Flow Collectors.
Import from File	Secure Network Analytics creates the domain and uses the appropriate configuration, based on the specific domain content you exported (host group, domain, or both). • XML files containing the domain configuration are not backwards compatible. These files are only compatible within the same system version number (for example, from Flow Collector v7.0 to Manager v7.0). • You can also import the entire host group configuration using the Host Group Management page.

If you select this method...	Then...
	• If you need to import interface groups in the Network Devices branch of the Host Group tree from another domain, use this option. You must first export the groupings as an XML file to your local drive. • None of the Flow Collectors contained in the XML file is imported.
 If you add a Flow Collector to an existing domain, that domain's specific configuration (policy, alarm severity, services, exporter SNMP, etc.) is applied to this Flow Collector.	

3. If you want to add a Data Store domain, check the **Configure as a Data Store domain** check box. A Data Store domain is for Secure Network Analytics systems that are using a Data Store and a Non-Data Store domain is for Secure Network Analytics systems that are not using a Data Store.
4. To save your configurations, click **Add**.
5. Complete the settings on this page. (See next section.)

Creating a Data Store Domain for Non-Data Store Users by Adding a Data Store Domain and Synchronizing

If you are in the process of transitioning a Non-Data Store Flow Collector to a Data Store Flow Collector and want to keep your data synchronized between your Non-Data Store domain and your newly created Data Store domain, follow the process outlined in [Domain Configuration: Synchronizing Domains](#). For more information on transitioning, refer to the [Secure Network Analytics System Configuration Guide](#).

Follow the steps below to add a new Data Store that you will be synchronizing with your Non-Data Store domain.

1. From the main menu, choose **[Current domain name] > Add Domain**.
2. Configure the following fields:

Domain Name: The name to be assigned to the domain. This name is shown on the Host Group tree.

Select Method: Select the Default method.

If you want to make this a Data Store domain, make sure you check the **Configure as a Data Store domain** check box.

3. To save your configurations, click **Add**.
4. From the main menu, choose **[Current domain name] > the Non Data Store domain that you want to synchronize with your Data Store domain**.
5. From the main menu, choose **Configure > System > Domain Properties**.
6. Select the **Edit** button.
7. Choose the Data Store domain that you want to synchronize this domain with in the **Target Domain to Synchronize** drop-down menu.
8. Click the **Save** button to save your changes.
9. Click the Synchronize button  that appears next to your Non-Data Store domain in the main menu to synchronize your domains.

Creating a Data Store Domain for Non-Data Store Users by Importing an Existing Domain (Optional)

If you are currently on a Non-Data Store domain and you want to add a Data Store domain to your Secure Network Analytics system for a future expansion into Data Store, you can do so by importing a Non-Data Store configuration into a new Data Store domain. When you import a existing domain, you won't have to re-configure items like alarms, host groups, and so on. Importing from an existing domain is like creating a new domain but with an existing configuration.

If the domain is newly created, then you will have to re-configure your Secure Network Analytics settings.

Follow the steps below to add a new Data Store domain and import all of its configuration from your Non-Data Store domain.

1. From the main menu, choose **[Current domain name] > the Non- Data Store domain that you want to synchronize with your Data Store domain**.
2. From the main menu, choose **Configure > System > Domain Properties**.
3. Make sure the **Export All configuration** radio button is selected. Refer to the [Domain Configuration: Configure Domain Settings](#) topic to view a list of the data that is exported.
4. Click the **Export** button to download the XML file.
5. From the main menu, choose **[Current domain name] > Add Domain**.
6. Enter a name for your new domain in the **Domain Name** field.

7. From the **Select Method** drop-down menu, select the **Import from File** option.
8. Select the XML file you downloaded in [step 4](#).
9. Check the **Configure as a Data Store domain** check box.
10. To add your new domain, click **Add**.

Related Topics

[Domain Configuration: Overview](#)


[Domain Configuration: Synchronizing Domains](#)

[Managing and Configuring Host Groups](#)

[Investigating Host Groups](#)

Domain Configuration: Edit a Domain

1. If you first need to access the domain, from the main menu, choose **[Current domain name]**.



2. From the main menu, choose **Configure > System > Domain Properties**.
The Domain Properties page opens.
3. In the upper right corner of the Domain Properties page, click **Edit**.
4. When finished editing, click **Save**.

Related Topics

[Domain Configuration: Overview](#)

[Managing and Configuring Host Groups](#)

[Investigating Host Groups](#)

Domain Configuration: Delete a Domain

! When you delete a domain, you will lose access to all data which has been collected for that domain. Make sure you only delete domains for which you no longer want to access the data collected by them.

Before you Begin


If your domain includes Flow Collectors, remove them from Central Management before you delete the domain. You can add the Flow Collectors to another domain, but the procedure includes resetting them to their factory defaults (RFD). For instructions, go to:

1. [Removing an Appliance from Central Management](#)
2. [Adding an Appliance to Central Management](#) (Optional)

! If you remove Flow Collectors from Central Management and delete the domain, you will lose the associated Flow Collector data.

Delete a Domain

1. If you first need to access the domain, from the main menu, choose **[Current domain name]**.



2. From the main menu, choose **Configure > System > Domain Properties**.
3. Click **Delete Domain**.

Related Topics

[Domain Configuration: Overview](#)

[Managing and Configuring Host Groups](#)

Investigating Host Groups

Domain Configuration: Synchronizing Domains

If you are in the process of transitioning a Non-Data Store Flow Collector to a Data Store Flow Collector, you may want to keep your configurations and tuning synchronized between your Non-Data Store domain and your Data Store domain. This topic describes the process for synchronizing your Non-Data Store domain with its associated Data Store domain.

Before You Begin

Ensure that you have already created a Data Store domain that you will be synchronizing with your Non-Data Store domain. If you have already followed the process described in [Adding a Data Store to a Non-Data Store Deployment and Transitioning Your Flow Collectors](#), your Data Store domain should already be created. Refer to [Domain Configuration: Add a Domain](#) for instructions on adding a domain.

Procedure

Follow these steps to synchronize your Non-Data Store domain (Source) with your Data Store domain (Target).

1. From the main menu, choose **[Current domain name] > the Non Data Store domain that you want to synchronize with your Data Store domain**.
2. From the main menu, choose **Configure > Domain Properties**.
3. Select the **Edit** button.
4. Choose the Data Store domain that you want to synchronize this domain with in the **Target Domain to Synchronize** drop-down menu.



You can only synchronize your target Data Store domain with one source Non-Data Store domain. If you attempt to synchronize your target Data Store domain with more than one source Non-Data Store domain, you will receive the following error, "This target domain is already in use by another Non-Data Store domain."

5. Click the **Save** button to save your changes.

Removing a Domain Synchronization Target Domain

Follow the steps below to remove a target domain.

1. From the main menu, choose **[Current domain name] > the Non Data Store domain that you want to synchronize with your Data Store domain**.

2. From the main menu, choose **Configure > Domain Properties**.
3. Select the **Edit** button.
4. Click the **Clear Target Domain** button.
5. Click the **Save** button to save your changes.

Configure Flow Collectors: Overview

Use this page to define the settings for a specific Flow Collector.

You need to use the Flow Collector Properties dialog in the Desktop Client to do the following:

- Create a Mitigation Allow List.
- Use non-deduplication groups to group interfaces (e.g., an exporter IP address and if-index combination) together that have data for which you want to prevent deduplication.

Open Flow Collector Configuration

From the main menu, choose **Configure > System > Flow Collectors**.

Flow Collector Configuration: Table of Contents

Configure Flow Collectors: Main Tab

Open Flow Collector Configuration: Main Tab

From the main menu, choose **Configure > System > Flow Collectors**.

The Main tab opens by default.

Global fields

The Flow Collector drop-down list contains all the names of the Flow Collectors configured for your Secure Network Analytics system.

The row underneath displays the name, IP address, model number, and serial number of the associated Flow Collector.

Data Collection

Use the Monitor port field to specify the network port number on the Flow Collector that receives data from devices that you configured to send traffic data to the Flow Collector. You can designate a single monitor port for a Flow Collector.

If you want the specified Flow Collector to accept flow traffic from any exporter, check the "Accept flows from any exporter" check box. In this case, new exporters automatically appear in the Exporters folder under the corresponding Flow Collector in the Enterprise tree. If you are using Interface groups, you must check this check box on each exporter associated with the Interface group.

Once all the exporters have been discovered and you do not want the Flow Collector to accept flow traffic from additional exporters, deselect the "Accept flows from any exporter" check box.



When the Flow Collector is placed outside a firewall, it is recommended that you turn off this setting.

System Alarms

Use this section to enable or disable the following alarms:

- Flow Collector Data Deleted
- Flow Collector Database Channel Down
- Flow Collector Database Unavailable
- Flow Collector Flow Data Lost**
- Flow Sensor Management Channel Down *

- Flow Sensor Time Mismatch *
- Flow Sensor Traffic Lost *
- Interface Utilization Exceeded Inbound**
- Interface Utilization Exceeded Outbound**

* These system alarms are available only for Flow Collector (NetFlow) appliances.

** These alarms are categorized as Traffic alarms.

You can change the severity of these system alarms on the **Configuration > Alarms** page.

Store Flow Interface Data

Use this section to specify the amount of flow interface data that the Flow Collector stores. Refer to the following descriptions of the available options:

- **None** The Flow Collector will not store flow interface data.
- **As much as possible** The Flow Collector will store enough flow interface data to consume all of the available disk space (e.g., in the /lancope/var partition). The amount of available disk space is displayed in the **Appliance Administration** (i.e., the Web UI) on the **Home** page. Information on how this space is consumed by the database is displayed on the **Support > Database Storage Statistics** page.
- **Up to (days)** The Flow Collector will store the flow interface data for the specified number of days. If this value causes the Flow Collector to consume all the available disk space, the oldest days of data are dropped on a first-in, first-out basis.

Number of days is based on calendar days, not days in the database.

 If you have deployed the Data Store on your system, the Store Flow Interface Data section on this page won't be displayed. You must configure the retention settings for your Data Store using the related API.

Alarm Configurations

Use this section to specify various settings for Flow Collector (NetFlow) and Flow Collector (sFlow) appliances. Secure Network Analytics uses these settings as defaults for newly discovered interfaces when the setting has not been manually defined or populated via SNMP for the interface on the Exporter Properties: Interfaces Page (in the Enterprise tree right-click an exporter and choose **Configuration > Properties**). Changes to these settings do not affect existing interfaces.

Refer to the following table for a description of the options in this section.

Setting	Description
Duration of export to trigger Longest Export Exceeded alarm	An alarm will be raised when the Flow Collector observes traffic from any exporter when the flow's duration exceeds the designated number of seconds.
Interface speed	
Inbound	The maximum speed for inbound traffic on the interface. This speed is used to calculate the inbound utilization percentages. The Manager uses this setting only if the setting has not been manually defined or populated via SNMP for the interface on the Exporter Properties: Interfaces Page (in the Enterprise tree right-click an exporter and choose Configuration > Properties).
Outbound	The maximum speed for outbound traffic on the interface. This speed is used to calculate the outbound utilization percentages. The Manager uses this setting only if the setting has not been manually defined or populated via SNMP for the interface on the Exporter Properties: Interfaces Page (in the Enterprise tree right-click an exporter and choose Configuration > Properties).
Utilization to trigger an Interface Utilization Exceeded alarm	
Inbound	The threshold (in percent) of the Inbound speed that the inbound utilization must exceed (for the minimum time specified) for the Interface Utilization Exceeded In alarm to be triggered. The Manager uses this setting only if the setting has not been manually defined or populated via SNMP for the interface on the Exporter Properties: Interfaces Page (in the Enterprise tree right-click an exporter and choose Configuration > Properties).

Setting	Description
Outbound	The threshold (in percent) of the Outbound speed that the outbound utilization must exceed (for the minimum time specified) for the Interface Utilization Exceeded Out alarm to be triggered. The Manager uses this setting only if the setting has not been manually defined or populated via SNMP for the interface on the Exporter Properties: Interfaces Page (in the Enterprise tree right-click an exporter and choose Configuration > Properties)
Duration of high utilization to trigger an Interface Utilization Exceeded alarm	
Inbound	The amount of time (in minutes) that the inbound utilization must exceed the specified threshold to trigger the Interface Utilization Exceeded In alarm. This threshold is checked every minute. The alarm becomes inactive if the inbound utilization falls below the threshold.
Outbound	The amount of time (in minutes) that the outbound utilization must exceed the specified threshold to trigger the Interface Utilization Exceeded Out alarm. This threshold is checked every minute. The alarm becomes inactive if the outbound utilization falls below the threshold.

Related Topics

[Configure Flow Collectors: Advanced Tab](#)

[Configure Flow Collectors: Overview](#)

Configure Flow Collectors: Advanced Tab

Open Flow Collector Configuration: Advanced Tab

1. From the main menu, choose **Configure > System > Flow Collectors**.

The Main tab opens by default.

2. Click the Advanced tab.

Broadcast List

Use this page to specify an unlimited number of IP addresses that are authorized to send broadcasts. Secure Network Analytics will not perform mitigation against hosts that are on the Broadcast List.

Ignore List

Use this page to specify an unlimited number of hosts or subnets for which the Flow Collector will completely ignore traffic. If the Flow Collector sees any traffic attributable to these IP addresses, it excludes that traffic from any graph or table.

Be certain that you can trust all traffic traveling to or from the hosts to be ignored. Secure Network Analytics will not analyze this traffic nor any that is spoofed to include any of these hosts. If an attack is launched on your network involving one of these hosts/subnets, the Flow Collector cannot report it.

Watchlist

Use this page to specify an unlimited number of hosts or subnets for which the Flow Collector will issue an alarm if those hosts or subnets become active. If the Flow Collector sees any traffic attributable to these IP addresses, it generates a Watch Host Active alarm.

Synchronize



- You must be logged in as the **admin** user to access this dialog.
- This feature is not available when pending Flow Collector property changes have not been applied yet.

At times you may need to synchronize the Flow Collectors on your network with your Manager. Some reasons may be as follows:

- You observe that your Flow Collectors are not working properly.
- You have added a new configuration (e.g., role policy, host group), but while you were doing so your Manager stopped communicating with any or all of your Flow Collectors. In this case, the policy will have been saved to your Manager, but it may not have propagated to any or all of your Flow Collectors.
- A Flow Collector is rebuilt, restored from a backup file, or reset to factory defaults.

During synchronization, Secure Network Analytics overwrites the configuration settings that exist on the Flow Collector with the configuration settings that exist on the Manager. If the Manager is restarted (with a restart services or reboot) in a regular, non-restore situation, the Manager will not send the configuration down to the Flow Collector.

Flow Collector Security Thresholds

Use this page to determine whether or not the Flow Collector ignores certain flows.

Refer to the following table for a description of the options in this section.

Setting	Description
Ignore flows between inside hosts	Enable this setting when it is not necessary to monitor data flow between inside hosts. For example, if the Flow Collector is placed outside your firewall.
Ignore flows between outside hosts	Enable this setting when it is not necessary to monitor data flow between outside hosts. For example, if the Flow Collector is placed inside your firewall.
Ignore flows to and from non-routable addresses	Enable this setting if you have network devices that allow data flows between non-routable addresses.
Ignore flows between inside hosts when calculating File Sharing Index	Enable this setting if you want the Flow Collector to ignore data flows between inside hosts when calculating the File Sharing Index (FSI).
Ignore Null0 flows	Enable this setting if you want the Flow Collector (NetFlow) to ignore Null0 flows. Null0 flows are sent by NetFlow-enabled devices after Null0 mitigation (i.e., the router sent the flow to an internal logical interface or ACL [access control list] mitigation).



This setting is available

Setting	Description
 only for Flow Collector (NetFlow) appliances.	
Seconds required to qualify a flow as long duration	The number of seconds that a flow must exist before the Flow Collector considers it a Suspect Long Flow. Lower settings increase the sensitivity for the Suspect Long Flow alarm. The range is 0-4,294,967,296 seconds, with a default setting of 32400 seconds (9 hours). You can exclude Instant Messaging (IM) technologies such as AOL AIM (port 5190), Yahoo Instant Messenger (TCP port 5050), and MSN Messenger (TCP port 1863) from Suspect Long Flow alarms by creating an Outside Hosts host group for the AOL, Yahoo, and MSN server ranges and then disabling the Suspect Long Flow alarm in that host group.
Suspect Long Duration Flow trust threshold	The number of long flow connections that are considered normal for hosts in your system. The threshold counter is counted per host, and is kept for each Inside and Outside host. The threshold range is 0-65534, with a default setting of 6. Setting the value to 0 disables trust threshold processing, and any long duration flow is considered suspect. This threshold does not tune or change. The host's behavior is learned, or more accurately, the system learns to ignore long flows for the host. It works this way: If this threshold is not 0, then if a host is seen with more long flows than the threshold during the first 7 days, any long flows for the host are ignored. After the 7 days of learning, if a host is seen with more long flows than the threshold, then the Suspect Long Flow

Setting	Description
	alarm is issued. The idea is that the system learns whether the normal behavior of the host is above or below this threshold. Then if a host that doesn't normally have a large number of long flows suddenly has more, an alarm is issued. However, hosts that normally have a large number of long flows will be ignored and do not cause false positives. The Suspect Long Flow alarm is designed to detect long duration connections between hosts that do not normally carry on long duration connections. For example, consider a user who connects his or her laptop to his or her home PC via a covert VPN for 3 days, but neither the home PC nor the laptop normally carry on long flow connections. The Flow Collector detects this abnormality and issues an alarm. You can exclude Instant Messaging (IM) technologies such as AOL AIM (port 5190), Yahoo Instant Messenger (port 1863), and MSN (port 5050) from Suspect Long Flow alarms by creating an Outside Hosts host group for the AOL, Yahoo, and MSN server ranges.
Seconds required to qualify a flow as Suspect Quiet Long Flow	The number of seconds that a flow must exist before the Flow Collector considers it a Suspect Quiet Long Flow Type, which is an IP communication between an Inside and an Outside host (with traffic in both directions). It exceeds the seconds required to qualify a flow as long duration setting and is suspiciously small.
Maximum number of bytes transferred to trigger a Suspect Quiet Long Flow alarm	The maximum number of bytes that are transferred between inside and outside hosts before the Flow Collector considers it a Suspect Quiet Long Flow Type.

Setting	Description
Minimum number of asymmetric flows per 5 minute period to trigger an Asymmetric Route alert	When the Flow Collector sees flows that indicate an IP conversation is happening, but sees only one side of the conversation (no packets from one host), it marks this flow as asymmetric. This is generally caused by visibility placement issues with spans or taps for the Flow Collector, such as dropped data, router cache problems, and actual asymmetric flows where a router needing to complete data collection is not sending flows. You can adjust the threshold to compensate for this data loss so that this behavior does not affect the collection process.
Minimum number of /24 subnets an infected host must contact before a Worm Activity or Worm Propagation alarm is triggered	Type the lowest number of /24 subnets that an infected host must contact to cause Secure Network Analytics to display a Worm Activity or Worm Propagation alarm. For example, if you type 4 in this field, Secure Network Analytics will not display either of these alarms until an infected host has contacted at least 4 Class C subnets. You can disable these alarms for specific services at the domain level.

Related Topics

[Configure Flow Collectors: Main Tab](#)

[Configure Flow Collectors: Overview](#)

Exporters and Interfaces: Overview

Exporters are network devices like routers, switches, and firewalls that send NetFlow data to a Flow Collector, and endpoint devices that send NVM telemetry to a Flow Collector. In some cases, Secure Network Analytics can use Flow Sensors to generate NetFlow for devices that don't support NetFlow, or where you want enhanced visibility, which can be accomplished using encrypted traffic analytics.

Use this page to view and define settings for exporters and Flow Sensors that are available for the Flow Collectors in your system.



- Only users assigned the role of Admin or Configuration Manager have access to this page and can configure exporters and Flow Sensors.
- Beginning in Secure Network Analytics (formerly Stealthwatch) v7.3.0, Exporter Hybrid Mode will no longer be available.

Open Exporters and Interfaces

From the main menu, choose **Configure > System > Exporters**.

The Exporters tab opens by default.

Exporters and Interfaces: Table of Contents

Exporters and Interfaces: Filter Columns

At the top of each of the columns there is a filter field in which you can enter the criteria by which you want to filter that column. As you start to type your entry, the column dynamically filters to display a list of entries that contains any of the characters you have entered.

You can click the  (**Filter**) icon to see a list of available filter options (for example, Is equal to; Is not equal to; Contains) from which you can choose to help you filter the list according to a specific rule. The "Contains" rule is the default filter mode. When you select an option, the  (**Clear Filter**) icon is displayed.

Clearing entries

To clear the entry, click the **Clear Filter** icon.

Related Tasks

[Exporters and Interfaces: Sort Columns](#)

Related Topics

[Exporters and Interfaces: Manage Exporters Overview](#)

[Exporters and Interfaces: Manage Interfaces Overview](#)

[Exporters and Interfaces: Overview](#)

Exporters and Interfaces: Sort Columns

To reverse the sort order for a column, click the column name. Use the column name to toggle between ascending and descending order. You can also click the arrow that is displayed to the right of the column name (after you click the column name). An up arrow indicates that the column is sorted in ascending order, and a down arrow indicates descending order.

Related Tasks

[Exporters and Interfaces: Filter Columns](#)

Related Topics

[Exporters and Interfaces: Manage Exporters Overview](#)

[Exporters and Interfaces: Manage Interfaces Overview](#)

[Exporters and Interfaces: Overview](#)

Exporters and Interfaces: Manage Exporters Overview

You can perform the following actions for exporters:

- [View exporter configurations.](#)
- [Add an exporter configuration.](#)
- [Edit a single exporter configuration.](#)
- [Bulk edit multiple exporter configurations.](#)
- [Delete an exporter configuration.](#)
- [Import a single exporter configuration.](#)
- [Export a single exporter configuration.](#)
- [Export all exporter configurations.](#)

Related Topics

[Exporters and Interfaces: Overview](#)

Exporters and Interfaces: View Exporter Configurations

By default, this table is sorted by IP address.

 You cannot view exporters in a failover Manager.

Field	Description
Flow Collector Name	Select from the drop-down list the Flow Collector whose list of associated exporters you want to view.
Exporter Name	Exporter name associated with the selected Flow Collector.
Exporter IP Address	IP address of the associated exporter that is sending data to the Flow Collector.
DNS Name	DNS Name of the associated exporter.
Type	Exporter type, which can be any of the following: • ASA Enhanced Firewall • ASA Firewall • Endpoint • Firewall • Flow Sensor • Virtual Flow Sensor • Exporter (There are certain exporters that Secure Network Analytics cannot categorize into any of the previously listed types because their exports do not contain exporter-specific fields unique to the exporter.) The Flow Collector auto-detects the

	exporter type. If you manually add an exporter, its type will be designated as <i>Exporter</i>. In addition to filtering exporter types to limit the number of exporters to view, you can also filter exporters within a Flow Collector instead of selecting each Flow Collector individually. Select from the drop-down list the Flow Collector whose list of associated exporters you want to view. Each exporter associated with that Flow Collector will be listed in a separate row with its related details. You can then filter on the exporters withing that Flow Collector.
SNMP Configuration	• Custom • Default • None • SNMP Profile Name
Actions	Use this context menu to edit or delete an exporter.

Related Tasks

[Exporters and Interfaces: Filter Columns](#)

[Exporters and Interfaces: Sort Columns](#)

Related Topics

[Exporters and Interfaces: Manage Exporters Overview](#)

[Exporters and Interfaces: Overview](#)

Exporters and Interfaces: Add an Exporter Configuration

1. Click the **⋮ (Ellipsis)** icon in the Action column and choose **Add New**.

The Add New Exporter page opens in which you need to define your configuration.

2. Refer to the following table as you complete your configuration.

Field or Option	Description
Select Flow Collector	The Flow Collector with which you want the exporter you are adding to be associated.
Exporter IP Address	IP address of the associated exporter. Use IPv4 formatting only. This name is unique within a Flow Collector.
Exporter Name	(Optional) Exporter name. You can use alpha or numeric characters and a maximum of 255 characters.
Ignore Netflow Egress	Check the check box if you want to specify that the Flow Collector (NetFlow) ignores egress traffic so that interface traffic statistics are not computed twice. This could result in erroneous reporting and consumption of additional flow rate licenses. Best practice NetFlow configuration recommends exporting only Ingress NetFlow for full traffic visibility.
SNMP Configuration section Select Type (choose one of the following methods):	The configured settings of the SNMP. From the drop-down list, choose one of the following options: • Custom Enables you to set up a customized SNMP configuration for this exporter. When you choose this option, additional fields appear for you to

	complete. A custom SNMP configuration can be used only by the specific exporter for which it is configured. • Default Assigns the SNMP configuration that is specified as the default profile (Default Config field on the Exporter SNMP Profiles tab on the Exporters page). • None Disables the SNMP polling for this exporter. • SNMP Profile Name Choose the SNMP configuration that you want enabled for this exporter. You can choose from any of the SNMP profiles you have configured (these profiles are listed on the Exporter SNMP Profiles tab on the Exporters page). An SNMP profile can be used by multiple exporters.

3. When finished, click **Save**.

You are returned to the main Exporter Configuration page. You can view the configuration you just added (and any others you have ever added) in the Configurations table, located in the bottom half of this page.

Related Tasks

[Exporters and Interfaces: Filter Columns](#)

[Exporters and Interfaces: Sort Columns](#)

Related Topics

[Exporters and Interfaces: Manage Exporters Overview](#)

[Exporters and Interfaces: Overview](#)

Exporters and Interfaces: Edit a Single Exporter Configuration

1. In the row that contains the exporter configuration you want to edit, click the **⋮ (Ellipsis) icon** in the Action column and choose **Edit**.

The Edit My Exporter page opens, displaying the settings you have previously configured.

2. Make any necessary edits. You can edit all fields for the exporter as well as any SNMP Configuration types or versions.
3. When finished, click **Save** when finished.



- You cannot edit the Exporter IP Address or any exporters that are of the following types: *Flow Sensor*, *Virtual Flow Sensor*, or *Endpoint*. You can edit these only in your Central Manager.
- You cannot edit exporters in a failover Manager.

Related Tasks

[Exporters and Interfaces: Filter Columns](#)

[Exporters and Interfaces: Sort Columns](#)

Related Topics

[Exporters and Interfaces: Manage Exporters Overview](#)

[Exporters and Interfaces: Overview](#)

Exporters and Interfaces: Bulk Edit Multiple Exporter Configurations

You can edit up to 10 exporters at one time.

1. Click within any cell in the row that contains the exporter you want to edit.

The Name and SNMP Configuration fields switch to Edit mode. These are the only two fields you can bulk edit.

2. Make any necessary edits to either one or both of these fields. You cannot choose the Custom option when editing the SNMP Configuration field. If you want to edit a custom configuration, you must do this through the editing process for a single exporter configuration. See [Edit a single exporter configuration](#).

Secure Network Analytics inserts a blue vertical bar at the beginning of each row that contains edits.

3. When finished editing the applicable exporters, click **Save**.
4. To edit another batch of exporters, repeat steps 1-3.



- You cannot bulk edit any exporters that are of the following types: *Flow Sensor*, *Virtual Flow Sensor*, or *Endpoint*. You can edit these only in your Central Manager.
- You cannot bulk edit exporters in a failover Manager.

Related Tasks

[Exporters and Interfaces: Filter Columns](#)

[Exporters and Interfaces: Sort Columns](#)

Related Topics

[Exporters and Interfaces: Manage Exporters Overview](#)

[Exporters and Interfaces: Overview](#)

Exporters and Interfaces: Delete an Exporter Configuration

In the row that contains the exporter configuration you want to delete, click the **Ellipsis** icon in the Action column and choose **Delete**.

All the interfaces under that exporter are removed. If any interfaces in this exporter are associated with any interface groups, then those interfaces are removed from those groups.



You cannot delete any exporters that are of the following types: *Flow Sensor*, *Virtual Flow Sensor*, or *Endpoint*. You can delete these only in your Central Manager.

Related Tasks

[Exporters and Interfaces: Filter Columns](#)

[Exporters and Interfaces: Sort Columns](#)

Related Topics

[Exporters and Interfaces: Manage Exporters Overview](#)

[Exporters and Interfaces: Overview](#)

Exporters and Interfaces: Import a Single Exporter Configuration

To import a single exporter configuration, in the row that contains the exporter configuration for which you want to overwrite with an imported configuration, click the context menu in the Actions column, and click **Import**.

Secure Network Analytics imports the following information for the requested exporter in .xml format:

The IP address in the XML file must match the exporter IP address for which you are importing the configuration information.

Secure Network Analytics imports the following information for the requested exporter in .xml format:

- Exporter configuration
- All associated interfaces and their configurations
- DSCP configuration associated with each interface, including the custom or DSCP profile name
- SNMP configuration associated with each exporter, including the custom or SNMP profile name

Related Tasks

[Exporters and Interfaces: Filter Columns](#)

[Exporters and Interfaces: Sort Columns](#)

Related Topics

[Exporters and Interfaces: Manage Exporters Overview](#)

[Exporters and Interfaces: Overview](#)

Exporters and Interfaces: Export a Single Exporter Configuration

To export a single exporter configuration, in the row that contains the exporter configuration you want to export, click the context menu in the Actions column, and click **Export**.

Secure Network Analytics exports the following information for the requested exporter in .xml format:

- Exporter configuration
- All associated interfaces and their configurations
- DSCP configuration associated with each interface, including the custom or DSCP profile name
- SNMP configuration associated with each exporter, including the custom or SNMP profile name

Related Tasks

[Exporters and Interfaces: Filter Columns](#)

[Exporters and Interfaces: Sort Columns](#)

Related Topics

[Exporters and Interfaces: Manage Exporters Overview](#)

[Exporters and Interfaces: Overview](#)

Exporters and Interfaces: Export all Exporter Configurations

Use this function if you want to export all the exporters that are observed within Secure Network Analytics. At times you may want to send this list of exporters to your network team so they can reconcile the devices that are missing or those that do not have NetFlow correctly configured.

If you have filtered for entries that contain specific settings and, for example, only five entries are displayed on the Exporters tab, but 400 exporters exist, Secure Network Analytics exports all of the 400 existing exporters.

Click the *** **(Ellipsis) icon** in the Action column and choose **Export to Excel**.

The Excel file (in .xlsx format) is automatically downloaded to your Downloads folder.

Related Tasks

[Exporters and Interfaces: Filter Columns](#)

[Exporters and Interfaces: Sort Columns](#)

Related Topics

[Exporters and Interfaces: Manage Exporters Overview](#)

[Exporters and Interfaces: Overview](#)

Exporters and Interfaces: Manage Interfaces Overview

You can perform the following actions for interfaces:

- [View interface configurations.](#)
- [Add an interface configuration.](#)
- [Edit an interface configuration.](#)
- [Bulk edit multiple interface configurations.](#)
- [Delete an interface configuration.](#)

Related Topics

[Exporters and Interfaces: Overview](#)

Exporters and Interfaces: View Interface Configurations

Use this table to view the settings for specific exporter and interface. You can determine the speed of inbound and outbound traffic passing through the interfaces of a particular exporter. You can also separately configure the threshold (in percent) for both the inbound and outbound traffic speed that must be surpassed for a particular interface to trigger an alarm.

1. For the applicable exporter, choose **View Interfaces** from the Actions column.

The Interfaces table associated with that exporter opens.

The number at the bottom right corner of the page shows the total number of exporters for the chosen Flow Collector.

2. To return to the Exporters page, click **Return to Exporters page** located above the top left corner of the table.

 You cannot view interfaces in a failover Manager.

Field or Option	Description
ID (ifIndex)	Unique identifier for the interface on the exporter. You can use this ID to perform flow analysis queries. It can be determined by any of the following ways: • Obtained from flow data. • Generated by the Flow Collector if missing from the flow data. • Added by the user. Enter a value matching the interface index ID from the exporter.
Name (optional)	The name assigned to the interface. You can enter it or allow Secure Network Analytics to discover it using SNMP polling.
Description (optional)	Text description of the interface. You can enter it or allow Secure Network Analytics to discover it.

Type (optional)	Type of interface. You can enter it or allow Secure Network Analytics to discover it.
DSCP Configuration	Indicates the type of template used to configure DSCP for the exporter. Manager settings can be inherited, which could be overwritten at the interface level. Options are as follows: • None - No template is applied. • Default - The default template for the domain is applied. • Select DSCP Profile - Select a template (that you have previously configured) for the domain. • Custom - Create a custom template for the domain.
Inbound Speed (in bps)	The maximum speed (line) for inbound traffic on the interface. You can enter it or allow Secure Network Analytics to discover it. If you have configured the Manager with the exporter's SNMP settings, the Manager calculates Utilization percentages for the interface using the line speed provided by the exporter. If not, the Manager uses the line speed configured here. If this value is blank, the Manager uses the "Inbound speed" configured on the Flow Collector Properties dialog.
Outbound Speed (in bps)	The maximum speed (line) for outbound traffic on the interface. You can enter it or allow Secure Network Analytics to discover it. If you have configured the Manager with the exporter's SNMP settings, the Manager calculates Utilization percentages for the interface using the line speed provided by the exporter. If not, the Manager uses the line speed configured here. If this value is blank, the Manager uses the "Outbound speed" configured on the Flow Collector Properties dialog.

Inbound Threshold (by %)	Secure Network Analytics triggers an alarm if the total inbound utilization of the interface exceeds the specified percentage. When you configure this setting, the Manager uses the threshold configured here (to determine when to raise alarms for the interface) rather than the default threshold configured on the Flow Collector Properties dialog.
Outbound Threshold (by %)	Secure Network Analytics triggers an alarm if the total outbound utilization of the interface exceeds the specified percentage. When you configure this setting, the Manager uses the threshold configured here (to determine when to raise alarms for the interface) rather than the default threshold configured on the Flow Collector Properties dialog.
Locked	A check mark in this column indicates that the SNMP polling feature is prevented from changing the configured values for the interface. If no check is displayed in this column, then the SNMP polling feature will update the configured interface values.

Related Tasks

[Exporters and Interfaces: Filter Columns](#)

[Exporters and Interfaces: Sort Columns](#)

Related Topics

[Exporters and Interfaces: Manage Interfaces Overview](#)

[Exporters and Interfaces: Overview](#)

Exporters and Interfaces: Add an Interface Configuration

Use this page to add new interfaces for the associated exporter that you may want to monitor.

1. For the applicable exporter, choose **View Interfaces** from the Actions column.

The Interfaces table associated with that exporter opens.

2. Click **Add New**.

The Add Interface page opens.

3. For descriptions of the fields in the upper section of this page, see [View Interfaces](#).

4. In the DSCP Configuration drop-down list, choose the desired option. You can choose from the following:

- Custom** When you choose this option, a grid is displayed that contains all the DSCP codes for which you can configure for the selected interface. Also, additional fields appear for you to complete (refer to the table at the end of this step to see descriptions for the fields that appear). A custom DSCP configuration can be used only by the specific interface for which it is configured.
- Default** Secure Network Analytics will use the default DSCP configuration, which is defined on the DSCP Configuration tab on the Manager Configuration page (the DSCP configuration for which you check the Make Default check box).
- None** This disables the DSCP configuration.
- Select DSCP Profile** Choose the DSCP profile that you want enabled for this interface. You can choose from any of the DSCP profiles you have configured on the DSCP Configuration tab on the Manager Configuration page. A DSCP profile can be used by multiple interfaces.

Field	Description
 Configure these settings when you choose Custom from the DSCP Configuration drop-down list.	
Inbound Allocation / Outbound	The amount of inbound (or outbound) traffic or the percentage of total bandwidth allocated to the associated code type. It can be determined by any of the following ways:

Allocation	• The sum of all the values specified by percentage in the Inbound Allocation column or Outbound Allocation column for a DSCP configuration cannot exceed 100%. • You can use the following units of measure when specifying a byte value: K, M, G, T, P, and E. • 1k = 1000. When you enter a number over 999, Secure Network Analytics changes that number to the nearest unit after you have clicked away from that field. <i>Example: 1000 changes to 1K; 1024 changes to 1.02K; 2,000,000,000 = 2 G</i> • You can enter the number of units (e.g., 1) with the unit of measurement immediately after it (no space). <i>Example: 10M, not 10 M</i> • You can enter decimal numbers only if you use a unit with it (e.g., 1.2M, not 1.2). • The sum of all the values specified by bps (bits per second) in the Inbound Allocation column or Outbound Allocation column for a DSCP configuration cannot exceed 9.22 E. • The sum of the inbound traffic cannot exceed the inbound line speed (the maximum speed for inbound traffic on the interface), and the sum of the outbound traffic cannot exceed the outbound line speed (the maximum speed for outbound traffic on the interface). ◦ If you have configured the Manager with the exporter's SNMP settings, the Manager calculates Utilization percentages for the interface using the line speed provided by the exporter. If not, the Manager uses the line speed configured here. If this value is blank, the Manager uses the "Inbound speed" value specified on the Flow Collector Properties dialog.
Inbound	Use one of the following units when specifying either the

Allocation Type / Outbound Allocation Type	Inbound or Outbound Allocation value: • % (percent) • bps (bits per second)

5. Click **Save**.

Related Tasks

[Exporters and Interfaces: Filter Columns](#)

[Exporters and Interfaces: Sort Columns](#)

Related Topics

[Exporters and Interfaces: Manage Interfaces Overview](#)

[Exporters and Interfaces: Overview](#)

Exporters and Interfaces: Edit an Interface Configuration

Use this page to edit a single interface's configuration settings.

1. For the applicable exporter, choose **View Interfaces** from the Actions column.
The Interfaces table associated with that exporter opens.
2. In the row that contains the interface you want to edit, click the **Ellipsis** icon in the Action column and choose **Edit**.
3. Make the necessary edits, and when you have finished, click **Save**.



- You cannot edit the name.
- You cannot edit interfaces in a failover Manager.

Related Tasks

[Exporters and Interfaces: Filter Columns](#)

[Exporters and Interfaces: Sort Columns](#)

Related Topics

[Exporters and Interfaces: Manage Interfaces Overview](#)

[Exporters and Interfaces: Overview](#)

Exporters and Interfaces: Bulk Edit Multiple Interface Configurations

You can edit up to 10 interfaces at one time.

1. For the applicable exporter, choose **View Interfaces** from the Actions column.

The Interfaces table associated with that exporter opens.

2. Click within any cell in the row that contains the interface you want to edit.

The Name, Description, Inbound Speed, Outbound Speed, Inbound Threshold, Outbound Threshold, and Locked fields switch to Edit mode. These are the fields you can bulk edit.

3. Make any desired edits to these fields.

Secure Network Analytics inserts a blue vertical bar at the beginning of each row that contains edits.

4. Repeat steps 2-3 for each interface you want to edit.
5. When finished, click **Save**.
6. To edit another batch of interfaces for the same exporter, repeat steps 2-5.



You cannot bulk edit interfaces in a failover Manager.

Related Tasks

[Exporters and Interfaces: Filter Columns](#)

[Exporters and Interfaces: Sort Columns](#)

Related Topics

[Exporters and Interfaces: Manage Interfaces Overview](#)

[Exporters and Interfaces: Overview](#)

Exporters and Interfaces: Delete an Interface Configuration

1. For the applicable exporter, choose **View Interfaces** from the Actions column.
The Interfaces table associated with that exporter opens.
2. In the row that contains the interface you want to delete, click the **Ellipsis** icon in the Action column and choose **Delete**.

Related Tasks

[Exporters and Interfaces: Filter Columns](#)

[Exporters and Interfaces: Sort Columns](#)

Related Topics

[Exporters and Interfaces: Manage Interfaces Overview](#)

[Exporters and Interfaces: Overview](#)

Exporter SNMP Profiles: Overview



- Only users assigned the role of Admin or Configuration Manager can manage exporter SNMP configurations.
- For optimal system performance, set SNMP polling to a 12-hour interval.
- Polling more frequently does not make your utilization metrics more up to date and may cause your system to run slower.

Use this page to view your existing exporter SNMP configurations as well as to create new ones. Secure Network Analytics uses the SNMP configuration that you select in the Default Config field to query the exporter(s) in the associated domain and obtain the names and speeds of the interfaces in your network. It does not obtain utilization information (utilization information comes from the NetFlow records that Secure Network Analytics collects).

Open Exporter SNMP Profiles

1. From the main menu, choose **Configure > System > Exporters**.

The Exporters tab opens by default.

2. Click the **Exporter SNMP Profiles** tab.

Exporter SNMP Profiles: Table of Contents

Exporter SNMP Profiles: Filter Columns

At the top of each of the columns there is a filter field in which you can enter the criteria by which you want to filter that column. As you start to type your entry, the column dynamically filters to display a list of entries that contains any of the characters you have entered.

You can click the  (**Filter**) icon to see a list of available filter options (for example, Is equal to; Is not equal to; Contains) from which you can choose to help you filter the list according to a specific rule. The "Contains" rule is the default filter mode. When you select an option, the  (**Clear Filter**) icon is displayed.

Clearing entries

To clear the entry, click the **Clear Filter** icon

Related Tasks

[Exporter SNMP Profiles: Sort Columns](#)

Related Topics

[Exporter SNMP Profiles: Manage SNMP Profile Configurations Overview](#)

[Exporter SNMP Profiles: Overview](#)

Exporter SNMP Profiles: Sort Columns

To reverse the sort order for a column, click the column name. Use the column name to toggle between ascending and descending order. You can also click the arrow that is displayed to the right of the column name (after you click the column name). An up arrow indicates that the column is sorted in ascending order, and a down arrow indicates descending order.

Related Tasks

[Exporter SNMP Profiles: Filter Columns](#)

Related Topics

[Exporter SNMP Profiles: Manage SNMP Profile Configurations Overview](#)

[Exporter SNMP Profiles: Overview](#)

Exporter SNMP Profiles: Manage SNMP Profile Configurations Overview

You can perform the following actions for SNMP exporters:

- [View exporter SNMP Configurations.](#)
- [Add an exporter SNMP Configuration.](#)
- [Edit an exporter SNMP Configuration.](#)
- [Delete an exporter SNMP Configuration.](#)

Related Topics

[Exporter SNMP Profiles: Overview](#)

Exporter SNMP Profiles: View Exporter SNMP Configurations

When you access this page, you can view the list of existing exporter SNMP configurations in the Configurations section. In the Global Options section, you can view the following currently configured settings. These settings apply to each of the exporter SNMP configurations.

Field	Description
Timeout (ms)	The maximum amount of time in milliseconds that the Manager waits when querying an exporter before aborting the attempt. • Minimum value: 1000 ms • Maximum value: 60,000 ms • Default value: 5000 ms This value applies globally to all exporters.
Retries	The number of attempts that the Manager makes after the first attempt to query an exporter. • Minimum value: 1 • Maximum value: 10 • Default value: 3 This value applies globally to all exporters.
Default Config	The configuration the Manager uses by default when querying an exporter. From the drop-down list, you can change this setting to another configuration or choose <i>None</i>. If you choose <i>None</i>, the default configuration no longer exists and the exporter SNMP polling is disabled.

For a description of the additional fields, see [Exporter SNMP Profiles: View Exporter SNMP Configurations](#).

Related Topics

[Exporter SNMP Profiles: Manage SNMP Profile Configurations Overview](#)

[Exporter SNMP Profiles: Overview](#)

Exporter SNMP Profiles: Add an Exporter SNMP Configuration

1. Click **Add New**.

The Exporter SNMP Configuration page opens in which you need to define your configuration.

2. Refer to the following table as you complete your configuration:

Field	Description
Name	The unique name you assign to the configuration.
Port	The port over which SNMP polling is performed. The default port is 161.
Polling every (minutes)	• The default interval is 720 minutes (12 hours). • When you edit this field, the polling interval restarts immediately for all interfaces. • You can enter a value between 0 and 525,600. • If you enter 0 (zero), polling is disabled.
SNMP Query Method (the following three methods are available): You can choose one or more of these methods.	The query method that the SNMP management system uses.
• Use ifxTable columns drop-down option	When you choose this option, the Manager's SNMP management system queries exporters for their interface configuration using the ifXTable. We recommend this setting if you have 10 Gbps interfaces.

• Use CatOS MIB drop-down option	When you choose this option, the Manager's SNMP management system queries exporters for their interface configuration using the CatOS MIB.
• Use PanOS MIB drop-down option	When you choose this option, the Manager SNMP management system decodes the Palo Alto Networks Pan-OS logical interfaces.
Version	The SNMP version used by the exporter.
Community (version 1 and 2c only)	The text string that the exporter uses to authenticate SNMP requests. The default SNMP community is <i>Public</i> .
If you are adding a configuration for version 3, you also need to configure the fields described below.	
Username	The unique name you assign to the user (the receiver).
Security Type	You can choose from the following options: • Authentication: Secure Network Analytics determines that messages were sent from a valid source. • Authentication & Encryption: Secure Network Analytics determines that messages were sent from a valid source and scrambles the packet contents to prevent access by unauthorized sources. • None: Secure Network Analytics does not perform any of the actions described in the previous two options.
Authentication Password	This password enables you to authenticate a user (8 characters minimum).

Authentication Protocol	Use this protocol to require user authentication. This protocol DOES NOT encrypt messages.
Encryption Password	This password enables you to encrypt and decrypt SNMP.
Encryption Protocol	Use this protocol to require encryption. This protocol DOES encrypt messages.

3. When finished, click **Save**.

You are returned to the main Exporter SNMP Configuration page. You can view the configuration you just added (and any others you have ever added) in the Configurations table, located in the bottom half of this page.

Related Topics

[Exporter SNMP Profiles: Manage SNMP Profile Configurations Overview](#)

[Exporter SNMP Profiles: Overview](#)

Exporter SNMP Profiles: Edit an Exporter SNMP Configuration

1. In the row that contains the exporter SNMP configuration you want to edit, click the **⋮ (Ellipsis)** icon in the Action column and choose **Edit**.

The Exporter SNMP table expands to display the settings you have previously configured.

2. Make the necessary edits, and when you have finished, click **Save**.

 You cannot edit the name.

Related Topics

[Exporter SNMP Profiles: Manage SNMP Profile Configurations Overview](#)

[Exporter SNMP Profiles: Overview](#)

Exporter SNMP Profiles: Delete an Exporter SNMP Configuration

In the row that contains the exporter SNMP configuration you want to delete, click the **Ellipsis** icon in the Action column and choose **Delete**.

Related Topics

[Exporter SNMP Profiles: Manage SNMP Profile Configurations Overview](#)

[Exporter SNMP Profiles: Overview](#)

SNMP Configuration Fields for Custom Option

Field	Description
Port	The port that the exporter uses when sending SNMP traps. The default port is 161.
Polling every (minutes)	- The default interval is 720 minutes (12 hours). - When you edit this field, the polling interval restarts immediately for all interfaces. - You can enter a value between 0 and 525,600. - If you enter 0 (zero), polling is disabled.
SNMP Query Method (the following three methods are available): You can choose one or more of these methods.	The query method that the SNMP management system uses.
Use ifxTable columns drop-down option	When you choose this option, the Manager's SNMP management system queries exporters for their interface configuration using the ifXTable. We recommend this setting if you have 10 Gbps interfaces.
Use CatOS MIB drop-down option	When you choose this option, the Manager's SNMP management system queries exporters for their interface configuration using the CatOS MIB.
Use PanOS MIB drop-down option	When you choose this option, the Manager's SNMP management system decodes the Palo Alto Networks Pan-OS logical interfaces.

Version	The SNMP version used by the exporter.
Community (version 1 and 2c only)	The text string that the exporter uses to authenticate SNMP requests. The default SNMP community is <i>Public</i> .
If you are adding a configuration for version 3, you also need to configure the fields described below.	
Username	The unique name you assign to the user.
Security	You can choose from the following options: • Authentication Secure Network Analytics determines that messages were sent from a valid source. • Authentication & Privacy Secure Network Analytics determines that messages were sent from a valid source and scrambles the packet contents to prevent access by unauthorized sources. • None Secure Network Analytics does not perform any of the actions described in the previous two options.
Authentication Password	This password enables you to authenticate a user.
Authentication Protocol	Use this protocol to require user authentication. This protocol DOES NOT encrypt messages. You can choose from the following protocols: • HMAC_MD5 • HMAC_SHA
Privacy Password	This password enables you to authenticate a

	user.
Privacy Protocol	Use this protocol to require user authentication. This protocol DOES encrypt messages. You can choose from the following protocols: • CBC_DES • CFB_AES_128

Related Topics

[Exporter SNMP Profiles: Manage SNMP Profile Configurations Overview](#)

[Exporter SNMP Profiles: Overview](#)

About Central Management

When you log in to an appliance for the first time, you will use the appliance console (SystemConfig) to configure the appliance and assign it to a Central Manager. The Central Management application is your Cisco Secure Network Analytics Manager (formerly Stealthwatch Management Console). For details, refer to the [System Configuration Guide](#).

Multiple Central Managers: You can configure more than one Central Manager in your system. However, each appliance can be managed by only one Central Manager.

Access: You need administrator privileges to access Central Management.

Open Central Management

1. Log in to your primary Manager.
(In your browser address field, type https:// and the appliance IP address. Press Enter.)
2. From the main menu, select **Configure > Global > Central Management**.

Inventory: Central Management opens to the [Inventory](#) tab.

Manage your Appliances: When your appliances are managed by Central Management, you can review their status and manage the following:

- [edit appliance configuration](#)
- [update software](#)
- [create backup configuration files](#)
- [view audit logs](#)
- [reboot](#)
- [shut down](#)
- [review license status](#)
- [update your flow collector database password](#)

Open Individual Appliances

Also, you can open an appliance from Central Management or by logging in to the appliance directly.

Use the appliance interface for the following:

- view system statistics
- additional appliance configuration options (not shown in Central Management)
- additional support options (not shown in Central Management)

Open an Appliance through Central Management

1. From the **Central Management Inventory**, click the **⋮ (Ellipsis)** icon for the appliance.
2. Select **View Appliance Statistics**.
3. Log in to the Appliance Administration interface.

Open Appliance Administration through Direct Login

To access the Appliance Admin through direct login, or if the appliance is disconnected from Central Management, use the following instructions.

1. In your browser address bar, type the appliance IP address as follows:

https://<IPAddress>

- **Manager:** add **/manager/index.html** after the IP address.
- **For example:** **https://1.1.1.1/manager/index.html**

2. Press Enter.
3. Log in to the appliance.

Related Topics

Central Management Inventory

Appliance Configuration

Appliance Support

Data Store Tab

Update Manager


App Manager

Database

Central Management Inventory

The Central Management inventory shows a summary of the appliances managed by your Manager in Central Management. You can review the [appliance status](#) for each appliance, as well as the appliance type, host name, and IP address.

- **Actions:** Use the [Actions menu](#) to access additional options for each appliance.
- **Filter:** Enter numbers, letters, or symbols to filter your inventory.
- **Sort:** Click a heading arrow to sort your list by the selected column. The arrow indicates the sort order.
- **Type:** If a Flow Collector has a **Data Store** tag, it is configured to send flows to your Data Store database. If a Flow Collector has a **Data Store Transition** tag, it is in the process of being transitioned from a non-Data Store Flow Collector to a Data Store Flow Collector. Refer to the [System Configuration Guide](#) for details.

Appliance Status	Host Name	Type
Connected	nflow-  -1	Flow Collector 

Appliance Status

Review the appliance status to make sure each appliance is shown as Connected or to take action if needed. For more information about each appliance status, hover your pointer over the status and refer to the following details.

Appliance Status	Description
Config Channel Down	If the appliance status is Config Channel Down, check the following: • Communication Settings: Confirm your network communication settings between the appliance and your Manager. • Trust Stores: Make sure your appliance identity certificates are saved to the correct Trust Stores. For instructions, refer to the SSL/TLS Certificates for Managed Appliances Guide. • Certificates: If you've changed the appliance identity certificate, check the procedure and

	confirm your certificates are saved to the correct Trust Stores. For instructions, refer to the SSL/TLS Certificates for Managed Appliances Guide. • Removing an Appliance: If you remove an appliance from Central Management while the configuration channel is down, make sure you also remove the appliance from System Configuration. For instructions, refer to Troubleshooting.
Connected	The appliance is connected to the Manager.  Before you change appliance configurations or install software updates and patches, make sure all appliances are shown as Connected.
Data Store Not Initialized	You need to finish your Secure Network Analytics system configuration. After you add all Managers and Data Nodes to your Central Management inventory, log in to your Manager appliance console (SystemConfig) and initialize the Data Store (Data Store > Initialization). Follow the instructions in the System Configuration Guide.
Data Store Not Configured	You need to finish your Secure Network Analytics system configuration. After you add all Managers, Flow Collectors, and Data Nodes to your Central Management inventory, log in to your Manager appliance console (SystemConfig) and configure your new appliances for secure communication with your Data Store (Data Store > New Appliances or Data Store > New Data Nodes). Follow the instructions in the System Configuration Guide.

Actions Menu

Click the **⋯ (Ellipsis)** icon for an appliance to access additional management options. Review each section below for details.

- [Edit the appliance configuration](#)
- [View appliance statistics](#)
- [Back up the appliance configuration \(Support\)](#)
- [View audit logs \(Support\)](#)
- [Reboot](#)
- [Shut down](#)
- [Remove the appliance from Central Management](#)
- [Open an Individual Appliance](#)

 If a menu is not available, check the appliance status.

Open the Central Management Inventory

When you log in to Central Management, the Inventory tab opens.

1. Log in to your primary Manager.
(In your browser address field, type `https://` and the IP address of the virtual appliance. Press Enter.)
2. Select **Configure > Global > Central Management**.

Open an Individual Appliance

You can open an appliance from Central Management for the following:

- view appliance statistics
 - additional appliance configuration options (not shown in Central Management)
 - additional appliance support options (not shown in Central Management)
1. Click the **⋯ (Ellipsis)** icon for the appliance.
 2. Select **View Appliance Statistics**.
 3. Log in to the appliance.

 To log in to an appliance directly, refer to [Open Appliance Administration through Direct Login](#).

Review Appliance Statistics

To see system statistics, services, disk usage, and docker services, log in to the appliance:

1. Click the **⋮ (Ellipsis)** icon for the appliance.
2. Select **View Appliance Statistics**.
3. Log in to the appliance.

Edit Configuration Settings

1. Click the **⋮ (Ellipsis)** icon for the appliance.
2. Select **Edit Appliance Configuration**.
3. See [Appliance Configuration](#) for details.

Back up Configuration

1. Click the **⋮ (Ellipsis)** icon for the appliance.
2. Select **Support**.
3. See [Backup Configuration Files](#) for details.

View Audit Logs

1. Click the **⋮ (Ellipsis)** icon for the appliance.
2. Select **Support**.
3. See [Audit Logs](#) for details.

Reboot

Use the following instructions to select an appliance from the Central Management inventory and reboot it. When you reboot an appliance, there is a brief period during the restart where data is not collected.



Make sure the appliance status is shown as Connected before you reboot an appliance. Do not reboot an appliance while configuration changes are pending.

1. Click the **⋮ (Ellipsis)** icon for the appliance.
2. Select **Reboot Appliance**.
Wait a few minutes before you use the appliance again.

Shut Down

Use the following instructions to select an appliance from the Central Management inventory and power it off. Make sure you have access to the appliance to start it again. Data cannot be collected while the appliance is shut down.



Refer to the [System Configuration Guide](#) for information about shutting down the Data Store database before you reboot or shut down a Data Node.



If you select Shut Down, you will turn off the appliance. To start a hardware appliance again, you must have physical access to it. To start a virtual appliance again, you must have administrative access to it.

1. Click the **⋮ (Ellipsis)** icon for the appliance.
2. Select **Shut Down Appliance**.

Related Tasks

[Adding an Appliance to Central Management](#)

[Removing an Appliance from Central Management](#)

Related Topics

[About Central Management](#)

[Domain Configuration: Overview](#)

Adding an Appliance to Central Management

Use the appliance console (SystemConfig) to add a stand-alone appliance to Central Management. It is important to review the following:

- **Appliance Identity Certificates:** When you add an appliance to Central Management using the appliance console (SystemConfig>Recovery>Add Appliance), the appliance identity certificate is automatically added to the Manager trust store.
- **Custom Certificates:** If your appliance has custom certificates, make sure you save the identity certificate and certificate chain (root and intermediate). Note that if you make network changes, your custom certificate will be replaced with a Cisco self signed certificate. Follow the instructions in the [SSL/TLS Certificates for Managed Appliances Guide](#) or contact [Cisco Support](#) for assistance.
- **Manager Administration Credentials:** You need the Manager user ID and password to add an appliance to Central Management.
- **RFD:** If you reset the factory defaults on an appliance, you will enter the host name, domain name, and additional configuration information as part of the appliance console (SystemConfig>Recovery>Factory Defaults). After you reset the factory defaults on the appliance, you will need to follow the First Time Setup process to add a new appliance. For more information, refer to the [System Configuration Guide](#).
- **Host Names:** A unique host name is required for each appliance. We cannot configure an appliance with the same host name as another appliance. Also, make sure each appliance host name meets the Internet standard requirements for Internet hosts.
- **Flow Collectors:** If you removed a Flow Collector from Central Management, it is also removed from the domain. You need to reset the factory defaults (RFD) if you plan to add it to a different domain (go to [Resetting Factory Defaults](#) for instructions). After you reset the factory defaults on the Flow Collector, you will need to follow the First Time Setup process to add a new Flow Collector. Follow the instructions in [Add an Appliance to Central Management](#) to select a Central Manager and a domain.
- **New Installations:** If this is a new installation and the appliance hasn't been configured, use the [System Configuration Guide](#) to install and configure your appliances correctly.



If your appliance has custom certificates, make sure you save the identity certificate and certificate chain (root and intermediate). Note that if you make network changes, your custom certificate will be replaced with a Cisco self signed certificate. Follow the instructions in the [SSL/TLS Certificates for Managed Appliances Guide](#) or contact **Cisco Support** for assistance.

Add an Appliance to Central Management

Use the following instructions to add a stand-alone appliance to your primary Manager/Central Manager. For more information, refer to the [System Configuration Guide](#).

1. Log in to your appliance console (SystemConfig).
2. Select the **Recovery** menu.
3. Select **Add Appliance**.
4. Follow the on-screen prompts to enter the Manager administration credentials and finish the configuration. Depending on the type of appliance, you may need to enter additional information (such as selecting domain or entering the Flow Collection port).



For troubleshooting, refer to **Central Management Inventory**. For more information about the appliance console (SystemConfig), refer to the [System Configuration Guide](#).

Related Tasks

Removing an Appliance from Central Management

Related Topics

About Central Management

Domain Configuration: Overview

Removing an Appliance from Central Management

Use the following instructions to remove an appliance from your Central Manager.

1. From the main menu, select **Configure > Global > Central Management**.
2. Click the **⋮ (Ellipsis)** icon for the appliance.
3. Select **Remove This Appliance**.
 - **RFD:** If you need to reset factory defaults (RFD) on an appliance before you add it to another Central Manager, go to [Resetting Factory Defaults](#).
 - **Flow Collectors:** If you removed a Flow Collector from Central Management, it is also removed from the domain. You need to reset the factory defaults (RFD) if you plan to add it to a different domain. Go to [Resetting Factory Defaults](#) and [Adding an Appliance to Central Management](#) for instructions.
 - **Central Management:** To add the appliance to a different Central Manager, log in to the appliance console (SystemConfig>Recovery>Add Appliance). For instructions, go to [Adding an Appliance to Central Management](#).



If your appliance has custom certificates, make sure you save the identity certificate and certificate chain (root and intermediate). Note that if you make network changes, your custom certificate will be replaced with a Cisco self signed certificate. Follow the instructions in the [SSL/TLS Certificates for Managed Appliances Guide](#) or contact [Cisco Support](#) for assistance.

Troubleshooting

If you log in to the appliance and it is not removed from Central Management, or the configuration channel is down, remove the appliance from System Configuration.

1. Log in to the appliance console as sysadmin.
2. Type **SystemConfig**.
3. Select **Recovery > Remove Appliance**.
4. Follow the on-screen prompts to remove the appliance from Central Management.

Resetting Factory Defaults

Use the following instructions to reset an appliance to its factory defaults (RFD). To completely erase data, make sure you reset factory defaults twice.

- **RFD twice:** To completely erase data, make sure you reset factory defaults twice.
- **Back up Configuration:** If you plan to restore the appliance configuration, make sure you save the backup configuration and database backup files. Refer to [Backup Configuration Files](#) and [Backup/Restore Database \(Appliance Admin Interface\)](#) for details. To restore the backup after RFD, contact [Cisco Support](#).



If you reset factory defaults (RFD) on an appliance, all existing data and configuration information will be deleted and can only be restored if you've made a backup.



If you reset an appliance to factory defaults, you cannot restore the configuration using Central Management. For assistance, please contact [Cisco Support](#).

1. Log in to the appliance console as sysadmin.
2. Select **Recovery > Factory Defaults**.
3. Follow the on-screen prompts to reset factory defaults and restart the appliance.



Make sure you RFD each appliance twice to completely erase data.

4. Log in to the appliance console as **sysadmin** and follow the on-screen prompts to configure the appliance IP address, host name, and domain. For instructions, refer to your [System Configuration Guide](#). This step is required even if you preserve network settings when you RFD.

Host Names: A unique host name is required for each appliance. We cannot configure an appliance with the same host name as another appliance. Also, make sure each appliance host name meets the Internet standard requirements for Internet hosts.

5. Log in to the appliance console (SystemConfig>Recovery) and add the appliance to Central Management. For instructions, refer to [Adding an Appliance to Central Management](#).

Related Tasks

[Adding an Appliance to Central Management](#)

Related Topics

[About Central Management](#)

[Domain Configuration: Overview](#)

Appliance Configuration

The Appliance Configuration page includes three tabs for configuring the appliance you selected in the [Central Management Inventory](#).

Best Practices

We recommend the following:

- **Before you make configuration changes**, review the [Central Management Inventory](#) to confirm the Appliance Status is **Connected**.
- **During configuration changes**, some changes require a system reboot. If you prefer to wait, you can revert your changes and edit your configuration settings and reboot later. Do not force the appliance to reboot while configuration changes are pending.
- **After your changes are applied**, make sure the appliance finishes the configuration changes and the Appliance Status returns to **Connected** (review the [Central Management Inventory](#)).
- **One at a time**: Configure one appliance at a time. Confirm the Appliance Status is **Connected** before you start configuring the next appliance in your cluster (review the [Central Management Inventory](#)).

Open the Appliance Configuration Page

1. From the main menu, select **Configure > Global > Central Management**.
2. From the Inventory tab, click the **⋮ (Ellipsis)** icon for the appliance.
3. Select **Edit Appliance Configuration**.
4. Select the **Appliance** tab, **Networks Services** tab, or the **General** tab.

Appliance Tab

[Advanced Intrusion Detection Environment \(AIDE\)](#)

[Network Interfaces](#)

[SSH](#)

[Sysadmin User](#)

[Host Naming](#)

[SSL/TLS Appliance Identity](#)

[Additional SSL/TLS Client Identities](#)

Network Services Tab

[SNMP Agent](#)

[NTP Server](#)

[Internet Proxy](#)

[Audit Log Destination](#)

[Local Resolution](#)

[DNS Server](#)

[DNS Cache](#)

General Tab

[Opening Message](#)

[Protected Sessions Time-Out](#)

[Compliance Mode](#)

[SMTP Configuration](#)

[Backup Configuration Encryption](#)

[Password Policy](#)

[Session Settings](#)

[DoDIN Notifications](#)

[Trust Store](#)

External Services

[Cisco Security Cloud Control \(SCC\)](#)

[Customer Success Metrics](#)

[Threat Feed](#)



For more appliance configuration options, log in to the [Appliance Administration](#) interface.

AIDE

The Advanced Intrusion Detection Environment (AIDE) is a host baselining system that detects modifications of critical files on a system.

When it is enabled, AIDE runs an audit of the current system once a day. It compares the hash sum, permissions, and time accessed of each monitored file on the current file system against the values stored in the appliance database.

Monitored Directories and Files

AIDE monitors static scripts, binaries, libraries, and information files that are only to be executed or read and never modified. It will monitor all files in these directories:

- /bin
- /boot
- /lib
- /lib64
- /sbin
- /usr/bin
- /usr/lib
- /usr/sbin
- /lancope/admin/bin
- /lancope/admin/lib
- /lancope/admin/snmp
- /lancope/bin
- /lancope/info
- /opt/napatech3
- /opt/vertica/agent
- /opt/vertica/bin
- /opt/vertica/config
- /opt/vertica/include
- /opt/vertica/install_params
- /opt/vertica/java
- /opt/vertica/lib
- /opt/vertica/lib64

- /opt/vertica/oss
- /opt/vertica/packages
- /opt/vertica/sbin
- /opt/vertica/scripts
- /opt/vertica/sdk
- /opt/vertica/share
- /opt/vertica/spread

The following files are not monitored because they can be modified:

- /lancope/info/patch
- /lib/init/rw

Open AIDE

Use the following instructions to open AIDE for a selected appliance.

1. From the main menu, select **Configure > Global > Central Management**.
2. Click the **⋮ (Ellipsis)** icon for the appliance.
3. Select **Edit Appliance Configuration**.
4. Select the **Appliance** tab.

Enable AIDE

1. Locate the **Advanced Intrusion Detection Environment** section.
2. Check the **Enable AIDE** check box.
3. Click **Apply Settings**.
4. The appliance initializes a database within 1 hour.

Review AIDE Audit Logs

Use the Audit Logs to review your AIDE data.

1. From the main menu, select **Configure > Global > Central Management**.
2. Click the **⋮ (Ellipsis)** icon for the appliance.
3. Select **Support**.
4. Select the **Audit Logs** tab.
5. In the [Category](#) drop-down, select **Management**.
6. Click **Search**.

Audit Log Destination

Use Audit Log Destination to configure a secure destination for audit log messages using syslog over TLS (Transport Layer Security).

Open Audit Log Destination

Use the following instructions to open Audit Log Destination for a selected appliance.

1. From the main menu, select **Configure > Global > Central Management**.
2. Click the **⋯ (Ellipsis)** icon for the appliance.
3. Select **Edit Appliance Configuration**.
4. Select the **Network Services** tab.

1. Configure Audit Log Destination

1. From the main menu, select **Configure > Global > Central Management**.
2. From the Inventory tab, click the **⋯ (Ellipsis)** icon for the appliance.
3. Select **Edit Appliance Configuration**.
4. Select the **General** tab.
5. Scroll to the **Trust Store** section.
6. Click **Add New** to add the root Certificate Authority (CA) SSL certificate that anchors your syslog server certificate chain, which supports **Syslog over TLS**. (For example: a syslog server running Splunk.)



For more information about adding certificates to the trust stores, refer to the [SSL/TLS Certificates for Managed Appliances Guide](#).

7. If your syslog server identity certificate does not include the Subject or Subject Alternative Name field that matches the value you will configure in Step 10, add the certificate to the trust store.
8. On the **Network Services** tab, locate the **Audit Log Destination** section.
9. Click **Add New** to enter the server name or IP address.



A maximum of two servers can be added.

10. In the **Server Name or IP Address** field, enter the server name (FQDN), IPv4, or IPv6 address of the remote syslog server.

11. In the **Destination Port** field, enter the port number that the appliance uses to provide audit trail information.

Destination Port Default: 6514



Entering the same IP address and server name for the same port leads to configuration issues.

12. In the **Certificate Revocation** field, select one of the following:
 - **Disabled:** Do not check the Syslog server revocation status. Always allow the TLS connection.
 - **Soft Fail:** Connect Syslog over TLS if the Syslog server revocation status is unavailable.
 - **Hard Fail:** Do not connect Syslog over TLS if the Syslog server revocation status is unavailable.

13. Click **Apply Settings**.

If the delivery fails, check the [Audit Log](#). Also, check the trust store and confirm the syslog SSL/TLS certificate is shown.

2. Enable SyslogCompliance

For Common Criteria (CC) compliance information, refer to the *Common Criteria Administrative Guide*.

If Syslog is enabled, all logs related to compliance will be saved in the remote Syslog server. Use the following instructions to establish Syslog server compliance for the selected appliance.

1. SSH into the appliance, or log into the appliance console.
2. Log in as **sysadmin**.
3. Type **SystemConfig**. Press Enter.
4. Select **Security** menu.
5. Select **Syslog Compliance**. Press Enter.



Make sure to read the status description at the top of the dialog box confirming whether you are enabling or disabling.

6. Select **Enable**. You may need to press the space bar on your keyboard to select it.
7. Select **OK**.

Backup Configuration Encryption

Backup Configuration Encryption is disabled by default. If you choose to enable Backup Configuration Encryption, make sure to enable it on every appliance in the system. Also, make sure to save your passwords to a secure location.

Open Backup Configuration Encryption

Use the following instructions to open Backup Configuration Encryption for a selected appliance. For more information about backup configuration files, refer to [Backup Configuration Files](#).

1. From the main menu, select **Configure > Global > Central Management**.
2. From the Inventory tab, click the **⋮ (Ellipsis)** icon for the appliance.
3. Select **Edit Appliance Configuration**.
4. Select the **General** tab.

Configure Encryption for Backups

Use the following instructions to enable Backup Configuration Encryption on each appliance.

Default: Backup Configuration Encryption is disabled by default.

1. From the main menu, select **Configure > Global > Central Management**.
 2. From the Inventory tab, click the **⋮ (Ellipsis)** icon for the appliance.
 3. Select **Edit Appliance Configuration**.
 4. On the **General** tab, locate the Backup Configuration Encryption section.
 5. Check the **Enable Encryption** check box to enable backup configuration encryption.
 6. In the **Backup Configuration Password** field, enter your password. This password is used to encrypt and to then decrypt backup configuration files.
 7. In the **Confirm Password** field, re-enter your password.
- **Encrypted Files** The file names of encrypted backup files contain the gpg extension.
 - **Restoring Files:** For details about restoring backup files, refer to [Restore an Appliance Backup Configuration File](#).



The password used to encrypt a file is required to restore it. While you can change the password that will apply to new backups, you will need to know the original password used to encrypt a file in order to decrypt and restore the file. Make sure to save your passwords to a secure location.

Cisco Security Cloud Control

Enabling this option registers your device in the Cisco Security Cloud Control (SCC). Smart Licensing is required for automatic registration, or you can manually register on the [Cisco XDR configuration page](#).

The following integrations require Security Cloud Control to be enabled:

- [Customer Success Metrics](#)
- [Cisco XDR](#)

DoDIN Notifications

Use the following instructions to configure SMTP and enable DoDIN Notifications. Before you configure these settings, follow the instructions in the *DoDIN Military Unique Deployment Guide* or the *Common Criteria Administrative Guide*.

The compliance notifications include the following:

- User management changes (creation, updating, deleting, enabling/disabling, password changes, etc.)
- Syslog failures
- AIDE detection information



To configure Response Management for non-compliance-related notifications (system alarms, etc.), refer to [Configure SMTP for Response Management](#).

Open DoDIN Notifications

Use the following instructions to open DoDIN Notifications for a selected appliance.

1. From the main menu, select **Configure > Global > Central Management**.
2. From the Inventory tab, click the **⋮ (Ellipsis)** icon for the appliance.
3. Select **Edit Appliance Configuration**.
4. Select the **General** tab.

Configure and Enable DoDIN Notifications

Before you configure these settings, follow the instructions in the *DoDIN Military Unique Deployment Guide* or the *Common Criteria Administrative Guide* and note the following:

- Your SMTP server identity certificate from your Certificate Authority (CA) needs to have a Subject or Subject Alternative Name that matches the IP address or host name you configure in step 7.
 - Make sure you configure these settings on every appliance in your Secure Network Analytics cluster.
1. From the main menu, select **Configure > Global > Central Management**.
 2. From the Inventory tab, click the **⋮ (Ellipsis)** icon for the appliance.
 3. Select **Edit Appliance Configuration**.
 4. Select the **General** tab.

5. Add the SMTP server certificate chain (root and intermediate certificates) to the appliance [trust store](#). Upload each part of the chain individually.
6. If you are using a host name to configure the SMTP server in step 7, configure the [DNS server](#).
7. In the SMTP Configuration section, complete the following fields:

Field	Details
SMTP Server	Enter the IP address or the host name (fully qualified domain name).
Port	Enter the port designated for secure SMTP communication.
From Email	Enter the source email address for your SMTP notifications.
User Name (optional)	Enter user name from your SMTP server.
Password (optional)	Enter password from your SMTP server.

8. Select an **Encryption Type**.

Encryption Type	Details
SMTPS	SMTP over SSL: Do not send notifications if the encryption is unavailable.
STARTTLS	Start TLS: Do not send notifications if the encryption is unavailable.
UN-ENCRYPTED	Do not check the encryption type. Always allow connection.

9. In the DoDIN Notifications section, click the **Enable** check box.
10. In the **To Email** field, enter the recipient email address.

DoDIN and Common Criteria notifications will be sent to this email address.

11. Click **Apply Settings**.
12. Repeat [these instructions](#) on every appliance in your cluster.

Compliance Mode

Use Compliance Mode to enable encryption libraries for FIPS and/or Common Criteria. To enable Common Criteria Encryption Libraries, follow the instructions in the *Common Criteria Administrative Guide* or the *Department of Defense Information Network (DoDIN) Military Unique Deployment Guide*.



Do not enable Common Criteria Encryption Libraries without following the instructions in the *Common Criteria Administrative Guide* or the *DoDIN Military Unique Deployment Guide*.

FIPS

The United States Federal Information Processing Standard (FIPS) defines security and interoperability requirements for computer systems that are used by the following:

- United States federal government agencies
- Federal contractors
- Other organizations that process information using a computer or telecommunications system on behalf of the federal government to accomplish a federal function

When you enable FIPS Encryption Libraries, the appliance is configured to comply with FIPS standards, ensuring the use of secure algorithms and meeting encryption requirements.



By default, FIPS Encryption Libraries is disabled. Make sure you enable FIPS Encryption Libraries on every appliance in the system so they can communicate.

Common Criteria

Common Criteria (CC) is used with FIPS. CC is an international standard for computer security certification. Refer to ISO/IEC 15408 for details.

Before you enable Common Criteria Encryption Libraries in Secure Network Analytics, follow the instructions in the *Common Criteria Administrative Guide*. It is important to follow the instructions to set up Common Criteria and enable encryption libraries correctly.



Do not enable Common Criteria Encryption Libraries without following the instructions in the *Common Criteria Administrative Guide*.

Open Compliance Mode

Use the following instructions to access Compliance Mode for a selected appliance.

1. From the main menu, select **Configure > Global > Central Management**.
2. From the Inventory tab, click the ***** (Ellipsis)** icon for the appliance.
3. Select **Edit Appliance Configuration**.
4. Select the **General** tab.

Enable FIPS Encryption Libraries

Use the following instructions to enable FIPS Encryption Libraries on every appliance.

Default: FIPS Encryption Libraries is disabled by default.

Recommendation: Before you get started, review the instructions for enabling FIPS Encryption Libraries. You will enable FIPS Encryption Libraries on every appliance in the system. The procedure includes updating the appliance software, installing FIPS 140 compliant certificates, and restarting the appliances. It is important to enable FIPS Encryption Libraries at a time that will cause the least amount of disruption.

Common Criteria: If you are enabling FIPS and Common Criteria Encryption Libraries, there are additional instructions to meet FIPS and Common Criteria requirements. Refer to the *Common Criteria Administrative Guide* for details.

The following features are not available when FIPS Encryption Libraries is enabled: TACACS, Radius authentication, Internet Proxy NTLM authentication, and External Services.

1. Prepare for Configuration

Update the software on every appliance in the system to version 7.1.0 or later. It is important to follow the update order and steps shown in the [Cisco Secure Network Analytics Update Guide](#).

2. Replace the Appliance Identity Certificates

Replace your appliance identity certificates with FIPS 140 compliant certificates using the instructions in the [SSL/TLS Certificates for Managed Appliances Guide](#). Make sure you update the trust stores and replace the appliance identity certificate on every appliance in your Secure Network Analytics cluster.



Make sure the appliance identity certificates include embedded extensions and meet current FIPS 140 compliance standards.

3. Enable FIPS Encryption Libraries

1. From the main menu, select **Configure > Global > Central Management**.
2. From the Inventory tab, click the **⋮ (Ellipsis)** icon for the appliance.
3. Select **Edit Appliance Configuration**.
4. Select the General tab.
5. Locate the **Compliance Mode** section.
6. **Confirm Prerequisites:** Read the prerequisites and confirm you have completed all steps and are ready to enable FIPS Encryption Libraries.

If you have completed all steps in the Certificates procedure, check the Trust Stores and appliance identity replacement check boxes.
7. If you are ready to enable FIPS Encryption Libraries, check the **Enable FIPS Encryption Libraries** check box.
8. Read the confirmation message and click **Confirm**.
9. Click **Apply Settings**.
10. Follow the on-screen prompts. The system reboots automatically.
11. If you want to have SSHD enabled on the appliance, confirm it is running.
 - Log in to the Appliance Admin interface.
 - Click **Home**. In the **Services** section, review the Status column for SSHD.
 - If the SSHD Status is shown as stopped, proceed to **Compliance Mode**.
12. **Connected:** Review the **Central Management Inventory** and make sure the appliance finishes the configuration changes and the Appliance Status returns to **Connected**.
13. Repeat the instructions in this section to enable FIPS Encryption Libraries on the next appliance in your cluster. Make sure you enable FIPS Encryption Libraries on every appliance in your Secure Network Analytics cluster.

Enable Common Criteria Encryption Libraries

Before you enable Common Criteria Encryption Libraries in Secure Network Analytics, follow the instructions in the *Common Criteria Administrative Guide*. It is important to follow the instructions to set up Common Criteria and enable encryption correctly.



Do not enable Common Criteria Encryption Libraries without following the instructions in the *Common Criteria Administrative Guide*.

Customer Success Metrics



This option is turned on by default. To turn off this function, on the General tab in Central Management, scroll down to the External Services section and uncheck the **Enable Customer Success Metrics** check box.

Customer Success Metrics (CSM) enables Secure Network Analytics system data to be sent to the cloud so that Customer Success can access vital information regarding the deployment, health, performance, and usage your system.

For more information, refer to the [Customer Success Metrics Configuration Guide](#).

DNS Server and DNS Cache

Use **DNS Server** to add or delete a DNS Server and set the Global DNS Time-Out value.

Use **DNS Cache** to change the cache size for the selected appliance.

Open DNS Server and DNS Cache

Use the following instructions to access DNS configuration settings for a selected appliance.

1. From the main menu, select **Configure > Global > Central Management**.
2. From the Inventory tab, click the **⋮ (Ellipsis)** icon for the appliance.
3. Select **Edit Appliance Configuration**.
4. Select the **Network Services** tab.

DNS Server

Use DNS Server to add or delete the IP address for DNS servers used by the selected appliance.

Add a New DNS Server

1. In the DNS Server section, click **Add New**.
2. Enter the IP address.
3. Click **Add**.
4. Click **Apply Settings**.

Delete a DNS Server IP address

Use the following instructions to delete the DNS server IP address for the selected appliance.

1. On the DNS Server list, locate the IP address you want to delete.
2. Click **Delete**.

Configure the Global DNS Time-Out Value

Use the Global DNS Time-Out Value to set the amount of time the host waits for a DNS request to be resolved before timing out.

1. In the **Global DNS Timeout Value** field, enter the number of seconds.
2. Click **Apply Settings**.

DNS Cache

Use DNS Cache to configure the amount of memory allocated for storing names resolved through DNS and failed attempts to resolve names.

Configure the DNS Cache Size

Use the following instructions to define the DNS cache size in memory.

1. In the **Cache Size in Memory field**, enter the amount of memory to allocate for caching DNS entries.
2. Click **Apply Settings**.

Clear the Cache

1. **Open an Individual Appliance** for the selected appliance.
2. Select **Configuration > Naming and DNS**.
 - To clear all cached DNS entries from the appliance, select **Clear entire cache**.
 - To clear all cached DNS lookup failure entries from the appliance, select **Clear failed cache entries**.

Host Naming

The appliance host name and network domain name are configured as part of the installation process using First Time Setup. The Host Naming section of Central Management shows this information as read-only.

- **Requirements:** A unique host name and a fully qualified domain name are required for each appliance. We cannot configure an appliance with the same host name as another appliance. Also, make sure each appliance host name meets the Internet standard requirements for Internet hosts.
- **Instructions:** To change the host name or network domain name of an existing appliance, follow the instructions in the [SSL/TLS Certificates for Managed Appliances Guide](#). As part of this procedure, you will remove the appliance from Central Management temporarily, and the appliance identity certificate is replaced automatically.

The appliance identity certificate is replaced automatically as part of this procedure.



If your appliance uses a custom certificate, please contact [Cisco Support](#) to change these settings. Make sure you have a copy of the custom certificate and private key.

Review Host Naming

Use the following instructions to review the host name and network domain name for a selected appliance.

1. From the main menu, select **Configure > Global > Central Management**.
2. From the Inventory tab, click the **⋯ (Ellipsis)** icon for the appliance.
3. Select **Edit Appliance Configuration**.
4. Select the **Appliance** tab.

Internet Proxy

Enable and configure Internet Proxy settings if you're using a proxy server for your Secure Network Analytics Internet connections. An HTTPS Internet connection is required for the following:

- **Threat Feed**
- **Smart Software Licensing**
- **Customer Success Metrics**



When the internet proxy is configured, a test ping to cisco.com is conducted through the proxy during the initial setup to check firewall accessibility.

Open Internet Proxy

Use the following instructions to open Internet Proxy for a selected appliance.

1. From the main menu, select **Configure > Global > Central Management**.
2. From the Inventory tab, click the **⋮ (Ellipsis)** icon for the appliance.
3. Select **Edit Appliance Configuration**.
4. Select the **Network Services** tab.

Enable Internet Proxy

1. Configure the DNS server using the [DNS Server](#) procedure.
2. Under Internet Proxy > Proxy Setup, check the **Enable** check box.
3. In the **Server Name Or IP Address** field, enter the proxy server FQDN or IP address.
4. In the **Port** field, enter the port number the Secure Network Analytics appliance uses to communicate with the proxy server.

Port Number Default: 8080.

5. **If the proxy server requires login credentials**, enter the following information:
 - **User Name:** Enter the user name required to log into the proxy server.
 - **Password:** Enter the password required to log into the proxy server.
 - **Authentication Type:** Select basic or ntlm.

NTLM authentication is not supported when FIPS encryption libraries are enabled. Refer to [Compliance Mode](#) for details.

- **Domain:** If you selected ntlm for the Authentication Type, enter the proxy server domain.

If the proxy server does not require login credentials, skip this step.

6. Click **Apply Settings**.
7. Follow the on-screen prompts. The appliance reboots automatically.

Local Resolution

Use the Local Resolution section to configure the local hosts file to resolve IP addresses to host names. This feature is useful when:

- Your network does not use a DNS server.
- The appliance is installed without access to a DNS server.
- You have names to resolve locally.

Open Local Resolution

Use the following instructions to open Local Resolution for a selected appliance.

1. From the main menu, select **Configure > Global > Central Management**.
2. From the Inventory tab, click the **⋯ (Ellipsis)** icon for the appliance.
3. Select **Edit Appliance Configuration**.
4. Select the **Network Services** tab.

Add a New Host

Use the following instructions to add a local host to the selected appliance.

1. In the Local Resolution section, click **Add New**.
2. In the **Host Name** field, enter the name of the host.
3. In the **Host IP Address** field, enter the host IP address.
4. Click **Add**.
5. Click **Apply Settings**.

Delete a Host

Use the following instructions to delete a local host from the selected appliance.

1. On the Local Resolution list, locate the host you want to delete.
2. Click **Delete**.
3. Click **Apply Settings**.

Network Interfaces

The appliance network interfaces are configured as part of the installation process using First Time Setup. You can change **selected network interfaces** in Central Management, or you can change the **appliance IP address** (eth0 network interface) using the appliance console (SystemConfig>Network>Interface).

Follow the instructions in the [SSL/TLS Certificates for Managed Appliances Guide](#). As part of this procedure, you will remove the appliance from Central Management temporarily, and the appliance identity certificate is replaced automatically.

The appliance identity certificate is replaced automatically as part of this procedure.



If your appliance uses a custom certificate, please contact [Cisco Support](#) to change these settings. Make sure you have a copy of the custom certificate and private key.

Review Network Interfaces

Use the following instructions to review Network Interfaces for a selected appliance.

1. From the main menu, select **Configure > Global > Central Management**.
2. From the Inventory tab, click the ***** (Ellipsis)** icon for the appliance.
3. Select **Edit Appliance Configuration**.
4. Select the **Appliance** tab.

NTP Server

The network time protocol (NTP) synchronizes the clocks of the appliances configured in Central Management, so the server operation time is within milliseconds across the network.

Use the instructions in this section to configure the NTP communication settings for the selected appliance. You need at least 1 NTP server in your appliance configuration. We also recommend that you configure a backup NTP server.

 Add at least 1 NTP server to your configuration.

- **7.5.x Installation:** If you've installed this appliance using Cisco Secure Network Analytics (formerly Stealthwatch) version 7.5.x or higher, the NTP server is added as part of the installation in First Time Setup.
- **Time Zone:** Starting in v7.0.x, Cisco Secure Network Analytics (formerly Stealthwatch) defaults to UTC. This setting cannot be changed.
- **Server Address:** Use an IP address or fully qualified domain name.
- **Disabled:** If NTP communications are disabled, contact your system administrator to manually configure the system time in System Configuration.

Open NTP Server

Use the following instructions to open NTP Server for a selected appliance.

1. From the main menu, select **Configure > Global > Central Management**.
2. From the Inventory tab, click the ***** (Ellipsis)** icon for the appliance.
3. Select **Edit Appliance Configuration**.
4. Select the **Network Services** tab.

Add an NTP Server

If you've updated the appliance, make sure you have at least 1 NTP server configured. Use these instructions to add an NTP server to the selected appliance configuration.

1. In the NTP Server section, click **Add New**.
2. In the **NTP Servers** field, click the drop-down arrow. Select an NTP server from the list.
3. Enter the server name (FQDN) or IP address.
4. Click **Add**.

5. Click **Apply Settings**.
6. Follow the on-screen prompts. The appliance reboots automatically.

Add a Custom NTP Server

Use these instructions to add a custom NTP server to the selected appliance configuration.

1. In the NTP Server section, click **Add New**.
2. In the **NTP Servers** field, click the drop-down arrow.
3. Select **Add Custom Server**.
4. Enter the server name or IP address.
5. Click **Add**.
6. Click **Apply Settings**.
7. Follow the on-screen prompts. The appliance reboots automatically.
8. Configure the [DNS Server](#) or [Local Resolution](#) to confirm the NTP custom server is resolved.

Add Authentication

Use these instructions to authenticate the connection to the selected NTP server.

Preparation: Make sure you have the NTP server key ID and key value.

1. Configure Authentication

1. In the NTP Server section, click the **⋮ (Ellipsis)** icon for the NTP server.
2. Select **Authenticate Connection**.
3. Enter the key ID and key value.
4. Click Apply Authentication.
5. Click **Apply Settings**.
6. Follow the on-screen prompts. The appliance reboots automatically.

2. Confirm Authentication

If you add authentication to a server, the key icon indicates authentication is configured. Make sure you review the audit log to confirm the authentication is successful.

1. Go to the [Central Management Inventory](#).
2. Click the **⋮ (Ellipsis)** icon for the appliance.

3. Select **Support**.
4. Select the **Audit Logs** tab.
5. In the **Category** field, select **Management**.
6. Click **Search**.
7. Confirm the NTP communication status and system time changes are shown as successful. (Check the Success column to confirm the event is shown as Yes).
 - If authentication failed (Success = No), you do not have a connection with that server. Delete the NTP server and add it again.
 - Refer to [Audit Logs](#) for details.

Delete an NTP Server

Use the following instructions to delete the NTP server IP address for the selected appliance.

1. On the NTP Server list, locate the server name or IP address you want to delete.
2. Click **Delete**.
3. Follow the on-screen prompts.
4. Click **Apply Settings**.
5. Follow the on-screen prompts. The appliance reboots automatically.

Opening Message

Use the Opening Message field to enter text that will be displayed before users log in to the selected appliance.

Open Opening Message

Use the following instructions to open Opening Message for a selected appliance.

1. From the main menu, select **Configure > Global > Central Management**.
2. From the Inventory tab, click the **⋮ (Ellipsis)** icon for the appliance.
3. Select **Edit Appliance Configuration**.
4. Select the **General** tab.

Create an Opening Message



The opening message is displayed before users log in to the appliance. Do not include confidential information in this field.

1. In the **Security Label** field, enter your message.

Language: The message can be in any language and contain any supported UNICODE character except for the ; (semicolon).

Character Limit: 10,000 alphanumeric characters and spaces maximum.

2. Click **Apply Settings**.

Password Policy

Use Password Policy to define the requirements for creating or changing a password on the selected appliance. Password-guessing programs used by attackers are continually improving. For maximum security, make sure your password requirements include a combination of uppercase letters, lowercase letters, numbers, and symbols.

Open Password Policy

Use the following instructions to open Password Policy for a selected appliance.

1. From the main menu, select **Configure > Global > Central Management**.
2. From the Inventory tab, click the **⋮ (Ellipsis)** icon for the appliance.
3. Select **Edit Appliance Configuration**.
4. Select the **General** tab, and locate the **Password Policy** section.

Define the Password Policy

Use the following instructions to define password complexity for users.

1. Configure your password policy to meet the following compliance requirements:

General

Field	Criteria	Description
Password must be at least	characters	The minimum number of characters. Defaults to 16.
Password must be at most	characters	The maximum number of characters. Defaults to 256.
Password expires after	days	The number of days before expiration. If you enter 0, the password will never expire. Note: Any values greater than 0 entered in this field must be greater than the value entered in the "Minimum

		number of days allowed between password changes" field.
Number of previous passwords disallowed	passwords	The number of previous passwords disallowed. Enter a value between 3 and 24. Defaults to 12.
Minimum number of days allowed between password changes	days	The minimum number of days allowed between password changes. Note: The value entered in this field must be lower than the value entered in the "Password Expires After:" field.

Password Complexity

 When the Password Complexity fields are set to zero, these parameters do not apply to the password policy (password complexity is disabled).

 The sum of the values entered in the PASSWORD COMPLEXITY section ("Password must contain a minimum of "uppercase letters" field, "lowercase letters" field, "numbers" field and "symbols" field) must be less than or equal to the value entered in the "Password must be at least" field from the GENERAL section.

Field	Criteria	Description	Characters
Password must contain a minimum of	uppercase letters	The minimum number of uppercase letters.	A B C D E F G H I J K L M N O P Q R S T U V W X Y Z
Password must contain a minimum of	lowercase letters	The minimum number of	a b c d e f g h i j k l m n o p q r s t u v

		lowercase letters.	w x y z
Password must contain a minimum of	numbers	The minimum number of numbers.	1 2 3 4 5 6 7 8 9 0
Password must contain a minimum of	symbols	The minimum number of symbols.	< > . , ? / ' " \ : ; ` ~ ! @ # \$ % ^ & * () - _ + = { } []
Number of characters different from previous password	characters	The number of characters which must be different from those within the previous password.	

2. Click **Apply Settings**

Protected Sessions Time-Out

Use the Protected Sessions Time-Out section to manage user sessions.

- **Administrator-Only Functions:** You can enter the number of minutes administrator privileges can be used before the session times out. When the time expires, the admin user is prompted to re-authenticate the session. If they do not re-authenticate the session, they can continue to use parts of the appliance that do not require administrator access. If you change this setting, the new values are applied immediately to all user sessions.
- **User Inactivity:** You can specify how long a user session can be inactive before users are automatically logged out. If you change this setting, the new values are applied after log out.



User sessions have a maximum of 12 hours. Although you may set the Administrator-Only Functions or User Inactivity time to exceed 12 hours, all users need to log in again after 12 hours.

Open Protected Sessions Time-Out

Use the following instructions to open Protected Sessions Time-Out for a selected appliance.

1. From the main menu, select **Configure > Global > Central Management**.
2. From the Inventory tab, click the **⋮ (Ellipsis)** icon for the appliance.
3. Select **Edit Appliance Configuration**.
4. Select the **General** tab.

Set the Protected Sessions Time-Out

1. Locate the **Protected Sessions Time-Out** section.
2. In the **Request User Re-Authentication for Administrator-Only Functions After** field, enter the number of minutes before re-authentication is required for an administrator session.

Default: The default setting is 1440 minutes (24 hours). The default is implemented when you install the appliance or upgrade the software from a major version (6.x to 7.x, etc.).

Range: 2 to 1440 minutes

0: If you enter 0, the administrator session never expires.

3. In the **Log Out User Due to Inactivity After** field, enter the number of minutes a session can be open with no activity.

Default: The default setting is 0

Range: 2 to 1440 minutes

0: If you do not change the default of 0, the session never expires.



If you change the **Log Out User Due to Inactivity After** field, make sure the **Maximum Number of Concurrent Sessions** field in the Session Settings section is set to 0. These features cannot be used at the same time.

4. Click **Apply Settings**

Session Settings

Use Session Settings to limit concurrent sessions and failed sign-on attempts. Session Settings also allows you to enable or disable the popup that provides information about a user's last login.

- **User Sessions Expiration:** User sessions have a maximum of 12 hours. After 12 hours, users need to log in again. This setting cannot be changed.
- **SSO SAML Only:** If [SSO SAML Only](#) is enabled, refer to [Session Limits](#) for more information.

Open Session Settings

Use the following instructions to open Session Settings for a selected appliance.

1. From the main menu, select **Configure > Global > Central Management**.
2. From the Inventory tab, click the ***** (Ellipsis)** icon for the appliance.
3. Select **Edit Appliance Configuration**.
4. Select the **General** tab, and locate the **Session Settings** section.

Define Session Settings

Use the following instructions to define session settings for users.

1. Complete the following fields:

Field	Criteria	Description
Number of unsuccessful attempts	number	The number of unsuccessful sign-on attempts before security lockout occurs. Defaults to 0, which allows for an unlimited number of attempts.
Duration of lockout	minutes	The duration in minutes of the security lockout. Defaults to 1.
Enable last login info popup	check box	Click the check box to enable/disable the last login info popup.
Maximum number of concurrent sessions	number	The number of concurrent sessions allowed per user. Defaults to 0, which allows for an unlimited number of concurrent sessions.



- The Security Lockout rules do not apply to admin and sysadmin users.
- If you change the **Maximum Number of Concurrent Sessions** field, make sure the **Log Out User Due to Inactivity After** field in the Protected Sessions Time-Out section is set to 0. These features cannot be used at the same time.

2. Click **Apply Settings**.

Click **Revert** to reverse changes.

Related Tasks

Session Limits

SMTP Configuration

Use SMTP Configuration for email notifications used in Response Management or DoDIN Notifications.

- **Response Management:** For more information, refer to [Response Management: Action Types](#) and follow the instructions on this page.
- **DoDIN or Common Criteria Compliance:** To configure Secure Network Analytics for the Department of Defense Information Network (DoDIN) or Common Criteria (CC) compliance, refer to [DoDIN Notifications](#) and follow the instructions in the *DoDIN Military Unique Deployment Guide* or the *Common Criteria Administrative Guide*.

Open SMTP Configuration

Use the following instructions to open SMTP Configuration for a selected appliance.

1. From the main menu, select **Configure > Global > Central Management**.
2. From the Inventory tab, click the **⋯ (Ellipsis)** icon for the appliance.
3. Select **Edit Appliance Configuration**.
4. Select the **General** tab.

Configure SMTP for Response Management

Follow these instructions to configure SMTP Configuration for Response Management email notifications.



To configure SMTP for DoDIN email notifications, follow the instructions in [DoDIN Notifications](#).

1. From the main menu, select **Configure > Global > Central Management**.
2. From the Inventory tab, click the **⋯ (Ellipsis)** icon for the appliance.
3. Select **Edit Appliance Configuration**.
4. Select the **General** tab.
5. Add the SMTP server certificate chain (root and intermediate certificates) to the Manager [trust store](#). Upload each part of the chain individually.
6. Use **one of the following options** for SMTP server identity verification:

- Confirm your SMTP server identity certificate from your Certificate Authority (CA) has a Subject or Subject Alternative Name that matches the IP address or host name you configure in step 8, **or**
 - Add the SMTP server identity certificate to the Manager [trust store](#).
7. If you are using a host name to configure the SMTP server in step 8, configure the [DNS server](#).
8. In the SMTP Configuration section, complete the following fields:

Field	Details
SMTP Server	Enter the IP address or the host name (fully qualified domain name).
Port	Default: Enter 25 . Secure SMTP Communication: Enter 587 .
From Email	Enter the source email address for your SMTP notifications.
User Name (optional)	Enter user name from your SMTP server.
Password (optional)	Enter password from your SMTP server.

9. Select an **Encryption Type**.

Encryption Type	Details
SMTPS	SMTP over SSL: Do not send notifications if the encryption is unavailable.
STARTTLS	Start TLS: Do not send notifications if the encryption is unavailable.
UN-ENCRYPTED	Do not check the encryption type. Always allow connection.

10. Click **Apply Settings**.

SNMP Agent

If you have a large networked environment that uses Simple Network Management Protocol (SNMP) as a method of system monitoring, enable SNMP Agent to access your system status information. When SNMP Agent is enabled, you can access the SNMP agent via the network if the client system is one of the Trusted Hosts.

- **Compatibility:** SNMP client V2c and V3
- **Configuration Maximum:** You can have only 1 SNMP configuration saved at a time. If you change the configuration from V2c to V3, or vice versa, your previous configuration will be deleted. For example, if you configure SNMP V3 with a custom user name, the user name will be lost, and the V3 configuration will be deleted if you change the configuration to SNMP V2c.

Open SNMP Agent

Use the following instructions to open SNMP Agent for a selected appliance.

1. From the main menu, select **Configure > Global > Central Management**.
2. From the Inventory tab, click the **⋮ (Ellipsis)** icon for the appliance.
3. Select **Edit Appliance Configuration**.
4. Select the **Network Services** tab.

Enable SNMP Agent

Use the following instructions to enable the SNMP agent.



You can have only 1 SNMP configuration saved at a time. If you change the configuration from V2c to V3, or vice versa, your previous configuration will be deleted.

1. Locate the **SNMP Agent** section.
2. Check the **Enable** check box.
3. Complete the contact information in the following fields:

Field	Description
Read Only Community:	Enter a text string that the agent will use to authenticate

	incoming SNMP requests.
SysObjectID:	Enter an object identifier that network management entities use to determine the monitoring capabilities of a device
SysName:	Enter an alphanumeric string (A-Z, a-z, and 0-9) that is 0 to 255 characters long, indicating the name assigned to the system. This string can also include special characters: period (.), hyphen (-), underscore (_), and space ().
SysLocation:	Enter an alphanumeric string (A-Z, a-z, and 0-9) that is 0 to 255 characters long, indicating where the system is physically located. This string can also include special characters: period (.), hyphen (-), underscore (_), and space ().
SysDescription:	Enter an alphanumeric string (A-Z, a-z, and 0-9) that is 0 to 255 characters long to describe the system. This string can also include special characters: period (.), hyphen (-), underscore (_), and space ().
SysServices:	Enter a number that indicates the quantity of SNMP services available on the system.
SNMP Port:	Enter the port number that the appliance uses to provide system status information. The default is 162.
SysContact:	Enter contact information about the system. The contact information needs to be an alphanumeric string (A-Z, a-z, and 0-9) that is 0 to 255 characters long. This string can also include special characters: period (.), hyphen (-), underscore (_), symbol (@), and space ().

4. In the **SNMP Version** field, select either **2c** or **3**, depending on the SNMP version.
 If you select **3**, the compatibility mode is V2c with security extensions.
 If you select **3**, complete the following required information:

- **Username:** Enter the user name for your SNMP manager. (Changing this field is optional. The default user name is read-only). The username needs to be an alphanumeric string (A-Z, a-z, and 0-9) that is 0 to 255 characters long. This string can also include special characters: period (.), hyphen (-), and underscore (_).
- **Authentication Password:** This password enables you to authenticate a user (8 characters minimum).
- **Authentication Protocol:** Use this protocol to require user authentication. This protocol DOES NOT encrypt messages.
- **Encryption Protocol:** This protocol enables you to encrypt and decrypt SNMP.
- **Encryption Password:** Use this password to require encryption. This protocol DOES encrypt messages.

We recommend that you do not use any of the following combinations of authentication and privacy protocols as the hashed password produced by the authentication algorithm is too small to be used as the key by the privacy algorithm.



- HMAC_SHA + AES_192
- HMAC_SHA + AES_256
- HMAC128_SHA224 + AES_256

If you decide to use any of the above combinations, we are using the method described in the following RFC to accommodate their use:

<https://datatracker.ietf.org/doc/html/draft-reeder-snmpv3-usm-3desede-00>

5. Click **Apply Settings**.

Review SNMP Object Identifiers

The SNMP agent supports a standard set of over 50 MIBs containing more than 2,700 OIDs.

SNMP Walk for Version V2c: If your SNMP client supports it, run an SNMP Walk or you can use the Unix program **snmpwalk** for version V2c as follows:

```
snmpwalk -m ALL -v 2c Stealthwatch_Host:port -c Community_String
```

The “-m ALL” option ensures that snmpwalk retrieves a full list of OIDs from all MIBs supported by the systems, for example OIDs in the UCD MIBs.

Without this option, snmpwalk by default retrieves only a subset of the available MIBs, so some OIDs (such as those in UCD MIBs) may not be retrieved.

SNMP Walk for Version V3: If your SNMP client supports it, run an SNMP Walk or you can use the Unix program **snmpwalk** for version V3 as follows:

```
snmpwalk -v3 Stealthwatch_Host:port -l <security level> -u <username> -a  
<authentication algorithm> -A <authentication password> -x <encryption  
algorithm> -X <encryption password>
```

Auth/Encryption algorithms and passwords are only valid when they are in the correct security level (For example, if you are in authPriv you need to include auth and encryption algorithms/password. If you are in authNoPriv, you need the auth algorithm/password but not the encryption algorithm/password).

Valid values for security level are:

- noAuthNoPriv
- authNoPriv
- authPriv

For auth/encryption algorithms, snmpwalk uses different names than we use in the user interface. Here is what each protocol in the user interface maps to for snmpwalk:

Auth:

'HMAC_SHA': 'SHA'

'HMAC128_SHA224': 'SHA224'

'HMAC192_SHA256': 'SHA256'

'HMAC256_SHA384': 'SHA384'

'HMAC384_SHA512': 'SHA512'

Encryption:

'CFB_AES_128': 'AES'

'CFB_AES_192': 'AES192'

'CFB_AES_256': 'AES256'

Example:

If you want to run an snmpwalk command when you are in V3, authPriv, using HMAC384_SHA512 and CFB_AES_256, you would use the following command:

```
snmpwalk -v3 Stealthwatch_Host:port -l authPriv -u myUsername -a SHA512 -A
myAuthPassword -x AES256 -X myEncryptionPassword
```

Object Identifiers (OIDs): The table below lists some of the Object Identifiers (alphanumeric OIDs) supported by SNMP Agent.

Alpha OID	Data Type	Description
SNMPv2-MIB::sysName	STRING	System Host Name
IF-MIB::ifDescr.#	STRING	System Description of Network Interface (e.g. eth0, lo0)
IF-MIB::ifIndex.#	INTEGER	Index number reference
IF-MIB::ifInOctets.#	Counter32	Number of Octets Received Inbound on this Interface
IF-MIB::ifOutOctets.#	Counter32	Number of Octets Sent Outbound on this Interface
IF-MIB::ifOperStatus.#	INTEGER	Status of Interface (up = 1, down=2)
HOST-RESOURCES-MIB::hrSystemProcesses	Gauge32	Total Number of Processes Running on System
HOST-RESOURCES-MIB::hrStorageDescr.#	STRING	Description of Device or Partition Name (e.g. Swap,/usr)
HOST-RESOURCES-MIB::hrStorageIndex.#	INTEGER	Numeric ID of this Device for reference in other OIDs
HOST-RESOURCES-MIB::hrStorageType.#	OID	Type of storage (Ram, Virtual Memory, Fixed Disk)
HOST-RESOURCES-MIB::hrStorageSize.#	INTEGER	Maximum Storage Capacity of Device

Alpha OID	Data Type	Description
HOST-RESOURCES-MIB::hrStorageUsed.#	INTEGER	Storage Capacity Currently being used
HOST-RESOURCES-MIB::hrStorageAllocationUnits.#	INTEGER	Unit of storage For Device in Bytes
UCD-SNMP-MIB::dskAvail.#	INTEGER	Space Available on Partition
UCD-SNMP-MIB::dskDevice.#	STRING	Partition Device Name (e.g. /dev/sda1)
UCD-SNMP-MIB::dskErrorFlag.#	INTEGER	Flag designating whether partition is “full”
UCD-SNMP-MIB::dskErrorMsg.#	STRING	Error message reported if partition fills up
UCD-SNMP-MIB::dskIndex.#	INTEGER	Index Mapping for Disk Partitions
UCD-SNMP-MIB::dskMinimum.#	INTEGER	Configured Minimum Disk KiB Avail before “full” (unused)
UCD-SNMP-MIB::dskMinPercent.#	INTEGER	Configured Minimum Disk Avail % before “full” (2=2%)
UCD-SNMP-MIB::dskPath.#	STRING	Name of mount point for partition
UCD-SNMP-MIB::dskPercent.#	INTEGER	Current Usage Percentage for partition
UCD-SNMP-MIB::dskPercentNode.#	INTEGER	Current Inode Percentage Usage for partition
UCD-SNMP-MIB::dskTotal.#	INTEGER	Maximum Capacity of this partition (KiB)
UCD-SNMP-MIB::dskUsed.#	INTEGER	Current Usage of this partition (KiB)

Alpha OID	Data Type	Description
UCD-SNMP-MIB::laConfig.#	STRING	Load Average before SNMP error is set
UCD-SNMP-MIB::laErrorMessage.#	STRING	Error message if configured load average is exceeded
UCD-SNMP-MIB::laErrorFlag.#	INTEGER	Flag designating whether load average has been exceeded
UCD-SNMP-MIB::laIndex.#	INTEGER	Index Mapping for Load Average Stats
UCD-SNMP-MIB::laLoad.#	STRING	Current Load Average for the Interval
UCD-SNMP-MIB::laLoadFloat.#	FLOAT	Current Load Average for the Interval
UCD-SNMP-MIB::laLoadInt.#	INTEGER	Current Load Average * 100 rounded to nearest integer
UCD-SNMP-MIB::laNames.#	STRING	Load Interval (Load-15 is 15 minute load average)
UCD-SNMP-MIB::prCount.#	INTEGER	Number of Processes running for the specified process name
UCD-SNMP-MIB::prErrFix.#	INTEGER	Unused
UCD-SNMP-MIB::prErrFixCmd.#	STRING	Unused
UCD-SNMP-MIB::prErrorMessage.#	STRING	Error Message if process is outside configured limits
UCD-SNMP-MIB::prErrorFlag.#	INTEGER	Flag Designating whether process is outside limits

Alpha OID	Data Type	Description
UCD-SNMP-MIB::prIndex.#	INTEGER	Index Mapping Process Names
UCD-SNMP-MIB::prMax.#	INTEGER	Maximum Number of Processes before Error (0=no limit)
UCD-SNMP-MIB::prNames.#	STRING	Name of process being monitored
UCD-SNMP-MIB::prMin.#	INTEGER	Minimum Number of Processes before error (0=no Min)

SSH

Use this section to control the ability to access the appliance using SSH (secure shell).

Default: disabled



When SSH is enabled, the system's risk of compromise increases. It is important to enable SSH only when you need it. When you are finished using SSH, disable it.

Open SSH

Use the following instructions to open SSH for a selected appliance.

1. From the main menu, select **Configure > Global > Central Management**.
2. From the Inventory tab, click the ***** (Ellipsis)** icon for the appliance.
3. Select **Edit Appliance Configuration**.
4. Select the **Appliance** tab.

Enable SSH

1. Locate the SSH section.
2. To allow SSH access on the appliance, check the **Enable SSH** check box.
3. Click **Apply Settings**.
4. Follow the on-screen prompts.

Disable SSH

1. To remove SSH access on the appliance, click the **Enable SSH** check box to clear it.
2. Click **Apply Settings**.
3. Follow the on-screen prompts.

SSL/TLS Appliance Identity and Additional SSL/TLS Client Identities

Use SSL/TLS Appliance Identity and Additional SSL/TLS Client Identities to manage your Secure Socket Layer (SSL) and Transport Layer Security (TLS) Certificates for the selected appliance. Follow the instructions in the [SSL/TLS Certificates for Managed Appliances Guide](#) for all certificate-related changes.



Your certificates are critical for your system's security. Improperly modifying your certificates can stop Secure Network Analytics appliance communications and cause data loss. Follow the instructions in the [SSL/TLS Certificates for Managed Appliances Guide](#) for all certificate-related changes.

- **Appliance Identity:** Each Secure Network Analytics version 7.x appliance is installed with a unique, self-signed appliance identity certificate. To replace the appliance identity certificate, follow the instructions in the [SSL/TLS Certificates for Managed Appliances Guide](#).

The appliance uses the SSL certificate to verify its identity to other appliances. For example, when a Manager generates a flow query and communicates to a Flow Collector, the Manager is authenticated by presenting its server identity certificate. The Flow Collector checks if this presented server identity certificate is a trusted certificate.

- **Client Identity:** The client identity is used for communication between external services. For details, follow the instructions in the [SSL/TLS Certificates for Managed Appliances Guide](#).

Review Certificates

Use the following instructions to review the appliance identity certificate or client certificates for the selected appliance.

1. From the main menu, select **Configure > Global > Central Management**.
2. From the Inventory tab, click the **⋮ (Ellipsis)** icon for the appliance.
3. Select **Edit Appliance Configuration**.
4. Select the **Appliance** tab.
5. **To review the appliance identity certificate**, go to the SSL/TLS Appliance Identity section.

To review the client identity certificates (Manager only), go to the Additional SSL/TLS Client Identities section.



Your certificates are critical for your system's security. Improperly modifying your certificates can stop Secure Network Analytics appliance communications and cause data loss. Follow the instructions in the [SSL/TLS Certificates for Managed Appliances Guide](#) for all certificate-related changes.

Certificate Expiration

If you begin receiving email notifications and seeing system alarms indicating that appliance identity certificates are expiring, refer to the [SSL/TLS Certificates for Managed Appliances Guide](#) for more information.

Threat Feed

The Cisco Secure Network Analytics Threat Feed (formerly Stealthwatch Threat Intelligence Feed) provides data from the global Threat Feed about threats to your network. The feed updates frequently and includes IP addresses, port numbers, protocols, host names, and URLs known to be used for malicious activity. The following host groups are included in the feed: command-and-control servers, bogons, and Tors.

Cisco Talos IP Blocklist Integration

Talos IP Blocklist is included in the Threat Feed as a host group under Command and Control Servers. The Talos dataset provides a list of IP addresses of the Talos untrusted threat level.

Failover Configuration

If you have a [failover configuration](#), make sure you enable Threat Feed on your primary Manager and secondary Manager.

Enabling the Threat Feed

Use the following instructions to enable the Threat Feed on the Manager.

1. License and Configure the Manager

To enable the Threat Feed, make sure your Manager is prepared as follows:

1. Add the Threat Feed license to your Cisco Smart Software Manager. For instructions, refer to the [Smart Software Licensing Guide](#).
2. Configure the [DNS server](#).
3. If your network uses an [Internet proxy](#), make sure it is configured on the Manager.
4. Configure your firewall to allow the following:

IP Ranges:

- **IPv4:** 185.146.155.64/27
- **IPv6:** 2620:122:f005::/56

If you have previously added firewall rules for the following IPv4 addresses, you can remove them:

- 64.14.29.0/24

- 64.27.162.0/24

Fully Qualified Domain Name:

- lancope.flexnetoperations.com
- esdhttp.flexnetoperations.com

2. Enable the Threat Feed

1. From the main menu, select **Configure > Global > Central Management**.
2. From the Inventory tab, click the **⋮ (Ellipsis)** icon for the Manager.
3. Select **Edit Appliance Configuration**.
4. On the **General** tab, locate the External Services section.
5. Check the **Enable Threat Feed** check box.
6. To adjust the **Feed Confidence Level**, click the drop-down.
 - Select a high number to see less potential threats (and less false positives).
 - Select a low number to see more potential threats (and more false positives).



If you have a [failover configuration](#), make sure you enable Threat Feed on your primary Manager and secondary Manager.



Disabling Threat Feed can invalidate custom security events and cause missing or false alarms. For more information, see the "When you delete a host group, service, or application" section in [Configuring Custom Policies \(Custom Security Events\)](#).

Review Alarms and Security Events

When the Threat Feed is enabled, the Threat Feed icon is shown in the Desktop Client Enterprise tree with an alarm status, and threats are displayed in their respective host group branches. For more information, refer to the [Desktop Client User Guide](#) or the Desktop Client Help.

Trust Store

When you add a certificate to an appliance trust store, you are allowing communication with that identity, whether it is another Secure Network Analytics appliance or an external service.

- **Instructions:** Follow the instructions in the [SSL/TLS Certificates for Managed Appliances Guide](#) for all trust store changes.
- **Upload Individual Files:** If your file includes more than one certificate, upload each certificate individually to the trust store.



When you add a certificate to your appliance Trust Store, your appliance trusts that identity and allows communication with it. Follow the instructions in the [SSL/TLS Certificates for Managed Appliances Guide](#) for all trust store changes.

Review Trust Store Certificates

Use the following instructions to review the certificates saved to the selected appliance Trust Store.

1. From the main menu, select **Configure > Global > Central Management**.
2. From the Inventory tab, click the **⋯ (Ellipsis)** icon for the appliance.
3. Select **Edit Appliance Configuration**.
4. Select the **General** tab.
5. Review the **Trust Store** list.

Appliance Support

The Appliance Support page includes two tabs for supporting the appliance you selected in the **Central Management Inventory**.

Open the Appliance Support Page

1. From the main menu, select **Configure > Global > Central Management**.
2. From the Inventory tab, click the **⋮ (Ellipsis)** icon for the appliance.
3. Select **Support**.
4. Select the **Configurations Files** tab or the **Audit Logs** tab. For details, see the following:
 - [Backup Configuration Files](#)
 - [Audit Logs](#)
 - [Diagnostics Packs](#)



For more appliance support options, log in to the [Appliance Administration](#) interface.

Audit Logs

The Audit Logs page provides a secure, independent audit record of events for the selected appliance, including:

- user login activities
- configuration changes
- system errors

Open Audit Logs

Use the following instructions to open the Audit Logs page for a selected appliance.

1. From the main menu, select **Configure > Global > Central Management**.
2. From the Inventory tab, click the **⋯ (Ellipsis)** icon for the appliance.
3. Select **Support**.
4. Select the **Audit Logs** tab.

Review the Audit Log Data

Each audit log provides the following appliance event information:

Column	Description
Date/Time	The date and time the data was recorded.
Category	Management: basic management and configuration changes. IDS: intrusion detection functions. Audit : auditing and system level functions. I&A: user identification and authentication. Security: user maintenance and lockout.
Event	The ID assigned to the event. You may need the Event ID if you contact Cisco Support .
Event Details	The description of the event.
User Name	The user login name associated with the event. The user ID is displayed in parentheses.
User Location	The IP address of the device used to perform the action.

Column	Description
Process Name (PID)	The component that issued the log message. The process ID reported by the component is displayed in parentheses.
Success	Yes: The action was completed. No : The action was not completed.

Sort by Column

1. Click a heading to sort your list by the selected column.
2. The heading arrows show the sort order (high-to-low or low-to-high).

Filter Search Criteria

Use the filters to search your audit log records.



You can use the filters individually or combined to narrow your search.

Filter by Date

1. Click the **Date/Time** field.
2. **Start:** Click the start date on the calendar, or enter the date in the **Start** field.
3. **End:** Click the end date on the calendar, or enter the date in the **End** field.

Date/Time

04/12/2024 10:26:00 am - 04/14/2024 10:26:00 pm

<

Apr 2024

>

Mo

Tu

We

Th

Fr

Sa

Su

25

26

27

28

29

30

31

1

2

3

4

5

6

7

8

9

10

11

12

13

14

15

16

17

18

19

20

21

22

23

24

25

26

27

28

29

30

1

2

3

4

5

10

26

00

AM

10

26

00

PM

Mo

Tu

We

Th

Fr

Sa

Su

29

30

1

2

3

4

5

6

7

8

9

10

11

12

13

14

15

16

17

18

19

20

21

22

23

24

25

26

27

28

29

30

31

1

2

3

4

5

6

7

8

9

10

26

00

PM

04/12/2024 10:26:00 am - 04/14/2024 10:26:00 pm
Cancel
Apply

- Time:** Use the drop down arrows to select the start time and end time (under the **Start** and **End** fields).
- Click **Apply**.
- Click **Search**.

Filter by Name, Location, Category

- Enter your search criteria:

User Name: Enter the complete user name and user ID. Use the following format: username(ID).

User Location: Enter the complete IP address.

Category: Select a **Category** from the list.

- Click **Search**.

Audit Log Destination

Use Audit Log Destination to configure a secure destination for audit log messages using syslog over TLS (Transport Layer Security).

Open Audit Log Destination

Use the following instructions to open Audit Log Destination for a selected appliance.

1. From the main menu, select **Configure > Global > Central Management**.
2. Click the **⋯ (Ellipsis)** icon for the appliance.
3. Select **Edit Appliance Configuration**.
4. Select the **Network Services** tab.

1. Configure Audit Log Destination

1. From the main menu, select **Configure > Global > Central Management**.
2. From the Inventory tab, click the **⋯ (Ellipsis)** icon for the appliance.
3. Select **Edit Appliance Configuration**.
4. Select the **General** tab.
5. Scroll to the **Trust Store** section.
6. Click **Add New** to add the root Certificate Authority (CA) SSL certificate that anchors your syslog server certificate chain, which supports **Syslog over TLS**. (For example: a syslog server running Splunk.)



For more information about adding certificates to the trust stores, refer to the [SSL/TLS Certificates for Managed Appliances Guide](#).

7. If your syslog server identity certificate does not include the Subject or Subject Alternative Name field that matches the value you will configure in Step 10, add the certificate to the trust store.
8. On the **Network Services** tab, locate the **Audit Log Destination** section.
9. Click **Add New** to enter the server name or IP address.



A maximum of two servers can be added.

10. In the **Server Name or IP Address** field, enter the server name (FQDN), IPv4, or IPv6 address of the remote syslog server.

11. In the **Destination Port** field, enter the port number that the appliance uses to provide audit trail information.

Destination Port Default: 6514



Entering the same IP address and server name for the same port leads to configuration issues.

12. In the **Certificate Revocation** field, select one of the following:
 - **Disabled:** Do not check the Syslog server revocation status. Always allow the TLS connection.
 - **Soft Fail:** Connect Syslog over TLS if the Syslog server revocation status is unavailable.
 - **Hard Fail:** Do not connect Syslog over TLS if the Syslog server revocation status is unavailable.

13. Click **Apply Settings**.

If the delivery fails, check the [Audit Log](#). Also, check the trust store and confirm the syslog SSL/TLS certificate is shown.

2. Enable SyslogCompliance

For Common Criteria (CC) compliance information, refer to the *Common Criteria Administrative Guide*.

If Syslog is enabled, all logs related to compliance will be saved in the remote Syslog server. Use the following instructions to establish Syslog server compliance for the selected appliance.

1. SSH into the appliance, or log into the appliance console.
2. Log in as **sysadmin**.
3. Type **SystemConfig**. Press Enter.
4. Select **Security** menu.
5. Select **Syslog Compliance**. Press Enter.



Make sure to read the status description at the top of the dialog box confirming whether you are enabling or disabling.

6. Select **Enable**. You may need to press the space bar on your keyboard to select it.
7. Select **OK**.

Backup Configuration Files

Use the Appliance Support, Configuration Files tab for the following:

- **Back Up Configuration Settings:** Secure Network Analytics creates backup configuration files of your managed appliances every 24 hours. The backup configuration files are saved up to 30 days. If a file is older than 30 days, it will be deleted. For details, refer to [Back Up Configuration Files](#).

Make sure you back up all appliances (Manager/Central Manager and each managed appliance) after installation, before and after software updates and upgrades, and whenever you make major changes to an appliance.



If you back up a Manager or a Flow Collector, you also have to back up the databases. You need both backups to restore these appliances completely. To back up the database, refer to the following instructions [Backup/Restore Database \(Appliance Admin Interface\)](#).

- **Restore Configuration Settings:** You can restore the configuration settings to the selected appliance or all appliances in Central Management. If the [backup configuration file is encrypted](#), you need to enter the password to decrypt and restore it.

Requirements: Before you get started, refer to [Restore Configuration Files](#) to review the requirements for restoring configuration files.



If you reset an appliance to its factory defaults, you cannot restore the configuration using Central Management. For assistance, please contact [Cisco Support](#).

Open Backup Configuration

1. From the main menu, select **Configure > Global > Central Management**.
2. From the Inventory tab, click the **⋮ (Ellipsis)** icon for the appliance.
3. Select **Support**.
4. Select the **Configuration Files** tab.

Back Up Configuration Files

When you back up your primary Manager/Central Manager, it creates an Manager backup configuration file and a Central Management backup configuration file. You can also back up the configuration of an individual, managed appliance.

When: Central Management creates a backup configuration file of your managed appliances every 24 hours. You should back up all appliances (Manager/Central Manager and each managed appliance) as follows:

- after you install the appliance
- before and after updates and upgrades
- whenever you make major changes to the appliance

Time: The backup process generally takes less than five minutes to complete.

Collection: The appliance continues to collect data during the backup.

Content: The configuration file includes all the files needed to restore an appliance's unique settings.

Included	Not Included
Trusted Hosts	Alarm Data
ARP Caches	Flow Data
Subnet Mask	Mitigation Data
Gateway Address	Schedule Reports
Host Name	Secure Network Analytics apps
Domain Name	Traffic Data
DNS Servers	Inter-Data Node Encryption
Device Type	
IP Address	
Individual Host Policies	

Secure Network Analytics Appliance Users	
Maps	



The content of the configuration file is determined by the appliance type. For example, the Manager configuration file includes information about the Flow Collectors sending data to it. Other appliances do not include this information.

Create a Backup Configuration File

Use the following instructions to select an appliance from the **Central Management Inventory** and create a backup file of the configuration settings.

Enable Encryption: To encrypt your backup configuration files and require a password to restore them, refer to **Backup Configuration Encryption** for details.

1. From the main menu, select **Configure > Global > Central Management**.
2. From the Inventory tab, click the **⋮ (Ellipsis)** icon for the appliance.
 - **All Managed Appliances:** To back up the configuration of all appliances managed by the Central Manager, select your primary Manager.
 - **Individual Managed Appliance:** To back up the configuration of an individual appliance in Central Management, select the Actions menu for the appliance. For example, if you only need to back up your Flow Sensor, select the Flow Sensor Actions menu.
3. Select **Support**.
4. Select the **Configuration Files** tab.
5. Click the **Backup Actions** drop-down.
6. Select **Create Backup**.

Manager/Central Manager: When you back up your primary Manager/Central Manager, it creates an Manager backup configuration file and a Central Management backup configuration file.

7. If you backed up an Manager or a Flow Collector, log in to the appliance and back up the database. For instructions, refer to **Backup/Restore Database (Appliance Admin Interface)**.



If you are backing up a Manager or Flow Collector, you also have to back up the databases. You need both backups to restore these appliances completely. To back up the database, refer to [Backup/Restore Database \(Appliance Admin Interface\)](#).

Download a Backup Configuration File

The backup configuration files are saved in Central Management up to 30 days. If a file is older than 30 days, it will be deleted from Central Management. If you download a backup configuration file, make sure you save it to a secure, remote location.

1. On the Backup Configuration Files list, locate the file you want to save.
2. Click **Download**.
3. Save the file to a secure, remote location.



Your backup configuration file includes confidential data. Make sure you save it to a secure, remote location.

Restore Configuration Files

Use this section to restore configuration settings to the selected appliance or all appliances in Central Management.

Time: The restoration process typically takes less than 5 minutes to complete. However, if a reboot is required, it may take an additional 10 minutes. Do not use the appliances during the restoration and reboot.

Compatibility: Make sure the backup configuration file matches the appliance type and software version of the appliance you are restoring.

Appliance Details: The backup configuration file contains the appliance host name and IP address. When you select a backup configuration file, make sure this information matches the appliance you are restoring.

Data Store: As part of the procedure, you will stop the database before you start the restoration. The database will restart automatically when the restoration is completed.

Collection: The appliance does not collect data during the restoration or reboot.

Encryption: If the [backup configuration file is encrypted](#), you need to enter the password to decrypt and restore it. The file names of encrypted backup files contain the **gpg** extension.

Details: When you restore a backup configuration, note the following:

- The backup file will change the appliance configuration.
- You can restore the configuration of an individual appliance or all appliances in Central Management.

Central Management: If you select a Central Management file, it restores the configuration for the primary Manager and all appliances in Central Management.

Manager: If you select a Manager file, it restores the configuration for the individual Manager.

If you are restoring the Central Management file and an individual Manager file, restore the files in the following order:



1. Central Management file
2. Primary Manager file
3. Secondary Manager file (Failover configuration only)

- **Reboots:** The appliance will reboot automatically if there are configuration changes in the backup file that require a reboot. Not all restorations require a reboot. Follow the on-screen prompts, and confirm the appliance is shown as **Connected** in the Central Management inventory.
- If the appliance IP address or host name is changed in the restoration, generate a new certificate after the reboot. Follow the instructions in the [SSL/TLS Certificates for Managed Appliances Guide](#).



If you reset an appliance to its factory defaults, you cannot restore the configuration using Central Management. For assistance, please contact [Cisco Support](#).

Restore an Appliance Backup Configuration File

1. From the main menu, select **Configure > Global > Central Management**.
2. Review the appliance status and confirm all appliances are shown as Connected.
3. **Data Store:** Select the **Data Store** tab. Confirm the Database Status is shown as Up. Click the *** (**Ellipsis**) icon. Select **Stop**.
4. From the Inventory tab, click the *** (**Ellipsis**) icon for the appliance.

- **All Managed Appliances:** To restore the configuration of all appliances managed by the Central Manager, select your primary Manager.
- **Individual Managed Appliance:** To restore the configuration of an individual appliance in Central Management, select the Actions menu for the appliance.

5. Select **Support**.
6. Select the **Configuration Files** tab.
7. On the Backup Configuration Files list, locate the file you want to restore.

If you selected the primary Manager, you can restore the individual Manager configuration or the Central Management configuration:

If you are restoring the Central Management file and an individual Manager file, restore the files in the following order:



1. Central Management file
2. Primary Manager file
3. Secondary Manager file (Failover configuration only)

- **Central Management:** To restore configuration for the primary Manager and all appliances in Central Management, select a Central Management file (**backup-cm-SMC____.tgz**). Make sure you select a file that starts with **backup-cm**.
- **Manager:** To restore configuration for the Manager exclusively, select an SMC file (**backup-appliance-SMC____.tgz**).

8. Click **Download**.
9. Save the file to a secure, remote location.



Your backup configuration file includes confidential data. Make sure you save it to a secure, remote location.

10. Click the **Backup Actions** drop-down.
11. Select **Restore Backup**.
12. Select the backup configuration file. Make sure the following information in the backup configuration file matches the appliance you are restoring:

- Host Name
- IP Address
- Appliance Type
- Software Version

13. Follow the on-screen prompts to open the file.

- **Restore:** Proceed with the restoration.
- **Cancel:** Cancel the restoration.



Restore one appliance at a time. Make sure the appliance status for all appliances is shown as Connected in the **Central Management Inventory** before you restore another backup configuration file.

14. Check the following:

- **IP Address:** If the appliance IP address or host name was changed in the restoration, generate a new certificate. Follow the instructions in the [SSL/TLS Certificates for Managed Appliances Guide](#).
- **Manager Failover Pair:** If you restore a primary Manager that is part of a failover pair, the secondary Manager will sync to the primary Manager after the restoration is completed.
- **New Appliance:** If a new appliance was added to Central Management after the backup configuration file date, you can [add it to Central Management](#) again using the appliance console (SystemConfig>Recovery>Add Appliance).

Review a Restoration Failure

If the appliance restoration reports a failure, try the [Restore](#) procedure again. If you continue to have problems, please contact **Cisco Support**.

Diagnostics Packs

 For instructions, please refer to the [System Configuration Guide](#).

Data Store Tab

Use the Data Store tab in Central Management to:

- View the status of your database or any Data Node
- View the current storage usage statistics for your database (see the [Data Store - Database Retention Tab](#) for details)
- View the status of all Data Nodes during updates (see the [Data Store - Database Update Status Tab](#) for details)
- Update the validity period of self-signed Data Store identity certificates (see the [Data Store - Database Security Tab](#) for details)

The Data Store tab also allows you to:

- Start or stop the database or any Data Node
- [Modify retention status](#) for flow interface data

Opening the Data Store Tab

1. From the main menu, select **Configure > Global > Central Management**.
2. Click the **Data Store** tab.

Viewing the Data Store Database Status

The Database Control tab page opens when you click the Data Store tab on the Central Management page. This tab displays the status of the database and each Data Node. The Data Nodes on this tab are sorted by their Private LAN IPs by default. You can re-sort the Data Nodes by clicking the column header by which you want to sort. Under normal conditions, your database and all Data Nodes will show a status of **Up**. Your database may be Up but one of your Data Nodes could be down. After recovering a failed Data Node, you may see your database showing as Up but your newly recovered Data Node will be in a recovering state.

Starting a Data Node

Follow the steps below to start a Data Node.

1. Ensure that the Database Control tab is selected.
2. Find the Data Node you want to start. Click the **⋮ (Ellipsis)** icon in the Actions column.
3. Select **Start** to start the Data Node.

Stopping a Data Node

Follow the steps below to stop a Data Node.

1. Ensure that the Database Control tab is selected.
2. Find the Data Node you want to start. Click the **⋮ (Ellipsis)** icon in the Actions column.
3. Select **Stop** to stop the Data Node.

Reviewing Last Action Results

Only one action may be in progress at any time regardless of the number of users. When an action is in progress, no other actions can be taken. Once an action has completed, the completion status will be displayed for all users in a banner at the top of the screen.

Follow the steps below to review last action results.

1. Ensure that the Database Control tab is selected.
2. Click the **Last Action Results** link at the bottom of the screen. The Action Results banner will remain on screen until you dismiss it.

Related Topics

[Data Store - Database Retention](#)

[Data Store - Database Update Status](#)

[Data Store - Database Security](#)

[About Central Management](#)

Changing the Data Store Database Passwords

Use System Configuration to change your Data Store database passwords (dbadmin and readonlyuser). You need to enable SSH temporarily as part of this procedure.

1. Log in to your Manager appliance console (SystemConfig) as sysadmin.
2. From the main menu, select **Data Store**.
3. Select **SSH**. Follow the on-screen prompts to enable SSH.
4. Select **Passwords** from the Data Store menu.
5. Follow the on-screen prompts to change the passwords.

When you exit System Config, the system restores your previous SSH settings.

Related Tasks

Changing your Flow Collector Database Password

Related Topics

Data Store Tab

Adding a Data Store to a Non-Data Store Deployment and Transitioning Your Flow Collectors

This process allows you to transition your existing Non-Data Store Flow Collectors to use the Data Store database without loss of pre-transition data or visibility. Once you have completed the steps outlined in the [System Configuration Guide](#), you can preserve your pre-existing data until you no longer need it. Transitioning Non-Data Store Flow Collectors to Data Store Flow Collectors also allows you to take advantage of features only available in Data Store such as:

- **Increased ingest capacity:** Data Store deployments are scalable up to 3 million flows per second and may alleviate some of your current ingest capacity limitations. Flow Collectors in Data Store mode may exhibit up to a 200% increase in performance.
- **Multi-Telemetry Support:** Data Store deployments are capable of handling NetFlow, Remote worker/Endpoint (NVM), and Firewall connection and security event telemetry.
- **Long-term data retention:** Data Store deployments provide scalable storage, enabling long-term data retention (up to 2 years) without adding Flow Collectors.
- **Enterprise-class data resiliency:** Telemetry data is stored redundantly across Data Nodes. This ensures no service interruption during single node failures.
- **Greatly improved query and reporting response times:** The Data Store provides drastically improved query performance and reporting response times, in some cases 10x faster or more compared with a Non-Data Store deployment model.

Follow the instructions in the [System Configuration Guide](#) to transition your Non-Data Store Flow Collectors to Data Store Flow Collectors.

Data Store – Database Retention

The Database Retention tab answers questions such as:

- How full is my database?
- How much is each telemetry type (NetFlow, NVM, firewall log) contributing to this fullness?
- How much data was newly stored in my database yesterday?
- What is the total capacity of my database?



All of the charts as well as the Data Storage Statistics section on this page are updated once per day.

Database Fullness Chart

The Database Fullness chart displays the amount of used and free space that exists in your Data Store database.

Per Telemetry Contribution Chart

The Per Telemetry Contribution chart displays a breakdown of the data that exists in your Data Store database.

Daily Storage

The Daily Storage section displays the incremental amount of data that was added to your database on the previous day. By monitoring your daily storage rate, you can evaluate how quickly your database is filling as well as how much each telemetry type is contributing to your daily storage accumulation.

Oldest Data in Data Store

This table shows the date and number of days since the oldest record was written to the Data Store. This data is updated once per day.

Data stored locally in a Flow Collector (or Flow Collector database) is not included in this table. If you are transitioning a non-Data Store Flow Collector to a Data Store Flow Collector and have a data retention policy, you can use this table to understand how much new data is in your Data Store and to know when it is an ideal time to complete your transition.

Opening the Data Store - Database Retention Tab

1. From the main menu, select **Configure > Global > Central Management**.
2. Click the **Data Store** tab.
3. Click the **Database Retention** tab.

Changing the Flow Interface Data Storage Period

On this tab, you can modify the retention days for flow interface data. This option only impacts the NetFlow portion of the data that is occupying space in your system. The default is 7 days. You may increase or decrease this number as needed.

Related Topics

[Data Store](#)

[Data Store - Database Update Status Tab](#)

[About Central Management](#)

Data Store - Database Update Status

After initiating an update of your Data Nodes from your [Central Management Update Manager](#), use the Update Status tab to monitor the progress of the database services update on each Data Node.

Opening the Data Store - Database Update Status Tab

1. From the main menu, select **Configure > Global > Central Management**.
2. Click the **Data Store** tab.
3. Click the **Database Update Status** tab.

Monitoring the Database Update Status


Each Data Node progresses through a series of states during an update. Click the Data Store Update Workflow link to see a visual representation of the update process (shown below).



For a successful update, follow the update order and instructions in the [Cisco Secure Network Analytics System Update Guide](#).



Some of the state transitions shown in the image below happen very quickly during the update process so you may not see them occur during a screen refresh.



Related Topics

[Data Store](#)

[Data Store - Database Retention Tab](#)

[Update Manager](#)

[About Central Management](#)

Update Manager

Use the Update Manager to apply software updates and patches. You can download the files from your Cisco Smart Account on Cisco Software Central at <https://software.cisco.com>.

Open Update Manager

1. From the main menu, select **Configure > Global > Central Management**.
2. Select the **Update Manager** tab.

Updates

Plan to update the entire system when your appliances will be experiencing relatively low volumes of traffic.

To update your appliances, follow the instructions in the [Cisco Secure Network Analytics System Update Guide](#). The guide includes details that are critical for a successful update, including:

- Planning for the update
- Installing patches
- Creating backup files and diagnostic packs
- Updating your appliances in order



For a successful update, follow the update order and instructions in the [Cisco Secure Network Analytics System Update Guide](#).

Updating Data Nodes

Click **Update all Data Nodes** to update your Data Nodes at the same time.

- **Status:** The Update Manager shows the overall update status, and the new installed version is shown after the update is completed. If a Flow Collector has a Data Store Transition tag, it is in the process of being transitioned from a non-Data Store Flow Collector to a Data Store Flow Collector. The appliances reboot during the update process, so the status information may be delayed when the appliance goes offline. To monitor the progress of the database services update on each Data Node, go to the **Data Store > Database Update Status** tab. Also, refresh each page to see the most recent status.

- **Instructions:** For a successful update, follow the update order and instructions in the [Cisco Secure Network Analytics System Update Guide](#) or [patch readme file](#).
- **Best Practices:** Although you can update each Data Node individually, we recommend using the **Update all Data Nodes** button to update your Data Nodes at the same time.

Patches

Keep Secure Network Analytics up-to-date by installing the latest patches for your software version. You can download the files from your Cisco Smart Account on Cisco Software Central at <https://software.cisco.com>.

For each patch, make sure you follow the instructions in the patch readme file. A link to a patch readme file is provided with the patch file on Cisco Software Central. You can also download patch readme files directly from the [Release Notes page](#) on Cisco.com.

Types of Patches

There are generally two types of patches:

- appliance-specific rollout patches for a version
- common update patches for all versions and/or appliances

Common update patches for CIMC firmware version updates and Cisco Bundles are released periodically. For details, refer to the [latest readmes](#). You can download the SWU files from Software Central at <https://software.cisco.com>.

CIMC Firmware Version Common Update Patch

When made available, make sure to apply the CIMC firmware version update patches for your UCS C-Series hardware. These are common update patch SWU files that apply to all Cisco Secure Network Analytics (formerly Stealthwatch) appliances.

Cisco Bundles Common Update Patch

Cisco releases bundles of pre-validated digital certificates of a select number of root certificate authorities (CAs), when necessary. We provide Cisco Bundles common update patch SWU files for you to install on of your Cisco Secure Network Analytics (formerly Stealthwatch) appliances.

Each patch includes a Core certificate bundle and an External certificate bundle, which are used for connecting to Cisco services and to non-Cisco services.



- Make sure you have the latest Cisco Bundles common update patch installed on all of your appliances.

- If you do a RefreshImage for an appliance, make sure to install the most recent Cisco Bundles common update patch. Doing a RefreshImage reverts to the CAs shipped with an appliance.

Detections Packs

As new detections become available, we provide an updated detections pack you'll install as a SWU file for your Manager. The first pack is included with Secure Network Analytics. Go to the **Configure > Detection > Pack Management** to access the Detections Pack Management page where you can select the newly released pack or to return to a previously released pack. Refer to the [Detections Pack Guide](#) for more information.

Related Topics

[Data Store Tab](#)

[Data Store - Database Update Status](#)

[Central Management Inventory](#)

App Manager

We introduced apps in v7.0.0 of Cisco Secure Network Analytics (formerly Stealthwatch). Secure Network Analytics apps are similar in concept to the apps you install on a smartphone. They are optional features that enhance and extend the capabilities of Secure Network Analytics. The release schedule for the apps is independent from the normal Secure Network Analytics upgrade process. Due to this, we can update apps as needed without having to link them with a core Secure Network Analytics release, and you can install apps without having to update your Secure Network Analytics system.

Use the App Manager page to manage your installed Secure Network Analytics apps. From this page you can install, update, uninstall, or view the status of an app. After installing an app, you can access it from the appropriate option on the dashboard in the Secure Network Analytics Web App. Your user permissions determines which apps you can view.

When you update Secure Network Analytics, the app that is currently installed is retained; however, some apps may require you to upgrade to the latest version of Secure Network Analytics. In addition, when you upgrade your Secure Network Analytics system, you may need to upgrade some or all of the apps. To learn which app version is supported by a particular version of Secure Network Analytics, see the [Secure Network Analytics Apps Version Compatibility Matrix](#).



Only a Primary Admin can install or uninstall an app.



When you update to a later version of an app, simply install the newer version over the existing version. You do not need to uninstall your existing app. If you uninstall an app, all files associated with it, including temporary files, are removed.

App Compatibility Notice

To simplify the Cisco Secure Network Analytics customer experience, only one version of a Secure Network Analytics app will be available to install at any point in time (similar to the app store model). Although we strive for maximum app compatibility, not all versions of an app will be compatible with all versions of Secure Network Analytics.

Cisco reserves the right to discontinue a Secure Network Analytics app at any time. There may be many reasons for doing so, including but not limited to the following:

1. The equivalent capabilities provided by the app are now provided elsewhere, either via a new version of the app, a new app, or via a feature in Secure Network

Analytics.

2. The capabilities provided by the app are no longer considered relevant or useful to our customer base.

If the decision is made to discontinue a Secure Network Analytics app, advance notice will be provided at least sixty days prior to the discontinuation date. Although Secure Network Analytics apps are currently included with your Secure Network Analytics license, Cisco reserves the right to charge license fees for certain Secure Network Analytics apps in the future.

How to find this page

1. From the main menu, select **Configure > Global > Central Management**.
2. Click the App Manager tab.

The App Manager page opens. A list of installed apps are listed in the Apps table.

What should you do on this page

View a list of your installed Secure Network Analytics apps

The table on this page displays any Secure Network Analytics apps you have installed, sorted by name, and their version. If an app is no longer supported or if it requires the user to upgrade to a newer version of Secure Network Analytics, this will be noted here as well. See [How do I know which version of an app I can use](#) in this topic.

How do I know which version of an app I can use

Use the App Manager page to manage your installed apps. From this page you can install, update, uninstall, or view the status of an app. Refer to the following table to learn about the possible app statuses.



Since it is possible that a newer version of an app exists and is not listed in App Manager, always check to see if a newer version is available in on Cisco Software Central.

Status	Definition	Action to Take
UpToDate	Your installed app is the most current version.	No action is required.

Status	Definition	Action to Take
UpdateAvailable	You have upgraded to a new version of Secure Network Analytics. Your existing app is supported by this version of Secure Network Analytics, but a new version of this app is available.	If you desire, go to Cisco Software Central to download and install the latest version (this will replace your existing version).
UpgradeRequired	You have upgraded to a new version of Secure Network Analytics, and your existing app is not supported by the Secure Network Analytics version you are now using.	To continue using this app, go to Cisco Software Central to download and install the latest version (this will replace your existing version).
AppNotSupported	You have upgraded to a new version of Secure Network Analytics. This app may no longer be supported by the version of Secure Network Analytics you are now using. It could be that this app has been deprecated or a newer version of this app has not yet been released.	Go to Cisco Software Central to see if a new version has been released.
NewApp	This is a new app.	If you desire, install this new app using Central Manager.
Error	The installation, upgrade, or removal process for the associated app has not successfully completed.	Contact Cisco Support . A partial installation, upgrade, or removal of this app may have occurred. If

Status	Definition	Action to Take
		so, this must be corrected.

Install a Secure Network Analytics app

When an app is available for a particular Secure Network Analytics version, it is posted in [Cisco Software Central](#). To install a Secure Network Analytics app, complete the following processes:

Download the SWU file from Cisco Software Central

1. Go to [Cisco Software Central](#).
2. In the Download and Upgrade section, click **Access downloads**.
3. In the **Select a Product** search bar, enter **Secure Network Analytics** and press **Enter**.
4. Choose **Secure Network Analytics Manager 2210** from the list.
5. Choose the applicable app from the list.
6. In the window on the right, click the  (**Download**) icon for the app SWU file and download to your choice of location.

Upload and install the app on Central Manager



- It usually takes a few minutes to upload and install an app.
- Only the system administrator can upload and install apps.

1. Verify that you are installing an app that is compatible with your current version of Secure Network Analytics. See the [Secure Network Analytics Apps Version Compatibility Matrix](#).
2. On the App Manager tab, click **Browse** to select the SWU file.
3. Select the app file.
The upload and installation process automatically begins.
4. (Conditional) If you need to cancel the upload process, click **Cancel** in the Upload dialog.

After you install the app, you can access it from the main menu under the **Dashboards** menu.

Uninstall an app

Click **Uninstall** in the row that contains the app you want to uninstall.

Sort columns

You can only sort by name or status. To reverse the sort order for either column, click the column name. Use the column name to toggle between ascending and descending order. You can also click the arrow that is displayed to the right of the column name (after you click the column name). An up arrow indicates that the column is sorted in ascending order, and a down arrow indicates descending order.

Access help

To access the help for a Secure Network Analytics app when you are logged in to the app,

click the  (**Help**) icon located next to the topic heading in the upper left corner of the page.

Related Topics

About Central Management

Smart Software Licensing

Use Smart Software Licensing to see an overview of your license status and usage.

- **Instructions:** For instructions and details about Cisco Smart Software Licensing, please refer to the [Cisco Secure Network Analytics Smart Software Licensing Guide](#).
- **Manage Licenses:** To add or update licenses on your account, log in to your [Cisco Smart Software Manager](#) or contact your administrator.

Open Smart Software Licensing

1. From the main menu, select **Configure > Global > Central Management**.
2. Select the **Smart Licensing** tab.

Product Instance Name

The Product Instance Name is the identifier we use for your Secure Network Analytics product instance, which includes your Manager and managed appliances. This identifier is shown on your Secure Network Analytics Smart Software Licensing page (in the Smart Software Licensing Status section) and in your [Cisco Smart Software Manager](#).

Evaluation and Registration

When you use Secure Network Analytics in Evaluation mode, your appliances are licensed to use selected features for 90 days. To use Secure Network Analytics with maximum default functionality, and to add licenses and features to your account, register your product instance for Smart Software Licensing.

- **Expiration:** Make sure you register before the 90-day evaluation period expires. When the evaluation period expires, flow collection will stop. To start flow collection again, register your product instance.
- **Registration:** To register, follow the instructions in the [Cisco Secure Network Analytics Software Licensing Guide](#). As part of the registration process, you will log in to your [Cisco Smart Software Manager](#) at <https://software.cisco.com> to generate the registration token and add licenses.

Transport Settings

Use Transport Settings to configure Secure Network Analytics to communicate with the Cisco Smart Software Manager. For configuration instructions and troubleshooting, please refer to the [Cisco Secure Network Analytics Smart Software Licensing Guide](#).

Database

Use the Database tab on the Central Management page to update your flow collector database password for all Flow Collector databases in a Non-Data Store domain. To change the Data Store database password, refer to [Changing the Data Store Database Passwords](#).



Make sure you change the default password. When a new Flow Collector is added to Central Management, the database password automatically updates to match the current password.

Opening the Database Tab

1. From the main menu, select **Configure > Global > Central Management**.
2. Click the Database tab.

Changing your Flow Collector Database Password

Use the following instructions to change your Flow Collector database password.

1. From the main menu, select **Configure > Global > Central Management**.
2. Click the **Database** tab.
3. To generate a random password, click the **Generate Password** button, otherwise enter your password in the Password and Confirm Password fields.
4. Click the **Show Password** check box to view your chosen password.
5. Click the **Apply Settings** button to save your changes.



When you change a database password, only Non-Data Store Flow Collectors and Transition Flow Collectors will receive the new password. For information on changing your Data Store passwords, refer to the [System Configuration Guide](#).

Related Tasks

Changing the Data Store Database Passwords

User Management: Users Tab Overview

Use this page to view existing Secure Network Analytics user accounts.



For information about configuring user accounts, see [User Management: User Configuration Overview](#).

Open the Users Tab

From the main menu, choose **Configure > Global > User Management**.

The User Management: Users tab opens.

Related Topics

[User Management: User Configuration Overview](#)

[User Management: User Configuration: Table of Contents](#)

Users Tab: Use the Menu

If you want to open a page in a new tab when pivoting from the context menu, do one of the following:

- On a Windows-based PC, perform a Ctrl+Click.
- On a Mac, perform a Cmd+Click.



This functionality does not work if you are connected to a virtual machine.

Related Topics

[User Management: Users Tab Overview](#)

Users Tab: Filter Columns

At the top of each of the filterable columns, there is a filter field in which you can enter the criteria by which you want to filter that column. As you start to type your entry, the column dynamically filters to display a list of entries that match what you have begun to type. This match can be for any word in the entry (beginning, middle, or last word). Type the complete entry into the field to return results that are an exact match.

Filtering with multiple entries

You can type multiple entries within one column as well as enter values across multiple columns. When you type multiple values within one column, these values are ORed together; therefore, all returned results need only match one of the entries within the column. Remember to leave a space between each entry. When you filter more than one column, filters are ANDed together; therefore, all returned results must match the search criteria in all of the columns. Enclose an entry within quotation marks to return results that exactly match your entry.

Clearing entries

Click the X in the field to clear the filter criteria and return to the previous view for that column. To quickly clear multiple filter criteria, refresh the page.

Related Topics

[User Management: Users Tab Overview](#)

Users Tab: Manage Passwords

In order to help ensure the security of Secure Network Analytics, you should change your password on a regular basis.

Refer to the following guidelines:

- The Admin (default Secure Network Analytics administrative user, username **admin**) can reset passwords for all users, including Primary Admins.
- A Primary Admin can reset passwords for all users, including other Primary Admins, but not for themselves or for the Admin.
- Only the Admin can change the Admin password.

Refer to the following table to view directions for changing your password depending on the circumstance.

If a user's...	Then...
Password is not expiring within 10 days and the user wants to change it anyway	<u>If you are NOT the Admin or the Primary Admin</u> Click the  (User) icon and select Profile > Change Password. <u>If you are the Admin or the Primary Admin</u> You can do either of the following: • In the applicable row click the ... (Ellipsis) icon in the Actions column and click Change Password from the menu.

If a user's...	Then...
	Click the User icon and select Profile > Change Password.
Password will be expiring within 10 days	The numeral 1 is displayed beside the  (Alert) icon . Click the Alert icon to view a message that informs you the number of days that remain until your password expires. When you click the message, the notification message and Alert icon disappear, and the Change Password dialog opens.
Password is not FIPS-compliant	An  (Information) icon with the words <i>Change Password</i> beside it is displayed under the applicable user's name in the User column. The user can still log in to Secure Network Analytics. There is no time limit as to when the Admin or a Primary Admin has to change the user's password, however, your organization cannot be

If a user's...	Then...
	certified as FIPS-compliant unless all passwords comply to FIPS requirements.
Authentication service setting has been changed from a non-local to a local service	An Information icon with the words <i>Enter Password</i> beside it is displayed under the applicable user's name in the User column. The user cannot log in to Secure Network Analytics until his or her password is changed by either the Admin or a Primary Admin.

Related Topics

User Management: Users Tab Overview

Users Tab: Enable and Disable User Accounts

If you want to...	In the Status column...
Enable a user to log in to the Web App	Slide the switch to the right (turn ON).
Disable a user	Slide the switch to the left (turn OFF).

You do not have to click **Save** after enabling or disabling. Your change is saved automatically.

Related Tasks

[Users Tab: Edit User Accounts](#)

[Users Tab: Delete User Accounts](#)

Related Topics

[User Management: Users Tab Overview](#)

Users Tab: Edit User Accounts

In the applicable row, do the following:

1. Click the **Ellipsis** icon in the Actions column.
2. Click **Edit** from the menu.

The [User page](#) opens with all the user's account details populated in the applicable fields.

If this user...	Is editing this user type...	He can edit...
Admin (the default administrative user that is built in to Secure Network Analytics and whose user name is <i>admin</i>)	Admin	Full Name and Email
	Primary Admin and all other users	All user information except for User Name
Primary Admin	Primary Admin and all other users	All user information except for User Name

The Admin has full access to Secure Network Analytics and is configured as follows:

- Authentication type is *System*.
- Data role is *All Data (Read & Write)*.
- Web role is *Primary Admin*.
- Desktop Client role is *All*.

Related Tasks

[Users Tab: Enable and Disable User Accounts](#)

[Users Tab: Delete User Accounts](#)

Related Topics

[User Management: Users Tab Overview](#)

Users Tab: Delete User Accounts

In the applicable row, click the **Ellipsis** icon in the Actions column and click **Delete** from the menu.

The user account is deleted.

You cannot delete the following users:



- The Admin user account (user name *admin*), which is the default administrative user built in to Secure Network Analytics.
- Your own user account.

Related Tasks

[Users Tab: Enable and Disable User Accounts](#)

[Users Tab: Edit User Accounts](#)

Related Topics

[User Management: Users Tab Overview](#)

Data Role Management: Data Roles Tab Overview



For information about configuring data roles, see [Data Role Management: Data Role Configuration Overview](#).

Use this page to view existing Secure Network Analytics data roles. This page is sorted by name. Data roles control which domains, host groups, appliances, and devices a user can view (read) and make configurations to (write), including the following:

- Domains
- Host groups
- Flow Collectors
- Exporters
- Flow Sensors
- External Devices
- Identity Devices

The following data roles are included by default in Secure Network Analytics:

This role...	Allows the user to...
All Data (Read Only)	View data in any domain for any host group, appliance, or device, but cannot make any configurations.
All Data (Read & Write)	View AND configure data in any domain for any host group, appliance, or device.

This role...	Allows the user to...
The specific menu items (i.e., functionality) that the user can view and/or configure (e.g., host groups, reports, active directory, policies) is determined by the user's web role.	

Open the Data Roles Tab

1. From the main menu, choose **Configure > GLOBAL User Management**.

The User Management: Users tab opens.

2. Click the Data Roles tab.

The User Management: Data Roles tab opens.

Related Topics

[Data Role Management: Data Role Configuration Overview](#)

Data Role Management: Data Roles Tab Overview: Table of Contents

Data Roles Tab: Use Context Menus

If you want to open a page in a new tab when pivoting from the context menu, do one of the following:

- On a Windows-based PC, perform a Ctrl+Click.
- On a Mac, perform a Cmd+Click.



This functionality does not work if you are connected to a virtual machine.

Related Topics

[Data Role Management: Data Roles Tab Overview](#)

Data Roles Tab: Filter Columns

At the top of each of the filterable columns, there is a filter field in which you can enter the criteria by which you want to filter that column.

For the Name field, as you start to type your entry in any of these fields, a drop-down list appears displaying a list of entries that match what you have begun to type. You can then chose the appropriate entry from the list so that you don't have to type the complete entry. The remaining columns dynamically filter to display only the entries that match your entry.

For the Description field, as you start to type your entry, the column dynamically filters to display a list of entries that match what you have begun to type. This match can be for any word in the entry (beginning, middle, or last word). Type the complete entry into the field to return results that are an exact match.

Filtering with multiple entries

You can type multiple entries within one column. Additionally, the Name column is a type-ahead field. As you start to type your entry in this field, a drop-down list appears displaying a list of entries that match what you have begun to type. You can then chose the appropriate entry from the list so that you don't have to type the complete entry.

Editing a role with filtered criteria

If you edit a data role after having filtered criteria, the entries in the table do not change until you save your edits. In this scenario, all columns filter only entries displayed in the table.

Clearing entries

Click the X in the field to clear the filter criteria and return to the previous view for that column. To quickly clear multiple filter criteria, refresh the page.

Related Topics

[Data Role Management: Data Roles Tab Overview](#)

Data Roles Tab: Sort Columns

To reverse the sort order for a column, click the column name.

Use the column name to toggle between ascending and descending order. You can also click the arrow that is displayed to the right of the column name (after you click the column name). An up arrow indicates that the column is sorted in ascending order, and a down arrow indicates descending order.

Related Topics

[Data Role Management: Data Roles Tab Overview](#)

Data Roles Tab: Edit Data Roles

In the applicable row, do the following:

1. Click the **Ellipsis** icon in the Actions column.
2. Click **Edit** from the context menu.

The Data Roles page opens with all the duplicated data role details populated in the applicable fields.

When you change a data role, users assigned that role are logged out as the changes are applied.

Related Topics

[Data Role Management: Data Roles Tab Overview](#)

Data Roles Tab: Delete Data Roles

In the applicable row, click the **⋮ (Ellipsis)** icon in the Actions column and click **Delete** from the context menu.

The data role is deleted.

Note the following guidelines when deleting a data role:

- Only the Admin (the default administrative user that is built in to Secure Network Analytics and whose user name is *admin*) and the Primary Admin can delete a data role.
- You cannot delete the All Data (Read & Write) or the All Data (Read Only) data roles.
- You cannot delete a data role that is already assigned to a user(s). You must first assign another data role to each user that is currently assigned the data role you want to delete.

Related Topics

[Data Role Management: Data Roles Tab Overview](#)

Data Roles Tab: Duplicate Data Roles

In the applicable row, do the following:

1. Click the **Ellipsis** icon in the Actions column.
2. Click **Duplicate** from the context menu.

The Data Roles page opens with all the duplicated data role details populated in the applicable fields.

Related Topics

[Data Role Management: Data Roles Tab Overview](#)

Data Role Management: Data Role Configuration Overview

Use this page to create and manage all custom data roles assigned to the users in your Secure Network Analytics system. For example, one scenario where you might want to do this is if you wanted to create a data role that allowed users to view (Read Only) and configure (Read & Write) information only for certain host groups and specific appliances or devices within one particular domain.

For information about viewing data roles, see [Data Role Management: Data Roles Tab Overview](#).

Open the Data Roles Page

Do one of the following to access this page

- From the main menu:
 1. From the main menu, from the main menu, choose **Configure > Global > User Management**.
The User Management: Users tab opens.
 2. Click the Data Roles tab.
 3. Click **Create** from the drop-down menu and select **Data Role**.
The Data Roles page opens.
- On the User Management: Data Roles tab:
 1. In the row of the applicable custom security event, click the **⋮ (Ellipsis)** icon in the Actions column.
 2. From the context menu, click **Edit**.
The Data Roles page opens.

What to do Next

[Data Role Management: Create Data Roles](#)

Related Topics

[Data Role Management: Data Roles Tab Overview](#)

Data Role Management: Data Role Configuration: Table of Contents

Data Role Management: Create Data Roles

Click **Create > Data Role**.

The Data Roles page opens.

Refer to the following parameters when configuring data roles:

- Only the Admin (default Secure Network Analytics administrative user, username **admin**) and the Primary Admin can create and modify a data role.
- You cannot edit the following data roles:
 - All Data (Read Only)
 - All Data (Read & Write)
- You can select a child without selecting a parent.
- When you select a host group, the associated domain is automatically selected.
- When you change a data role, users assigned that role are logged out as the changes are applied.

To create or edit a data role, do the following:

1. Complete the fields at the top of the page.
2. Set the Read and Read & Write permissions (that you want associated with the user) for each appropriate component in the enterprise tree.

After you have saved the information on this page, the User Manager: Data Roles tab opens and the account you just added is displayed in the Users table.

Refer to the following creating a data role on this page.

Name (required)

This must be a unique name.

Description

You can enter whatever content is pertinent to your needs.

Data Access for this Role

Type the name of the component (e.g., domain, host group, appliance or device name) for which you want to make edits, and it moves to the top of the enterprise tree.

Data Role configuration tree

Expand the tree to access the components for which you want to configure view and configure permissions that are associated with a user. When you select a parent host group, its children and any host groups under that parent host group are automatically

selected. If you want to select children or sub host groups of a parent host group, you must first deselect the parent host group.

If you give a user Read & Write access to a child host group, the user still sees the parent host group of that child in the enterprise tree but cannot edit that parent host group.

The following default data roles are displayed as check boxes within the Read Only column and the Read & Write column.

If you check this check box for the component...	Then the user can...
Read Only	View the data for that component but cannot make any configurations.
Read & Write	View AND configure the data for that component.
The specific menu items (i.e., functionality) that the user can view and/or configure (e.g., host groups, reports, active directory, policies) is determined by the user's web role.	

Related Topics

[Data Role Management: Data Role Configuration Overview](#)

Managing Authentication and Authorization

Use this page to manage an authentication and authorization service which has already been configured. Authentication is the process of verifying who a user is, and authorization is the process of verifying what the user should be able to access. Secure Network Analytics has an authentication service, local, which is automatically enabled. You can manage your authentication and authorization services centrally from your Manager.

 Only a Primary Admin can manage authentication or authorization services.

Authentication Services

Confirm which of the following authentication services to manage:

- Remote Authentication Dial-In User Service (RADIUS) - Authentication service only
- Terminal Access Controller-Access Control System (TACACS+) - For detailed instructions about configuring TACACS+, refer to the [Secure Network Analytics TACACS+ Configuration Guide](#).
- Lightweight Directory Access Protocol (LDAP)



If you've configured your Managers as a failover pair, the external authentication services are only available on the primary Manager. Before you can use configured external authentication services on a secondary Manager, you need to promote the secondary Manager to primary.

Authorization Services

Confirm which of the following authorization services to manage:

- Terminal Access Controller-Access Control System (TACACS+) - For detailed instructions about configuring TACACS+, refer to the [Secure Network Analytics TACACS+ Configuration Guide](#).
- Lightweight Directory Access Protocol (LDAP)



If you've configured your Managers as a failover pair, the external authorization services are only available on the primary Manager. Before you can use configured external authorization services on a secondary Manager, you need to promote the secondary Manager to primary.

How to find this page

Use this page to edit or delete services, and to enable or disable remote authorization.



Any authentication or authorization service which has been configured is listed in the table on this page. For field descriptions, go to [Configuring Authentication and Authorization](#).

1. From the main menu, choose **Configure > Global > User Management**.

The User Management: Users tab opens.

2. Click the Authentication and Authorization tab.

The User Management: Authentication and Authorization tab opens.



Only a Primary Admin can manage authentication or authorization services. For more information about configuring data roles, see [Data Role Management: Data Role Configuration Overview](#). For more information about managing data roles, see [Data Role Management: Data Roles Tab Overview](#).

Edit an authentication or authorization service

1. From the main menu, choose **Configure > Global > User Management**.

The User Management: Users tab opens.

2. Click the Authentication and Authorization tab.

The User Management: Authentication and Authorization tab opens with the table displayed.

3. Locate the Authentication Type column and identify the service (RADIUS, TACACS+, or LDAP) you'd like to edit.
4. Click the **⋮ (Ellipsis)** icon in the Actions column for the service, and select **Edit** from the drop-down menu.
5. Make changes as needed.
6. Click **Save**.

Delete an authentication or authorization service

-  Make sure to remove the authentication or authorization service from each associated user's profile before you attempt to delete the service.

1. From the main menu, choose **Configure > Global > User Management**.
The User Management: Users tab opens.
2. Click the Authentication and Authorization tab.
The User Management: Authentication and Authorization tab opens with the table displayed.
3. Locate the Authentication Type column and identify the service (RADIUS, TACACS+, or LDAP) you'd like to delete.
4. Click the **⋮ (Ellipsis)** icon in the Actions column for the service, and select **Delete** from the drop-down menu.
5. Click **Delete** when the dialog box displays if you are sure you want to delete the service.

-  Click **Cancel** if you don't want to delete the service.

Enable remote authorization

-  RADIUS is not supported for remote authorization.

1. From the main menu, choose **Configure > Global > User Management**.
The User Management: Users tab opens.
2. Click the Authentication and Authorization tab.
The User Management: Authentication and Authorization tab opens with the table displayed.
3. Locate the Authentication Type column and identify the service (RADIUS, TACACS+, or LDAP) you'd like to enable.
4. Click the **⋮ (Ellipsis)** icon in the Actions column for the service, and select **Enable Remote Authorization** from the drop-down menu.

5. Click **Enable** when the dialog box displays if you are sure you want to enable the service.

 Click **Cancel** if you don't want to enable the service.

A check mark displays in the Remote Authorization column next to the service you enabled.

Disable remote authorization

 If you disable a remote authorization service, all those using the service are automatically logged off and are unable to regain access.

1. From the main menu, choose **Configure > Global > User Management**.
The User Management: Users tab opens.
2. Click the Authentication and Authorization tab.
The User Management: Authentication and Authorization tab opens with the table displayed.
3. Locate the Authentication Type column and identify the service (RADIUS, TACACS+, or LDAP) you'd like to disable.
4. Click the **⋮ (Ellipsis)** icon in the Actions column for the service, and select **Disable Remote Authorization** from the drop-down menu.
5. Click **Disable** when the dialog box displays if you are sure you want to disable the service.

 Click **Cancel** if you don't want to disable the service.

Related topics

[Configuring Authentication and Authorization](#)

[Data Role Management: Data Role Configuration Overview](#)

[Data Role Management: Data Roles Tab Overview](#)

Configuring Authentication and Authorization

Authentication is the process of verifying who a user is, and authorization is the process of verifying what the user should be able to access. Secure Network Analytics has an authentication service, local, which is automatically enabled.

 Only a Primary Admin can configure authentication or authorization services.

Authentication Services

You have the option to add and configure one of the following authentication services:

- Remote Authentication Dial-In User Service (RADIUS): Authentication service only
- Terminal Access Controller-Access Control System (TACACS+): For detailed instructions about configuring TACACS+, refer to the [TACACS+ Configuration Guide](#).
- Lightweight Directory Access Protocol (LDAP)
- Security Assertion Markup Language Single Sign-On (SAML SSO): Authentication service only



If you've configured your Managers as a failover pair, the external authentication services are only available on the primary Manager. Before you can use configured external authentication services on a secondary Manager, you need to promote the secondary Manager to primary.

Authorization Services

You have the option to add and configure one of the following authorization services:

- Terminal Access Controller-Access Control System (TACACS+): For detailed instructions about configuring TACACS+, refer to the [TACACS+ Configuration Guide](#).
- Lightweight Directory Access Protocol (LDAP)

You can then manage your authentication and authorization services centrally from the Manager.



If you've configured your Managers as a failover pair, the external authorization services are only available on the primary Manager. Before you can use configured external authorization services on a secondary Manager, you need to promote the secondary.

RADIUS Message-Authenticator

A Message-Authenticator attribute is used to authenticate and ensure the integrity of RADIUS packets to prevent spoofing.



Ensure you configure the RADIUS server to validate the Message-Authenticator attribute.

If you have RADIUS server configured, complete the following instructions:

1. Log in to the RADIUS server.
2. Enable the Message-Authenticator.

How to find this page

From the main menu:

1. Choose **Configure > Global > User Management**.

The User Management: Users tab opens.

2. Click the **Authentication and Authorization** tab.

The User Management: Authentication and Authorization tab opens.



Only a Primary Admin can manage authentication or authorization services. For more information about configuring data roles, see [Data Role Management: Data Role Configuration Overview](#). For more information about managing data roles, see [Data Role Management: Data Roles Tab Overview](#).

Add an authentication or authorization service

To add an Authentication or Authorization Service:

1. Click **Create** and select **Authentication Service**.
2. Select the type of service from the drop-down menu.

The Authentication Service page opens. Depending on your selection, fields specific to RADIUS, TACACS+, LDAP or SSO display.

3. Complete the fields, as required.

Add a RADIUS server

If you're using Windows Internet Authentication Service as your RADIUS authentication server, your Active Directory user account must have "dial-in" access enabled. If it is not enabled, you can not log in to the Manager.

 RADIUS is not supported for remote authorization.

To add a RADIUS server:

1. Complete the following fields.

Field	Notes
Name	Enter a unique name to identify the server.
Description	Enter a description which specifies how or why the server is being used.
Cache Timeout (Seconds)	The amount of time (in seconds) that a user name or password is considered valid before Secure Network Analytics requires re-entry of the information.
Prefix	This field is optional. The prefix string is placed at the beginning of the user name when the name is sent to the RADIUS server. For example, if the user name is zoe and the realm prefix is <i>DOMAIN-A\</i> , the user name DOMAIN-A\zoe is sent to the server. If you do not configure the Prefix field, only the user name is sent to the server.

Suffix	This field is optional. The suffix string is placed at end of the user name. For example, if the suffix is <i>@mydomain.com</i> , the username <i>zoe@mydomain.com</i> is sent to the RADIUS server. If you do not configure the Suffix field, only the user name is sent to the server.

2. Click **Add New** in the Servers section, then complete the following fields.

Field	Notes
IP Address	Use either IPv4 or IPv6 addresses when configuring authentication services.
Port	Enter any numbers from 0 to 65535 which correspond to the applicable port.
Secret Key	Enter the secret key that was configured for

	the applicable server.

3. Click **Add**.

4. Click **Save**.

The new RADIUS server is added, and information for the server displays.

Add a TACACS+ server

For instructions about configuring TACACS+, refer to the [TACACS+ Configuration Guide](#).

 You can add only one TACACS+ server to the TACACS+ authentication service.

To add a TACACS+ server:

1. Complete the following fields.

Field	Notes
Authentication Service	
Name	Enter a unique name to identify the server.
Description	Enter a description which specifies how or why the server is being used.
Cache Timeout (Seconds)	The amount of time (in seconds) that a user name or password is considered valid before Secure Network Analytics requires re-entry of the information.
Prefix	This field is optional. The prefix string is placed at the beginning of the user name when the name is sent to the TACACS+ server. For example, if the user name is <i>zoe</i> and the realm prefix is <i>DOMAIN-A\</i> , the user name <i>DOMAIN-A\zoe</i> is sent to the server. If you do not configure the Prefix field, only

	the user name is sent to the server.
Suffix	This field is optional. The suffix string is placed at end of the user name. For example, if the suffix is <i>@mydomain.com</i> , the username <i>zoe@mydomain.com</i> is sent to the TACACS+ server. If you do not configure the Suffix field, only the user name is sent to the server.
Server	
IP Address	Use either IPv4 or IPv6 addresses when configuring authentication services.
Port	Enter any numbers from 0 to 65535 which correspond to the applicable port.
Secret Key	Enter the secret key that was configured for the applicable server.

2. Click **Save**.

The new TACACS+ server is added, and information for the server displays.

Add an LDAP server



Secure Network Analytics only supports LDAP on Microsoft servers, which are not load-balanced.

To add an LDAP server:



Make sure you install the LDAP server certificate to the [Trust Store](#) before configuring LDAP.

1. Complete the following fields.

Field	Notes
Friendly Name	Enter a name for the LDAP server.
Description	Enter a description for the LDAP server.
Server Address	Enter the fully qualified domain name as specified in the Subject Alternative Name (SAN) field of the LDAP server certificate. • If the SAN field contains only the IPv4 address, enter the IPv4 address in the Server Address field. • If the SAN field contains the DNS

	name, enter the DNS name in the Server Address field. • If the SAN field contains both DNS and IPv4 values, use the first value listed. • If the SAN field is empty, use the Common Name (CN).
Port	Enter the port designated for secure LDAP communication (LDAP over TLS). A commonly used TCP port for LDAP is 636.

2. For **Certificate Revocation**, select one of the following:

- **Disabled:** Do not check the LDAP server revocation status. Always allow users to log in.
- **Soft Fail:** Users can log in to the system if the LDAP certificate revocation status is unavailable.
- **Hard Fail:** Users cannot log in to the system if the LDAP certificate revocation status is unavailable.

3. Complete the following fields.

Field	Notes
Bind User	Enter the administrative user ID used to connect to the LDAP server. Secure Network Analytics does not support added or repeated comma characters in the Bind User field. Make sure you do not enter additional or repetitive commas if using a DN as part of the Bind User name.
Password	Enter the password for the Bind User.

Confirm Password	Re-enter the exact characters of the Password you just entered.
Base Accounts	Enter the Distinguished Name (DN). The DN applies to the branch of the directory in which searches for users should begin. It is often the top of directory tree (your domain), but you can also specify a sub-tree within the directory. The Bind User and the users you are authenticating should be accessible from Base Accounts. For example: DC=example,DC=com

4. Click **Save**.

The new LDAP server is added, and information for the server displays.

Define LDAP Authorization Group Mappings

For LDAP authorization, users are defined remotely in Microsoft Active Directory. The cisco-stealthwatch roles and passwords are also defined in Active Directory.



The Primary Admin is assigned to the `cisco-stealthwatch-master-admin` group and can't be a member of any other cisco-stealthwatch groups.

Each user, other than the Primary Admin, must be assigned to at least one of each group, as follows.

Data Role: Choose only one

- `cisco-stealthwatch-all-data-read-and-write`
- `cisco-stealthwatch-all-data-read-only`
- `cisco-stealthwatch-<custom>`

 Make sure custom data roles begin with `cisco-stealthwatch-`.

Web Functional Role: Choose at least one

- `cisco-stealthwatch-configuration-manager`
- `cisco-stealthwatch-power-analyst`
- `cisco-stealthwatch-analyst`

Desktop Functional Role: Choose at least one

- `cisco-stealthwatch-desktop-stealthwatch-power-user`
- `cisco-stealthwatch-desktop-configuration-manager`
- `cisco-stealthwatch-desktop-network-engineer`
- `cisco-stealthwatch-desktop-security-analyst`

 Make sure custom desktop functional roles begin with `cisco-stealthwatch-desktop-`.

Custom Roles and Groups:

Custom roles and groups must be defined on both the Manager and in Active Directory.

For example, if you define a custom role, `custom1`, for a desktop client role, it maps to `cisco-stealthwatch-desktop-custom1` in Active Directory. Non-custom roles and groups don't need to be added to the Manager.

Configuring SAML SSO

Use the following instructions to configure Security Assertion Markup Language Single Sign-On (SAML SSO). SSO is an authentication process that allows a user to access multiple applications with one set of credentials.

Support Details

The following identity provider configurations are supported.

- Microsoft Active Directory Federation Services (ADFS) for SAML/SSO
- Okta

- Login.gov
- Microsoft Entra ID

The following configurations are supported.

- Additional Proxies – Support of additional proxies give you the ability to customize service provider FQDNs. Note that if you access Secure Network Analytics through a proxy, you will want to set the service provider FQDN of the proxy address.

The following configurations are not supported.

- SAML Request Signing and Encrypted SAML responses

1. Prepare for Configuration

You need one of the following requirements to configure SSO:

Requirement	Details
Identity Provider Metadata URL	The URL must use the fully qualified domain name or IPv4 address.  If you are using an Identity Provider Metadata URL, you first need to add the root signing certificate for your identity provider to the Manager Trust Store. Make sure that your URL address matches the address that is in the certificate.
Identity Provider Metadata XML File	XML file obtained from your identity provider.

2. Create an SSO Authentication Service

Use the following instructions to create an SSO authentication service.

Select the SSO Authentication Service

1. Log in to the Manager (Web App).
2. Select **Configure > Global > User Management**.

3. Select the **Authentication and Authorization** tab.
4. Select **Create > Authentication Service > SSO**.

Enter Your Basic Information

1. Enter a description for the SSO service in the **Description** field.
2. Select an identity provider from the **Identity Provider Type** drop-down menu.

Enter Your Identity Provider Information

Choose one of the following options to enter your identity provider information.

- Click the **Identity Provider Metadata URL** radio button and enter your identity provider metadata URL in the field below.



If you are using an Identity Provider Metadata URL, you first need to add the root signing certificate for your identity provider to the Manager Trust Store. Make sure that your URL address matches the address that is in the certificate.

- Click the **Upload Identity Provider Metadata XML File** radio button. Click the **Add File** button and navigate to the location of your identity provider metadata XML file on your hard drive. Select the file and click **Open** to upload the file.

Disable Requested Authentication Context

If you encounter authentication context errors while connecting to your Identity Provider (IdP), you can try disabling authentication context to resolve the issue.

1. Check the **Disable Requested Authentication Context** check box.
2. Test the connection to ensure the authentication context errors are resolved.

If issue persists even after disabling authentication context, contact [Cisco Support](#) for further assistance.

Add Your General Configuration Information

1. Select a name identifier format from the **Name Identifier Format** drop-down menu. For additional information on Name Identifier Formats and the identity providers they support, refer to the SAML 2.0 core standard documentation. You can alternatively contact your identity provider to find out which Name Identifier formats they support.
2. Add a log in screen label in the **Log in Screen Label** field. The label you enter here will be visible on the log in screen.
3. If you have a custom service provider address, enter it into the **Custom Service Provider Address** field. The service provider entity ID will be overwritten with this address. This is the address users will use to log in to the appliance.
4. Click **Save** to save your changes.

3. Add a Local SSO User

Use the following instructions to add a local SSO user. Users are authenticated through/by the Identity Provider.

1. Log in to the Manager (Web App).
2. Select **Configure > Global > User Management**.
3. Select **Create > User**.

For instructions, click the  (**Help**) icon. Select **Help**. For details about adding users, refer to "Configuring Users."

4. Complete the fields to create a new user. Configure the user as follows:
 - **Authentication Service:** Select the SSO service you created in [2. Create an SSO Authentication Service](#).
 - **User Name:** Enter a user name for the IdP account. For example, you can use the first part of the user's email address. For name@cisco.com, you would enter "name" in this field.

- **SAML ID:** Enter the Name ID for this user. You can obtain the Name ID from your identity provider. If your Name ID is dynamic or if you can't look it up, log in as that user (expecting the login to fail). Once the login fails, review the audit log to see the Name ID that was sent when the login attempt failed.

5. Click **Save**.

6. Confirm the SSO User is shown in User Management.

4. Test SAML Login

1. On the Web UI login page, select **Log in with SSO**.
2. You will be redirected to your Identity Provider URL to provide your log in credentials.
3. After being authenticated by your Identity Provider, you will be redirected back to your Manager.

Troubleshooting

Scenario	Notes
SSO Only is enabled and the SSO account is unavailable	Access the System Configuration utility and select Advanced > SSO Only to disable the SSO Only option. This enables local login.
SSO only is enabled, the sysadmin account is disabled, and the SSO account is unavailable	Reboot the Manager into the Grub menu and contact Cisco Support to get a Consent Token. Access the System Configuration utility and select Advanced > SSO Only to disable the SSO Only option. This enables local login.
Cannot save the Identity Provider Metadata URL	Confirm that your Manager can reach the Identity Provider, that your Identity Provider certificate has a name that matches the host portion of the configured URL, and that your Manager Trust Store includes the IdP root CA certificate. If issues persist, contact Cisco Support.
Cannot save the Identity Provider Metadata XML file	Confirm that the metadata file conforms to the SAML 2.0 schema. Review the IdP configuration and make sure the data you entered is accurate and doesn't include any extra spaces.

	Also, review the IdP event logs.
SSO User is authenticated but log in to the Manager fails	Verify that a local user with SSO authentication service and a SAML ID has been added.
The authentication service was unable to fulfill your request	Verify that a local user exists with a valid SAML ID that matches what the IdP presents as the Name ID for the user logging in.
Login issues using SSO with ADFS on Windows machines across Chrome, Firefox, and Edge browsers	Change the browser's proxy settings from "Use system proxy settings" to "Auto-detect proxy settings for this network."
Additional Issues	Download a SAML tracer for your browser. Repeat the SSO log in for the unable to look up Name ID scenario to review the exchanges between the IdP and the service provider.

Related topics

Managing Authentication and Authorization

Data Role Management: Data Role Configuration Overview

Data Role Management: Data Roles Tab Overview

Session Limits

Use Session Limits to set the maximum number of concurrent sessions and failed sign-on attempts for a role when SSO SAML Only is enabled.



If you've set session limits for a specific role, it overrides the [Session Settings](#) limit configured in Central Management.

Accessing Session Limits

1. [Configure and enable SAML SSO and enable SSO Only](#).
2. From the main menu, choose **Configure > Global > User Management**.
3. Select the **Session Limits** tab.

Configuring Session Limits

1. Locate a role in the list.
2. In the **Maximum Sessions** column, click the up or down arrows to set the maximum number of concurrent sessions for the role.

Default: The number of concurrent sessions allowed per user defaults to 0, which allows for an unlimited number of concurrent sessions. For more information, refer to [Session Settings](#).

Related Tasks

[Configuring Authentication and Authorization](#)

[Session Settings](#)

[Users Tab: Enable and Disable User Accounts](#)

[Users Tab: Delete User Accounts](#)

Related Topics

[User Management: Users Tab Overview](#)

User Management: Sysadmin Password

Use this page to change the **sysadmin** user password for your Manager and all managed appliances.

 You can change the **sysadmin** user password only after the system administrator (**sysadmin** user) enables **Sysadmin Password** option in the Manager appliance console (SystemConfig). For detailed instruction about enabling this option in the Manager appliance console refer to the [System Configuration Guide](#).

Open the Sysadmin Password Tab

 The **Sysadmin Password** tab is only accessible to **admin** users.

1. From the main menu, choose **Configure > Global > User Management**.
2. Select the **Sysadmin Password** tab.

Change Sysadmin Password

 If the password fields are disabled (read-only), it indicates that the **sysadmin** user has not enabled the **Sysadmin Password** option in the Manager appliance console (SystemConfig).

1. In the **New Password** field, enter the new password.
2. In the **Confirm Password** field, re-enter the new password.
3. (Optional) Click **Generate Password** to automatically create a strong password.
4. (Optional) Check the **Show Password** check box to view the entered passwords.
5. Click Save.

Related Topics

[User Management: Users Tab Overview](#)

[Sysadmin User](#)

User Management: User Configuration Overview

Use this page to create and manage the user accounts in your Secure Network Analytics system, specifically for your primary Manager.

- **User Accounts:** For information about viewing and maintaining user accounts, see [User Management: Users Tab Overview](#).
- **All Other Appliances:** To manage user accounts for an individual Data Node, Flow Collector, Flow Sensor, or UDP Director, refer to [Add/Edit/Delete Users](#) in the Appliance Administration interface.

Refer to the following guidelines when using this page:

- By default, a user is assigned the Analyst web role in the Manager.
- Beginning with Cisco Secure Network Analytics (formerly Stealthwatch) v7.0,
 - Existing and new users in the Desktop Client are assigned the Analyst web role.
 - Existing users in the Appliance Administration (Admin) interface are assigned the Primary Admin web role.

Open the User page

Do one of the following to access this page

- From the main menu:
 1. From the main menu, choose **Configure > Global > User Management**.
The User Management: Users tab opens.
 2. Click **Create** from the drop-down menu and select **User**.
The User page opens.
- On the User Management: Users tab:
 1. In the row of the applicable custom security event, click the **⋮ (Ellipsis)** icon in the Actions column.
 2. From the menu, click **Edit**.
The User page opens.

Related Topics

User Management: Users Tab Overview

User Management: User Configuration: Table of Contents

User Configuration: Create User Accounts



All users have access to all domains, however, you can modify that access by editing the user's [data role](#).

To create a user account, do the following:

1. Complete the fields at the top of the page.
2. Select the roles you want the user to have.

After you have saved the information on this page, the User Manager: Users page opens and the account you just added is displayed in the Users table.

User Name (required)

This must be a unique name. The value in this column is the user's login user ID. You cannot change this field once you have created the user account.

Full Name

The user's first and last name.

Email

The user's email address.

Authentication Service

Authentication is the process by which the identity of a client or user is verified by a server or host. The process can take many forms, but usually it involves the client presenting a password to the server for verification in order to obtain access.

If you have one of the following authentication services, then you can configure the Manager to leverage one of these services for system access capability. Secure Network Analytics supports the following authentication services:

- RADIUS (Remote Authentication Dial-in User Service)



If your system is using Windows Internet Authentication Service as the RADIUS authentication server, the user's Active Directory user account must have "Dial-in" access enabled. Otherwise, the user cannot log in to the Manager.

- TACACS+ (Terminal Access Controller Access Control System)
- LDAP (Lightweight Directory Access Protocol)

The options listed in the drop-down list reflect the authentication service types that you have created in the Desktop Client in the User Management feature, on the Authentication

Service tab. By default, *local* is selected when you add a new user, regardless of which application in which they are created.

- Local means that it is authenticated within the Manager.
- All services other than local are authenticated based on how the service is defined in the Desktop Client in the User Management feature, on the Authentication Service tab.

If you're using Windows Internet Authentication Service as your Radius authentication server, your Active Directory user account must have "Dial-in" access enabled. Otherwise, you cannot log in to the Manager.

Password (conditional)

This field is required if you select *local* for the authentication service.

Use the Password Policy section in the Central Management > Edit Appliance Configuration > General tab to define the requirements for creating or changing a password on the selected appliance. Password-guessing programs used by attackers are continually improving. For maximum security, make sure your password requirements include a combination of uppercase letters, lowercase letters, numbers, and symbols.

For more information about passwords, see [Password Policy](#).

Confirm Password (conditional)

This field is required if you select *local* for the authentication service.

Related Tasks

[User Configuration: Assign Data Roles to Users](#)

[User Configuration: View Web Role Permissions](#)

[User Configuration: Assign Web Roles to Users](#)

[User Configuration: Create Desktop Client Roles](#)

[User Configuration: Assign Desktop Client Roles to Users](#)

Related Topics

[User Management: User Configuration Overview](#)

User Configuration: Assign Data Roles to Users

Data roles control which domains, host groups, appliances, and devices a user can view (read) and make configurations to (write). The list of appliances and devices is as follows:

- Flow Collectors
- Exporters
- Flow Sensors
- External Devices
- Identity Devices

To assign a data role, in the Role Settings section near the bottom of the page, use the Data Role drop-down list.

The All Data (Read Only) is selected by default.

You must select a data role. You cannot select more than one. When you assign a new data role to a user, the user is logged out as the change is applied.

Secure Network Analytics contains the following default data roles.

This role...	Allows the user to...
All Data (Read only)	View data in any domain or host group, or on any appliance or device, but cannot make any configurations.
All Data (Read & Write)	View AND configure data in any domain or host group, or on any appliance or device.

This role...	Allows the user to...
The specific functionality (e.g., flow search, policy management, network classification, etc.) that the user can view and/or configure is determined by the user's web role.	

Related Tasks

[User Configuration: Create User Accounts](#)

[User Configuration: View Web Role Permissions](#)

[User Configuration: Assign Web Roles to Users](#)

[User Configuration: Create Desktop Client Roles](#)

[User Configuration: Assign Desktop Client Roles to Users](#)

Related Topics

[User Management: User Configuration Overview](#)

User Configuration: View Web Role Permissions

For a complete list of permissions assigned to each web role, at the bottom of the User page, click the **Compare** link next to a Web Roles heading.

Related Tasks

[User Configuration: Create User Accounts](#)

[User Configuration: Assign Data Roles to Users](#)

[User Configuration: Assign Web Roles to Users](#)

[User Configuration: Create Desktop Client Roles](#)

[User Configuration: Assign Desktop Client Roles to Users](#)

Related Topics

[User Management: User Configuration Overview](#)

User Configuration: Assign Web Roles to Users

Web roles control which functionality (e.g., flow search, policy management, network classification, etc.) in a domain or related to a host group, appliance, or device users can view and/or configure in the Manager. Remember that the specific domains, host groups, appliances, and devices within which a user can access and configure functionality are determined by the user's data role.



Only Web roles are available on Data Store domains. You will not see separate tabs for Web or Desktop roles if you are working in a Data Store domain.

Refer to the following guidelines when using this page:

- You can assign and view web roles only in the Manager.
- A user must be assigned at least one web role, though you can select as many as you want.
- You cannot modify or delete any web roles.
- You cannot create custom web roles.
- When you assign a new web role to a user, the user is logged out as the change is applied.

To assign a web role, in the Role Settings section at the bottom of the page, click the Web tab to access the web role options.

The following table lists the default web roles that Secure Network Analytics uses.

This role...	Allows the user to do this in the Manager...
Primary Admin	View all functionality and change anything. It is automatically assigned to all web roles (all other check boxes are dimmed), and its data role is <i>All Data (Read & Write)</i>. You <i>cannot</i> change the data role for a Primary Admin. The check box for the Primary Admin is located directly under the "Settings" heading. All other web role check boxes are located on the Web tab.
Power Analyst	Perform the initial investigation into traffic and flows as

This role...	Allows the user to do this in the Manager...
	well as configure policies and host groups. This role is automatically assigned the <i>All Data (Read & Write)</i> data role. You <i>cannot</i> change the data role for a Power Analyst.
Configuration Manager	View configuration-related functionality. You can change the data role for a Configuration Manager.
Analyst	Perform the initial investigation into traffic and flows. You can change the data role for an Analyst.
For a complete list of permissions assigned to each web role, click Compare (located next to the Web Roles heading on the Web tab).	

Related Tasks

[User Configuration: Create User Accounts](#)

[User Configuration: Assign Data Roles to Users](#)

[User Configuration: View Web Role Permissions](#)

[User Configuration: Create Desktop Client Roles](#)

[User Configuration: Assign Desktop Client Roles to Users](#)


Related Topics

[User Management: User Configuration Overview](#)

User Configuration: Create Desktop Client Roles

You can create Desktop Client roles (previously referred to as function roles in the Desktop Client only in the Desktop Client. If you have done so, they are listed in the Manager, on the Desktop tab on the User Management: User page, (which this current topic is about) and in the Desktop Client, as described below.

- When you change a Desktop Client role, users assigned that role are logged out as the changes are applied.
- In the Desktop Client, the Desktop Client role(s) assigned to a user are listed in the Desktop Client Role dialog (see image below). For information about how to access this dialog, see the "Add/Edit Desktop Client Roles" topic in the Desktop Client help.



Related Tasks

[User Configuration: Create User Accounts](#)

[User Configuration: Assign Data Roles to Users](#)

[User Configuration: View Web Role Permissions](#)

[User Configuration: Assign Web Roles to Users](#)

[User Configuration: Assign Desktop Client Roles to Users](#)

Related Topics

User Management: User Configuration Overview

User Configuration: Assign Desktop Client Roles to Users

Desktop Client roles control which functionality (e.g., flow search, policy management, reports, etc.) users can view and configure in the Desktop Client. To assign Desktop Client roles, click the Desktop tab near the bottom of the page to access the Desktop Client role options.



Any edits you make to a role in the Desktop Client are also reflected in the Manager. When you assign a new Desktop Client role to a user, the user is logged out as the change is applied.

Role Settings

☐ Master Admin

DATA ROLE

All Data (Read Only) ▼

Web Desktop ←

☒ Desktop Client Manager

DESKTOP CLIENT ROLES +

☐ All ☐ Configuration Manager ☐ Network Engineer ☐ Security Analyst ☐ StealthWatch Power User

Secure Network Analytics comes with the following default Desktop Client roles (previously referred to as function roles in the Desktop Client).

Click one or more Desktop Client roles that you want to assign to the associated user. If you select Desktop Client Manager, you do not have to select any more roles, since that role encompasses all of the other roles.



Desktop roles are only available on Non-Data Store domains. For Data Store domains, only Web roles are available. You will not see separate tabs for Web or Desktop roles if you are working in a Data Store domain.

Select this option...	To allow the user to view...
Desktop Client Manager	All menu items and change anything within the Desktop Client.
Configuration Manager	All menu items and configure all appliances, devices, and domain settings.

Select this option...	To allow the user to view...
Network Engineer	All traffic-related menu items within the Desktop Client, append alarm and host notes, and perform all alarm actions, except mitigation.
Security Analyst	All security-related menu items, append alarm and host notes, and perform all alarm actions, including mitigation.
Secure Network Analytics Power User	All menu items, acknowledge alarms, and append alarm and host notes, but without the ability to change anything.

To give a user access to the Desktop Client, assign the user one or more Desktop Client roles by doing the following:

1. From the Authentication Service: drop-down list (in the top section of this page), select the applicable option.
2. In the Role Settings section near the bottom of the page, click the Desktop tab.
3. You must select either the Desktop Client Manager check box or select one or more roles in the Desktop Client Roles section. See the table above for a list of the default Desktop Client roles.
 - By default, the All Data (Read Only) option is selected.
 - If you select the Desktop Client Manager check box, the following occurs:
 - All the Desktop Client Roles check boxes are selected.
 - The user is assigned the All Data (Read & Write) option data role.

Related Tasks

[User Configuration: Create User Accounts](#)

[User Configuration: Assign Data Roles to Users](#)

[User Configuration: View Web Role Permissions](#)

[User Configuration: Assign Web Roles to Users](#)

[User Configuration: Create Desktop Client Roles](#)

Related Topics

[User Management: User Configuration Overview](#)

Data Retention: Overview



- Only users assigned the role of Admin or Configuration Manager can manage data retention.
- Data retention configurations cannot be deleted.

Use this page to change the settings that control how long Secure Network Analytics retains data for various documents.

Secure Network Analytics stores data using different methods. In each case, data storage capacity depends on many factors, such as the total space available on the appliance (the disk space varies by the model of the appliance), the FPS (flow per second), and the average flow size.

- **Traffic, schedules, and most documents**

Secure Network Analytics stores this information using a database on the Manager. Each data type retains its associated data according to the settings on the Data Retention page.

- **Flow and Security Event List data**

Your Secure Network Analytics Flow Collectors store this information.

Data retention varies between documents, depending on how data is provided to Secure Network Analytics. For example, your system receives traffic data every several minutes. The data's resolution decreases as time passes, though maximum daily values are retained. Therefore, traffic data that your system originally recorded every five minutes converts into hourly data, and eventually it converts into daily data.

When you print a document that contains traffic data spanning multiple days, the information is displayed using the most reduced (i.e., coarsest) resolution. For example, if you request a week of data, the data displays as daily values, even though hourly data may be available for a more recent time period. If your appliance's percentage of used disk space reaches a certain value, it begins deleting old data one day at a time.

Open Data Retention

From the main menu, choose **Configure > Global > Manager**.

The Data Retention tab is selected by default. You can view detailed data retention information associated with the specific IP address and appliance referenced at the top of the page.

Data Retention: Table of Contents

Data Retention: Sort Columns

To reverse the sort order for a column, click the column name. Use the column name to toggle between ascending and descending order. You can also click the arrow that is displayed to the right of the column name (after you click the column name). An up arrow indicates that the column is sorted in ascending order, and a down arrow indicates descending order.

Related Topics

[Data Retention: Overview](#)

Data Retention: View Details

View details for the following sections:

Manager application configuration

This section displays the name, IP address, model number, and serial number of the associated Manager.

Data retention

Refer to the following information about the fields in the Data Retention table.

Name The name associated with the information category you are viewing.

Configuration The length of time and type of data that is retained for each information category (represented by a single row in the table).

This setting overrides any maximum setting the user might choose on the Add Schedule: Schedule page or Edit Schedule: Schedule page in the Desktop Client. For more information, see the Add Schedule: Schedule Page topic and the Edit Schedule: Schedule page in the Desktop Client online help.

Example: If this setting is 90 days, and the user chooses 120 days on the Add Schedule: Schedule page or Edit Schedule: Schedule page for a particular Manager document, that document is archived no longer than 90 days.

For the Alarm Table, active alarms are not deleted. Also, the number of retention days is relative to the last active time, NOT the time the alarm condition was reported as active.

When you change the retention to a longer period, you must wait for the difference of time to expire before the data being stored corresponds exactly to the retention settings. Until that time, the data is displayed using the most reduced (i.e., coarsest) resolution available.

Example: You changed the retention for 5-min data from 3 days to 10 days. You must now wait 7 days before the data being stored corresponds exactly to the retention settings.

Related Topics

[Data Retention: Overview](#)

Data Retention: Edit Data Retention Values

1. In the Actions column for the applicable document, click the  (**Edit**) icon.

The Edit Data Retention page opens. Depending on the configuration details of the entry you are editing, one or more data retention configuration fields are displayed at the top of the Data Retention table.

2. Make edits using any or all of the displayed fields to designate the default length of time that Secure Network Analytics retains a specified data type. You can either type your entry in the field or click the field up and down arrows to choose a value. Following are the various configuration fields that you might see (depending on what you chose in the Name column):

- Days of 5-minute data
- Days of hourly data
- Days of daily data
- Number of days

[Click here to view the default data retention values for each data category.](#)

3. To save your edits, click **Save**.



We encourage administrators to decrease the resolution for storing less valuable data so that the system has more room to store more valuable data.

Related Topics

[Data Retention: Overview](#)

Default Data Retention Values

Name	Configuration Default Value	Range
AlarmTable	60 days	Enter a value between 1 and 3650
Autonomous SystemTraffic	7 days of 5 minute-data; 1 day of hourly data; 365 days of daily data	Enter a value between 1 and 3650
DHCP Lease	60 days	Enter a value between 1 and 3650
Domain Packet Size - TCP	60 days	Enter a value between 1 and 3650
Domain Packet Size - UDP	60 days	Enter a value between 1 and 3650
Domain PortUsage - TCP	60 days	Enter a value between 1 and 3650
Domain Port Usage - UDP	60 days	Enter a value between 1 and 3650
Domain Traffic	15 days of 5-minute data; 60 days of hourly data; 365 days of daily data	Enter a value between 1 and 3650
Domain Traffic Statistics	15 days of 5-minute data; 60 days of hourly data; 365 days of daily data	Enter a value between 1 and 3650

DSCP Traffic	15 days of 5-minute data; 60 days of hourly data; 365 days of daily data	Enter a value between 1 and 3650
Email Alerts	60 days	Enter a value between 1 and 3650
External Events	3 days	Enter a value between 1 and 3650
Flow Collection Summary	365 days	Enter a value between 1 and 3650
Flow Collection Trend	30 days	Enter a value between 1 and 3650
Host Group Alarm Severity Per Minute	8 days	Enter a value between 1 and 3650
Host Group Application Traffic	15 days of 5-minute data; 60 days of hourly data; 365 days of daily data	Enter a value between 1 and 3650
Host Group Information	60days	Enter a value between 1 and 3650
Host Group Relationship Application Traffic	15 days of 1-minute data; 60 days of hourly data; 365 days of daily data	Enter a value between 1 and 3650
Host Group Relationship Information	60 days	Enter a value between 1 and 3650

Host Group Relationship Service Traffic	15 days of 1-minute data; 60 days of hourly data; 365 days of daily data	Enter a value between 1 and 3650
Host Group Relationship Status	1 days	Enter a value between 1 and 3650
HostGroup Service Traffic	15 days of 5-minute data; 60 days of hourly data; 365 days of daily data	Enter a value between 1 and 3650
Host Group Traffic	15 days of 5-minute data; 60 days of hourly data; 365 days of daily data	Enter a value between 1 and 3650
Host Hourly Alarm Severity	8 days	Enter a value between 1 and 3650
Host Hourly Application Severity	8 days	Enter a value between 1 and 3650
Host Information	60 days	Enter a value between 1 and 3650
Identity and Device	60 days	Enter a value between 1 and 3650
Interface Application Traffic	15 days of 1-minute data; 60 days of hourly data; 365 days of daily data	Enter a value between 1 and 3650
Interface Traffic	15 days of 1-minute data; 60 days of hourly data; 365 days of daily data	Enter a value between 1 and 3650

Scheduled Document Archives	60 days	Enter a value between 1 and 3650
Security Analytics and Logging (OnPrem) Daily Usage	365 days Note: Retention policy for the Manager database (Vertica) used for licensing .	Enter a value between 1 and 3650
Security Analytics and Logging (OnPrem) Incremental Usage	30 days Note: Retention policy for the Manager database (Vertica) used by Report Builder .	Enter a value between 1 and 3650
SMC Performance (Manager Performance)	60 days	Enter a value between 1 and 3650
Touched Hosts	60 days	Enter a value between 1 and 3650
User Identity	60 days	Enter a value between 1 and 3650

Related Topics

Data Retention: Overview

Data Retention: Configure Disk Space

If the free disk space percentage on the Manager falls to 20 percent, the Manager Disk Space Low alarm is generated. This alarm becomes inactive when the free disk space rises above 20 percent. This enables you to adjust the retention values yourself so that you keep the data that is most valuable to you.

If the free disk space percentage on the Manager drops below 15 percent, the Manager automatically begins to free up disk space by deleting and condensing data beyond your retention settings. An alarm note is appended to the alarm to list these changes to the data retention settings. You can view alarm notes in the Alarm Table or the Manager Alarm Table. For more information, see the Alarm Notes topic in the Desktop Client.

Related Topics

[Data Retention: Overview](#)

DSCP Configuration: Overview



Only users assigned the role of Admin or Configuration Manager can manage DSCP configurations.

Use this page to manage the Differentiated Services Code Point (DSCP) configurations for your network. You can set up different configurations, each with different percentages or maximum bandwidth, and you can then choose which configurations to use for your various exporters and interfaces. You can use these configurations within any domain. Your goal is to reduce packet loss and latency on your network by increasing bandwidth for functionality or applications that require more resources than others.

Open DSCP Configurations

1. From the main menu, choose **Configure > Global > Manager**.

The Data Retention tab is selected by default. You can view detailed data retention information associated with the specific IP address and appliance referenced at the top of the page.

2. Click the DSCP Configuration tab.

DSCP Configurations: Table of Contents

DSCP Configuration: Sort Columns

By default, the columns in the DSCP Configuration table are sorted in ascending order. To reverse the sort order for a column, click the column name. Use the column name to toggle between ascending and descending order. You can also click the arrow that is displayed to the right of the column name (after you click the column name). An up arrow indicates that the column is sorted in ascending order, and a down arrow indicates descending order.

Related Topics

[DSCP Configuration: Manage DSCP Configurations Overview](#)

[DSCP Configuration: Overview](#)

DSCP Configuration: Manage DSCP Configurations Overview

You can perform the following actions for DSCP configurations:

- [DSCP Configuration: View Configuration Details](#)
- [DSCP Configuration: Add a DSCP Configuration](#)
- [DSCP Configuration: Edit a DSCP Configuration](#)
- [DSCP Configuration: Delete a DSCP Configuration](#)

Related Topics

[DSCP Configuration: Overview](#)

DSCP Configuration: View Configuration Details

Manager application configuration

This section displays the name, IP address, model number, and serial number of the associated Manager.

DSCP configuration

Name In the Name drop-down list at the top of the DSCP Configuration table, choose the name of the DSCP configuration for whose details you want to view.

Make Default Check this check box if you want Secure Network Analytics to use a particular configuration by default whenever you configure your exporters and interfaces. The term *Default* will be displayed in the Name field (at the end of the name) when this DSCP configuration name is selected.



A label is displayed beside the name if it is the default configuration that Secure Network Analytics uses whenever you configure your exporters and interfaces. If you want to make a particular DSCP configuration the default, select the Make Default check box when editing.

DSCP The name of the code point type (for example, af11) assigned to the configuration. By default, this column is sorted alphabetically in ascending order.

Inbound Allocation / Outbound Allocation The amount of inbound (or outbound) traffic or the percentage of total bandwidth allocated to the associated code type.

- The sum of all the values specified by percentage in the Inbound Allocation column or Outbound Allocation column for a DSCP configuration cannot exceed 100%.
- You can use the following units of measure when specifying a byte value: K, M, G, T, P, and E.
- 1k = 1000. When you enter a number over 999, Secure Network Analytics changes that number to the nearest unit after you have clicked away from that field.

Example: 1000 changes to 1K; 1024 changes to 1.02K; 2,000,000,000 = 2 G

- You can enter the number of units (e.g., 1) with the unit of measurement immediately after it (no space).

Example: 10M, not 10 M

- You can enter decimal numbers only if you use a unit with it (e.g., 1.2M, not 1.2).
- The sum of all the values specified by bps (bits per second) in the Inbound Allocation column or Outbound Allocation column for a DSCP configuration cannot exceed 9.22 E.
- The sum of the inbound traffic cannot exceed the inbound line speed (the maximum speed for inbound traffic on the interface), and the sum of the outbound traffic cannot exceed the outbound line speed (the maximum speed for outbound traffic on the interface).
 - If you have configured the Manager with the exporter's SNMP settings, the Manager calculates Utilization percentages for the interface using the line speed provided by the exporter. If not, the Manager uses the line speed configured here. If this value is blank, the Manager uses the "Inbound speed" value specified on the Flow Collector Properties dialog.

Inbound Allocation Type / Outbound Allocation Type Use one of the following units when specifying either the Inbound or Outbound Allocation value:

- % (percent)
- bps (bits per second)

Related Tasks

[DSCP Configuration: Sort Columns](#)

Related Topics

[DSCP Configuration: Manage DSCP Configurations Overview](#)

[DSCP Configuration: Overview](#)

DSCP Configuration: Add a DSCP Configuration

1. From the Actions drop-down menu, choose **Add New**.

The DSCP Configuration table opens.

2. Enter a unique name for your configuration.
3. (Optional) If you want to designate a particular DSCP configuration as the default, select the Make Default check box. Secure Network Analytics uses this configuration whenever you configure your exporters and interfaces.
4. To edit any of the predefined code point types for this configuration, in the Actions column for the applicable type, click the  (**Edit**) icon.
5. Complete the applicable fields for the chosen code type. You can choose to configure just inbound or just outbound values.
6. Click **Apply** to save your changes.
7. Repeat Step 4 through Step 6 for each code point type you want to edit.
8. To save the DSCP configuration, click **Save**.

Related Tasks

[DSCP Configuration: Sort Columns](#)

Related Topics

[DSCP Configuration: Manage DSCP Configurations Overview](#)

[DSCP Configuration: Overview](#)

DSCP Configuration: Edit a DSCP Configuration

1. From the Name drop-down list, choose the applicable DSCP configuration name.

The DSCP Configuration table opens.

2. From the Actions drop-down menu, choose **Edit**.
3. To edit a code point type, in the Actions column for the applicable type, click the **Edit** icon.
4. After you edit a code point, make sure you click **Apply** to save your changes.
5. (Optional) If you want to designate this DSCP configuration as the default, select the Make Default check box. This is now be the configuration used whenever you configure your exporters and interfaces.
6. To save your changes, click **Save**.

Related Tasks

[DSCP Configuration: Sort Columns](#)

Related Topics

[DSCP Configuration: Manage DSCP Configurations Overview](#)

[DSCP Configuration: Overview](#)

DSCP Configuration: Delete a DSCP Configuration

1. From the Name drop-down list, choose the applicable DSCP configuration name.

The DSCP Configuration table opens.

2. From the Actions drop-down menu, choose **Delete**.

Related Tasks

[DSCP Configuration: Sort Columns](#)

Related Topics

[DSCP Configuration: Manage DSCP Configurations Overview](#)

[DSCP Configuration: Overview](#)

Failover Configuration

Use Failover Configuration to establish a failover pair between two Managers so that one of them serves as a backup console to the other.

We've provided a summary here. For a successful configuration and operation, review the requirements and follow the instructions in the [Cisco Secure Network Analytics Failover Configuration Guide](#).



If your primary Manager goes offline, please note that the Managers do not swap roles automatically. Make sure you change the Manager roles in the order shown in the [Cisco Secure Network Analytics Failover Configuration Guide](#).

Open Failover Configuration

Use the following instructions to open Failover Configuration.

1. Log in to your Manager as admin.
2. From the main menu, choose **Configure > Global > Manager**.
3. Click the **Failover Configuration** tab.

Configuring Failover

To configure your Managers as a failover pair, follow the instructions in the [Cisco Secure Network Analytics Failover Configuration Guide](#).

The guide includes details that are critical for a successful configuration, including:

- **Licenses:** Make sure both Managers have the same licenses installed. To review your licenses, go to **Smart Software Licensing** and refer to the [Cisco Secure Network Analytics Smart Software Licensing Guide](#).
- **Certificates:** To set up trust between appliances so they can communicate, make sure you save the correct certificates to the required appliance Trust Stores.
- **Backup Files:** Back up the appliances before you start the failover configuration.
- **Configuration Order:** Configure the secondary Manager before the primary Manager.
- **Changing Roles:** If your primary Manager goes offline, make sure you change the Manager roles in the order shown in the guide. The order is critical, and they do not swap roles automatically.
- **Troubleshooting:** Refer to the [Cisco Secure Network Analytics Failover Configuration Guide](#) for solutions.



For a successful configuration and operation, follow the instructions in the [Cisco Secure Network Analytics Failover Configuration Guide](#).

Primary and Secondary Roles

As part of the configuration, you will assign a primary Manager and a secondary Manager. When you save the configuration, the following occurs:

- **Primary Manager:** The primary Manager pushes its domain configuration, user settings, and policies to the secondary Manager. Use the primary Manager to manage your appliances, change appliance configurations, change passwords, define alarms, apply policies, and more.
- **Secondary Manager:** The secondary Manager deletes its configuration, so it can synchronize with the primary Manager configuration and settings. Also, the secondary Manager changes to read-only for all users, which means that you will not have access to sections of the secondary Manager and you cannot retrieve files from the secondary Manager.

For details, refer to the [Cisco Secure Network Analytics Failover Configuration Guide](#).

Redundant Site

Site Redundancy allows you to establish near-redundancy across clusters in two Cisco Secure Network Analytics sites that contain separate deployments with similar appliances. Site Redundancy enables you to maintain your domain and analytics configuration in your primary site and manually synchronize it with the redundant site. It also provides high availability protection in the event a data center loses power. With site redundancy, you will be able to log into either of the redundant clusters and see nearly the same data.

 This feature is only available to Admin and Configuration Manager roles.

Requirements

Review the following requirements before you begin your redundant site configuration.

- Create redundant Data Store domains in both your primary and redundant site using identical names. Make sure that both sites have the same number of Data Store domains and that the Data Store domain names are identical in both sites. For more information on domains, refer to [Domain Configuration: Overview](#).

 Only Data Store domains are synchronized for site redundancy. Non-Data Store domains are not synchronized.

- Ensure that the Secure Network Analytics software version is the same at both sites.
- Add your redundant Manager certificates to the primary Manager Trust Store. See [Adding Certificates to Trust Stores](#) for more information.
- Add your primary Manager certificates to the redundant Manager Trust Store. See [Adding Certificates to Trust Stores](#) for more information.

Once you have completed the requirements, you can proceed to the [Configuring a Redundant Site](#) procedure.

Adding Certificates to Trust Stores

Use the following instructions to save the required appliance identity certificates and chains to the Trust Stores.

Trust Store Requirements

The instructions will guide you through the following requirements:

- Adding the redundant Manager certificates to the primary Manager Trust Store.
- Adding the primary Manager certificates to the redundant Manager Trust Store.

Certificate Chain

If your appliance identity certificate includes a certificate chain, make sure you add the certificate chain (root and intermediate) to the Trust Stores.

Uploading Certificates to the Trust Store

Upload each file individually.

1. Download the Appliance Identity Certificates

Use the following instructions to download and save your appliance identity certificates. The steps vary based on the browser you are using.

If your certificates are already saved, you can skip this procedure. Go to **2. Add Certificates to the Manager Trust Stores**.



You can also click the lock/security icon in your browser. Follow the on-screen prompts to download your certificates. The steps vary based on the browser you are using.

1. In the browser address bar, replace the path after the IP address with the following: **/secrets/v1/server-identity**

For example: `https://<IPaddress>/secrets/v1/server-identity`

2. Follow the on-screen prompts to save the certificate.

Open: To view the file, select a text file format.

Troubleshooting: If you do not see the prompt to download the certificate, check your Downloads folder in case it was downloaded automatically, or try a different browser.

3. Repeat steps 1 and 2 on each Manager.

2. Add Certificates to the Manager Trust Stores

Use the following instructions to save your redundant Manager appliance identity certificate and chain (if applicable) to the primary Manager Trust Store.

1. Log in to your Manager.
2. Select **Configure > Global > Central Management**.
3. Confirm the Appliance Status is shown as Connected.
4. Click the **Actions** menu for the Manager.
5. Select **Edit Appliance Configuration**.
6. On the **Central Management Inventory > General** tab, locate the **Trust Store** section.
7. Click **Add New**.



Make sure you upload each appliance identity certificate and chain (root and intermediate) certificate individually.

8. In the **Friendly Name** field, enter a name for the certificate.
9. Click **Choose File**. Select the certificate.
10. Click **Add Certificate**. Confirm the certificate is shown in the Trust Store list.
11. Repeat steps 6 through 9 to add any other required certificates to the Trust Store.
 - If you are logged in to the redundant Manager, add the primary Manager certificates.
 - If you are logged in to the primary Manager, add the redundant Manager certificates.
12. Click **Apply Settings**. Follow the on-screen prompts.
13. **Connected:** On the Central Management Inventory page, confirm the Appliance Status returns to Connected.
14. Repeat steps 1 through 13 on the other Manager.

Open Site Redundancy Configuration

Use the following instructions to open Site Redundancy Configuration.

1. Log in to your Manager as admin or configuration manager.
2. From the main menu, choose **Configure > Global > Manager**.
3. Click the **Site Redundancy Configuration** tab.

Configuring a Redundant Site

Follow these steps to configure a redundant site.

1. Select the **Enable Configuration** check box.
2. Enter the Fully Qualified Domain Name (FQDN) or IP address for the Manager at your redundant site in the **Name of Manager at Redundant Site** field. Note that the Manager name must match the Common Name or Subject Alternative Name in the Manager identity certificate.
3. Click the **Save** button to save your changes.
4. Click the **Synchronize** button to synchronize your primary site with your remote site. This will synchronize your domain configuration and analytics configuration between the two sites. Site Redundancy configuration synchronization includes Data Store domain specific configuration as well as alert configuration (if enabled). Domain configuration includes Host Group Management, Alarm Severity, Policy Management, Services, Applications, Exporter SNMP profiles (not including passwords), and Domain AS Numbers. Analytics configuration includes Priorities, Country Watchlist, and Alert Expiration.

 This process should not be repeated on the redundant site.

5. A dialog box appears asking you to confirm that you want to synchronize your changes. Click **Synchronize** to continue.

You will see the "in progress" ellipsis icon indicating the synchronization is in progress. When it is complete you will see a success or failure banner.

 When you perform a synchronization, the Redundant Site Flow Collector Engine configuration is overwritten in the process. It is not recommended to synchronize more than once per hour.

Disabling a Redundant Site

Perform the following steps to disable your redundant site.

1. To disable a redundant site, de-select the **Enable Configuration** check box.
2. Click the **Save** button to save your changes. This will disable the redundant site as well as the Synchronize button.
3. (optional) Removing the site certificates of a disabled redundant site can add an additional layer of protection to your Secure Network Analytics system. If you want to remove the site certificates that you added during the [Configuring a Redundant Site](#) procedure, you can do so by performing the following steps.
 1. Log in to your Manager.
 2. Select **Configure > Global > Central Management**.

3. Confirm the Appliance Status is shown as Connected.
4. Click the **Actions** menu for the Manager.
5. Select **Edit Appliance Configuration**.
6. On the **Central Management Inventory** > **General** tab, locate the **Trust Store** section.
7. Under the **Actions** column, click **Delete** for each of the certificates you want to remove.

Troubleshooting

In the event that you encounter an issue with your site redundancy configuration, ensure the following:

- Certificates must be in place. Refer to [Adding Certificates to Trust Stores](#) for more information.
- The Secure Network Analytics software version needs to be the same at both sites.
- The number and names of your Data Store domains at both sites needs to match.

To review the log file for errors, navigate to `/lancope/var/smc/log/smc-configuration.log`

Interfaces and Exporters Cleaner

Use the Interfaces and Exporters Cleaner to remove inactive interfaces and exporters to maintain system performance and resource usage.

Only users assigned the role of Admin have access to this page and can configure the Interfaces and Exporters Cleaner.

Open the Interfaces and Exporters Cleaner

To open the Interfaces and Exporters Cleaner, do the following:

1. Log in to your Manager as admin.
2. From the main menu, choose **Configure > Global > Manager**.
3. Click the **Interfaces and Exporters Cleaner** tab.

Configure the Interfaces and Exporters Cleaner

Follow the instructions in the [System Configuration Guide](#) for configuration and troubleshooting.

Configuring the UDP Director

Use this page to configure your UDP Director, including the following:

- Configure Forwarding Rules
- Configure High Availability
- Export Forwarding Rules

For additional UDP Director configuration menus, refer to [Central Management](#). To configure High Availability, refer to the [System Configuration Guide](#).



If your UDP Director is not shown on this page, [add it to Central Management](#) using the appliance console (SystemConfig>Recovery>Add Appliance). Make sure you follow the instructions and review the requirements.

How to find this page

Do the following to access this page

From the main menu, choose **Configure> Global > UDP Director**.

What you can do on this page

Do the following on this page

Note the following when using this page:

- Make sure your UDP Director is licensed. To review license status and usage, go to [Smart Software Licensing](#) in Central Management.
- SSL is used to do the following:
 - To send messages from the Active Directory (AD) Server to the UDP Director.
 - To send messages from the UDP Director to the Manager.

Export forwarding rules

To export the forwarding rules XML file:

1. Click the **Actions** menu.
2. Select **Export Forwarding Rules**
3. On the dialog, click **Save**, and then save the file to your preferred location.

Configure forwarding rules or High Availability (HA)

 UDP does not forward over IPv6.

Make sure you configure at least one forwarding rule if you're planning to set up High Availability.

 High Availability is only available on UDP Director hardware appliances. High Availability is not available on virtual appliances. To configure High Availability, refer to the [System Configuration Guide](#).

Check connection status of a UDP Director

The connection status indicator in the Connection Status column of the UDP Directors table indicates the connection status for each added UDP Director.

- A green status indicator signifies that a connection to the UDP Director has been established.
- An orange status indicator signifies that a connection to the UDP Director has *not* been established.

Refresh the management channel status

To refresh status on the management channel, click the  (**Refresh icon**) in the Management Channel Status column.

- The green circle indicates that the management channel is connected.
- The orange circle indicates that the management channel is not connected.

Troubleshoot, if necessary

The following are the possible reasons for a failed connection

- The IP address entered is not for a UDP Director or is an IPv6 address.
- The UDP Director is a version previous to the Manager version.
- UDP Director is not the same version as the Manager.
- There is no network connection.

- The client SSL certificate cannot be found on the UDP Director.
- The Manager server certificate does not match any client certificates on the UDP Director.
- The accept connection option on the Managing Systems page of the UDP Director is turned off (and the IP address is not in the list).

Related topics

[Adding an Appliance to Central Management](#)

[Appliance Configuration](#)

[Forwarding Rules: Overview](#)

Forwarding Rules: Overview

Use this page to do the following:

- Add, edit, delete forwarding rules
- Import or export a forwarding rules XML file
- Search for forwarding rules

Open Forwarding Rules

1. From the main menu, choose **Configure > Global > UDP Director**.
2. On the UDP Director Configuration page, in the applicable row, click the **...** (**Ellipsis**) icon in the Actions column.
3. From the menu, click **Actions > Configuring Forwarding Rules**.

Forwarding Rules: Table of Contents

Forwarding Rules: Add a Rule

Complete the following steps to add a forwarding rule for the UDP Director:

1. From the main menu, choose **Configure > Global > UDP Director**.
2. On the UDP Director Configuration page, in the applicable row, click the **...** (**Ellipsis**) icon in the Actions column.
3. From the menu, click **Actions > Configure Forwarding Rules**.
4. Click **Add New Rule**.
5. In the Description field, enter a brief description that identifies the rule.
6. In the Source IP Address:Port List field, type the IP address of the device that sends data to the UDP Director followed by the port number through which the data will be sent.

Some examples are as follows:

- 10.11.16.38:5322
- 192.168.0.0/16:9000
- All:2055



- Use the syntax [IP address]:[Port Number], as shown in the previous examples.
- You can use Classless Inter-Domain Routing (CIDR) notation to enter a range of IP addresses.
- You can type "All" to accept data from any source IP address on this port.
- You can add Source IP Address:Port combinations within a rule by adding them to a new line.

7. In the Destination IP Address field, enter the IP address of the device receiving data from the UDP Director.
8. In the Destination Port Number, enter the port number for the receiving device.
9. Click **Save**. The new rule is added to the table on Forwarding Rules page.
10. Do you want to sync the changes?
 - a. If yes, click the Sync button at the top of the page. The new rule is saved.
 - b. If no, click the Discard Edits button at the top of the page. When the Configuration dialogue opens, click **Yes**.

Related topics

[Forwarding Rules: Overview](#)

[Configuring the UDP Director](#)

Forwarding Rules: Edit a Rule

Complete the following steps to edit a forwarding rule for the UDP Director:

1. From the main menu, choose **Configure > Global > UDP Director**.
2. On the UDP Director Configuration page, in the applicable row, click the **...** (**Ellipsis**) icon in the Actions column.
3. From the menu, click **Actions > Configuring Forwarding Rules**.
4. In the applicable row, from the Actions menu, click **Edit**.
5. Change the items on the Forwarding Rule page.
6. Click **Save**.

Related topics

[Forwarding Rules: Overview](#)

[Configuring the UDP Director](#)

Forwarding Rules: Delete a Rule

Complete the following steps to delete a forwarding rule for the UDP Director:

1. From the main menu, choose **Configure > Global > UDP Director**.
2. On the UDP Director Configuration page, in the applicable row, click the **...** (**Ellipsis**) icon in the Actions column.
3. From the menu, click **Actions > Configuring Forwarding Rules**.
4. In the applicable row, from the Actions menu, click **Delete**.

Related topics

[Forwarding Rules: Overview](#)

[Configuring the UDP Director](#)

Forwarding Rules: Import a Rule

Save your existing rules to a different location before you import another rules file. Otherwise, the existing settings will be permanently lost.

To import a forwarding rules XML file (This XML file must follow a specific structure to be accepted.), complete the following steps:

1. From the main menu, choose **Configure > Global > UDP Director**, and from the menu choose **UDP Configuration**.
2. On the UDP Director Configuration page, in the applicable row, click the **...** (**Ellipsis**) icon in the Actions column.
3. From the menu, click **Actions > Configuring Forwarding Rules**.
4. Click **Import/Export > Import Rules**.
5. Locate the rule file you want to import.
6. Click **Import**.

Related topics

[Forwarding Rules: Overview](#)

[Configuring the UDP Director](#)

Forwarding Rules: Export a Rule

To export and save a forwarding rules file, complete the following steps:

1. From the main menu, choose **Configure > Global > UDP Director**.
2. On the UDP Director Configuration page, in the applicable row, click the **...** (**Ellipsis**) icon in the Actions column.
3. From the menu, click **Actions > Configuring Forwarding Rules**.
4. Click **Import/Export > Export Rules**.
5. Do one of the following:
 - To review the file before saving it, click **Open**.
 - To save the file, click **Save** and choose the preferred location.

Related topics

[Forwarding Rules: Overview](#)

[Configuring the UDP Director](#)

Managing External Lookup

The External Lookup feature allows you to launch a Web application (or internal asset database) to view additional information about an IP address. You can launch this Web application or database directly from the Desktop Client or the Web App.

You can also use the External Lookup feature to create shortcuts that enable you to jump quickly from the Desktop Client to the Web App.

Cisco Secure Network Analytics (formerly Stealthwatch) includes the following default Web applications (lookup options) for use with the External Lookup feature; you do not have to add them to your system:

- DShield.org
- Host Report
- OpenDNS Investigate
- Talos Reputation

Some examples of Web applications that the Admin (the default administrative user that is built into Secure Network Analytics and whose user name is *admin*) can add to view additional information about an IP address include the following:

- BigFix
- CiscoWorks
- Cisco Identity Services Engine (ISE)
- Splunk
- Tripwire
- Ziften

To add a non-default lookup option, you must use the External Lookup Configuration tool in the Web App. For more information about how to do this, refer to [Configuring External Lookup](#).

Related topics

[Configuring External Lookup](#)

[Performing an External Lookup](#)

Configuring External Lookup

Why you use this page

Use this page to do the following:

- View the lookup options you have added.
- Add, edit, delete, enable, or disable a lookup option.
- Configure the specific parameters that you want to send to a Web application. The parameters you configure are only sent if they are available for the IP address on which you are performing the lookup.

How to find this page

From the main menu, choose **Configure > Global > External Lookup**.

What you can do on this page



- DShield.org, Host Report, OpenDNS Investigate, and Talos Reputation and are included by default for use with the External Lookup feature; you do not have to add them to Secure Network Analytics. To use any other Web application with this feature, you must add it to Secure Network Analytics.

View lookup options

Do the following as they apply:

- View the list of Web applications (lookup options) to determine if this list contains the lookup options you need and that they are enabled for use with the External Lookup feature.
- To disable a lookup option so that it is not available for use with the External Lookup feature (but to retain its configuration for later use), click **Enabled** in the applicable row. The button toggles to show the status of *Disabled*. To enable this vendor in the future, click **Disabled**. The button toggles to show the status of *Enabled*.
- To edit or delete a lookup option, in the Actions column, click the **⋮ (Ellipsis)** icon to open the context menu, and then select the appropriate option.

Add lookup options and configure parameters

Click **Add External Lookup** in the External Lookup section. Refer to the following for information about configuring parameters:

- To view information about internal IP addresses in a Web application, ensure you select the "Enable lookup of internal IP addresses" check box.
- The parameters you configure appear in a Web application only if they are available for the IP address on which you are performing the lookup.
- You can map up to 20 query parameters for each lookup option.
- If you want a parameter to be required when performing a lookup using a specific Web application, select the Required check box. Every parameter you designate to be required for a specific Web application must be available for the IP address on which you are performing a lookup. If one or more of the required parameters are not available for the relevant IP address, this lookup option will not be enabled in the pop-up menu.
- The URL script builder file contains the script that configures the query parameters into the URL format that the Web application requires in order to run a query.

 You can upload scripts that do not exceed 100 KB.

- If you do not upload a script builder file, Secure Network Analytics will use the default standard query parameters shown below.

 BaseURL?[ParameterName1]=[ParameterValue1]&[ParameterName2]=[ParameterValue2]&[ParameterName3]=[ParameterValue3] (and so on for each attribute you add)
- If your query parameters do not match the standard query parameters shown previously, then you must upload your customized script builder configuration. Below are some script examples to refer to when configuring a customized script builder file.

URL and script examples

Example 1

The following URL and script examples are used for Web applications that use values without parameter names (e.g., Splunk).

```
https://splunk-ip-or-url/en-US/app/search/flashtimeline
?q=search index=* {0}&earliest=-1d&latest=now
```

```
import java.util.ArrayList;
import java.util.List;
import java.text.*;

def List<String> values = new ArrayList<String>();

vendorValues.each { valueOperand ->
    values.add(valueOperand.getFromValue().toString());
};

MessageFormat messageFormat = new MessageFormat(baseUrl);
return messageFormat.format(values.toArray());
```

Click the following link to download a .txt file of the script shown in the previous image.

[Splunk script](#)

To build the script that configures the query parameters into the URL format shown previously in this example, use the Parameter Name field entry highlighted in the image below.

External Lookup : Splunk Lookup (Source IP) ⓘ

Name: *

Splunk Lookup (Source IP)

☒ Enable lookup of internal IP addresses

Base URL: *

https://splunk-ip-or-url/en-US/app/search/flashtimeline?q=search index=" {0} &earliest=-1&la

QUERY PARAMETER MAPPING:

Parameter Name:	Stealthwatch Attribute Name:	
{0}	Source IP Address	☐ Required

URL SCRIPT BUILDER FILE UPLOAD: ⓘ

You can configure as many attributes as you need; however, ensure that you configure the same number of parameters.

Example 2

The following URL and script examples are used for Web applications that use rest-like parameters (e.g., Secure Network Analytics Host Report).

`https://lancope-smc/lc-landing-page/smc.html#/host/172.21.114.17`

```
def String query = "";
vendorValues.each { valueOperand ->

    query += valueOperand.getName() + "/";
    def String convertedStr = "";
    if (valueOperand.getFromValue() instanceof String || valueOperand.getFromValue() instanceof Integer) {
        convertedStr = valueOperand.getFromValue().toString();
    } else if (valueOperand.getFromValue() instanceof Date) {
        convertedStr = new SimpleDateFormat("yyyy-MM-dd'T'HH:mm:ss").format(valueOperand.getFromValue().time);
    }
    query += URLEncoder.encode(convertedStr, "UTF-8");
};

def char lastChar = baseUrl.charAt(baseUrl.length() - 1);
if (lastChar != '?' && lastChar != '/' && lastChar != '&') {
    baseUrl = baseUrl + "?";
};

query = baseUrl + query;
return query;
```

Click the following link to download a .txt file of the script shown in the previous image.

[Secure Network Analytics Host Report script](#)

To build the script that configures the query parameters into the URL format shown previously in this example, use the Parameter Name field entry highlighted in the image below.

External Lookup : Host Report (Source IP) ⓘ

Name: *

Host Report (Source IP)

☒ Enable lookup of internal IP addresses

Base URL: *

https://lancope-smc/lc-landing-page/smc.html/

QUERY PARAMETER MAPPING:

Parameter Name:

host

Stealthwatch Attribute Name:

Source IP Address

☒ Required

+

URL SCRIPT BUILDER FILE UPLOAD: ⓘ

Browse

Cancel

Save

What you can do next

To launch a vendor's Web application or internal asset database to view additional information about an IP address, perform an external lookup. For information about how to do this, see [Performing an External Lookup](#).

Related topics

[Managing External Lookup](#)

[Performing an External Lookup](#)

Performing an External Lookup

Why you use this feature

Use this feature to query a Web application to view additional information about an IP address.

How to use this feature

1. Open any page in either the Web App or the Desktop Client that contains the relevant IP address. Do one of the following:
 - In the Web App, do one of the following to access the menu located on most pages throughout the Web App:
 - Click the (**Ellipsis**) icon beside the applicable IP address.
 - Click the **Ellipsis** icon in the Actions column of a data table or configuration table.
 - Click a point in a graph. (In the Traffic by Peer Host Group graph on the Host Report page and the Top Host Groups by Traffic graph on the Host Group Report page, you must click a host group, a column, or the line between two host groups.)
 - In the Desktop Client, right-click the relevant IP address (there are a few locations where you cannot access the External Lookup option from an IP address).
2. In the pop-up menu that appears, click **External Lookup**. A secondary pop-up menu appears.
3. Click the desired lookup option from the secondary pop-up menu that appears in step 3. The Web application for the lookup option you selected opens (you may be prompted to log in to the Web application) and displays the query results for the IP address on which you are performing the lookup.

Every parameter you designate to be required for a specific Web application must be available for the IP address on which you are performing a lookup. If one or more of the required parameters are not available for the relevant IP address, that lookup option will not be enabled in the pop-up menu. For more information, see [Configuring External Lookup](#).

What you can do next

Depending on the information that is returned by the vendor, you might want to do one or more of the following:

- Add this IP address to a specific host group for monitoring or isolation.
- Flag the IP for additional research by an analyst or investigator.
- Launch a mitigation action against the IP address. (You must configure Secure Network Analytics so that it can do this.)

Related Topics

[Managing External Lookup](#)

[Configuring External Lookup](#)

Managing Cisco® ISE or ISE-PIC

The Secure Network Analytics system offers two different versions of the Cisco Identity Services Engine (ISE).

Cisco Identity Services Engine (ISE) enables you to do the following in Secure Network Analytics:

- Receive active user and device information from the Cisco ISE cluster.
- Request ISE to apply or remove ANC policies for a host.
- View the mitigation actions (in the audit log) that have been requested for a host.

Cisco ISE Passive Identity Connector (ISE-PIC) enables you to do the following in Secure Network Analytics:

- Receive passive user and device information from the Cisco ISE cluster.

For more information, refer to the [ISE Configuration Guide](#).

Related topics

[Configuring Cisco® ISE or ISE-PIC](#)

[Applying Mitigation Actions to a Host Using Cisco® ISE](#)

[Viewing Mitigation Actions for a Host \(Host History\)](#)

Applying Mitigation Actions to a Host Using Cisco® ISE

Use this feature to apply ANC policies to the endpoint and change the endpoint authorization status on the network according to the rules and policies configured on ISE. For more information, refer to the [ISE Configuration Guide](#).

You can use this feature if you have met the following requirements:

- You are using Cisco Identity Services Engine (ISE). Cisco Identity Services Engine Passive Identity Connector (ISE-PIC) does not support mitigation.
- You have configured a Cisco ISE mitigation node.
- You have exchanged the required certificates.
- You have enabled Adaptive Network Control (ANC) for the ISE cluster.
- You have configured ANC policies on ISE.
- You have configured authorization policies on ISE to use ANC policy classifications applied to the endpoint to determine the authorization status for the endpoint on the network.
- You have an Admin, Configuration Manager, or Power Analyst role in Secure Network Analytics.

How to find Mitigation Actions links

On the Host Report, go to the Host Summary component. Click **Edit** next to ISE ANC policies.

What you can do on this page

Do the following as necessary:

Apply ANC Policies to a Host

To apply an ANC policy to a host:

1. Click **Edit** next to ISE ANC policies on Host Summary component. The ANC policies panel will open, which shows all ISE clusters configured in the host domain, along with information about available ANC policies. The ANC policies drop-down will be available for ISE clusters where applying an ANC policy is possible.
2. Choose the ISE cluster where you want to apply the policy and select the applicable ANC policy in the ANC policy drop-down.
3. Click **Save**.

Removing ANC Policies from a Host

ANC Policies applied to the host will be shown in Host Summary component next to ISE ANC Policy.

To remove an ANC Policy from the host:

1. Click **Edit** next to ISE ANC policies. The ANC policies panel will open.
2. In the ANC policy drop-down for the ISE cluster that has the applicable ANC policy, select **No Policy Applied**.
3. Click **Save**.

Secure Network Analytics sends requests to ISE for assigning or removing ANC policies for the host based on the host IP. For the request to be successfully completed on ISE, it should have information about the session associated with the given IP in its database.

What you can do next

To [view a log of all mitigation actions](#) that have been requested for this host, in the Host Summary section click **History**.

Related Topic

Managing Cisco® ISE or ISE-PIC

Configuring Cisco® ISE or ISE-PIC

Viewing Mitigation Actions for a Host (Host History)

Configuring Cisco® ISE or ISE-PIC

Use this page to add, edit, or delete the following components:

- Cisco Identity Services Engine (ISE) and Cisco Identity Services Engine Passive Identity Connector (ISE-PIC) clusters.
- PxGrid nodes – Configure these nodes to allow Secure Network Analytics to do the following:
 - Receive user and device information from the Cisco ISE or ISE-PIC cluster.
 - Interact with Adaptive Network Control (provided by the Cisco ISE device) so you can use the mitigation functionality. For information about mitigation actions, refer to [Applying Mitigation Actions to a Host Using Cisco® ISE](#).
- If a user session has no network activity for 24 hours, it will no longer be considered an active session in Secure Network Analytics.
- For detailed instructions about how to configure ISE for Secure Network Analytics, refer to the [ISE Configuration Guide](#).

How to find this page

From the main menu, choose **Configure > Integrations > Cisco® ISE**.

Only users with Admin privileges can access this page.

What you can do next

Do the following as necessary

- To [view a log of all mitigation actions](#) that have been requested for a host, go to the Host Summary section on the Host Report and click **History**.
- To [apply mitigation actions](#) to a host, go to the Host Summary component section on the Host Report and click **Edit** next to ISE ANC Policies.

Related topics

[Managing Cisco® ISE or ISE-PIC](#)

[Viewing Mitigation Actions for a Host \(Host History\)](#)

[Applying Mitigation Actions to a Host Using Cisco® ISE](#)

Viewing Mitigation Actions for a Host (Host History)

Use this page to view a log of all mitigation actions that have been requested for a host. For more information, refer to the [ISE Configuration Guide](#).

How to find this page

Do the following to access this page

On the Host Report, go to the Host Summary component and click **History**.

What you can view on this page

View a log of all mitigation actions that have been requested for a host

The Host History page displays a table that contains the following information:

- Date and time of request
- Type of request (assign or remove ANC policies)
- User who made the request
- IP address of host
- MAC address of host
- User ID logged in as host

Entries are listed by date and time order with the most recent request listed first. A limit of 100 records is displayed. You can sort on any column.

View additional details on all mitigation actions that have been requested

1. Log in to your Manager of the primary PAN node.
2. From the Main menu, click **Operations > Reports**.
3. From the left navigation pane that opens, click **ISE Reports > Audit > Adaptive Network Control Audit**.

What you can do next

Do the following as necessary

To [manage ANC policies](#) for the host, return to the Host Summary component on this page and click **Edit** next to ISE ANC policies.

Related Topic

[Configuring Cisco® ISE or ISE-PIC](#)

[Applying Mitigation Actions to a Host Using Cisco® ISE](#)

ISE ANC Policy Assignments Report

Use this page to see a list of the hosts to which the ISE (Identity Services Engine) ANC (Adaptive Network Control) policy was applied for the last 30 days.

The advantage of configuring and applying ANC policies in [Response Management](#) is so that Secure Network Analytics can automatically assign an ANC policy to a host associated with alarms you have designated in a rule. You don't have to manually apply ANC policies from the [ANC Policies panel](#). (Note that you can only apply ANC policies from this panel, not configure them.) For more information, refer to the [ISE Configuration Guide](#).



To configure and apply ANC policies, you must have [configured Cisco® ISE](#) for your Secure Network Analytics system.

How to find this page

From the main menu, choose **Monitor > Integrations > ISE ANC Assignments**.

View the ISE ANC Policy Assignments List

The report lists up to a maximum of 1000 rows, with each row representing a single host. If a policy has been applied to a host more than once in the last 30 days, then you will see only the last occurrence in this report.

Report column definitions

Column	Description
Edit check box	To assign an ANC policy to several hosts at one time, check the check box beside each host to which you want to assign the policy. From the ANC Policy drop-down list, choose the applicable policy. These are policies that have been created in the ISE management interface. • Only policies that are available for ISE clusters where applying an ANC policy is possible are included in this list. • When applying an ANC policy to several hosts at one

Column	Description
	time, the hosts must belong to the same ISE cluster. When you choose multiple hosts that don't all belong to the same ISE cluster, you can only clear an ANC policy (by choosing the None policy).
Host IP	The host to which an ANC policy was applied. Click this entry to access the Host Report and check the history of the associated ANC policy for the host. The history shows each change that someone made to the policy.
ISE Cluster	The ISE cluster to which the host belongs.
MAC Address	The MAC address of the host. The MAC address is the unique identifier of a device.
Assignment Mode	Indicates the mode via which the policy was applied. Automatic Indicates that the policy was applied by a rule configured in Response Management. Manual Indicates that a user applied the policy from the ANC Policies panel (accessed via the Host Report).
Requested By	- If the policy was applied by a rule in Response Management, the entry is <i>Response Management</i>. - If a user applied the policy, the name of the user is displayed.
Time	The date and time the policy was applied.
Requested ANC Policy	The ANC policy that was requested at the time/date that is displayed in the Time column. If you assign a different policy using the context menu in the Assign ANC Policy column, this policy is then displayed in the Requested ANC Policy column.
Effective	Click Retrieve to see the ANC Policy that the host is

Column	Description
ANC Policy	currently assigned to in the ISE management interface.
Assign ANC Policy	To assign a different policy to a single host, click the context menu in this column and choose the ANC policy (these are policies that have been created in the ISE management interface). The policy is then displayed in the Requested ANC Policy column. Only policies that are available for ISE clusters where applying an ANC policy is possible are included in this list.

Related Topic

[Response Management: Overview](#)

[Applying Mitigation Actions to a Host Using Cisco® ISE](#)

[Viewing Mitigation Actions for a Host \(Host History\)](#)

Configuring Secure Network Analytics to Use Active Directory

Use this page to do the following:

- Configure integration with Active Directory (AD) so that user information is displayed on the [User Report](#) page.

The information is retrieved from an [identity management device](#) such as Cisco Identity Services Engine (ISE), Cisco Identity Services Engine Passive Identity Connector (ISE-PIC), or Cisco ASA. As part of a security investigation, this information can identify users and the devices they are using on the network. For this feature to work properly, the identity management device must be configured separately.

You must have one of the previously mentioned identity management devices configured for use with Secure Network Analytics; otherwise, the User Report page does not display any data.

- View any previously configured Active Directory instances (which are displayed in alphabetical order).

If you have added your users to a built-in AD container (e.g., Users), then when you configure an Active Directory server in the Web App, the Bind DN must have the canonical name (CN) set to the built-in folder (e.g., CN=username, CN=Users, DC=domain, DC=com). However, if you have added your users to a new container, then the Bind DN must have the organizational unit (OU) set to the new container name (e.g., CN=username, OU=Corporate Users, DC=domain, DC=com). These names are available from your Admin (the default administrative user that is built in to each Secure Network Analytics system and whose user name is *admin*), as well as the password for the Bind DN user.

You must have admin privileges to configure integration with Active Directory.

How to find this page

Do the following to access this page

From the main menu, choose **Configure > Integrations > Active Directory**.

The Active Directory page opens, displaying AD lookup configuration information, if any has been set up. A link is available to add a new configuration.

What you can do on this page

Add a new AD connection

When you add a new AD connection, a new firewall rule is created. The rule specifies the destination port that the new AD connection uses. When you delete an AD connection that is the only one using a specified port, the corresponding firewall rule and destination port are deleted.

When you add a new AD connection with a self-signed SSL certificate that has not been uploaded to the Manager, a warning message is displayed. Then you must quit the process, take corrective action, and resume the AD connection process. Although the AD connection was not added, any information you entered was saved on the form.

If you have more than one AD instance (e.g., if you have multiple domains), the system searches through them in the order they are listed. You can reorder them if necessary.

If you want the connection between the Secure Network Analytics product and the Active Directory to be secure, you need to select the SSL check box. In addition, the SSL must be configured on the AD server and certificates must be valid.

What you can do next

Do the following as necessary

- Edit or delete an Active Directory instance. To do this, in the applicable row, click the menu in the Actions column and select the appropriate option.
- Add another Active Directory instance. There is no limit to the number of AD instances you can add, but adding a large number may affect system performance. To do this, click **Add new configuration** in the Active Directory Lookup Configuration table.

Managing the Cisco XDR Integration

The Cisco XDR integration enables you to do the following:

- Send Secure Network Analytics security findings (alarms, alerts, and observations) to Cisco XDR for correlation with other security findings. To send alerts and observations, enable [Analytics](#).
- Use Secure Network Analytics tiles on the Cisco XDR dashboards to monitor key operational metrics, such as Top Alarming Hosts, Top Alarms By Count, Top Inside Host Groups by Traffic, and more.
- Allow Cisco XDR to request top security events from Secure Network Analytics to enrich the investigation context in Cisco XDR Threat Response workflows.

For more information about Cisco XDR, refer to the [Cisco XDR website](#).

Access the Cisco XDR Integration

From the main menu, choose **Configure > Integrations > Cisco XDR**.

Use this page to do the following:

- View your Cisco XDR integration settings.
- Add, edit, or delete Cisco XDR integration settings.
- Add or remove device registration in Cisco Security Cloud Control (SCC).

Configure the Cisco XDR Integration

Follow the instructions in the [Cisco XDR Integration Guide](#).

Send Alerts and Observations to Cisco XDR

Enable [Analytics](#). Follow the instructions in the [Cisco XDR Integration Guide](#).

Secure Network Analytics Help

Get answers now

[Learn More!](#)

Resources



[Products](#)

Get to know our products



[Documentation](#)

Official product documentation and guides



[Support](#)

Have an issue? Let us know



[Use Cases](#)

Test use cases for your needs



[Information Hub](#)

A wealth of resources to help you succeed



[Training](#)

Get the most out of Secure Network Analytics

Get help when you need it

Explore Secure Network Analytics Help to find the answers you need. We have downloadable guides, user interface help, and links to support and training resources to get you back on track.

Tips

Explore the Products area to get started.

Use the search bar to perform a full text search on any topic.

Having trouble? Check out our Support area.

Did you know?

Secure Network Analytics is the only solution that can detect malware in encrypted traffic without decryption.

XDR Analytics helps you extend your visibility to the public cloud without the need for software agents.

Secure Network Analytics integrates with many other Cisco solutions to further protect your network.

About the Cisco Secure Network Analytics Data Store

The Cisco Secure Network Analytics Data Store provides a central repository to store your network's flow data, collected by your Cisco Secure Network Analytics Flow Collector. The Data Store database is comprised of a cluster of Data Nodes, each containing a portion of your flow data, and a backup of a separate Data Node's data. Because all of your flow data is in one centralized database, as opposed to spread across multiple Flow Collectors, your Cisco Secure Network Analytics Manager (formerly Stealthwatch Management Console) can retrieve query results from the Data Store more quickly than if it queried all of your Flow Collectors separately. The Data Store database cluster provides improved fault tolerance, improved query response, and quicker graph and chart population.

In a Cisco Secure Network Analytics deployment with a Data Store, the Data Store database cluster sits between your Manager and Flow Collectors. One or more Flow Collectors ingests and deduplicates flow data, performs analysis, and reports flow data and results directly to the Data Store. The Data Store collates the information from the different data sets, storing it within the Data Store database cluster, and distributing it roughly equally to all of the Data Nodes within the Data Store database cluster. The Data Store facilitates flow data storage, keeps all of your traffic in that centralized location as opposed to spread across multiple Flow Collectors, and offers greater storage capacity than a single Flow Collector, or even multiple Flow Collectors.

Note that the Manager and Flow Collectors are configured to work with a Data Store, due to the architectural changes required, including connecting to and communicating with the Data Store's centralized database.

To resolve user-submitted queries, including graphs and charts, the Manager queries the Data Store database. The Data Store database finds matching results in the columns relevant to the query, then retrieves the matching rows and returns the query results to the Manager. The Manager generates the graph or chart without needing to collate multiple result sets from multiple Flow Collectors. This mitigates the cost of querying, as compared to querying multiple Flow Collectors, and it improves query performance.

Because of the Data Store database architecture, the Manager and all Flow Collectors must communicate with the Data Store, and they must be configured during deployment to work with the Data Store. You can have a hybrid environment with some Flow Collectors storing data locally and other Flow Collectors storing data in the Data Store.

[Refer to the System Configuration Guide for details.](#)

A Data Store can be comprised of 1, 3, or more Data Nodes (in sets of 3). Each hardware Data Node is its own hardware chassis. Virtual Edition Data Nodes are available for Data Store Virtual Edition deployments. Multiple Data Nodes can be clustered as part of your Data Store database cluster, up to a maximum of 36 Data Nodes. Refer to the [Hardware or Virtual Edition Appliance Installation Guides](#) for details.

Note the following about the Data Store:

- You need a Data Store-compatible Manager and Data Store-compatible Flow Collector to use the Data Store. See the [Secure Network Analytics Hardware and Software Version Support Matrix](#) for compatibility details. For installation instructions, refer to the [Hardware or Virtual Edition Appliance Installation Guide and the System Configuration Guide](#).
- To process and store NVM telemetry, see the [Network Visibility Module \(NVM\) Configuration Guide](#).
- If you configure a Flow Collector for Data Store compatibility, the Flow Collector interface hides certain functionality. Use Central Management to perform Flow Collector configuration and other related tasks. To monitor storage statistics, refer to the [Data Store - Database Retention](#) tab and **Report Builder**.
- Use the Web App to monitor and configure your Secure Network Analytics installation if you deploy a Data Store. The Desktop Client is incompatible with a Data Store.

About Cisco Secure Network Analytics Flow Collector NetFlow

The Cisco Secure Network Analytics Flow Collector NetFlow gathers NetFlow, cFlow, J-Flow, Packeteer 2, NetStream, and IPFIX data to provide cost-effective, behavior-based network protection. To achieve full network visibility with a traditional probe-based approach, you would need to install a probe for each router or switch on your network. This results in many costly hardware installations. Conversely, Secure Network Analytics's flow-based approach provides you with full network visibility at a fraction of the cost.

The Flow Collector aggregates high-speed network behavior data from multiple networks or network segments to deliver end-to-end protection and improve performance across geographically dispersed networks. As the Flow Collector receives data, it identifies known or unknown attacks, internal misuse, or misconfigured network devices, regardless of packet encryption or fragmentation. Once Secure Network Analytics identifies the behavior, the system can take any action you have configured it to take, if any, for that kind of behavior.

Each Flow Collector (NetFlow) can process data for as many as 1,000,000 hosts from up to 2,000 flow exporters, depending on the Flow Collector model and license restrictions.

NetFlow data is an excellent source of information for session statistics pertaining to conversational pairs, bit rates, and packet rates. However, NetFlow data lacks certain performance statistics, such as latency or response time. The Cisco Secure Network Analytics Flow Sensor and the Flow Sensor Virtual Edition (VE) provide this kind of performance information, feeding it into the Flow Collector (NetFlow) as IPFIX or NetFlow v9 packets.


Beginning with v6.2, the Flow Sensor automatically uses IPFIX to export flows. If you want it to use NetFlow v9 packets, you must change the setting.



You must have a Manager and a Flow Collector to deploy and run Secure Network Analytics.



If you configure a Flow Collector for Data Store compatibility, the Appliance Administration interface (Appliance Admin) hides certain functionality. Use Central Management to perform Flow Collector configuration and other related tasks. If you want to monitor storage statistics, refer to [Report Builder](#). See [About the Data Store](#) for more information.



About NetFlow Version 9

The [Cisco Secure Network Analytics Flow Collector NetFlow](#) serves as a central collector for flow data generated by NetFlow-enabled devices.

The main feature of the NetFlow Version 9 export format is that it is template-based. A template describes a NetFlow record format and attributes of the fields (such as type and length) within the record. The router assigns each template an ID, which is communicated to the Flow Collector (NetFlow) along with the template description. The template ID is used for all further communication from the router to the Flow Collector (NetFlow).

Templates greatly enhance the flexibility of the NetFlow record format, because they allow a NetFlow collector or display application to process NetFlow data without necessarily knowing the format of the data in advance. Templates are used to describe the type and length of individual fields within a NetFlow data record that match a template ID.

Recommended field types

The following field types are recommended for best results in an exported NetFlow V9 template in order for the Flow Collector (NetFlow) to accept the record:

Field Type	Description
LAST_SWITCHED*	System uptime at which the last packet of this flow was switched.
FIRST_SWITCHED*	System uptime at which the first packet of this flow was switched.
IN_BYTES*	Incoming counter with length N x 8 bits for number of bytes associated with an IP Flow.
IN_PKTS*	Incoming counter with length N x 8 bits for the number of packets associated with an IP Flow.
INPUT_SNMP	Input interface index; default for N is 2 but higher values could be used.
OUTPUT_SNMP	Output interface index; default for N is 2 but higher values could be used.
IPV4_SRC_ADDR/IPV6_SRC_ADDR*	IPv4 source address or IPv6 source address (use the one that applies to you).

Field Type	Description
IPV4_DST_ADDR/IPV6_DST_ADDR*	IPv4 destination address or IPv6 destination address (use the one that applies to you).
PROTOCOL*	IP protocol byte.
L4_SRC_PORT*	TCP/UDP source port number ie : FTP, Telnet, or equivalent.
L4_DST_PORT*	TCP/UDP destination port number ie: FTP, Telnet, or equivalent.
TCP_FLAGS	Cumulative of all the TCP flags seen for this flow.

* These are the minimum required fields.

About IPFIX

IPFIX is a protocol that was developed by the IETF (i.e., Internet Engineering Task Force) to provide a standard for the exporting of flow information from exporters to allow for measurement, accounting, and billing. Because the IPFIX protocol is based upon NetFlow v9, the information listed above still applies.

Required field types

The following list shows sets of IPFIX field types. One of these sets is required in an exported IPFIX template for the Flow Collector (NetFlow) to accept the record.

- NF_F_LAST_SWITCHED
- NF_F_FIRST_SWITCHED
- NF_F_SYSTEM_INIT_TIME_MILLISECONDS

Or

- NF_F_FLOW_END_MILLISECONDS
- NF_F_FLOW_START_MILLISECONDS

Or

- NF_F_FLOW_END_SECONDS
- NF_F_FLOW_START_SECONDS

Or

- NF_F_FLOW_END_MICROSECONDS
- NF_F_FLOW_START_MICROSECONDS

Or

- NF_F_FLOW_END_NANOSECONDS
- NF_F_FLOW_START_NANOSECONDS

Or

- NF_F_FLOW_END_DELTA_MICROSECONDS
- NF_F_FLOW_START_DELTA_MICROSECONDS
- NF_F_IN_BYTES NF_F_IN_PKTS
- NF_F_SRC_ADDR_IPV4
- NF_F_DST_ADDR_IPV4

Or

- NF_F_SRC_ADDR_IPV6
- NF_F_DST_ADDR_IPV6
- NF_F_PROTOCOL
- NF_F_L4_SRC_PORT
- NF_F_L4_DST_PORT

The [Cisco Secure Network Analytics Flow Collector NetFlow](#) also supports the IPFIX template withdrawal management commands.

About the Cisco Secure Network Analytics Flow Collector sFlow

The Cisco Secure Network Analytics Flow Collector sFlow gathers data from existing sFlow-enabled routers and switches from network infrastructure vendors, such as Brocade, Extreme, or HP ProCurve. This Flow Collector monitors traffic from the very edge of the network.

Generally, sFlow-enabled devices send samples based on physical ports in their hardware. This capability allows for very granular statistics all the way down to the access ports to which hosts attach. With the ability to monitor every physical port on up to 1,000 sFlow-enabled devices, Secure Network Analytics offers a tremendous amount of visibility into the network and the hosts that are generating traffic.

In addition, some sFlow-enabled devices can send a payload sample when a session initiates. This feature allows for lightweight payload analysis without the need for SPANs or TAP ports.

Each Flow Collector (sFlow) can process data for as many as 1,000,000 hosts from up to 2,000 flow exporters, depending on the Flow Collector model and license restrictions.




Flow Sensors are not compatible with Flow Collectors that are configured to collect sFlow.



You must have a Manager and a Flow Collector to deploy and run Secure Network Analytics.




If you configure a Flow Collector for Data Store compatibility, the Appliance Administration interface (Appliance Admin) hides certain functionality. Use Central Management to perform Flow Collector configuration and other related tasks. If you want to monitor storage statistics, refer to [Report Builder](#). See [About the Data Store](#) for more information.



About the Cisco Secure Network Analytics Flow Sensor

The Cisco Secure Network Analytics Flow Sensor is a network appliance that operates similarly to a traditional packet capture appliance or IDS in that it plugs into a switch port analyzer (SPAN), mirror port, or Ethernet test access port (TAP).



In a traditional NetFlow environment, packets travel across the network through a NetFlow-enabled device. This device then creates a NetFlow record for the communications and sends the record to a NetFlow collector, such as the Cisco Secure Network Analytics Flow Collector NetFlow.

Some access-layer network devices do not support NetFlow exports. These devices create a “blind spot” in the network that hinders troubleshooting and may even lead to compliance/regulatory issues in some network environments.

Check your switch vendor's documentation for the flow export capabilities of your switch.

The Flow Sensor augments visibility into the following network areas:

- Where NetFlow is not available
- Where NetFlow is available, but you want deeper visibility into performance metrics and packet Data Loss (bytes)


Performance/Latency Metrics

The Flow Sensor passively captures Ethernet frames and calculates various performance metrics for each flow. It then exports them as IPFIX or NetFlow v9 records using our own flow template. Because the Flow Sensor has packet-level visibility, it can calculate round trip time (RTT), server response time (SRT), and packet loss for TCP sessions. It includes all of these additional fields in the NetFlow records that it sends to the Flow Collector (NetFlow).

By directing the Flow Sensor toward any NetFlow v9-capable flow collector, you can derive valuable detailed traffic statistics from NetFlow. When combined with the Flow Collector (NetFlow), the Flow Sensor also provides deep insight into performance metrics

and behavioral indicators. These flow performance indicators provide insight into the latency introduced by the network, as well as the server-side application. As with traditional NetFlow, the Flow Collector (NetFlow) analyzes and deduplicates the flows as they arrive.

As SYN, SYN/ACK, and ACK packets are received by the Flow Sensor, its engine times the arrival of subsequent packets to deduce the minimum, maximum, and average RTT and SRT for each flow. The RTT and SRT values are encoded into a NetFlow record and sent to the Flow Collector (NetFlow) for analysis, storage, and reporting. The following illustration depicts how these two measurements work.




The RTT measures the network, providing the round-trip time from the client to the server and back. This number is the same value you would obtain if you pinged another machine on the network from your computer.

The SRT measures the application, providing the amount of time it took the server side of a flow to answer incoming requests. In the case of an HTTP server, the SRT would reflect how long it takes the server to begin sending hypertext markup language (HTML) and images back to the client's browser.

The Flow Sensor allows you to quickly assess if the source of perceived performance latency is due to RTT or SRT. Due to the way the RTT/SRT algorithm works, it does not matter where the Flow Sensor is connected into your network, as long as it can see both sides of the communications.

About the Cisco Secure Network Analytics Flow Sensor Virtual Edition

The Cisco Secure Network Analytics Flow Sensor Virtual Edition does for the virtual environment what the Flow Sensor hardware version does for the physical environment. Using the same basic technology as the Flow Sensor, the Flow Sensor Virtual Edition is a virtual image that you download, rather than a physical appliance.



The Flow Sensor Virtual Edition installs inside a VM server (i.e., ESX host). Once installed, the Flow Sensor Virtual Edition connects promiscuously to the virtual switches that are configured in the VM server.

Working with the Cisco Secure Network Analytics Flow Collector NetFlow, the Flow Sensor Virtual Edition enables you to see traffic going into and out of your VM servers, as well as traffic traveling between the VMs within those servers. The Flow Sensor Virtual Edition sees this traffic and gathers information about those communications, just like the Flow Sensor does in the physical environment.

The Flow Sensor Virtual Edition calculates performance (RTT and SRT) and records which VM is being used, the VM name, its media access control (MAC) address, which virtual switch ports are being used, and all other information about the activities going on within the VM server. The Flow Sensor Virtual Edition then puts everything together into a NetFlow record and sends it to the Flow Collector (NetFlow) where it is deduplicated and processed like any other NetFlow record.



- One Flow Sensor Virtual Edition must be installed on each VM server (i.e., ESX host) that needs to be monitored.
- Refer to the [Virtual Edition Appliance Installation Guide](#) for details.

About Secure Network Analytics Manager

Cisco Secure Network Analytics (formerly Stealthwatch) provides continuous, real-time monitoring and pervasive views into all network traffic. It dramatically improves visibility across the extended network and accelerates response times for suspicious incidents.

When you install and configure Secure Network Analytics, you will configure your appliances to be managed by your primary Secure Network Analytics Manager (formerly Stealthwatch Management Console or SMC).

Use the Manager for managing your environment, investigating behavior, and responding to threats. You can also use Central Management to edit appliance configurations, update software, reboot, shut down, and more.

If you have more than one Manager, you can set up a Manager failover pair so that one of them serves as backup console to the other.

Related Topics

Tasks You Can Perform

Tasks You Can Perform

Log in to your Manager and use the web UI (also referred to as Web App) for the following tasks. You can also use the Desktop Client for certain tasks. Refer to [Tasks you can perform only in the Desktop Client](#) for more information.

• Active directory	• Flow collector configuration
• Alarm severity configuration	• Flow search
• Appliance configuration	• Host and host group monitoring
• Audit log	• Host group management and configuration
• Authentication services configuration	• Host search
• Central Manager	• Interface monitoring
• Cisco ISE	• Job management (flow search, packet query, top report)
• Copyright infringement	• Network classification
• Core policies	• Network diagrams (formerly maps)
• Custom applications	• Network monitoring
• Custom (security event) policies	• Relation management
• Data retention configuration	• Response management

• Data role configuration	• Response policies
• DSCP configuration	• Services configuration
• Domain properties configuration	• Manager failover
• External lookup	• Top reports
• Exporter SNMP configuration	• UDP configuration
• Flow analysis	• User Management

Tasks you can perform only in the Desktop Client

- Acknowledge alarms
- Build dashboards
- Configure mitigation
- Interface group configuration and reporting
- Manage Desktop Client documents (share, schedule)

About the Cisco Secure Network Analytics UDP Director


The Cisco Secure Network Analytics FlowReplicator has been renamed as the Cisco Secure Network Analytics UDP Director™. For the near future, the Help will use both names when referring to this appliance.

The UDP Director™ (also known as FlowReplicator) is a high-speed, high-performance UDP packet replicator. The UDP Director is very helpful in redistributing NetFlow, sFlow, syslog, or Simple Network Management Protocol (SNMP) traps to various collectors.

The UDP Director collates UDP messages from multiple sources, and then forwards them to multiple destinations.

The data is sent to the appropriate destination based on rules that the user configures. You can specify destinations based on source IP address, destination IP address, and destination port.

All devices across the enterprise can have a single, standardized destination (the UDP Director). Flow, syslog, and SNMP trap information can be received wherever it is needed.



A note about load balancers

If your UDP Director™ (also known as FlowReplicator) uses a load balancer, the script for checking the status or health of the UDP Director should be:

```
wget -q -O - --no-check-certificate --http-user=admin --http-passwd=[admin  
password] https://IP/fr/health
```

If you updated your UDP Director from v5.6, you will need to change the health status script to the one shown above.


Integrations

Cisco Secure Network Analytics (formerly Stealthwatch) integrates with many other Cisco Security solutions.

Cisco Encrypted Traffic Analytics

Traditional flow monitoring provides a high-level view of network communications by reporting the addresses, ports and byte and packet counts of a flow. In addition, intraflow metadata, or information about events that occur inside of a flow, can be collected, stored and analyzed within a flow monitoring framework. This data is especially valuable when traffic is encrypted, because deep-packet inspection is no longer viable. This intraflow metadata, called Cisco encrypted traffic analytics, is derived by using new types of data elements or telemetry that are independent of protocol details, such as the lengths and arrival times of messages within a flow. These data elements have the attractive property of applying equally well to both encrypted and unencrypted flows.

Using these data elements or intraflow telemetry to identify malware communication in encrypted traffic means encrypted traffic analytics can maintain the integrity of the encrypted flow without the need for bulk decryption. For more information, refer to the [ETA Cryptographic Audit Overview Help topic](#) or the [ETA Cryptographic Audit Release Notes](#).



Cisco Identity Services Engine (ISE)

Gain awareness of everything hitting your network. Provide access consistently and efficiently. Relieve the stress of complex access management. For more information, refer to the [ISE Configuration Guide](#).

Cisco XDR

Cisco XDR connects the breadth of Cisco's integrated security portfolio and your entire security infrastructure for a consistent experience that unifies visibility, enables automation, and strengthens your security across network, endpoint, cloud, and applications. The result is simplified security, built into the solutions you already have. For more information about Cisco XDR, refer to the [Cisco XDR Integration Guide](#) or the [Cisco XDR website](#).

App Manager

We introduced apps in v7.0.0 of Cisco Secure Network Analytics (formerly Stealthwatch). Secure Network Analytics apps are similar in concept to the apps you install on a smartphone. They are optional features that enhance and extend the capabilities of Secure Network Analytics. The release schedule for the apps is independent from the normal Secure Network Analytics upgrade process. Due to this, we can update apps as needed without having to link them with a core Secure Network Analytics release, and you can install apps without having to update your Secure Network Analytics system.

Use the App Manager page to manage your installed Secure Network Analytics apps. From this page you can install, update, uninstall, or view the status of an app. After installing an app, you can access it from the appropriate option on the dashboard in the Secure Network Analytics Web App. Your user permissions determines which apps you can view.

When you update Secure Network Analytics, the app that is currently installed is retained; however, some apps may require you to upgrade to the latest version of Secure Network Analytics. In addition, when you upgrade your Secure Network Analytics system, you may need to upgrade some or all of the apps. To learn which app version is supported by a particular version of Secure Network Analytics, see the [Secure Network Analytics Apps Version Compatibility Matrix](#).



Only a Primary Admin can install or uninstall an app.



When you update to a later version of an app, simply install the newer version over the existing version. You do not need to uninstall your existing app. If you uninstall an app, all files associated with it, including temporary files, are removed.

App Compatibility Notice

To simplify the Cisco Secure Network Analytics customer experience, only one version of a Secure Network Analytics app will be available to install at any point in time (similar to the app store model). Although we strive for maximum app compatibility, not all versions of an app will be compatible with all versions of Secure Network Analytics.

Cisco reserves the right to discontinue a Secure Network Analytics app at any time. There may be many reasons for doing so, including but not limited to the following:

1. The equivalent capabilities provided by the app are now provided elsewhere, either via a new version of the app, a new app, or via a feature in Secure Network

Analytics.

2. The capabilities provided by the app are no longer considered relevant or useful to our customer base.

If the decision is made to discontinue a Secure Network Analytics app, advance notice will be provided at least sixty days prior to the discontinuation date. Although Secure Network Analytics apps are currently included with your Secure Network Analytics license, Cisco reserves the right to charge license fees for certain Secure Network Analytics apps in the future.

How to find this page

1. From the main menu, select **Configure > Global > Central Management**.
2. Click the App Manager tab.

The App Manager page opens. A list of installed apps are listed in the Apps table.

What should you do on this page

View a list of your installed Secure Network Analytics apps

The table on this page displays any Secure Network Analytics apps you have installed, sorted by name, and their version. If an app is no longer supported or if it requires the user to upgrade to a newer version of Secure Network Analytics, this will be noted here as well. See [How do I know which version of an app I can use](#) in this topic.

How do I know which version of an app I can use

Use the App Manager page to manage your installed apps. From this page you can install, update, uninstall, or view the status of an app. Refer to the following table to learn about the possible app statuses.



Since it is possible that a newer version of an app exists and is not listed in App Manager, always check to see if a newer version is available in on Cisco Software Central.

Status	Definition	Action to Take
UpToDate	Your installed app is the most current version.	No action is required.

Status	Definition	Action to Take
UpdateAvailable	You have upgraded to a new version of Secure Network Analytics. Your existing app is supported by this version of Secure Network Analytics, but a new version of this app is available.	If you desire, go to Cisco Software Central to download and install the latest version (this will replace your existing version).
UpgradeRequired	You have upgraded to a new version of Secure Network Analytics, and your existing app is not supported by the Secure Network Analytics version you are now using.	To continue using this app, go to Cisco Software Central to download and install the latest version (this will replace your existing version).
AppNotSupported	You have upgraded to a new version of Secure Network Analytics. This app may no longer be supported by the version of Secure Network Analytics you are now using. It could be that this app has been deprecated or a newer version of this app has not yet been released.	Go to Cisco Software Central to see if a new version has been released.
NewApp	This is a new app.	If you desire, install this new app using Central Manager.
Error	The installation, upgrade, or removal process for the associated app has not successfully completed.	Contact Cisco Support . A partial installation, upgrade, or removal of this app may have occurred. If

Status	Definition	Action to Take
		so, this must be corrected.

Install a Secure Network Analytics app

When an app is available for a particular Secure Network Analytics version, it is posted in [Cisco Software Central](#). To install a Secure Network Analytics app, complete the following processes:

Download the SWU file from Cisco Software Central

1. Go to [Cisco Software Central](#).
2. In the Download and Upgrade section, click **Access downloads**.
3. In the **Select a Product** search bar, enter **Secure Network Analytics** and press **Enter**.
4. Choose **Secure Network Analytics Manager 2210** from the list.
5. Choose the applicable app from the list.
6. In the window on the right, click the  (**Download**) icon for the app SWU file and download to your choice of location.

Upload and install the app on Central Manager



- It usually takes a few minutes to upload and install an app.
- Only the system administrator can upload and install apps.

1. Verify that you are installing an app that is compatible with your current version of Secure Network Analytics. See the [Secure Network Analytics Apps Version Compatibility Matrix](#).
2. On the App Manager tab, click **Browse** to select the SWU file.
3. Select the app file.
The upload and installation process automatically begins.
4. (Conditional) If you need to cancel the upload process, click **Cancel** in the Upload dialog.

After you install the app, you can access it from the main menu under the **Dashboards** menu.

Uninstall an app

Click **Uninstall** in the row that contains the app you want to uninstall.

Sort columns

You can only sort by name or status. To reverse the sort order for either column, click the column name. Use the column name to toggle between ascending and descending order. You can also click the arrow that is displayed to the right of the column name (after you click the column name). An up arrow indicates that the column is sorted in ascending order, and a down arrow indicates descending order.

Access help

To access the help for a Secure Network Analytics app when you are logged in to the app,

click the  (**Help**) icon located next to the topic heading in the upper left corner of the page.

Related Topics

About Central Management

Home

The appliance **Home** page, which is refreshed every minute, displays statistics about the appliance.

The available sections and information on the Home page are determined by the appliance. Your appliance may not have all the sections and fields described here.

Display this page

From the main menu, select **Home**.

System

System

Use this section to find information about the appliance, such as the IP address, host name, version, or serial number. For more information, refer to [Field Descriptions](#) below.

Engine Status

Manager Engine Status

The Engine Status section indicates the amount of flow data this appliance has collected and/or dropped over the last five minutes and over the course of the day.

For more information, refer to [Field Descriptions](#) below.

Flow Collector Engine Status

The Engine Status section indicates the amount of flow data this appliance has collected and/or dropped over the last five minutes and over the course of the day.

For more information, refer to [Field Descriptions](#) below.

Flow Sensor VE Engine Status

If you are using the Flow Sensor Virtual Edition (VE), configure the virtual server. The Engine Status section indicates the amount of traffic data this appliance has captured and/or dropped, the packets it has processed, and the flow data it has exported over the last one minute.

It also displays any current Traffic Lost alarm that occurs. This alarm indicates that the appliance cannot handle the current amount of traffic flowing through it. This alarm is also shown on the Desktop Client.

If authentication for the Flow Sensor VE to communicate with the VM server has not been configured or has been configured incorrectly, an error message is shown above the table. The message provides a link to the virtual. Use the page to correct the authentication configuration.

For more information, refer to [Field Descriptions](#) below.

UDP Director Engine Status

The Engine Status section displays how much data (in packets per second) is flowing into and out of this UDP Director.

You can click the **More details** link to open the [Status Report](#).

Detailed Flow Statistics

The Detailed Flow Statistics table provides information that you can use to estimate your flows per second (FPS) rate. It shows current data as well as data for the previous seven days.

Using the Detailed Flow Statistics table can consume up to 30 percent of the appliance's central processing unit (CPU). Make sure you only run this process when you need it. Do not run it continuously.

You can export information in the table to a comma-separated value (CSV) file to a desired location.

For more information, refer to [Field Descriptions](#) below.

[Go to procedure.](#)

Services

Services

The Services section displays the status of the processes that are running on this appliance and how much memory they are using. For more information, refer to [Field Descriptions](#) below.

[Go to procedure.](#)

Docker Services

The Docker Services section displays the status of the Docker processes that are running on this appliance. For more information, refer to [Field Descriptions](#) below.

Network Interfaces

Network Interfaces

The Network Interfaces section displays the data for the hardware interfaces on the appliance.

For more information, refer to [Field Descriptions](#) below.

[Go to procedure](#)

Disk Usage

Disk Usage

The Disk Usage section displays information about the hard drive space on this appliance. For more information, refer to [Field Descriptions](#) below.

- **RAID:** If your appliance has a RAID array, any associated status messages appear above the table. This information is refreshed once per hour.
- **Database:** If your appliance includes a database, you can see statistics about how the database consumes the disk space on the **Support > [Database Storage Statistics](#)** page.



- **Expansion:** The data storage on a Virtual Edition (VE) appliance expands automatically when the appliance reboots. If you increase the virtual machine memory using another method on your Hypervisor host, restart the appliance after you save your changes. Refer to the [Virtual Edition Appliance Installation Guide](#) for more information.

Procedures

UDP Director Engine Status

To enable or disable the Detailed Flow Statistics table:

In the **UDP Director** Detailed Flow Statistics section, do one of the following:

- To enable the Detailed Flow Statistics table, click the Enable check box to insert a check.
- To disable the Detailed Flow Statistics, click the Enable check box to remove the check.

Export detailed flow statistics data

In the Detailed Flow Statistics section, click **Export** and choose a location.

The data is saved as a CSV file. The exported data does not include the current day's statistics.

Services



You should only stop or restart a service at the direction of **Cisco Support**. If you restart a service that is running, there may be a short period in which data is not captured.

Stop, start, reload, or restart a service

In the table row for a specific service, do one of the following:

- Click **Stop** to stop a service that is running.
- Click **Start** to start a service that is not running.

- Click **Reload** to gracefully stop the service workers and start them again (keeping the same process ID).
- Click **Restart** to stop a service that is running and then start it again (giving it a new process ID).

Network Interfaces

Configure hardware interfaces on the appliance

Please go to [Central Management Help](#) to configure these settings.

Field Descriptions

Refer to the following for information about the fields in each section.

System

Item	Description
Build	An alphanumeric code that identifies the build version of the appliance.
Domain name	The domain name that is assigned to the appliance. To change this setting after you've installed the appliance, refer to the SSL/TLS Certificates for Managed Appliances Guide .
Free Memory	The amount of memory that is currently available for use.
Host name	The host name that is assigned to the appliance. To change this setting after you've installed the appliance, refer to the SSL/TLS Certificates for Managed Appliances Guide .
IP Address	The IP address that is assigned to the appliance. To change this setting after you've installed the appliance, refer to the SSL/TLS Certificates for



Item	Description
	Managed Appliances Guide .
Load Average	A measure of the average system load over the past one-, five-, and fifteen-minute periods, respectively.
Platform	The hardware device or virtual platform on which the appliance software is installed.
Serial No.	The serial number or ID of the appliance.
Total Memory	The total amount of memory installed on the appliance.
Uptime	The amount of time since the appliance was last rebooted.
Version	The version of the software that is installed on the appliance.
VM Server CPU*	The virtual server resource allocation for the central processing unit, in the format: minimum allocated to this VM, maximum allowed for this VM.
VM Server Memory*	The virtual server resource allocation for memory, in the format: minimum allocated to this VM, maximum allowed for this VM.

* Only applies to virtual editions (VE) of the product.

Flow Collector Engine Status

Item	Description
Capture	
Flow rate	The average rate at which flows have been received

Item	Description
(fps)	from all exporters in the corresponding time period.
Flows	The total number of flows received from all exporters in the corresponding time period.
Dropped Flows	The total number of flows that were lost, usually due to network congestion, in the corresponding time period.
Dropped Flows (%)	The percentage of Flows that the Dropped Flows represents.
Process	
Flow Rate (fps)	The average rate at which flows have been processed by the appliance in the corresponding time period.
Flows	The total number of flows that have been processed by the appliance in the corresponding time period.

Flow Sensor Engine Status

Item	Description
Interface	The name of the promiscuous interface that provided the data.
Alarms	Any alarms reported by the Flow Sensor: Traffic Lost – The Flow Sensor has not seen packets in a 30-minute period from the corresponding interface.
Engine Uptime	The amount of time that has elapsed since the engine program was last restarted.



Item	Description
Capture	
Bandwidth (bps)	The rate (in bits per second) at which the Flow Sensor is receiving data on the corresponding interface.
Packet Rate (pps)	The rate (in packets per second) at which the Flow Sensor is receiving data on the corresponding interface.
Packets	The number of packets that the Flow Sensor has received on the corresponding interface since the last one-minute period boundary.
Dropped Packets	The number of packets that the Flow Sensor has dropped (as reported by the kernel) on the corresponding interface since the last one-minute period boundary.
Dropped Packets (%)	The percentage of Packets that the Dropped Packets represents.
Process	
Packet Rate (pps)	The rate (in packets per second) at which the Flow Sensor is processing data on the corresponding interface.
Packets	The number of packets that the Flow Sensor has processed on the corresponding interface since the last one-minute period boundary.
Export	
Flow Rate (fps)	The rate (in flows per second since the last one-minute period boundary) at which the Flow Sensor

Item	Description
	is exporting NetFlow flows to the collector (for example, Flow Collector (NetFlow)).
Flows	The number of NetFlow flows that the Flow Sensor has exported to the collector (for example, Flow Collector (NetFlow)) since the last one-minute period boundary.

The Total row provides the sum for each column.

UDP Director Engine Status

Item	Description
Packet Rate (pps)	The rate (in packets per second) at which the UDP Director received data since the last one-minute period.
Packets	The total number of packets that the UDP Director has received today.

Detailed Flow Statistics

Item	Description
95th FPS	Shows the 95th percentile of the data for the corresponding row.
Average FPS	Shows the average flows per second of the data for the corresponding row.
Date	Date when flows per second (FPS) occurred, in order from newest data at the top to oldest data at the bottom. Each day at midnight, the data for each day is rolled to the next day. (Any data that may exist for



Item	Description
	the seventh day is deleted.). The Current row presents today's FPS, is updated every minute, and is highlighted to differentiate it from the previous days.
Errors	Shows the number of flow errors for the corresponding row.
Maximum FPS	Shows the peak flows per second of the data for the corresponding row.

Services

Item	Description
Name	The name of each service: <i>org.apache.catalina</i> is on every appliance and is the application container. Restarting this restarts the Java application. Note that there is no Stop action on this service. If this service stops, you can no longer access the appliance. <i>/usr/bin/sshd</i> is the daemon that provides remote SSH access to the appliance. This is common to all appliances, but will only appear if SSH has been enabled.
Status	The current status of the corresponding service.
PID	The process identifier of the corresponding service.
Started	The amount of time that the corresponding service has been running. If it has been less than 24 hours, the amount of time is shown. If it has been more than 24 hours, the date that the service started is shown.

Item	Description
CPU Usage	The percentage of the CPU's (central processing unit) total capacity that is being used by the corresponding service.
Memory Usage	The percentage of the appliance's total memory that is being used by the corresponding service.
Action	You should only stop or restart a service at the direction of Cisco Support. If you restart a service that is running, there may be a short period in which data is not captured. This column contains buttons that allow you to manually execute actions for the corresponding service. The following actions may be available: • Stop • Start • Reload • Restart

Docker Services

Manager Docker Services

Item	Description
Name	The name of each service: • <i>The ISE-Client</i> is a service that receives user session information and sends mitigation actions to Cisco's Identity Services Engine (ISE) through the Cisco Platform Exchange Grid (PxGrid). • <i>Flow-Aggregation-Stream</i> is a service that populates the First Seen and Last Seen fields on



Item	Description
	the Host List page. • <i>Legacy-Auth</i> is an authentication service that validates user credentials. • <i>Secure Network Analytics Policy</i> allows you to configure policies for various types of hosts that interact with your system to enable you to monitor different areas of your network. • <i>Secure Network Analytics Reporting</i> is a function that offers consolidated and customizable reports. • <i>Token-Authority</i> is a service that manages tokens (JWT) securely.
Status	The current status of the corresponding service.
Started	The amount of time that the corresponding service has been running. If it has been less than 24 hours, the amount of time is shown. If it has been more than 24 hours, the date that the service started is shown.
Action	Before you stop or restart a service, please contact Cisco Support . If you restart a service that is running, there may be a short period in which data is not captured and all logged in users will need to re-authenticate.

Flow Collector Docker Services

Item	Description
Name	The name of each service: • <i>Global Threat Analytics Data Uploader</i> forwards NetFlow to the Global Threat Analytics cloud for

Item	Description
	analysis by their supervised machine-learning algorithms to detect malware. • <i>Data Exporter</i> exports a restricted set of NetFlow for use by third-party applications that are written by professional services or customers. To enable this service, please contact Cisco Support. • <i>Flow-aggregation-stream</i> aggregates the information for the First Seen and Last Seen fields and sends it to the flow aggregation receiver. • <i>Legacy-Auth</i> is an authentication service that validates user credentials. • <i>Token-Authority</i> is a service that manages tokens (JWT) securely.
Status	The current status of the corresponding service.
Started	The amount of time that the corresponding service has been running. If it has been less than 24 hours, the amount of time is shown. If it has been more than 24 hours, the date that the service started is shown.
Action	Before you stop or restart a service, please contact Cisco Support . If you restart a service that is running, there may be a short period in which data is not captured and all logged in users will need to re-authenticate.

UDP Director and Flow Sensor Docker Services

Item	Description
Name	The name of each service: • <i>Legacy-Auth</i> is an authentication service that



Item	Description
	validates user credentials. <i>Token-Authority</i> is a service that manages tokens (JWT) securely.
Status	The current status of the corresponding service.
Started	The amount of time that the corresponding service has been running. If it has been less than 24 hours, the amount of time is shown. If it has been more than 24 hours, the date that the service started is shown.
Action	Before you stop or restart a service, please contact Cisco Support . If you restart a service that is running, there may be a short period in which data is not captured and all logged in users will need to re-authenticate.

Network Interfaces

Item	Description
Name	The text identifier for the interface.
Up	The state of the interface, as follows: - True – Enabled - False – Disabled
IP Address	The IP address associated with the interface. Only the management port (eth0) will show an IP address.
MAC Address	The MAC address associated with the interface.
Auto Neg	The state of the auto-negotiation setting, either:

Item	Description
	• True – The setting is On. • False – The setting is Off.
Link Speed (bps)	The speed at which the interface is sending and receiving data.
Duplex	The duplex setting for the interface, either: • Full – for Full Duplex • Half – for Half Duplex
MTU	The maximum packet size the interface is allowed to transmit per transaction. • Eth0: To configure the MTU for the eth0 network interface on an appliance, refer to the System Configuration Guide. • Monitoring Interface: The appliance automatically sets the monitoring interface MTU to the maximum allowed by each interface.
RX	Statistics for the received data according to the following characteristics: • Bytes – The amount of undamaged data received, in bytes. • Packets – The amount of undamaged data received, in packets. • Errors – The number of damaged packets in the received data. • Dropped – The number of dropped packets in the received data.
TX	Statistics for the transmitted data according to the



Item	Description
	following characteristics: • Bytes – The amount of undamaged data transmitted, in bytes. • Packets – The amount of undamaged data transmitted, in packets. • Errors – The number of damaged packets in the transmitted data. • Dropped – The number of dropped packets in the transmitted data.
Action	Includes an Edit link to another page that allows you to manually configure the interface.

Disk Usage

Item	Description
Name	The location in the appliance's directory structure where the corresponding data exists.
Used	The amount of available space (in percentage) occupied by the corresponding data. When the percentage reaches 75%, the row turns red. If this occurs: • On a physical Manager – you need to adjust your data retention settings. Refer to Data Retention: Edit Data Retention Values. • On a physical Flow Collector – the data is dropped. • On a virtual appliance – the data is dropped.
Size (byte)	The amount of space (in bytes) allocated for the

Item	Description
	corresponding data.
Used (byte)	The amount of space (in bytes) occupied by the corresponding data.
Available (byte)	The amount of unused space (in bytes) for the corresponding data.

Related Topics

Network Interface Configuration



Network Interface Configuration



Please go to [Central Management Help](#) to configure these settings.

Status Report

Use the Status Report to see the following data:

- **Inbound Sources:** Displays traffic received from exporters (traffic that is sent to the UDP Director).
- **Outbound Destinations:** Displays traffic sent by the UDP Director to other devices.

To display this page, do the following:

From the main menu, select **Home** then click **More details** in the Engine Status section.

Field Descriptions

Refer to the following for information about the fields in each section.

Inbound Sources

This table displays traffic received from exporters using the following columns:

Item	Description
Source	DNS name of the device (if it is available) and the IP address of the device that sent the information to the UDP Director.
Packet Rate for Last Minute (pps)	Number of UDP packets received by the UDP Director per second from each input device during the last minute.
Packets Today	Total number of UDP packets received by the UDP Director from each input device since the last archive hour.



Outbound Destinations

This table displays traffic sent by the UDP Director to other devices using the following columns:

Item	Description
Destination	DNS name of the device (if it is available) and the IP address of the device to which the UDP Director has sent the information.
Packet Rate for Last Minute (pps)	Number of UDP packets sent by the UDP Director per second to each receiving device during the last minute.
Packets Today	Total number of UDP packets sent by the UDP Director to each receiving device since the last archive hour.

Naming and DNS

Use the Naming and DNS page to review the cache and IP lookup settings for the host where you are installing the appliance.

 Please go to [Central Management Help](#) to configure these settings.

Clear the cache

- To clear all cached DNS entries from the appliance, select **Clear entire cache**.
- To clear all cached DNS lookup failure entries from the appliance, select **Clear failed cache entries**.

Network Host and IP Lookup

Use the Network Host and IP Lookup section to test the settings on this page by resolving a known host (IP address or domain).

Test the settings

1. In the Host name or IP Address field, enter either the name or IP address of a host or enter a domain name (e.g., yahoo.com).
2. Click **Resolve**.

The appliance attempts to resolve the host name or IP address and displays results on the DNS Lookup page.

If the "Clear entry from local cache" check box is enabled when you click **Resolve**, the appliance replaces the previous entry in the cache for this host with current information and displays the results on the DNS Lookup page.

Naming and DNS

Use the Naming and DNS page to review the cache and IP lookup settings for the host where you are installing the appliance.

 Please go to [Central Management Help](#) to configure these settings.

Clear the cache

- To clear all cached DNS entries from the appliance, select **Clear entire cache**.
- To clear all cached DNS lookup failure entries from the appliance, select **Clear failed cache entries**.

Network Host and IP Lookup

Use the Network Host and IP Lookup section to test the settings on this page by resolving a known host (IP address or domain).

Test the settings

1. In the Host name or IP Address field, enter either the name or IP address of a host or enter a domain name (e.g., yahoo.com).
2. Click **Resolve**.

The appliance attempts to resolve the host name or IP address and displays results on the DNS Lookup page.

If the "Clear entry from local cache" check box is enabled when you click **Resolve**, the appliance replaces the previous entry in the cache for this host with current information and displays the results on the DNS Lookup page.

System Time and NTP



Please go to [Central Management Help](#) for information and instructions.



Services



Please go to [Central Management Help](#) for information and instructions.

Alpha OIDs



If your appliance is managed by a Central Manager, go to [Central Management Help](#) for information and instructions.

The table below lists some of the Object Identifiers (alphanumeric OIDs) supported by the [SNMP Agent](#) feature:

The SNMP Agent supports a standard set of over 50 MIBS containing more than 2,700 OIDs.

Alpha OID	Data Type	Description
SNMPv2-MIB::sysName	STRING	System Hostname
IF-MIB::ifDescr.#	STRING	System Description of Network Interface (e.g. eth0, lo0)
IF-MIB::ifIndex.#	INTEGER	Index number reference
IF-MIB::ifInOctets.#	Counter32	Number of Octets Received Inbound on this Interface
IF-MIB::ifOutOctets.#	Counter32	Number of Octets Sent Outbound on this Interface
IF-MIB::ifOperStatus.#	INTEGER	Status of Interface (up =1, down=2)
HOST-RESOURCES-MIB::hrSystemProcesses	Gauge32	Total Number of Processes Running on System
HOST-RESOURCES-MIB::hrStorageDescr.#	STRING	Description of Device or Partition Name (e.g. Swap,/usr)
HOST-RESOURCES-MIB::hrStorageIndex.#	INTEGER	Numeric Id of this Device for reference in other OIDs
HOST-RESOURCES-	OID	Type of storage (Ram,



Alpha OID	Data Type	Description
MIB::hrStorageType.#		VirtualMemory, FixedDisk)
HOST-RESOURCES-MIB::hrStorageSize.#	INTEGER	Maximum Storage Capacity of Device
HOST-RESOURCES-MIB::hrStorageUsed.#	INTEGER	Storage Capacity Currently being used
HOST-RESOURCES-MIB::hrStorageAllocationUnits.#	INTEGER	Unit of storage For Device in Bytes
UCD-SNMP-MIB::dskAvail.#	INTEGER	Space Available on Partition
UCD-SNMP-MIB::dskDevice.#	STRING	Partition Device Name (e.g. /dev/sda1)
UCD-SNMP-MIB::dskErrorFlag.#	INTEGER	Flag designating whether partition is “full”
UCD-SNMP-MIB::dskErrorMsg.#	STRING	Error message reported if partition fills up
UCD-SNMP-MIB::dskIndex.#	INTEGER	Index Mapping for Disk Partitions
UCD-SNMP-MIB::dskMinimum.#	INTEGER	Configured Minimum Disk KiB Avail before “full” (unused)
UCD-SNMP-MIB::dskMinPercent.#	INTEGER	Configured Minimum Disk Avail % before “full” (2=2%)
UCD-SNMP-MIB::dskPath.#	STRING	Name of mount point for partition
UCD-SNMP-MIB::dskPercent.#	INTEGER	Current Usage Percentage for partition
UCD-SNMP-MIB::dskPercentNode.#	INTEGER	Current Inode Percentage Usage for partition

Alpha OID	Data Type	Description
UCD-SNMP-MIB::dskTotal.#	INTEGER	Maximum Capacity of this partition (KiB)
UCD-SNMP-MIB::dskUsed.#	INTEGER	Current Usage of this partition (KiB)
UCD-SNMP-MIB::laConfig.#	STRING	Load Average before snmp error is set
UCD-SNMP-MIB::laErrorMessage.#	STRING	Error message if configured load average is exceeded
UCD-SNMP-MIB::laErrorFlag.#	INTEGER	Flag designating whether load average has been exceeded
UCD-SNMP-MIB::laIndex.#	INTEGER	Index Mapping for Load Average Stats
UCD-SNMP-MIB::laLoad.#	STRING	Current Load Average for the Interval
UCD-SNMP-MIB::laLoadFloat.#	FLOAT	Current Load Average for the Interval
UCD-SNMP-MIB::laLoadInt.#	INTEGER	Current Load Average * 100 rounded to nearest integer
UCD-SNMP-MIB::laNames.#	STRING	Load Interval (Load-15 is 15 minute load average)
UCD-SNMP-MIB::prCount.#	INTEGER	Number of Processes running for the specified process name
UCD-SNMP-MIB::prErrFix.#	INTEGER	Unused
UCD-SNMP-MIB::prErrFixCmd.#	STRING	Unused
UCD-SNMP-MIB::prErrorMessage.#	STRING	Error Message if process is outside configured limits



Alpha OID	Data Type	Description
UCD-SNMP-MIB::prErrorFlag.#	INTEGER	Flag Designating whether process is outside limits
UCD-SNMP-MIB::prIndex.#	INTEGER	Index Mapping Process Names
UCD-SNMP-MIB::prMax.#	INTEGER	Maximum Number of Processes before Error (0=no limit)
UCD-SNMP-MIB::prNames.#	STRING	Name of process being monitored
UCD-SNMP-MIB::prMin.#	INTEGER	Minimum Number of Processes before error (0=no Min)

SSL Certificate

Follow the instructions in the [SSL/TLS Certificates for Managed Appliances Guide](#) for all certificate-related changes.



Your certificates are critical for your system's security. Improperly modifying your certificates can stop Secure Network Analytics appliance communications and cause data loss. Follow the instructions in the [SSL/TLS Certificates for Managed Appliances Guide](#) for all certificate-related changes.



Certificate Authority Certificates

Follow the instructions in the [SSL/TLS Certificates for Managed Appliances Guide](#) for all certificate-related changes.



Your certificates are critical for your system's security. Improperly modifying your certificates can stop Secure Network Analytics appliance communications and cause data loss. Follow the instructions in the [SSL/TLS Certificates for Managed Appliances Guide](#) for all certificate-related changes.

Forwarding Rules



For instructions, please refer to the [System Configuration Guide](#).



Enable the High Availability Cluster



High Availability is only available on UDP Director hardware appliances. High Availability is not available on virtual appliances. To configure High Availability, refer to the [System Configuration Guide](#).

Use this page to configure settings for redundant UDP Directors. Both nodes are fully redundant; however, only one node is online at a time. The online node is known as the Primary in the pair, while the offline node is the Secondary. If the Primary node in the pair should fail, the Secondary node takes over and becomes the Primary.



Single Node Data Store systems do not support High Availability. If you have a single node Data Store system and lose power to your Data Node, the Single Node database will restart automatically when power is restored and the Data Node boots.

This page includes the following sections:

- High Availability Cluster Status
- High Availability Settings
- Advanced Performance

Make sure you configure at least one [forwarding rule](#) if you're planning to set up High Availability.

Field Descriptions

Refer to the following for information about the fields in each section.

High Availability Cluster Status

All of these fields are required for configuration.

Item	Description
Node Name	The name of the UDP Director in the cluster.

Item	Description
Online Status	The Primary has a status of “online.” The Secondary has a status of “offline.”
Health Status	This indicates the status of the disk partitions. The statuses are healthy, degraded, failed, and unknown.
Primary/Secondary	This indicates whether the UDP Director is the primary or secondary one
Eth0 IP Address	The dedicated unicast IP address of the Eth0 interface on the appliance
Eth0 Virtual IP	The IP address given for the virtual IP. The Primary and Secondary must use the same Virtual IP address.
Sync Address Ring 1	The IP address for the synchronization ring 1  High availability synchronization is the process by which configurations are kept identical between the appliances.
Sync Address Ring 2	The IP address for the synchronization ring 2
Promote to Primary button	Click this button to force a failover in order to switch the roles of the appliances in the HA configuration. This button will only be enabled on the Primary node.



High Availability Settings

Item	Description
Enable High Availability check box	Check the check box to enable the corresponding fields.
Node ID	Select a Node ID to configure. Select Node 1 to configure the primary node and Node 2 to configure the secondary node.
Virtual IP Address	This must be an IP address that is unused, and may be in the same subnet as the dedicated IP address for Eth0 or in its own subset. This address allows other appliances to connect to an address that applies to either UDP Director in the cluster.
Subnet Mask	Assign a subnet mask for the Virtual IP Address. This must be entered in dotted notation (e.g. 255.255.255.0) and cannot be entered in as CIDR notation.
Shared Secret	This can be a word or phrase, and it may contain spaces, special characters, and is case sensitive. Be sure to make a note of this Secret as you will have to enter the exact same string for the shared secret in the corresponding field on the second appliance in the pair when you configure it.
Sync Ring #1(Eth2) Unicast	This is the unicast address for Sync

Item	Description
IP Address	Ring 1 for the Primary. This address should be different for the Primary from the Secondary, and they should be on the same subnet. This is a dedicated IP address.
Sync Ring #1(Eth2) Subnet IP Address	Primary and Secondary, which should be on the same subnet
Sync Ring #2(Eth3) Unicast IP Address	This is the unicast address for Sync Ring 2 for the secondary.
Sync Ring #2(Eth3) Subnet IP Address	The secondary subnet IP Address (same as Primary)
Paired Node Host Name	Enter your paired node host name address here. If you selected a Node ID of 1, enter your secondary host name here. If you selected a Node ID of 2, enter your primary host name here.
Paired Node Sync Ring #1 (eth2) IP Address	Enter the Eth 2 IP address for your paired node.
Paired Node Sync Ring #2 (eth3) IP Address	Enter the Eth3 IP address for your paired node.

Advanced Parameters

Item	Description
Enable Advanced Parameters	Check the check box to allow changes to the Advanced parameters.



Item	Description
Sync Ring #1 Multicast Address	This is the multicast address for Sync Ring 1. This address is identical on both the Primary and Secondary.
Sync Ring #2 Multicast Address	This is the multicast address for Sync Ring 2. This address is identical on both the Primary and Secondary.
Token Timeout Value	Measured in milliseconds – a token is passed back and forth between the two appliances on a sync ring. If a token is sent, but not received back by the time this timeout value has expired, the token is declared lost.
Token Retries	The number of times the cluster will retry to send a receive a token before the other node is declared dead. The default is 3.
Cluster Recheck Interval	How often the cluster will check for new configuration changes and update the cluster status (barring the Token Retries value being exceeded).

Remote File System

Use the Remote File System (RFS) page to store database backups at another location from the appliance. The remote file system process uses the Common Internet File System (CIFS) protocol to store backup files. It also uses CIFS protocol when restoring a database or configuration settings.

This menu is not available on the Flow Sensor.

Display this page

From the main menu, select **Configuration > Remote File System**.

Configure a remote file system

1. In the IP Address field, enter the IP address of the remote server.
2. In the Port Number field, enter the port number that the appliance uses to communicate with the remote server. This port is configurable. The default port number is 445.
3. In the Share Name field, enter the location on the remote server in which files will be stored.
4. In the Username field, enter the user name required to access the remote server.
5. In the Password field, enter the password associated with the user name.
6. Do one of the following:
 - Click **Test** to verify the configured settings.
 - Click **Clear Configuration** to remove the configured settings and restore the default port.
 - Click **Reset** to discard any changes.
 - Click **Apply** to apply the settings.



Proxy Ingest

Use the Flow Collector Proxy Ingest page to do the following:

- Configure the connection to a proxy server so that the proxy data can be ingested by the Flow Collector.
- Exclude IP addresses from triggering alarms. (These addresses may be related to the proxy records and you want to exclude them.)

Important:

- Make sure to use the same NTP server for the Flow Collector and for the proxy used for ingestion so that the times are synchronized.
 - If you have a Cisco proxy server, you must use the proper port number so that the proxy syslog will not be dropped.
-

To display this page:

From the main menu, select **Configuration > Proxy Ingest**.

This page contains the following sections:

Proxy Servers

This section lists the proxy servers that have been configured to collect syslog data.

Add New Proxy

Use this section to add a proxy server so that it can send proxy data to the Flow Collector.

Excluded from Alarming List

This section displays the IP addresses that have been configured to prevent the triggering of alarms.

Add Entry to Exclude

Use this section to exclude any IP address from triggering alarms.

Procedures

Add Proxy Server

Follow the instructions in the [Proxy Log Configuration Guide](#) for details.

To delete a proxy server:

1. From the Proxy Server table at the top of the page, check the Delete check box for the server you want to delete.
2. Click **Apply**. The proxy server will no longer forward syslog to the Flow Collector.

Add an Entry to Exclude

This section provides the ability to exclude IP addresses so that they do not trigger alarms. Proxies may see IP addresses that are not observed by exporters sending NetFlow to Flow Collectors. When proxy data is enabled for Flow Collectors and then forwarded to the Manager, new IP addresses could show up that were not excluded from alarming during prior configuration of the Manager. If these alarms become too numerous, then you can exclude the IP addresses.

To exclude an IP address from triggering alarms:

1. In the IP Address field, type an IP address.
2. Click **Add**. The IP Address appears in the Exclude from Alarming list.



Global Settings



Please go to [Central Management Help](#) for information and instructions.

License Management



To review license status and usage, go to **Smart Software Licensing** in Central Management.



Add/Edit/Delete Users

Use the following instructions to manage users for an individual Data Node, Flow Collector, Flow Sensor, or UDP Director.

- **Changing Users on Each Appliance:** When you add, edit, or delete a user, the changes only apply to the individual appliance. For example, if you need to add a user to all Flow Collectors, you need to log in to each Flow Collector and add the user.
- **Manager:** To manage users for the Manager, go to [User Management: User Configuration Overview](#).

Display this page

1. Log in to the Appliance Administration interface as admin.

You can log in through **Central Management Inventory > Actions > View Appliance Statistics** or through direct login:

- **URL:** https://<ApplianceIPAddress>
- **Login:** admin
- **Default Password:** lan411cope

2. Select **Manage Users > Add/Edit/Delete Users**.

Add a User

1. In the **User Name** and **Full Name** fields, enter the user information.
2. In the **Password** field, enter a password or click **Generate Password**.
3. In the **Confirm Password** field, enter the new password.
4. Click **Add User**.

Edit a User

Use the following instructions to edit the full name shown for a user.

1. In the **Current Users** section, locate the user.
2. Click **Edit**.

3. In the **Full Name** field, enter your changes.
4. Click **Save**.

Delete a User

1. In the **Current Users** section, locate the user.
2. Click **Delete**.
3. Follow the on-screen prompts to delete the user.



Change Password

Use the following instructions to change the admin password for an individual Data Node, Flow Collector, Flow Sensor, or UDP Director.

- **Manager:** To change the admin password on the Manager, follow the instructions in the [System Configuration Guide](#).
- **System Instructions:** For all system password instructions, such as changing default passwords or previous passwords, refer to the [System Configuration Guide](#).

Display the Change Password Page

1. Log in to the Appliance Administration interface as admin.

You can log in through **Central Management Inventory > Actions > View Appliance Statistics** or through direct login:

- **URL:** https://<ApplianceIPAddress>
- **Login:** admin
- **Default Password:** lan411cope

2. Select **Manage Users > Change Password**.

Change the Admin Password

 When you log in to an appliance and change the admin password, the change only applies to the individual appliance. For example, if you need to change the admin password on all Flow Collectors, you need to log in to each Flow Collector and change the password.

1. Enter the current password and new password in the fields.

- **Password Format:** Make sure the new password meets your password policy and the requirements listed on the Change Password page.

- **Generate Password:** To generate a password, click **Generate Password**.
2. In the **Confirm new password** field, enter the new password.
 3. Click **Apply**. Follow the on-screen prompts to change the password.



Backup/Restore Database (Appliance Admin Interface)

Use the Backup/Restore Database page to back up an appliance database. During the backup process, the database contents (and the system configuration) are automatically copied to the remote file system.



Before you can backup the database, configure the [remote file system](#) (the file share).

- **Backing up Databases:** If you back up the Flow Collector or Manager database, make sure you also back up the [appliance configuration](#). You will need both backups to restore these appliances completely. Also, backing up the Flow Collector database includes trimming the Flow Collector database. [Review the instructions](#) before you get started.
- **Backing up Appliance Configuration:** Go to [Central Management Help](#) for instructions to back up the appliance configuration.



Restoring Databases: Contact [Cisco Support](#) to restore a database. The Restore Database operation will overwrite your current database and configuration with the contents of the previous backup. Existing network settings are not overwritten.

Do not interrupt the restoration process after it has begun.

Display this page

From the main menu in the Appliance Admin interface, select **Support > Backup/Restore Database**.

Back up a Database

Use this section to back up the Secure Network Analytics database for the Manager, the Flow Collector (NetFlow), or the Flow Collector (sFlow). The Manager continues to monitor your system during the backup process.

The following information is backed up:

- Traffic data gathered from the monitored Secure Network Analytics appliances
- Alarm and probe suppression settings for host policies
- Most of the data collected for display in standard appliance documents

The database does not include scheduled reports.

Time

The time required to perform a database backup is approximately between 0.5 GB and 2 GB per minute. Subsequent backups will only take as long as the "differential" since the last backup.

For example, for a 2 TB system, an initial backup would take approximately 3 days. Subsequent backups (on a significantly busy system storing 60 GB per day) would take approximately 2 hours.

Starting and Stopping the Backup

If you click **Cancel** (or if the process cannot continue for some reason), the backup process stops. If the process is stopped, the process begins where it left off the next time you click the **Create Backup** button. In this case, be sure that the previously stopped backup file is not locked.

Backup Display Information

At the bottom of the page, the following information about the backup is shown:

- Time and date of the last backup
- Status of the last backup
- Duration of the last backup

Back up a Database

This process involves completing the following procedures:

Preparation: Before you back up the database, configure the [remote file system](#) (the file share).

1. Disable SNMP Polling for a Manager
2. Trim the Flow Collector Database
3. Back up the Databases
4. Re-enable SNMP Polling in the Manager



Make sure you follow the instructions and complete all procedures for the database backup. For assistance, please contact [Cisco Support](#).

1. Disable SNMP Polling for a Manager

Backing up the database can take a long time. To prevent the SNMP process from interrupting the backup, turn off SNMP polling. Then, re-enable SNMP polling after the backup finishes.

To disable SNMP polling, complete the following steps:

1. Launch the Desktop Client as the admin user (but do not close the Appliance Admin interface).
2. In the Enterprise tree, right-click an exporter.
3. Select **Configuration > Exporter SNMP Configuration**.
4. Note the entry in the **Default** field. You will re-enter this information after you back up the databases.
5. In the **Default** drop-down list, select None. SNMP polling for this domain is now off.
6. Click **OK**.
7. Repeat steps 2 through 6 for each domain on your system.

2. Trim the Flow Collector Database

The Flow Collector database backup may take multiple days to finish and will slow your network speed if the database is large. Before you back up your databases, we recommend trimming the Flow Collector database. This will free the available disk space for storing flows and reduce the amount of time it takes to back up the database.



The Flow Collector stores the maximum number of days based on the disk space and the amount of data collected per day. When the maximum (75% of the /var partition) is hit, the database will start to delete the oldest data first to allow new data to come in.

1. Review your Database Storage Statistics

Use the following instructions to check your database storage.

1. Log in to the Flow Collector Appliance Admin interface.
2. Select **Support > Database Storage Statistics**.
3. Review the days stored in Capacity, Flow Data Summary, and CI Event Data Summary (or Security Event Data Summary).

Database Storage Statistics

Capacity

	Average	Warning
Capacity in Days	50	49
Remaining Days	22	21
Bytes Per Day	549.46M	563

Flow Data Summary

Data	Days	Containers	Rows			Bytes
			Total	Average Per Day	Largest Day	
Flow Details	28	32	148.75M	5.31M	5.49M	3.4
Flow Interface Details	14	20	213.3M	15.24M	15.65M	5.9
Total	28	52	362.05M	20.55M	21.15M	9.4

CI Event Data Summary

Data	Days	Containers	Rows			Bytes
			Total	Average Per Day	Largest Day	
CI Events	28	29	351.17k	12.54k	12.85k	8.53M
CI Event Details	28	29	351.17k	12.54k	12.85k	4.06M
Total	28	58	702.34k	25.08k	25.71k	12.59M

2. Trim the Interface Details

The Flow Interface Data is the data related to the interfaces of exporters. Secure Network Analytics saves flow interface data and flow data. The Flow Interface default setting causes the system to push out the flow data, so it can keep all the interface statistics it can.

Exporter	Exporter Type	Interface	Direction	TTL	DSCP	Flow Action
Cisco	Cisco	#Index-2	Outbound			Permitted
Cisco	Cisco	#Index-3	Inbound			Permitted

Exporter	Exporter Type	Interface	Direction	TTL	DSCP	Flow Action
Cisco	Cisco	#Index-2	Inbound			Permitted
Cisco	Cisco	#Index-3	Outbound			Permitted

Backing up this data takes time. If you don't need all of it, shorten the storage limit (for example: 7 days). Any data older than the limit will be lost.

Use the following instructions to purge the database of the interface statistics data older than the limit you set, so you can free up the available disk space for storing flows.

1. Log in to your Secure Network Analytics Desktop Client as the admin user.
2. Locate the Flow Collector in the Enterprise Tree. Click the plus (+) sign to expand the container.
3. Right-click the Flow Collector. Select **Configuration > Properties**.
4. In the Flow Collector Properties dialog box, click **Advanced**.
5. Select the **Store flow interface data**.
6. Shorten the storage limit.
For example, if you set the limit to **Up to 7 days**, anything older than 7 days will be lost.
7. Click **OK**.
8. Wait 5 minutes to proceed to the next steps.

3. Trim Flow Details and CI Event Data

To reduce the size of the Flow Details & CI Event/Details in the Flow Collector database, please contact [Cisco Support](#). This step is optional, and the trimming process takes only a few minutes to complete, but the process requires guidance.




When you trim the NetFlow, you will specify the number of days to keep Flow Details & CI Event/Details in the Flow Collector database. Two things will occur with this configuration:

- The database is trimmed down to the number of days you enter.
- The database starts rolling the older data out based on the oldest day but without trying to save as much as possible.

3. Back up the Databases

To back up a Flow Collector database or Manager database to a remote file system, complete the following steps:

- **Space:** Make sure the [remote file system](#) has enough space to store the database backup.
 - **Time:** After you back up the database once, subsequent backups will be quicker because the process backs up only what has changed since the last backup. This process backs up approximately 0.5 GB to 2 GB of data per minute.
1. Return to the Appliance Admin interface (but do not close the Desktop Client).
 2. Determine how much space you will need on the [remote file system](#) to store the database backup as follows:
 - Click **Home**.
 - Locate the **Disk Usage** section.
 - Review the **Used (byte)** column for the **/lancopelvar** file system. You will need at least this much space plus 15% more on the remote file system to store the database backup.



Name	Used	Size (byte)	Used (byte)	Available (byte)
/	37%	4.92G	1.68G	2.99G
/lancopelvar	68%	37.03G	24.48G	11.79G

3. Click **Support > Backup/Restore Database**.
4. Click **Create Backup**. This process may take a long time.

- After the backup process starts, you can mouse away from the page without interrupting the process. However, if you click **Cancel** while the backup is in progress, you may not be able to resume the backup without restarting the appliance.
- Follow the on-screen prompts until the backup is completed.
- To view details of the backup process, click **View Log**.

5. Click **Close** to close the progress window.

4. Re-enable SNMP Polling in the Manager

To re-enable SNMP polling, complete the following steps:

1. Return to the Desktop Client (but do not close the Appliance Admin interface).
2. Right-click the appropriate domain and select **Configuration > Exporter SNMP Configuration**. The Exporter SNMP Configuration page for that domain opens.
3. From the Default drop-down list, select the original entry for the selected domain (refer to step 4 in [Disabling SNMP Polling](#)). SNMP polling for this domain is now re-enabled.
4. Click **OK**.
5. Repeat steps 2 through 4 in this procedure for each domain on your system.
6. Close the Desktop Client.

Restore a Database



Contact [Cisco Support](#) to restore a database.



Advanced Settings

The Advanced Settings page shows a list of settings that can be used to change the behavior of your Secure Network Analytics system. Use this page to add new settings or change or delete existing settings.



These settings should only be changed under direct instruction from [Cisco Support](#). Misconfiguration may seriously impact the performance of this Flow Collector and/or the loss of monitoring capabilities.

To display this page, do the following:

From the main menu, select **Support > Advanced Settings**.

Procedures



- After applying any changes, refresh the page to make sure your changes were accepted by the Flow Collector engine. If your changes were reverted, they might not be valid in combination with other Advanced Settings. Review your settings and try again. Contact [Cisco Support](#) for additional assistance.
- Make sure your telemetry ports are unique. If you configure duplicate telemetry ports, the ports will be reset to their internal defaults to avoid loss of flow data. For example, if NetFlow and NVM are exported to the same telemetry port, each device exporting NVM data will create an exporter on the Flow Collector and exhaust the exporter resources in the Flow Collector engine, resulting in loss of flow data.

Adding a New Advanced Setting

1. In the **Add New Option** field at the bottom of the page, enter the name that will appear in the Option Label column.
2. In the **Option** value field, enter the value to apply to the setting.

3. Click **Add**.
4. Do one of the following:
 - Click **Reset** to discard any changes.
 - Click **Apply** to apply the settings.

Changing an Advanced Setting

1. Click in the option value field for the setting you want to change.
2. Enter the new value.
3. Click outside the field to activate the Apply button.
4. Click **Apply**.
5. When the confirmation message appears, click **OK**.

Deleting an Advanced Setting

1. Click the **Delete** check box for the setting you want to remove.
2. Click **Apply**.
3. When the confirmation message appears, click **OK**.

Defining a Threshold Setting for Worm Activity and Worm Propagation Alarms



Using the worm-subnet_threshold option, you can set the minimum number of 24 subnets that can be used before the Worm Activity and Worm Propagation alarms are triggered.

1. Right-click the Flow Collector in the Enterprise tree.
2. Select **Configuration > Appliance Administration**.
3. Click **Support > Advanced Settings**.
4. Scroll to the worm-subnet_threshold row and enter the appropriate value.
5. Click **Apply**.
6. When the confirmation message appears, click **OK**.



Defining a Threshold Setting for Inactive Interfaces



There are additional steps required for configuring an interface as inactive. Make sure to follow the instructions in the [System Configuration Guide](#).

1. Log in to the Flow Collector.
2. Click **Support > Advanced Settings**.
3. Scroll to the interface_age row and enter the appropriate value (default is 2678400 seconds, which is equal to 31 days).
4. Click **Apply**.
5. When the confirmation message appears, click **OK**.

Changing the Flow Settings in a Flow Collector



Starting in the 7.4.1 release, sFlow will not be released as a separate ISO image. You can toggle a Flow Collector NetFlow to sFlow by following the steps below.

1. Log in to the Flow Collector.
2. Click **Support > Advanced Settings**.
3. In the engine_startup_mode field, enter one of the following values:
 - Default value from the model file - 0
 - NetFlow - 1
 - sFlow - 2



If the engine_startup_mode field does not appear in the Advanced Settings list, you can add it at the bottom of the page by using the Add New Option and Option Value fields.

4. Click **Apply** and then click **OK**.
5. Log in to your Manager.
6. Select **Configure > Flow Collectors**.

7. Enter one of the following numeric values in the Monitor Port field (these are industry standard default port numbers for NetFlow and sFlow. If your exporters are configured to use a non-standard port, you must use that port number instead):

- 2055 - NetFlow
- 6343 - sFlow

8. Click **Save** to save your changes.

Once the mode switch (NetFlow to sFlow or sFlow to NetFlow) completes, the following items that are based on flows from the previous mode are cleared:

- Caches: host cache, flow cache, security event cache
- Saved baseline files

You can confirm the mode switch by checking the flow trend graph on the dashboard to see if flows are being processed under the new mode.



Firewall event logs and Network Visibility Module (NVM) flows are not supported on a Flow Collector configured for sFlow.

Enabling Data Compression in the Data Store

Data compression is an option you can use to reduce bandwidth usage between a Flow Collector and the Data Store. It is especially helpful in scenarios where the network bandwidth from a Flow Collector to the Data Store is limited. By enabling compression, you may reduce this bandwidth by up to 90%. Make the following configuration changes in the Flow Collector interface to enable compression of data sent to the Data Store.



Compression is enabled by default in all new deployments/fresh installations. Upgrades from any prior version will have compression disabled by default unless the ingest_enable_compression Advanced Setting is set to 1.

1. Log in to the Flow Collector interface.
2. Click **Support > Advanced Settings**.



3. In the `ingest_enable_compression` field, enter one of the following:
 - 1 - Enable data compression.
 - 0 - Disable data compression.
4. Click **Apply** and then Click **OK** in the information window that appears. While many of the settings on this page could negatively impact performance if set incorrectly, enabling data compression can only improve system performance in regards to data transfer between a Flow Collector and the Data Store.

Enabling Network Visibility Module (NVM)

 There are additional steps required for configuring NVM. Make sure to follow the instructions in the [Network Visibility Module \(NVM\) Configuration Guide](#).

1. Log in to the Flow Collector.
2. Click **Support > Advanced Settings**.
3. In the `enable_nvm` field, enter one of the following:
 - 1 - Enable ingest of NVM flows.
 - 0 - Disable ingest of NVM flows.
4. If you want to change the port for NVM flows, enter the new value in the `nvm_netflow_port` field (default port is 2030).
5. Click **Apply** and then click **OK**.

Enabling Firewall Logs

 There are additional steps required for configuring Firewall Logs. Make sure to follow the instructions in the [Cisco Security Analytics and Logging \(On Premises\): Firewall Event Integration Guide](#).

1. Log in to the Flow Collector.
2. Click **Support > Advanced Settings**.
3. In the `enable_sal` field, enter one of the following:

- 1 - Enable ingest of Firewall event logs.
- 0 - Disable ingest of Firewall event logs.

4. Configure the following fields:

Description	Field Name	Value	Details
Port Number (required)	sal_ syslog_ port	8514 or new value	The default port is 8514. • Do not enter a port that is already reserved such as NetFlow (2055) or Firewall/Syslog (514). • Enter a port number between 1025 to 65535.
Network Detections (optional)	sal_to_ flow_ cache	1	The default is 0 (disabled). • To enable, enter 1. • If you enable Firewall Logs for network detections, you can query the data in flow searches, custom security events, and Report Builder.

5. Click **Apply** and then click **OK**.



Advanced Settings for Flow Sensors

Use the Advanced Settings page to add, change, or delete settings that affect the gathering and exporting of data. Changing these settings can affect the performance of your system.

You can click **Apply** at the end of a procedure, or you can click **Apply** after performing multiple procedures on this page.

To display this page:

From the main menu, select **Configuration > Advanced Settings**.

This page includes the following sections:

Export Packet Payload

Use this section to specify whether the Flow Sensor includes the first 26 bytes of binary payload data in the data that it sends to the Flow Collector.



This information appears in the Application Details columns of the Flow Table document within the Desktop Client.

Procedures

Enable

1. Check the **Export Packet Payload** check box.
2. Click **Apply**.

Disable

1. Uncheck the **Export Packet Payload** check box.
2. Click **Apply**.

Export Application Identification

Use this section to specify whether the Flow Sensor attempts to identify applications before sending data to the Flow Collector. After you enable this setting, you can define additional settings.

Procedures

Enter or change advanced settings

1. To specify that the Flow Sensor attempts to identify applications before sending data to the Flow Collector, select the **Export Application Identification** check box.
2. Choose one or more of the following settings:

Enabling any of these settings causes the Flow Sensor to consume more memory.

Setting	Description
Include IPv6	Specify whether or not the Flow Sensor analyzes both IPv4 and IPv6 packets. • When this setting is enabled, the Flow Sensor analyzes both IPv4 and IPv6 packets. • When this setting is disabled, the Flow Sensor analyzes only IPv4 packets.
Include HTTPS Header Data	Specify whether the Flow Sensor includes header data from HTTPS flows in the data that it sends to the Flow Collector. The data includes the SSL common name and SSL organization name. This information appears



Setting	Description
	in the Server Application Details column of the Flow Table document within the Desktop Client. This setting requires that the Flow Type is set to IPFIX.
Include HTTP Header Data	Specify whether the Flow Sensor includes header data from HTTP flows in the data that it sends to the Flow Collector. When this setting is selected, a secondary field allows you to specify the maximum length of the HTTP path (in bytes) that the Flow Sensor includes as part of the flow data. This information appears in the Client Application Details column and Server Application Details column of the Flow Table document within the Desktop Client. This setting requires that the Flow Type is set to IPFIX. The maximum setting is 256 bytes.

3. Click **Apply**.

VXLAN Decapsulation

This option enables the Flow Sensor to use Virtual Extensible Local Area Network (VXLAN) decapsulation capabilities. Without VXLAN decapsulation, the Flow Sensor simply detects VXLAN encapsulated traffic as flows between two Virtual Tunnel Endpoints (VTEPs). Decapsulation allows for much richer content by being able to analyze the tunneled traffic and thus gain greater insight into the traffic patterns in the network.



The Flow Sensor will only decapsulate VXLAN traffic which was originally sent to the standard VXLAN port (4789).

Procedures

Enable

1. Check the **Enable VXLAN Decapsulation** check box.
2. Click **Apply**.

Disable

1. Uncheck the **Enable VXLAN Decapsulation** check box.
2. Click **Apply**.

GENEVE Decapsulation

This option enables the Flow Sensor to decapsulate any Generic Network Virtualization Encapsulation (GENEVE) traffic received on its monitoring ports.

Procedures

Enable

1. Check the **Enable GENEVE Decapsulation** check box.
2. Click **Apply**.



Disable

1. Uncheck the **Enable GENEVE Decapsulation** check box.
2. Click **Apply**.

ERSPAN Decapsulation

This option enables the Flow Sensor to use Encapsulated Remote Switching Port Analyzer (ERSPAN) decapsulation capabilities, which allows the Flow Sensor to detect the ERSPAN header in packets, and then decapsulate the header and process the inner packet contents.



- You are required to assign the monitoring interface an IP address to allow termination of the ERSPAN tunnel on the Flow Sensor.
- ERSPAN decapsulation is not supported on the Flow Sensor 4210.
- Flow Sensor models may have kernel controlled or Data Plane Development Kit (DPDK) controlled interfaces. DPDK controlled interfaces support ping responses, but kernel controlled interfaces will not respond to ping due to its firewall rules.
- If Nexus switches are used to generate and send ERSPAN packets, the switches could truncate or drop jumbo packets. We recommend you configure your Nexus switch to support jumbo packets over ERSPAN.

Procedures

Enable

1. Check the **Enable ERSPAN Decapsulation** check box.
2. Click **Apply**.
3. Verify **enable_erspan_decapsulation(1)** is listed in `/lancope/var/flowsensor/log/flowsensor*.log`.

Assign the monitoring interface an IP addresses



We recommend assigning the interface's IP address to a different subnet than eth0. If you need to use eth0, follow the instructions in the [Add parameter information to flowsensor.xml](#) section.

1. Open the [Central Management Inventory](#).
2. Click the **⋮ (Ellipsis)** icon for the Flow Sensor.
3. Select **Edit Appliance Configuration**.
4. On the **Appliance** tab, scroll to the **Network Interfaces** section.
5. Add the IP address that represents the destination of the ERSPAN tunnel.
For more information about adding Network Interfaces, refer to the [Network Interfaces](#) help topic.
6. Click **Apply Settings**.
7. Reboot the Flow Sensor.
8. Wait 15 minutes for the Flow Sensor to configure the monitoring interfaces.
9. Verify the added interfaces are listed on the **Network Interfaces** section of the Flow Sensor Admin UI Home page.

Add parameter information to flowsensor.xml



Only perform these steps if you configured the interface's IP address to the eth0 subnet.

1. Edit `/lancope/var/flowsensor/config/flowsensor.xml`.
2. Modify `sysctl_arp_fix` value="1" from 0 to 1 or add one more variable if it does not exist:

```
<sysctl_arp_fix value="1" min="0" max="1"
default="0"/>

systemctl restart engine
```



Disable

1. Uncheck the **Enable ERSPAN Decapsulation** check box.
2. Click **Apply**.

X-Forwarded-For Processing

This option enables the Flow Sensor to use X-Forwarded-For (XFF) processing to identify the originating IP address of a client connecting to a web server through an HTTP proxy or a load balancer.



ETA and X-Forwarded-For Processing cannot be configured together.

Procedures

Enable

1. Check the **Enable X-Forwarded-For Processing** check box.
2. Click **Apply**.

Disable

1. Uncheck the **Enable X-Forwarded-For Processing** check box.
2. Click **Apply**.

ETA Processing

This options enables the Flow Sensor to use ETA processing to generate and transmit IDP and SPLT fields to your Manager.



- Enabling ETA increases NetFlow bandwidth usage, especially when using v9. We recommend using IPFIX for the Flow Export Format.
- ETA and X-Forwarded-For Processing cannot be configured together.
- ETA cannot be enabled on Dell or PowerEdge Flow Sensor models.

Procedures

Enable

1. Check the **Enable ETA Processing** check box.
2. Click **Apply**.

Disable

1. Uncheck the **Enable ETA Processing** check box.
2. Click **Apply**.

Load Balancing

This option enables the Flow Sensor 4000 series to distribute flow data to more than one Flow Collector. Use this option if the flow data from the Flow Sensor exceeds the capacity of one Flow Collector.

Procedures

Enable

1. Check the **Enable Load Balancing** check box.
2. Click **Apply**.

Disable

1. Uncheck the **Enable Load Balancing** check box.
2. Click **Apply**.

Monitoring Interface Selection

Use this section to specify the types of interfaces the Flow Sensor 4240 and the Flow Sensor 4300 uses. The Flow Sensor 4240 uses 2 x 40G or 4 x 10G (SFP) interfaces while the Flow Sensor 4300 uses 2 x 40G/100G or 4 x 10G (SFP).

You must be using multiple Flow Collectors and have Load Balancing enabled for this setting to work properly. Refer to the [Flow Sensor and Load Balancer Integration Guide](#) for more information on these settings.





- This option is only available on the Flow Sensor 4240 and Flow Sensor 4300.
- The default setting for the Flow Sensor 4240 is 2 x 40G while the default setting for the Flow Sensor 4300 is 2 x 40G/100G.

Procedures

Specify which mode the Flow Sensor 4240 uses

1. Select one of the following:
 - 2 x 40G
 - 4 x 10G (SFP)
2. Click **Apply**.



The Flow Sensor engine will restart if you change this option.

Specify which mode the Flow Sensor 4300 uses

1. Select one of the following:
 - 2 x 40G/100G
 - 4 x 10G (SFP)
2. Click **Apply**.



The Flow Sensor engine will restart if you change this option.

Cache Mode

Use this section to define which cache mode this Flow Sensor uses.

Procedures

To select the type of cache mode

1. Refer to the following table:

If any of the following applies...	Then select...
• Asymmetric routing is present. • Single state table for application and latency calculations. • Uses less memory. • Lower overall pps processing rates. • Results in one NetFlow event created across multiple interfaces.	Use single, shared cache mode for multiple interfaces (1K, VE models) and pairs of interfaces (3K, 4K).
• Uses more memory. • Higher overall pps processing rates. • Each interface maintains its own latency and application database. • Results in a unique NetFlow record for each interface that sees a given packet.	Use independent caches for each monitoring port

2. Click **Apply**.

Changing the IP Mode

You can change the IP mode (IPv4 or IPv6) for your Flow Sensor in one of two ways:


- Logging into your Flow Sensor, removing the old IP address, and adding the new IP address.
- Remove the Flow Sensor from Central Management, then add it back to Central Management and provide the new Flow Collector IP address.

Updating the IP Mode by Logging into the Flow Sensor

Follow these steps to update the IP mode by logging into the Flow Sensor.




1. Log in to your Flow Sensor.
2. Navigate to **Configuration > NetFlow Collectors**.



3. Add the new IP address in the **IP Address** field under the Add New Collector banner.
4. Add the port number in the **Port** field.
5. Click **Add** to add the new address.
6. Select the **Delete** check box next to the IP address you want to remove.
7. Click **Apply** to save your changes.

Updating the IP Mode by Removing and Re-Adding the Flow Sensor to Central Management

1. Log in to your Manager and select **Configure > Global > Central Management**.



2. Click the three dots next to the Flow Sensor under Actions, and then select **Remove This Appliance**.
3. Log in to your Flow Sensor appliance console (SytemConfig).
4. Select **Recovery > Add Appliance**.
5. Follow the steps to add your Flow Sensor back to Central Management.



Database Storage Statistics

Use the Database Storage Statistics page to review information about the use of space by the database on the appliance.

To display this page, do one of the following:

From the main menu, select **Support > Database Storage Statistics**.

Field Descriptions

Refer to the following for information about the fields in each section.

Capacity

The Capacity section displays statistics for the following data:

- Capacity in Days – Describes amount of space available to the database according to the number of days that the database can store data at the current rate of disk usage.
- Remaining Days – Describes the unused amount of space available to the database, according to the number of days remaining at the current rate of disk usage.
- Bytes Per Day – Lists the current rate of disk usage.

Data is displayed in a table using the following columns:

Column	Description
Average	The average amount or rate of disk space usage. The total disk capacity is a constant, but when described in terms of 'number of days', the capacity (i.e., the Capacity in Days column) increases

	as the Bytes Per Day (for example the rate of disk usage) decreases, and vice-versa.
Worst Case	The maximum amount or rate of disk space usage. The total disk capacity is a constant, but when described in terms of 'number of days', the capacity (i.e., the Capacity in Days column) increases as the Bytes Per Day (for example the rate of disk usage) decreases, and vice-versa.

Flow Data Summary

The Flow Data Summary section displays statistics *for multiple days* for the following data.

Table rows flagged with "(To Be Migrated)" indicate that the data is being moved into the new data structure. Historical queries cannot be made against any data that hasn't yet been migrated.

- Flow Details: The 1-minute statistics for the flow.
- Flow Interface Details: The 1-minute statistics for each Interface that a flow traverses.
- Total: A summary of each column.

Data is displayed in a table using the following columns:

Column	Description
Data	The type of data stored in the database.



Column	Description
Days	The amount of data stored, according to the number of days that the data represents.
Containers	A number that is used for troubleshooting database performance.
Rows	
Total	The total number of rows in the database consumed by the corresponding data.
Average Per Day	The average number of rows in the database consumed by the corresponding data over the number of days indicated in the Days column.
Largest Day	The largest number of rows in the database consumed by the corresponding data on any day over the number of days indicated in the Days column. Note that the Largest Day for each data type will not necessarily occur on the same day.
Bytes	
Total	The total number of bytes in the database consumed by the corresponding data.
Average Per Day	The average number of bytes in the database consumed by the corresponding data over the number of days indicated in the Days column.
Largest Day	The largest number of bytes in the database consumed by the corresponding data on any day over the number of days indicated in the Days column. Note that the Largest Day for each data type will

Column	Description
	not necessarily occur on the same day. You can use this data to plan for a worst-case scenario.

Security Event Data Summary

The Security Event Data Summary section displays statistics (which represent multiple days) for the following tables:

- Flex Security Events – The table of security events and their details. (This row is displayed only if the Flow Collector is receiving security event data.)
- CI Events – The (deprecated) table of security events.
- CI Event Details – The (deprecated) table of inactive security event types, ports, and protocols.
- Total – A summary of each column.

Data is displayed in a table using the following columns:

Column	Description
Data	The type of data stored in the database.
Days	The amount of data stored, according to the number of days that the data represents.
Containers	A number that is used for troubleshooting database performance.
Rows	
Total	The total number of rows in the database consumed by the corresponding data.
Average Per Day	The average number of rows in the database



Column	Description
	consumed by the corresponding data over the number of days indicated in the Days column.
Largest Day	The largest number of rows in the database consumed by the corresponding data on any day over the number of days indicated in the Days column. Note that the Largest Day for each data type will not necessarily occur on the same day.
Bytes	
Total	The total number of bytes in the database consumed by the corresponding data.
Average Per Day	The average number of bytes in the database consumed by the corresponding data over the number of days indicated in the Days column.
Largest Day	The largest number of bytes in the database consumed by the corresponding data on any day over the number of days indicated in the Days column. Note that the Largest Day for each data type will not necessarily occur on the same day. You can use this data to plan for a worst-case scenario.

Flow Data By Day

The Flow Data By Day section displays *daily statistics* for the flow data.

Data is displayed in a table using the following columns:

Column	Description
Date	The date for which the corresponding data applies.
Containers	A number that is used for troubleshooting database performance.
Rows	The total number of rows in the database consumed by the data on the corresponding date.
Used Bytes	The total number of bytes of data stored in the database on the corresponding date.

Security Event Data By Day

The Security Event Data By Day section displays *daily statistics* for the security event data.

Data is displayed in a table using the following columns:

Column	Description
Date	The date for which the corresponding data applies.
Containers	A number that is used for troubleshooting database performance.
Rows	The total number of rows in the database consumed by the data on the corresponding date.
Used Bytes	The total number of bytes of data stored in the database on the corresponding date.



Backup/Restore Database (Appliance Admin Interface)

Use the Backup/Restore Database page to back up an appliance database. During the backup process, the database contents (and the system configuration) are automatically copied to the remote file system.



Before you can backup the database, configure the [remote file system](#) (the file share).

- **Backing up Databases:** If you back up the Flow Collector or Manager database, make sure you also back up the [appliance configuration](#). You will need both backups to restore these appliances completely. Also, backing up the Flow Collector database includes trimming the Flow Collector database. [Review the instructions](#) before you get started.
- **Backing up Appliance Configuration:** Go to [Central Management Help](#) for instructions to back up the appliance configuration.



Restoring Databases: Contact [Cisco Support](#) to restore a database. The Restore Database operation will overwrite your current database and configuration with the contents of the previous backup. Existing network settings are not overwritten.

Do not interrupt the restoration process after it has begun.

Display this page

From the main menu in the Appliance Admin interface, select **Support > Backup/Restore Database**.

Back up a Database

Use this section to back up the Secure Network Analytics database for the Manager, the Flow Collector (NetFlow), or the Flow Collector (sFlow). The Manager continues to monitor your system during the backup process.

The following information is backed up:

- Traffic data gathered from the monitored Secure Network Analytics appliances
- Alarm and probe suppression settings for host policies
- Most of the data collected for display in standard appliance documents

The database does not include scheduled reports.

Time

The time required to perform a database backup is approximately between 0.5 GB and 2 GB per minute. Subsequent backups will only take as long as the "differential" since the last backup.

For example, for a 2 TB system, an initial backup would take approximately 3 days. Subsequent backups (on a significantly busy system storing 60 GB per day) would take approximately 2 hours.

Starting and Stopping the Backup

If you click **Cancel** (or if the process cannot continue for some reason), the backup process stops. If the process is stopped, the process begins where it left off the next time you click the **Create Backup** button. In this case, be sure that the previously stopped backup file is not locked.

Backup Display Information

At the bottom of the page, the following information about the backup is shown:

- Time and date of the last backup
- Status of the last backup
- Duration of the last backup

Back up a Database

This process involves completing the following procedures:

Preparation: Before you back up the database, configure the [remote file system](#) (the file share).

1. Disable SNMP Polling for a Manager



2. Trim the Flow Collector Database

3. Back up the Databases

4. Re-enable SNMP Polling in the Manager



Make sure you follow the instructions and complete all procedures for the database backup. For assistance, please contact [Cisco Support](#).

1. Disable SNMP Polling for a Manager

Backing up the database can take a long time. To prevent the SNMP process from interrupting the backup, turn off SNMP polling. Then, re-enable SNMP polling after the backup finishes.

To disable SNMP polling, complete the following steps:

1. Launch the Desktop Client as the admin user (but do not close the Appliance Admin interface).
2. In the Enterprise tree, right-click an exporter.
3. Select **Configuration > Exporter SNMP Configuration**.
4. Note the entry in the **Default** field. You will re-enter this information after you back up the databases.
5. In the **Default** drop-down list, select None. SNMP polling for this domain is now off.
6. Click **OK**.
7. Repeat steps 2 through 6 for each domain on your system.

2. Trim the Flow Collector Database

The Flow Collector database backup may take multiple days to finish and will slow your network speed if the database is large. Before you back up your databases, we recommend trimming the Flow Collector database. This will free the available disk space for storing flows and reduce the amount of time it takes to back up the database.

The Flow Collector stores the maximum number of days based on the disk space and the amount of data collected per day. When the maximum (75%

of the /var partition) is hit, the database will start to delete the oldest data first to allow new data to come in.

1. Review your Database Storage Statistics

Use the following instructions to check your database storage.

1. Log in to the Flow Collector Appliance Admin interface.
2. Select **Support > Database Storage Statistics**.
3. Review the days stored in Capacity, Flow Data Summary, and CI Event Data Summary (or Security Event Data Summary).

Database Storage Statistics

Capacity

	Average	W
Capacity in Days	50	49
Remaining Days	22	21
Bytes Per Day	549.46M	563

Flow Data Summary

	Days	Containers	Total	Average Per Day	Largest Day	Total
Flow Details	28	32	148.75M	5.31M	5.49M	3.3
Flow Interface Details	14	20	213.3M	15.24M	15.65M	5.9
Total	28	52	362.05M	20.55M	21.15M	9.4

CI Event Data Summary

	Days	Containers	Total	Average Per Day	Largest Day	Total
CI Events	28	29	351.17k	12.54k	12.85k	8.53M
CI Event Details	28	29	351.17k	12.54k	12.85k	4.06M
Total	28	58	702.34k	25.08k	25.71k	12.59M

2. Trim the Interface Details

The Flow Interface Data is the data related to the interfaces of exporters. Secure Network Analytics saves flow interface data and flow data. The Flow

Interface default setting causes the system to push out the flow data, so it can keep all the interface statistics it can.

Exporter	Exporter Type	Interface	Direction	TTL	DSCP	Flow Action
Cisco	Cisco	#Index-2	Outbound			Permitted
Cisco	Cisco	#Index-3	Inbound			Permitted

Exporter	Exporter Type	Interface	Direction	TTL	DSCP	Flow Action
Cisco	Cisco	#Index-2	Inbound			Permitted
Cisco	Cisco	#Index-3	Outbound			Permitted

Backing up this data takes time. If you don't need all of it, shorten the storage limit (for example: 7 days). Any data older than the limit will be lost.

Use the following instructions to purge the database of the interface statistics data older than the limit you set, so you can free up the available disk space for storing flows.

1. Log in to your Secure Network Analytics Desktop Client as the admin user.
2. Locate the Flow Collector in the Enterprise Tree. Click the plus (+) sign to expand the container.
3. Right-click the Flow Collector. Select **Configuration > Properties**.
4. In the Flow Collector Properties dialog box, click **Advanced**.
5. Select the **Store flow interface data**.
6. Shorten the storage limit.

For example, if you set the limit to **Up to 7 days**, anything older than 7 days will be lost.

7. Click **OK**.
8. Wait 5 minutes to proceed to the next steps.

3. Trim Flow Details and CI Event Data

To reduce the size of the Flow Details & CI Event/Details in the Flow Collector database, please contact [Cisco Support](#). This step is optional, and the trimming process takes only a few minutes to complete, but the process requires guidance.

When you trim the NetFlow, you will specify the number of days to keep Flow Details & CI Event/Details in the Flow Collector database. Two things will occur with this configuration:

- The database is trimmed down to the number of days you enter.
- The database starts rolling the older data out based on the oldest day but without trying to save as much as possible.

3. Back up the Databases

To back up a Flow Collector database or Manager database to a remote file system, complete the following steps:

- **Space:** Make sure the [remote file system](#) has enough space to store the database backup.
 - **Time:** After you back up the database once, subsequent backups will be quicker because the process backs up only what has changed since the last backup. This process backs up approximately 0.5 GB to 2 GB of data per minute.
1. Return to the Appliance Admin interface (but do not close the Desktop Client).
 2. Determine how much space you will need on the [remote file system](#) to store the database backup as follows:
 - Click **Home**.
 - Locate the **Disk Usage** section.
 - Review the **Used (byte)** column for the **/lancope/var** file system. You will need at least this much space plus 15% more on the remote file system to store the database backup.

Name	Used	Size (byte)	Used (byte)	Available (byte)
/	37%	4.92G	1.68G	2.99G
/lancope/var	68%	37.03G	24.48G	11.79G

3. Click **Support > Backup/Restore Database**.
4. Click **Create Backup**. This process may take a long time.



- After the backup process starts, you can mouse away from the page without interrupting the process. However, if you click **Cancel** while the backup is in progress, you may not be able to resume the backup without restarting the appliance.
- Follow the on-screen prompts until the backup is completed.
- To view details of the backup process, click **View Log**.

5. Click **Close** to close the progress window.

4. Re-enable SNMP Polling in the Manager

To re-enable SNMP polling, complete the following steps:

1. Return to the Desktop Client (but do not close the Appliance Admin interface).
2. Right-click the appropriate domain and select **Configuration > Exporter SNMP Configuration**. The Exporter SNMP Configuration page for that domain opens.
3. From the Default drop-down list, select the original entry for the selected domain (refer to step 4 in [Disabling SNMP Polling](#)). SNMP polling for this domain is now re-enabled.
4. Click **OK**.
5. Repeat steps 2 through 4 in this procedure for each domain on your system.
6. Close the Desktop Client.

Restore a Database



Contact [Cisco Support](#) to restore a database.

Browse Files

The Browse Files page displays the file system within the /lancope/var folder on the appliance. Looking at the files may be useful for troubleshooting purposes.

For better viewing on Windows computers, we recommend that you associate the .LOG file type (and other file types that you want to view) with WordPad.

Display this page

From the main menu, select **Support > Browse Files**.

Procedures

Browse the file system

- Click a folder name to see the contents of that folder.
- Click the Parent Directory link to move back up the structure.
- Click a file name to see that file displayed in a separate window in your browser.
- Click a table header to sort the corresponding column.

The page lists the size of each file and the date and time the file was last saved. The time is localized to the appliance.

Open the lancope/var folder

To display the file system within the /lancope/var folder, enter the following URL in your browser as shown for the appliance type:

- **Manager:** https://your_manager_ip_address/smc/files
- **Flow Sensor:** https://your_fs_ip_address/fs/files
- **UDP Director or FlowReplicator:** https://your_fr_ip_address/fr/files

Manager: To display the file system within the /lancope/var/smc/config folder on the appliance, enter the following URL in your browser:

- https://your_Manager_ip_address/Manager/files/smc/config



Packet Capture

Use the following instructions to create a packet capture file for either Manager or other appliances. You can use packet captures to view data from the [tcpdump](#) utility, which monitors network traffic by capturing and displaying packets. You can also match them against a set of criteria.

Generate a Packet Capture

You will need to log into the appliance console (SystemConfig) as a sysadmin user to create a packet capture.

1. Log in to the appliance as **sysadmin**.
2. Select **Advanced > Packet Capture**. If there are no existing packet captures, you will advance directly to the packet capture configuration menu. If you have existing packet captures, select **Create** to create a new packet capture.
3. Select an interface from the list of available interfaces for the **Interface** field. If you only have one interface, it will be shown by default.



If you use Flow Sensor appliance you can generate packet captures for DPDK and NAPATECH interfaces.

4. (optional) In the Host IP Address field, specify a single host on which to filter. Enter a host using dotted decimal format (Example: 10.0.0.1).

The resulting capture file will include both sides of communication rather than only one direction.

5. Enter a TCP or UDP port number in the **Port Filter** field. The port may be either source or destination. You can leave this field blank to set the filter to "Any".
6. In the **Duration (900 Seconds Max)** field, specify the number of seconds to perform the packet capture.



Be careful to enter a reasonable amount of time because an extremely large capture file can potentially consume all the free space on the appliance's hard disk.
This field works in conjunction with the **Packets** setting. The



capture will run until either the **Duration** or number of **Packets** is reached.

7. In the **Packets (100,000 max)** field, specify the number of packets to capture before ending the capture. This field works in conjunction with the Duration setting. The capture will run until either the Duration or number of Packets is reached.
8. Click **OK** to start the packet capture.



To stop a packet capture abruptly, you can typically press **Ctrl+C**. However, this method will not work for DPDK interfaces.

To view packet capture files, download them and view them locally. Files are stored in `/lancope/var/tcpdump`. You can log directly into the appliance or access it through the Central Manager inventory > View Appliance Statistics. Go to Support > Browse Files to download the packet capture.



Tcpdump Expressions

The sections below provide reference information on tcpdump expressions used on the [Packet Capture](#) screen.

Tcpdump Template Examples

Use the following examples to verify that your Flow Sensor is properly exporting NetFlow templates:

tcpdump to get all templates (v9 and v10):

```
tcpdump -s 100 '(udp[24:2] = 2 and udp[8:2] = 10) or  
(udp[8:2] = 9 and udp[28:2] = 0)'
```

tcpdump for IPFIX templates:

```
tcpdump -s 100 'udp[24:2] = 2'
```

If your Flow Sensor is exporting NetFlow templates, either tcpdump command should give you results similar to the following:

```
listening on eth0, link-type EN10MB (Ethernet), capture size 100 bytes  
22:01:00.341609 IP flow-sensor-hostname.16384 > flow-collector-  
hostname.2055: UDP, length 436  
22:01:00.341818 IP flow-sensor-hostname.16384 > flow-collector-  
hostname.2055: UDP, length 484  
22:01:00.341895 IP flow-sensor-hostname.16384 > flow-collector-  
hostname.2055: UDP, length 256  
22:01:00.341966 IP flow-sensor-hostname.16384 > flow-collector-  
hostname.2055: UDP, length 396  
22:01:00.342034 IP flow-sensor-hostname.16384 > flow-collector-  
hostname.2055: UDP, length 228  
22:01:00.342107 IP flow-sensor-hostname.16384 > flow-collector-  
hostname.2055: UDP, length 436  
22:01:00.342384 IP flow-sensor-hostname.16384 > flow-collector-  
hostname.2055: UDP, length 484  
22:01:00.342459 IP flow-sensor-hostname.16384 > flow-collector-  
hostname.2055: UDP, length 256  
22:01:00.342531 IP flow-sensor-hostname.16384 > flow-collector-  
hostname.2055: UDP, length 396  
9 packets captured  
9 packets received by filter  
0 packets dropped by kernel
```

Tcpdump Expressions

Available expressions for use in constructing [bpf filters](#):

dst host *host*

True if the IPv4/v6 destination field of the packet is *host*, which may be either an address or a name.

src host *host*

True if the IPv4/v6 source field of the packet is *host*.

host *host*

True if either the IPv4/v6 source or destination of the packet is *host*. Any of the above host expressions can be prepended with the keywords, *ip*, *arp*, *rarp*, or *ip6* as in: *ip host host* which is equivalent to: *ether proto \ip and host host*. If *host* is a name with multiple IP addresses, each address will be checked for a match.

ether dst *ehost*

True if the Ethernet destination address is *ehost*. *Ehost* may be either a name from */etc/ethers* or a number (see *ethers(3N)* for numeric format).

ether src *ehost*

True if the Ethernet source address is *ehost*.

ether host *ehost*

True if either the Ethernet source or destination address is *ehost*.

gateway *host*

True if the packet used *host* as a gateway. I.e., the Ethernet source or destination address was *host* but neither the IP source nor the IP destination was *host*. *Host* must be a name and must be found both by the machine's host-name-to-IP-address resolution mechanisms (*host name file*, *DNS*, *NIS*, etc.). (An equivalent expression is *ether host ehost* and not *host host* which can be used with either names or numbers for *host* / *ehost*.) This syntax does not work in IPv6-enabled configuration at this moment.

dst net *net*



True if the IPv4/v6 destination address of the packet has a network number of net. Net may be either a name from /etc/networks or a network number (see networks(4) for details).

src net *net*

True if the IPv4/v6 source address of the packet has a network number of net.

net *net*

True if either the IPv4/v6 source or destination address of the packet has a network number of net.

net net mask *netmask*

True if the IP address matches net with the specific netmask. May be qualified with src or dst. Note that this syntax is not valid for IPv6 net.

net net/len

True if the IPv4/v6 address matches net with a netmask len bits wide. May be qualified with src or dst.

dst port *port*

True if the packet is ip/tcp, ip/udp, ip6/tcp or ip6/udp and has a destination port value of port. The port can be a number or a name used in /etc/services (see tcp(4P) and udp(4P)). If a name is used, both the port number and protocol are checked. If a number or ambiguous name is used, only the port number is checked (e.g., dst port 513 will print both tcp/login traffic and udp/who traffic, and port domain will print both tcp/domain and udp/domain traffic).

src port *port*

True if the packet has a source port value of port. True if either the source or destination port of the packet is port. Any of the above port expressions can be prepended with the keywords, tcp or udp, as in: tcp src port port which matches only tcp packets whose source port is port.

less *length*

True if the packet has a length less than or equal to length. This is equivalent to: `len <= length`.

greater *length*

True if the packet has a length greater than or equal to length. This is equivalent to: `len >= length`.

ip proto *protocol*

True if the packet is an IP packet (see `ip(4P)`) of protocol type protocol. Protocol can be a number or one of the names `icmp`, `icmp6`, `igmp`, `igrp`, `pim`, `ah`, `esp`, `vrp`, `udp`, or `tcp`. Note that the identifiers `tcp`, `udp`, and `icmp` are also keywords and must be escaped via backslash (`\`), which is `\\` in the C-shell. Note that this primitive does not chase the protocol header chain.

ip6 proto *protocol*

True if the packet is an IPv6 packet of protocol type protocol. Note that this primitive does not chase the protocol header chain.

ip6 protochain *protocol*

True if the packet is IPv6 packet, and contains protocol header with type protocol in its protocol header chain. For example, `ip6 protochain 6` matches any IPv6 packet with TCP protocol header in the protocol header chain. The packet may contain, for example, authentication header, routing header, or hop-by-hop option header, between IPv6 header and TCP header. The BPF code emitted by this primitive is complex and cannot be optimized by BPF optimizer code in `tcpdump`, so this can be somewhat slow.

ip protochain *protocol*

Equivalent to `ip6 protochain protocol`, but True if the packet is an Ethernet broadcast packet. The `ether` keyword is optional.

ip broadcast

True if the packet is an IP broadcast packet. It checks for both the all-zeroes and all-ones broadcast conventions, and looks up the local subnet mask.

ether multicast



True if the packet is an Ethernet multicast packet. The `ether` keyword is optional. This is shorthand for ``ether[0] & 1 != 0'`.

ip multicast

True if the packet is an IP multicast packet.

ip6 multicast

True if the packet is an IPv6 multicast packet.

ether proto *protocol*

True if the packet is of `ether type protocol`. Protocol can be a number or one of the names `ip`, `ip6`, `arp`, `rarp`, `atalk`, `aarp`, `decnet`, `sca`, `lat`, `mopdl`, `moprc`, `iso`, `stp`, `ipx`, or `netbeui`. Note these identifiers are also keywords and must be escaped via backslash (`\`). [In the case of FDDI (e.g., ``fddi protocol arp'`) and Token Ring (e.g., ``tr protocol arp'`), for most of those protocols, the protocol identification comes from the 802.2 Logical Link Control (LLC) header, which is usually layered on top of the FDDI or Token Ring header. When filtering for most protocol identifiers on FDDI or Token Ring, tcpdump checks only the protocol ID field of an LLC header in so-called SNAP format with an Organizational Unit Identifier (OUI) of 0x000000, for encapsulated Ethernet; it doesn't check whether the packet is in SNAP format with an OUI of 0x000000. The exceptions are `iso`, for which it checks the DSAP

(Destination Service Access Point) and SSAP (Source Service

fields of the LLC header, stp and netbeui, packet with an OUI of 0x080007 and the Appletalk etype.

In the case of Ethernet, tcpdump checks the Ethernet type field for most of those protocols; the exceptions are iso, sap, and netbeui, for which it checks for an 802.3 frame and then checks the LLC header as it does for FDDI and Token Ring, atalk, where it checks both for the Appletalk etype in an Ethernet frame and for a SNAP-format packet as it does for FDDI and Token Ring, aarp, where it checks for the Appletalk ARP etype in either an Ethernet frame or an 802.2 SNAP frame with an OUI of 0x000000, and ipx, where it checks for the IPX etype in an Ethernet frame, the IPX DSAP in the LLC header, the 802.3 with no LLC header encapsulation of IPX, and the IPX etype in a SNAP frame.]

decnet src *host*

True if the DECNET source address is *host*, which may be an address of the form ``10.123'', or a DECNET host name. [DECNET host name support is only available on Ultrix systems that are configured to run DECNET.]

decnet dst *host*

True if the DECNET destination address is *host*.

decnet host *host*

True if either the DECNET source or destination address is *host*.

ip, ip6, arp, rarp, atalk, aarp, decnet, iso, stp, ipx, netbeui

Abbreviations for: ether proto p where p is one of the above protocols.

lat, moprc, mopdl

Abbreviations for: ether proto p where p is one of the above protocols. Note that tcpdump does not currently know how to parse these protocols.

vlan [*vlan_id*]

True if the packet is an IEEE 802.1Q VLAN packet. If [*vlan_id*] is specified, only encountered in expression changes the



decoding offsets for the remainder of expression on the assumption that the packet is a VLAN packet.

tcp, udp, icmp

Abbreviations for: ip proto p or ip6 proto p where p is one of the above protocols.

iso proto protocol

True if the packet is an OSI packet of protocol type protocol. Protocol can be a number or one of the names clnp, esis, or isis.

clnp, esis, isis

Abbreviations for: iso proto p where p is one of the above protocols. Note that tcpdump does an incomplete job of parsing these protocols.

expr relop expr

True if the relation holds, where relop is one of >, <, >=, <=, =, !=, and expr is an arithmetic expression composed of integer constants (expressed in standard C syntax), the normal binary operators [+ , - , * , / , & , |], a length operator, and special packet data accessors. To access data inside the packet, use the following syntax: proto [expr : size] Proto is one of ether, **fddi**, **tr**, **ip**, **arp**, **rarp**, **tcp**, **udp**, **icmp** or **ip6**, and indicates the protocol layer for the index operation. Note that tcp, udp and other upper-layer protocol types only apply to IPv4, not IPv6 (this will be fixed in the future). The byte offset, relative to the indicated protocol layer, is given by expr. Size is optional and indicates the number of bytes in the field of interest; it can be either one, two, or four, and defaults to one. The length operator, indicated by the keyword len, gives the length of the packet. For example, ``ether[0] & 1 != 0`` catches all multicast traffic. The expression ``ip[0] & 0xf != 5`` catches all IP packets with options. The expression ``ip[6:2] & 0x1fff = 0`` catches only unfragmented datagrams and frag zero of fragmented datagrams. This always means the first byte of the TCP header, and never means the first byte of an intervening fragment. Some offsets and field values may be expressed as names rather than as numeric values. The following protocol header field offsets are available: **icmptype** (ICMP type field), **icmpcode** (ICMP

code field), and **tcpflags** (TCP flags field). The following ICMP type field values are available: **icmp-echoreply**, **icmp-unreach**, **icmp-sourcequench**, **icmp-redirect**, **icmp-echo**, **icmp-routeradvert**, **icmp-routersolicit**, **icmp-timxceed**, **icmp-paramprob**, **icmp-tstamp**, **icmp-tstampreply**, **icmp-ireq**, **icmp-ireqreply**, **icmp-maskreq**, **icmp-maskreply**. The following TCP flags field values are available: **tcp-fin**, **tcp-syn**, **tcp-rst**, **tcp-push**, **tcp-push**, **tcp-ack**, **tcp-urg**. Primitives may be combined using: A parenthesized group of primitives and operators (parentheses are special to the Shell and must be escaped). Negation (**!** or **not**). Concatenation (**&&** or **and**). Alternation (**||** or **or**). Negation has highest precedence. Alternation and concatenation have equal precedence and associate left to right. Note that explicit and tokens, not juxtaposition, are now required for concatenation. If an identifier is given without a keyword, the most recent keyword is assumed. For example, **not host vs and ace** is short for **not host vs and host ace** which should not be confused with **not (host vs or ace)** Expression arguments can be passed to **tcpdump** as either a single argument or as multiple arguments, whichever is more convenient. Generally, if the expression contains Shell metacharacters, it is before being parsed.



Advanced BPF Filters

The tables below provide a reference for advanced users using BPF (Berkeley Packet Filter) filters on the [Packet Capture](#) screen. The Packet Capture screen provides an interface to tcpdump, which accepts the filters in the tables below.

Notable Parts of an IP datagram	
ip[0] & 0xf0	high nibble: IP version. almost always 4
ip[0] & 0x0f	low nibble: header length in 4 octet words. should be 5
ip[1]	type of service/ QoS / DiffServ
ip[2:2]	total length of datagram in octets
ip[4:2]	IP ID number
ip[6] & 0x80	reserved bit (possibly used for ECN)
ip[6] & 0x40	DF bit

Notable Parts of an IP datagram	
ip[6] & 0x20	MF bit
ip[6:2] & 0x1fff	fragment offset (number of 8octet blocks)
ip[8]	ttl
ip[9]	protocol
ip[10:2]	header checksum
ip[12:4]	source IP
ip[16:4]	destination IP

Notable Parts of an ICMP message	
icmp[0]	type
icmp[1]	code
icmp[2:2]	checksum
icmp[4...]	payload

Notable Parts of a UDP header	
udp[0:2]	source port
udp[2:2]	destination port



Notable Parts of a UDP header	
udp[4:2]	datagram length
udp[6:2]	UDP checksum

Notable Parts of a TCP header	
tcp[0:2]	source port
tcp[2:2]	destination port
tcp[4:4]	sequence number
tcp[8:4]	acknowledgement number
tcp[12]	header length
tcp[13]	tcp flags
tcp[14:2]	window size
tcp[16:2]	checksum
tcp[18:2]	urgent pointer
tcp[20..60]	options or data

Flags	Numeric	Description	bpf Filter (only specified flag set)	bpf Filter (specified flag set plus any other flag)
---- - --F	0000 0001 =	FIN	tcp[13]=1	tcp [13]&0x01=0x01

Flags	Numeric	Description	bpf Filter (only specified flag set)	bpf Filter (specified flag set plus any other flag)
	0x01			
---- - -S-	0000 0010 = 0x02	SYN	tcp[13]=2	tcp [13]&0x02=0x02
---- - R--	0000 0100 = 0x04	RST	tcp[13]=4	tcp [13]&0x04=0x04
---- P---	0000 1000 = 0x08	PSH	tcp[13]=8	tcp [13]&0x08=0x08
---A - ---	0001 0000 = 0x10	ACK	tcp [13]=0x10	tcp [13]&0x10=0x10
--U- - ---	0010 0000 = 0x20	URG	tcp [13]=0x20	tcp [13]&0x20=0x20
---A - -S-	0001 0010 = 0x12	SYN-ACK	tcp [13]=0x12	tcp [13]&0x12=0x12
--UA P---	0011 1000 = 0x38	PSH-urg- ACK	tcp [13]=0x38	tcp [13]&0x38=0x38
---A - R--	0001 0100 = 0x14	RST-ACK	tcp [13]=0x14	tcp [13]&0x14=0x14



Flags	Numeric	Description	bpf Filter (only specified flag set)	bpf Filter (specified flag set plus any other flag)
---- - -SF	0000 0011 = 0x03	SYN-FIN	tcp [13]=0x03	tcp [13]&0x03=0x03
--U- P--F	0010 1001 = 0x29	URG-PSH- FIN	tcp [13]=0x29	tcp [13]&0x29=0x29
XY-- - ---	1100 0000 = 0xC0	both reserved bits set	tcp [13]=0xC0	tcp [13]&0xC0=0xC0
XYUA PRSF	1111 1111 = 0xFF	all flags set	tcp [13]=0xFF	N/A

ICMP Types and Codes	
0	ECHOREPLY
3	UNREACHABLE
03:00	NET
03:01	HOST
03:02	PROTOCOL
03:03	PORT
03:04	NEEDFRAG
03:05	SRC_ROUTE_

ICMP Types and Codes	
	FAILED
03:06	NET_UNKNOWN
03:07	HOST_UNKNOWN
03:08	SRC_HOST_ISOLATED
03:09	NET_PROHIB
03:10	HOST_PROHIB
03:11	BAD_TOS_FOR_NET
03:12	BAD_TOS_FOR_HOST
03:13	FILTER_PROHIB
03:14	HOST_PRECEDENCE_VIOLATION
03:15	PRECEDENCE_CUTOFF
4	SOURCEQUENCH
5	REDIRECT
05:00	NET
05:01	HOST



ICMP Types and Codes	
05:02	TOSNET
05:03	TOSHOST
8	ECHO
9	ROUTERADVERT
10	ROUTERSOLICIT
11	TIME_EXCEEDED
11:00	IN_TRANSIT
11:01	DURING_FRAG_ REASSEMBLY
12	PARAMETER_ PROBLEM
12:01	MISSING_OPT_ FOR_REQUEST
13	TSTAMP_REQ
14	TSTAMP_REPLY
15	INFO_REQ
16	INFO_REPLY
17	NETMASK_REQ
18	NETMASK_REPLY

Update Manager

Use the Update Manager to apply software updates and patches. You can download the files from your Cisco Smart Account on Cisco Software Central at <https://software.cisco.com>.

Open Update Manager

1. From the main menu, select **Configure > Global > Central Management**.
2. Select the **Update Manager** tab.

Updates

Plan to update the entire system when your appliances will be experiencing relatively low volumes of traffic.

To update your appliances, follow the instructions in the [Cisco Secure Network Analytics System Update Guide](#). The guide includes details that are critical for a successful update, including:

- Planning for the update
- Installing patches
- Creating backup files and diagnostic packs
- Updating your appliances in order



For a successful update, follow the update order and instructions in the [Cisco Secure Network Analytics System Update Guide](#).

Updating Data Nodes

Click **Update all Data Nodes** to update your Data Nodes at the same time.

- **Status:** The Update Manager shows the overall update status, and the new installed version is shown after the update is completed. If a Flow Collector has a Data Store Transition tag, it is in the process of being transitioned from a non-Data Store Flow Collector to a Data Store Flow Collector. The appliances reboot during the update process, so the status information may be delayed when the appliance goes



offline. To monitor the progress of the database services update on each Data Node, go to the **Data Store** > [Database Update Status](#) tab. Also, refresh each page to see the most recent status.

- **Instructions:** For a successful update, follow the update order and instructions in the [Cisco Secure Network Analytics System Update Guide](#) or [patch readme file](#).
- **Best Practices:** Although you can update each Data Node individually, we recommend using the **Update all Data Nodes** button to update your Data Nodes at the same time.

Patches

Keep Secure Network Analytics up-to-date by installing the latest patches for your software version. You can download the files from your Cisco Smart Account on Cisco Software Central at <https://software.cisco.com>.

For each patch, make sure you follow the instructions in the patch readme file. A link to a patch readme file is provided with the patch file on Cisco Software Central. You can also download patch readme files directly from the [Release Notes page](#) on Cisco.com.

Types of Patches

There are generally two types of patches:

- appliance-specific rollup patches for a version
- common update patches for all versions and/or appliances

Common update patches for CIMC firmware version updates and Cisco Bundles are released periodically. For details, refer to the [latest readmes](#). You can download the SWU files from Software Central at <https://software.cisco.com>.

CIMC Firmware Version Common Update Patch

When made available, make sure to apply the CIMC firmware version update patches for your UCS C-Series hardware. These are common update patch SWU files that apply to all Cisco Secure Network Analytics (formerly Stealthwatch) appliances.

Cisco Bundles Common Update Patch

Cisco releases bundles of pre-validated digital certificates of a select number of root certificate authorities (CAs), when necessary. We provide Cisco Bundles common update patch SWU files for you to install on of your Cisco Secure Network Analytics (formerly Stealthwatch) appliances.

Each patch includes a Core certificate bundle and an External certificate bundle, which are used for connecting to Cisco services and to non-Cisco services.



- Make sure you have the latest Cisco Bundles common update patch installed on all of your appliances.
- If you do a RefreshImage for an appliance, make sure to install the most recent Cisco Bundles common update patch. Doing a RefreshImage reverts to the CAs shipped with an appliance.

Detections Packs

As new detections become available, we provide an updated detections pack you'll install as a SWU file for your Manager. The first pack is included with Secure Network Analytics. Go to the **Configure > Detection > Pack Management** to access the Detections Pack Management page where you can select the newly released pack or to return to a previously released pack. Refer to the [Detections Pack Guide](#) for more information.

Related Topics

[Data Store Tab](#)

[Data Store - Database Update Status](#)

[Central Management Inventory](#)



Backup/Restore Configuration



Please go to [Central Management Help](#) for information and instructions.

Diagnostics Pack



Please go to [Central Management Help](#) for information and instructions.



Audit Log



Please go to [Central Management Help](#) for information and instructions.

Operations



If your appliance is managed by a Central Manager, go to [Central Management Help](#) for information and instructions.

Use the Operations menu for access to the following administrative functions:

Item	Description
Restart Appliance	Restarts (i.e., reboots) the appliance. There is a brief period during the restart when data is not being collected.
Shutdown Appliance	Forces the appliance to power off. Once this command is issued, the appliance must be turned back on at its physical location. Data cannot be collected while the appliance is shut down. For a physical appliance, you must have physical access to the appliance in order to start it again. For a virtual appliance, you must have administrator access to it.

Access these functions

To issue a command, place your cursor over the Operations item in the main menu, and select one of the following from the drop-down menu:

- Restart Appliance
- Shutdown Appliance

About Central Management

When you log in to an appliance for the first time, you will use the appliance console (SystemConfig) to configure the appliance and assign it to a Central Manager. The Central Management application is your Cisco Secure Network Analytics Manager (formerly Stealthwatch Management Console). For details, refer to the [System Configuration Guide](#).

Multiple Central Managers: You can configure more than one Central Manager in your system. However, each appliance can be managed by only one Central Manager.

Access: You need administrator privileges to access Central Management.

Open Central Management

1. Log in to your primary Manager.
(In your browser address field, type https:// and the appliance IP address. Press Enter.)
2. From the main menu, select **Configure > Global > Central Management**.

Inventory: Central Management opens to the [Inventory](#) tab.

Manage your Appliances: When your appliances are managed by Central Management, you can review their status and manage the following:

- [edit appliance configuration](#)
- [update software](#)
- [create backup configuration files](#)
- [view audit logs](#)
- [reboot](#)
- [shut down](#)
- [review license status](#)
- [update your flow collector database password](#)

Open Individual Appliances

Also, you can open an appliance from Central Management or by logging in to the appliance directly.

Use the appliance interface for the following:

- view system statistics
- additional appliance configuration options (not shown in Central Management)
- additional support options (not shown in Central Management)

Open an Appliance through Central Management

1. From the **Central Management Inventory**, click the **⋯ (Ellipsis)** icon for the appliance.
2. Select **View Appliance Statistics**.
3. Log in to the Appliance Administration interface.

Open Appliance Administration through Direct Login

To access the Appliance Admin through direct login, or if the appliance is disconnected from Central Management, use the following instructions.

1. In your browser address bar, type the appliance IP address as follows:

https://<IPAddress>

- **Manager:** add **/manager/index.html** after the IP address.
- **For example:** https://1.1.1.1/manager/index.html

2. Press Enter.
3. Log in to the appliance.

Related Topics

[Central Management Inventory](#)

[Appliance Configuration](#)

[Appliance Support](#)

[Data Store Tab](#)

[Update Manager](#)

[App Manager](#)


[Database](#)



About the Desktop Client

As the control center for Cisco Secure Network Analytics, the Desktop Client manages, coordinates, configures, and organizes all of the different components of the system. The Desktop Client allows you to access the Cisco Secure Network Analytics Manager's (formerly Stealthwatch Management Console's) user-friendly graphical user interface (GUI) from any local computer with access to a Web browser. Through this GUI, you can easily access real-time security and network information about critical segments throughout your enterprise.

i Use the Web App to monitor and configure your Secure Network Analytics installation if you deploy a Data Store. The Desktop Client is incompatible with a Data Store.



Featuring Java-based platform independence, the Manager enables:

- Centralized management, configuration, and reporting for your Cisco Secure Network Analytics Flow Collectors
- Graphical charts for visualizing traffic
- Drill-down analysis for troubleshooting
- Consolidated and customizable reports
- Trend analysis
- Performance monitoring
- Immediate notification of security breaches



You must have a Manager and a Flow Collector to deploy and run Secure Network Analytics.



Smart Software Licensing

Use Smart Software Licensing to see an overview of your license status and usage.

- **Instructions:** For instructions and details about Cisco Smart Software Licensing, please refer to the [Cisco Secure Network Analytics Smart Software Licensing Guide](#).
- **Manage Licenses:** To add or update licenses on your account, log in to your [Cisco Smart Software Manager](#) or contact your administrator.

Open Smart Software Licensing

1. From the main menu, select **Configure > Global > Central Management**.
2. Select the **Smart Licensing** tab.

Product Instance Name

The Product Instance Name is the identifier we use for your Secure Network Analytics product instance, which includes your Manager and managed appliances. This identifier is shown on your Secure Network Analytics Smart Software Licensing page (in the Smart Software Licensing Status section) and in your [Cisco Smart Software Manager](#).

Evaluation and Registration

When you use Secure Network Analytics in Evaluation mode, your appliances are licensed to use selected features for 90 days. To use Secure Network Analytics with maximum default functionality, and to add licenses and features to your account, register your product instance for Smart Software Licensing.

- **Expiration:** Make sure you register before the 90-day evaluation period expires. When the evaluation period expires, flow collection will stop. To start flow collection again, register your product instance.
- **Registration:** To register, follow the instructions in the [Cisco Secure Network Analytics Software Licensing Guide](#). As part of the registration process, you will log in to your [Cisco Smart Software](#)

Manager at <https://software.cisco.com> to generate the registration token and add licenses.

Transport Settings

Use Transport Settings to configure Secure Network Analytics to communicate with the Cisco Smart Software Manager. For configuration instructions and troubleshooting, please refer to the [Cisco Secure Network Analytics Smart Software Licensing Guide](#).



Update Manager

Use the Update Manager to apply software updates and patches. You can download the files from your Cisco Smart Account on Cisco Software Central at <https://software.cisco.com>.

Open Update Manager

1. From the main menu, select **Configure > Global > Central Management**.
2. Select the **Update Manager** tab.

Updates

Plan to update the entire system when your appliances will be experiencing relatively low volumes of traffic.

To update your appliances, follow the instructions in the [Cisco Secure Network Analytics System Update Guide](#). The guide includes details that are critical for a successful update, including:

- Planning for the update
- Installing patches
- Creating backup files and diagnostic packs
- Updating your appliances in order



For a successful update, follow the update order and instructions in the [Cisco Secure Network Analytics System Update Guide](#).

Updating Data Nodes

Click **Update all Data Nodes** to update your Data Nodes at the same time.

- **Status:** The Update Manager shows the overall update status, and the new installed version is shown after the update is completed. If a Flow Collector has a Data Store Transition tag, it is in the process of being transitioned from a non-Data Store Flow Collector to a Data Store Flow Collector. The appliances reboot during the update process, so the status information may be delayed when the appliance goes

offline. To monitor the progress of the database services update on each Data Node, go to the **Data Store** > [Database Update Status](#) tab. Also, refresh each page to see the most recent status.

- **Instructions:** For a successful update, follow the update order and instructions in the [Cisco Secure Network Analytics System Update Guide](#) or [patch readme file](#).
- **Best Practices:** Although you can update each Data Node individually, we recommend using the **Update all Data Nodes** button to update your Data Nodes at the same time.

Patches

Keep Secure Network Analytics up-to-date by installing the latest patches for your software version. You can download the files from your Cisco Smart Account on Cisco Software Central at <https://software.cisco.com>.

For each patch, make sure you follow the instructions in the patch readme file. A link to a patch readme file is provided with the patch file on Cisco Software Central. You can also download patch readme files directly from the [Release Notes page](#) on Cisco.com.

Types of Patches

There are generally two types of patches:

- appliance-specific rollup patches for a version
- common update patches for all versions and/or appliances

Common update patches for CIMC firmware version updates and Cisco Bundles are released periodically. For details, refer to the [latest readmes](#). You can download the SWU files from Software Central at <https://software.cisco.com>.

CIMC Firmware Version Common Update Patch

When made available, make sure to apply the CIMC firmware version update patches for your UCS C-Series hardware. These are common update patch SWU files that apply to all Cisco Secure Network Analytics (formerly Stealthwatch) appliances.



Cisco Bundles Common Update Patch

Cisco releases bundles of pre-validated digital certificates of a select number of root certificate authorities (CAs), when necessary. We provide Cisco Bundles common update patch SWU files for you to install on of your Cisco Secure Network Analytics (formerly Stealthwatch) appliances.

Each patch includes a Core certificate bundle and an External certificate bundle, which are used for connecting to Cisco services and to non-Cisco services.



- Make sure you have the latest Cisco Bundles common update patch installed on all of your appliances.
- If you do a RefreshImage for an appliance, make sure to install the most recent Cisco Bundles common update patch. Doing a RefreshImage reverts to the CAs shipped with an appliance.

Detections Packs

As new detections become available, we provide an updated detections pack you'll install as a SWU file for your Manager. The first pack is included with Secure Network Analytics. Go to the **Configure > Detection > Pack Management** to access the Detections Pack Management page where you can select the newly released pack or to return to a previously released pack. Refer to the [Detections Pack Guide](#) for more information.

Related Topics

[Data Store Tab](#)

[Data Store - Database Update Status](#)

[Central Management Inventory](#)

Navigating the Manager

The topics in this section explain how to access general information (such as the Help and resources) and how to accomplish interface-related tasks (such as editing your profile and logging out).

System Navigation: Table of Contents



Compatibility

Supported Browsers

Secure Network Analytics supports the following browsers (recommended resolution of 1440 px wide):

- Chrome (latest rapid release)
- Firefox (latest rapid release)
- Edge (latest rapid release)

Supported Operating Systems

- Windows® 10
- Windows 11
- Windows Server 2012 R2
- Windows Server 2016 Remote Desktop Services (formerly Terminal Services) in 2008 R2
- Windows Server 2016
- Mac OS X (latest version)

Use Context Menus

If you want to open a page in a new tab when pivoting from a context menu (...), do one of the following:

- On a Windows-based PC, perform a Ctrl+Click .
- On a Mac, perform a Cmd+Click.

This functionality does not work if you are connected to a VPN.



Install the Desktop Client Overview

Refer to the following information before installing the Desktop Client.



If your Manager is configured with a Cisco Secure Network Analytics Data Store, then you do not have the option to download the Desktop Client. The Desktop Client is not supported by this deployment type.

- You can locally install different versions of Desktop Client.
- If you want to access multiple versions of Desktop Client, you need a different executable file for each Manager.
- You can have different versions of Desktop Client open simultaneously.
- When you update to a later version of Cisco Secure Network Analytics, you need to install the new version of Desktop Client.
- If you are using both a primary and a secondary Manager, you need to log off one Manager before you can log in to the other Manager.

What to do Next

[Install the Desktop Client Using Mac OS](#)

[Install the Desktop Client Using Windows](#)

Install the Desktop Client Using Mac OS



- If you notice that the Desktop Client appears to "hang" frequently, try increasing this value.
- If you receive an error message involving Java, try choosing a lower memory allocation.

1. If you have previously downloaded the same version of the Desktop Client that you are now downloading, delete it before advancing to Step 2.
2. On any page in the Secure Network Analytics Manager, do one of the following to open the Desktop Client, depending on whether or not you have a data store-enabled system:
 - For data store-enabled system: Click the Download arrow on any page.
 - For non-data store-enabled system: Click the [Download Desktop Client](#) icon at the bottom right corner of any page.
An .exe file is downloaded.
3. Click the .dmg file to begin the installation process.

The  (**Desktop Client**) icon and folder are displayed on your monitor, as shown below.




4. Drag the **Desktop Client** icon into the Application folder.
The icon is added to the Launchpad.
5. Open the Desktop Client by clicking this icon.
 - In addition to a user name and password, you must also enter a Manager server name.



- (optional) If you want Secure Network Analytics to verify the Hostname, check the Enable Strict Hostname Verification check box.

Secure Network Analytics verifies the identity of the server to ensure that the client is talking to the correct server. Secure Network Analytics also verifies that the DNS name in the subjectAltName field on the certificate that the server sent matches the host portion of the URL that was used to make the request.

6. The first time you open the Desktop Client, the following dialog opens (the version number displayed in this dialog will change depending on which version you are downloading):



7. To open the Desktop Client, on your computer go to **Finder > Applications**, right-click the **Desktop Client** icon, and click **Open**.

Change the memory size

You can change how much Random Access Memory (RAM) to allocate on your client computer to run the Desktop Client. Consider a larger memory allocation if you work with many open documents or large data sets (such as flow queries with over 100k records).

1. In Finder, go to your home directory.
2. Open the Secure Network Analytics folder (formerly Stealthwatch).
3. In the Secure Network Analytics directory, open the folder that contains the desired Secure Network Analytics version.

4. Open the application.vmoptions file using an appropriate editing application to begin editing. (This file is created after you open the Desktop Client for the first time.)

- For the minimum memory size, we recommend that you allocate no less than 512 MB. This number is listed in the third line of the file. For editors that display the content in one continuous line, refer to the number highlighted in the image below to see which number represents the minimum memory size.

```
Enter one VM parameter per line
# Use -Xms to specify the initial Java heap size and Use -Xmx to specify the maximum heap
-Xms512m
-Xmx2048m
```

- You can allocate up to half the size of your computer's RAM for the maximum memory size. This number is listed in the fourth line of the file. For editors that display the content in one continuous line, refer to the number highlighted in the image below to see which number represents the maximum memory size.

```
Enter one VM parameter per line
# Use -Xms to specify the initial Java heap size and Use -Xmx to specify the maximum heap
-Xms512m
-Xmx2048m
```

- Use whole numbers. For example, enter Xmx512m, not Xmx0.5m.



- If you notice that the Desktop Client appears to "hang" frequently, try increasing this value.
- If you receive an error message involving Java, try choosing a lower memory allocation.



Install the Desktop Client Using Windows



- You must have sufficient rights to install the Desktop Client.
- The Desktop Client requires a 64-bit operating system. It cannot run on a 32-bit operating system or Linux.
- Please clear the cache before launching the new Desktop Client.

1. On any page in the Secure Network Analytics Manager, do one of the following to open the Desktop Client, depending on whether or not you have a data store-enabled system:
 - For data store-enabled system: Click the Download arrow on any page.
 - For non-data store-enabled system: Click the  icon at the bottom right corner of any page.
An .exe file is downloaded.
2. Click the .exe file to begin the installation process.
A setup wizard opens.
3. Proceed through the steps in the wizard to install the Desktop Client.
The  (Desktop Client) icon is added to your desktop.
4. Open the Desktop Client by clicking this icon.
 - In addition to a user name and password, you must also enter a Manager server name.
 - (optional) If you want Secure Network Analytics to verify the Hostname, check the "Enable Strict Hostname Verification" check box.

Secure Network Analytics verifies the identity of the server to ensure that the client is talking to the correct server. Secure Network Analytics also verifies that the DNS name in the subjectAltName field

on the certificate that the server sent matches the host portion of the URL that was used to make the request.

Change the memory size

You can change how much Random Access Memory (RAM) to allocate on your client computer to run the Desktop Client. Consider a larger memory allocation if you work with many open documents or large data sets (such as flow queries with over 100k records).

1. In Windows Explorer, go to your home directory.
2. Open the Secure Network Analytics folder using this path: AppData > Roaming > Secure Network Analytics.
3. In the Secure Network Analytics directory, open the folder that contains the desired Secure Network Analytics version.
4. Open the application.voptions file using an appropriate editing application to begin editing. (This file is created after you open the Desktop Client for the first time.)

- For the minimum memory size, we recommend that you allocate no less than 512 MB. This number is listed in the third line of the file. For editors that display the content in one continuous line, refer to the number highlighted in the image below to see which number represents the minimum memory size.

```
# Enter one VM parameter per line
# Use -Xms to specify the initial Java heap size and Use -Xmx to specify the maximum heap
-Xms512m
-Xmx2048m
```

- You can allocate up to half the size of your computer's RAM for the maximum memory size. This number is listed in the fourth line of the file. For editors that display the content in one continuous line, refer to the number highlighted in the image below to see which number represents the maximum memory size.

```
# Enter one VM parameter per line
# Use -Xms to specify the initial Java heap size and Use -Xmx to specify the maximum heap
-Xms512m
-Xmx2048m
```

- Use whole numbers. For example, enter Xmx512m, not Xmx0.5m.




- If you notice that the Desktop Client appears to "hang" frequently, try increasing this value.




- If you receive an error message involving Java, try choosing a lower memory allocation.

Use Filters


Several pages in Secure Network Analytics provide filters that you use to focus on the information you need. If you have used the Desktop Client, you are familiar with the filters opening as separate pages in a dialog. In the Web App, however, the categories of filters are displayed on the left side of the page you are viewing. You use them in a manner similar to that of commercial Web sites.



- You can choose a filter from one or more categories, which are then listed in the Current Filters window. Selecting a filter from multiple categories *decreases* the number of results returned by the query.



- Depending on the filter category, you may be able to choose more than one filter from that category. The query retrieves results that match any of the chosen filters. Some of the results may match more than one of the chosen filters within that category.



- If you want to apply several filters from a category, click **Select Multiple**. When the Select Multiple Attributes dialog opens, choose the filters you want to use and click **Apply**. Using multiple filters within the same category *increases* the number of results returned by the query.

Select Multiple Attributes

Alarms

☒ Target Index (2)	☐ Data Hoarding (0)
☒ Concern Index (1)	☐ Exfiltration (0)
☐ Anomaly (0)	☐ DDoS Source (0)
☐ Exploitation (0)	☐ DDoS Target (0)
☐ Command & Control (0)	☐ Policy Violation (0)
	☐ Recon (0)

Cancel Apply

- To quickly remove all filters, click **Clear All** in the Current Filters window.
- A filter category displays a maximum of 15 filters by default. If the category includes additional filters, a link labeled "More" appears. When you click the link, the list expands to show the first 50 filters. If more than 50 filters are provided, a scroll bar is available. A link labeled "Less" is displayed so you can collapse the list.
- Some filters include numbers in parentheses to show how many items in your system match that filter. If you choose multiple filters within a category, the numbers by the filters may not match the number of hosts or users in your query results. Remember that if you choose multiple filters within the same category, some of the results may match more than one of the chosen filters within that category.

*Example: You want to find out which hosts are generating Exploitation and DDOS Source alarms. On the Host List page, in the Filter Results By: Alarms section, you click **Select Multiple**, and then you choose and apply those two alarms from the Select Multiple Attributes dialog. The Filter Results By: Alarms section now shows that there are 6 Exploitation and 1 DDOS Source alarm, for a total of 7 alarms. However, your query results, listed in the table in the Host section on the right side of the page, show only 6 hosts with alarms. The difference in total numbers (7 vs. 6) is because one host is generating one each of these alarm category types.*



Example:

Inside Hosts (6)

Current Filters

Alarms

Exploitation ×

DDoS Source ×

[Clear All](#)

Filter Results By:

ALARMS ▼

Exploitation (6)

Concern Index (6)

Command & Control (5)

DDoS Source (1)

Anomaly (0)

Data Hoarding (0)

Exfiltration (0)

DDoS Target (0)

Policy Violation (0)

Recon (0)

Target Index (0)

[Select Multiple](#)

Host

Sort by overall severity ?

Host Address

0.30 ...

0.40 ...

0.20 ...

0.70 ...

0.30 ...

.255.253 ...

First

Previous

1

Next

Choose a Theme

You can choose a theme that will define a look of the Manager either before or after you have logged in to Cisco Secure Network Analytics.

The theme switcher provides 3 options:

- Light
- Dark
- Auto (default)

The **Auto** theme automatically changes the appearance of the Manager according to your operating system preferences.

 Selecting the **Light** or **Dark** theme options will update the theme regardless of your operating system preferences. Your selected theme will be saved in the browser and automatically applied in future sessions. Please note that clearing your browser cache will reset the theme to the default setting.

Before you log in

Click the  (**Theme**) icon on the Login page and choose a theme from the drop-down menu.

After you log in

From the toolbar, choose the  (**User**) icon > **Theme** > **[theme name]**.



Select a Language

You can choose a different language in which to read the online text in the Manager either before or after you have logged in to Cisco Secure Network Analytics.

Before you log in

Click the  (**Global View**) icon on the Login page and choose a language from the drop-down menu.

After you log in

From the toolbar, choose the  (**User**) icon > **Language** > [language name].

Log Out

To log out of Cisco Secure Network Analytics, from the toolbar, choose the

 (**User**) icon > **Logout**.



Change Your Profile Name

To change your profile name, from the toolbar, choose the  (**User**) icon > **Profile**. Your profile name is the name displayed at the top of the context menu after you click the **User** icon.

If you delete your email address, it is no longer attached to your user name.



Your profile name is displayed at the top of the context menu when you click the User icon.

Change Your Password

The first time you log into Secure Network Analytics or when the Admin resets your password, you will be prompted to change your password.

 Remote users will not be prompted to change their passwords.

To change your password, choose the  (**User**) icon > **Profile** > **Change Password** from the drop-down menu.

For more information about changing a password, see [Users Tab: Manage Passwords](#).



View System Notifications

Notifications are indicated by a numeral displayed beside the  (**Alert**) icon , which is located on any page. The numeral indicates how many messages you have.

Notifications alert you to issues such as critical alarms or a soon-to-expire password.

For information about system notifications, see [System Notifications: Overview](#).

Search for a Host, IP Range, or User

To find a specific host, IP range, or user, click the  (**Search**) icon in the toolbar on any page. Type your entry in the search box and press **Enter** to run the search; clicking the **Search** icon does not run the search.

If you perform a search for a specific topic or term and try to print the results using the  (**Print**), note that if you are viewing the Secure Network Analytics Manager in either Chrome or Firefox, the page that is printed is the help topic associated with the page that was active when you performed the search. If you are in Chrome or Firefox, refer to the following guidelines:

- In Chrome, either right-click on the page containing the results and choose Print from the pop-up menu, or click the vertical ellipsis and click Print from the pop-up menu.
- In Firefox, click the  (**List Menu**) icon and click **Print** from the pop-up menu.

Secure Network Analytics Help

Get answers now

[Learn More!](#)

Resources



[Products](#)

Get to know our products





[Documentation](#)

Official product documentation and guides



[Support](#)

Have an issue? Let us know



[Use Cases](#)

Test use cases for your needs



[Information Hub](#)

A wealth of resources to help you succeed



[Training](#)

Get the most out of Secure Network Analytics

Get help when you need it

Explore Secure Network Analytics Help to find the answers you need. We have downloadable guides, user interface help, and links to support and training resources to get you back on track.

Tips

Explore the Products area to get started.

Use the search bar to perform a full text search on any topic.

Having trouble? Check out our Support area.

Did you know?

Secure Network Analytics is the only solution that can detect malware in encrypted traffic without decryption.

XDR Analytics helps you extend your visibility to the public cloud without the need for software agents.

Secure Network Analytics integrates with many other Cisco solutions to further protect your network.



Get answers now

Lancop Company Confidential

Useful Information

You can use Secure Network Analytics more effectively by using the following information:

- Allocation of client memory
- Conventions for entering IP addresses
- Tips for using the Desktop Client and Manager Help
- Browsers that Secure Network Analytics supports
- Known issues about the Manager or the Manager Help
- Glossary of terms used in Secure Network Analytics

Search Tips

- Enter keywords for your search. For example: "host group"
- Do not search with natural language. For example: "How do I configure host groups?"

Related Topics

[Desktop Client Help](#)

Access Help

To access the Help, from the toolbar, choose the  **(Help)** icon > **Help**.

You can search words or phrases.

Use buttons on the Help toolbar to do the following:

- Change the language.
- Print a Help topic.
- Go to the previous help page or next help page.

Search Tips

- Enter keywords for your search. For example: "host group"
- Do not search with natural language. For example: "How do I configure host groups?"

Related Topics

[Desktop Client Help](#)



Desktop Client Help

To access the Desktop Client Help, log in to your client and click **Help** from the main menu.

Glossary

Use this glossary to find definitions of terms used in Secure Network Analytics. Click the alphabetical category below to access a topic of terms.

A - F	G - L	M - Q	R - W	X - Z

Secure Network Analytics Help

Get answers now

[Learn More!](#)

Resources



[Products](#)

Get to know our products



[Documentation](#)

Official product documentation and guides



[Support](#)

Have an issue? Let us know



[Use Cases](#)

Test use cases for your needs



[Information Hub](#)

A wealth of resources to help you succeed



[Training](#)

Get the most out of Secure Network Analytics

Get help when you need it

Explore Secure Network Analytics Help to find the answers you need. We have downloadable guides, user interface help, and links to support and training resources to get you back on track.

Tips

Explore the Products area to get started.

Use the search bar to perform a full text search on any topic.

Having trouble? Check out our Support area.

Did you know?

Secure Network Analytics is the only solution that can detect malware in encrypted traffic without decryption.

XDR Analytics helps you extend your visibility to the public cloud without the need for software agents.

Secure Network Analytics integrates with many other Cisco solutions to further protect your network.



Glossary

A

Active Host

Host that is known to the Stealthwatch Flow Collector and that has transmitted a data packet in the last 5 minutes.

Alert

Indication of unusual network activity that has not necessarily resulted in an alarm. (This term is not used in the SMC Web App interface.)

Archive Hour

User-defined time when each Stealthwatch Flow Collector in the domain resets itself. Acceptable values are 0 to 23, where 0 is midnight local time. At the defined time, the Flow Collector resets all index counts to zero, saves the log and Web files it has gathered during the preceding 24 hours, and then begins a new day (24 hours) of data collection. Defined on the Domain page of the Domain Properties dialog.

ASN

Autonomous System Number; network of connected IP routing prefixes under the control of one or more network operators that presents a common, clearly defined routing policy to the Internet.

B

Batch target

A special target that lets you build and/or publish multiple other targets in a single group (or "batch"). You can schedule batches to run at any time.

Beaconing

Communication between an internal and outside host (with traffic in only one direction) that exceeds the number of seconds required for a flow to qualify as long duration.

Block snippet

A snippet that is created out of one or more paragraphs.

Botnet

One or more hosts that have been infected by software giving a remote attacker full administrative control through use of a command and control server.

C**CI**

Concern Index; primary means by which the system notifies you of suspicious activity; the Stealthwatch Flow Collector adds points to the CI whenever it observes possible attacks on one or more hosts; the greater the CI, the greater the level of concern over that behavior; when activity surpasses 100% of the CI threshold set for a particular host, the Flow Collector generates a High Concern Index alarm against that host.

CI Event

(This term is used in the SMC client interface.) Occurs when a packet is sent intentionally to evoke a response from a network host, such as during a DoS attack or scanning activity. CI events could indicate a security breach, a mis-configured device, malfunctioning server, or other source of networking issues. Every time a host is seen engaging in a CI event, Stealthwatch increases the number of CI points for that host. If the number of points exceeds the set threshold, Stealthwatch raises a High Concern Index alarm against the host that is the source of the activity.

Condition tag

A marker that you can apply to different areas of your content so that some sections show up in some of your outputs but not in others.

Cross-reference

A navigation link that lets you connect text in one topic to another topic (or a bookmark within a topic). Cross-references let you create "automated" links that are based on commands you provide. This allows you to keep links consistent and change them in just one place by using the "xref" style.

CSV

Comma-Separated Value

D

Deduplication

Process by which Stealthwatch automatically selects the best exporter for a given source-IP-to-destination-IP pair, thereby ensuring that any flows that might have traveled through more than one router are counted only once.

DHCP

Dynamic Host Configuration Protocol; communications protocol that allows system administrators to centrally manage and automate the assignment of IP addresses in a network; especially useful in education and other environments where users change frequently; supports static addresses for computers that need permanent IP addresses, such as Web servers.

DoS

Denial of Service; instances of servers being prevented from fulfilling requests for services, such as during an ICMP flood.

Drop-down text

A feature that lets you collapse content in your topic. The content is expanded (and therefore displayed) when the end user clicks a link.

DSCP

Differentiated Services Code Point; computer networking architecture that specifies a simple, scalable, and coarse-grained mechanism for classifying and managing network traffic and for providing QoS guarantees on modern IP networks; essentially, a field in the header of IP packets that is used for packet classification purposes. For example, DSCP can be used to provide low-latency, guaranteed service to critical network traffic, such as voice or video, while providing simple best-effort traffic guarantees to non-critical services such as Web traffic or file transfers.

E

External Event

Data received from technologies other than flow export technologies and associated with behavior-based, flow-driven events.

F

Flow

In general, traffic traveling in one direction (unidirectional) for a single, specific purpose, such as a client requesting access to a server; the server's response to the client is considered another flow. However, in Stealthwatch, a flow is the two-way (bidirectional) "conversation" between these two hosts, consisting of the client, the server, and a service port.

Footnote

A comment that is used to explain a specific area of the text. Both the area in the text and the comment contain a number or symbol that ties the two together. A footnote (or endnote) comment can be placed at the end of a page, document, chapter, section, or book.

FSI

File Sharing Index; tracks atypical file transfers, which could indicate P2P activity, that could place an organization at risk; once the file-sharing activity for a particular host causes its FSI points to surpass 100% of its FSI threshold, the Stealthwatch Flow Collector generates a High File Sharing Index alarm.

G

GMT

Greenwich Mean Time; see UTC.

H

Host

Computer, router, switch, or any other network element that has an IP address and is creating TCP/IP traffic on the network.

I

ICMP

Internet Control Message Protocol; used by IP to transmit maintenance information and error reports, such as unreachable host or incorrect parameter.

IDS

Intrusion Detection System

Inactive Host

Host that is known to the Stealthwatch Flow Collector, but has that not transmitted a data packet in the last 5 minutes.

IPFIX

IP Flow Information Export; open source version of NetFlow v9 and follows similar methodologies, such as flow templates.

ISP

Internet Service Provider

M

MAC Address

Media Access Control address; a unique identifier assigned to most network interface cards by the manufacturer.

MPLS

Multiprotocol Label Switching; highly scalable, protocol agnostic, data-carrying mechanism used in high-performance telecommunications networks that directs and carries data from one network node to the next, making it easy to create "virtual links" between distant nodes. Data packets are assigned labels and packet-forwarding decisions are made solely on the contents of this label, without the need to examine the packet itself; allows the creation of end-to-end circuits across any type of transport medium, using any protocol; eliminates dependence on a particular data link layer technology, such as frame relay or Ethernet; eliminates the need for multiple Layer 2 networks to satisfy different types of traffic.

MyTerm

This is a test. Delete me.

N

NAT

Network Address Translation

NBA

Network Behavior Analysis

NBAR

Network Based Application Recognition is a mechanism used by some Cisco equipment to identify applications by performing deep packet inspection.

NetFlow

Flow technology that started as proprietary software developed by Cisco Systems and has been adopted and adapted by other hardware vendors; a NetFlow-enabled router sends all flow information to the Stealthwatch Flow Collector for NetFlow; tracks every packet, allowing you to see everyone who is communicating through each router and fairly accurately determine the amount of traffic going through it.

New Host

Host that was previously unknown to the Stealthwatch Flow Collector (i.e., not in the host database until now) and that has completed a valid flow.

NIC

Network Interface Card

O

OPSEC

Open Platform for Security

OS

Operating System

OVF

Open Virtualization Format

P

P2P

Peer-to-Peer; software that enables commercial or private Internet users to communicate or share resources.

Phantom Host

Host that was previously unknown to the Stealthwatch Flow Collector and that was sent a data packet, but did not respond within 10 seconds. A large number of phantom hosts can be an indication of network scanning.

Ping_Oversized_Packet

Indicates that a source host has sent an ICMP echo request or reply that is larger than 90 data bytes. These events may be harmless network health checks or may contain a covert data channel.

Q

QoS

Quality of Service; resource reservation control mechanism that assigns different priority levels to different applications, users, or data flows to guarantee a certain level of performance. QoS guarantees are important if the network capacity is limited, especially for real-time streaming multimedia applications such as voice-over-IP, online games, and IP-TV.

R

RADIUS

Remote Authentication Dial-In User Service

RTM

Response Time Management

RTT

Round-Trip Time; measures the network, providing the round-trip time from the client to the server and back. This number is the same value you would

obtain if you pinged another machine on the network from your computer.

S

Security Event

(This term is used in the SMC Web App interface.) Occurs when a packet is sent intentionally to evoke a response from a network host, such as during a DoS attack or scanning activity. Security events could indicate a security breach, a misconfigured device, malfunctioning server, or other source of networking issues. Every time a host is seen engaging in a security event, Stealthwatch increases the number of index points for that host. If the number of points exceeds the set threshold, Stealthwatch raises an alarm against the host that is the source of the activity.

Single-Sourcing

Reusing content and producing multiple outputs from the same set of source files. Flare lets you single-source your projects in many ways, using various features. This includes features such as topic-based authoring, conditions, snippets, variables, multiple tables of contents, and more.

Snippet

A pre-set chunk of content that you can use in your project over and over. Snippets are similar to variables, but snippets are used for longer chunks of content that you can format just as you would any other content in your topic. In snippets, you can also insert tables, pictures, and whatever else can be included in a normal topic.

SNMP

Simple Network Management Protocol; UDP-based network protocol used mostly in network management systems to monitor network-attached devices for conditions that warrant administrative attention; exposes management data in the form of variables on the managed systems, which describe the system configuration. These variables can then be queried (and sometimes set) by managing applications.

Span

A tag that is used to group inline elements to format them with styles. A span tag doesn't perform any specific action; it simply holds the attributes (e.g.,

font size, color, font family) that you apply to inline content.

SRT

Server Response Time; measures the application, providing the amount of time it took the server side of a flow to answer incoming requests. In the case of an HTTP server, the SRT would reflect how long it takes the server to begin sending HTML and images back to the client's browser.

Style

An element to which you assign a certain look and/or behavior. You can then apply that style to your content. Different kinds of styles are available in a stylesheet, to be used for various purposes in your content.

T

Table

A group of intersecting columns and rows that you can add to a topic for various purposes, such as comparing one thing with another or giving field descriptions for a software dialog.

TACACS

Terminal Access Controller Access Control System

Target

One "instance" of an output type. When you build your final output, you are essentially building one or more of the targets in your project.

Text snippet

A snippet that is created out of a portion of one paragraph.

TI

Target Index; tracks hosts that appear to be victims of the suspicious behavior of other hosts. Every time a Stealthwatch Flow Collector sees that a host is the target of suspicious behavior, it adds points to the TI for that host. When the number of points exceeds 100% of the acceptable TI threshold for that host, the Flow Collector generates a High Target Index alarm, alerting you to the possibility that the host is under attack.

Topic

A chunk of information about a particular subject. Topics are the most important part of a project. Everything else is contained within topics (e.g., hyperlinks, text, pictures) or points toward topics (e.g., table of contents, index, browse sequences). The very reason end users open a Help system is to find information, a little direction. They find that help within individual topics.

U**UTC**

Universal Time Coordinated; the 24-hour time standard by which the world regulates time; based on the 0° longitude meridian in Greenwich, England; also known as Greenwich Mean Time (GMT) or Zulu (Z) Time. The Stealthwatch system uses UTC as a basis for calculating time. However, the SMC automatically converts UTC to display your local time throughout the user interface.

V**Variable**

A pre-set term or content that you can use in your project over and over. Variables are similar to snippets, but variables are used for brief, non-formatted pieces of content (such as the name of your company's product or your company's phone number).

VM

Virtual Machine; environment, usually a program or operating system, which does not physically exist, but is created within another environment. One host environment (VM server/ESX server) often can run multiple VMs at once. In Stealthwatch, each VM provides traffic reporting. Use the VM Status document to view all VMs in the system and their current operational status.

VM Server

ESX server; physical device running VMware ESX that hosts one or more VMs. Each VM server provides traffic reporting. The traffic shown is an aggregation of all traffic found within its current list of VMs. Use the VM Server Status document to view all VM servers in the system and their current operational statistics.

VPN

Virtual Private Network; service that allows users who are connected to an outside service provider to access restricted services and applications.

X

XML Editor

The window in the Flare interface where you can add content and formatting to elements such as topics and snippets.

Z

Z Time

Zulu Time; see UTC.

Index

A

advanced settings 1646, 1652
agent community 1617
audit log 1700
auto negotiate 1610

B

backup 1639, 1670, 1698
banner 1632
broadcast 1610
browse files 1677
bytes 1664

C

capacity 1664
Cisco Secure 1569
configuration, backup or restore 1698
cpu 1593

D

database storage statistics 1664
database, backup or restore 1639, 1670
date time 1615
description of SMC 1704
diagnostic 1699
disk usage 1593, 1664
dns 1613-1614
duplex 1593, 1610

E

engine status 1593

F

file sharing 1629

files 1677

flow collector 1571

FlowCollector 1571

flowreplicator 1584

flowsensor, description 1578, 1580

G

gateway 1610

global settings 1632

H

HTTP 1593

I

interface speed 1610

interfaces 1593

IPFIX 1573

L

log files 1677

M

mac address 1593, 1610

mask 1610

memory 1593

mount point 1593

mtu 1610

N

naming 1613-1614

NetFlow 1571, 1573
network interface 1610
ntp server 1615

O

oid 1617
opening message 1632
operations 1701

P

password 1632
pid 1593
power off 1701

R

reboot 1701
remote file system 1629
resolution 1613-1614
restart 1701
restore 1639, 1670, 1698

S

services 1593
session timeout 1632
sFlow 1576
shutdown 1701
smc description 1704
snmp agent 1617
statistics, database 1664
subnet mask 1610
system settings 1593

T

Tcpdump 1680

time out 1632

time zone 1615

timeout 1632

Tomcat 1593

troubleshooting 1699

V

ve 1580

virtual flowsensor 1580

X

Xe 1571

Glossary

A

Active Host

Host that is known to the Stealthwatch Flow Collector and that has transmitted a data packet in the last 5 minutes.

Alert

Indication of unusual network activity that has not necessarily resulted in an alarm. (This term is not used in the SMC Web App interface.)

Archive Hour

User-defined time when each Stealthwatch Flow Collector in the domain resets itself. Acceptable values are 0 to 23, where 0 is midnight local time. At the defined time, the Flow Collector resets all index counts to zero, saves the log and Web files it has gathered during the preceding 24 hours, and then begins a new day (24 hours) of data collection. Defined on the Domain page of the Domain Properties dialog.

ASN

Autonomous System Number; network of connected IP routing prefixes under the control of one or more network operators that presents a common, clearly defined routing policy to the Internet.

B

Batch target

A special target that lets you build and/or publish multiple other targets in a single group (or "batch"). You can schedule batches to run at any time.

Beaconing

Communication between an internal and outside host (with traffic in only one direction) that exceeds the number of seconds required for a flow to qualify as long duration.

Block snippet

A snippet that is created out of one or more paragraphs.

Botnet

One or more hosts that have been infected by software giving a remote attacker full administrative control through use of a command and control server.

C**CI**

Concern Index; primary means by which the system notifies you of suspicious activity; the Stealthwatch Flow Collector adds points to the CI whenever it observes possible attacks on one or more hosts; the greater the CI, the greater the level of concern over that behavior; when activity surpasses 100% of the CI threshold set for a particular host, the Flow Collector generates a High Concern Index alarm against that host.

CI Event

(This term is used in the SMC client interface.) Occurs when a packet is sent intentionally to evoke a response from a network host, such as during a DoS attack or scanning activity. CI events could indicate a security breach, a mis-configured device, malfunctioning server, or other source of networking issues. Every time a host is seen engaging in a CI event, Stealthwatch increases the number of CI points for that host. If the number of points exceeds the set threshold, Stealthwatch raises a High Concern Index alarm against the host that is the source of the activity.

Condition tag

A marker that you can apply to different areas of your content so that some sections show up in some of your outputs but not in others.

Cross-reference

A navigation link that lets you connect text in one topic to another topic (or a bookmark within a topic). Cross-references let you create "automated" links that are based on commands you provide. This allows you to keep links consistent and change them in just one place by using the "xref" style.

CSV

Comma-Separated Value

D

Deduplication

Process by which Stealthwatch automatically selects the best exporter for a given source-IP-to-destination-IP pair, thereby ensuring that any flows that might have traveled through more than one router are counted only once.

DHCP

Dynamic Host Configuration Protocol; communications protocol that allows system administrators to centrally manage and automate the assignment of IP addresses in a network; especially useful in education and other environments where users change frequently; supports static addresses for computers that need permanent IP addresses, such as Web servers.

DoS

Denial of Service; instances of servers being prevented from fulfilling requests for services, such as during an ICMP flood.

Drop-down text

A feature that lets you collapse content in your topic. The content is expanded (and therefore displayed) when the end user clicks a link.

DSCP

Differentiated Services Code Point; computer networking architecture that specifies a simple, scalable, and coarse-grained mechanism for classifying and managing network traffic and for providing QoS guarantees on modern IP networks; essentially, a field in the header of IP packets that is used for packet classification purposes. For example, DSCP can be used to provide low-latency, guaranteed service to critical network traffic, such as voice or video, while providing simple best-effort traffic guarantees to non-critical services such as Web traffic or file transfers.

E

External Event

Data received from technologies other than flow export technologies and associated with behavior-based, flow-driven events.

F

Flow

In general, traffic traveling in one direction (unidirectional) for a single, specific purpose, such as a client requesting access to a server; the server's response to the client is considered another flow. However, in Stealthwatch, a flow is the two-way (bidirectional) "conversation" between these two hosts, consisting of the client, the server, and a service port.

Footnote

A comment that is used to explain a specific area of the text. Both the area in the text and the comment contain a number or symbol that ties the two together. A footnote (or endnote) comment can be placed at the end of a page, document, chapter, section, or book.

FSI

File Sharing Index; tracks atypical file transfers, which could indicate P2P activity, that could place an organization at risk; once the file-sharing activity for a particular host causes its FSI points to surpass 100% of its FSI threshold, the Stealthwatch Flow Collector generates a High File Sharing Index alarm.

G

GMT

Greenwich Mean Time; see UTC.

H

Host

Computer, router, switch, or any other network element that has an IP address and is creating TCP/IP traffic on the network.

I

ICMP

Internet Control Message Protocol; used by IP to transmit maintenance information and error reports, such as unreachable host or incorrect parameter.

IDS

Intrusion Detection System

Inactive Host

Host that is known to the Stealthwatch Flow Collector, but has that not transmitted a data packet in the last 5 minutes.

IPFIX

IP Flow Information Export; open source version of NetFlow v9 and follows similar methodologies, such as flow templates.

ISP

Internet Service Provider

M

MAC Address

Media Access Control address; a unique identifier assigned to most network interface cards by the manufacturer.

MPLS

Multiprotocol Label Switching; highly scalable, protocol agnostic, data-carrying mechanism used in high-performance telecommunications networks that directs and carries data from one network node to the next, making it easy to create "virtual links" between distant nodes. Data packets are assigned labels and packet-forwarding decisions are made solely on the contents of this label, without the need to examine the packet itself; allows the creation of end-to-end circuits across any type of transport medium, using any protocol; eliminates dependence on a particular data link layer technology, such as frame relay or Ethernet; eliminates the need for multiple Layer 2 networks to satisfy different types of traffic.

MyTerm

This is a test. Delete me.

N

NAT

Network Address Translation

NBA

Network Behavior Analysis

NBAR

Network Based Application Recognition is a mechanism used by some Cisco equipment to identify applications by performing deep packet inspection.

NetFlow

Flow technology that started as proprietary software developed by Cisco Systems and has been adopted and adapted by other hardware vendors; a NetFlow-enabled router sends all flow information to the Stealthwatch Flow Collector for NetFlow; tracks every packet, allowing you to see everyone who is communicating through each router and fairly accurately determine the amount of traffic going through it.

New Host

Host that was previously unknown to the Stealthwatch Flow Collector (i.e., not in the host database until now) and that has completed a valid flow.

NIC

Network Interface Card

O

OPSEC

Open Platform for Security

OS

Operating System

OVF

Open Virtualization Format

P

P2P

Peer-to-Peer; software that enables commercial or private Internet users to communicate or share resources.

Phantom Host

Host that was previously unknown to the Stealthwatch Flow Collector and that was sent a data packet, but did not respond within 10 seconds. A large number of phantom hosts can be an indication of network scanning.

Ping_Oversized_Packet

Indicates that a source host has sent an ICMP echo request or reply that is larger than 90 data bytes. These events may be harmless network health checks or may contain a covert data channel.

Q

QoS

Quality of Service; resource reservation control mechanism that assigns different priority levels to different applications, users, or data flows to guarantee a certain level of performance. QoS guarantees are important if the network capacity is limited, especially for real-time streaming multimedia applications such as voice-over-IP, online games, and IP-TV.

R

RADIUS

Remote Authentication Dial-In User Service

RTM

Response Time Management

RTT

Round-Trip Time; measures the network, providing the round-trip time from the client to the server and back. This number is the same value you would

obtain if you pinged another machine on the network from your computer.

S

Security Event

(This term is used in the SMC Web App interface.) Occurs when a packet is sent intentionally to evoke a response from a network host, such as during a DoS attack or scanning activity. Security events could indicate a security breach, a misconfigured device, malfunctioning server, or other source of networking issues. Every time a host is seen engaging in a security event, Stealthwatch increases the number of index points for that host. If the number of points exceeds the set threshold, Stealthwatch raises an alarm against the host that is the source of the activity.

Single-Sourcing

Reusing content and producing multiple outputs from the same set of source files. Flare lets you single-source your projects in many ways, using various features. This includes features such as topic-based authoring, conditions, snippets, variables, multiple tables of contents, and more.

Snippet

A pre-set chunk of content that you can use in your project over and over. Snippets are similar to variables, but snippets are used for longer chunks of content that you can format just as you would any other content in your topic. In snippets, you can also insert tables, pictures, and whatever else can be included in a normal topic.

SNMP

Simple Network Management Protocol; UDP-based network protocol used mostly in network management systems to monitor network-attached devices for conditions that warrant administrative attention; exposes management data in the form of variables on the managed systems, which describe the system configuration. These variables can then be queried (and sometimes set) by managing applications.

Span

A tag that is used to group inline elements to format them with styles. A span tag doesn't perform any specific action; it simply holds the attributes (e.g.,

font size, color, font family) that you apply to inline content.

SRT

Server Response Time; measures the application, providing the amount of time it took the server side of a flow to answer incoming requests. In the case of an HTTP server, the SRT would reflect how long it takes the server to begin sending HTML and images back to the client's browser.

Style

An element to which you assign a certain look and/or behavior. You can then apply that style to your content. Different kinds of styles are available in a stylesheet, to be used for various purposes in your content.

T

Table

A group of intersecting columns and rows that you can add to a topic for various purposes, such as comparing one thing with another or giving field descriptions for a software dialog.

TACACS

Terminal Access Controller Access Control System

Target

One "instance" of an output type. When you build your final output, you are essentially building one or more of the targets in your project.

Text snippet

A snippet that is created out of a portion of one paragraph.

TI

Target Index; tracks hosts that appear to be victims of the suspicious behavior of other hosts. Every time a Stealthwatch Flow Collector sees that a host is the target of suspicious behavior, it adds points to the TI for that host. When the number of points exceeds 100% of the acceptable TI threshold for that host, the Flow Collector generates a High Target Index alarm, alerting you to the possibility that the host is under attack.

Topic

A chunk of information about a particular subject. Topics are the most important part of a project. Everything else is contained within topics (e.g., hyperlinks, text, pictures) or points toward topics (e.g., table of contents, index, browse sequences). The very reason end users open a Help system is to find information, a little direction. They find that help within individual topics.

U**UTC**

Universal Time Coordinated; the 24-hour time standard by which the world regulates time; based on the 0° longitude meridian in Greenwich, England; also known as Greenwich Mean Time (GMT) or Zulu (Z) Time. The Stealthwatch system uses UTC as a basis for calculating time. However, the SMC automatically converts UTC to display your local time throughout the user interface.

V**Variable**

A pre-set term or content that you can use in your project over and over. Variables are similar to snippets, but variables are used for brief, non-formatted pieces of content (such as the name of your company's product or your company's phone number).

VM

Virtual Machine; environment, usually a program or operating system, which does not physically exist, but is created within another environment. One host environment (VM server/ESX server) often can run multiple VMs at once. In Stealthwatch, each VM provides traffic reporting. Use the VM Status document to view all VMs in the system and their current operational status.

VM Server

ESX server; physical device running VMware ESX that hosts one or more VMs. Each VM server provides traffic reporting. The traffic shown is an aggregation of all traffic found within its current list of VMs. Use the VM Server Status document to view all VM servers in the system and their current operational statistics.

VPN

Virtual Private Network; service that allows users who are connected to an outside service provider to access restricted services and applications.

X

XML Editor

The window in the Flare interface where you can add content and formatting to elements such as topics and snippets.

Z

Z Time

Zulu Time; see UTC.