

Network Security Platform

7.3.0

Product Guide

Revision A

© 2025 Forcepoint
Forcepoint and the FORCEPOINT logo are trademarks of Forcepoint.
All other trademarks used in this document are the property of their respective owners.

Published 10 June 2025

Every effort has been made to ensure the accuracy of this document. However, Forcepoint makes no warranties with respect to this documentation and disclaims any implied warranties of merchantability and fitness for a particular purpose. Forcepoint shall not be liable for any error or for incidental or consequential damages in connection with the furnishing, performance, or use of this manual or the examples herein. The information in this documentation is subject to change without notice.

Contents

Preface 21
Open the online help 21

Introduction to the Forcepoint Network Security Platform solution 23

1 The Forcepoint Network Security Platform solution 25
Security Management Center (SMC) 25
Components in the Forcepoint Network Security Platform solution 26
Benefits of the SMC 28
Licensing components and features 29

2 Introduction to Forcepoint Network Security Platform in the Engine/VPN role 31
Overview of Forcepoint Network Security Platform in the Engine/VPN role 31
Forcepoint Network Security Platform benefits 32

3 Introduction to Forcepoint Network Security Platform in the IPS and Layer 2 Engine roles 41
What IPS engines and Layer 2 Engines do 41
How IPS engines and Layer 2 Engines respond to incidents 42
Main benefits of IPS engines and Layer 2 Engines 43
IPS Cluster load balancing 44
Disconnect mode for IPS engines and Layer 2 Engines and how it works 44

Deployment 47

4 Deploying the SMC 49
Overview of SMC deployment 49
Security considerations for SMC deployment 50
Positioning the Management Server 51
Positioning Log Servers 52
Positioning SMC Clients 52
Alternative methods for accessing the SMC Client 52
Example: SMC deployment 53
Post-installation steps for the SMC 53

5 Deploying Forcepoint Network Security Platform in the Engine/VPN role 55
Supported platforms for Security Engine deployment 55
Deploying Security Engines on cloud-based virtualization platforms 56
Running Security Engines as Master Engines 57
Hardware requirements for installing Security Engine on third-party hardware 58
Hardware for Engine Cluster nodes 58
Guidelines for deploying Forcepoint Network Security Platform in the Engine/VPN role 59
Positioning Engines 59
Post-installation steps for Forcepoint Network Security Platform in the Engine/VPN role 64
Cable connection guidelines for Engines 65

6 Deploying Forcepoint Network Security Platform in IPS and Layer 2 Engine roles 67
Supported platforms for Security Engine deployment 67

Running Security Engines as Master Engines ... 68
Hardware requirements for installing Security Engine on third-party hardware ... 69
Guidelines for deploying IPS engines and Layer 2 Engines ... 69
Positioning IPS engines and Layer 2 Engines ... 70
Deploying IPS engines in IDS or IPS mode ... 73
IPS deployment examples ... 81
Post-installation steps for Forcepoint Network Security Platform in the IPS role ... 85
Deploying Layer 2 Engines in IPS or Passive Engine mode ... 85
Layer 2 Engine deployment example ... 88
Post-installation steps for Forcepoint Network Security Platform in the Layer 2 Engine role ... 89
Cable connection guidelines for IPS and Layer 2 Engines ... 90
Speed and duplex settings for Security Engines ... 93

Setting up ... 95

7 Using the SMC Client ... 97
SMC Client and how it works ... 97
Log on to the SMC ... 100
Log on to the SMC using certificate-based authentication ... 101
Customize the SMC Client layout ... 102
Bookmark SMC Client views ... 103
Change the logon view ... 106
Centralized management of global system settings ... 106
View, approve, and commit pending changes ... 107
Change the default language of the SMC Client ... 108
Using search features ... 108
Communicating with other administrators ... 113
Use Tools Profile elements to add commands to elements ... 114
Using the online help locally ... 115

8 Network address translation (NAT) and how it works ... 119
Network address translation and how it works ... 119
Static source translation ... 120
Dynamic source translation ... 121
Static destination translation ... 121
Destination port translation ... 122
IPv6 transition mechanisms ... 122

9 Configuring system communications ... 125
Considerations for setting up system communications ... 125
Define contact IP addresses ... 127
Select the Location for the SMC Client ... 142
Create HTTP Proxy elements ... 143
Configuring NTP for the SMC Appliance and Security Engine ... 144
Configuring SMC Appliance communications ... 146
Considerations for Multi-Link system communication ... 148

10 Managing certificates for system communications ... 149
How certificates work ... 149
Types of internal certificate authorities ... 150
Using certificates to secure communications to external components ... 154
Creating certificates ... 158

Configure settings for certificate validation 161
Renewing certificates 162
Renewing external certificates 166

11 Managing elements 179
Introduction to elements 179
Benefits of exporting or importing elements 189
Restore elements from Snapshots 193
Lock elements 197
Unlock elements 198
How the Trash works 198
How Categories help you view only certain elements 201
Legacy elements and options 207

Monitoring 209

12 Monitoring Forcepoint Network Security Platform components 211
Getting started with monitoring the system 211
System monitoring tools in the SMC Client 212
Overviews and how they work 227
Monitoring users on the Dashboard 233
Monitoring connections, block lists, VPN SAs, users, routing, SSL VPNs, and neighbors 236
View and compare Element Snapshot elements 243
Monitoring connections using Geolocation elements 244
Monitoring configurations and policies 247
Monitor administrator actions 248
Monitor tasks 248
Traffic captures and how they work 249
Checking maintenance contract information 250
Upcoming event notification 252

13 Application Health Monitoring 255
Enable Application Health Monitoring 255
View status of Network Applications 256
Disable Application Health Monitoring 260

14 Monitoring third-party devices 261
Getting started with monitoring third-party devices 261
Converting logs from third-party devices 262
Methods for monitoring third-party device status 273
Configuring Log Servers to monitor third-party devices 276
Activate monitoring of third-party devices 277
Configuring third-party devices for monitoring 278
Changing the ports for third-party device monitoring 279
Activate or deactivate third-party device monitoring alerts 280

15 Viewing and exporting logged data 281
Getting started with the Logs view 281
Browsing log data 288
Changing how log entries are displayed 299
Exporting data from the Logs view 301
Save data in PDF or HTML format 304

Create rules from log entries....306
Forwarding log data to an Elasticsearch cluster....307

16 Reports....311
Getting started with reports....311
Restricting a report's scope....313
Designing reports....315
Generate and view reports....322
Exporting reports....328
Status reporting....332
Example of reports....332

17 Filtering data....333
Getting started with filtering data....333
Creating filters....334
Organizing Filter elements....344
Examples of filters....346

Controlling Security Engines....347

18 Controlling Security Engine operation....349
Commanding Security Engines remotely....349
Set Security Engine options....356
Change a NetLink state manually....359
Disable cluster nodes temporarily....360
Re-enable disabled cluster nodes....361
Editing Security Engine configurations....362

19 Working on the Security Engine command line....363
Considerations for working on the Security Engine command line....363
Access the Security Engine command line....364
Reconfigure Security Engine settings....365
Create Security Engine scripts....367
Send commands to Virtual Security Engines....368

SMC configuration....371

20 Administrator accounts....373
Getting started with administrator accounts....373
How administrator accounts work....374
Default administrator account elements....375
Administrator account configuration overview....377
Creating Administrator Role and Access Control List elements....377
Add administrator accounts....379
Enforce an approval workflow....383
Restrict the log data an administrator can view....383
Customize log colors for administrators....384
Replicate administrator accounts....385
Enable and define password policy settings....386
Change administrator passwords....387
Authenticate administrators using OpenID authentication method....389

Authenticate administrators using RADIUS or TACACS+ methods....392
Authenticate administrators on engines using the Radius authentication method....394
Authenticate administrators using SAML v2 authentication method....396
Using LDAP authentication for administrators....398
Authenticate administrators using certificate-based authentication....399
Disable administrator accounts....405
Delete administrator accounts....406
API client accounts and how they work....407
Configure SMC API....407

21 Alert escalation....411
Alert escalation and how it works....411
Creating Alert elements....414
Configure notifications for alerts....416
Alert Chain elements and how they work....417
Creating Alert Policy elements....419
Install Alert Policy elements....420
Acknowledge active alerts....421
How custom scripts for alert escalation work....422
Create SMTP Server elements....423
Use a script for SMS notification....424
SNMP for the SMC Appliance....425
Test alerts....428
Examples of alert escalation....429

22 Domain elements....433
Getting started with Domain elements....433
How Domain elements work....434
Create and modify Domain elements....436
Log on to a Domain....438
Log off from all Domains....441
Move elements between Domains....442
View Domain status....443
Delete a Domain....443
Examples of Domain elements....444

23 Getting Started with the Web Portal....447
Enabling the Web Portal Client....448
Create Web Portal User accounts....449

24 Using the SMC Client in a web browser....451
Configuration overview....452
Start the SMC Client in a web browser....453

25 SMC Client downloads from the Management Server....455
Enable and configure SMC Client downloads on the Management Server....455
Download the SMC Client from the Management Server....456

26 Configuring the Log Server....457
Modify Log Server elements....457
Select backup Log Servers for high availability....458
Forwarding log data from Log Servers to external hosts....459
Edit Log Server configuration parameters....464

Certify Log Servers ... 465

27 Configuring SMC servers for high availability ... 467
Using additional SMC servers for high availability ... 467
Management Server HA configuration overview ... 468
Log Server HA configuration overview ... 472
Manage HA Management Servers and Log Servers ... 476

28 Reconfiguring the SMC and Security Engines ... 481
Modify Management Server elements ... 481
Configure SMC API ... 483
Forward audit data from Management Servers to external hosts ... 486
Change the Management Server database password ... 489
Change the Management Server or Log Server platform ... 490
Change Management Server and Log Server IP addresses ... 491
Troubleshooting connecting to Management Servers ... 495
Things to consider when changing the Security Engine role ... 496

Security Engine configuration ... 501

29 Creating and modifying Security Engines ... 503
Getting started with Security Engines ... 503
Creating Security Engine elements ... 504
Editing existing Security Engines ... 509
Configure a backup unit or a connection synchronization for external high availability ... 524
Configure global contact policy settings for node-initiated contact to the Management Server ... 524
Synchronizing the time on Security Engines ... 526

30 Creating and modifying Master Engine and Virtual Engine elements ... 527
Configuration of Master Engines and Virtual Engines ... 527
Create Master Engines or Virtual Engines ... 530
Create Virtual Resource elements ... 531
Add additional nodes to Master Engines ... 532
Moving a Virtual Engine to a different Master Engine ... 533
Convert Engines to Master Engines and Virtual Engine elements ... 533
Example: deploying Virtual Engines for MSSP customers ... 541

31 Network interface configuration ... 543
Network interfaces for Security Engine ... 543
Configuring interfaces for Engines ... 547
Configuring interfaces for IPS engines ... 572
Configuring interfaces for Layer 2 Engines ... 585
Configuring interfaces for Master Engines ... 598
Configuring interfaces for Virtual Engines ... 609
Add manual ARP entries to Security Engine ... 615
Examples of interface configurations ... 615

32 Connecting Security Engine to the SMC ... 629
Management connections for Security Engines and how they work ... 629
Configuration overview of connecting Security Engines to the SMC ... 630
Connect Security Engines to the SMC ... 631

33 Element-based network address translation (NAT) ... 633

Element-based NAT and how it works ... 633
Add NAT definitions for element-based NAT ... 634
Edit or remove NAT definitions for element-based NAT ... 635

34 Configuring the Security Engine tester ... 637
Getting started with the Security Engine tester ... 637
Specify global tester settings ... 638
Add Security Engine tests ... 638
Configuring additional test-specific settings for Security Engine tests ... 639
View configured Security Engine tests ... 643
Remove Security Engine tests ... 643
Disable or enable Security Engine tests ... 644

35 Engine permissions ... 647
Getting started with permissions ... 647
Define administrator permissions for Security Engines ... 648
Select the allowed policies ... 649

36 DNS Relay ... 651
Getting started with DNS relay ... 651
Enable DNS relay ... 654

37 Setting up SNMP for Security Engines ... 657
Getting started with SNMP configuration for Security Engines ... 657
Create an SNMP Agent for SNMP version 1 or 2c ... 659
Create an SNMP Agent for SNMP version 3 ... 660
Configure what triggers SNMP traps ... 661
Activate the SNMP agent on Security Engines ... 662
Configure Engine access rules to allow SNMP queries from trusted SNMP managers ... 663

38 Setting up LLDP for Security Engines ... 665
Getting started with LLDP for Security Engines ... 665
LLDP configuration overview ... 666
Create custom LLDP Profile elements ... 666
Enable LLDP on Security Engines ... 667

39 Alias element translations for Security Engines ... 669
Getting started with Alias element translations ... 669
Add Alias element translation values ... 670
Remove Alias element translation values ... 671

40 Add-on features for Security Engines ... 673
Getting started with add-on features ... 673
Edit add-on settings for Security Engines ... 674

41 Advanced Security Engine settings ... 675
Getting started with advanced Security Engine settings ... 675
Open the advanced settings ... 676
Adjusting Engine clustering options ... 676
Adjust IPS clustering options ... 679
Adjust Layer 2 Engine clustering options ... 680
Adjust Master Engine clustering options ... 681
Configure inspection of tunneled traffic ... 683
Set connection timeouts ... 684

Configure SYN rate limits ... 684
Configure log handling settings ... 685
Configure DoS protection settings ... 686
Configure scan detection settings ... 687
Using custom properties profiles to upload custom scripts ... 687

Routing ... 691

42 Configuring routing and antispoofing ... 693
Getting started with routing ... 693
Routing configuration overview ... 694
Add routers ... 696
Add or view the default route ... 697
Add static routes ... 697
Using metrics or ECMP on multiple routes to the same destination ... 698
Configure Forced Next Hop routing ... 701
Configure policy routing ... 703
Configuring multicast routing ... 703
Configure DHCP message routing ... 707
Check routes using the Route Query tool ... 709
Remove static routes ... 709
Modifying antispoofing ... 710
Examples of routing configuration ... 713

43 Configuring dynamic routing ... 717
Getting started with dynamic routing ... 717
Dynamic routing configuration overview ... 718
Creating elements for dynamic routing ... 718
Configure BGP ... 722
Configure OSPFv2 ... 724
Configure PIM ... 726
Preview the dynamic routing configuration in FRR syntax ... 727
Using the command line to configure dynamic routing ... 728
Restart dynamic routing processing ... 730

44 Outbound traffic management ... 731
Getting started with outbound traffic management ... 731
Defining Multi-Link routes ... 735
Enable outbound traffic management using element-based NAT ... 739
Manually configuring outbound traffic management ... 739
Monitoring and testing outbound traffic management ... 744
Examples of manually configuring Multi-Link ... 745

45 Inbound traffic management ... 749
Getting started with inbound traffic management ... 749
Create Server Pool elements ... 753
Configure server availability monitoring ... 755
Create Access rules to allow the type of traffic that is handled by the Server Pool ... 756
Enable Server Pool load balancing using NAT rules ... 757
Enable Server Pool load balancing using Access rules ... 758
Configuring dynamic DNS updates for Server Pools ... 759
Using Server Pool Monitoring Agents ... 762

Examples of Server Pools ... 780

46 Dynamic link selection ... 783
Getting started with dynamic link selection ... 783
Using dynamic link selection with QoS Class elements ... 785
Create Connection Type elements ... 788
Create Link Usage Profile elements ... 789
Select a Link Usage Profile element for an Security Engine ... 790
Define exceptions to the Link Usage Profile for an Security Engine ... 790
Select a Link Usage Profile for a policy-based VPN ... 791
Select a Link Usage Profile for a Route-based Tunnels group ... 791
Select a Link Usage Profile for a VPN Broker domain ... 792

Traffic inspection policies ... 793

47 Creating and managing policy elements ... 799
Getting started with policies ... 799
Create template policies or policies ... 808
How sub-policies work ... 810
Install policies ... 813
Using policy elements and rules ... 816
Deleting policies templates or sub-policies ... 819
Engine Policy elements examples ... 819
IPS Policy example ... 821
Local alternative policies ... 822

48 Ethernet rules ... 825
Getting started with Ethernet rules ... 825
Configuration of Ethernet rules ... 826
Examples of Ethernet rules ... 829

49 Access rules ... 831
Getting started with Access rules ... 831
Overview of Access rules ... 831
Configuring Access rules ... 833
Using Access rules ... 838
Examples of engine Access rules ... 845
Examples of IPS Access rules ... 848

50 NAT rules ... 851
Getting started with NAT rules ... 851
Configuring NAT rules ... 853
NAT and system communications ... 856
Outbound load-balancing NAT ... 859
Proxy ARP and NAT ... 859
Protocols and NAT ... 859
Examples of NAT configuration ... 860
Examples of NAT rules ... 862

51 Inspection Policy elements ... 867
Inspection Policy elements and how they work ... 867
How Inspection Policy elements are designed ... 868
Set default options for Exception rules in Inspection Policy elements ... 874

Example: Tuning an Inspection Policy element to eliminate false positives for a engine....875

52 Snort inspection on Security Engines....877
Getting started with Snort inspection on Security Engines....877
Prepare Snort configuration files....878
Import Snort configuration files globally for all Security Engines....879
Enable Snort inspection for Security Engines....880
Override settings in the global Snort configuration for individual Security Engines....881
Add Access rules for Snort inspection....881
Logging for Snort inspection....882

53 Editing policies....885
Getting started with editing policies....885
The different parts of the policy editing view....887
Editing rules in a policy....889
Add Insert Points in Policy Templates....896
Automatic rules and how they work....896
Configure settings for Automatic rules....897
Add Ethernet rules....897
Add Access rules....899
Add NAT rules....903
Add Inspection rules....904
Add Exception rules....908
Specify rule validity times....912
Validate rules automatically....912
How default rules can be changed....917

54 Defining IP addresses....919
Defining IP addresses as elements....919
Access and modify network elements....924
Edit Expression elements....924
Using SMC elements to represent IP addresses in policies....925

55 Working with Service elements....927
Getting started with Service elements....927
Creating Service elements....928
Protocol elements and how they work....931
Defining Protocol parameters....931

56 Defining Situation elements....941
Getting started with Situation elements....941
Situations configuration overview....942
Create custom Situation elements....943
Context options for Situation elements....944
Defining Context Options for Correlation Situation elements....946
Default elements for Situation elements....952
Using Tags with Situation elements....953
Vulnerability elements and how they work....955
Using Situation elements....957
Examples of custom Situation elements....957

57 Using Network Application elements....961
Getting started with Network Application elements....961

Create TLS Match elements for network application detection....963
Access rules for network application detection....964
Example: blocking network application use....967

58 Defining User Response elements....969
User Response elements and how they work....969
Create User Response elements....971

59 Quality of Service....973
Quality of Service (QoS) and how it works....973
Create QoS Class elements....981
Define QoS Policy elements....982
Apply QoS to traffic....985
Examples of bandwidth management and traffic prioritization scenarios....986

60 Anti-malware scanning....989
Getting started with anti-malware scanning....989
Selecting traffic for anti-malware scanning....990
Enable anti-malware on the Security Engine....990
Manually update the anti-malware database....991
View the status of the anti-malware database....992

61 File filtering....993
How file filtering works....993
Integrate on-premises DLP servers with Forcepoint Network Security Platform....995
Integrate file reputation services and sandboxes....997
Integrate McAfee GTI file reputation with Forcepoint Network Security Platform....997
Forcepoint Advanced Malware Detection and how it works....998
Integrate a Forcepoint Advanced Malware Detection appliance with Forcepoint Network Security Platform....1003
Define Sandbox Service elements....1003
Connect Forcepoint Network Security Platform to a sandbox service....1004
View sandbox analysis reports....1006
Restrict file types with file filtering....1007
Support for McAfee Advanced Threat Defense....1009

62 Integrating Forcepoint One Endpoint with Forcepoint Network Security Platform....1011
Forcepoint One Endpoint and how it works....1011
Create ECA Configuration elements....1014
Enable Forcepoint Endpoint Context Agent (ECA) on the Security Engine....1015
Define Endpoint Application elements....1015
Create Endpoint Settings elements....1016
Use endpoint information in Access rules....1017
Enable logging of endpoint information....1017

63 Filtering URLs....1019
URL filtering and how it works....1019
Enable ThreatSeeker....1022
Use an HTTP proxy to connect to the ThreatSeeker Intelligence Cloud server....1023
Add Access rules for category-based URL filtering....1024
Add URL List Application elements to manually block or allow URLs....1025
Add Access rules for custom URL List Applications....1027
Examples of URL filtering....1028

64 Protocol Agents on Security Engines 1031
Protocol Agents overview 1031
Configuring Protocol Agents 1033
Using Protocol Agents 1034
Examples of Protocol Agents 1040

65 Sidewinder Proxies 1043
Sidewinder Proxies and how they work 1043
Using Sidewinder Proxies 1048
Change logging options for Sidewinder Proxies 1050
Enable Sidewinder Proxy 1050
Configure Sidewinder SSH Proxy 1051
Create custom Service elements for Sidewinder Proxies 1056
Add rules for Sidewinder Proxy 1057
Advanced settings for Sidewinder Proxies 1057
Supported advanced Sidewinder Proxy settings 1058

66 Setting up TLS inspection 1063
TLS inspection and how it works 1063
Configure TLS inspection for server protection 1066
Configure TLS inspection for client protection 1068
Define trusted certificate authorities for TLS inspection 1075
Exclude traffic from decryption for TLS inspection 1076
Active destination server certificate probing 1079
Examples of TLS inspection 1080

67 Setting up QUIC inspection 1083
QUIC inspection and how it works 1083
Verify QUIC inspection settings on Security Engine 1083
Examples of QUIC Inspection 1085

68 Forward traffic to a proxy service for external inspection 1087
Getting started with forwarding traffic 1087
Create a Proxy Server element 1089
Add Access rules to forward traffic 1089
Add NAT rules to forward traffic 1090
Example: Using Access rules to forward traffic 1091
Example: Using NAT rules to forward traffic 1092
Configuring Explicit HTTP Proxy (Experimental) 794
Create a service element for HTTP Explicit proxy 795
Configuring Integrated Windows Authentication 796
Add access rules for Explicit HTTP Proxy 797

69 Block listing IP addresses 1093
Block listing traffic and how it works 1093
Add Access rules for block listing 1096
Configure automatic block listing of traffic 1097
Block list traffic manually 1098
Monitoring Block listing 1099

Users and authentication 1101

70 Setting up directory servers 1103
Getting started with directory servers 1103
Integrating external directory servers 1106
Enabling access control by user 1113
Defining user accounts 1119
Add Users to User Group elements 1122
Remove Users from User Group elements 1122
Import user information 1123
Export user information 1123
Change user passwords 1124
Remove the authentication settings of a user 1124
Reset a engine's local user database 1125
Set user database replication on or off for Engines and Master Engines 1125
Examples of Directory Servers 1126

71 Setting up user authentication 1127
Getting started with user authentication 1127
Integrating external authentication services 1130
Define Access rules for authentication 1136
Enable browser-based user authentication 1137
Configure client certificate authentication for browser-based user authentication 1145
Authenticate to the Security Engine 1147
Customize the User Authentication Pages for browser-based user authentication 1148
Monitoring and testing user authentication 1150
Examples of user authentication 1150

Virtual private networks 1155

72 VPNs in Forcepoint Network Security Platform 1157
Types of VPNs in Forcepoint Network Security Platform 1157
Pre-shared key (PSK) authentication in VPNs 1160
Certificate-based authentication in VPNs 1161
Configuring VPNs with external gateway devices 1162
Adaptive Forward Erasure Correction (FEC) for VPN tunnels 1163
Logs related to VPNs 1164
Clustering and VPNs 1164
VPNs and Multi-Link for VPN 1165
VPN Broker 1167

73 Configuring VPNs 1169
VPN configuration overview 1169
Define a custom Gateway Profile element 1173
Defining VPN gateways 1173
Defining Site elements for VPN gateways 1180
Define VPN Traffic Selector elements 1184
Defining VPN profiles 1185
Defining Policy-Based VPN elements 1193
Configuring route-based VPNs 1209
Examples of policy-based VPNs 1217
Examples of route-based VPNs 1219

74 Example VPN configurations 1223
Getting started with example VPN configurations 1223
Example VPN configuration 1: Basic VPN between Security Engines 1224
Example VPN configuration 2: Basic VPN with a partner gateway 1228
Example VPN configuration 3: Basic VPN for remote clients 1236
Example VPN configuration 4: Basic VPN hub 1244

75 Managing VPN certificates 1251
VPN certificates and how they work 1251
Define additional VPN certificate authorities 1254
Create an internal ECDSA certificate authority for VPN gateways 1256
Select the default internal certificate authority 1257
Create a VPN certificate or certificate request for a VPN Gateway element 1257
Import an externally signed VPN gateway certificate 1259
Sign external VPN certificate requests with an internal certificate authority 1260
Select which internal certificate authority signs each certificate 1261
Replacing expired VPN certificates 1262
Export signed VPN gateway certificates or VPN certificate authority certificates 1264
Check when VPN gateway certificates expire 1266
Check when VPN certificate authorities expire 1266

76 Reconfiguring existing VPNs 1269
Changing tunnels in a VPN 1269
Add gateways to an existing VPN 1270
Changing gateway IP addresses in an existing VPN 1270
Give VPN access to more hosts in policy-based VPNs 1271
Route all Internet traffic through policy-based VPNs 1272
Redirect traffic between VPN tunnels 1273
Replace pre-shared keys for VPNs 1274
Adjusting gateway settings for Security Engines in existing VPNs 1275

77 VPN client settings 1279
VPN client settings and how they work 1279
Defining IP addresses for VPN clients 1279

78 Configuring the SSL VPN Portal 1285
Getting started with the SSL VPN Portal 1285
Make services available in the SSL VPN Portal 1286
Allow access to services using the SSL VPN Portal 1286
Define an SSL VPN Portal element 1287
Enable the SSL VPN Portal for an Security Engine 1288

Maintenance and upgrades 1289

79 Configuration of automatic updates and upgrades 1291
Getting started with automatic updates and upgrades 1291
Configure automatic updates and upgrades 1292

80 Backing up and restoring system configurations 1295
Backups and how they work 1295
Back up system configurations 1297
Copy backup files to a storage location 1298

Import backup files into the SMC ... 1298
Restoring backups ... 1299
Restore system configurations after a hardware failure ... 1302
Managing SMC Appliance backups ... 1303

81 Managing log data ... 1307
Log data management and how it works ... 1307
Reducing unnecessary log generation ... 1310
Archive log data ... 1311
Delete log data ... 1312
Export log data ... 1315
View history of completed Log Tasks ... 1316
Overwrite old log or audit entries when log storage is full ... 1317
Examples of log management ... 1317

82 Managing and scheduling Tasks ... 1319
Getting started with Tasks ... 1319
Task configuration overview ... 1320
Task types ... 1321
Creating Task Definitions ... 1322
Schedule Tasks ... 1327
Start Tasks manually ... 1328
Pause the scheduled execution of Tasks ... 1329
Remove Tasks from schedules ... 1329
Stop running Tasks ... 1330

83 Managing licenses ... 1331
Getting started with licenses ... 1331
Generate licenses ... 1333
Upgrading licenses manually ... 1335
Changing license binding details ... 1336
Install licenses for unlicensed components ... 1339
Replacing licenses of previously licensed components ... 1340
Check that all components are licensed ... 1342
Check validity and status of licenses ... 1343

84 Upgrading the SMC ... 1345
Getting started with upgrading the SMC ... 1345
Upgrading the SMC configuration overview ... 1346
Obtain SMC installation files ... 1347
Upgrade SMC servers ... 1348
Default SMC installation directories ... 1349

85 Upgrading Security Engines ... 1351
Getting started with upgrading Security Engines ... 1351
Upgrading the Security Engines configuration overview ... 1352
Obtain and import Security Engine upgrade files ... 1353
Upgrade Security Engines remotely ... 1354

86 Manual dynamic updates ... 1357
Getting started with manual dynamic updates ... 1357
Dynamic update configuration overview ... 1358
Import dynamic update packages ... 1358

Activate dynamic update packages ... 1359

87 SMC Appliance maintenance ... 1361
Getting started with SMC Appliance maintenance ... 1361
Patching and upgrading the SMC Appliance ... 1362
Roll back the SMC Appliance to the previous version on the command line ... 1365

Troubleshooting ... 1367

88 General troubleshooting tips ... 1369
If your problem is not listed ... 1369
Tools for further troubleshooting ... 1369

89 Troubleshooting Administrator accounts and passwords ... 1371
Replace forgotten passwords ... 1371
Troubleshoot user accounts ... 1372
Create an emergency administrator account ... 1372

90 Messages for troubleshooting ... 1375
Alert log messages for troubleshooting ... 1375
Log messages for troubleshooting ... 1377
Error messages for troubleshooting ... 1382

91 Troubleshooting Security Engine operation ... 1385
Troubleshoot Security Engines that do not go or stay online ... 1385
Troubleshoot errors when commanding Security Engines ... 1386
Troubleshoot heartbeat and synchronization errors ... 1387
Troubleshoot contact between Security Engines and the Management Server ... 1387

92 Troubleshooting licenses ... 1389
Problems with licenses ... 1389
Troubleshoot licenses that are shown as retained ... 1389
Troubleshoot licenses that are shown as unassigned ... 1390

93 Troubleshooting logging ... 1393
Troubleshoot the Logs view ... 1393
Troubleshoot log storage ... 1394
Troubleshoot Log Server operation ... 1395

94 Troubleshooting the SMC Client ... 1397
Troubleshoot disabled options in the SMC Client ... 1397
Troubleshoot slow SMC Client startup and use ... 1398
Troubleshoot logging on to the SMC Client ... 1398
Troubleshoot SMC Client layout and views ... 1399
Troubleshoot missing or incomplete statistics ... 1400
Troubleshoot status monitoring ... 1400
Troubleshoot Management Server commands ... 1401

95 Troubleshooting NAT ... 1403
Problems with NAT and possible causes ... 1403
Troubleshoot NAT that is not applied correctly ... 1404
Troubleshoot NAT that is applied when it should not be ... 1405

96 Troubleshooting policies ... 1407

Troubleshoot policy installation ... 1407
Troubleshoot rules ... 1410
Troubleshooting packets incorrectly dropped by antispoofing ... 1413
Troubleshooting unsupported definitions in IPv6 Access rules ... 1414

97 Troubleshooting reporting ... 1415
Error messages for reports ... 1415
Troubleshoot report generation ... 1415
Troubleshoot reports with empty sections or incomplete data ... 1416

98 Troubleshooting upgrades ... 1419
Troubleshoot upgrade failure because of running services ... 1419
Respond to Security Management Center installation failed messages ... 1419

99 Troubleshooting VPNs ... 1421
Find VPN issues ... 1421
Information about IPsec tunnels in logs ... 1422
VPN certificate issues ... 1423
Problems with VPNs with external gateways ... 1423
Forcepoint VPN Client connection issues ... 1424
If traffic is not sent into route-based VPNs ... 1425

Appendices ... 1427

A Command line tools ... 1429
Forcepoint Security Management Center commands ... 1429
Security Engine commands ... 1445
Server Pool Monitoring Agent commands ... 1454

B Default communication ports ... 1457
Forcepoint Security Management Center ports ... 1457
Security Engine ports ... 1460

C Working with expressions ... 1465
Introduction to expressions ... 1465
Using operands ... 1466
Expression processing order ... 1468
Grouping operands using parentheses ... 1468
Nesting expressions ... 1469

D Predefined Aliases ... 1471
Predefined User Aliases ... 1471
System Aliases ... 1471

E Situation Context parameters ... 1473
Correlation Context parameters ... 1473
Other Context parameters ... 1476

F Regular expression syntax ... 1477
SMC regular expression syntax ... 1477
Special character sequences ... 1479
Pattern-matching modifiers ... 1480
Variable expression evaluation ... 1482
Stream operations ... 1484

System variables 1485
Independent subexpressions 1487
Parallel matching groups 1487
Tips for working with regular expressions 1488

G Schema updates for external LDAP servers 1489
Schema updates for external LDAP servers 1489

H Log fields 1491

I Keyboard shortcuts 1493

J Multicasting 1495
Multicasting vs. unicasting or broadcasting 1495
Overview of IP Multicasting 1496
Internet Group Management Protocol and how it works 1497
Ethernet multicasting 1497
Multicasting and Forcepoint Network Security Platforms 1498

Preface

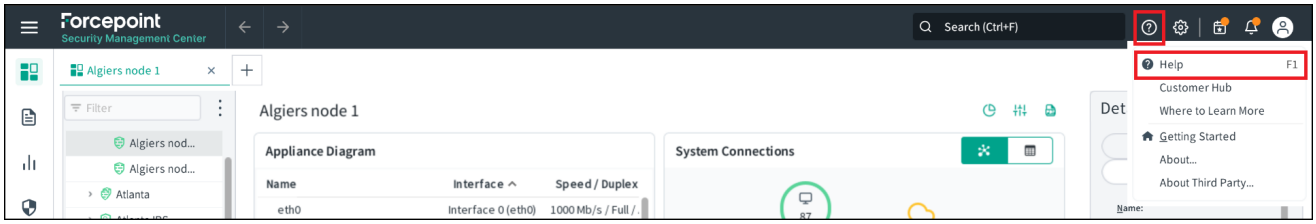
This guide provides the information you need to work with your Forcepoint product.

Open the online help

The SMC Client provides context-sensitive online help.

You can open the online help in the following ways:

- Click the help icon in the toolbar, then select **Online Help**.



- Click the **Help** button in a dialog box.
- Press **F1**.

Find product documentation

In the Forcepoint Customer Hub, you can find information about a released product, including product documentation, technical articles, and more.

You can get additional information and support for your product in the Forcepoint Customer Hub at https://support.forcepoint.com. There, you can access product documentation, release notes, Knowledge Base articles, downloads, cases, and contact information.

You might need to log on to access the Forcepoint Customer Hub. If you do not yet have credentials, create a customer account. To create a customer account, navigate to the Customer Hub Home page, and then click the **Create Account** link.

Links to downloads

Security Engine upgrades and dynamic update packages are available at these websites.

- Security Engine upgrade downloads: https://support.forcepoint.com/s/download
- Dynamic update package downloads: https://autoupdate.ngfw.forcepoint.com

Conventions

The following typographical conventions and icons are used.

Book title, *term*, *emphasis*	Title of a book, chapter, or topic; a new term; emphasis.
Bold	Text that is strongly emphasized.
`User input, code, message`	Commands and other text that the user types; a code sample; a displayed message.
Interface text	Words from the product interface like options, menus, buttons, and dialog boxes.
Hypertext	A link to a topic or to an external website.
	**Note:** Additional information, like an alternate method of accessing an option.
	**Tip:** Suggestions and recommendations.
	**Important/Caution:** Valuable advice to protect your computer system, software installation, network, business, or data.
	**Warning:** Critical advice to prevent bodily harm when using a hardware product.

Part I

Introduction to the Forcepoint Network Security Platform solution

Contents

- The Forcepoint Network Security Platform solution on page 25
- Introduction to Forcepoint Network Security Platform in the Engine/VPN role on page 31
- Introduction to Forcepoint Network Security Platform in the IPS and Layer 2 Engine roles on page 41

Before setting up Forcepoint Network Security Platform, it is useful to know what the different components do and what engine roles are available.

Chapter 1

The Forcepoint Network Security Platform solution

Contents

- Security Management Center (SMC) on page 25
- Components in the Forcepoint Network Security Platform solution on page 26
- Benefits of the SMC on page 28
- Licensing components and features on page 29

The Forcepoint Network Security Platform solution consists of one or more Forcepoint Network Security Platforms and the Forcepoint Security Management Center (SMC). The SMC is the management component of the Forcepoint Network Security Platform solution.

Security Management Center (SMC)

The Forcepoint Security Management Center is the centralized management component of the Forcepoint Network Security Platform solution. The SMC makes the Forcepoint Network Security Platform solution especially well-suited to complex and distributed network environments.

The SMC configures and monitors all components in the Forcepoint Network Security Platform solution. The centralized management system provides a single point of contact for many geographically dispersed administrators.

The unified management platform provides major benefits for organizations of all sizes:

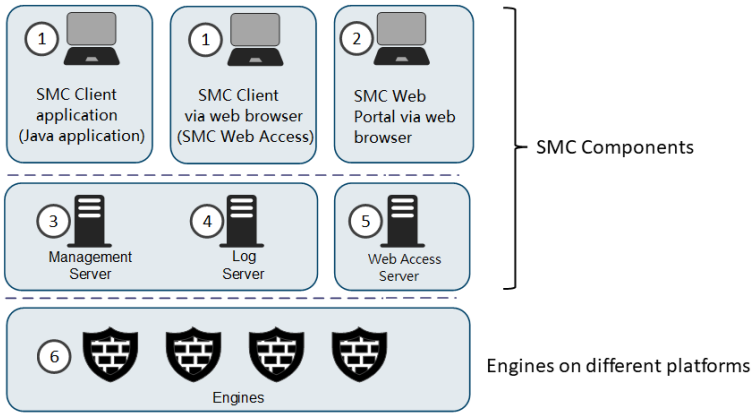
- By allowing automatic coordinated responses when a security threat is detected, interaction between components managed by the same Management Server creates security benefits. Automatic coordinated responses provide instant blocking of unwanted traffic, and reduce the need for immediate human intervention.
- Multiple administrators can log on at the same time to efficiently configure and monitor all Security Engines. The SMC provides a single user interface. This interface allows unified configuration, monitoring, and reporting of the whole Forcepoint Network Security Platform solution with the same tools and within the same user session.
- The reuse of configuration information across components in the system allows you to:
  - Avoid the laborious and error-prone duplicate work of configuring the same details for all components individually.
  - Export and import the configurations between multiple separate systems.

The SMC is designed to manage large installations and geographic distribution. The SMC design creates flexibility and allows scaling up the existing components and adding new types of components to the system without sacrificing its ease-of-use.

Components in the Forcepoint Network Security Platform solution

The Forcepoint Network Security Platform solution includes Security Engines, SMC server components, and SMC user interface components.

Components in the Forcepoint Network Security Platform solution



1. The **SMC Client** is the user interface for the SMC, used for all configuration and monitoring tasks. You can install the SMC Client as an application on your local workstation or access it on Management Server or Web Access Server using a web browser. You can have an unlimited number of SMC Clients.
2. The **SMC Web Portal** is a browser-based, read-only user interface that provides limited access to configuration, logs, and reports on the Management Server. Configuring and using the Web Portal is optional.
3. The **Management Server** is the central component for system administration. A single Management Server can manage different types of Security Engines and Log Servers.
4. The **Log Server** stores traffic logs that can be managed and compiled into reports. In addition, Log Servers also:
   - Correlate events
   - Monitor the status of Security Engines
   - Displays real-time statistics
   - Forward logs to third-party devices

   In smaller installations, the Log Server can run on the same server as the Management Server. For larger environments, it is recommended to install the Log Server on a separate server. There can be multiple Log Server instances.
5. The **Web Access Server** provides additional scalability for administrators using SMC Client via web browser. Configuring and using this component is optional.
6. **Security Engines** process and inspects traffic. A single Security Engine installation can used for multiple purposes, depending on its configuration:
   - Layer 3 Firewall with VPN and SD-WAN capabilities
   - IPS or Layer 2 Firewall with Layer 2 Interfaces
   - IDS with a capture interface

SMC Clients

The SMC Client is the tool for all day-to-day configuration and management tasks, including network interface configuration and remote upgrades.

All commands and configuration changes are relayed through the Management Server, so the SMC Clients never connect to the Security Engines directly. SMC Clients also connect to Log Servers to fetch log entries for administrators to view. Many SMC Clients can be deployed anywhere in the network.

Management Server

The Management Server is the central component for system administration. One Management Server can manage many different types of engines.

The Management Server provides the following types of services:

- Administration and system commands: The Management Server is the central point of all administration tasks (accessed through the SMC Client).
- Configuration database: The Management Server stores all configuration information for Engine/VPN role, IPS, and Layer 2 engines and other system components.
- Monitoring: The Management Server tracks the operating state of the system components and relays this information to the administrators.
- Alert notifications: The Management Server can notify administrators about new alerts in the system, for example, by sending out an email or an SMS text message.
- Certificate authorities (CAs): The Management Server installation includes two basic CAs:
  - An Internal CA that issues all certificates that system components need for system communications.
  - VPN CA that can be used to issue certificates for VPN authentication.

Log Server

Log Servers store traffic logs that can be managed and compiled into reports. Log Servers also correlate events.

Multiple Log Servers can be deployed, which is useful in geographically distributed systems. Log Servers provide the following types of services:

- Log data: Log Servers receive and store logs from other system components and make the data available for viewing and generating reports.
- Statistics and status data: Log Servers receive, relay, and store information about the operation of other system components and keep a record available for generating reports.
- Event correlation: Log Servers detect patterns of events in traffic inspected by multiple Security Engines.

Web Access Server

The Web Access Server is an optional, separate component that provides web-based access.

It enables users to access the UI through their web browsers. To access this interface, users must be standard administrators.

Benefits of the SMC

The SMC has three main benefits: centralized remote management of system components, support for large-scale installations, and server high availability.

Centralized remote management

A centralized point for managing all system components simplifies the system administration significantly.

Ease of administration is central to the SMC. The centralized management system:

- Provides administrators with visibility into the whole network.
- Simplifies and automates system maintenance tasks.
- Reduces the work required to configure the system.

You can also combine information from different sources without having to integrate the components with an external system.

The centralized management system is not an add-on; the system has been designed from the start to be centrally managed.

The main centralized management features in the Security Management Center include the following:

- Sharing configuration data in different configurations eliminates the need for duplicate work, which reduces the complexity of configurations and the amount of work required for changes. For example, an IP address used in the configurations of several different Security Engines has to be changed only one time in one place. It has to be changed only once because it is defined as a reusable element in the system.
- Remote upgrades can be downloaded and pushed automatically to several components. A single remote upgrade operation updates all necessary configuration details on the Security Engines, including operating system patches and updates.
- Fail-safe policy installation with automatic rollback to prevent policies that prevent management connections from being installed.
- The integrated backup feature allows saving all system configurations stored on the Management Server in one manually or automatically run backup.
- Central access point for administrators with centralized access control. The SMC Client requires no separate installation, because it can be made available centrally and be started through a web browser. Several administrators can be logged on at the same time and simultaneously change the system. Conflicting changes are automatically prevented. Administrator rights can be easily adjusted in a highly granular way.

Support for large-scale installations

The Security Management Center is scalable from managing a single Security Engine up to a system consisting of hundreds of components.

Several Log Servers are required in larger systems, but a single Management Server can still effectively manage large installations. Features that make large-scale installations easy to manage include:

- The possibility to separate configurations into isolated Domains.
- To filter configuration definitions in and out of view based on user-defined categories.

High availability

You can optionally install one or more additional Management Servers or Log Servers.

If the active Management Server is damaged, loses power, or becomes otherwise unusable, additional Management Servers allow system control without delays and loss of configuration information. A special Management Server license for multiple Management Servers is required.

Note

The Forcepoint Network Security Platform Security Management Center Appliance (SMC Appliance) does not support high availability for the Management Server or the Log Server.

Log Servers can also be used as backups for each other to allow continued operation when a Log Server goes offline. When a Log Server becomes unavailable, engines can automatically start sending new logs and monitoring data to another pre-selected Log Server. Log Servers do not automatically synchronize their data, but you can set up automatic tasks in the system for backing up important records.

Licensing components and features

License files provide your system a proof of purchase. The Management Server maintains license files.

You receive most licenses as proof-of-license (POL) codes. The proof-of-serial (POS) license code for Forcepoint Network Security Platform appliances is printed on a label attached to the appliances. Using your license code, you can log on to the License Center and view and manage your licenses at: https://stonesoftlicenses.forcepoint.com.

Generally, each SMC server and each Engine, IPS, Layer 2 Engines, and Master Engine node must be separately licensed in your SMC. Virtual Engines do not require their own licenses.

- The SMC components must always be licensed by importing a license file that you create at the Forcepoint website.
- Licenses for Forcepoint Network Security Platform appliances can be generated automatically. You might also need to generate these licenses manually at the Forcepoint website, depending on the appliance model and Management Server connectivity.

The use of some individual features is also limited by license.

All licenses indicate the latest version for which they are valid and are valid on all earlier software versions up to the version indicated. Licenses are by default automatically updated to the newest version possible for the component. If automatic license updates are not possible or disabled, you must generate new licenses manually before upgrading to a new major release.

License upgrades are included in maintenance contracts. If the maintenance contract of a component expires, it is not possible to upgrade the license to any newer version. Evaluation licenses are valid for 30 days.

Chapter 2

Introduction to Forcepoint Network Security Platform in the Engine/VPN role

Contents

- Overview of Forcepoint Network Security Platform in the Engine/VPN role on page 31
- Forcepoint Network Security Platform benefits on page 32

Security Engine in the Engine/VPN role provides access control and VPN connectivity.

Overview of Forcepoint Network Security Platform in the Engine/VPN role

Forcepoint Network Security Platform in the Engine/VPN role provides access control and VPN connectivity.

The term *Engine* refers to the combination of the Forcepoint Network Security Platform software in the Engine/VPN role, and the hardware device or the virtual machine that the software runs on.

Engines have the following representations in the SMC:

- The Single Engine and Engine Cluster elements are containers for the main configuration information directly related to the Engine.
- The individual physical engine devices are shown as one or more Nodes under the main Engine element in some views of the SMC Client.

The Forcepoint Network Security Platform software includes an integrated operating system (a specially hardened version of Linux). There is no need for separate operating system patches or upgrades. All software on the engines is upgraded during the software upgrade.

Forcepoint Network Security Platform benefits

In addition to standard engine features, the Forcepoint Network Security Platform provides additional advanced features.

Support for multi-layer inspection

Multi-layer inspection combines access control, application identification, deep inspection, and file filtering flexibly to optimize security and system performance.

Access control includes packet filtering, connection tracking, URL categories, network application detection, user identification, authentication, and endpoint context information. Forcepoint Network Security Platform in the Engine/VPN role uses state tables to track connections and check whether a packet is a part of an established connection. Forcepoint Network Security Platform in the Engine/VPN role can also act as a packet filter for types of connections that do not require stateful access control. By default, all Engine Access rules implement stateful access control.

Deep inspection checks the actual data being transferred. Deep inspection detects harmful patterns in network traffic. Traffic normalization is used to prevent advanced evasion methods, which are intended to allow harmful traffic to bypass network security devices.

File filtering includes file reputation, anti-malware, and sandbox scans.

Forcepoint Network Security Platform in the Engine/VPN role can apply application level inspection with or without proxying the connections. Protocol Agents provide protocol validation for specific protocols. Protocol Agents are also used to handle protocols that generate complex connection patterns, to redirect traffic to proxy services, and to change data payload if necessary.

Related concepts

Protocol Agents overview on page 1031

Layer 2 interfaces for Forcepoint Network Security Platform in the Engine/VPN role

Layer 2 interfaces on Security Engines in the Engine/VPN role allow the Security Engine to provide the same kind of traffic inspection that is available for Security Engines in the IPS and Layer 2 Engine roles.

Layer 2 interfaces on Security Engines in the Engine/VPN role provide the following benefits:

- When the same Security Engine has both layer 2 and layer 3 interfaces, administration is easier because there are fewer Security Engine elements to manage in the SMC.
- It is more efficient and economical to use one Security Engine hardware device that has both layer 2 and layer 3 interfaces because a smaller number of Security Engine appliances can provide the same traffic inspection.
- When you use layer 2 interfaces on Security Engines in the Engine/VPN role, the Security Engine can use options and features that are not available on Security Engines in the IPS or Layer 2 Engine roles.
  For example, an Security Engine in the Engine/VPN role can use Forcepoint Endpoint Context Agent (ECA), Forcepoint User ID service, NetLinks for communication with the SMC, and dynamic control IP addresses,

while also providing the same kind of traffic inspection that is available for Security Engines in the IPS and Layer 2 Engine roles.

Note

When you use layer 2 interfaces on Security Engines in the Engine/VPN role, follow the same cable connection guidelines as for IPS and Layer 2 Engines.

Related concepts

Cable connection guidelines for IPS and Layer 2 Engines on page 90

Advanced traffic inspection

The Engine's traffic inspection process is designed to ensure a high level of security and throughput. The Engines' policies determine when to use stateful connection tracking, packet filtering, or application-level security.

The Engine uses the resources necessary for application-level security only when the situation demands it, and without unnecessarily slowing or limiting network traffic.

Some types of connections can be selected for inspection of the data content against harmful or otherwise unwanted patterns in connections. The deep packet inspection features provide IPS-type capabilities right on the Engine, and help in finding and stopping malicious or suspicious network activities. You can even inspect the content of encrypted HTTPS connections using the built-in deep packet inspection features.

An anti-malware scanner and a sandbox complement the standard traffic inspection features.

Built-in clustering for load balancing and high availability

The Engine provides innovative built-in clustering and load-balancing features that provide several benefits over traditional solutions.

Traditionally, to achieve high availability on the engine itself, additional hardware switches, software clustering products, or special load-balancing devices have been added and maintained. This often results in the transfer of a *single point of failure* to another network component — typically the network link.

Forcepoint Network Security Platforms have built-in support for clustering, which allows operating up to 16 physical Engine devices as a single unit. All units can actively handle traffic at the same time. No special configuration is required in the surrounding network as the whole implementation is achieved through basic networking standards.

The engines dynamically load-balance individual connections between the cluster nodes, transparently transferring connections to available nodes in case a node becomes overloaded or experiences a failure. The processing of network traffic is automatically balanced between the cluster nodes. This way, the performance of the Engine upgrades by simply adding new nodes to the cluster when necessary. You can also take individual nodes offline during business hours for maintenance purposes. Connections handled by that particular engine are transparently redistributed to other online nodes.

The Forcepoint Network Security Platform also comes with built-in technology for high availability and load balancing between different network connections.

Benefits of clustering

Clustering engine nodes can significantly reduce the risk of problems with availability and maintenance.

A Single Engine can be a single point of failure. This can affect the availability of business critical applications and complicate the maintenance of the engine equipment. Clustering engine nodes can significantly reduce the risk of these problems.

The Forcepoint Network Security Platform solution uses built-in clustering technology. No additional software or hardware is needed to cluster several nodes. If a node itself or the surrounding network equipment malfunctions, the other nodes in the cluster take over the traffic processing, minimizing any disruptions to the traffic flow. Similarly, maintenance is easier with a cluster, because individual nodes can be taken offline and even exchanged for new hardware without causing service outages.

Engine Clusters also balance the load of traffic processing between the engine nodes. You can flexibly add nodes to scale up the Engine Cluster, improving the throughput and performance.

Communication between Engine Cluster nodes

Information between Engine Clustered nodes is synchronized through selected interfaces via a heartbeat network that uses multicast transmissions.

The Engine Cluster nodes exchange information constantly. The state tables that list open connections (state sync) and the operating state of the other nodes (heartbeat) are exchanged. This exchange of information guarantees that all nodes have the same information about the connections. If an engine node becomes unavailable, the other nodes of the cluster immediately notice the change. The exchange of information between clustered Engine nodes is synchronized through selected interfaces via a heartbeat network using multicast transmissions. The heartbeat messages are authenticated, and can also be encrypted if necessary. Authentication is enabled by default.

Load balancing

In load-balanced clustering, traffic is balanced between the nodes dynamically.

In a Engine Cluster configuration, the recommended way to cluster the nodes is load-balanced clustering, where traffic is balanced between the nodes dynamically. Load-balanced clustering provides both fault tolerance and performance benefits.

The traffic arriving at the Engine Cluster is balanced across the nodes according to the settings of the cluster's load-balancing filter. This filtering process distributes packets between the engine nodes and keeps track of packet distribution. The Engine determines the packet ownership of the nodes by comparing the incoming packet with node-specific values based on the packet headers. The load-balancing filter is preconfigured for optimal performance and is not meant to be adjusted independently by the system administrators.

The Engine Cluster keeps track of which node is handling each ongoing connection. As a result, all packets that are part of a given connection can be handled by the same node. Some protocols use multiple connections, which are sometimes handled by different nodes, but this distribution does not usually affect the processing of the traffic.

Standby operation

In standby clustering, only one node at a time processes traffic, and other nodes wait on standby.

Nodes that wait on standby are ready to take over when the currently active node goes offline. Nodes that should not take over automatically can be set offline. The drawback with standby mode is that there is no performance gain in clustering the engines.

Clustering modes for engines

You can configure traffic to be directed to the cluster using several modes.

There are several modes for how traffic can be directed to the cluster. The modes are explained in the following table. If necessary, see the documentation for the router, hub, or switch you are using for information about which mode is best in your environment:

Clustering modes

Mode	Description
Packet dispatch	Packet dispatch is the recommended clustering mode. One node per physical interface is the dispatcher that handles the distribution of traffic between the different nodes for all CVIs on that physical interface. The assigned node handles the traffic processing. No additional switch configuration is needed. This mode can also be used with hubs but it is not the optimal clustering mode with hubs.
Unicast MAC	Unicast MAC is the recommended mode when hubs are used. This mode cannot be used with most switches. All nodes in the cluster share unicast MAC address for the CVI. All nodes in the cluster see all packets
Multicast MAC	The nodes share multicast MAC address for the CVI. All nodes in the cluster see all packets. Do not use this mode instead of the packet dispatch mode except in special cases, for example, if MAC address of the network interface cards cannot be changed.
Multicast MAC with IGMP	The clustering works otherwise the same as in the Multicast MAC mode except that the engine answers to IGMP membership queries. This mode allows limiting multicast flooding when the switch does not support static MAC address forwarding tables.

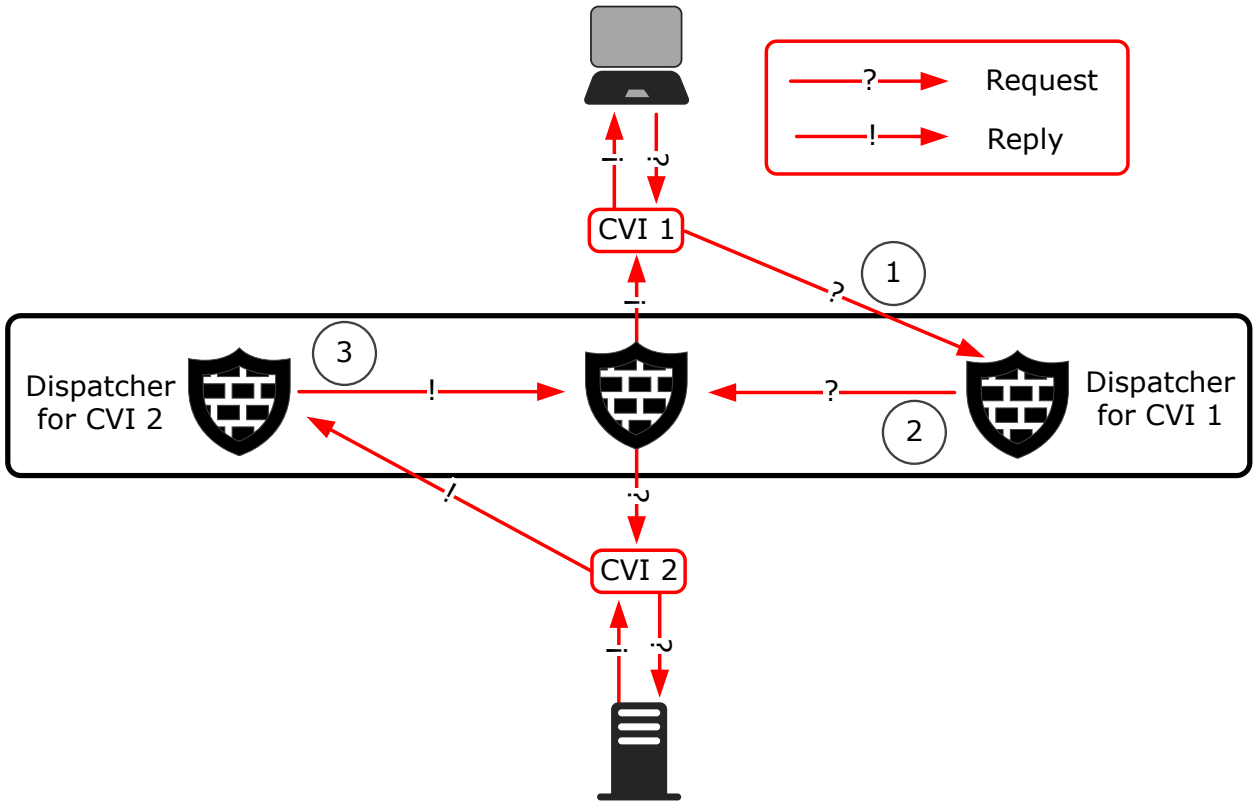
All CVIs on the same physical interface must use the same mode. It is possible to set different cluster modes for CVIs that are defined for different physical interfaces.

Packet Dispatch mode and how it works

In Packet Dispatch mode, one node acts as the dispatcher and assigns packets to the other nodes.

In Packet Dispatch mode, the node selected as the dispatcher on the physical interface assigns the packets to itself or to some other node. The assigned node then handles the actual resource-intensive traffic processing. The dispatcher attempts to balance the nodes' loads evenly, but assigns all packets that belong to the same connection to the same node. The node that acts as the packet dispatcher can be different for CVIs on different physical interfaces. The following illustration shows an example of how packet dispatch handles a connection.

Packet Dispatch CVI Mode



1. The dispatcher node for CVI 1 receives a new packet.
2. The dispatcher node either handles the packet itself or dispatches the packet to one of the other engine nodes for processing according to the load-balancing filter. The packet is sent to the other node through the interface the packet arrived from.
3. The dispatcher node for CVI 2 forwards the replies within the open connection to the same node.

One node is responsible for handling each connection. The node responsible for the connection handles all resource-consuming tasks: it determines if the connection is allowed to continue, translates addresses as necessary, and logs the connection.

The dispatcher node controls the CVI's IP address and MAC address. The other nodes use their own physical interface's MAC address for the same CVI. When the dispatcher node goes offline, one of the other nodes becomes the dispatcher node. The new dispatcher node changes its interface's MAC address to the address defined for the Packet Dispatch CVI.

The network switch must update its address table without significant delay when the packet dispatcher MAC address is moved to another engine node. This operation is a standard network addressing operation where the switch learns that the MAC address is located behind a different switch port. Then, the switch forwards traffic destined to the CVI address to this new packet dispatcher.

VPN Capability

Forcepoint Network Security Platform supports software-defined wide area networks (VPNs).

VPN features include:

- Multi-Link technology

- Clustered Multi-Link VPNs
- Dynamic link selection for Multi-Link VPNs
- Quality of Service (QoS)
- Application routing

Multi-Link technology

Multi-Link provides redundant ISP connections for VPN.

Multi-Link allows you to configure redundant ISP connections using standard network connections, without the need for redundant external routers and switches. You can use any IP-based connection with a dedicated IP address range as part of a Multi-Link configuration. You can also define standby links that are used only when primary links fail.

Traffic is dynamically balanced across the different links based on a performance measurement or based on the links' relative bandwidths. New connections automatically start to use other links when the Engine detects that one of the links fails. The Engine uses NAT to direct the traffic through the different links to make the source IP address valid for the link used.

Standby NetLinks act as backup Internet connections that are only activated if all primary NetLinks fail. Using standby NetLinks provides high availability of Internet connectivity, but is less expensive than having multiple NetLinks active at the same time. Using Multi-Link for load balancing can also help reduce costs. Traffic can be balanced between two or more slower, less expensive, Internet connections instead of one faster connection. Most often, multiple network links are used to guarantee continuity of Internet access, but you can also use Multi-Link to provide redundant links for internal networks.

Multi-Link technology provides highly available network connections for the following scenarios:

- Outbound connections — Multi-Link routing makes sure that outbound traffic always uses the optimal link toward its destination and allow you to configure standby links as backups. The traffic can be distributed across the links in several different ways.
- Inbound connections — The built-in inbound traffic management feature can use Multi-Link to guarantee continuity of the services that your company offers to external users.
- VPN connections — The Multi-Link tunnel selection for VPN traffic is done independently from other types of traffic. Standby links can also be selected independently for a VPN. Connections that use Multi-Link VPN tunnels are transparently moved to other NetLinks even if the NetLink that they are using fails.

Related concepts

Getting started with outbound traffic management on page 731
Using Multi-Link with Server Pools in inbound traffic management on page 750
VPNs and Multi-Link for VPN on page 1165

Clustered Multi-Link VPNs

Forcepoint Network Security Platform in the Engine/VPN role provides fast, secure, and reliable VPN connections. The added benefits of the clustering and Multi-Link technologies provide load balancing and failover for both the VPN gateways and the network connections.

The system's scalability allows you to control how many tunnels are created and used.

The VPN links can be in three different modes:

- Active — When there are multiple links in active mode, traffic is dynamically balanced across the different links based on a performance measurement or based on the relative bandwidths of the links.

- Aggregate — When there are multiple links in aggregate mode, each connection is balanced between all the aggregate links in round robin fashion.
- Standby — Standby links are only used if the active or aggregate links fail.

Related concepts
Types of VPNs in Forcepoint Network Security Platform on page 1157

Dynamic link selection for Multi-Link traffic

Some traffic is affected more easily by changes in the quality of the connection. Forcepoint Network Security Platform in the Engine/VPN role can dynamically select the ISP link that best matches the quality requirements of traffic.

Dynamic link selection has the following benefits:

- Using the connection that best matches the quality requirements of the traffic maximizes the performance of the applications that use the connection.
- Specifying which connection types are preferred, avoided, or not used allows you to use more expensive standby connections only when necessary.
  For example, when all connections are working normally, you can configure business-critical traffic to use one link and all other traffic to use another link.

Dynamic link selection is supported on Security Engines, Master Engines, and Virtual Security Engines in the Engine/VPN role role.

Dynamic link selection is only supported for layer 3 physical interfaces.

Related concepts
Getting started with dynamic link selection on page 783

Quality of Service (QoS) and bandwidth management

Quality of Service (QoS) Policies are interface-specific rules on a Engine that help you ensure that important network services are given priority over less important traffic.

With QoS rules, you can set up a minimum guaranteed bandwidth and maximum bandwidth limit for traffic, and set a priority value for the traffic. You can optionally define settings for Active Queue Management (AQM) to queue and send traffic according to a scheduling algorithm. Sending traffic reduces the volume of dropped or retransmitted packets when there is network congestion.

You can also create DSCP Match/Mark rules that read or write DiffServ Code Point (DSCP) type of service (ToS) field values. Creating DSCP Match/Mark rules allows you to integrate the Engine with other network equipment that implements QoS management in your own or your ISP's network.

Related concepts
Quality of Service (QoS) and how it works on page 973

Application routing

Application routing allows you to apply different NAT rules to traffic, select which VPN traffic uses, and redirect traffic to different proxy servers depending on the network applications detected in the traffic.

Application routing provides the following benefits:

- Many network applications are offered from data centers around the world. Traffic related to specific network applications can be routed to the data center that is geographically closest to the source of the traffic.
- Application routing allows you to optimize the use of limited bandwidth.

Application detection works best on protocols where the client initiates communication. Protocols in which this happens are typically HTTP and HTTPS. To ensure that a Network Application element can be used for matching in application routing, only use elements that have the Application Routing tag. The tag might be added to additional Network Application elements in dynamic update packages.

Some examples of use cases for application routing include the following:

- You can route traffic from specific network applications through the local Internet connection, and route other business traffic to a data center using another connection, such as MPLS.
- You can exclude specific network applications from being redirected to proxies.
- You can direct some network applications to one proxy, and direct the rest of the web traffic to another proxy.
- You can direct all traffic related to a specific network application to one ISP connection, and reserve the other ISP connection for more important traffic. For example, you could direct YouTube traffic to a low-cost ISP connection, and direct business-critical traffic to a faster, but more expensive ISP connection.

Built-in inbound traffic management

The built-in Server Pool feature allows Engines to monitor a pool of alternative servers that offer the same service to the users.

If one of the servers becomes unavailable or overloaded, the Engine automatically redirects new connections to the alternative servers. Server pools can also interact with the Multi-Link feature for high availability of the incoming network connection.

Related concepts

Getting started with inbound traffic management on page 749

Integrating Engine/VPN role with IPS and Layer 2 Engines

You can use Forcepoint Network Security Platform in the Engine/VPN role, IPS, and Layer 2 Engine roles together for traffic inspection.

IP address Block Listing is a shared feature for Security Engine in the Engine/VPN role, IPS, and Layer 2 Engine roles. Block Listing allows blocking harmful traffic not just at the component that detects it, but also on other engines on the connection path.

Instead of using Security Engines in the IPS or Layer 2 Engine role, you can also use layer 2 interfaces on Security Engines in the Engine/VPN role for traffic inspection.

Related concepts

Layer 2 interfaces for Forcepoint Network Security Platform in the Engine/VPN role on page 32

Chapter 3

Introduction to Forcepoint Network Security Platform in the IPS and Layer 2 Engine roles

Contents

- What IPS engines and Layer 2 Engines do on page 41
- How IPS engines and Layer 2 Engines respond to incidents on page 42
- Main benefits of IPS engines and Layer 2 Engines on page 43
- IPS Cluster load balancing on page 44
- Disconnect mode for IPS engines and Layer 2 Engines and how it works on page 44

The Security Engines in the IPS and Layer 2 Engine roles are part of the Forcepoint Network Security Platform solution. The IPS component provides intrusion detection and prevention, and the Layer 2 Engines provide access control and deep inspection of traffic.

What IPS engines and Layer 2 Engines do

An IPS engine or a Layer 2 Engine picks up and examines network traffic in real time. Layer 2 Engines and IPS engines perform event correlation and analysis for traffic they inspect.



Note

Layer 2 Engines are basic engines with a limited set of features. They provide access control and deep inspection of traffic. More advanced engine features such as VPNs and authentication are not supported on Layer 2 Engines.

Layer 2 Engines and IPS engines detect known attacks using attack signatures that are augmented with protocol awareness to form attack *fingerprints*. Protocol awareness decreases the number of false positives compared to simple signatures. Each pattern is applied only to the correct type of traffic. For example, an attack that uses HTTP can be detected when the pattern is seen in HTTP traffic. The HTTP pattern does not falsely match an email message header transported over SMTP.

While fingerprinting accurately detects known attacks, it does not detect attacks that are not yet known. IPS and Layer 2 Engines provide two types of anomaly detection to complement fingerprinting:

- *Protocol analysis* identifies violations in network communications, such as unexpected data, incorrect connection states, and additional or invalid characters. Detecting such violations is useful because many attacks purposely violate standards to trigger abnormal operating responses in vulnerable target systems.

- *Statistical anomaly detection* gathers traffic statistics to detect events such as slow scans and unusual number of connections. This method tracks patterns based on frequency and sequence of events, or the occurrence of sets of related events within a specified time range. For example, many connection attempts from one host to many ports and IP addresses indicates a network scan.

Layer 2 Engines and IPS engines can also initiate immediate responses to any threats that they detect. Depending on how they are installed, engines can also block traffic based on commands that other components send.

How IPS engines and Layer 2 Engines respond to incidents

There are various responses that an IPS engine and a Layer 2 Engine can take when it detects traffic of interest. For example, they can log the connection or actively filter out the traffic.

Several responses are available:

- As the mildest response, an event can be logged. The log entries can be used, for example, for generating statistical reports. Generating statistical reports might be appropriate, for example, for tracking trends in normal network traffic patterns.
- A step up from a log entry is to generate an alert entry that can be escalated to administrators through multiple configurable alert channels. Alert channels include email, mobile phone text messaging (SMS), and SNMP, in addition to being used like log entries.
- Also, logs and alerts can record the full packet headers and data payload for further analysis.



Note

Storing or viewing the packets' payload can be illegal in some jurisdictions due to laws related to the privacy of communications.

- Block Listing makes it possible to block unwanted network traffic for a specified time. IPS engines and Layer 2 Engines can add entries to their own Block Lists based on events in the traffic they inspect. They can also send Block List requests to other Security Engines. Connections that match the Block List are mainly stopped (depending on the enforcing component's policy).

The available responses on an IPS engine or Layer 2 Engine depend on the engine's physical configuration.

Main benefits of IPS engines and Layer 2 Engines

IPS engines and Layer 2 Engines have four key benefits: accuracy, manageability, scalability, and high availability.

Accuracy of IPS engines and Layer 2 Engines

To provide the best possible accuracy, the IPS and Layer 2 Engines provide multiple detection methods that complement each other.

Effective response to network security incidents requires the capability to recognize an enormous number of possible threats. The IPS system must not produce a high number of false alarms that:

- Engage the system administrators in needless investigations.
- Automatically stop legitimate business communications.

Attack signatures are supplemented with protocol-specific matching to produce accurate fingerprints of attacks. The observations on network traffic are not passed on to administrators directly, but instead collected together for further analysis and combined presentation.

What is considered to be a serious threat to a crucial system in one environment might not be considered an event at all in another network. There is more than one set of traffic inspection policies that would work ideally in every environment. So IPS and Layer 2 Engine provides detailed customization possibilities for the entire inspection process. The efficient configuration tools provide default policies that can be edited using drag and drop, while still allowing highly detailed controls for advanced configuration.

With accurate detection results, efforts can be concentrated on countering real threats instead of working on analyzing an endless stream of false alarms.

Manageability of IPS engines and Layer 2 Engines

IPS engines and Layer 2 Engines provide network administrators the tools to save time, reduce mistakes, and get a network overview.

While ease-of-use is one of the main goals for the product, IPS engines and Layer 2 Engines do not achieve it by cutting the available features. The system provides extensive inspection process tuning possibilities, detailed information for monitoring, advanced automation, and tools for complete remote management (including all software upgrades). The distributed architecture allows components to be on separate computers and in different networks. Components can even be in different countries and continents – and still be easily managed as a single system.

An easy-to-use system helps the administrators concentrate on investigating the security threats instead of configuring the security systems.

Scalability and high availability of IPS engines and Layer 2 Engines

Scalability and high availability guarantee that the system can adapt to growing needs, simplify planned maintenance, and protect against hardware failure. IPS engines and Layer 2 Engines can be flexibly scaled up to form clusters of up to 16 devices that work as a single virtual entity.

Clustering IPS engines improves performance and provides high availability for the traffic inspection service.

In Layer 2 Engine Clusters, only one Layer 2 Engine node is active at a time. The other Layer 2 Engine nodes remain in standby mode. If the active Layer 2 Engine node fails, one of the standby nodes automatically starts processing traffic.

IPS Cluster load balancing

In a load-balanced cluster, traffic is dynamically balanced between the nodes.

The recommended way to cluster the nodes in an IPS Cluster is load-balanced clustering, where traffic is balanced between the nodes dynamically. Load-balanced clustering provides both fault tolerance and performance benefits.

When load-balanced clustering is used, the traffic arriving at the IPS Cluster is balanced across the nodes with a load-balancing filter. This filtering process distributes packets between the IPS Cluster nodes and tracks packet distribution. The IPS Cluster determines the packet ownership of the nodes by comparing the incoming packet with node-specific values based on the packet headers.

The IPS Cluster tracks which node is handling each ongoing connection. As a result, the same node can handle all packets that are part of a given connection. Some protocols use multiple connections, which are sometimes handled by different nodes, but this usually does not affect the processing of the traffic.

Disconnect mode for IPS engines and Layer 2 Engines and how it works

IPS engines and Layer 2 Engines support disconnect mode, which enables constant monitoring of link connections and minimizes delays caused by link failures.



Note

Disconnect mode is supported only on modular Forcepoint Network Security Platform appliance models that have full-sized bypass interface modules (not mini modules).

When IPS engines or Layer 2 Engines are deployed in inline mode, link failures cause significant traffic transfer delays if the link failure is undetected. Failure to detect link failures can be prevented in disconnect mode.

If a link fails on one side of an Inline Interfaces pair, the IPS engine or Layer 2 Engine:

- Detects the failure
- Simulates cable disconnection on the other side
- Takes down the other side's link transmitter (TX)

The IPS engine or Layer 2 Engine continues to monitor the receiver (RX) side of a pair of Inline Interfaces. It detects when the link is up again and brings the transmitter (TX) backup accordingly.

Part II
Deployment

Contents

- Deploying the SMC on page 49
- Deploying Forcepoint Network Security Platform in the Engine/VPN role on page 55
- Deploying Forcepoint Network Security Platform in IPS and Layer 2 Engine roles on page 67

Before you can set up the system and start configuring elements, you must consider how the different SMC components should be positioned and deployed.

Chapter 4

Deploying the SMC

Contents

- Overview of SMC deployment on page 49
- Security considerations for SMC deployment on page 50
- Positioning the Management Server on page 51
- Positioning Log Servers on page 52
- Positioning SMC Clients on page 52
- Alternative methods for accessing the SMC Client on page 52
- Example: SMC deployment on page 53
- Post-installation steps for the SMC on page 53

When deploying the SMC, there are some general guidelines for positioning components to guarantee the security of the system.

Overview of SMC deployment

The positioning of SMC components depends on the size and complexity of the network environment.

Supported platforms for SMC deployment

SMC server components can be installed on third-party hardware or they are available as a dedicated Forcepoint Network Security Platform Security Management Center Appliance (SMC Appliance).

Third-party hardware

CAUTION

Do not install the SMC components on Security Engine hardware.

- You can install the SMC on third-party hardware that meets the hardware and operating system requirements. For information about hardware requirements, see the Release Notes.
- You can install all SMC server components on the same computer, or install separate components on different computers.
- In a large or geographically distributed deployment, we recommend installing the Management Server, Log Server, and optional Web Access Server on separate computers.

Physical Forcepoint SMC Appliance

The Management Server and a Log Server are integrated with the hardware operating system as a dedicated server appliance.

Virtual Forcepoint SMC Appliance

The Management Server and a Log Server are integrated with the operating system and provided as `.iso` image for installation to your own hypervisor as a virtual machine.

General SMC deployment guidelines

The basic SMC installation consists of a Management Server, a Log Server, and the SMC Client.

It is possible to run the Management Server and the Log Server on the same computer in low-traffic environments. In larger environments, the components are run on dedicated servers. Several Log Servers might be needed in large or geographically distributed organizations. The SMC Client connects to the Management Server for configuring and monitoring the system and to Log Servers for browsing the log entries.

General guidelines for SMC components in an SMC deployment

Component	General guidelines
Management Server	Position on a central site where it is physically accessible to the administrators responsible for maintaining its operation.
Log Servers	Place the Log Servers centrally and locally on sites as needed based on log data volume and administrative responsibilities.
Web Access Server	The Web Access Server can be deployed in any location that has network access to the Management Server and the Log Servers.
SMC Client	The SMC Client can be used from any location that has network access to the Management Server and the Log Servers.

Security considerations for SMC deployment

The information stored in the SMC Manager is highly valuable to anyone conducting or planning malicious activities in your network. Someone who gains administrator rights to the Management Server can change the configurations.

An attacker can gain access by exploiting operating system weaknesses or other services running on the same computer to gain administrator rights in the operating system.

Important

Secure the Management Server computer. Anyone who has administrator rights to the operating system can potentially view and change any SMC configurations.

Consider at least the following points to secure the Management Server and Log Server:

- Prevent any unauthorized access to the servers. Restrict access to the minimum required both physically and with operating system user accounts.
- We recommend allowing access only to the required ports.
- Never allow SMC Client connections from insecure networks.
- Take all necessary steps to keep the operating system secure and up to date.
- We recommend that you do not run any third-party server software on the same computer with the SMC servers.
- We recommend placing the servers in a separate, secure network segment without third-party servers and limited network access.

You can optionally install the SMC with external certificate management. Using certificates issued by an external CA allows you to use your own established internal CA infrastructure to generate certificates for internal TLS communication between system components. Certificate revocation checking is also supported. If any devices are compromised, the certificates associated with them can be revoked and replaced centrally using the external certificate management system.

Related tasks

Change the type of the internal certificate authority on page 151

Related reference

Forcepoint Security Management Center ports on page 1457
Security Engine ports on page 1460

Positioning the Management Server

The Management Server is positioned on a corporate headquarters or data center central site where it can reach all other SMC Manager components.

The Management Server does not need to be close to administrators. The SMC Clients connect to the Management Server and Log Servers over the network using an encrypted connection.

We recommend using the same SMC to manage all your engines. This unified approach simplifies managing physically distributed network environments and allows closer integration, for example, sending Block List requests from IPS engines to Engines. The configuration information and log data can then be shared and used efficiently together. A single Management Server can manage many components efficiently. You can optionally install one or more additional Management Servers for a high availability setup. Only one Management Server is active at a time. The additional Management Servers function as standby Management Servers.



Note

The SMC Appliance does not support high availability for the Management Server or the Log Server.

The Management Server also handles active alerts and alert escalation to inform the administrators of critical events. In an environment with multiple Management Servers, all active alerts are replicated between the Management Servers.

Positioning Log Servers

Log Servers store engine-generated logs and traffic captures. Several Log Servers can be located both on a central site as well as at remote sites.

The transferred amounts of data can be substantial, so the primary concern for Log Server deployment is the number and throughput of the engine components that send data to the Log Server. A single shared Log Server can be sufficient for a number of remote sites with low traffic volumes, whereas a large office with very high volumes of network traffic might require even several Log Servers for efficient use.

Positioning SMC Clients

The SMC Client provides an interface for managing and monitoring the system.

The SMC Client can be used anywhere for system administration and for browsing logs and alerts. The Security Engines are managed through the Management Server, so the SMC Client never connects directly to the engines.

Alternative methods for accessing the SMC Client

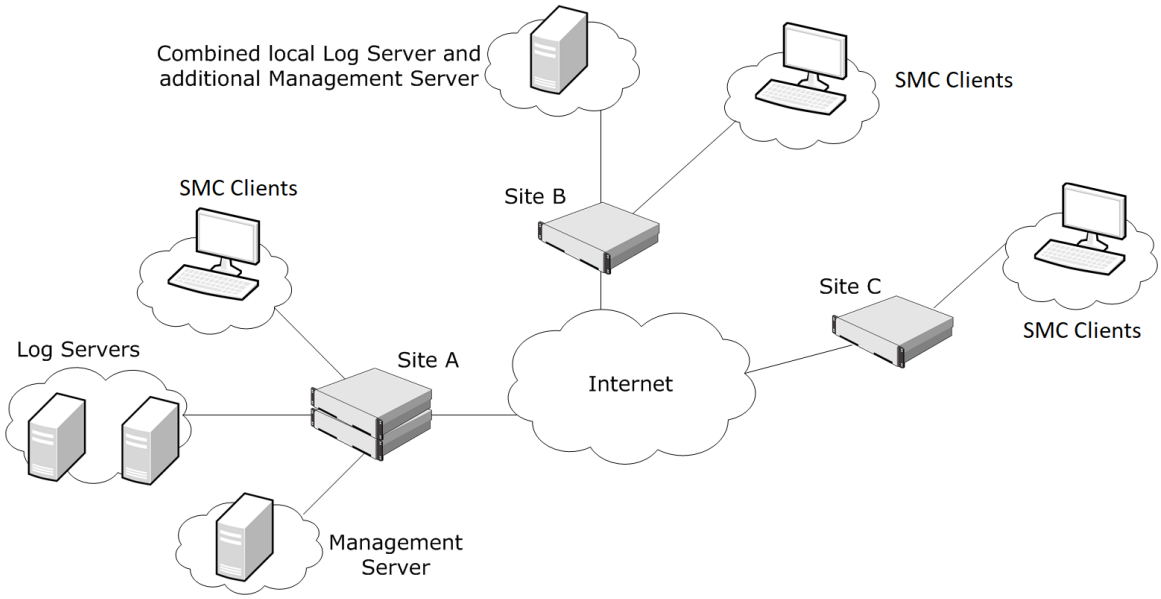
As an alternative to downloading and installing the SMC Client on each workstation that an administrator uses, you can run the SMC Client in a web browser, or you can download the SMC Client from the Management Server.

- Web Access — You can enable the feature on the Management Server or Web Access Server. Administrators log on to the SMC Client on a web page, and the SMC Client runs as an HTML5 application in the web browser. The web browser is the only requirement on the workstation.
- SMC downloads — You can enable the feature on the Management Server. Administrators download and install the SMC Client from the SMC Downloads web page.

Example: SMC deployment

In this example deployment, a company has operations in three different locations. There are some Security Engines and administrators who are responsible for managing the local equipment at each site.

Example of a distributed SMC Manager deployment



Site A is the main site of the company. The active Management Server that manages all local and remote components is at Site A. The main administrators responsible for maintaining the server are also stationed there. There are also two separate Log Servers at Site A. There are a high number of Security Engines at this site, producing a high volume of logs. The Log Servers also work as backup servers for each other.

Site B is a large branch office that is also designated as the disaster recovery site for the main site. The most important services are duplicated. This site has a moderate number of Security Engines. A separate Log Server is installed at Site B to ensure swift log browsing for the local administrators.

Site C is a small branch office that has only a few Security Engines. There is a single local administrator who is an infrequent user of the SMC. There are no SMC components at Site C; the local Security Engines send their data to the Log Servers at Site A.

Post-installation steps for the SMC

After installing the SMC, you must complete some configuration tasks to guarantee the efficient management and security of the system.

The basic administration tasks you must complete after installation include the following:

- Schedule automatic **Backup Tasks** to back up the essential configuration information stored on the Management Server.
- Set up automated tasks to manage the gathered log data and prevent the Log Server storage space from filling up with logs.

We also highly recommend that you set up the following features:

- Define additional administrator accounts and delegating administrative tasks.
- Review settings for automatic updates and making sure the feature works to keep your system current.
- Define custom alerts and alert escalation policies.

To efficiently manage the system, you must also familiarize yourself with monitoring system operation.

Chapter 5

Deploying Forcepoint Network Security Platform in the Engine/VPN role

Contents

- Supported platforms for Security Engine deployment on page 55
- Deploying Security Engines on cloud-based virtualization platforms on page 56
- Running Security Engines as Master Engines on page 57
- Hardware requirements for installing Security Engine on third-party hardware on page 58
- Hardware for Engine Cluster nodes on page 58
- Guidelines for deploying Forcepoint Network Security Platform in the Engine/VPN role on page 59
- Positioning Engines on page 59
- Post-installation steps for Forcepoint Network Security Platform in the Engine/VPN role on page 64
- Cable connection guidelines for Engines on page 65

The positioning of an engine depends on the network environment and the function of the Security Engine.

Supported platforms for Security Engine deployment

You can run Security Engines on various platforms.

The following general types of platforms are available for Security Engines:

- Purpose-built Security Engine appliances

  Note

  For information about supported appliance models, see Knowledge Base article 9743.

- VMware ESX and KVM virtualization platforms
- Microsoft Hyper-V virtualization platform (Security Engine with Layer 3 Interfaces only)
- Microsoft Azure cloud-based virtualization platform (Security Engine with Layer 3 Interfaces only)
- Amazon Web Services (AWS) cloud-based virtualization platform (Security Engine with Layer 3 Interfaces only)
- Third-party hardware that meets the hardware requirements

Note

When Security Engine is running on a virtualization platform or a cloud virtualization platform, Master Engines and Virtual Engines are not supported.

For supported versions of virtualization platforms, see the Release Notes.

The Security Engine software includes an integrated, hardened Linux operating system. The operating system eliminates the need for separate installation, configuration, and patching.

Related concepts

Hardware requirements for installing Security Engine on third-party hardware on page 58
Configuration of Master Engines and Virtual Engines on page 527

Deploying Security Engines on cloud-based virtualization platforms

You can deploy Security Engines on cloud-based virtualization platforms, such as the Amazon Web Services (AWS) cloud and the Microsoft Azure cloud.

Security Engines on cloud-based virtualization platforms provide VPN connectivity, access control, and inspection for services hosted on cloud-based virtualization platforms.

For information about deploying Security Engines in the AWS cloud, see the document *How to deploy Forcepoint Network Security Platform in the Amazon Web Services cloud* and Knowledge Base article 10156.

For information about deploying Security Engines in the Azure cloud, see the document *How to deploy Forcepoint Network Security Platform in the Azure cloud* and Knowledge Base article 14485.

After deployment, you can manage Security Engines on cloud-based virtualization platforms using the SMC Client in the same way as other Security Engines. If you deploy Security Engines that use the scaling feature, you can only preview the Security Engines and make changes to the Engine policies.

Note

Only Single Security Engine with Layer 3 Interfaces are supported. Master Engines and Virtual Engines are not supported.

Licensing

Two licensing models are supported.

- **Bring Your Own License** — You pay only the AWS or Azure standard runtime fee for the Security Engine instance. You must install a license for the Security Engine in the SMC.
- **Hourly** (pay as you go license) — You pay the AWS or Azure standard runtime fee for the Security Engine instance plus an hourly license fee based on the runtime of the Security Engine. No license installation is needed for the Security Engine in the SMC.

For features that require separate licenses, the SMC automatically detects which licensing model the Security Engine uses.

Support for scaling in cloud-based virtualization platforms

When Security Engines are deployed from the Microsoft Azure or AWS cloud environment, additional instances can be created and removed, depending on traffic load.

You deploy the Cloud Auto-Scaled Engines from the cloud environment, and in the SMC Client, the Cloud Auto-Scaled Engine are automatically added to Cloud Auto-Scaled Group elements. You can monitor the Cloud Auto-Scaled Engines on the dashboard, for example.

Limitations

- Cloud Auto-Scaled Engines cannot be edited in the SMC Client.
- The automatic scaling feature is only supported in the Azure cloud. In the AWS cloud, you must add and remove instances manually.

Running Security Engines as Master Engines

There are some hardware requirements and configuration limitations when you use an Security Engine as a Master Engine.

Running the Security Engine as a Master Engine does not require a third-party virtualization platform. When you run Forcepoint Network Security Platform as a Master Engine, the Forcepoint Network Security Platform hardware provides the virtual environment and resources for the hosted Virtual Engines.



Note

You must always install the Forcepoint Network Security Platform software on a hardware device to run the Security Engine as a Master Engine.

You can run Master Engines on the following types of hardware platforms:

- Purpose-built Forcepoint Network Security Platform appliances with 64-bit architecture
- Third-party hardware with 64-bit architecture that meets the hardware requirements

For information about system requirements, see the Release Notes.

The following limitations apply when you use an Security Engine as a Master Engine:

- Each Master Engine must run on a separate 64-bit physical device.
  When Security Engine is running on a virtualization platform or a cloud virtualization platform, Master Engines and Virtual Engines are not supported.
- All Virtual Engines hosted by a Master Engine or Master Engine cluster must have the same role and the same Failure Mode (*fail-open* or *fail-close*).
- Master Engines can allocate VLANs or interfaces to Virtual Engines. If the Failure Mode of the Virtual IPS engines or Virtual Layer 2 Engines is *Normal* (fail-close) and you want to allocate VLANs to several engines, you must use the Master Engine cluster in standby mode.

Related concepts

Hardware requirements for installing Security Engine on third-party hardware on page 58
Configuration of Master Engines and Virtual Engines on page 527

Hardware requirements for installing Security Engine on third-party hardware

There are some basic hardware requirements when you run Security Engine on third-party hardware.

For more information, see the Release Notes.

For information about supported Ethernet interface types and adapters, see Knowledge Base article 9721.

CAUTION

Check that the Automatic Power Management (APM) and Advanced Configuration and Power Interface (ACPI) settings are disabled in BIOS. Otherwise, the engine might not start after installation or can shut down unexpectedly.

CAUTION

The hardware must be dedicated to the Security Engine. No other software can be installed on it.

Hardware for Engine Cluster nodes

You can run different nodes of the same cluster on different types of hardware.

The hardware the cluster nodes run on does not need to be identical. Different types of equipment can be used as long as all nodes have enough network interfaces for your configuration. Engine Clusters can run on a Security Engine appliance, on a standard server with an Intel-compatible processor, or as a virtual machine on a virtualization platform.

If equipment with different performance characteristics is clustered together, the load-balancing technology automatically distributes the load so that lower performance nodes handle less traffic than the higher performance nodes. However, when a node goes offline, the remaining nodes must be able to handle all traffic on their own to ensure high availability. For this reason, it is usually best to cluster nodes with similar performance characteristics.

Guidelines for deploying Forcepoint Network Security Platform in the Engine/VPN role

There are some general deployment guidelines for Engines, Master Engines, and the SMC.

Guidelines for deploying Forcepoint Network Security Platform in the Engine/VPN role

Component	General Guidelines
Management Server	Position on a central site where it is physically accessible to the administrators responsible for maintaining its operation.
Log Servers	Place the Log Servers centrally and locally on sites as needed based on log data volume and administrative responsibilities.
SMC Clients	SMC Clients can be used from any location that has network access to the Management Server and the Log Servers.
Management Server	Position on a central site where it is physically accessible to the administrators responsible for maintaining its operation.
Engines	Position Engines at each location so that all networks are covered. Engines can be clustered. Functionally, the Engine Cluster is equal to a single high-performance Engine. Cluster deployment sets up a heartbeat link between the Engines. The heartbeat link allows the devices to: ■ Track each others' operating status. ■ Agree on the division of work. ■ Exchange information on traffic.
Master Engines	Position Master Engines where Virtual Engines are needed. For example, at a hosting location for MSSP services or between networks that require strict isolation. Master Engines can be clustered. A clustered Master Engine provides scalability and high availability. In a Master Engine Cluster, the Virtual Resource is active in one Master Engine at a time. Cluster deployment sets up a heartbeat link between the Engines. The heartbeat link allows the devices to: ■ Track each others' operating status. ■ Agree on the division of work. ■ Exchange information on traffic.

Positioning Engines

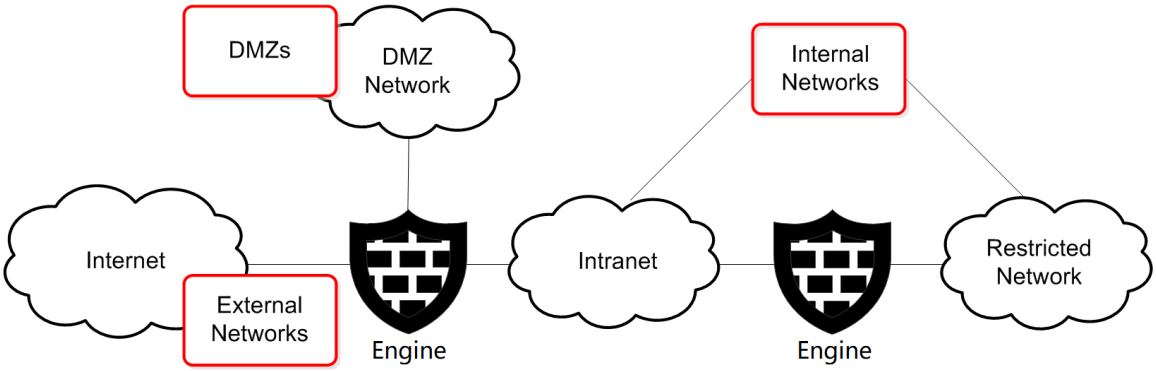
The Engine is a perimeter defense, positioned between networks with different security levels.

Engines generally control traffic between:

- External networks (the Internet) and your internal networks.
- External networks (the Internet) and DMZ (demilitarized zone) networks.
- Between internal networks (including DMZs).

Engines separate the different networks by enforcing rules that control access from one network to another.

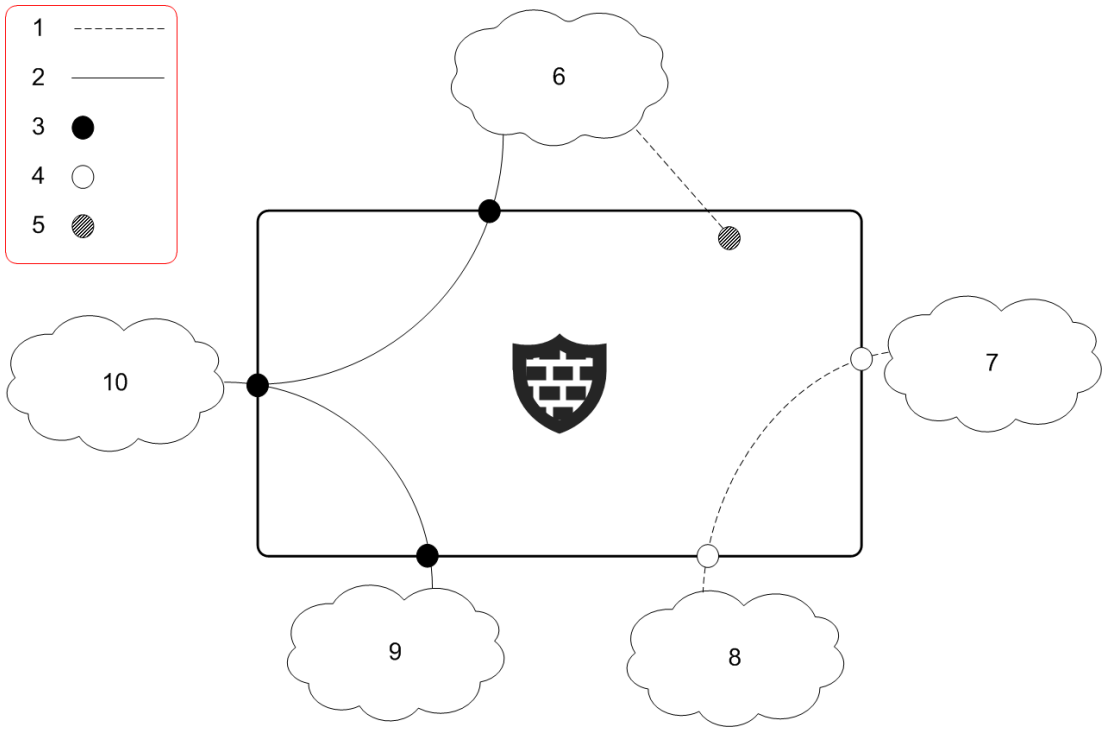
The Engine in different types of network segments



Not all organizations necessarily have all types of networks that are shown here. One Engine can cover all enforcement points simultaneously if it is practical in the network environment and compatible with the organization's security requirements.

In multi-layer deployment, a Engine can have both layer 2 physical interfaces and layer 3 physical interfaces. Layer 2 interfaces on Engines allow the engine to provide the same kind of traffic inspection that is supported on IPS engines and Layer 2 Engines.

The Engine in a multi-layer deployment



1	Traffic inspection only
2	Routed traffic and traffic inspection
3	Layer 3 physical interface These interfaces can route traffic.
4	Layer 2 physical interface of the inline IPS interface or inline Layer 2 Engine interface type These interfaces cannot route traffic. They can only provide traffic inspection.
5	Layer 2 physical interface of the capture interface type These interfaces cannot route traffic. They can only provide traffic inspection.
6	DMZ network
7	Department A internal network
8	Department B internal network
9	Internal network
10	External networks

Using engines to separate internal and external networks

The most common and most important use for a engine is to separate internal networks from the public Internet.

External network considerations for engines

	Description	Implications on Engines
Main purpose	Connectivity between the protected and public networks.	The Engine selects which traffic is permitted into and out of the internal networks and translates addresses between internal IP addresses and public IP addresses. The Engine is typically also a VPN endpoint.
Hosts	Only equipment directly connected to the public network, such as routers and the Engine.	The communicating hosts in external networks are unknown in many cases. IP address spoofing is a possibility. External hosts can be trusted if they are identified using VPN authentication mechanisms.
Users	Access to this network is open, but local access to the hosts is restricted to the administrative staff only.	Internal users are known and trusted. Users in public networks are unknown and untrusted. VPN authentication and encryption can be used to allow specific users access from external networks to internal resources.
Traffic volume	Varies from low to high, generally the full bandwidth of all Internet links combined.	Hardware requirements vary depending on the environment. Clustering allows flexible engine throughput adjustments. Multi-Link allows high availability and load balancing for outbound connections. QoS Policies can control the bandwidth use.
Traffic type	Any type of traffic can be encountered, especially in incoming traffic. Some filtering is done by the Internet service provider.	The Engine controls which traffic is allowed into your networks. It is beyond the Engine's control what and how much traffic it receives from the public networks. Advanced inspection checks can be activated on the Engine and traffic can be redirected to a proxy service, depending on the protocol.
Network security	Little or no access controls to pre-filter traffic arriving from the Internet. Ensure that the hosts in this network are security-hardened and actively patched against known vulnerabilities.	Ensure the Engine's policy is as restrictive as possible. Generally, new connections are not allowed from the external to the internal networks (servers for external services are placed in DMZs). After use, disable SSH access to the Engine's command line from external networks.

Using engines to separate internal networks

Internal networks are mixed environments with servers and end-user computers. Engines restrict traffic between the different internal networks, but traffic within each network is often not secured in any significant way.

Internal network considerations for engines

	Description	Implications on Engines
Main purpose	Network services and connectivity for authorized end users. Back-end servers that serve other networks and user groups.	Internal networks transfer confidential data but can be permissive for the traffic within the network. Engines can control access between different internal networks to enforce different security levels and prevent some types of network threats.
Hosts	Mixed environment consisting of servers, laptops, desktops, network printers, and copiers.	Network communications of the servers and the end-user computers differ in characteristics. Hosts can be actively maintained and patched to reduce some types of risks. Access between networks can be restricted based on the type of host. Engine logs provide a record of network use and alerts can be configured for unusual connection attempts.
Users	Authorized personnel.	Users can be considered trusted, but on various levels. The Engine authenticates users for access between internal networks that have different security levels.
Traffic volume	Varies from low to high. Grows highest at network choke-points in large environments.	Installation at network choke-points often requires high-performance hardware. Clustering can provide load balancing and high availability in critical locations.
Traffic type	Diverse, with many different applications communicating within and in/ out of the network.	The Engine policy must balance users' demands for a wide range of different services with the need to keep the internal networks safe. Advanced inspection features further inspect permitted communications.
Network security	A "trusted network" where the users and the traffic are considered to be authorized.	The Engine establishes boundaries between networks to protect sensitive data and essential services. Availability of network services sometimes overrides security.

Using engines to separate DMZ networks

DMZ networks (demilitarized zone networks, also known as perimeter networks) are isolated environments for servers that offer services mainly for external access.

DMZ considerations for engines

	Description	Implications on Engines
Main purpose	DMZs provide a limited number of services, mostly for external users. The services are often business-critical and open for public access.	The Engine selects which traffic is permitted into and out of the DMZs. The Engine typically also translates IP addresses from public IP addresses that are routable in the external networks to private addresses used in internal networks. VPNs can be used to provide services for partner-type users.

	Description	Implications on Engines
Hosts	A uniform environment consisting mainly of servers that often provide public or semi-public services.	A limited number of services are provided to an often large number of hosts. Some types of administrative access are allowed to a few specific trusted hosts.
Users	Mostly unknown, but some services can be for specific users. Administrators have wider permissions.	Users are often unknown or authenticated by the target servers themselves. Engine authentication can be useful for restricting administrator rights from internal networks.
Traffic volume	Low to medium, generally the full bandwidth of all Internet links combined (shared with other local networks). Traffic to other local networks can be high in volume.	Hardware requirements vary depending on the environment. Clustering allows flexible adjustments to throughput. The inbound traffic management features can balance traffic between redundant servers.
Traffic type	Rather uniform traffic, with only specific applications and servers communicating within, into, and out of the networks.	The Engine controls which traffic is allowed access in and out of each DMZ from external and internal networks. Usually, only a few specific services have to be allowed. Advanced inspection checks can be activated on the Engine and traffic can be redirected to a proxy service, depending on the protocol.
Network security	A network between the trusted and untrusted security zones allowing access for authorized and public use.	External access to services makes the servers in a DMZ a target for attacks. Connections between the DMZ networks and to other internal networks facilitate further attacks, so these connections must be strictly controlled.

Post-installation steps for Forcepoint Network Security Platform in the Engine/VPN role

There are some steps to follow after you have completed the installation, installed a basic policy, and turned the Engines online.

Note

The configuration information is stored on the Management Server. Most changes are transferred to the engines only when you install or refresh the Engine Policy.

The basic administration tasks you must learn or complete next include the following:

- Read and control the operating state of Engines.
- Adjust the automatic tester that monitors the operation of the Engines and the surrounding network.
- Develop your Engine Policies further.

The most typical customization steps include:

- Configure multiple network connections for load-balanced, highly available networking.
- Configure traffic management for incoming connections to groups of servers.

- Set up bandwidth management and traffic prioritization policies.
- Configure the engine to redirect traffic to proxy services.
- Configure secure connectivity between different locations and for traveling users.

Cable connection guidelines for Engines

The cabling of Engines depends on the engine type and the installation.

Make sure that all Ethernet cables are correctly rated (CAT 5e or CAT 6 in gigabit networks).

If you have a two-node Engine Cluster, it is recommended to use a crossover cable without any intermediary devices between the nodes. If you use an external switch between the nodes, follow these guidelines:

- Make sure that portfast is enabled on the external switches.
- Make sure that the speed/duplex settings of the external switches and the Engine devices are set to Auto.
- Configure the external switches to forward multicast traffic.

For layer 2 physical interfaces on Engines, follow these cable connection guidelines:

- Capture interfaces — Follow the cable connection guidelines for IPS and Layer 2 Engines.
- Inline IPS interfaces — Follow the cable connection guidelines for IPS.
- Inline Layer 2 Engine interfaces — Follow the cable connection guidelines for Layer 2 Engines.

Related concepts

Cable connection guidelines for IPS and Layer 2 Engines on page 90

Chapter 6

Deploying Forcepoint Network Security Platform in IPS and Layer 2 Engine roles

Contents

- Supported platforms for Security Engine deployment on page 67
- Running Security Engines as Master Engines on page 68
- Hardware requirements for installing Security Engine on third-party hardware on page 69
- Guidelines for deploying IPS engines and Layer 2 Engines on page 69
- Positioning IPS engines and Layer 2 Engines on page 70
- Deploying IPS engines in IDS or IPS mode on page 73
- IPS deployment examples on page 81
- Post-installation steps for Forcepoint Network Security Platform in the IPS role on page 85
- Deploying Layer 2 Engines in IPS or Passive Engine mode on page 85
- Layer 2 Engine deployment example on page 88
- Post-installation steps for Forcepoint Network Security Platform in the Layer 2 Engine role on page 89
- Cable connection guidelines for IPS and Layer 2 Engines on page 90
- Speed and duplex settings for Security Engines on page 93

The positioning of an IPS engine or Layer 2 Engine depends on the network environment and the function of the IPS engine or Layer 2 Engine.

Supported platforms for Security Engine deployment

You can run Security Engines on various platforms.

The following general types of platforms are available for Security Engines:

- Purpose-built Security Engine appliances

  Note

  For information about supported appliance models, see Knowledge Base article 9743.

- VMware ESX and KVM virtualization platforms
- Microsoft Hyper-V virtualization platform (Security Engine with Layer 3 Interfaces only)
- Microsoft Azure cloud-based virtualization platform (Security Engine with Layer 3 Interfaces only)

- Amazon Web Services (AWS) cloud-based virtualization platform (Security Engine with Layer 3 Interfaces only)
- Third-party hardware that meets the hardware requirements

Note

When Security Engine is running on a virtualization platform or a cloud virtualization platform, Master Engines and Virtual Engines are not supported.

For supported versions of virtualization platforms, see the Release Notes.

The Security Engine software includes an integrated, hardened Linux operating system. The operating system eliminates the need for separate installation, configuration, and patching.

Related concepts

Hardware requirements for installing Security Engine on third-party hardware on page 58
Configuration of Master Engines and Virtual Engines on page 527

Running Security Engines as Master Engines

There are some hardware requirements and configuration limitations when you use an Security Engine as a Master Engine.

Running the Security Engine as a Master Engine does not require a third-party virtualization platform. When you run Forcepoint Network Security Platform as a Master Engine, the Forcepoint Network Security Platform hardware provides the virtual environment and resources for the hosted Virtual Engines.

Note

You must always install the Forcepoint Network Security Platform software on a hardware device to run the Security Engine as a Master Engine.

You can run Master Engines on the following types of hardware platforms:

- Purpose-built Forcepoint Network Security Platform appliances with 64-bit architecture
- Third-party hardware with 64-bit architecture that meets the hardware requirements

For information about system requirements, see the Release Notes.

The following limitations apply when you use an Security Engine as a Master Engine:

- Each Master Engine must run on a separate 64-bit physical device.
  When Security Engine is running on a virtualization platform or a cloud virtualization platform, Master Engines and Virtual Engines are not supported.
- All Virtual Engines hosted by a Master Engine or Master Engine cluster must have the same role and the same Failure Mode (*fail-open* or *fail-close*).
- Master Engines can allocate VLANs or interfaces to Virtual Engines. If the Failure Mode of the Virtual IPS engines or Virtual Layer 2 Engines is *Normal* (fail-close) and you want to allocate VLANs to several engines, you must use the Master Engine cluster in standby mode.

Related concepts
Hardware requirements for installing Security Engine on third-party hardware on page 58
Configuration of Master Engines and Virtual Engines on page 527

Hardware requirements for installing Security Engine on third-party hardware

There are some basic hardware requirements when you run Security Engine on third-party hardware.

For more information, see the Release Notes.

For information about supported Ethernet interface types and adapters, see Knowledge Base article 9721.

CAUTION

Check that the Automatic Power Management (APM) and Advanced Configuration and Power Interface (ACPI) settings are disabled in BIOS. Otherwise, the engine might not start after installation or can shut down unexpectedly.

CAUTION

The hardware must be dedicated to the Security Engine. No other software can be installed on it.

Guidelines for deploying IPS engines and Layer 2 Engines

There are some general deployment guidelines for IPS engines, Layer 2 Engines, and the SMC Manager.

Naturally, there are valid reasons to make exceptions to these general rules depending on the actual network environment.

General guidelines for IPS and Layer 2 Engine deployment

Component	General Guidelines
Management Server	Position on a central site where it is physically accessible to the administrators responsible for maintaining its operation.
Log Servers	Place the Log Servers centrally and locally on sites as needed based on log data volume and administrative responsibilities.
SMC Clients	SMC Clients can be used from any location that has network access to the Management Server and the Log Servers.

Component	General Guidelines
IPS engines	Position IPS engines at each location so that traffic in all appropriate networks can be inspected. IPS engines can be clustered. Functionally, the IPS Cluster is equal to a single high-performance IPS engine. Cluster deployments set up heartbeat links between the IPS engines. The heartbeat links allow the devices to track each others' operating status and agree on the division of work.
Layer 2 Engines	Position Layer 2 Engines at each location so that traffic in all appropriate networks can be inspected. Layer 2 Engines can be clustered for high availability. Only one Layer 2 Engine node in the Layer 2 Engine Cluster is active at a time. If the active Layer 2 Engine node goes offline, another Layer 2 Engine node automatically starts processing traffic.
Master Engines	Position the Master Engines where Virtual Engines are needed. For example, at a hosting location for MSSP services or between networks that require strict isolation. Master Engines can be clustered. A clustered Master Engine provides scalability and high availability. In a Master Engine Cluster, the Virtual Resource is active in one Master Security Engine at a time. Cluster deployments set up heartbeat links between the engines. The heartbeat links allow the devices to track each others' operating status, agree on the division of work, and exchange information on traffic.

Positioning IPS engines and Layer 2 Engines

IPS and Layer 2 Engines pick up passing network traffic for inspection in real time. The positioning of the engines is the most critical part of the deployment.

Each engine can inspect the network traffic of one or more network segments in IDS and IPS configurations.

The following table describes the modes for IPS engines and Layer 2 Engines.

Modes for IPS engines and Layer 2 Engines

Role	Default Policy	Mode	Description
IPS	Allows everything that is not explicitly denied in the policy.	Inline	In inline (IPS) mode, an IPS engine actively filters traffic. The IPS engine is connected as a "smart cable" between two network devices, such as routers and a switch. The IPS engine itself does not route traffic: packets enter through one port, are inspected, and exit through the other port that makes up the pair of Inline Interfaces. Failover network interface cards (NICs) are recommended on the IPS engine to allow network connectivity when the IPS engine is offline. An inline IPS engine can also transparently segment networks and control network access.
		Capture	In capture (IDS) mode, an IPS engine listens to network traffic that is replicated to the IPS engine through: ■ Port mirroring (switch SPAN ports) ■ Dedicated network TAP devices

Role	Default Policy	Mode	Description
Layer 2 Engine	Denies everything that is not explicitly allowed in the policy.	Inline	In inline (IPS) mode, a Layer 2 Engine actively filters traffic. The engine is connected as a “smart cable” between two network devices, such as routers. The engine itself does not route traffic: packets enter through one port, are inspected, and exit through the other port that makes up the pair of Inline Interfaces. Fail-open network interface cards (NICs) can only be used on the Layer 2 Engine if the Failure Mode of the pair of Inline Interfaces is Normal. An inline Layer 2 Engine can also transparently segment networks and control network access.
		Capture (Passive Engine)	In capture (Passive Engine) mode, a Layer 2 Engine listens to network traffic that is replicated to the Layer 2 Engine through port mirroring (switch SPAN ports).
		Passive Inline	A Layer 2 Engine installs inline between two network devices, such as routers and a switch, but does not filter traffic. An inline Layer 2 Engine can be set to Passive Engine mode by configuring the Layer 2 Engine to only log connections.

The same IPS engine can be used for both IPS and IDS operation simultaneously. For example, an IPS engine can be deployed inline to examine traffic from one network to another and capture traffic that stays within each network.

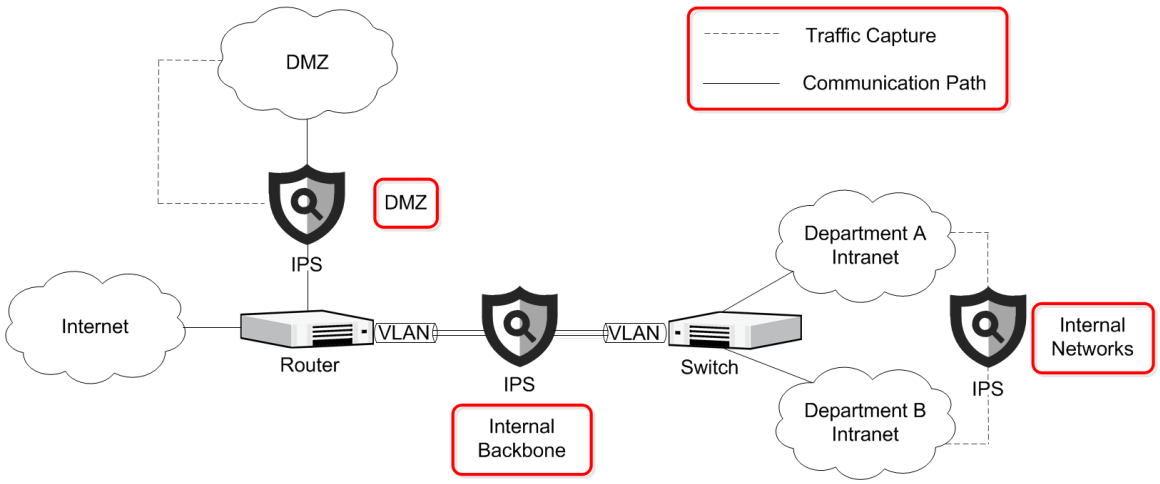
Take the following into consideration when you decide where to install the engines:

- The critical assets to be protected and the potential attack paths.
- The most suitable locations along the attack path for detecting and responding to attack attempts to protect the assets.
- The volume and profile of traffic to be inspected at each location.

Select the engine role based on the way the engine handles inspected traffic:

- Use a Layer 2 Engine if traffic must be denied unless it is explicitly allowed.
- Use an IPS engine if traffic must be allowed unless it is denied.

Example of positioning Security Engines in different network segments



The illustration outlines common deployment scenarios for IPS engines in general internal networks and in DMZ networks. Layer 2 Engines can be used in similar scenarios. IPS engines and Layer 2 Engines are not

necessarily needed at each of these points in all environments. A single IPS engine or a single Layer 2 Engine can also cover several or even all scenarios simultaneously if the physical setup makes it practical.

Positioning IPS engines and Layer 2 Engines in internal networks

In internal networks, access is permissive for purely internal communications, but there are strict controls at the perimeter engine that separates the internal network from public networks.

Inbound traffic from public networks to internal networks is forbidden with few exceptions.

Internal network considerations for IPS engines and Layer 2 Engines

	Description	Considerations for IPS engines and Layer 2 Engines
Main purpose	Network services and connectivity for authorized users. Back-end servers that serve other networks and user groups.	IPS engines and Layer 2 Engines can be used within internal networks and for strengthening the perimeter defense with additional layers of inspection.
Hosts	Mixed environment consisting of servers, laptops, desktops, network printers, and copiers.	IPS engines and Layer 2 Engines can control access between internal hosts uncontrolled by other devices. Connections between internal network zones are of particular interest for inspection.
Users	Authorized personnel. Access in and out of the network controlled by a Engine.	End-user-controlled devices can be distinguished from other hosts to create more accurate and fine-grained rules.
Traffic volume	Varies from low to high. Grows highest at network choke-points in large environments.	Installation at network choke-points where traffic levels are high requires high-performance hardware. Clustering and load balancing can be applied to increase performance and provide high availability in critical locations.
Traffic type	Diverse with many different applications communicating within and in/ out of the network.	A wide range of permitted applications means that the policy has a wide scope. Access control and inspection can be fine-tuned based on the security levels of the different network segments or zones. TLS inspection can be activated to inspect SSL/TLS encrypted traffic. The IPS engines and Layer 2 Engines can also detect and control Application use.
Network security	A “trusted network” where the users and the traffic are considered to be authorized.	The primary line of defense is at the perimeter. It is possible that authorized users in the trusted network become willingly or accidentally involved in a security incident.

Positioning IPS engines and Layer 2 Engines in DMZ networks

DMZ networks (demilitarized zone networks, also known as perimeter networks) allow inbound access to a wide range of users, but are unified environments in terms of devices.

The services offered are limited in number as well and their allowed usage is often strictly defined.

DMZ considerations for IPS engines

	Description	Considerations for IPS engines
Main purpose	DMZs provide a limited number of services for external users. The services are often business-critical and open for public access.	DMZs are a tempting target for attacks because of their accessibility, importance, and visibility. IPS engines provide crucial protection in DMZs, unless the DMZs are already protected by engines.
Hosts	Often a uniform environment consisting mainly of servers. No outbound communication is initiated from the DMZ to the public networks.	Most sources are not trusted and IP address spoofing is a possibility. Internal networks can be considered more trustworthy if there is a Engine that prevents IP address spoofing.
Users	Most services are public, but some services might also be offered to specific users. Administrators have wider permissions.	For recognized users, allowed and forbidden activities can be specified in great detail for each type of access.
Traffic volume	Low to medium, generally the full bandwidth of all Internet links combined (shared with other local networks). Traffic to other local networks can be high in volume.	Hardware requirements vary greatly depending on the environment. Clustering allows flexible adjustments to the inspection performance.
Traffic type	Rather uniform traffic, with only well-known applications and servers communicating within and into the networks.	The limited, well-defined set of protocols and applications means inspection can be tuned in great detail. If servers provide HTTPS services, decrypting the traffic for inspection might require heavy processing.
Network security	A network between the trusted and untrusted security zones allowing access for authorized and public use.	External access to services makes the servers in a DMZ a tempting target for attacks. Connections between the DMZs and other networks facilitate further attacks.

Deploying IPS engines in IDS or IPS mode

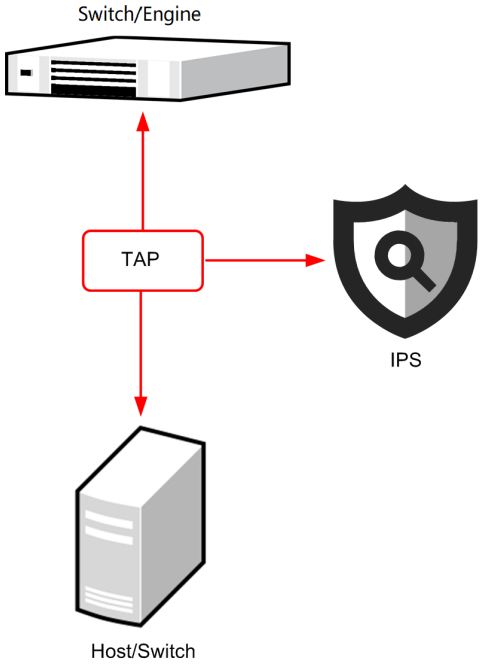
IPS engines can be configured in IPS mode or IDS mode.

IPS deployment in IDS mode

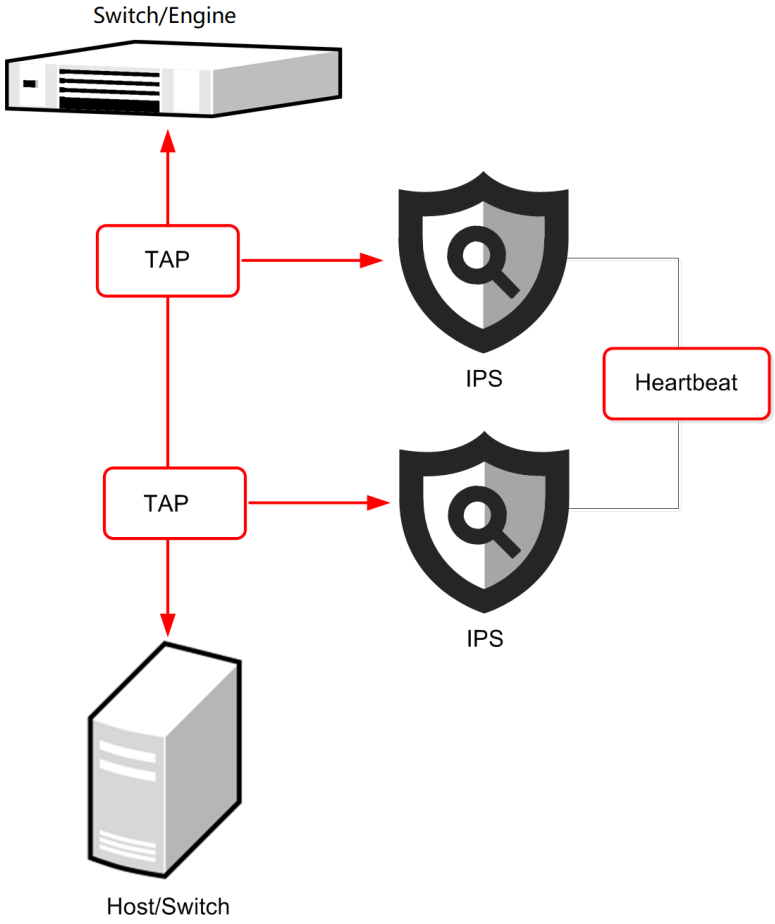
One of the options in IDS mode is to use network TAP devices that copy packets for the IPS engines.

In an IPS Cluster, all nodes must receive all packets. The nodes agree over the heartbeat link which node inspects which connections.

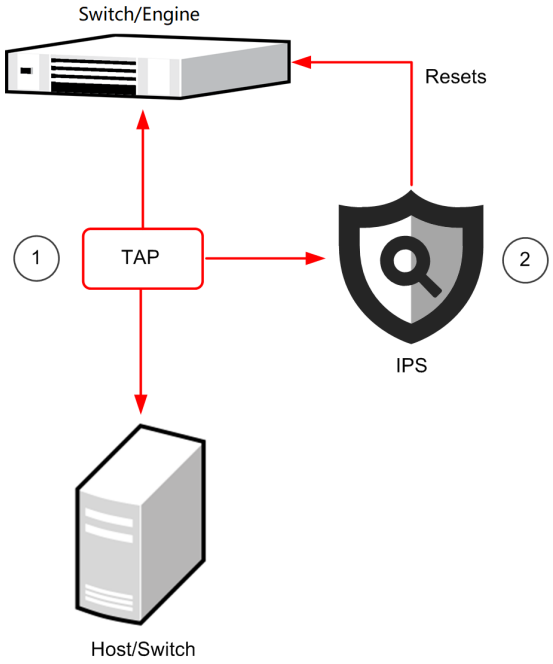
Single IPS in IDS mode with a network TAP



IPS Cluster in IDS mode with network TAPs

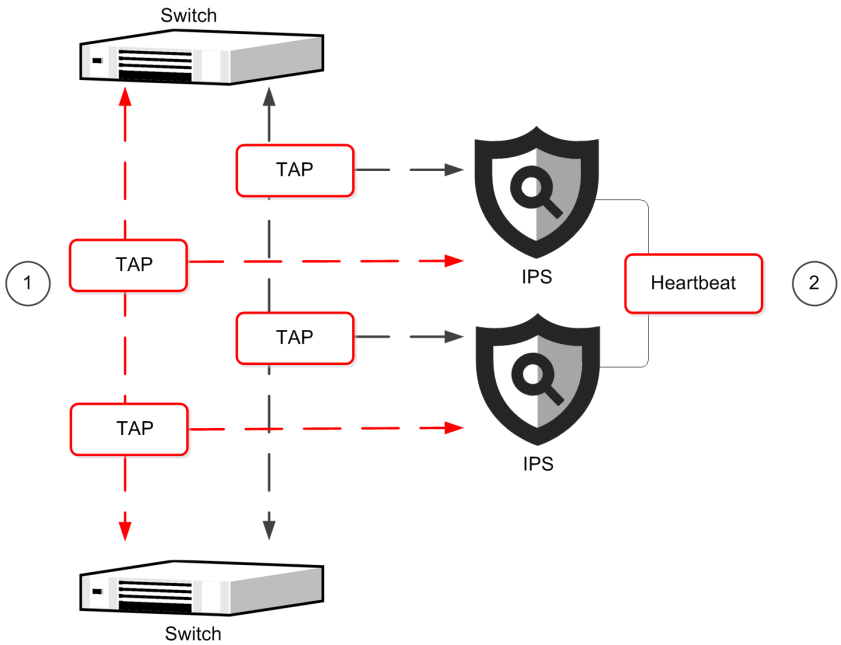


Single IPS in IDS mode with a network TAP and an interface for sending resets



1. A pattern in captured traffic triggers the reset.
2. IPS sends a reset within the same broadcast domain to each communicating host posing as the other host by using its IP address and MAC address.

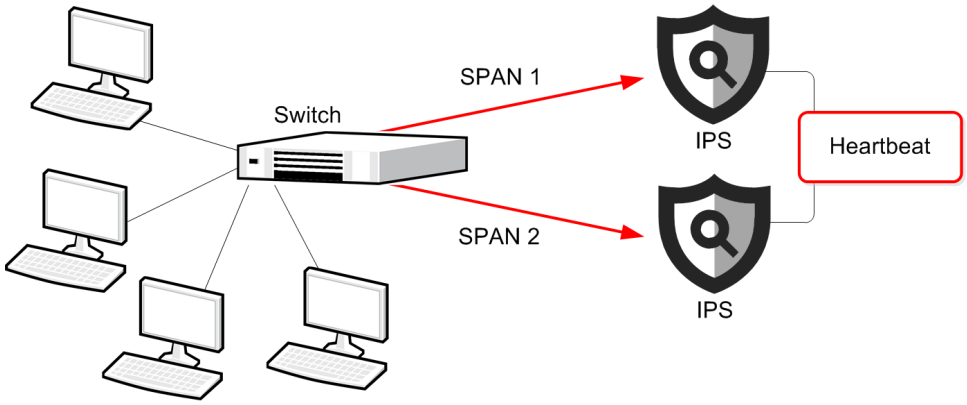
IPS Cluster in IDS mode with network TAPs on a redundant link



1. Switches balance traffic across redundant links.
2. Links are combined into a *Logical Interface* to inspect whole connections.

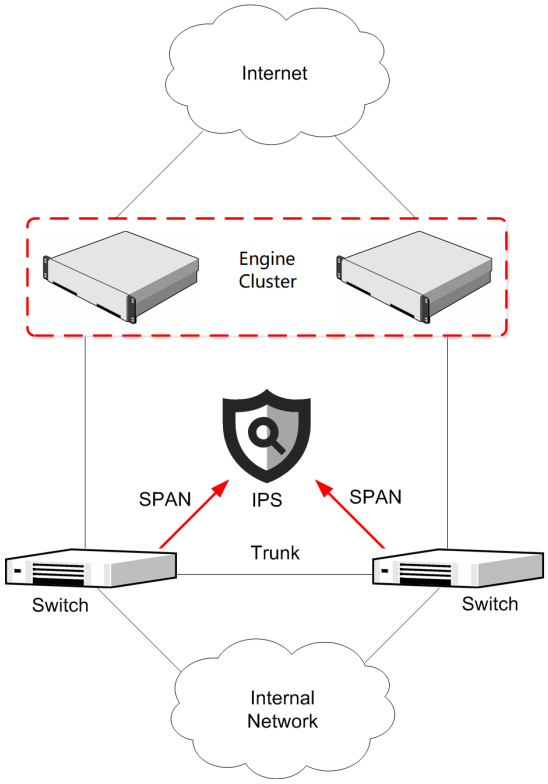
Packets can also be duplicated for inspection through a SPAN or mirror port on a switch/router. In an IPS Cluster, each node must be connected to a SPAN or mirror port of its own. Hubs are not recommended, but you can use hubs in configurations where the low performance of a hub is not an issue. For example, in a basic testing environment.

IPS Cluster in IDS mode with SPAN/mirror ports



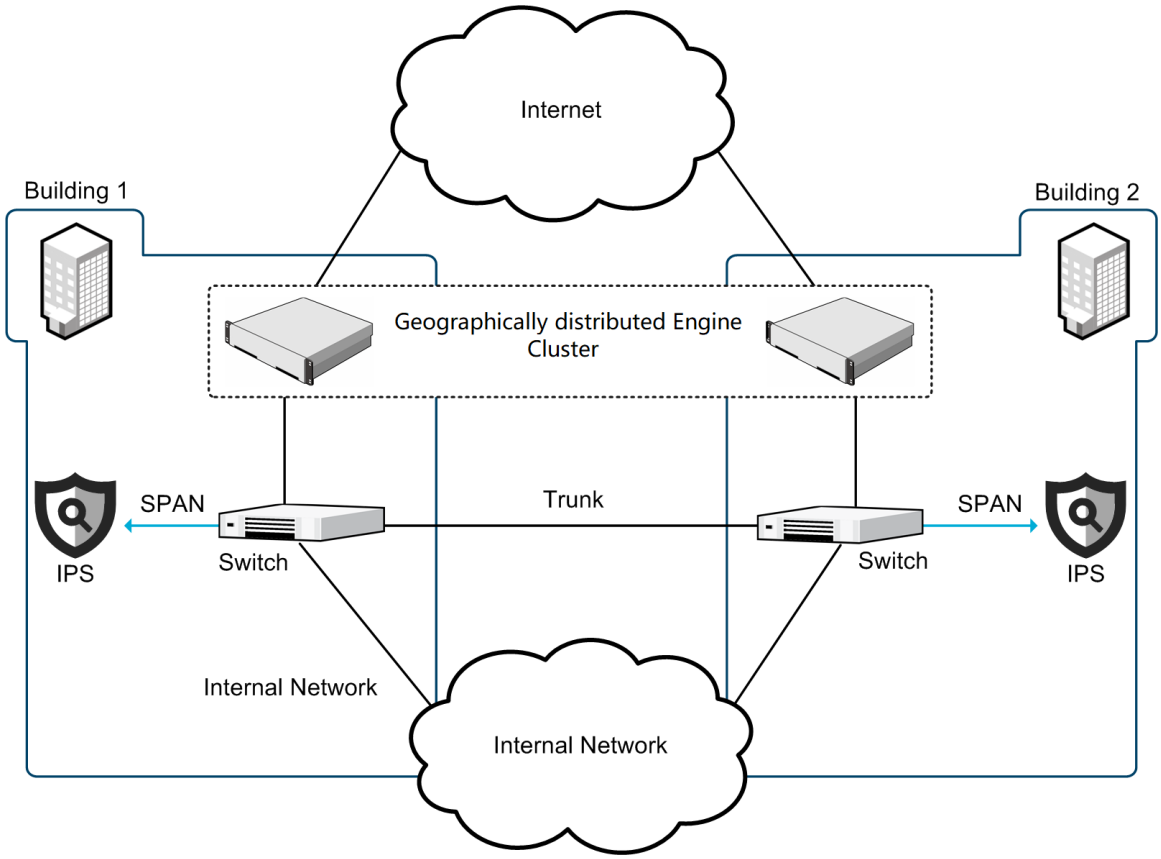
An IPS Cluster can be deployed alongside a Engine Cluster. In this configuration, the IPS Cluster is in the same broadcast domain as the Engine.

IPS connected to SPAN ports alongside redundant switches



In a redundant disaster-recovery setup, Engine Cluster nodes can be far apart. The IPS engines are not clustered in this configuration, but they have identical policies.

Single IPS engines in a distributed disaster-recovery environment



IPS deployment in IPS mode

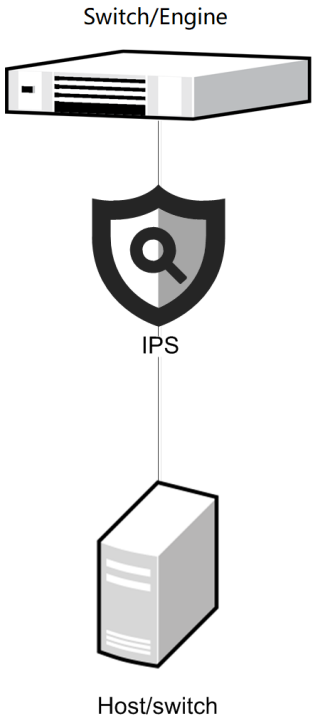
In an inline IPS configuration, the IPS engines are installed directly in the traffic path.

Fail-open network cards are recommended to allow traffic flow when the IPS engines are offline.

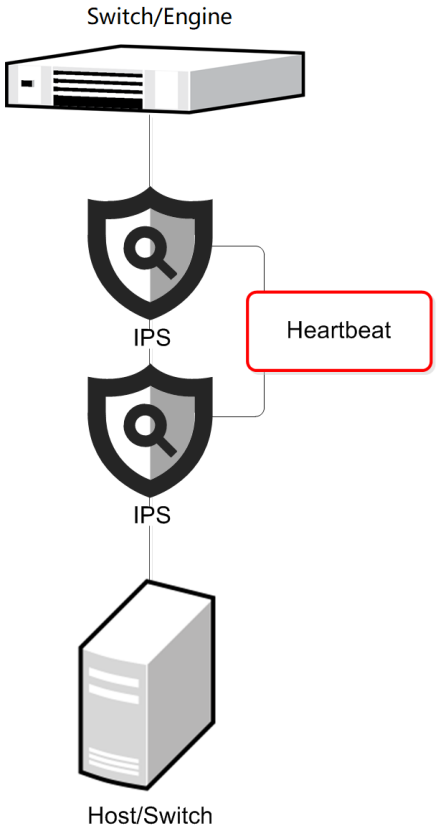
CAUTION

Always use standard cabling methods with an inline IPS engine. Use crossover cables to connect the appliance to hosts and straight cables to connect the appliance to switches.

Single inline IPS engine

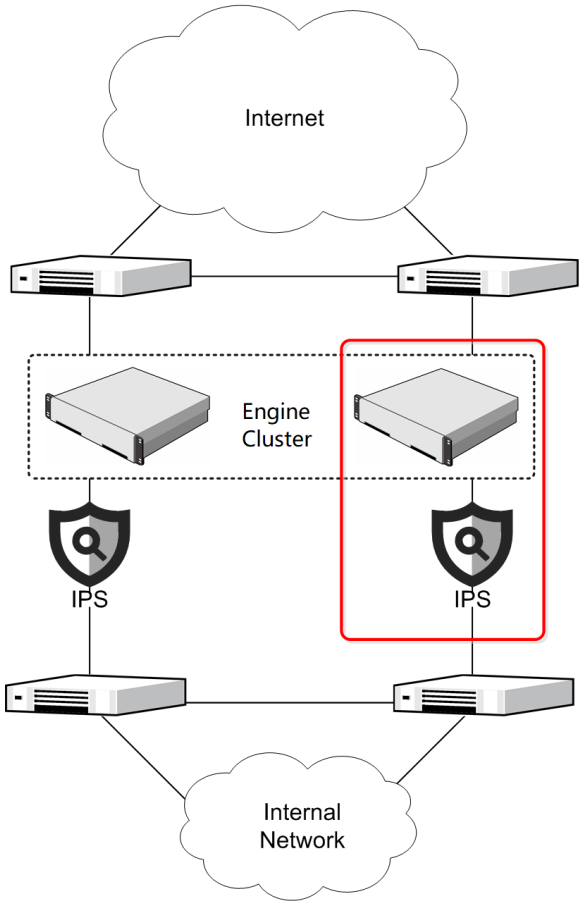


Serial IPS Cluster



The same node handles the packets within a connection.

Redundant single inline IPS engines alongside a Engine Cluster



IPS engines are connected alongside each individual Engine. The IPS engines have the same policy, but they are not clustered.

Note

In this deployment scenario, the Medium-Security Inspection Policy must be used on the IPS engines.

Related concepts

Cable connection guidelines for IPS and Layer 2 Engines on page 90

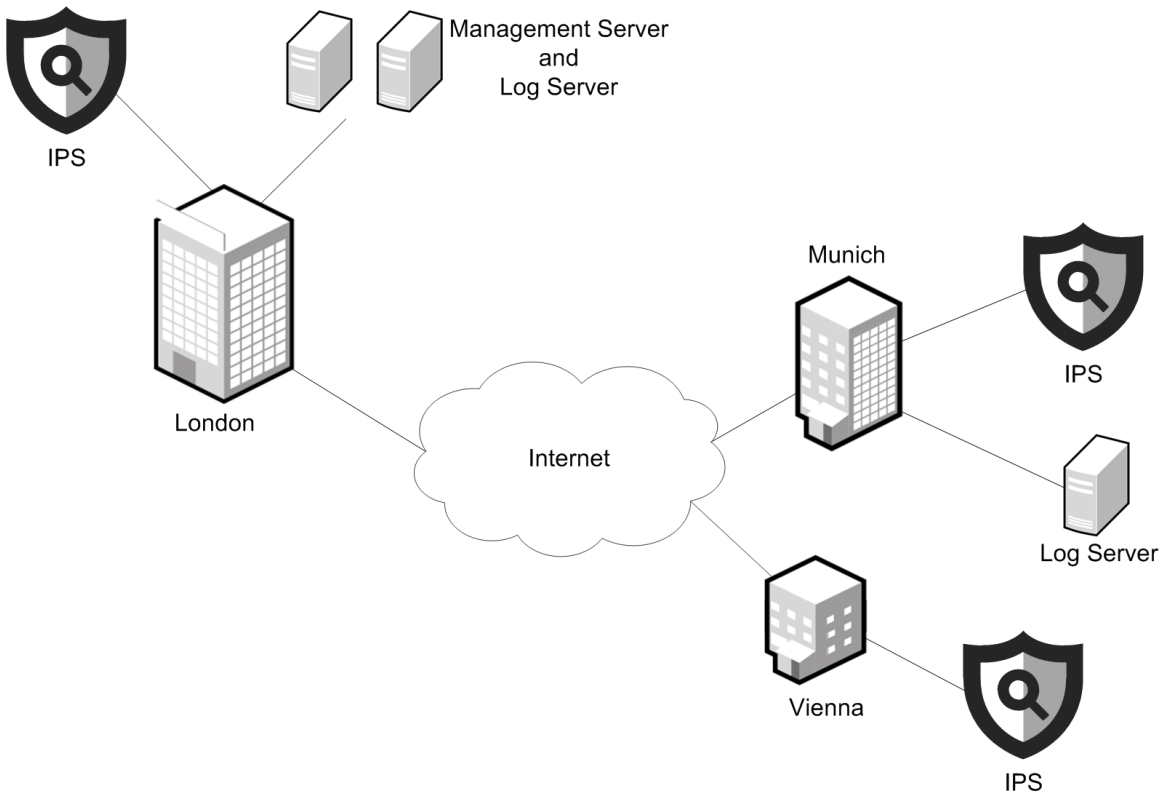
IPS deployment examples

These examples show how to deploy Forcepoint Network Security Platform in the IPS role in an organization.

The scenario presented here is not meant to be representative of a typical installation. The main focus here is to highlight some of the criteria that you can use when planning your deployment. The example covers considerations that affect most installations, but is not an exhaustive list of the factors you might need to consider. The IPS system could be deployed in alternative ways even in this example scenario, depending on issues that are not covered here, such as the physical layout of the individual local networks, the hardware available, and budget constraints.

This example explains the IPS deployment at a company that has three offices: headquarters in London, a branch office in Munich and a small satellite office in Vienna.

The example company's networks



All offices have IPS components. There are also SMC components at the two larger sites. The example company has some critical assets to protect and some of the networks experience a heavy traffic load. The example company has decided on a high availability solution for most locations and acquired the following components:

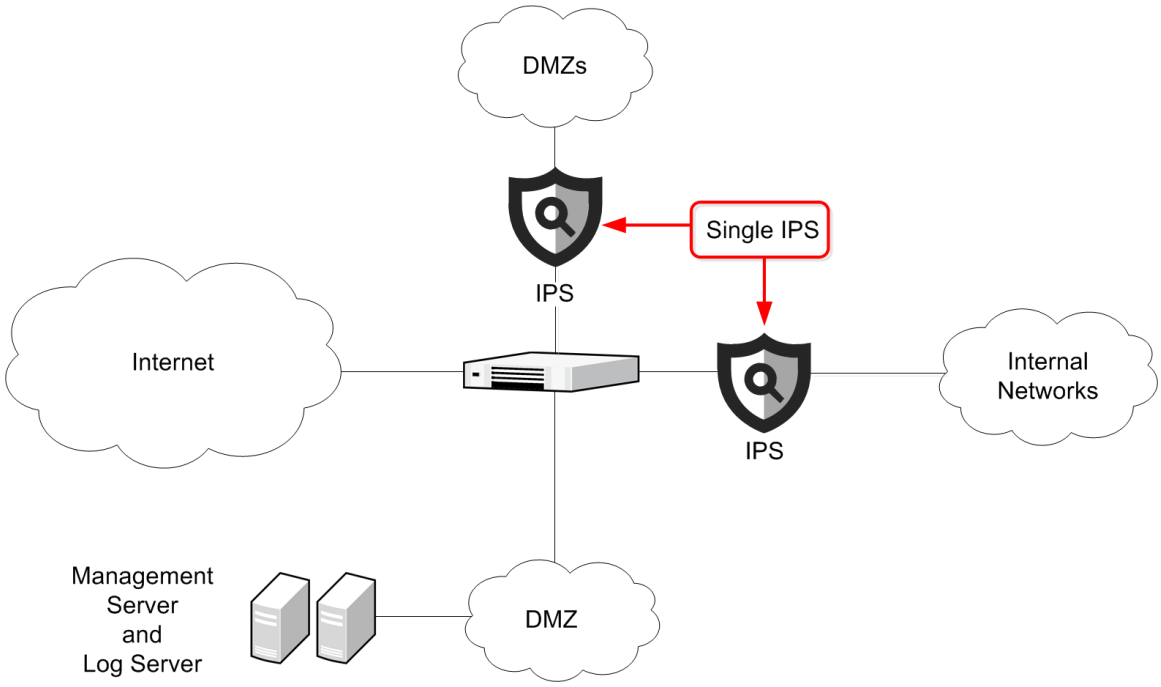
- Three IPS engines
- One Management Server
- Two Log Servers

Example: large-scale IPS installation

The following is an example of a large-scale installation with two Single IPS engines.

The example company's main office at London has many end users and servers. The servers host nearly all company external services and receive a high volume of traffic. The large end-user base generates a high volume of network traffic as well. There are many different applications and protocols in use, resulting in a diverse traffic pattern. The most important asset that the company wants to protect at its headquarters are the web servers hosting the company's online store. The main system administrators work at the main office site.

Headquarters in London



In this case, the company has made the following decisions:

- Because most of the administrators are at this site, the Management Server that controls the whole distributed system is located here.
- There are many administrators and components, so there is also a Log Server here.
- Several DMZs for different services handle a high total volume of traffic. Part of the traffic is encrypted HTTPS, which uses significant processing power to decrypt for inspection. As the overall load is heavy, the company decided to protect the DMZs using a dedicated high-performance Forcepoint Network Security Platform appliance.
- A separate single IPS is installed to protect the diverse high-volume communications of the internal networks.
- The Management Server and the Log Server are placed in a dedicated DMZ for security.

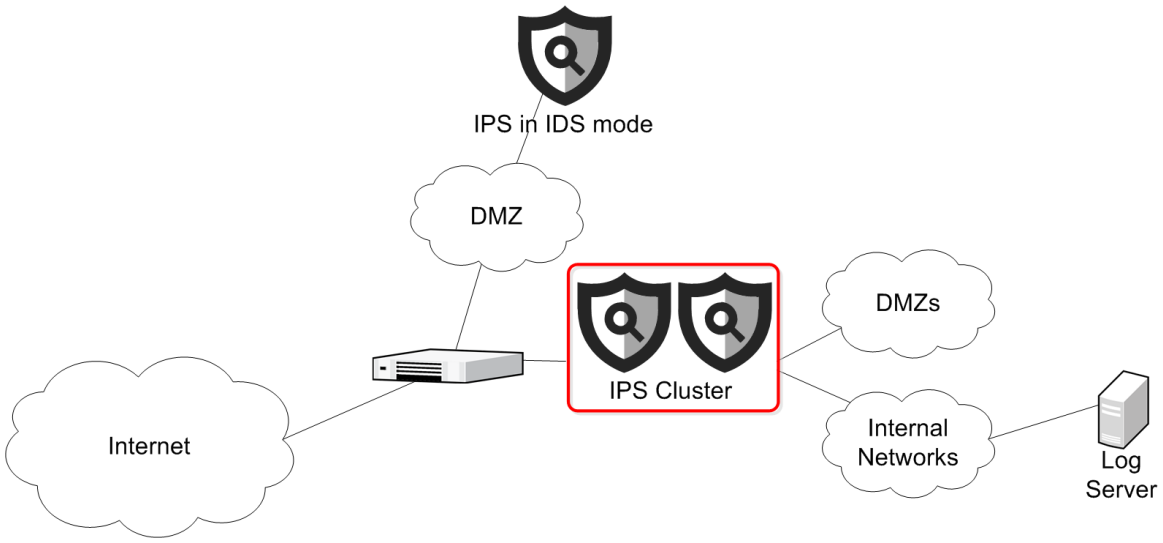
Example: medium-scale IPS installation

The following is an example of a medium-scale installation with an IPS Cluster.

The example company's branch office at Munich has a moderate number of end-user clients. Some services are only offered at the London headquarters and used remotely through a VPN. There are still many local servers, but mostly for internal and partner use. Also, there are some administrators at this location who are responsible for the:

- Daily upkeep of the office infrastructure
- Small satellite office in Vienna

Large branch office in Munich



In this case, the company has made the following decisions:

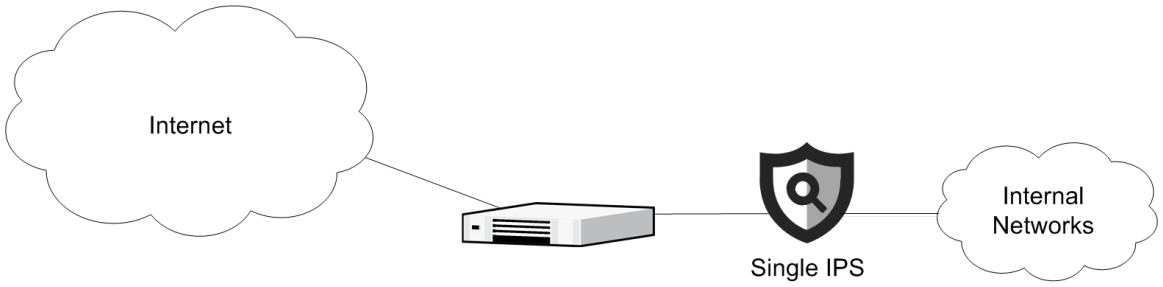
- Because there are administrators who browse logs extensively at this site, there is a dedicated Log Server here.
- One IDS is installed to inspect the network traffic in a DMZ that supports partner access.
- The IPS Cluster is placed in a dedicated DMZ for security.

Example: small-scale IPS installation

The following is an example of a small-scale installation with a Single IPS engine.

The example company's small satellite office at Vienna has a relatively low number of end-user clients, and there are no servers of any major significance. Users rely mostly on the services at the Munich office, which they access through a VPN. Also, the users have direct Internet access for general web browsing. There are no local administrators. Administrators in Munich manage the systems remotely.

Small satellite office in Vienna



To inspect the low-volume traffic that the end users' Internet and VPN bound communications generate, the company installs a Single IPS at the office. Because there are no local administrators and the traffic volumes are low, the logs are sent to the Munich Log Server. When the logs are sent to the Munich Log Server, it is quick and easy for the responsible administrators there to view and manage the data.

Post-installation steps for Forcepoint Network Security Platform in the IPS role

There are some steps to follow after you have completed the installation, installed a basic policy, and turned the IPS engines online.

Note

The configuration information is stored on the Management Server. Most changes are transferred to the engines only when you install or refresh the IPS policy.

The basic administration tasks you must learn or complete next include the following:

- How to read and control the operating state of IPS engines.
- Adjusting the automatic tester that monitors the operation of the IPS engines and the surrounding network.

After you have installed your first IPS policy, your next task is gathering information about the events detected in your networks during a “tuning period”. Once you have enough information on what kind of traffic — malicious and harmless — can be seen in your network, you can edit your policies to improve the detection accuracy and to get rid of false alarms. The most typical customization steps include:

- Creating your own policy or policy template.
- Editing the Ethernet rules, Access rules, and Inspection rules.
- Creating your own custom Situations.

Deploying Layer 2 Engines in IPS or Passive Engine mode

Layer 2 Engines can be configured in IPS mode or Passive Engine mode.

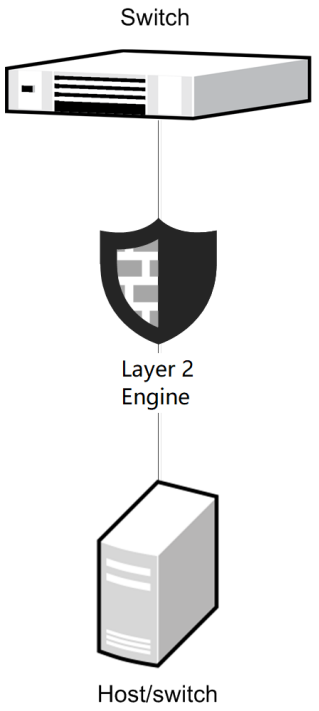
Layer 2 Engine deployment in IPS mode

In an IPS configuration, the Layer 2 Engines are installed inline directly in the traffic path.

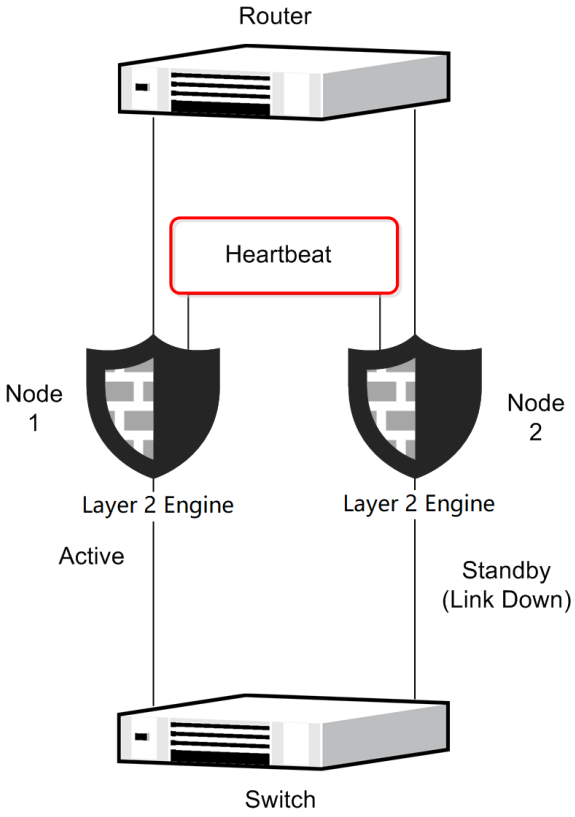
CAUTION

Always use standard cabling methods with an inline Layer 2 Engine. Use crossover cables to connect the appliance to hosts and straight cables to connect the appliance to switches.

Single inline Layer 2 Engine



Active/Standby Layer 2 Engine Cluster



Related concepts
Cable connection guidelines for IPS and Layer 2 Engines on page 90

Layer 2 Engine deployment in Passive Engine mode

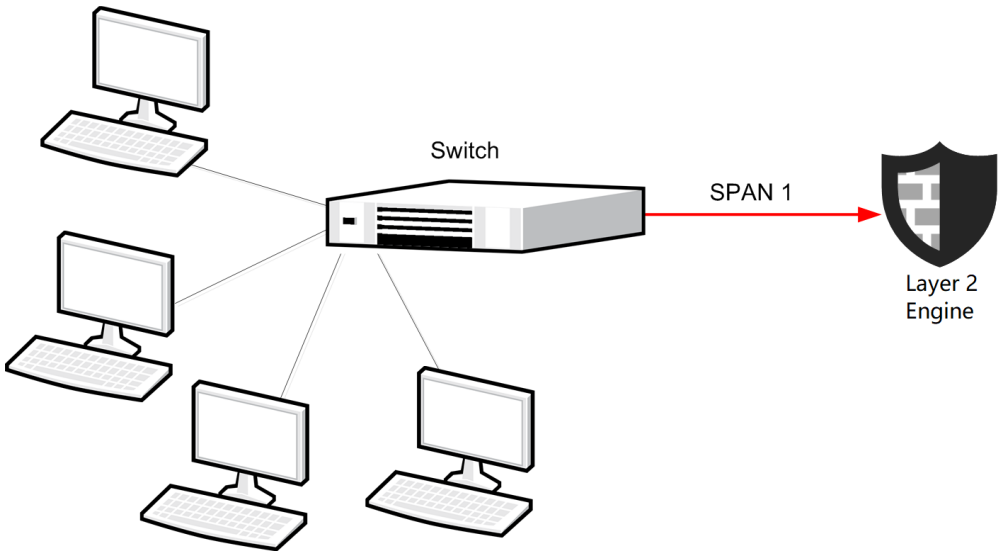
In Passive Engine mode, a Layer 2 Engine inspects but does not actively filter traffic.

Layer 2 Engines can be deployed in Passive Engine mode in two ways:

- In capture mode to inspect packets that have been duplicated for inspection through SPAN or mirror ports.
- In passive inline mode by setting the engine to only log connections by default.

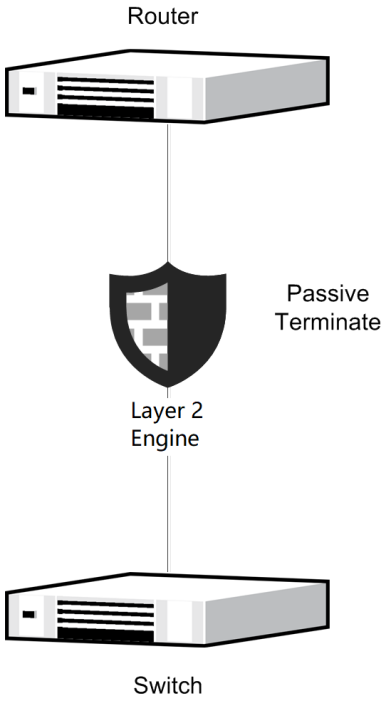
In a capture mode installation, packets are duplicated for inspection through a SPAN or mirror port on a switch/router. In a Layer 2 Engine Cluster, each node must be connected to a SPAN or mirror port of its own.

Passive Engine: a Single Layer 2 Engine in capture mode with SPAN/mirror ports



When you select Only Log Connection mode for the global Default Connection Termination, you can deploy Layer 2 Engines in Passive Engine mode in an inline configuration.

Passive Engine: a Single Layer 2 Engine in passive inline mode



Layer 2 Engine deployment example

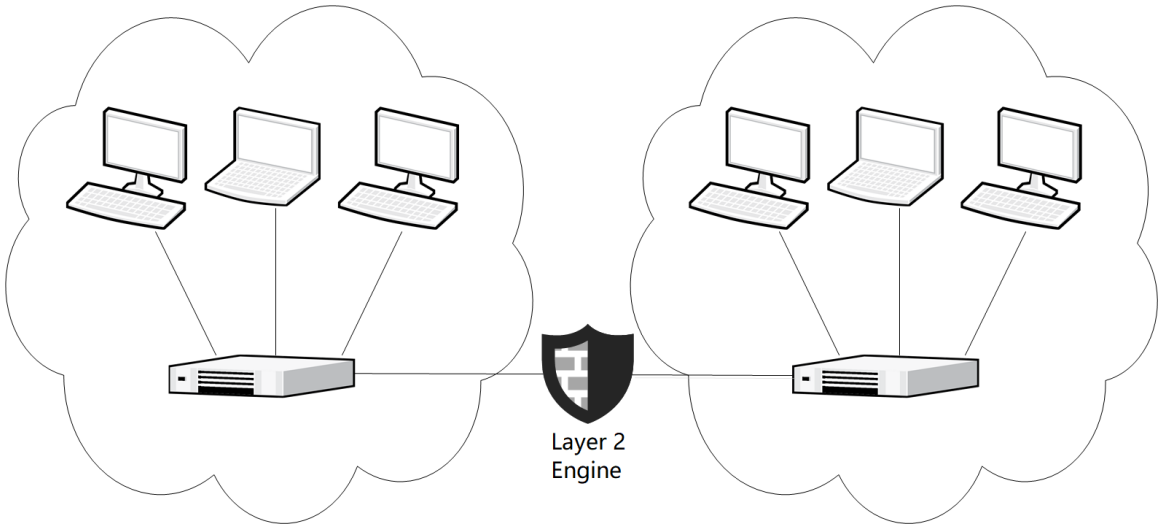
This example shows how to deploy Forcepoint Network Security Platform in the Layer 2 Engine role in an organization.

The scenario presented here is not meant to be representative of a typical installation. The main focus here is to highlight some of the criteria that can be used in planning deployment. The example covers considerations that affect most installations, but does not comprise an exhaustive list of all factors that you might need to consider. The Layer 2 Engines can be deployed in alternative ways in this example scenario. For example, depending on issues that are not covered here, such as the physical layout of the individual local networks, the hardware available, and budget constraints.

Single Layer 2 Engine example

This example uses a Single Layer 2 Engine in an organization that has a large internal network. Administrators want to prevent hosts connected to different switches in the same network segment from communicating directly at the protocol level. Using the Layer 2 Engine makes it possible to implement access control for any Ethernet protocols between switches within the same network segment. There is no need to change the network topology.

Single Layer 2 Engine in an intranet



Post-installation steps for Forcepoint Network Security Platform in the Layer 2 Engine role

There are some steps to follow after you have completed the installation, installed a basic policy, and turned the Layer 2 Engines online.

Note

The configuration information is stored on the Management Server. Most changes are transferred to the engines only when you install or refresh the Layer 2 Engine Policy.

The basic administration tasks you must learn or complete next include the following:

- How to read and control the operating state of Layer 2 Engines.
- Adjusting the automatic tester that monitors the operation of the Layer 2 Engines and the surrounding network.

After you have installed your first Layer 2 Engine Policy, your next task is gathering information about the events detected in your networks during a “tuning period”. Once you have enough information on what kind of traffic — malicious and harmless — can be seen in your network, you can edit your policies to improve the detection accuracy and to get rid of false alarms. The most typical customization steps include:

- Creating your own policy or policy template.
- Editing the Ethernet rules, Access rules, and Inspection rules.
- Creating your own custom Situations.

Cable connection guidelines for IPS and Layer 2 Engines

The cabling of IPS engines and Layer 2 Engines depends on the engine type and the installation.

Make sure that all Ethernet cables are correctly rated (CAT 5e or CAT 6 in gigabit networks).

Follow standard cable connections with inline IPS engines and Layer 2 Engines:

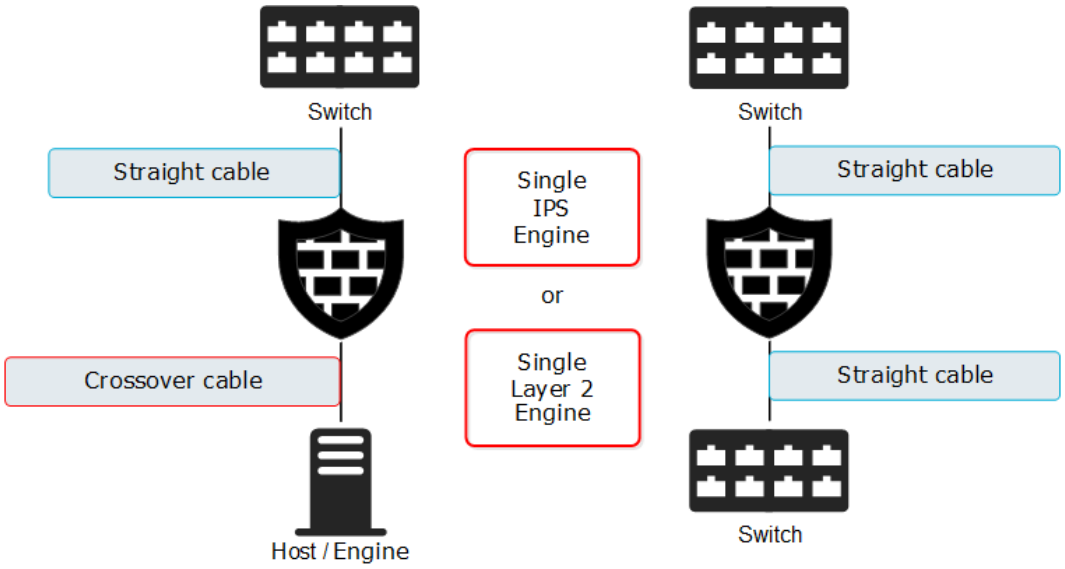
- Use straight cables to connect the IPS engines and Layer 2 Engines to external switches.
- Use crossover cables to connect the IPS engines and Layer 2 Engines to hosts (such as routers or Engines).

Note

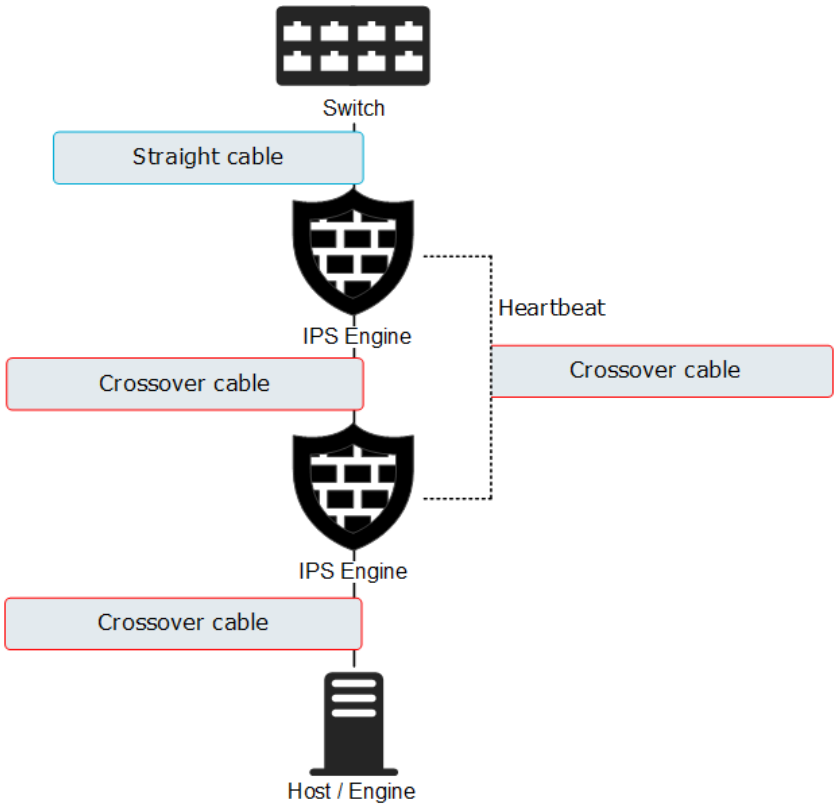
Fail-open network interface cards support Auto-MDIX, so both crossover and straight cables might work when the IPS engine is online. However, only the correct type of cable allows traffic to flow when the IPS engine is offline and the fail-open network interface card is in bypass state. It is recommended to test the IPS deployment in offline state to make sure that the correct cables are used.

Cable connections for Master Engines that host Virtual IPS engines or Virtual Layer 2 Engines follow the same principles as the connections for inline IPS engines and Layer 2 Engines.

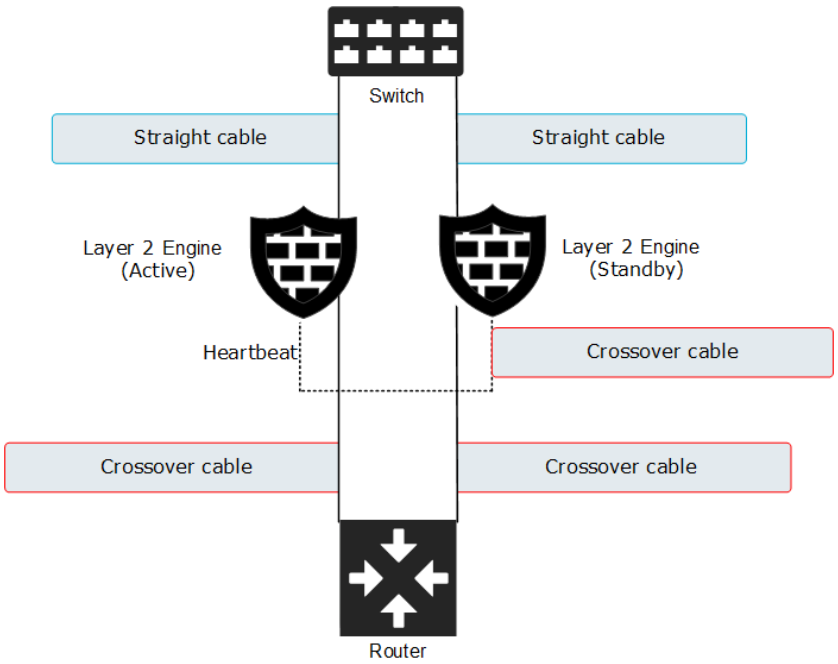
Correct cable types for Single IPS engines and Single Layer 2 Engines



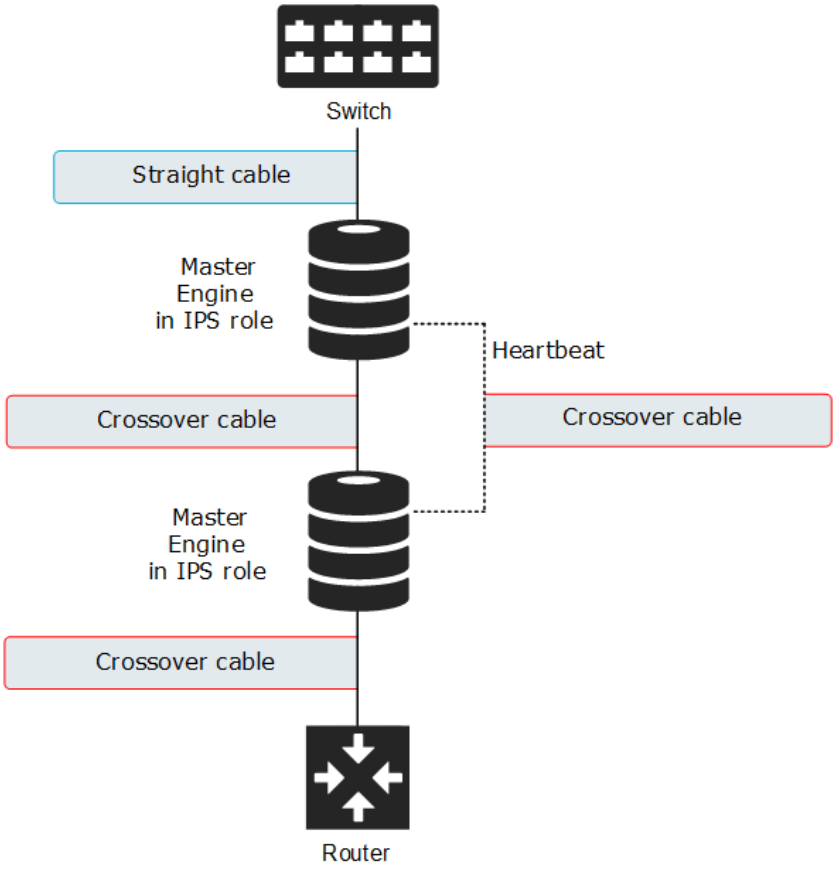
Correct cable types for Serial IPS Clusters



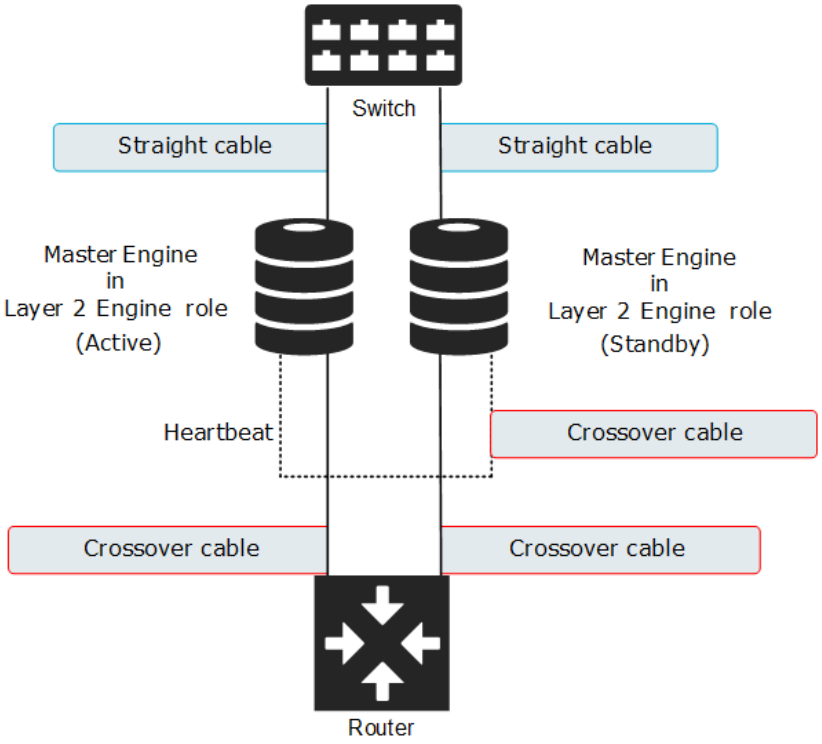
Correct cable types for Active/Standby Layer 2 Engine Clusters



Correct cable types for Serial Virtual IPS Clusters



Correct cable types for Active/Standby Virtual Layer 2 Engine Clusters



Speed and duplex settings for Security Engines

Mismatched speed and duplex settings are a frequent source of networking problems.

The basic principle for speed and duplex settings is that network cards at both ends of each cable must have identical settings. This principle also applies to the automatic negotiation setting: if one end of the cable is set to auto-negotiate, the other end must also be set to auto-negotiate and not to any fixed setting. Gigabit standards require interfaces to use auto-negotiation. Fixed settings are not allowed at Gigabit speeds.

For Inline Interfaces, the settings must be identical on both links within each Inline Interface pair. Use identical settings on all four interfaces, instead of just matching settings at both ends of each cable (two + two interfaces). If one of the links has a lower maximum speed than the other link, the higher-speed link must be set to use the lower speed.

Part III

Setting up

Contents

- Using the SMC Client on page 97
- Network address translation (NAT) and how it works on page 119
- Configuring system communications on page 125
- Managing certificates for system communications on page 149
- Managing elements on page 179

After deploying the SMC components, you are ready to start using the SMC Client and carrying out some of the first configuration tasks.

Chapter 7

Using the SMC Client

Contents

- SMC Client and how it works on page 97
- Log on to the SMC on page 100
- Log on to the SMC using certificate-based authentication on page 101
- Customize the SMC Client layout on page 102
- Bookmark SMC Client views on page 103
- Change the logon view on page 106
- Centralized management of global system settings on page 106
- View, approve, and commit pending changes on page 107
- Change the default language of the SMC Client on page 108
- Using search features on page 108
- Communicating with other administrators on page 113
- Use Tools Profile elements to add commands to elements on page 114
- Using the online help locally on page 115

The SMC Client provides the user interface for setting up, managing, and monitoring all features in the SMC.

SMC Client and how it works

Use the SMC Client for configuring, controlling, and monitoring the SMC Manager.

You can manage the Security Engines in the SMC Client. You can also use the SMC Client to monitor third-party devices.

The SMC Client offers several task-specific views. There are alternative ways to change between the different views:

- The main menu and the toolbar shortcuts are always available.
- More links are provided, for example, in the right-click menus of elements and in the logs. You can also bookmark your most frequently visited views.

You have several options for opening a new view:

- Clicking a link or main toolbar icon replaces the current view with the new one.
- Clicking while holding the **Ctrl** key opens the new view as a new tab.
- Clicking while holding the **Shift** key opens the new view as a new window.

To open a new empty tab, click **+ New Tab**. You can also use the keyboard shortcut **Ctrl+T**. From the list of views that opens, select the view to be shown in the new tab.

The bottom right corner of the SMC Client window shows the memory usage of the SMC Client. If the memory usage gets too high, the SMC Client automatically restarts, and an alert and an audit entry are generated.

Related concepts

Introduction to elements on page 179
Getting started with monitoring the system on page 211
Overviews and how they work on page 227
Getting started with the Logs view on page 281
System monitoring tools in the SMC Client on page 212

Related tasks

Change the logon view on page 106

What the Configuration view shows

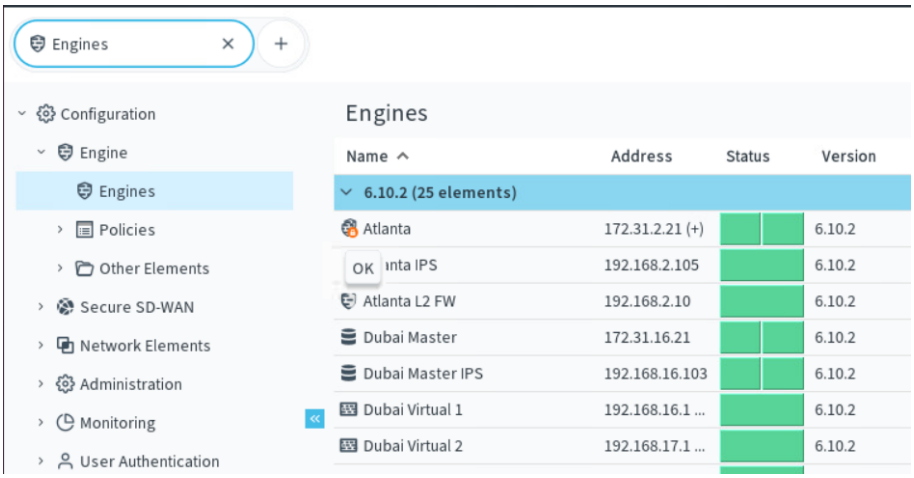
The **Configuration** view allows you to view, change, and add configuration information in the system.

There are branches in the **Configuration** view for Engine Configuration, Network Element, Secure SD-WAN Configuration, Administration, Monitoring Configuration, and User Authentication tasks. These configuration branch elements can be accessed directly from the SMC Client sidebar navigation.

The configurations are stored as elements, which are shown in a tree structure. Elements are created and changed through the right-click menus that open when you right-click a tree branch or an element. The main level of the branches contains the elements that change most often. Supporting and less frequently changed elements can be found under the **Other Elements** branch.

- The **Engine Configuration** branch allows you to manage Security Engine elements and configure Security Engine policies.
- The **Network Elements** branch allows you to manage various hosts, networks, and servers.
- The **Secure SD-WAN Configuration** branch allows you to configure ISP connections, inbound and outbound traffic management, VPNs, VPN Gateways, and SSL VPN Portals.
- The **Administration** branch allows you to manage the system, including access rights, updates, licenses, administrator accounts, and alert escalation.
- The **Monitoring Configuration** branch allows you to create statistical reports, diagrams, and configure more monitoring-related features (such as third-party device monitoring).
- The **User Authentication** branch allows you to configure user authentication and directory services, and manage user accounts.

Security Engine branch of the Configuration view



Related concepts
Introduction to elements on page 179

The Policy Editing view

The instructions for traffic handling are stored as rules in Policy elements. You can view and edit the rules in the **Policy Editing** view.

You can open policies in two modes. Any number of administrators can simultaneously check the rules in preview mode. When you open the policy in edit mode, the policy is locked for you exclusively.

The **Policy Editing** view has tabs for:

- Different types of rules in the policy
- Side pane for selecting and creating elements that you use in the rules

Related concepts
Getting started with policies on page 799

Reporting

The reporting feature allows you to create statistical summaries based on log data and stored statistical data. The Reports can be viewed in the SMC Client or exported automatically or manually.

Related concepts
Getting started with reports on page 311

Log on to the SMC

The SMC Client connects to the Management Server and to Log Servers.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select the Management Server in one of the following ways.
   - Select an existing Management Server IP address or DNS name.
   - Click **Add Management Server**, then enter the IP address or DNS name of the Management Server.

   In Demo mode, select 127.0.0.1.

2) Enter the user name and password for the Administrator you defined during the Management Server or SMC Appliance installation.

   In Demo Mode, both the user name and password are "demo".

   Note

   In FIPS mode, previously used user names are not shown in the logon dialog box.

3) Click **Log On**.

Result

After you log on to the SMC Client, the SMC Client shows the date and time when you last logged on to the SMC Client, and the IP address from which you last logged on. If your administrator permissions have been changed since the last time you logged on, you are notified that your permissions have been changed.

Log on to the SMC using certificate-based authentication

You can log on to the SMC using an X.509 certificate stored in the Windows certificate store or on a smart card, such as a Common Access Card (CAC).

Before you begin

To use smart cards for authentication, you must have smart card reader hardware and software.

To use certificate files for authentication, you must save the certificates in the Windows certificate store.

You must export the CA certificate that has signed the TLS Credentials element that is used by the Management Server, import the certificate on each administrator's computer, and configure the operating system to trust the certificate.

Note

Certificate-based authentication is only supported for SMC Clients installed in Windows 10. Certificate-based authentication is not supported for Web Portal Users.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) If you have a smart card, insert the smart card into the smart card reader.

2) Click ⋮, then select **Client Certificate** from the **Authentication Method** options.

3) Select the Management Server in one of the following ways.
   - Select an existing Management Server IP address or DNS name.
   - Click **Add Management Server**, then enter the IP address or DNS name of the Management Server.

4) Click **Log On**.

5) (First logon only) To accept the certificate chain for the Management Server, click **Accept**.

6) If there is more than one certificate on the smart card or in the Windows certificate store, select the certificate to use for authentication, then click **Select**.

7) (Smart card only) In the **PIN** field, enter then PIN for the smart card, then click **Login**

Result

After you log on to the SMC Client, the SMC Client shows the date and time when you last logged on to the SMC Client, and the IP address from which you last logged on. If your administrator permissions have been changed since the last time you logged on, you are notified that your permissions have been changed.

Customize the SMC Client layout

You can select different panes and view options through the ☰ **Menu** > **View** menu. You can also change the size of the text in various views.

Some layout options are specific to the currently active view and some options are global, that is the layout options are dispayed for views as applicable. The layout is saved as your preference for further use.

Tip

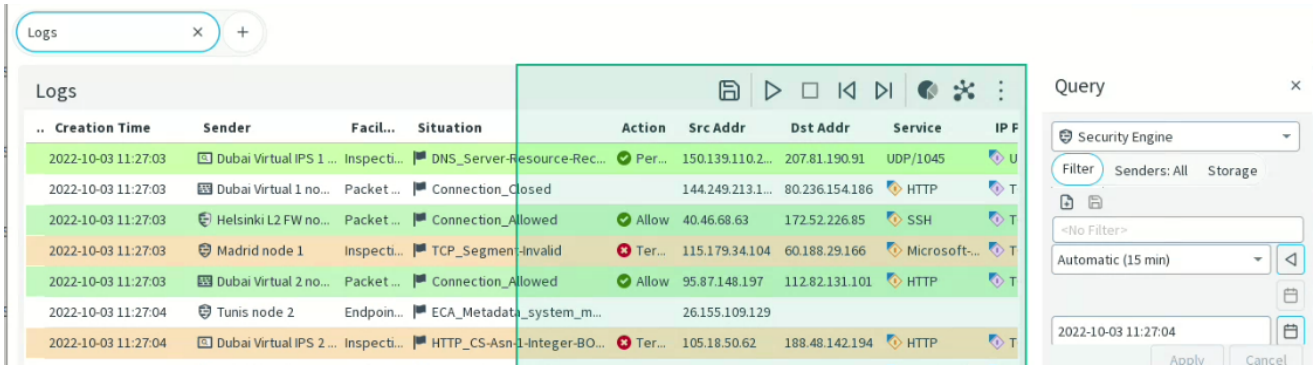
Bookmark alternative layouts to quickly return to a specific view and layout at any later time.

Steps

1) To resize a pane, drag by the outer edge of the pane as usual when resizing.

2) To move a pane, drag by the title bar at the top like you would move a window.

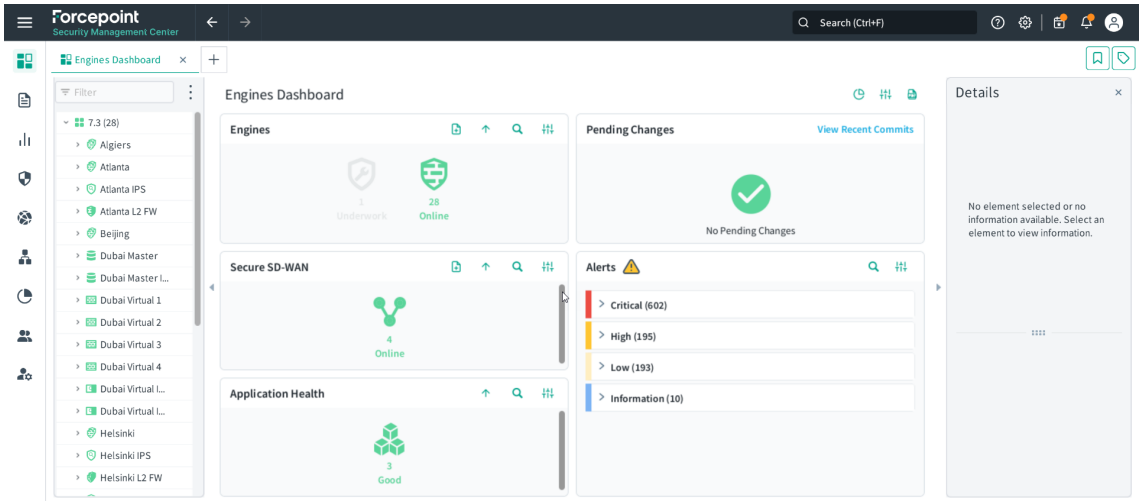
   You can move the panes in several positions that are highlighted as you drag the pane around. Drop the pane where you prefer to have it. If the highlighted area completely covers some other pane, the second pane adds a tab.



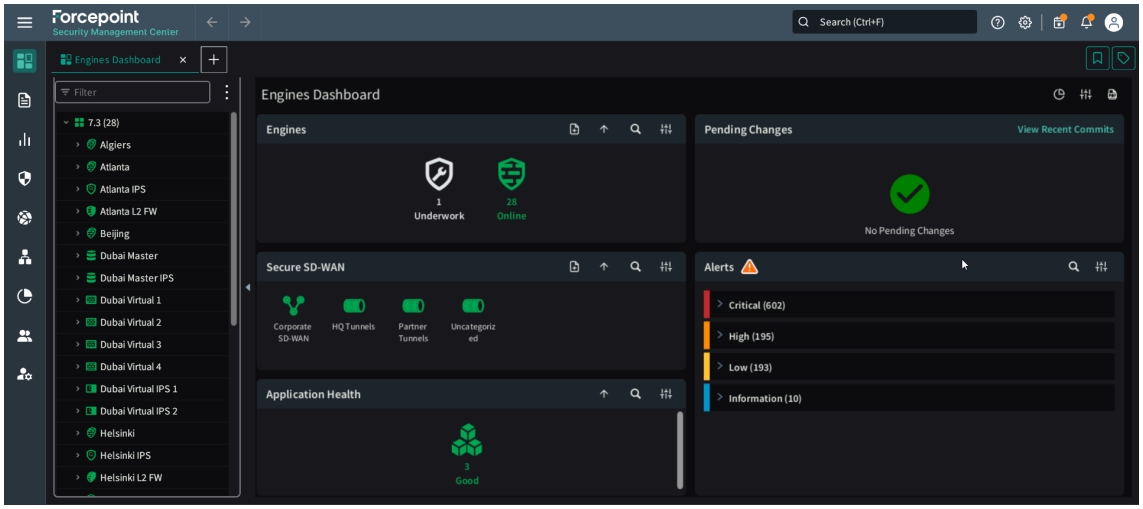
3) To change the size of the text in policies, Configuration views, or the **Logs** view, for example, select ⋮ **More actions** > **Text Size**.

4) To change the theme of the SMC, do the following:

- To apply the light theme, select **Menu** > **View** > **Theme** > **Light**.



- To apply the dark theme, select **Menu** > **View** > **Theme** > **Dark**.



Note

You must restart the SMC client session (not the SMC itself) to switch between the themes.

5) To restore the views to their defaults, select **Menu** > **View** > **Layout** > **Reset Layout**.
The text size is not affected.

Bookmark SMC Client views

Bookmark frequently used views and arrange bookmarks into folders to more easily find them.

Bookmarks can store many of the settings you have selected in views. For example, you can bookmark a **Logs** view with particular filtering settings. Several windows and tabs can be stored behind a single click when you combine bookmarks into *Bookmark Folders*.

Bookmarks in the default **Shared Bookmarks** folder are shown to all administrators that log on to the same Management Server. Other bookmarks are private to the SMC Clients of individual administrators.

Bookmark the current view in the SMC Client

You can create a bookmark for the currently active tab and other window-level elements in the configuration you select.

Some view-specific options are stored in bookmarks. For example, the:

- Currently active filter in the **Logs** view.
- Type of elements that are listed in a **Configuration** view at the time the bookmark is created.

Bookmarking is a main window action, so the properties dialog boxes for the various elements are never included in the bookmark.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Arrange the view as you would like to see it when the bookmark is opened.

2) Select **Menu** > **Bookmark** > **Add Bookmark**.

3) (Optional) The default name is taken from the bookmarked view's title, but you can change the name in the **Name** field. You can also add comments for your reference in the **Comment** field.

4) (Optional) Next to the **In Folder** field, click **Select** and select the folder where the bookmark is placed.
   - The default **Bookmarks** creates the folder at the top level of the bookmarks tree.
   - Select the **Shared Bookmarks** folder if you want other administrators to see this bookmark. All other folders are private to your SMC Client.
   - Select the **Toolbar** folder or one of its subfolders to add the bookmark folder to the toolbar. If the **Toolbar** folder is not available, activate the bookmarks toolbar.

5) (Optional) Deselect **Window Layout** if you prefer the bookmark to not change the layout of the window when you open it.

6) Click **OK**.

Bookmark all open tabs in the SMC Client

You can create a Bookmark Folder that contains bookmarks for all tabs you have open in the same window.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Open the tabs you want to bookmark in the same window.

2) Close any tabs you do not want to bookmark.

3) Select ☰ **Menu** > **Bookmark** > **Bookmark All Tabs**.

4) Fill in the Bookmark Folder properties and click **OK**.

Create bookmark folders in the SMC Client

Bookmark folders organize bookmarks and make it easier to open several bookmarks at once.

The folders you create are also added as items under the ☰ **Menu** > **Bookmark** menu.



Tip

You can bookmark all open tabs in a bookmark folder that is automatically created to contain new bookmarks.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select ☰ **Menu** > **Bookmark** > **Manage Bookmarks**.

2) Right-click **Bookmarks** in the tree, then select **New Bookmark Folder**.

3) Enter a name in the **Name** field. You can also add comments for your reference in the **Comment** field.

4) (Optional) Click **Select** next to the **In Folder** field, then select the folder where the bookmark is placed.
   - The default **Bookmarks** creates the folder at the top level of the bookmarks tree.
   - Select the **Shared Bookmarks** folder if you want other administrators to see this bookmark. All other folders are private to your SMC Client.
   - Select the **Toolbar** folder or one of its subfolders to add the bookmark folder to the toolbar. If the **Toolbar** folder is not available, activate the bookmarks folder.

5) Click **OK**.

Add bookmarks to the toolbar in the SMC Client

You can add your bookmarks to the toolbar under the shortcut icons.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select ☰ **Menu** > **View** > **Layout** > **Bookmark Toolbar**.
   The bookmark toolbar is shown under the toolbar icons.

2) Click the default **New Toolbar Folder** item.

3) Enter the name for the first folder to add to the toolbar and click **OK**.
The first folder appears in the toolbar. The **Toolbar** folder is added to the bookmark hierarchy, allowing you to add, remove, and edit the bookmarks in the toolbar.

Next steps

Add bookmarks to the toolbar by storing the bookmark in the **Toolbar** folder or one of its subfolders. Move existing bookmarks to the toolbar by dragging and dropping the bookmark or bookmark folder to the **Toolbar** folder in the ☰ **Menu** > **Bookmark** > **Manage Bookmarks** tree.

Change the logon view

You can choose which view opens when you log on to the SMC Client, replacing the default **Dashboard** view.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Arrange the screen with the windows, tabs, and views you want to have open at each logon.

2) Select ☰ **Menu** > **Bookmark** > **Save as Startup Session**.

Centralized management of global system settings

Use the Global System Properties dialog box to centrally manage global system settings and configure password policy settings.



Note

You can only change the settings when you are logged on to the Shared Domain.

Related tasks

Configure automatic updates and upgrades on page 1292

View, approve, and commit pending changes

You can view the configuration changes that you and other administrators have made before the new configurations and policies are transferred to the Security Engines. If an optional approval workflow is enforced, changes must be approved before they are transferred to the engines.

Pending changes are shown by default in the **Dashboard** view for all engines and on the engine-specific home pages.

You can view, commit, and transfer pending changes to the engines. You can use pending changes to prevent transferring configurations that are not complete. You cannot reject individual changes. If you want to override a change, edit the element or policy again, then approve both changes. If a later change overrides an earlier change, only the most recent change is transferred to the engine.

You can optionally enable an approval workflow in which an administrator must approve changes before they are committed. For example, a supervisor or senior administrator can view and approve the changes made by other administrators. Administrators with the following permissions can view the changes, approve the changes, and transfer the configurations to the engines:

- Administrators that have the Approve Changes permission
- Administrators with unrestricted permissions (superusers)

By default, the same administrator who made the changes cannot approve the changes. You can optionally allow administrators to approve their own changes.

You can approve pending changes individually, but you must commit all of the changes at the same time. You can commit pending changes only when all changes have been approved.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) (Optional) To compare the pending changes to the engine's current policy, click **View Changes** on the home page of the engine.

2) (Optional) To view recent policy uploads, click **View Recent Commits**. To return to the **Pending Changes** pane, click **View Pending Changes**.

3) If an approval workflow is enforced, approve some or all pending changes.
   - To approve an individual pending change, click the Approval Status checkmark icon for that change.
   - To approve all changes, click **Approve All**.

4) To commit all pending changes and transfer them to the engine or engines, click **Commit Changes**.

Result

The changes made to configurations and policies by all administrators are transferred to the engines.

Related concepts

What the Pending Changes pane shows on page 219

Related tasks
Enforce an approval workflow on page 383

Change the default language of the SMC Client

Each administrator can change the default language of the SMC Client during installation or manually after installation. The options are English and Japanese.

Before you begin

During a local installation of the SMC, you can set the default language of the SMC Client to English or Japanese in the installation wizard. Each administrator can select English or Japanese in the logon dialog box when logging on to the SMC Client.

Steps

1) To change the language after installation, locate the `locale.properties` file in the `<user>/.stonegate/user_locale` folder.

2) Change the smc.locale.default setting in the `locale.properties` file to the language you want:
   - `smc.locale.default=ja` for Japanese
   - `smc.locale.default=en` for English

3) Save the changes to the `locale.properties` file.

Using search features

You can search for elements, references to other elements, duplicate IP addresses, duplicate Service elements, unused elements, users, and elements in the Trash.

Note

Administrator permissions restrict which elements appear in the search results. Each administrator only sees elements for which they have permissions. If you use administrative Domains, the search results only show elements in the Domain that you are currently logged in to, and elements in the Shared Domain.

Use the search bar

The search bar is always shown at the top of the SMC Client window. You can search for elements, policies, and folders, for example.

The search results are grouped and ranked by relevance. When editing an element or a policy, you can drag-and-drop elements from the search results list. You can also perform tasks from the search results. For example, if you search for “password”, in the **Actions** category, you can change the Management Server database password. If you search for the name of an engine, you can open the Engine Editor in preview mode.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Enter a search word or phrase in the search bar.



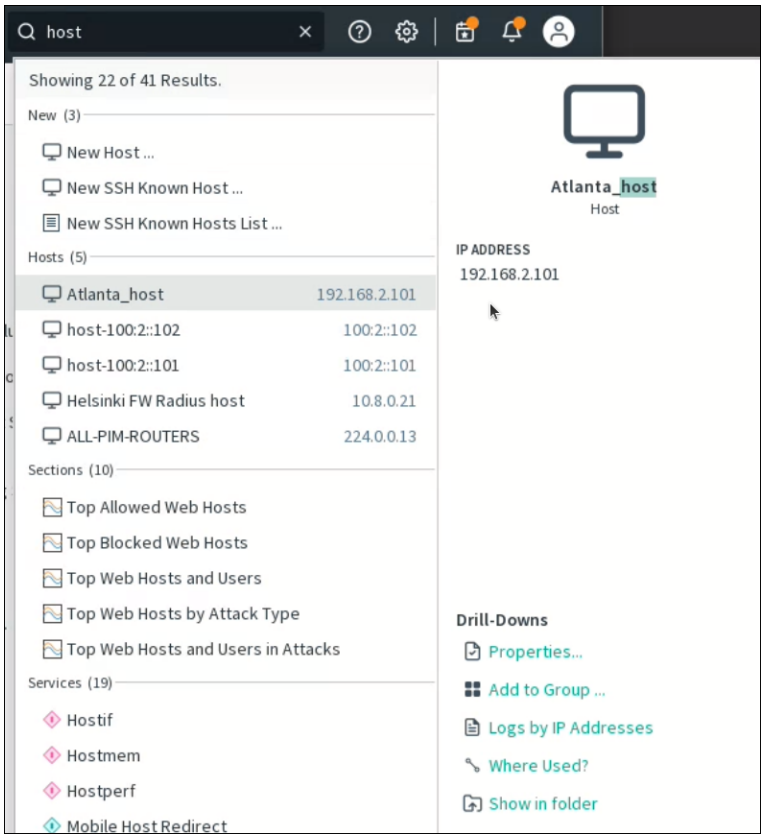
2) To open the properties of an element in the search results, click the element.

   If you click a policy or engine, the policy or engine opens in preview mode.

Tip

If you search for an element type, you can create a new element under the **New** category.

3) To preview an element in the search results list, place the cursor over the element.



Tip

You can also use the keyboard arrow keys to move up and down in the list.

You can see some of the properties of the element. You can also perform some drill-down actions, such as checking where an element is used.

4) To view a full list of search results, select **Show All Matching Elements**.
The element search view opens, and your search criteria is used automatically.

Use type-ahead search

Type-ahead search allows you to quickly filter elements.

You can use type-ahead search in most views in which element lists, tables, or trees are displayed. You can also use type-ahead search in rule cells in policies.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Activate type-ahead search in one of the following ways:
   - Select an element in a list, table, or tree.
   - Click the small **Search** icon in any view that includes the icon.

- In a policy editing view, click a rule cell in which elements can be used.

2) Type the information that you want to search for (for example, an element name, an IP address, a comment, or a TCP/UDP port number).

Result

The SMC Client filters the display to only show the elements that include what you typed.

Note

In policy editing views, when you have a Category Filter activated, only search results that match the Category Filter are displayed.

Use the search tool

Use the search tool to find elements in the SMC.

Available search options

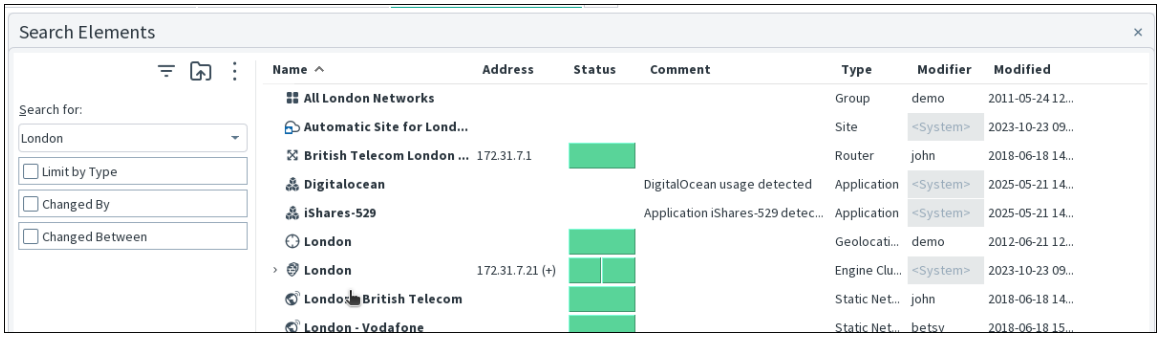
Search option	Description
Search	Search for elements based on an element property, such as a name, comment, or IP address.
Search References	Search for references to elements, to see where they are used. For example, you can find the references to elements you want to delete; referenced elements cannot be deleted until the references are removed. **Tip** You can also right-click an element, then select **Tools** > **Where Used?**.
Search DNS	Search for hosts by their DNS name. The DNS search queries a DNS server, and the hosts found on the DNS server are compared to the Host elements defined in the SMC. **Note** The Management Server must have Zone Transfer rights to be able to make queries on the DNS server.
Search Duplicate IPs	Search for elements that have the same IP address in the SMC.
Search Duplicate Services	Search for Service elements that have the same protocol and destination port number.
Search Unused Elements	Search for elements that are not referenced by any other element.
Search Users	Search for users from an LDAP domain.
Search Trash	Search for elements that have been moved to the Trash.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Menu** > **Search**, then a search option.

The Element Search



2) Enter the search criteria.

3) (Optional) Select options to limit the search.

4) Click **Search** to start the search.
The search results are displayed.

Tip

If the element you searched for does not exist, right-click the empty search results, then create a Host or Network element that has the name and IP address configured according to your search terms.

Create Host elements from DNS search results

If the DNS search results do not find an existing Host element for a host name, you can create Host elements based on the host names and IP address found by the search.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click one or more IP addresses in the search results, then select **Create Network Element(s)**.

2) (Optional) If you want to add the new Host elements to a Group, select **Create in a Group**, then enter a name in the **Group Name** field.
If the Group element does not yet exist, a new Group is automatically created.

3) Define a name for the new Host elements.

4) Click **OK**.

   The **Network Element Creation** view opens in a new tab to show the progress of the element creation process. The status of each new element is displayed in the **Info** column:

   - If the status is **Created**, the element was successfully created.
   - If the status is **Not created (name already in use)**, an element with the selected name exists. If you want to change the Host element's name, right-click the **Name** cell, then select **Properties** to open the element properties. Change the name and click **OK** to save the changes.

5) Click **Close** to close the **Network Element Creation** tab and return to the search results.

   The names of the new elements are shown in the **Network Element** column.

Communicating with other administrators

The administrator messaging feature allows you to communicate with other administrators who are currently logged on to the SMC Client.

For example, you can inform administrators about service breaks before upgrading the SMC. Administrator messaging is enabled by default.

Enable or disable messaging

An administrator with unrestricted permissions (superuser) can enable or disable administrator messaging.

If Domain elements have been configured, the setting is applied in all Domains.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.

2) Expand the **Access Rights** branch in the **Administration** tree.

3) Right-click **Administrators** and select or deselect **Administrator Messaging Enabled**.

Send messages to other administrators

Administrators can send messages to all other administrators.

Only administrators who have the Manage Administrators permission can send messages to individual administrators. Each administrator must be logged on to a unique administrator account for individual messages to be sent.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Click **Send Message** in the status bar at the bottom right of the SMC Client.

2) Select the **Administrators** to whom you want to send the message.
   - (Administrators with the Manage Administrators permission) Click **Select** to select individual administrators.
   - Click **Set All Administrators** to send the message to all administrators.

3) Enter your message, then click **Send**.
   The message is sent to the selected administrators.

Use Tools Profile elements to add commands to elements

You can add commands (for example, tracert, SSH, or ping) to an element's right-click menu with Tools Profiles. The commands are added in the **More actions** submenu.

Create Tools Profile elements

Tools Profile elements add commands to the right-click menus of other elements. You can include information dynamically from the element definition in the command.

Only one Tools Profile can be selected for each element, but each Tools Profile can include several commands.

The commands are run on the workstation that is running the SMC Client. Commands are operating system-specific. You must add a separate command for each operating system. Administrators see commands according to their operating system (for example, a Linux command is not shown if the SMC Client is running in Windows).

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.

2) Browse to **Other Elements** > **Tools Profiles**.

3) Right-click **Tools Profiles**, then select **New Tools Profile**.

4) Enter the **Name** for the new Tools Profile. This is added as an item under the **More actions** submenu of elements that the Tools Profile is attached to.

5) Click **Add**, then select the operating system.

6) Double-click the **Name** cell, then enter the item to add.

7) (Optional) To run the command in a console application, such as the command prompt in Windows or terminal in Linux, select **Run in console**.

8) Double-click the **Command** cell, then enter the command or the full path to the application.

9) (Optional) Double-click the **Parameters** cell, then enter the parameters for the command.
In addition to static parameters, the following two variables can be used:
   - `${IP}` — the primary IP address of the element that is right-clicked.
   - `${NAME}` — the name of the element that is right-clicked.

10) Click **OK**.

Attach Tools Profile elements to elements

You can attach a Tools Profile to elements that have a single primary IP address.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click an element and select **Properties**.

2) Select the Tools Profile in one of the following ways:
   - Select a Tools Profile from the list.
   - Select **Other** and browse to the Tools Profile.
   - Select **New** and create a Tools Profile.

3) Click **OK**.
The commands defined in the Tools Profile are added to the right-click menu of the element in the **More actions** submenu.

Using the online help locally

By default, the SMC Client's online help is accessed through the Internet. When you use the online help locally, the online help is available even when there is no Internet connectivity.



Note

When you use a local copy of the online help, you must manually update the online help when a new version is available.

Use the online help locally with an installed SMC Client

When you have installed the SMC Client on your own computer, you can configure the SMC Client to use a copy of the online help from your own computer or from a server in the local network.

Steps

1) Download the online help .zip file for your release from https://help.forcepoint.com/flexedge/sd-wan/en-us/7.3.0/onlinehelp/network_security_platform_730_pg_a_en-us_online_help.zip.

2) Extract the .zip file to a suitable location in your local network.
   You can also extract the Help file to a share or your local intranet server.

3) On the computer where you use the SMC Client, browse to the <user home>/.stonegate/data folder.
   Example: In Windows, browse to C:\Users\<user_name>\.stonegate.

   Note

   If the SMC Client is open, close it before editing the SGClientConfiguration.txt file.

4) Edit the SGClientConfiguration.txt file.

5) Add a parameter `HELP_SERVER_URL=` and enter the path to the main folder under which the online help files are stored as the value for the parameter.
   - If you extracted the online help on the same computer where you use the SMC Client, enter `file:///` and the path.
     Example: If you extracted the online help on your own computer to C:\help\network_security_platform_730_pg_a_en-us_online_help, enter `HELP_SERVER_URL=file:///C:/help/network_security_platform_730_pg_a_en-us_online_help`.

     Note

     Use only forward slashes (/) in the URL even if the operating system uses backslashes (\) in file paths. If the path contains spaces, replace them with %20 in the URL.

   - If you extracted the online help on a server in the local network, enter `http://` and the path.
     Example: If you extracted the online help to a folder called network_security_platform_730_pg_a_en-us_online_help on an intranet server, enter `HELP_SERVER_URL=http://<intranet.server>/network_security_platform_730_pg_a_en-us_online_help`.

6) Save the changes to the SGClientConfiguration.txt file.

7) Start the SMC Client.

Result

The SMC Client automatically uses the online help from the specified location.

Use the online help locally with Web Access

When you use Web Access, you can configure the Management Server or Web Access Server to use a local copy of the Online Help for SMC Client sessions in a web browser.

Steps

1) Download the online help .zip file for your release from https://help.forcepoint.com/flexedge/sd-wan/en-us/7.3.0/onlinehelp/network_security_platform_730_pg_a_en-us_online_help.zip.

2) On the Management Server or Web Access Server host where SMC Web Access is enabled, create the following folder: `<smc_installation_folder>/webserver/webswing/main/Help`.

3) Copy the online Help .zip file to the `<smc_installation_folder>/webserver/webswing/main/Help` folder and extract the content.

4) (Linux only) Recursively change the owner of the `<smc_installation_folder>/webserver/webswing/main/Help` directory to the sgadmin user.
Example:

```
# chown -R sgadmin:sgadmin /usr/local/forcepoint/smc/webserver/webswing/main/Help/
```

5) On the Management Server or Web Access Server host where SMC Web Access is enabled, browse to the `<smc_installation_folder>/data` folder.



Note

If you installed the Management Server in the `C:\Program Files\Forcepoint\SMC` folder in Windows, the SGClientConfiguration.txt file is located in the `C:\ProgramData` folder.

6) Edit the SGClientConfiguration.txt file.

7) Add a parameter `HELP_SERVER_URL=https://<SMC_Web_Access_IP>:<Web_Access_port>/Help/`.
Example: If the IP address of the Management Server is 192.168.100.55 and the listening port for SMC Web Access is 8085, enter `HELP_SERVER_URL=https://192.168.100.55:8085/Help/`.

8) Save the changes to the SGClientConfiguration.txt file.

9) Exit your SMC Client session and open new session.

Result

SMC Client sessions in a web browser now automatically use the Online Help from the specified URL.

Chapter 8

Network address translation (NAT) and how it works

Contents

- Network address translation and how it works on page 119
- Static source translation on page 120
- Dynamic source translation on page 121
- Static destination translation on page 121
- Destination port translation on page 122
- IPv6 transition mechanisms on page 122

Network address translation (NAT) means changing the IP address or port information in packets. Most often, NAT is used to allow internal hosts to communicate via networks where their actual address is not routable and to conceal the internal network structure from outsiders.

Network address translation and how it works

Network address translation (NAT) changes the source or destination IP address or port for packets traversing the engine.

NAT is most often used to hide internal networks behind a single or just a few routable IP addresses on the external network. NAT is also often used to translate an external, routable destination address into the private internal address of a server. For destination NAT, port translation (sometimes referred to as PAT) is also possible when the protocol in question uses ports. Port translation can be used to redirect a standard service, such as HTTP (port 80/TCP), to a non-standard port (for example, port 8080/TCP). The NAT rules are stored in policy elements.

NAT is applied to traffic that has been already been allowed by Access rules that have connection tracking enabled. If you have Access rules that turn off connection tracking for some traffic, you cannot use address translation with those connections.

There are five possible methods for network address translation (these methods are explained in more detail in the next sections):

- *Static source translation*, which translates each single IP address to some other single IP address (one-to-one relationship).
- *Dynamic source translation*, which translates several IP addresses to a single IP address or a small pool of IP addresses (many-to-one/many-to-some relationship) differentiated by port. This method is not supported with Multi-Link if the Loose connection tracking mode is used.
- *Static destination translation*, which translates each single IP address to some other single IP address (one-to-one relationship).
- *Destination port translation*, which translates a port to a different one (one-to-one relationship).

- Both source translation and destination translation for the same connection.

Dynamic destination translation is done automatically as part of the Server Pool feature.

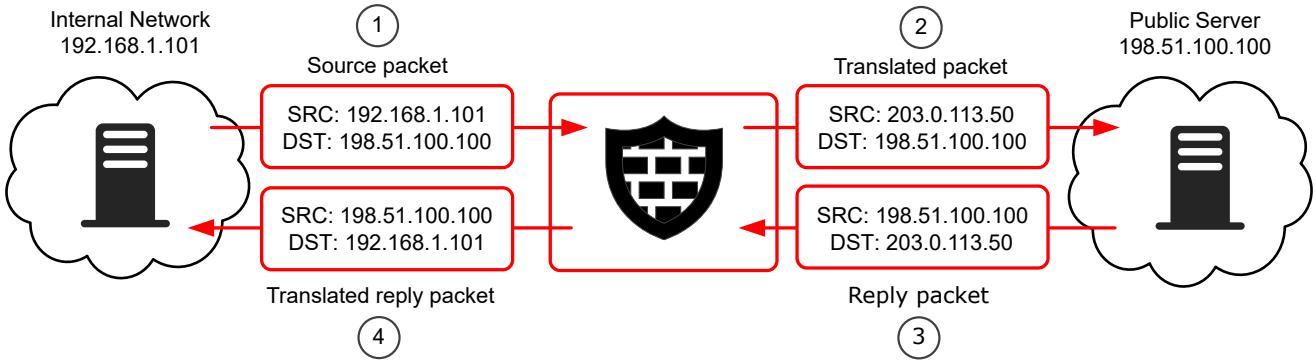
Also, when NAT is applied, *return address translation* is needed to allow reply packets to reach the correct sender or to show the source address that the destination host expects. However, return address translation does not normally need configuration as it is applied automatically with the help of connection tracking.

Related concepts
Getting started with inbound traffic management on page 749
Getting started with policies on page 799

Static source translation

In static source translation (one-to-one source translation), the source IP address of a certain host is always translated using the same specific IP address.

Static source translation provides one-to-one source translation. Often, the original source address is the actual assigned IP address for a device on an internal network or DMZ. The translation is then applied to a public IP address belonging to the public IP address range assigned by the Internet service provider (ISP).



1. The packet starts out with the original source (SRC) and destination (DST) IP addresses.
2. The engine replaces the source address of the packets with a translated source IP address.
3. The server responds, using the translated IP address as the destination of the response.
4. Connection tracking information is used to automatically translate the reply packets. The engine replaces the destination IP address in the server's response with the original address so that the responses find their way back to the host.

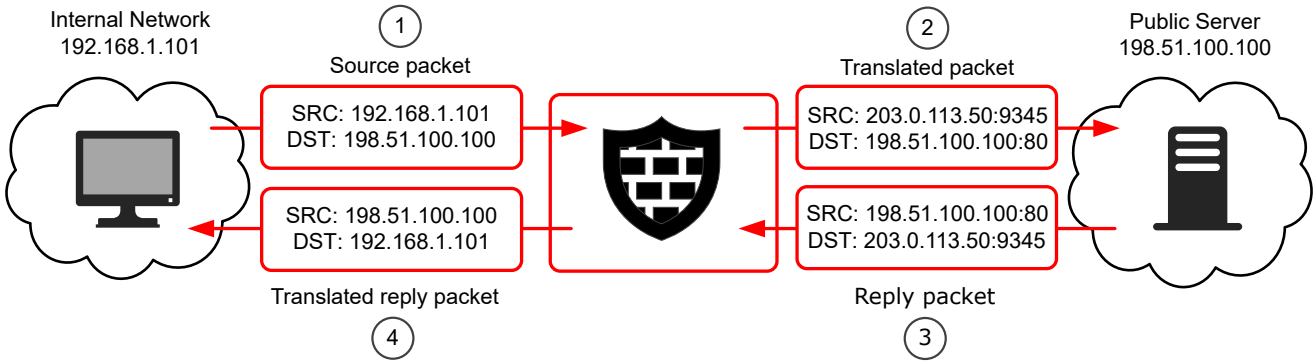
You can also define static translation using whole networks. There is still a fixed one-to-one relationship between each original and translated IP address, so the original and translated networks must be of the same size. The addresses map to their counterparts in the other network. For example, if you translate the network 192.168.10.0/24 to 203.0.113.0/24, the host 192.168.10.201 is always translated to 203.0.113.201.

Dynamic source translation

Dynamic source translation allows translating many original IP addresses to a much smaller pool of translated addresses, even a single IP address.

Dynamic source translation, sometimes referred to as hide NAT, is often used to mask the internal networks of a company behind one or a few public, routable IP addresses provided by an ISP.

This illustration shows the process for dynamic source translation. Because dynamic source translation involves multiple hosts using the same IP address (in a many-to-one or many-to-some relationship), the engine needs more information to differentiate the connections when the reply packets arrive. For this, the engine uses the source port.



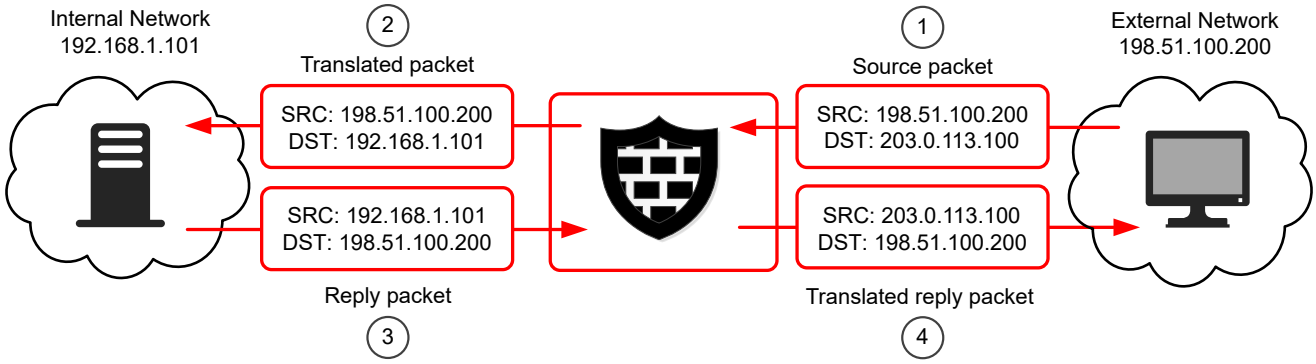
1. Hosts make connections.
2. Each host is assigned a unique port from one of the unreserved high ports to track its connections.
3. The reply packet is sent to the same unique port.
4. The destination is translated to the original source address and port

Static destination translation

Destination translation is typically needed to translate new incoming connections from a server's public IP address to the server's private IP address.

You can use static destination translation for both IP addresses and ports.

In this illustration, a host on the Internet connects to a server on the internal network.



1. The host connects to the external, public IP address.
2. The engine translates the destination address to the private IP address of the server on the internal network.
3. The server sends its response back.
4. The engine automatically translates the source address back to the external IP address.

You can also define static translation for whole same-size networks at once. This works in the same way as in static source translation.

Destination port translation

Destination NAT can also be used to translate ports.

For example, web traffic to the corporate web servers on a DMZ would typically come in on port 80. However, an administrator might want to run the web service on a non-standard port for security or network management reasons. The original destination port can be translated using static destination port translation with or without destination address translation.

IPv6 transition mechanisms

IPv6 transition mechanisms enable communication between devices that have only IPv4 addresses and devices that have only IPv6 address. They can also enable limited IPv4 top IPv4 communication over IPv6 only network connections. These translation mechanisms do the translation between IPv4 and IPv6 addresses.

Note

- The IPv6 transition mechanism feature is not supported on Virtual Engines.
- Only one translation mode can be activated at a time.
- In order to use these translation mechanisms Security Engine must have at least one IPv4 address and at least one IPv6 address.
- Protocols with separate control and data or media connections and connections that use IPv4 or IPv6 addresses in the payload will not work through these translation mechanisms.
- IPv6 prefixes used in the translations must not overlap with the engines own IPv6 addresses.
- Translation in the transition mechanisms support only unicast communication.
- Translated connections traverse through Security Engine as two separate connections. However, depending on the translation mode, either IPv4 or IPv6 packets to and from translation are handled without creating rules manually.
- Use of translation can lead to Path MTU discovery related problems if there are any faults in the way the networks in the path work.

NAT64

NAT64 is a stateful IPv6 to IPv4 translation mechanism defined in RFC 6146. When NAT64 mechanism is activated, access rules that control traffic are made in the IPv6 side. IPv6 traffic with destination address matching the configured prefix will go to NAT64 translation. There is no need to create any rules for this traffic in the IPv4 access rules. All traffic configured to the IPv4 pool will go to translation.

Note

- NAT64 supports only unicast UDP, TCP, and ICMP traffic.
- NAT64 can not be used with active-active clustering.
- In the RFC 6877, NAT64 is referred as 464XLAT PLAT.
- Due to RFC 6146 compliant NAT functionality, size of the IPv4 pool will set strict limit to maximum number of connections.
- Due to RFC 6146 compliant NAT implementation, there will be additional connection state for the NAT processing.
- Local IPv4 pool addresses must not overlap with the engines own addresses or NAT addresses used in IPv4 rules.

Configuration: See the section *Engine Editor > Add-Ons > IPv6 Transition Mechanism*.

464XLAT CLAT

Client side translation of RFC 6877. Used together with 464XLAT PLAT at the other edge of IPv6 only network in order to allow IPv4 traffic through IPv6 only network.

When used, local IPv4 addresses are detected based on the configured routing. Access rules controlling the traffic must be made in the IPv4 access rules. There is no need to create any rules for translated traffic in the IPv6 access rules. In the IPv4 side, traffic which does not have any route will be sent to translation. In the IPv6 side, the traffic from the remote IPv6 prefix to the local IPv6 prefix will be sent to translation.

Configuration: See the section *Engine Editor > Add-Ons > IPv6 Transition Mechanism*.

SIIT EAM

Stateless IPv4/IPv6 translation as defined in RFC 7757. When used without defining any explicit address mapping entries, same as SIIT, access rules controlling the traffic must be made in the IPv4 access rules.

There is no need to create any rules for translated traffic in the IPv6 access rules. In the IPv4 side, traffic which does not have any route will be sent to translation. In the IPv6 side, traffic which will have valid IPv4 route after translation are sent to translation.

Configuration: See the section *Engine Editor* > *Add-Ons* > *IPv6 Transition Mechanism*.

Enable communication between IPv4 and IPv6 devices

Enable IPv6 transition mechanisms that allow communication between devices that have only IPv4 addresses and devices that have only IPv6 address.

Steps

1) Select **Engine Configuration**.
2) Right-click a Security Engine, then select **Edit <element type>**.
3) Browse to **Add-Ons** > **IPv6 Transition Mechanism**.
4) From the **Type** drop-down menu, select translation mechanism, then configure the settings.
5) Click **Save**.

Chapter 9

Configuring system communications

Contents

- Considerations for setting up system communications on page 125
- Define contact IP addresses on page 127
- Select the Location for the SMC Client on page 142
- Create HTTP Proxy elements on page 143
- Configuring NTP for the SMC Appliance and Security Engine on page 144
- Configuring SMC Appliance communications on page 146
- Considerations for Multi-Link system communication on page 148

System communications involve traffic between SMC components, traffic between SMC components and external components that are a part of the system configuration, and external access into the system.

Considerations for setting up system communications

Engines and Layer 2 Engines do not automatically allow communications of other system components that pass through the engine. Make sure that all necessary traffic is allowed in the engine's policy.

The predefined **Firewall Template Policy** and **Layer 2 Firewall Template Policy** allow most types of system communications between the engine and the components it interacts with. You must create rules to allow any other communications through Engines or Layer 2 Engines.

System communications through a NAT device

If NAT is applied between two SMC components, you must define the translated IP address as a *contact address*. In NATed communications, the contact address is contacted instead of the component's real IP address. A single component can have several contact addresses.

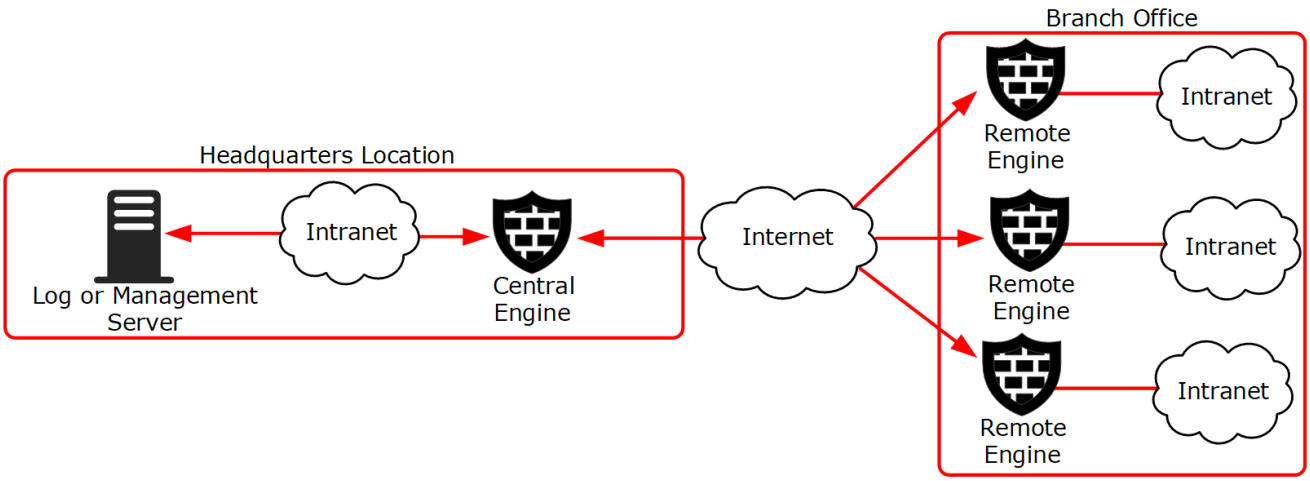
Location elements define when a contact address is used and which of the defined contact addresses is used. When NAT is applied between two communicating SMC components, you must separate them into different locations. Components that are in the same location use the primary IP address when communicating with each other and ignore all contact addresses. When components contact a component that belongs to a different location, they use the defined contact address.

For example, when a Management Server contacts an engine node through NAT, the Management Server uses the NATed contact address, not the engine's real IP address. The NAT device between the components translates the NATed address to the engine's real IP address as usual.

You can define a default contact address for contacting a component in the Properties dialog box of the element. When components that belong to another location contact the element and the element has no contact address defined for its location, the element's default contact address is used.

If you do not select a location for an element that has the location option, the element's location is set to Not Specified. When SMC and Security Engine components contact components for which the location is not specified, they use the element's default contact address. If you want components to use the primary IP address when communicating with each other, the elements must belong to the same location, or you must define a location-specific exception for the elements.

Example of using contact addresses and locations



In this example scenario, a Management Server and a Log Server manage SMC components both at a company's headquarters and at three branch offices.

- The SMC servers and the Central Engine are at the "Headquarters" location.
- The Remote Engines are all at the "Branch Office" location.

In this scenario, contact addresses are typically needed as follows:

- The Engine at the headquarters or an external router can provide the SMC servers external IP addresses on the Internet. The components at the branch offices contact the servers through the Internet. The external addresses of the SMC servers must be defined as contact addresses for the "Branch Office" location.
- The Branch Office Engine or an external router can provide external addresses for the SMC components at the branch office. The external IP addresses of the engines must be defined as contact addresses for the "Headquarters" location so that the Management Server can contact the components.
- Alternatively, the external address of each component can be defined as a Default contact address without adding a specific entry for "Headquarters" or "Branch Office". The Default contact address is used when a component does not have a specific contact address definition for the contacting component's location. The components must still be divided into separate locations for the contact address to be used.

If there are SMC Clients used at any of the branch offices, each administrator must also select "Branch Office" as their location in the SMC Client. Selecting the SMC Client location allows administrators to view logs from a remote Log Server that is behind a NAT device.

Related concepts

Network interfaces for Security Engine on page 543
Defining IP addresses as elements on page 919

Related tasks
Select the Location for the SMC Client on page 142

Related reference
Forcepoint Security Management Center ports on page 1457
Security Engine ports on page 1460

Define contact IP addresses

Contact addresses are required when NAT is applied between to SMC components. You can define contact addresses for Security Engines, Master Engines, and most types of server elements.

You can specify an IP address or Fully Qualified Domain Name (FQDN) as the contact address to enable Security Engine to contact SMC management server or log server. If FQDN is specified, then you must also specify a DNS server in the engine configuration. The DNS server is used to match server host names to their corresponding IP addresses.



Note

FQDN resolves to IPv4 or IPv6 address. If FQDN resolves to multiple addresses, then all the addresses are attempted and the first IP address that works is used.

The contact addresses are defined in the element properties. The contact addresses are based on Location elements. You can also define a Default contact address that is used whenever no contact address is defined for a certain Location.

You create the Locations and add elements to the Locations based on how your network is set up. Then you define the Contact Addresses for each element for each Location in the properties of the elements. All SMC components in other Locations then use the addresses defined for their Location for contact.

Related concepts
Considerations for Multi-Link system communication on page 148

Related tasks
Define endpoints for VPN Gateway elements on page 1174

Create Location elements

If network address translation (NAT) is applied between communicating SMC components, the components must be assigned to different Locations in the configuration. You create the Locations and add elements to them based on how your network is set up.

If a system has several Locations, but each component always has the same external IP address, each element only needs a Default contact address. When new system elements are added, they have to be assigned a specific Location, but they only need a Default contact address.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.
2) Browse to **Other Elements** > **Locations**.
3) Right-click **Locations**, then select **New Location**.
4) Enter a **Name** and an optional **Comment.**
5) Browse to the type of elements you want to assign to the Location element in the **Resources** pane.
6) Select one or more elements and click **Add**.
   The selected elements are added to the **Content** pane.
7) Click **OK**.

Define Management Server or Log Server contact addresses

If NAT is used between SMC components or if SMC servers are contacted by external servers, you can configure contact addresses for Management Servers and Log Servers.

You can configure multiple contact addresses for each type of server.

If you use Multi-Link, we recommended defining a separate contact address for each NetLink for the Management Server and the Log Server. This way, if a NetLink goes down, the engines can still be managed and can still send status and log data to the Log Server.

Note

If the IP addresses at which the server can be reached change, you must manually update the server contact addresses.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Network Elements**.
2) Browse to **Servers**.
3) Right-click the server element for which you want to define a contact address and select **Properties**.
   The **Properties** dialog box for that server opens.
4) Select the **Location** of this server.

5) If necessary, edit the contact addresses.
   - A default contact address is automatically entered based on the element properties.
   - If the server has multiple default contact addresses, separate the addresses with commas.

6) (Optional) Click **Exceptions** to define further contact addresses for contacting the server from specific Locations:
   a) Click **Add** and select a Location. A new row is added to the table.
   b) Click the **Contact Address** column and enter the IP addresses or FQDN that the components belonging to this Location must use when they contact the Server.
   You can enter several contact addresses per Location for Management Servers and Log Servers. Separate the contact addresses with a comma.

   Note

   Elements that belong to the same Location element always use the primary IP address (defined in the element's properties) when contacting each other. Elements that do not belong to a specific Location belong to the Default Location.

   c) Click **OK** to close the **Exceptions** dialog box.

7) Click **OK** to close the **Server Properties** dialog box.

Related tasks

Select the Location for the SMC Client on page 142
Change the Management Server IP address on page 491

Management Server Properties dialog box

Use this dialog box to define Management Server properties.

Option	Definition
General tab	
Name	The name of the element.
Installation ID	Shows the unique installation identifier (UIID) for the SMC.
IPv4 Address	Specifies the IPv4 address or FQDN of the server. The server can have both an IPv4 and an IPv6 address.
IPv6 Address	Specifies the IPv6 address of the server. The server can have both an IPv4 and an IPv6 address.
Resolve	Automatically resolves the IP address of the server.
Location (Optional)	Specifies the location to which the server belongs if there is a NAT device between the server and other SMC components.
Contact Addresses section (All optional settings)	
Default	Used by default when a component that belongs to another Location connects to this server.

Option	Definition
Exceptions	Allows you to define exceptions to the default contact address. Opens the **Exceptions** dialog box.
Log Server	Specifies the Log Server to which the server sends its logs.
RADIUS Method (Optional)	Specifies a RADIUS authentication method for authenticating administrators. ■ **PAP** — Password Authentication Protocol. ■ **CHAP** — Challenge-Handshake Authentication Protocol. ■ **MSCHAP**, **MSCHAP 2** — Microsoft versions of the CHAP protocol. We recommend using MSCHAP 2 if the server supports it. ■ **EAP-MD5** — Extensible Authentication Protocol with an MD5 Hash. This option is selected by default.
TACACS Method (Optional)	Specifies a TACACS+ authentication method for authenticating administrators. ■ **ASCII** — American Standard Code for Information Interchange. **CAUTION** This authentication method transmits the user name and password as unencrypted plain text. ■ **PAP** — Password Authentication Protocol. ■ **CHAP** — Challenge-Handshake Authentication Protocol. ■ **MSCHAP** — Microsoft versions of the CHAP protocol. This option is selected by default.
TLS Credentials (Optional)	Specifies the TLS Credentials element that is used for certificate-based authentication of administrators.
TLS Profile (Optional)	Specifies the TLS Profile element that is used for certificate-based authentication of administrators.
Include in Database Replication (Multiple Management Servers only)	When selected, the Management Server is included in database replication between Management Servers for high availability. **CAUTION** Leave this option selected unless you have a specific reason to deselect it. Deselecting this option makes the Management Server's database incompatible with the databases of the other Management Servers.
Audit Storage Full	Specifies the action when the Management Server detects that the audit storage is full. ■ **Stop Management Server** — The Management Server writes an audit entry indicating that the audit storage is full, stops all processes, then shuts down. ■ **Overwrite Oldest Audit Entries** — The Management Server overwrites audit entries, starting with the oldest audit entries.
Category (Optional)	Includes the element in predefined categories. Click **Select** to select a category.
Tools Profile	Adds commands to the right-click menu for the element. Click **Select** to select an element.
Comment (Optional)	A comment for your own reference.

Option	Definition
Notifications tab	

Option	Definition
E-mail section — Specifies email notification details.	
SMTP Server	Select the SMTP Server that is used to send the alert notifications as email. Click **Select** to select an element.
Sender Name	Enter the name to be used in the From field of the email. If this setting is left blank, the **Default Sender Name** defined in the **SMTP Server Properties** is used.
Sender Address	Enter the email address to be used in the From field of the email. If this setting is left blank, the **Default Sender Address** defined in the **SMTP Server Properties** is used.
SMS section Click **Add** to add an element to the table, or **Remove** to remove the selected element. Click **Up** or **Down** to move the selected item up or down.	
Name	Shows the name of the channel.
Channel Type	Shows the type of the channel. ■ **Script** — SMS messages are sent using a custom script. ■ **SMTP** — SMS messages are sent using an SMTP server. ■ **HTTP** — SMS messages are sent using HTTP. You can add multiple SMS Channels Types. If the first SMS Channel fails, the subsequent SMS channels are used in the order in which they are listed. Use the **Up** and **Down** buttons to change the order of the channels if necessary.
Host/URL/Script	Shows the server, URL, or script used for SMS notification.
Edit	Opens the **Channel Properties** dialog box for the selected entry.
SNMP section	
Gateways	Enter the host name or IP address of the SNMP Gateways to which the alert notifications are sent as SNMP traps. You can specify a list of gateways separated by semicolons. If your SNMP gateway port is not the default port 162, specify the port number by typing a colon and the port after the host name (for example, `snmp-gw:4390`).
Custom Alert Scripts section	
Root Path	Enter the root path on the Management Server where custom alert scripts are executed. The default location is `<installation directory>/data/notification`. Do not define the script name here. Add the script name in the Alert Chain at each place you want to call a particular script. You can use multiple scripts.

Option	Definition
Web Access tab	
Enable	Enables the feature.
Host Name (Optional)	Enter the host name that the service uses. Leave the field blank to allow requests to any of the server's host names.

Option	Definition
Port Number	Enter the TCP port number that the service listens to. By default, port 8085 is used when Web Access is enabled on the Management Server and port 8083 when enabled on the Web Access Server. **Note** Make sure that the listening port is not in use on the server.
Listen Only on Address (Optional)	If the server has several addresses and you want to restrict access to one address, specify the IP address to use.
Session Timeout	Enter the timeout in seconds after which the session expires. While the session is still active, the administrator does not need to log on again if they close the web browser.
Server Credentials	Select the TLS Credentials element that is used for HTTPS connections. Click **Select** to select an element.
Server TLS Cryptography Suite Set	Select the TLS Cryptography Suite Set element that defines the allowed algorithms for HTTPS connections. Click **Select** to select an element.
Generate Server Logs (Optional)	Select if you want to log all file load events for further analysis with external web statistics software.
Use SSL for session ID (Optional)	Track sessions in your web application. Do not select this option if your network requires you to use cookies or URIs for session tracking.
Client Certificate Authentication	When selected, administrators can be authenticated using client certificates via SMC Web Access. **Note** - It is only supported for browser-based authentication. - To make this work, you must select the **Client Certificate** option as the authentication method in the administrator properties.

Option	Definition
Client API tab	
Enable	Enables the SMC API feature.
Host Name	Enter the host name that the service uses. Leave the field blank to allow requests to any of the server's host names. **Note** API requests are served only if the API request is made to this host name. To allow API requests to any host name, leave this field blank.
Port Number (Optional)	Enter the TCP port number that the service listens to. By default, port 8082 is used. In Linux, the value of this parameter must always be higher than 1024.

Option	Definition
Listen Only on Address (Optional)	If the server has several addresses and you want to restrict access to one address, specify the IP address to use.
Server Credentials	Select the TLS Credentials element that is used for HTTPS connections. Click **Select** to select an element.
Server TLS Cryptography Suite Set	Select the TLS Cryptography Suite Set element that defines the allowed algorithms for HTTPS connections. Click **Select** to select an element.
Generate Server Logs (Optional)	Select if you want to log all file load events for further analysis with external web statistics software.
Use SSL for session ID (Optional)	Track sessions in your web application. Do not select this option if your network requires you to use cookies or URIs for session tracking.
Web Portal Client section	
Enable	Enables the Web Portal Client feature. **Important** - You must have the SMC API setting configured and enabled before you enable the Web Portal Client feature. For more details on SMC API settings, refer to the *Configure SMC API* topic. - When you upgrade from SMC version 7.2.X or earlier to version 7.3 and the SMC API is not configured and enabled. The web portal will not function properly after the upgrade. To ensure the web portal works, you must first manually configure and enable the SMC API.

Option	Definition
Client Downloads tab	
Enable	Enables the feature.
Management Client Download	When selected, the Management Server provides the SMC Client for download on the **SMC Client Downloads** Downloads page.
ECA Evaluation	To easily deploy Forcepoint One Endpoint to a limited set of users for evaluation purposes, enable the ECA Evaluation feature. For more information, see Knowledge Base article 16193.
Host Name (Optional)	Enter the host name that the service uses. Leave the field blank to allow requests to any of the server's host names.
Port Number	Enter the TCP port number that the service listens to. By default, port 8080 is used for new SMC installations, and port 8084 is used when you upgrade the SMC. **Note** Make sure that the listening port is not in use on the server.

Option	Definition
Listen Only on Address (Optional)	If the server has several addresses and you want to restrict access to one address, specify the IP address to use.
Server Credentials	Select the TLS Credentials element that is used for HTTPS connections. Click **Select** to select an element.
Server TLS Cryptography Suite Set	Select the TLS Cryptography Suite Set element that defines the allowed algorithms for HTTPS connections. Click **Select** to select an element.
Generate Server Logs (Optional)	Select if you want to log all file load events for further analysis with external web statistics software.

Option	Definition
Connection tab	
Proxy Settings	
Use proxy server for HTTPS connection	Select if the connection from the Management Server to the Forcepoint servers requires a proxy server. **Note:** The **Proxy Address** field must contain only the proxy hostname, without http:// or https://.
Proxy address	Defines the address of the HTTP proxy.
Proxy port	Defines the port of the HTTP proxy.
Authenticate to the proxy server	Select if the proxy server requires user authentication.
Proxy user name	Enter the user name for the proxy user.
Proxy user password	Enter the password for the proxy user. By default, passwords and keys are not shown in plain text. To show the password or key, deselect the **Hide** option.

Option	Definition
Elasticsearch tab The Elasticsearch tab is only visible after you have created an Elasticsearch Cluster element. **Important** Forwarding log data to an Elasticsearch cluster is an advanced feature that requires knowledge of how to configure Elasticsearch. You must already have an Elasticsearch cluster deployed and configured in your environment.	
Elasticsearch Cluster	Shows the Elasticsearch cluster that receives log data from the SMC server.
Client Authentication Settings	Defines how the connection between the server and the Elasticsearch cluster is secured. ■ **Inherited** — The SMC server uses the settings defined in the Elasticsearch Cluster element. ■ **Override** — The SMC server uses custom settings.

Option	Definition
TLS Certificate	(*When **Override** is selected*.) Specifies the TLS certificate that is used to secure the connection between the SMC server and the Elasticsearch cluster. ■ **Use Internal Certificate** — The SMC server uses its own internal certificate. ■ **Use Imported Certificate** — The SMC server uses the specified external certificate. ■ **No Client Authentication** — The connection is not authenticated.

Option	Definition
Audit Forwarding or **Log Forwarding** tab Click **Add** to add a row to the table, or **Remove** to remove the selected row.	
Target Host	The Host element that represents the target host to which data is forwarded. Double-click to open the **Select Host** dialog box.
Service	Click the cell, then select the network protocol for forwarding data from the drop-down list. The following **network protocol** options are supported: ■ TCP ■ UDP ■ TCP with TLS ■ Kafka ■ Kafka with TLS **Note** ■ For log data in IPFIX or NetFlow v9 format, UDP is the only available network protocol. ■ You might have to define an Access rule that allows traffic to the target host. In this case, make sure that the Service you select is also used as the Service in the Access rule.
Port	The Port that is used for forwarding data. Double-click to edit the cell. The default port is 2055. For log data, the default port used by IPFIX/NetFlow data collectors is 2055. **Note** You might have to define an Access rule that allows traffic to the target host. In this case, make sure that the Port you select is also used as the Port in the Access rule.

Option	Definition
Format	Click the cell, then select the data forwarding format from the drop-down list. ■ **CSV** — Forwards in comma separated value format. ■ **Short CSV** — Forwards truncated data in comma separated value format. (Log Server only) ■ **XML** — Forwards in XML format. ■ **JSON** — Forwards log data in JSON format. ■ **CEF** — Forwards in common event format. ■ **LEEF** — Forwards in log extended event format. ■ **NetFlow v9** — Forwards in a format that is compatible with NetFlow v9. (Log Server only) ■ **IPFIX** — Forwards in a format that is compatible with IPFIX. (Log Server only) ■ **McAfee ESM** — Forwards in a format that is compatible with McAfee ESM. **Note:** You can customize the interval by which the NetFlow template is sent. This is done by modifying the parameter `NETFLOW_TEMPLATE_INTERVAL_SECONDS=<seconds>` in the SMC Log Server configuration file. The default interval is 300 seconds.
Data Type	Click the cell, then select the applicable data type option from the drop-down list. The log data for the selected data type only is forwarded to the target host.
Filter (Optional)	An optional local filter that defines which data is forwarded. The local filter is only applied to the data that matches the Audit Forwarding or Log Forwarding rule. Double-click to open the **Local Filter Properties** dialog box.
TLS Profile	Allows you to select a TLS Profile element that contains settings for cryptography, trusted certificate authorities, and the TLS version used in TLS-protected traffic. Double-click to open the **Select Element** dialog box. The TLS Profile is only available if you have selected TCP with TLS as the Service.
TLS Server Identity (Optional)	(*When a TLS Profile is selected*) Select the identity of a TLS server to secure the TLS-protected traffic from the Management Server or Log Server to an external syslog server. Double-click to open the **TLS Server Identity** dialog box.
TLS Certificate Used for Forwarding Logs	Select the certificate for TLS-protected data forwarding. ■ **Use Internal Certificate** — A Management Server or Log Server certificate (signed by the Internal CA) is used for TLS-protected syslog communication. ■ **Use Imported Certificate** — A certificate signed by an external CA is used. Click **Select** to select a certificate or to create a TLS Credentials element. ■ **No Client Authentication** — The Management Server or Log Server certificate is not authenticated.
Kafka Topic	Double-click the cell to specify a name for the kafka topic. Kafka topics are log groups that hold messages and events in a logical order, which in turn allow users to send and receive data between kafka servers with ease.

Option	Definition
NAT tab (All optional settings)	
Engine	Shows the selected engine.
NAT Type	Shows the NAT translation type: **Static** or **Dynamic**.
Private IP Address	Shows the Private IP Address.
Public IP Address	Shows the defined Public IP Address.
Port Filter	Shows the selected Port Filters.

Option	Definition
Comment	An optional comment for your own reference.
Add NAT Definition	Opens the **NAT Definition Properties** dialog box.
Edit NAT Definition	Opens the **NAT Definition Properties** dialog box for the selected definition.
Remove NAT Definition	Removes the selected NAT definition from the list.

Option	Definition
Certificate tab (All optional settings) **Note** This tab is only displayed if the **Use External Certificate Authority** feature is enabled during the SMC installation	
Current Certificate	Shows information about the current certificate of the server. Click **Export Certificate** to export the current certificate. Click **Renew Certificate** to renew the certificate.
Check Revocation	Checks against certificate revocation lists (CRLs or OCSP) whether the certificate of the new connection has been revoked. The certificate must be signed by a valid certificate authority.
Ignore Revocation Check Failures if There Are Connectivity Problems	When selected, the server ignores all CRL check failures if the server cannot connect to CRL or OCSP server. This is done based on what is specified in the certificate.
Organization (O)	The name of your organization as it appears in the certificate.
Organization Unit (OU)	The name of your department or division as it appears in the certificate.
Country (C)	Standard two-character country code for the country of your organization.
State/Province (ST)	The name of state or province as it appears in the certificate.
Locality (L)	The name of the city as it appears in the certificate.
Common Name (CN)	The value for the Common Name field in the certificate request. For server certificates, the value is typically the fully qualified domain name (FQDN).
Public Key Algorithm	The algorithm used for the public key.
Key Length	The length of the key in bits.
Signature Algorithm	The signature algorithm that was used to sign the certificate.
Subject Alternative Name [DNS]	Enter a unique subject alternative name for the identity to be certified.
Generate Certificate Request	Generates the certificate request so that you can sign it using an external certificate authority.

Define contact addresses for Security Engines

You must define a Location and a contact address if NAT is applied to the communications between the Security Engine and some other component that contacts the Security Engine.

If you use Multi-Link, add contact addresses for each NetLink.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click an engine element and select **Edit <element type>**.
   The Engine Editor opens.
3) In the **General** pane, select the **Location** for this element.
4) Browse to **Interfaces** in the navigation pane on the left.
5) In the tree view, expand the tree and double-click the Cluster Virtual IP Address (CVI), Node Dedicated IP Address (NDI), or the IP address for which you want to define a contact address.
   On Engine Clusters, the CVI contact address is used for VPNs and NDI contact addresses are used for other system communications.

Define contact addresses for a single Security Engine or a Cluster Virtual IP Address

If NAT is applied to communications between the Security Engine and some other component that contacts the Security Engine, define the Default contact address for the single Security Engine or the Cluster Virtual IP Address (CVI). You can also define location-specific contact addresses.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click an engine element and select **Edit <element type>**.
3) In the **General** pane, select the **Location** for this element.
4) Browse to **Interfaces** in the navigation pane on the left.
5) In the tree view, expand the tree and double-click the Cluster Virtual IP Address (CVI) or the IP address for which you want to define a contact address.
   On Engine Clusters, the CVI contact address is used for VPNs and NDI contact addresses are used for other system communications.

6) In the **IP Address Properties** dialog box, define the Default contact address. The Default contact address is used by default whenever a component that belongs to another Location connects to this interface.
   - If the interface has a static Default contact address, enter the Default contact address in the **Default** field.
   - If the interface has a dynamic Default contact address, select **Dynamic** (next to the **Default** field) before entering the Default contact address.

7) If components from some Locations cannot use the Default contact address to connect to the interface, click **Exceptions** to define Location-specific contact addresses.

8) Click **Add** and select the Location.

9) Click the **Contact Address** column and enter the IP address that the components in this Location use when they contact the interface or select **Dynamic** if the interface has a dynamic contact address.

   Note

   Elements that belong to the same Location element always use the primary IP address (defined in the element's properties) when contacting each other. Elements that do not belong to a specific Location belong to the Default Location.

10) Click **OK** to close the **Exceptions** dialog box.

11) Click **OK** to close the **IP Address Properties** dialog box.

Related tasks

Select the Location for the SMC Client on page 142

Define contact addresses for Node Dedicated IP Addresses

If NAT is applied to communications between the engine and some other component that contacts the engine, define the Default contact address for the Node Dedicated IP Address (NDI). You can also define location-specific contact addresses.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click an engine element and select **Edit <element type>**.

3) In the **General** pane, select the **Location** for this element.

4) Browse to **Interfaces** in the navigation pane on the left.

5) In the tree view, expand the tree and double-click the Node Dedicated IP Address (NDI) or the IP address for which you want to define a contact address.
On Engine Clusters, the CVI contact address is used for VPNs and NDI contact addresses are used for other system communications.

6) In the **IP Address Properties** dialog box, double-click the **Contact Address** cell and define the contact address for each node in the **Node Dedicated IP Address** section.

7) Enter the **Default** contact address.
The Default contact address is used by default whenever a component that belongs to another Location connects to this interface.

8) If components from some Locations cannot use the Default contact address, click **Add** to define Location-specific contact addresses.

9) Click the **Contact Address** column and enter the IP address that the components assigned to this Location must use when they contact the node.



Note

Elements that belong to the same Location element always use the primary IP address (defined in the element's properties) when contacting each other. Elements that do not belong to a specific Location belong to the Default Location.

10) Click OK to close the **Exceptions** dialog box.

11) Once you have defined the contact addresses for each node, click **OK** to close the **IP Address Properties** dialog box.

Related tasks

Select the Location for the SMC Client on page 142

Define contact addresses for an IPS Cluster or a Layer 2 Engine Cluster

If NAT is applied to communications between the IPS Cluster or Layer 2 Engine Cluster and some other component that contacts the cluster, define the Default contact address for the IPS Cluster or a Layer 2 Engine Cluster. You can also define location-specific contact addresses.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click an engine element and select **Edit <element type>**.
The Engine Editor opens.

3) In the **General** pane, select the **Location** for this element.

4) Browse to **Interfaces** in the navigation pane on the left.

5) In the tree view, expand the tree and double-click the Cluster Virtual IP Address (CVI) or the IP address for which you want to define a contact address.
On Engine Clusters, the CVI contact address is used for VPNs and NDI contact addresses are used for other system communications.

6) In the **IP Address Properties** dialog box, double-click the Contact Address cell.

7) Enter the **Default** contact address at the top of the dialog box.
The Default contact address is used by default whenever a component that belongs to another Location connects to this interface.

8) If components from some Locations cannot use the Default contact address, click **Add** to define Location-specific contact addresses.

9) Click the **Contact Address** column and enter the IP address that the components belonging to this Location must use when they contact the interface.



Note

Elements that belong to the same Location element always use the primary IP address (defined in the element's properties) when contacting each other. Elements that do not belong to a specific Location belong to the Default Location.

10) Click **OK** to close the **Exceptions** dialog box.

11) Click **OK** to close the **IP Address Properties** dialog box.

Related tasks

Select the Location for the SMC Client on page 142

Define contact addresses for an External VPN Gateway

Define a contact address for an External VPN Gateway if the IP address for contacting the gateway is different from the IP address of the gateway's interface (for example, because of NAT).

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.

2) Browse to **VPN Gateways**.

3) Right-click a Gateway element, then select **Properties**.

4) In the properties view, click the **Endpoints** tab.

5) Right-click an endpoint and select **Properties**.

6) Enter the Default contact address or select **Dynamic** if the Default contact address is dynamic.
   - The Default contact address is used by default whenever a component that belongs to another Location connects to this endpoint.

7) (Optional) If some components cannot use the Default contact address, click **Exceptions** to define contact addresses that the components use to connect to this endpoint.

8) Click **Add** and select a Location.
   A new row is added to the table.

9) Click the **Contact Address** column and enter the IP address that components belonging to this Location use when they contact the endpoint, or select **Dynamic**.

10) Click **OK** to close the **Exceptions** dialog box.

11) Click **OK** to close the **Endpoint Properties** dialog box.

Select the Location for the SMC Client

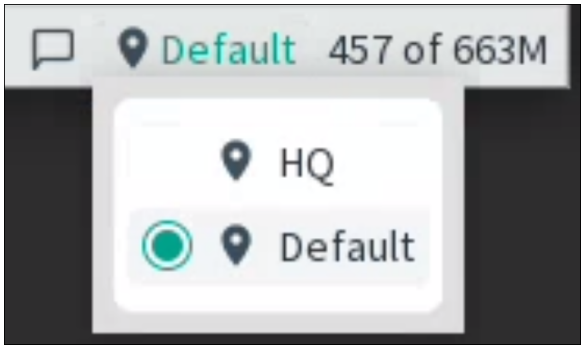
If NAT is applied between the SMC Client and a Log Server, you might need to change the Location of the SMC Client to view the logs.

The Location to select depends on the system configuration. The **Default** selection is appropriate if the Log Server has a specific Location and the Log Server's Default contact address is correct for your current network connection.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Click the **Location** name in the status bar at the bottom right corner of the SMC Client window, then select the Location.



Next steps

You might also need to add a Location and define a contact address for this specific Location in the Log Server's properties.

Create HTTP Proxy elements

You can send HTTP requests through an HTTP proxy so that the Security Engine does not need to access the external network directly.

You can use HTTP proxies when the Security Engine needs to communicate with file reputation services, sandbox services, URL categorization services, and certificate validation services.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Browse to **Other Elements** > **Engine Properties** > **HTTP Proxies**.
3) Right-click **HTTP Proxies**, then select **New** > **HTTP Proxy**.
4) Configure the settings.
5) Click **OK**.

Configuring NTP for the SMC Appliance and Security Engine

You can use external NTP servers to provide time synchronization for the SMC Appliance and Security Engines.

By default, the SMC Appliance uses the public Forcepoint NTP servers. The Management Server and default Log Server synchronize with the SMC Appliance time.

By default, Security Engines get time setting commands from the Management Server. If an Security Engine is configured to use NTP and it can successfully get the time from an external NTP server, the Security Engine ignores time setting commands from the Management Server.

You must have an SMC Appliance to configure NTP for SMC servers in the SMC Client.

Note

You cannot configure NTP in the SMC Client for SMC installations on third-party hardware or virtualization platforms. On third-party hardware or virtualization platforms, the SMC gets the time from the operating system. When you use third-party hardware or virtualization platforms, you must configure NTP at the operating system level on the third-party hardware or virtualization platform. For more information, see Knowledge Base article 9680.

Create NTP Server elements

NTP Server elements represent the external NTP servers to provide time synchronization for the SMC Appliance and Security Engines.

You can use the same NTP servers for the SMC Appliance and Security Engines.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Network Elements**.

2) Right-click **Servers**, then select **New** > **NTP Server**.

3) In the **Name** field, enter a unique name.

4) (Optional) In the **Host Name** field, enter the host name of the NTP server.

Note

If you do not enter a host name, you must enter an IPv4 address or an IPv6 address.

5) (Optional) In the **IP Address** or **IPv6 Address** field, specify the IP address in one of the following ways:
   - Enter the IPv4 or IPv6 address of the NTP server.
     The same NTP Server element can have both an IPv4 address and an IPv6 address.

- To automatically resolve the IP address from the host name in the **Name** field, click **Resolve**.

Note

If you do not enter an IPv4 address or an IPv6 address, you must enter a host name.

6) From the **Key Type** drop-down list, select the key type.

7) In the **Key ID** field, enter a unique numerical identifier for the key.
The value must be between 1—65534.

8) In the **Key** field, enter the key.

9) Click **OK**.

Enable NTP time synchronization on the SMC Appliance

You can configure the SMC Appliance to use external NTP servers so that network devices accurately log events and complete scheduled tasks.

Before you begin

Create an NTP Server element.
You must have an SMC Appliance to configure NTP in the SMC Client.

Note

You cannot configure NTP synchronization from both the command line of the SMC Appliance and the SMC Client. Command line changes persistently override any NTP changes that you configure in the SMC Client.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the SMC Client, select **Settings** > **Global System Properties**.

2) On the **NTP** tab, select **Enable time synchronization from NTP server**.

3) To add a row to the table, click **Add**.

4) To add an NTP server, right-click the **NTP Server** cell, select **Select**, then select an NTP Server element.

5) (Optional) If there is more than one NTP server, select the preferred NTP server.

6) Click **OK**.

Enable NTP time synchronization for Security Engines

You can configure Security Engines to use external NTP servers.

In environments where there are Master Engines and Virtual Engines, you can select NTP servers only for Master Engines. Virtual Engines do not communicate directly with NTP servers.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click an Security Engine, then select Edit <element type>.
3) In the navigation pane on the left, browse to **General** > **NTP**.
4) Select **Enable time synchronization from NTP server**.
5) To add a row to the table, click **Add**.
6) To add an NTP server, right-click the **NTP Server** cell, select **Select**, then select an NTP Server element.
7) (Optional) If there is more than one NTP server, select the preferred NTP server.
8) Click **Save and Refresh**.

Configuring SMC Appliance communications

You can configure SSH access to the SMC Appliance, configure serial console connections for the SMC Appliance, and set a BIOS password for the SMC Appliance.

Configuring SSH access to the SMC Appliance

You can access the SMC Appliance command line through SSH.



Note

In FIPS mode, SSH access to the SMC Appliance is not allowed.

Two authentication options are available:

- **Logon and password** — All administrators with unrestricted permissions (superusers) in SMC can use their password to log on to the SMC Appliance through SSH.
- **Public and private key** — You can generate a public and private RSA key pair for each administrator or server. For more information about using an RSA key pair to authenticate to the SMC Appliance, see Knowledge Base article 12503.

Related concepts
Getting started with administrator accounts on page 373

Connect to the SMC Appliance using a serial console

You can configure the serial port on the SMC Appliance to receive serial console connections.

Steps

1) From the command line, log on to the SMC Appliance.

2) To enable the serial console, enter the following command:

```
sudo smca-system toggle-console
```



Tip

To disable the serial console, enter the same command again.

The SMC Appliance is now listening for serial console connections.

3) To connect to the serial console, connect a console cable from your computer to the SMC Appliance, then open a terminal console program with these settings:
   - **Bits per second** — 9600
   - **Data bits** — 8
   - **Parity** — none
   - **Stop bits** — 1

Use the SMC Appliance to make outbound serial connections

You can configure the serial port on the SMC Appliance to make outbound serial connections, for example to Security Engine appliances.

Steps

1) From the command line, log on to the SMC Appliance.

2) If the serial console is enabled, enter the following command to disable the serial console:

```
sudo smca-system toggle-console
```

The SMC Appliance stops listening for serial console connections.

3) Connect a console cable from the SMC Appliance to the other device.

4) Enter the following command to start the screen utility:

```
screen /dev/ttyS0 9600
```

Considerations for Multi-Link system communication

If a remotely managed Engine has Multi-Link, we recommend adding a primary and a secondary control interface for different ISP connections. Adding a control interface for each ISP connection guarantees connectivity if one of the ISP connections fails.

Make sure that you configure these addresses consistently in the following parts of the configuration:

- For the interface address on the Engine.
- For the external contact addresses (if applicable).
- In the NAT rules of the Engine that protects the SMC Manager servers (as necessary).

If there is a Multi-Link connection between a Management Server or Log Server and the components that contact them, define a contact address for each network connection. Make sure that your NAT rules translate from each external address to the correct internal address of the SMC server.

Related concepts

Defining Multi-Link routes on page 735

Related tasks

Select system communication roles for engine interfaces on page 570

Chapter 10

Managing certificates for system communications

Contents

- How certificates work on page 149
- Types of internal certificate authorities on page 150
- Using certificates to secure communications to external components on page 154
- Creating certificates on page 158
- Configure settings for certificate validation on page 161
- Renewing certificates on page 162
- Renewing external certificates on page 166

Certificates are proof of identity that SMC components and Security Engines use to authenticate themselves in communications.

How certificates work

SMC servers and Security Engines use certificates to identify each other in system communications, and to secure communications to external components.

Note

Do not confuse certificates with licenses. Certificates are proof of identity that components use to authenticate themselves in communications. Licenses are a proof of purchase used for ensuring that your organization is a legal license holder of the software.

To be able to communicate with other SMC components, each SMC server and Security Engine must have a valid certificate.

Certificates can also be used:

- For communication with some external components.
- In VPNs for authentication between remote gateways.
- By Security Engines for TLS inspection.

By default, the certificates used in system communications are generated by the internal certificate authority (CA) that runs on the Management Server. You can optionally install the SMC with external certificate management to use certificates issued by an external CA.

Note

You can only configure the SMC to use external certificates when you install the SMC. It is not possible to change to using external certificates in an existing installation. In SMC 6.10, this feature is only available when you use the SMC Appliance.

For more information, see the *Forcepoint Network Security Platform Installation Guide*.

Related concepts
TLS inspection and how it works on page 1063
VPN certificates and how they work on page 1251

Default elements for certificates

There are several predefined elements for working with certificates.

The Management Server includes an Internal RSA Certificate Authority element. By default, Internal RSA Certificate Authority issues all certificates that SMC servers and Security Engines use for communication with other SMC components.

Predefined Trusted Certificate Authority elements represent the signing certificates of major certificate authorities. Default Trusted Certificate Authority elements are automatically added from dynamic update packages and cannot be edited or deleted. You can also create your own Trusted Certificate Authority elements to represent other certificate authorities that the SMC servers and Security Engines trust.

Limitations of certificates

Certificates used in system communications become invalid when the certificate authority changes.

The internal certificate authority can change if the Management Server is reinstalled and the configuration is recreated manually or by importing elements instead of importing a backup. Management Server backups contain certificate authority information. If backup restoration is not an option, all SMC components must receive a new certificate for system communications.

In some cases, restoring the Management Server backup might also cause the internal certificate authority to be different from the certificate authority that was used to create certificates for some components. The invalid certificates must be replaced with new ones.

Types of internal certificate authorities

The internal certificate authority can be either an internal Elliptic Curve Digital Signature Algorithm (ECDSA) certificate authority or an internal RSA certificate authority.

ECDSA is a digital signature algorithm that uses elliptic curve cryptography. Using an internal ECDSA certificate authority enables 256-bit encryption on the Management Server for connections between the Management Server and Security Engines.

You can only use one type of internal certificate authority at a time.

Change the type of the internal certificate authority

When you install the SMC, an internal Elliptic Curve Digital Signature Algorithm (ECDSA) certificate authority or an internal RSA certificate authority is automatically created. You can optionally change the type of the internal certificate authority.

Before you begin

CAUTION

Creating a new internal CA replaces the existing internal CA. We strongly recommend creating a Management Server backup before creating a new internal certificate authority.

When you create a new internal certificate authority, SMC components gradually start using the new internal CA to sign certificates. The state of the internal CA changes as the CA starts signing certificates.

Internal certificate authority states

State	Description
Created for Different Certificate Type	The new internal CA has been created, but it is not yet ready to begin signing certificates.
Ready to Use for Different Certificate Type	The new internal CA is ready to begin signing certificates. At first, only Management Server certificates are signed by the new internal CA. Certificates for other components are signed by whichever internal CA is currently used by the Management Server.
Active	Certificates for all components are signed by the new internal CA. In an environment with multiple Management Servers, the new internal CA changes to the **Active** state when all Management Servers use the new internal CA.

When you start using a new internal CA, you must recertify all SMC servers. You might also need to make initial contact between the Security Engines and the Management Server.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.

2) Browse to **Certificates** > **Internal Certificate Authorities**.

3) Right-click **Internal Certificate Authorities**, then select **Create New Internal RSA Certificate Authority** or **Create New Internal ECDSA Certificate Authority**.
   You are prompted to confirm that you want to create a new internal CA.

4) Click **Yes**.

The element creation process begins and a new tab opens to show the progress of the process. When the process is finished, the progress report shows the steps that you must take next. The status of the new internal CA is **Created for Different Certificate Type**.

5) Restart the Log Server and the Web Access Server.

6) Start the Renew Internal Certificate Authorities Task.

a) Select **Administration**.

b) Select **Tasks** > **Definition**.

c) Right-click the **Renew Internal Certificate Authorities Task**, then select **Start**.

When the task finishes running, the status of the new internal CA is **Ready to Use for Different Certificate Type**.

7) Check the progress report of the task to see what further steps are required.

a) Browse to **History**.

b) Right-click the **Renew Internal Certificate Authorities** Task, then select **Show Progress Report**.

The progress report shows which steps you must take next. Follow the instructions to resolve any issues. For example, you might be prompted to check the status or connectivity of some Security Engines.

8) Recertify the Management Server.

9) Start the **Renew Internal Certificate Authorities** Task again.

When the Task is finished, the status of the new internal CA is **Active**.

10) Recertify the Log Server and the Web Access Server.

Next steps

If you created a new internal ECDSA CA and Security Engines cannot communicate with the Management Server, make sure that 256-bit encryption is enabled on the Security Engines. Then make initial contact between the Security Engines and the Management Server.

Related tasks

Recertify SMC servers on page 162
Back up system configurations on page 1297
Start Tasks manually on page 1328

Related reference

Forcepoint Security Management Center commands on page 1429

Manually enable 256-bit security strength for Security Engines

When you start using a new internal ECDSA certificate authority, 256-bit encryption is automatically enabled for Security Engines. If an Security Engine cannot communicate with the Management Server, manually enable 256-bit encryption on the Security Engine, then make initial contact between the Security Engine and the Management Server.

Before you begin

Create a new internal ECDSA certificate authority.

Steps

1) On the command line of the Security Engine, enter one of the following commands to start the Security Engine Configuration Wizard:

   - `sg-reconfigure --no-shutdown`

     The Security Engine Configuration Wizard starts without shutting down the Security Engine. Network interface settings cannot be changed in this mode.

   - `sg-reconfigure`

     The Security Engine shuts down, then the Security Engine Configuration Wizard starts. All options are available if you have a local connection. If you have a remote SSH connection, you cannot change network interface settings because the Security Engine always uses the no-shutdown mode for SSH connections.

2) Select **Next** on each page until the **Prepare for Management Contact** page opens.

3) Select **Contact** or **Contact at Reboot**, then press the spacebar.

4) Enter the Management Server IP address and the one-time password.

   Note

   The one-time password is specific to each Security Engine and can be used only for one initial connection to the Management Server. After initial contact has been made, the Security Engine receives a certificate from the SMC for identification. If the certificate is deleted or expires, repeat the initial contact using a new one-time password.

5) Select **256-bit Security Strength**, then press the spacebar to use 256-bit encryption for the connection to the Management Server.

6) (Optional) Enter the fingerprint for the Management Server.

   a) Select **Edit Fingerprint**, then press **Enter**.

b) Enter the Management Server's certificate fingerprint.
The fingerprint is shown in the SMC Client when you save the initial configuration.

7) Select **Finish**, then press **Enter**.

Result

The Security Engine tries to make initial Management Server contact. The progress is shown on the command line.

Using certificates to secure communications to external components

You can use certificates to secure communications from the SMC servers or Security Engines to external components.

You can use certificates to secure the following types of communications:

- Forwarding log or audit data from the Management Server or Log Server to external syslog servers.
- LDAP connections between the Security Engine and external LDAP or Active Directory servers.
- Communication between Security Engines and the Forcepoint User ID Service server.
  For information about configuring the Forcepoint User ID Service server to communicate with Security Engines, see the document *How to integrate Forcepoint User ID Service with other Forcepoint products* and Knowledge Base article 14100.

The configuration consists of the following general steps:

1) Define the trusted certificate authority for securing communications with external components in one of the following ways:
   - Use one of the default Trusted Certificate Authority elements.
   - Create a Trusted Certificate Authority element and import an external CA's certificate.
   - Use the Management Server's internal certificate authority.
     Export the active internal CA's certificate, then configure the external component to trust the internal CA.

2) Create a TLS Profile element.
   TLS Profile elements define the following settings:
   - Settings for cryptography
   - Trusted certificate authorities
   - TLS version

3) To verify the identity of the TLS server to secure the TLS-protected traffic from the Log Server or the Management Server, configure TLS server identity.

For example, if you want to use the Forcepoint User ID Service server's certificate to secure communications from Forcepoint User ID Service to the Security Engine, you must create a Trusted Certificate Authority element to represent the CA, then select the CA as a trusted CA in the TLS Profile element that is used in the Forcepoint User ID Service configuration on the Security Engine.

Create Trusted Certificate Authority elements

If you want to use a certificate signed by a certificate authority that is not one of the default Trusted Certificate Authority elements, you must create a new Trusted Certificate Authority element.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.
2) Browse to **Certificates** > **Certificate Authorities** > **Trusted Certificate Authorities**.
3) Right-click **Trusted Certificate Authorities**, then select **New Trusted Certificate Authority**.
4) In the **Name** field, enter a unique name.

   Note

   No other fields on the **General** tab can be edited. The fields are filled in automatically based on the information contained in the certificate that you import.

5) On the **Certificate** tab, import a certificate.
   a) Click **Import**.
   b) Browse to the certificate, then click **Open**.
   c) Click **OK**.
6) Click **OK**.

Next steps

To use the Trusted Certificate Authority element in a TLS Profile element, create or modify the TLS Profile element.

Export certificate of the active internal certificate authority

You can use the Management Server's active internal certificate authority as the trusted certificate authority for securing communications with external components.

You must export the certificate of the active internal certificate authority, then configure the external component to trust the active internal certificate authority, and import the certificate of the active internal certificate authority on the external component.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.
2) Browse to **Certificates** > **Certificate Authorities** > **Internal Certificate Authorities**.
3) Right-click the active internal certificate authority, then select **Properties**.
4) On the **Certificate** tab, click **Export**.
5) Save the certificate.
6) Click **Cancel** to close the properties of the internal certificate authority.

Next steps

Configure the external component to trust the internal certificate authority, and import the certificate of the active internal certificate authority on the external component.

Create TLS Cryptography Suite Set elements

TLS Cryptography Suite Set elements define which cryptographic algorithms are allowed for encrypting TLS traffic.

Note

The options in TLS Cryptography Suite Set elements do not apply to TLS 1.3. By default, all supported cryptographic algorithms are enabled for TLS 1.3.

The default NIST (SP 800-52) Compatible SSL Cryptographic Algorithms element allows SSL cryptographic algorithms that are compatible with the following standard: *NIST SP 800-52 Rev. 1 Guidelines for the Selection, Configuration, and Use of Transport Layer Security (TLS) Implementations*. If the default cryptographic algorithms meet your needs, there is no need to create a custom TLS Cryptography Suite Set element.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.
2) Browse to **Certificates** > **Other Elements** > **TLS Cryptography Suite Sets**.
3) Right-click **TLS Cryptography Suite Sets**, then select **New TLS Cryptography Suite Set**.
4) In the **Name** field, enter a unique name.

5) Select one or more cryptographic algorithms.
   - Algorithms in the **Common** section are compatible with SSL 3.0, TLS 1.0, TLS 1.1, and TLS 1.2.
   - Algorithms in the **TLS 1.2 Only** section are only compatible with TLS 1.2.

6) Click **OK**.

Create TLS Profile elements

TLS Profile elements define the settings for cryptography, trusted certificate authorities, and the TLS version used in TLS-protected traffic.

You can use TLS Profile elements for the following purposes:

- Enabling TLS-protected audit or log data forwarding to an external syslog server
- Enabling TLS encryption for LDAP connections between the Security Engine and external LDAP or Active Directory servers
- Defining the TLS settings for HTTPS connections for browser-based user authentication
- Defining the trusted certificate authority for client certificate authentication for browser-based user authentication
- Authenticating connections between the Security Engine and the server on which Forcepoint User ID Service has been installed

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.

2) Browse to **Certificates** > **Other Elements** > **TLS Profiles**

3) Right-click **TLS Profiles**, then select **New TLS Profile**.

4) In the **Name** field, enter a unique name for the TLS Profile.

5) Click **Select** next to the **TLS Cryptography Suite Set** field, then select a TLS Cryptography Suite Set element.

6) Select the trusted Certificate Authorities.
   - Select **Trust Any** if you want to allow the use of any valid certificate authority.
   - Select **Trust Selected**, then click **Add** to specify the trusted Certificate Authorities.

7) Configure the other settings as needed.

8) Click **OK**.

Configure TLS server identity

TLS server identity determines how SMC servers or Security Engines verify the identity of the external servers with which they communicate.

You can configure TLS server identity in the following elements:

- Management Servers and Log Servers — Defines how the identity of the syslog server to which log data is forwarded from the Management Server or the Log Server is verified.
- Active Directory Server or LDAP Server — Defines how the identity of the Active Directory Server or LDAP Server is verified when the LDAPS or Start TLS protocols is used to secure the LDAP connection between the external server and the Management Server and Security Engines.
- Forcepoint User ID Service — Defines how the identity of the Forcepoint User ID Service that sends user identification information to the Security Engines is verified.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click the element for which you want to define the TLS server identity, then select **Properties**.

2) Define the TLS server identity depending on the type of element.
   - Management Server or Log Server — Click the **Audit Forwarding** or **Log Forwarding** tab, double-click the **TLS Server Identity** cell, then define the TLS server identity settings.
   - Active Directory Server or LDAP Server — On the **General** tab, select **LDAPS** or **Start TLS** from the **LDAP Protocol** drop-down list.
   - Forcepoint User ID Service — Click the **Certificate** tab.

3) From the **TLS Server Identity** drop-down list, select the server identity type field to be used.

4) (Optional) Click **Fetch from Certificate** to fetch the value of the server identity type field from a certificate.

   Note

   You can fetch the value of the server identity field from a certificate only if the server identity field is **Distinguished Name**, **SHA-1**, **SHA-256**, **SHA-512**, or **MD5**.

5) In the **Identity Value** field, enter the value of the server identity field.

6) Click **OK**.

Creating certificates

You can generate certificates in the SMC, then sign the certificate request with tools in the SMC or with an external certificate authority.

TLS Credentials elements represent both certificate requests and signed certificates in the SMC Client. When a certificate request has been signed, the TLS Credentials element represents a certificate. In the Configuration view, the **State** column for the TLS Credentials element shows whether the element represents a certificate request or a signed certificate.

There are three ways to sign certificate requests:

- Self-sign the certificate request.
- Sign the certificate request with the Management Server's internal certificate authority.
- Export the certificate request, sign the certificate request with an external certificate authority, then import the signed certificate.

TLS Credentials elements that represent signed certificates can be used in the properties of several types of elements to secure connections involving those elements.

Types of elements where TLS Credentials elements can be used

Element	Purpose
Web Access Server	The certificate is used to secure the server's connections using HTTPS.
Management Server	The certificate is used to secure communications between the SMC API client and the Management Server.
SSL VPN Portal	The private key and certificate are used to establish SSL connections to the SSL VPN Portal.
SSL VPN Portal Service	The private key and the certificate are used in DNS Mapping to establish SSL connections to the service.

Create a certificate request

To create a certificate request, you must create a TLS Credentials element.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.
2) Browse to **Certificates** > **TLS Credentials**.
3) Right-click **TLS Credentials**, then select **New TLS Credentials**.
4) Complete the certificate request details.
   a) Enter a name for the certificate.
   b) In the **Common Name** field, enter the IP address or domain name of the server.
   c) Complete the remaining fields as needed.
   d) Click **Next**.
5) Sign the certificate request or finish creating the certificate request.
   - To create a self-signed certificate, select **Self-Sign**, then click **Finish**.
   - To create a certificate signed by the Management Server's internal certificate authority, select **Sign with Internal CA**, then click **Finish**.

- To sign the certificate request with an external certificate authority select **Sign with External CA**, then click **Finish**.

Next steps

If you want to sign the certificate request with an external certificate authority, export the certificate request.

Export a certificate request

If you want to sign a certificate request with an external certificate authority, you must export the certificate request.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.
2) Select **Certificates** > **TLS Credentials**.
3) Right-click the certificate request, then select **Export Certificate Request**.
4) Copy or export the certificate request.
   - Copy the certificate request, then paste it in an external application to sign it externally.

     Note

     If you copy and paste the certificate request, include the “Begin Certificate Request” header and the “End Certificate Request” footer.

   - Click **Export**, browse to the location where you want to save the certificate request, then click **Save**.
5) Click **OK** to close the Export Certificate Request dialog box.

Next steps

Sign the certificate request in an external application, then import the signed certificate request into the SMC.

Import an externally signed certificate

If you signed a certificate request with an external certificate authority, you must import the signed certificate into the SMC.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.

2) Select **Certificates** > **TLS Credentials**.

3) Right-click the certificate request that has been signed by an external certificate authority, then select **Import Signed Certificate**.

4) Select one of the following options to import the signed certificate.
   - Select **From File**, then browse to the signed certificate file on your local workstation.
   - Select **As Text**, then copy and paste the content of the signed certificate into the dialog box.

   Note

   If you copy and paste the content of the signed certificate, include the “Begin Certificate Request” header and the “End Certificate Request” footer.

5) Click **OK**.

Result

The status information in the **State** column shows that the certificate request has now been signed.

Configure settings for certificate validation

Certificate validation settings allow you to define the settings that the Security Engine uses when it connects to a Certificate Revocation List (CRL) or Online Certificate Status Protocol (OCSP) server.

The Security Engine validates certificates and checks the certificate revocation status for features that have certificate validation and certificate revocation checks enabled, such as features that use a TLS Profile in the configuration.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Engine, IPS, or Layer 2 Engine element, then select **Edit <element type>**.

2) Browse to **Advanced Settings** > **Certificate Validation**.

3) (Optional) If the Security Engine cannot access external networks directly, select the HTTP proxy through which OCSP and CRL lookups are sent.

4) (Optional) Enter the timeout for communication from the Security Engine to the CRL or OSCP server.
   The default timeout is 120 seconds.

5) Click **Save and Refresh** to transfer the configuration changes.

Renewing certificates

You must renew certificates and certificate authorities when they expire.

All certificates have a validity start date (“not before”) and a validity end date (“not after”). In the SMC, internally generated certificates are valid for three years from their creation.

The SMC’s internal Certificate Authorities are valid for 10 years. A new internal RSA CA or a new internal ECDSA CA is automatically created six months before the expiration date. Components that use certificates signed by the internal CAs must receive new certificates that have been signed by the new internal CAs.

When the system has created a new internal CA, SMC components gradually start using the new internal CA to sign certificates. Initially, the new internal CA is in the **Ready to Use** state, and only Management Server certificates are signed by the new internal CA. Certificates for other components are signed by the internal CA that is used by the Management Server. In an environment with multiple Management Servers, the new internal CA changes to the “Active” state when all Management Servers are using the new internal CA.

Each component must receive a new certificate signed by the new internal CA. The SMC automatically creates new certificates for Security Engines. For other components, you must always manually create new certificates. If the automatic certificate creation fails, you must create new certificates manually for Security Engines.

Check the expiration date of internal certificates or CAs

You can check the status of internal certificates used in system communications and the status of the internal certificate authority that automatically signs the internal certificates.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.

2) Browse to **Certificates** > **Internal Certificates** or **Certificates** > **Certificate Authorities** > **Internal Certificate Authorities**.

3) Check the **Expiration Date** column.



Tip

To view detailed information, right-click a certificate or Internal Certificate Authority element, then select **Properties**.

Recertify SMC servers

You must renew SMC server certificates when the certificates are about to expire or have expired.

The following situations require you to renew SMC server certificates:

- A message indicates that the certificate of a Management Server, Log Server, or Web Access Server is about to expire or has expired.

- A message indicates that the certificate authority that signed the certificate of a Management Server, Log Server, or Web Access Server is about to expire. A new certificate authority has been created, and the server requires a new certificate.
- The SMC components refuse communication attempts with each other.

If the Management Server certificate expires, it is not possible to log on using the SMC Client. Log Server certificate expiration or loss prevents log browsing, reporting, and status monitoring from working correctly, and forces the engines to spool logs locally.

You can renew the certificates of any of the SMC servers without affecting the other components.

When administrators log on to the SMC Client or to the Web Portal for the first time after the server's certificate is changed, they receive a notification of the certificate fingerprint change on the Management Server. If you want to check the certificate fingerprint before accepting it, run the `sgShowFingerprint` command on the server.

Steps

1) Stop the SMC server you want to recertify.



Note

To certify a Log Server or a Web Access Server, the Management Server must be running and accessible through the network.

2) On the command line of the server that you want to certify, go to the `<installation directory>/bin` folder.



Note

If you installed the SMC in the `C:\Program Files\Forcepoint\SMC` directory in Windows, command-line scripts can be found in the `C:\Program Files\Forcepoint\SDWAN Manager\bin` directory.

3) To recertify a Management Server, run the following script:

```
sgCertifyMgtSrv.[bat|sh]
```

4) To certify an additional Management Server, follow these steps.

   a) Verify that the active Management Server is running and that the additional Management Server has a connection to the active Management Server.

   b) Stop the additional Management Server.

   c) Run the following script on the additional Management Server:

```
sgCertifyMgtSrv.[bat|sh] -standby
```

5) To recertify a Log Server, run the following script:

```
sgCertifyLogSrv.[bat|sh]
```

6) To recertify a Web Access Server, run the following script:

```
sgCertifyWebAccessSrv.[bat|sh]
```

7) If prompted in the recertification dialog box, authenticate using an SMC administrator account with unrestricted (superuser) permissions.



Note

Do not enter the credentials for the root account for command line access.

If there are multiple administrative Domains, you can also specify the **Domain** the Log Server or the Web Access Server belongs to. If you do not specify the Domain, the Shared Domain is used.

8) Make sure that the **Recertify an Existing Server** option is selected, and that the correct server is selected in the list.

9) Click **OK**, then wait for confirmation that the server certificate has been renewed.

10) Start the SMC server that you recertified.

When you restart the server, all other components accept the new certificate because it is issued by a certificate authority that they trust. SMC components only trust the internal certificate authority that issued their own certificate.

Related reference

Forcepoint Security Management Center commands on page 1429

Renew Security Engine certificates

Security Engine certificates are renewed automatically. You might have to renew Security Engine certificates manually in some cases.

The following situations might require you to manually renew Security Engine certificates:

- A message indicates that the certificate for an Security Engine has expired.
- A message indicates that the certificate authority that signed the component's certificate is about to expire or has expired. A new certificate authority has been created, and the engine requires a new certificate.
- Components refuse connection attempts with each other.
- You have created an ECDSA CA and the engine has lost connectivity to the Management Server. You might also have to manually enable 256-bit security strength for the engine.

If the certificate for system communications expires, the Security Engines continue processing traffic normally but all communications with other components stop. For clusters, traffic might be disrupted if expired certificates prevent nodes from synchronizing information. The same disruption can also happen if the internal certificate authority that signs the certificates for system communications is in the process of being renewed, and Security Engines do not have new certificates signed by the new internal certificate authority that the system has automatically created.

Security Engine certificates might expire if you have disabled automatic certificate renewal.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the SMC Client, save the initial configuration and generate a new one-time password for the Security Engine.

2) To renew contact between the engine and the Management Server using the new one-time password, run the following command on the command line of the Security Engine:

```
sg-reconfigure
```

3) Follow the prompts in the Security Engine Configuration Wizard until the **Prepare for Management Contact** page opens.

4) Select **Contact**, then press the spacebar.

5) Enter the Management Server IP address and the one-time password.

6) Highlight **Finish**, then press **Enter**.

Related tasks

Save the initial configuration and generate the one-time password on page 632
Reconfigure Security Engine settings on page 365

Related reference

Security Engine commands on page 1445

Renew certificates for SMC components and Security Engines when certificate authorities expire

If a certificate authority is about to expire, the components that use certificates signed by the certificate authority require new certificates that are signed by a valid certificate authority.

Messages in the SMC Client about expiring certificate authorities indicate that a certificate authority is about to expire, a new certificate authority has been automatically created, or a certificate authority has expired.

You might need to renew certificates for SMC components and Security Engines in the following cases:

- The certificate authority that signed the certificate of a component is about to expire.
- A certificate authority has been automatically renewed, and a new certificate must be generated for the component.
- Components refuse connection attempts with each other.
- Automatic certificate renewal for Security Engines fails.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Re-certify the SMC servers.
2) To use the new certificate on Security Engines after automatic certificate renewal, refresh the policy.
3) If the automatic certificate renewal for Security Engines fails, renew the Security Engine certificates manually.

Related tasks

Recertify SMC servers on page 162
Renew Security Engine certificates on page 164

Renewing external certificates

You must renew the external certificates of SMC components before they expire.

All certificates have a validity start date (“not before”) and a validity end date (“not after”).

Pre-requisites:

You must have the external Certificate Authority (CA) configured.

Note

- This is the certificate authority that is used to sign certificate requests.
- SMC does not control the certificate validity time for externally signed certificates. The certificate validity time is dependent on the external CA configuration.

To renew an externally signed certificate, do the following steps:

1) Check the expiration date of the external certificate.
2) Renew the external certificate for the following components:
   - Management server. For more details, refer to the *Renew external certificate for the Management Server* topic.
   - Log server. For more details, refer to the *Renew external certificate for the Log Server* topic.
   - Engine. For more details, refer to the *Renew external certificates for the Engine* topic.
   - Web Access Server. For more details, refer to the *Renew external certificates for the Web Access Server* topic.

Check the expiration date of the external certificate

You can check the status of external certificates that are used in system communications.

Steps

1) Select **Administration**.

2) Browse to **Certificates** > **Internal Certificates**.

3) Check the **Expiration Date** column.

Note

To view detailed information, right-click a certificate or Internal Certificate Authority element, then select **Properties**.

Note

You can also view the certificate details for a server from the server properties. Click the certificate hyperlink in the **Certificate** tab of the server properties. The **Certificate** tab is only displayed if the **Use External Certificate Authority** feature is enabled during the SMC installation.

Renew external certificate for the Log Server

You must renew the external certificates for the Log server manually before it expires.

Before you begin

- You must have the external Certificate Authority (CA) configured.

  Note

  This is the certificate authority that is used to sign certificate requests.

- The Log Server must be running before generating the new Certificate Signing Request (CSR).

Steps

1) Select **Dashboard** > **Servers / Devices Dashboard**.

2) Browser to **Log Server**.

3) Generate the certificate request by using one of the following ways:

- From the context menu of the Log Server.
  a) Right-click the Log Server, then select **Certificate** > **Renew Certificate**.
  b) Click the **Yes** button, and then click the **OK** button. The certificate request is generated.
- From the Log Server properties dialog-box:

  Note

  Use this option if the existing settings in the **Certificate Definition** section need to be updated.

  a) Right-click the Log Server, then select **Properties**.
  b) Click the **Certificate** tab.
  c) Click the **Renew Certificate** button.
  d) Configure the settings in the **Certificate Definition** section.
  e) Click the **Generate Certificate Request** button, and then click the **Yes** button.
  f) Click the **OK** button. The **Certificate Request** section is displayed.

4) Export the generated certificate request by using one of the following ways:

- From the Log Server properties dialog-box:
  a) Right-click the Log Server, then select **Properties**.
  b) Click the **Certificate** tab.
  c) In the **Certificate Request** section, click the **Export Certificate Request** button.
  d) Navigate to the desired location and click the **Export** button, and then click the **OK** button.
- From the context menu of the Log Server:
  a) Right-click the Log Server, then select **Certificate** > **Export Certificate Request**.
  b) Navigate to the desired location and click the **Export** button, and then click the **OK** button.

5) Sign the exported certificate request by using the external Certificate Authority (CA).

6) Import the signed certificate by using one of the following ways:

- From the Log Servers properties dialog-box:
  a) Right-click the Log Server, then select Properties.
  b) Click the Certificate tab.
  c) In the **Certificate Request** section, click the **Import Signed Certificate** button.
  d) Select one of the following options:
     - The **From File** option:
       i) Select the **From File** radio button, and then click the **Browse** button.
       ii) Navigate the location where the signed certificate is saved.
       iii) Select the signed certificate file, and then click the **Import** button.
       iv) Click the **OK** button.
     - The **As Text** option:
       i) Select the **As Text** radio button.
       ii) Paste the certificate details.
       iii) Click the **OK** button.

- From the context menu of the Log Server:
  a) Right-click the Log, then select **Certificate** > **Import Certificate**.
  b) Select one of the following options:
     - The **From File** option:
       i) Select the **From File** radio button, and then click the **Browse** button.
       ii) Navigate to the location where the signed certificate is saved.
       iii) Select the signed certificate file, and then click the **Import** button.
       iv) Click the **OK** button.
     - The **As Text** option:
       i) Select the **As Text** radio button.
       ii) Paste the certificate details.
       iii) Click the **OK** button.

7) Stop the Log Server.

8) Execute the following script:

```
sgCertifyLogSrv
```

9) Restart the Log Server.

10) Verify the expiration date of the renewed certificate. For more details, refer to the **Check the expiration date of the certificate** topic.

Renew external certificate for the Management Server

You must renew the external certificates for the Management server manually before it expires.

Before you begin

You must have the external Certificate Authority (CA) configured.

Note

This is the certificate authority that is used to sign certificate requests.

Steps

1) Select **Dashboard** > **Servers / Devices Dashboard**.

2) Browser to **Management Server**.

3) Generate the certificate request by using one of the following ways:

- From the context menu of the Management Server.
  a) Right-click the Management Server, then select **Certificate** > **Renew Certificate**.
  b) Click the **Yes** button, and then click the **OK** button. The certificate request is generated.
- From the Management Server properties dialog-box:

  Note

  Use this option if the existing settings in the **Certificate Definition** section need to be updated.

  a) Right-click the Management Server, then select **Properties**.
  b) Click the **Certificate** tab.
  c) Click the **Renew Certificate** button.
  d) Configure the settings in the **Certificate Definition** section.
  e) Click the **Generate Certificate Request** button, and then click the **Yes** button.
  f) Click the **OK** button. The **Certificate Request** section is displayed.

4) Export the generated certificate request by using one of the following ways:

- From the Management Server properties dialog-box:
  a) Right-click the Management Server, then select **Properties**.
  b) Click the **Certificate** tab.
  c) In the **Certificate Request** section, click the **Export Certificate Request** button.
  d) Navigate to the desired location and click the **Export** button, and then click the **OK** button.
- From the context menu of the Management Server:
  a) Right-click the Management Server, then select **Certificate** > **Export Certificate Request**.
  b) Navigate to the desired location and click the **Export** button, and then click the **OK** button.

5) Sign the exported certificate request by using the external Certificate Authority (CA).

6) Import the signed certificate by using one of the following ways:

- From the Management Servers properties dialog-box:
  a) Right-click the Management Server, then select Properties.
  b) Click the **Certificate** tab.
  c) In the **Certificate Request** section, click the **Import Signed Certificate** button.
  d) Select one of the following options:
     - The **From File** option:
       i) Select the **From File** radio button, and then click the **Browse** button.
       ii) Navigate the location where the signed certificate is saved.
       iii) Select the signed certificate file, and then click the **Import** button.
       iv) Click the **OK** button.
     - The **As Text** option:
       i) Select the **As Text** radio button.
       ii) Paste the certificate details.
       iii) Click the **OK** button.
- From the context menu of the Management Server:
  a) Right-click the Management, then select **Certificate** > **Import Certificate**.
  b) Select one of the following options:
     - The **From File** option:
       i) Select the **From File** radio button, and then click the **Browse** button.
       ii) Navigate to the location where the signed certificate is saved.
       iii) Select the signed certificate file, and then click the **Import** button.
       iv) Click the **OK** button.
     - The **As Text** option:
       i) Select the **As Text** radio button.
       ii) Paste the certificate details.
       iii) Click the **OK** button.

7) Stop the Management Server.

8) Execute the following script with the option mode=ext-pki-renew:

```
sgCertifyMgtSrv
```

9) Restart the Management Server.

10) Verify the expiration date of the renewed certificate. For more details, refer to the **Check the expiration date of the certificate** topic.

Renew external certificate for the Web Access Server

You must renew the external certificates for the Web Access Server manually before it expires.

Before you begin

- You must have the external Certificate Authority (CA) configured.

  Note

  This is the certificate authority that is used to sign certificate requests.

- The Web Access Server must be running before generating the new Certificate Signing Request (CSR).

Steps

1) Select **Dashboard** > **Servers / Devices Dashboard**.

2) Browser to **Web Access Server**.

3) Generate the certificate request by using one of the following ways:

- From the context menu of the Web Access Server.
  a) Right-click the Web Access Server, then select **Certificate** > **Renew Certificate**.
  b) Click the **Yes** button, and then click the **OK** button. The certificate request is generated.
- From the Web Access Server properties dialog-box:

> **Note**
>
> Use this option if the existing settings in the **Certificate Definition** section need to be updated.

  a) Right-click the Web Access Server, then select **Properties**.
  b) Click the **Certificate** tab.
  c) Click the **Renew Certificate** button.
  d) Configure the settings in the **Certificate Definition** section.
  e) Click the **Generate Certificate Request** button, and then click the **Yes** button.
  f) Click the **OK** button. The **Certificate Request** section is displayed.

4) Export the generated certificate request by using one of the following ways:

- From the Web Access Server properties dialog-box:
  a) Right-click the Web Access Server, then select **Properties**.
  b) Click the **Certificate** tab.
  c) In the **Certificate Request** section, click the **Export Certificate Request** button.
  d) Navigate to the desired location and click the **Export** button, and then click the **OK** button.
- From the context menu of the Web Access Server:
  a) Right-click the Web Access Server, then select **Certificate** > **Export Certificate Request**.
  b) Navigate to the desired location and click the **Export** button, and then click the **OK** button.

5) Sign the exported certificate request by using the external Certificate Authority (CA).

6) Import the signed certificate by using one of the following ways:

- From the Web Access Servers properties dialog-box:
  a) Right-click the Web Access Server, then select Properties.
  b) Click the **Certificate** tab.
  c) In the **Certificate Request** section, click the **Import Signed Certificate** button.
  d) Select one of the following options:
     - The **From File** option:
       i) Select the **From File** radio button, and then click the **Browse** button.
       ii) Navigate the location where the signed certificate is saved.
       iii) Select the signed certificate file, and then click the **Import** button.
       iv) Click the **OK** button.
     - The **As Text** option:
       i) Select the **As Text** radio button.
       ii) Paste the certificate details.
       iii) Click the **OK** button.

- From the context menu of the Web Access Server:
  a) Right-click the Web Access, then select **Certificate** > **Import Certificate**.
  b) Select one of the following options:
     - The **From File** option:
       i) Select the **From File** radio button, and then click the **Browse** button.
       ii) Navigate to the location where the signed certificate is saved.
       iii) Select the signed certificate file, and then click the **Import** button.
       iv) Click the **OK** button.
     - The **As Text** option:
       i) Select the **As Text** radio button.
       ii) Paste the certificate details.
       iii) Click the **OK** button.

7) Stop the Web Access Server.

8) Execute the following script:

```
sgCertifyWebAccessSrv
```

9) Restart the Web Access Server.

10) Verify the expiration date of the renewed certificate. For more details, refer to the **Check the expiration date of the certificate** topic.

Renew external certificate for the Engine

You must renew the external certificates for the Engine manually before it expires.

Before you begin

You must have the external Certificate Authority (CA) configured.

Note

This is the certificate authority that is used to sign certificate requests.

Steps

1) Select **Engine Configuration**.

2) Browser to **Engines**.

3) Generate the certificate request:
   a) Right-click the Engine, then select **Certificate** > **Renew Certificate**.
   b) Click the **Yes** button, and then click the **OK** button. The certificate request is generated.

4) Export the generated certificate request:
   a) Right-click the Engine, then select **Certificate** > **Export Certificate Request**.
   b) Navigate to the desired location and click the **Export** button, and then click the **OK** button.

5) Sign the exported certificate request by using the external Certificate Authority (CA).

6) Import the signed certificate:

   a) Right-click the Engine, then select **Certificate** > **Import Certificate**.

   b) Select one of the following options:

   - The **From File** option:

     i) Select the **From File** radio button, and then click the **Browse** button.

     ii) Navigate to the location where the signed certificate is saved.

     iii) Select the signed certificate file, and then click the **Import** button.

     iv) Click the **OK** button.

   - The **As Text** option:

     i) Select the **As Text** radio button.

     ii) Paste the certificate details.

     iii) Click the **OK** button.

7) Verify the expiration date of the renewed certificate. For more details, refer to the **Check the expiration date of the certificate** topic.

8) Refresh the policy for the engine. For more details, refer to the **Refresh the currently installed policy** topic.

Chapter 11

Managing elements

Contents

- Introduction to elements on page 179
- Benefits of exporting or importing elements on page 189
- Restore elements from Snapshots on page 193
- Lock elements on page 197
- Unlock elements on page 198
- How the Trash works on page 198
- How Categories help you view only certain elements on page 201
- Legacy elements and options on page 207

Certain tasks are common to most elements. Some of these tasks are not mandatory for defining an element, but are still helpful as you get your SMC up and running.

Introduction to elements

Apart from a few minor exceptions, all configurations are created in the SMC, where information is stored as reusable *elements*.

For example, the Security Engines, traffic inspection policies, IP addresses, log filters, backups, and the licenses for the system components are all displayed as elements.

Different element types are provided for different concepts. The elements in the system define information both for adjusting the traffic inspection policies and for managing the system. This chapter gives you a brief description of each type of element.

Elements used in the configuration of Security Engines

You can view the types of elements used for configuring Security Engines.

Types of elements in Security Engine configuration

Element Type	Explanation
Security Engines	Configurations particular to individual Security Engines, such as interface configurations.
Policies	The rules for inspecting and handling network traffic.
Network Elements	Represent IP addresses.

Element Type			Explanation
Other Elements	Endpoint Information		Endpoint Application and Endpoint Settings elements can be used for matching in Access rules. The elements can be used to identify applications used on endpoint clients, and also determine the operating system or status of the local anti-virus or engine.
	Ethernet Services		Definitions for protocols that can be used for traffic filtering on the Ethernet level.
	Event Bindings		Sets of log events that can be used in Correlation Situations to bind together different types of events in traffic.
	File Types		Elements that represent different types of files that can be allowed or blocked in Access rules.
	HTTPS Inspection Exceptions		Lists of domains that can be used to exclude some traffic from HTTPS decryption and inspection.
	Logical Interfaces		Interface reference that can combine several physical interfaces into one logical entity. Used for defining traffic handling rules.
	MAC Addresses		Represent MAC addresses in Ethernet-level traffic filtering.
	Network Applications		Provide a way to dynamically identify traffic patterns related to the use of a particular application.
	Policy Snapshots		Saved versions of the Security Engine configurations. Created each time you install or refresh a policy on an Security Engine.
	Protocols		Supported network protocols. Can be used to define new Services for matching traffic in policies. You cannot add, delete, or change the Protocol elements.
	Services		Network protocols and ports.
	Situations		Patterns that deep inspection looks for in traffic.
	TLS Matches		Define matching criteria for the use of the TLS (transport layer security) protocol in traffic, and specify whether TLS traffic is decrypted for inspection.
	Vulnerabilities		References that link some Situations to publicly available databases of known vulnerabilities in various software.
	Dynamic Routing Elements		Elements and Access Lists used for configuring dynamic routing. For more information on elements used in configuring dynamic routing, see the chapter about dynamic routing.
	Engine Properties	DNS Relay Profiles	Define the host name mappings, domain-specific DNS servers, fixed domain answers, and DNS answer translations that the engine uses when it provides DNS services to the internal network.

Element Type			Explanation
		Sandbox Services	Define the settings for connecting to a sandbox server for Forcepoint Advanced Malware Detection.
		SNMP agents	Configuration information for sending SNMP traps to external components about system events related to Security Engines.
		User Identification Services	Elements for the Forcepoint User ID Service or the Integrated User ID Service that associate IP addresses with users. The Integrated User ID Service is primarily meant for demonstration purposes and proof-of-concept testing of user identification services. **Note** For Forcepoint Network Security Platform version 6.4 or higher, we recommend that you use the Forcepoint User ID Service.
		User Authentication Pages	Define the look of the logon page, challenge page, and different status pages shown to end users who authenticate through a web browser.
		User Responses	Settings for notifying end users about different policy actions.
		ECA Configurations	Define settings for securing connections between endpoint clients and the Security Engine.

Related concepts
Network elements on page 182
Service elements on page 187
Situation elements on page 188

VPN elements

You can view the types of elements used for configuring VPNs, and inbound and outbound traffic management.

Types of elements in the VPN configuration

Element Type	Explanation
Policy-Based VPNs	Configurations particular to a VPN between two or more VPN gateways.
Route-based Tunnels	Configurations particular to VPN tunnels between Engine interfaces that are designated as tunnel endpoints.
Traffic Handlers	Configure outbound and inbound traffic management features (load balancing and high availability).

Element Type			Explanation
SSL VPN Portal	SSL VPN Portals		Configurations particular to Engines that act as SSL VPN Portals.
	SSL VPN Portal Policies		SSL VPN Portal Policy elements that define which services are available in the SSL VPN Portal and which users can access the services.
	SSL VPN Portal Services		SSL VPN Portal Services that are used through the SSL VPN Portal.
	SSL VPN SSO Domains		SSL VPN SSO Domains where the same user name and password are valid.
	SSL VPN Portal Service Profiles		Defines the settings for SSO and cookie protection in the SSL VPN Portal Services.
VPN Gateways			Configurations particular to individual VPN Gateways, such as IP address information.
Other Elements	Connection Types		Connection type elements that define how VPN endpoints are used in a Multi-Link configuration.
	Tunnel Groups		Tunnel Group elements for monitoring Route-based Tunnels.
	Profiles	Gateway Profiles	Information about the capabilities of particular types and versions of VPN gateway devices. Allow automatic configuration validation.
		Gateway Settings	Advanced global Engine settings related to VPN performance tuning.
		VPN Profiles	The main authentication, encryption, and integrity checking settings for VPNs.
	QoS Classes		An identifier that can be assigned to network traffic to define QoS policies for the traffic.
	TLS Cryptography Suite Sets		Settings that define which cryptographic algorithms are allowed to be used in the SSL VPN.
	VPN Certificates	Gateway Certificates	Certificate issuers whose signature is accepted as proof of identity on certificates in one or more VPNs.
		VPN Certificate Authorities	Certificate issuers whose signature is accepted as proof of identity on certificates in one or more VPNs.

Network elements

Network elements represent included IP addresses in Security Engine configurations.

Types of network elements

Element Type	Explanation
Address Ranges	Define a set of consecutive IP addresses between a start address and end address that you define.

Element Type	Explanation
Aliases	Context-dependent elements with no fixed value. The value is defined per engine and is determined when a policy containing the Alias is installed.
Countries	Country elements contain lists of IP addresses registered in a particular country.
Domain Names	The name of an Internet domain that is automatically resolved by an Security Engine to all IP addresses associated with the domain.
Expressions	Allow defining IP addresses using logical operators, which simplify the definition of complex sets of addresses.
Groups	Allow collecting together other Network Elements of any type. Represents all IP addresses defined in the included elements.
Hosts	Represent a single IP address.
IP Address Lists	Represent a list of IP addresses.
Networks	Represent a complete (sub)network of addresses.
Security Engines	Configurations particular to individual Security Engines, such as the interface configuration.
Routers	Represent a next-hop router in configurations where required. In policies, represent a single IPv4 or IPv6 address.
Servers	Represent an SMC server or an external server that provides a service to the system. In policies, represent a single IP address.
Zones	Interface reference that can combine several network interfaces of Security Engines into one logical entity. Used for defining interface matching requirements in traffic handling rules in policies.

Related concepts
Introduction to expressions on page 1465

System administration elements

You can view the types of elements used for system administration and their descriptions.

Types of elements for system administration

Element Type		Explanation
Access Rights	Access Control Lists	Sets of elements that you can grant to one or more administrator accounts when assigning administrator rights.
	Administrator Roles	Sets of actions that administrators are allowed to carry out both globally and specifically on some set of elements.
	Administrators	SMC administrator accounts.
	API Clients	Accounts for the users of the SMC Application Programming Interface (API).
	Web Portal Users	User accounts for the Web Portal.

Element Type			Explanation
Alert Configuration	Alert Chains		Lists of administrators and contact methods for escalating Alerts.
	Alert Policies		Rules for choosing which Alerts are escalated using which Alert Chain.
	Alert Senders		System components that can send Alerts.
	Alerts		Labels for Alerts that help in separating different Alerts from each other in Alert escalation.
	Policy Snapshots		Saved versions of the alert configuration. Created each time you install or refresh the Alert Policy on a Domain.
Bookmarks			User-created shortcuts to views in the SMC Client.
Certificates	Certificate Authorities	Client Protection Certificate Authorities	Certificates that are used in TLS inspection for client protection.
		ECA Evaluation Certificate Authorities	Certificates that are used in communications with ECA servers for evaluation of ECA.
		Internal Certificate Authorities	Certificates that are used in communications between the system components.
		Trusted Certificate Authorities	Certificates that identify certificate authorities that are trusted by the SMC and Security Engines.
		Trusted Update Certificates	Certificates that are used to verify the digital signatures of dynamic update packages and engine upgrades. A new Trusted Update Certificate is automatically added through a dynamic update package before the old one expires.
	Internal Certificates		Certificates that are used in communications between the system components.
	TLS Credentials		Represent both certificate requests and signed certificates in the SMC Client. When a certificate request has been signed, the TLS Credentials element represents a certificate. TLS Credentials elements that represent signed certificates can be used in the properties of several types of elements to secure connections involving those elements.
	Other Elements	TLS Cryptography Suite Sets	Define which cryptographic algorithms are allowed for encrypting TLS traffic.
		TLS Profiles	Define the settings for cryptography, trusted certificate authorities, and the TLS version used in TLS-protected traffic.
Engine Upgrades			Packages for remote upgrades that have been manually or automatically imported into the system.
Licenses			The components' licenses (proof of purchase).

Element Type		Explanation
Tasks	Definition	System maintenance Tasks and Task definitions.
	History	History of running and executed Tasks both started by users and generated by the system.
Trash		Stores elements that you have deleted. You can permanently delete elements that have been moved to the Trash.
Updates		Dynamic update packages that update definitions in your installation. Most of the content is Situations (used in deep packet inspection).
Other Elements	Backups	Management Server and Log Server backups.
	Categories	Allow filtering the view in the SMC Client to a subset of elements.
	Domains	Create boundaries for managing elements and configurations based on administrator configurations.
	Geolocations	Used for illustrating the geographical location of IP addresses (for example, in logs and diagrams).
	Locations	Used for defining contact addresses when NAT (IP address translation) is applied to communications between system components.
	Tools Profiles	Additional, user-configured commands and tools for components.

Monitoring elements

Use the monitoring elements to configure the monitoring features in the SMC.

Types of elements in monitoring

Element Type		Explanation
Diagrams		Allow you to visualize your network environment and monitor the system graphically.
Overviews		Customizable views for system status monitoring, statistics, and shortcuts to configuration tools.
Reports	Design	Define how log data and statistical data from engines are processed and displayed in reports.
	History	Saved statistical presentations of network traffic and the system.
	Sections	Define statistical items that are included in a report and the way that the items are displayed.
Third-Party Devices	Logging Profiles	Define the logging characteristics for a third-party device (what data from the third-party device logs is shown).

Element Type			Explanation
	MIBs		Allow you to import and browse management information bases to support third-party SNMP monitoring.
	Probing Profiles		Define how the Management Server tests if third-party devices are running.
Other Elements	Data contexts		Defines the log data types shown in the Logs view or the Reports view.
	Filters		Allow log data filtering in various tasks.
	Geolocations		Show where Hosts (for example, attackers) are on a world map and how much traffic they create.
	Monitoring Snapshots	Block List	Saved version of Block List entries for an Security Engine. Created when you save a Block List Snapshot in the Block List view.
		Connections	Saved version of connection entries for an Security Engine. Created when you save a Connections Snapshot in the Connections view.
		Logs	Saved version of log, alert, and audit entries for an Security Engine. Created when you save a Logs Snapshot in the Logs view.
		Routing	Saved version of routing entries for an Security Engine. Created when you save a Routing Snapshot in the Routing view.
		SSL VPNs	Saved version of active SSL VPN connections. Created when you save an SSL VPN snapshot in the SSL VPN Monitoring view.
		Users	Saved version of active users for an Security Engine. Created when you save a User Snapshot in the Users view.
		VPN SAs	Saved version of active VPN SAs for an Security Engine. Created when you save a VPN SA snapshot in the VPN SAs view.
	Network Elements		Represent IP addresses.
	Overview Templates		Templates that allow you to create a statistical overview with predefined information selected for the view.

Related concepts

Network elements on page 182

User authentication elements

You can view the types of elements used for configuring user authentication and directory services and the element descriptions.

Types of elements in user authentication configuration

Element Type		Explanation
Authentication Methods		Configured authentication methods for end-user and administrator authentication. Used in rules that require end-user authentication.
Servers		Active Directory Servers, LDAP Servers, RADIUS Authentication Servers, and TACACS+ Authentication Servers for end-user and administrator authentication and directory services.
Users		End users stored in the internal LDAP database or an external LDAP database. Used in rules that require end-user authentication.
Other Elements	SMTP Servers	SMTP servers that send email or SMS messages about changes to user accounts to end users. The same SMTP Servers can also be used to send Alerts to Administrators.

Service elements

Service elements are used in Access rules to match traffic and to set parameters for handling the traffic.

There are predefined system Service elements for official (IANA-reserved) and well-known protocols and services (such as DNS, FTP, and HTTP). You can also create your own custom Service elements to specify a port that is not predefined or to define custom options for handling some types of traffic.

Types of Services

Element Type	Explanation
Group	Groups of services containing the Service elements that together fulfill a certain role (for example, the services that allow IPsec VPN connections).
ICMP	Identifies the message by the ICMP Type and Code fields.
IP-proto	Identifies the protocol by the IP address header Protocol field.
SUN-RPC	Identifies the Sun remote procedure call (RPC) service by the program identifier.
TCP	Identifies the service by the TCP header *Source Port* or *Destination Port* fields.
UDP	Identifies the service by the UDP header *Source Port* or *Destination Port* fields.
With Protocol	Default Services containing Protocols that have default parameters set to typically used values.
With Proxy	Default Services containing SSM Proxy Protocol Agents.

Situation elements

Situation elements are used in Inspection rules to define patterns that deep packet inspection looks for in traffic.

The Situations tree is constructed differently compared to most other trees. The Situations tree contains several alternative groupings, so most Situations are shown in several places. The groupings allow you to easily find Situations that are specific to the task at hand. For example, Situations specific to the HTTP protocol (some of which are specific to particular web browsers) are stored at the following location in the Situations tree: **Situations** > **By Type** > **Traffic Identification** > **Browsers**.

Some branches are groupings that you can add to yourself. You can use most of these branches in Inspection rules. The Situation Type groupings are used as the basis for the tree-based Inspection rules configuration in Inspection Policy elements.

Situations and their groupings are updated in dynamic update packages. The following table lists the default branches at the time of writing this document.

Default groupings of Situations at the time of publishing this document

Tree branch		Explanation
All Situations		All Situations in the system without any grouping.
By Context	Anti-Malware	Events triggered in the malware scan.
	Correlations	Correlation Situations for detecting patterns in event data.
	DoS Detection	Situations for detecting DoS (denial-of-service) attacks.
	DXL	Legacy situations related to McAfee Threat Intelligence Exchange (TIE). McAfee Threat Intelligence Exchange (TIE) is no longer supported in Security Engine 6.10 and higher.
	Files	Situations based on identifying file types from traffic. Content identified based on file type fingerprints is redirected to appropriate file streams.
	Protocols	Situations that identify protocols from traffic.
	Scan Detection	Situations for detecting network scans.
	System	System-internal events.
By Tag	By Hardware	Situations that detect something specific to a particular hardware platform grouped by platform (for example, x86 (32-bit) or x86-64 (64-bit)). An example of something hardware specific is an attempt to exploit a known vulnerability that only exists on a particular platform.
	By Operating System	Situations that detect something specific to a particular operating system, grouped by operating system (for example, Windows (for all Windows versions) or Windows 2000).
	By Situation Tag	Free-form grouping for some special use cases. The Recent Updates branch is especially useful. The branches dynamically list Situations that have been recently added to the system in the 1–5 most recent dynamic update packages. (This list helps in tuning your policies.)
	By Software	Situations that detect something specific to a particular software, grouped by brand or product name (for example, Adobe Acrobat or Microsoft Office).
By Type		These Situations are shown as the main Rules tree in the Inspection rules.

Tree branch	Explanation
By Vulnerability	Situations that detect attempts to exploit known vulnerabilities grouped by vulnerability name.
Custom Situations	Custom Situations that the administrators create. Custom Situations can also appear in the other branches.

Benefits of exporting or importing elements

The ability to export and import most kinds of elements allows you to reuse or restore them without having to create them again.

- You can reuse elements in a different SMC.
- You can import old versions of elements or deleted elements by restoring them from a Policy Snapshot or from an Element Snapshot.
- You can restore elements that have been moved to the Trash.

You can export and import elements using the following interfaces:

- SMC Client
- Command Line Interface tools (sgExport or sgImport)

You can export elements to the following file types using the **Menu** > **File** > **Print** option:

- .html
- .pdf

Note

Exported files are meant for importing elements into the database of a Management Server. They are not meant to be viewed or edited in external applications.

You can import elements from the following file types:

- .csv — You can create a .csv file and import values from it.
- .tsv — You can create a .tsv file and import values from it.
- .zip — You can import elements from a .zip file of elements exported from the SMC Client.

Note

When you export and import elements that have been moved to the Trash, all references to the elements remain valid. An element that has been exported from the Trash remains in the Trash when imported to an environment with several Management Servers.

Related concepts

Restore elements from Snapshots on page 193

Related tasks

Save elements, log data, reports, and statistics on page 304
Export selected elements on page 190
Export all elements on page 191
Create .csv or .tsv files for importing elements on page 191
Restore elements from the Trash on page 199
Import user information on page 1123
Export user information on page 1123

Export selected elements

You can export most kinds of individual elements.

You cannot export some kinds of elements, such as administrator accounts and certificates. To export an element that references an element that cannot be exported, you must first manually create a corresponding element that has the same name as the referenced element. Otherwise, the export fails.

To protect sensitive data in exports, export .zip files are automatically encrypted in FIPS mode. If you have defined an export banner, the text of the banner is added at the beginning of each exported XML file to indicate that the export contains sensitive or classified data.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Menu** > **File** > **Export** > **Export Elements**.

2) Select an option:
   - Enter a file name for the export file.
   - Click **Browse** to select the location where you want to create the file.

3) (Optional) In the **Password** field, enter the password for the encrypted export file.

   Note

   In FIPS mode, you must enter a password.

4) (Optional) To view and export elements that have been moved to the Trash, select **More actions** > **Show Deleted Elements**.

5) Select the elements that you want to export, then click **Add**.

6) When you have finished adding elements to the Content list, click **Export**.
   A new tab opens to show the progress of the export.

Related tasks

Import elements from a file on page 193
Export user information on page 1123

Export all elements

You can export all elements as a group.

You cannot export some kinds of elements, such as administrator accounts and certificates. To export an element that references an element that cannot be exported, you must first manually create a corresponding element that has the same name as the referenced element. Otherwise, the export fails.

To protect sensitive data in exports, export .zip files are automatically encrypted in FIPS mode. If you have defined an export banner, the text of the banner is added at the beginning of each exported XML file to indicate that the export contains sensitive or classified data.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Menu** > **File** > **Export** > **Export All Elements**.

2) Select an option:
   - Enter a file name for the export file.
   - Click **Browse** to select the location where you want to create the file.

3) (Optional) In the **Password** field, enter the password for the encrypted export file.

   Note

   In FIPS mode, you must enter a password.

4) (Optional) To view and export elements that have been moved to the Trash, select **More actions** > **Show Deleted Elements**.

5) Click **Export**.
   A new tab opens to show the progress of the export.

Related tasks

Import elements from a file on page 193
Export user information on page 1123

Create .csv or .tsv files for importing elements

You can create .csv (comma-separated value) files or .tsv (tab-separated value) files for importing elements.

In a .csv file, commas separate all values. In a .tsv file, a tab character separates all values.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Create a new .csv or .tsv file. (For example, you can use a spreadsheet application.)

2) In the first row of the file, specify the header as follows:

- .csv files — Enter `ip,name,comment`.
- .tsv files — Enter `ip<tab>name<tab>comment`.

For example:

- .csv:
  - 10.10.10.10,host-10.10.10.10,Host abc
  - 10.20.30.10/24,net-10.20.30.40,Net xyz
- .tsv:
  - 10.10.10.10<tab>host-10.10.10.10<tab>Host abc
  - 10.20.30.10/24<tab>net-10.20.30.40<tab>Net xyz

Note

Only the IP address is mandatory in the header row. All data entered in the file must follow the format of the header row.

3) Enter the IP address of the element and optionally a name and a comment on the row below the header. Use the same format as in the header.

Note

Only the IP address is mandatory. If you have other parameters in the header row, enter a separator (a comma or tab) even if you do not enter a name or comment in the row.

Example: If the header row is `ip,name,comment` and you want to omit the name and the comment in the .csv file, enter `10.1.1.1,,`.

Example: If the header row is `ip<tab>name<tab>comment` and you want to omit the name and the comment in the .tsv file, enter `10.1.1.1<tab><tab>`.

If you omit the name, the SMC automatically generates a name for the element based on its IP address. The SMC detects the element type based on the syntax of the IP address as follows:

- 10.10.10.10 — Specifies a Host element.
- 10.10.10.0/24 — Specifies a Network element.
- 10.10.10.10–20.20.20.20 — Specifies an Address Range element.

4) (Optional) For each element, add another row to the file.

5) Save the file.

Related tasks

Import elements from a file on page 193

Import elements from a file

You can import elements from a .csv or .tsv file.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Menu** > **File** > **Import** > **Import Elements**.

2) Select the files you want to import, then click **Import**.
   The Management Server automatically checks if any elements in the file to be imported have the same name and XML content as any elements that exist. A new tab opens.

3) If any conflicts are found between elements in the import file and existing elements, select the **Action** for each conflict according to the conflict type.

   Tip

   If there is a conflict between existing elements and elements in the Policy Snapshot, the differences are shown in color in the XML format. To view the elements in XML format, select **More actions** > **Show XML**.

4) When there are no more conflicts, click **Continue** to start the import.

5) When the import is finished, click **Close**.

Related concepts
Benefits of exporting or importing elements on page 189

Restore elements from Snapshots

You can restore all elements from a Policy Snapshot or select the elements to be restored.

Related tasks
Restore all elements from Policy Snapshots on page 194
Restore selected elements from Policy Snapshots on page 195

Restore all elements from Policy Snapshots

You can restore all elements from a Policy Snapshot.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Browse to **Other Elements** > **Policy Snapshots**.
3) To open the list of Policy Snapshots, expand the Policy Snapshot type.
4) Restore elements in one of the following ways:
   - To restore all elements from a Policy Snapshot that is stored on the Management Server, right-click the Policy Snapshot from which you want to restore elements, then select **Restore**.
   - To restore all elements from a Policy Snapshot that you have backed up externally, click ⋮ **More actions**, select **Restore External Snapshots**, then browse to the Policy Snapshot backup that you want to restore.
5) If any conflicts are found between elements in the Policy Snapshot and the existing elements, resolve them by selecting the **Action**.

   Tip

   If there is a conflict between existing elements and elements in the Policy Snapshot, the differences are shown in color in the XML format. To view the elements in XML format, select ⋮ **More actions** > **Show XML**.
6) When there are no more conflicts, click **Continue**.
7) When the restoration is finished, click **Close**.

Related tasks

Restore selected elements from Policy Snapshots on page 195
Back up Policy Snapshots on page 196

Restore selected elements from Policy Snapshots

You can restore selected elements from a Policy Snapshot.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Browse to **Other Elements** > **Policy Snapshots**.
3) Expand the branch for the Policy Snapshot Type.
4) Right-click the Policy Snapshot from which you want to restore elements, then select **View Policy Snapshot**.
5) Select one or several elements to restore, right-click, then select **More actions** > **Restore**.
6) If any conflicts are found between the elements in the Policy Snapshot and the existing elements, resolve them by selecting the **Action**.

   Tip

   If there is a conflict between existing elements and elements in the Policy Snapshot, the differences are shown in color in the XML format. To view the elements in XML format, select ⋮ **More actions** > **Show XML**.

7) When there are no more conflicts, click **Continue** to start the restoring.
8) When the restoration is finished, click **Close**.

Related tasks

Restore all elements from Policy Snapshots on page 194

Restore elements from Element Snapshots

You can restore elements from Element Snapshots, which are stored in Audit logs.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Display Element Snapshots in the **Logs** view.

2) Right-click the audit entry of an element, then select **Compare to Current Element**.

Note

If the Element Snapshot properties differ from the properties of the existing element, a red border is displayed. You can view the red border around the Audit Log Version (snapshot) and the Current Version of the element.

Tip

If there is a conflict between existing elements and elements in the Policy Snapshot, the differences are shown in color in the XML format. To view the elements in XML format, select **Show: XML**.

3) To restore the properties of the Element Snapshot to the current element, click **Restore**.

4) If any conflicts are found between the elements in the Element Snapshot and existing elements, resolve them by selecting the **Action**:
   - **Import** — The element that exists is overwritten with the element in the Element Snapshot.
   - **Duplicate** — The element in the Element snapshot is renamed by adding a number to the end of the element's name and imported as a new element.
   - **Do not Import** — The element is not imported.

5) Click **Continue**.

6) When the import is done, click **Close**.

Related concepts

Restore elements from Snapshots on page 193

Related tasks

View and compare Element Snapshot elements on page 243

Back up Policy Snapshots

You can back up Policy Snapshot elements to store them externally.

You can restore all elements from Policy Snapshots that you have backed up externally in the same way as for Policy Snapshots that are stored on the Management Server.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Browse to **Other Elements** > **Policy Snapshots**.

3) To open the list of Policy Snapshots, expand the Policy Snapshot type.

4) Right-click the Policy Snapshot that you want to back up, then select **More actions** > **Backup**.

5) Browse to the location where you want to save the file, then click **Backup**.
The SMC Client shows the location where the backup is saved.

6) Click **OK**.
You are prompted to select whether to delete the Policy Snapshot element.

7) Click **Yes** or **No**.

Related tasks
Restore all elements from Policy Snapshots on page 194

Lock elements

An administrator who is allowed to edit an element can lock the element and add a comment to explain the reason for locking it.

Before you begin

To lock or unlock an element, you must be logged on to the Shared Domain or the Domain in which the element is stored.

Note

You cannot lock predefined system elements or elements that have been sent to the Trash.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click the element you want to lock, then select **More actions** > **Lock**.

2) Enter a **Comment** explaining the reason for locking the element.

3) Click **OK**.
The element is now locked and a lock symbol is displayed on its icon.

Unlock elements

Unlock an element so that you can edit or delete it.

Before you begin

To lock or unlock an element, you must be logged on to the Shared Domain or the Domain in which the element is stored.

Locked elements are displayed with a lock symbol. Unlock a locked element before editing or deleting it. The administrator who created the locked element or an administrator with unrestricted (superuser) permissions can unlock the element.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click the element you want to unlock, then select **More actions** > **Unlock**.

2) In the dialog box that opens, click **Yes** to confirm the unlock.

How the Trash works

Before deleting an element, you move it to the Trash.

In the SMC Client, you can view and search for elements in the Trash. An administrator with unrestricted permissions (superuser) can search for elements in the Trash in all administrative Domains.

Before working with the Trash feature, review the following considerations:

- If you want to delete an administrator account, first disable the account. Then you can delete the disabled Administrator element in the **Administration** branch of the **Configuration** view.
- Domains cannot be moved to the Trash. You can only permanently delete Domains.
- An element in the Trash is still valid in any previous configuration where the element was used before it was moved to the Trash. However, you cannot add an element that is in the Trash to any new configuration.
- When you export and import elements that have been moved to the Trash, all references to the elements remain valid. An element that has been exported from the Trash remains in the Trash when imported to an environment with several Management Servers.
- You can also restore elements that have been moved to the Trash. An element in the Trash is permanently deleted only when you delete it from the Trash or when you empty the Trash.

Related tasks

Delete elements from the Trash on page 200
Restore elements from the Trash on page 199
Disable administrator accounts on page 405
Delete a Domain on page 443

Move elements to the Trash

You move elements to the Trash before deleting them permanently.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click the element, then select **Delete**.
   A confirmation dialog box opens.

2) If the element you are moving to the Trash is used in any configuration, view and remove all references to the element:
   a) Click **Open References** to view the references.
   b) To remove the references, right-click each element that references the element that you want to remove, select **Edit**, and remove the element from the configuration.

3) Click **Yes**.
   The element is moved to the Trash.

Related tasks

Restore elements from the Trash on page 199
Delete elements from the Trash on page 200
Disable administrator accounts on page 405
Delete a Domain on page 443

Restore elements from the Trash

You restore elements from the Trash so that you can add them to new configurations.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Menu** > **View** > **Panels** > **Trash**.

2) Right-click the element that you want to restore, then select **Undelete**.



Note

You can view the references to the element you are restoring. In the confirmation dialog box, click **Open References**.

3) Click **Yes**.
   The element is restored.

Related tasks
Move elements to the Trash on page 199
Delete elements from the Trash on page 200

Delete elements from the Trash

You can permanently delete an element that you moved to the Trash.

If you select ⋮ **More actions** > **Show Deleted Elements** to view the elements that have been moved to the Trash, you can delete the element from the Trash in your current view. Otherwise, you can either delete a single element from the Trash branch or delete all elements in the Trash by emptying the Trash.

To permanently delete an element:

- Administrators must have sufficient rights.
- The element must not be used in any configuration, for example, in a policy.

CAUTION

Deletion is permanent. There is no undo. To recover a deleted element, you must either recreate it or restore it from a previously created backup or XML export file that contains the element.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select ☰ **Menu** > **View** > **Panels** > **Trash**.

2) Right-click the element that you are deleting, then select **Delete**.
   A confirmation dialog box opens.

3) If the element you are moving to the Trash is used in any configuration, view and remove all references to the element:
   a) Click **Open References** to view the references.
   b) To remove the references, right-click each element that references the element that you want to remove, select **Edit**, and remove the element from the configuration.

4) Click **Yes**.
   The element is permanently deleted.

Related tasks
Restore elements from the Trash on page 199

Empty the Trash

You can permanently delete all elements in the Trash at one time.

CAUTION

Deletion is permanent. There is no undo. To recover a deleted element, either recreate it or restore it from a previously created backup or XML export file that contains the element.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Menu** > **View** > **Panels** > **Trash**.

2) Select **More actions** > **Empty Trash**.
A confirmation dialog box opens.

3) Click **Yes**.
All elements in the Trash are permanently deleted.

Related tasks

Restore elements from the Trash on page 199

How Categories help you view only certain elements

Categories allow you to restrict which elements are displayed in the SMC Client. When you activate a Category Filter, elements that do not belong to one of the selected Categories are filtered out of your view.

Categories help you manage large networks by filtering the elements that are displayed. You can create separate Categories for elements that belong to a Engine, IPS, or Layer 2 Engine configuration and then select the category you want to configure. You can freely select how to assign the Categories, and quickly and flexibly change which combinations of Categories are shown according to your tasks.

In a large installation, there can be hundreds of elements, but you usually do not need to work with all elements at the same time. Category elements allow you to group related elements according to any criteria you want. Using Categories, you can quickly filter your SMC Client view. Elements that do not belong to the selected Category are filtered out so that only the relevant elements are visible. Categories allow you to manage many elements more efficiently by making it easier to find the elements you need.

There are two predefined Categories:

- The *System Elements* Category is assigned to all default elements in the SMC. You can use it to display all predefined elements in the system.
- The *Not Categorized* Category contains all elements that have not yet been assigned a Category.

Grouping Category elements with Category Tag elements

If you have many custom Category elements, you can group the Categories using Category Tag elements. Category Tags can also be used to filter elements in SMC Client views.

After you have created a Category Tag, you can select that Category Tag for a Category. You can also arrange Category Tags into groups by selecting a parent Category Tag for Category Tag elements.

When Category Filters are available, Category tags can be used as filtering criteria in the SMC Client.

Category configuration overview

You can create and combine Categories in a custom Category Filter.

Follow these general steps to configure a custom Category Filter:

1) Create a Category.

2) Associate elements with the Category.

3) Select a Category as the active Category Filter.

4) Combine Categories in a custom Category Filter.

Related tasks
Create new Category elements on page 202
Select Categories for elements on page 203
Activate Category Filters on page 204
Combine Category elements in custom filters on page 206

Create new Category elements

You can create as many Categories as you need, and you can base the Categories on any criteria.

For example, you can create separate Categories for elements related to different geographic locations. The same element can belong to several Categories. Categories are stored as elements and they are visible to other administrators as well.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.

2) Browse to **Other Elements** > **Categories**.

3) Right-click **Categories**, then select **New** > **Category**.

4) Give the Category a unique name.

5) (Optional) Enter a **Comment** for your own reference.

6) Click **OK**.

Related tasks
Select Categories for elements on page 203
Activate Category Filters on page 204

Add Category elements to groups using Category Tag elements

Add Categories to groups using Category Tags for easier management of categorized elements.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.

2) Browse to **Other Elements** > **Categories**.

3) Right-click **Categories**, then select **New** > **Category** to create a Category, or right-click an existing Category to open the Category properties.

4) Select a Category Tag for the Category.
   a) To select a category, click **Add**.
   b) To create a Category Tag, select **More actions** > **New** > **Category Tag**.

5) Click **OK**.

Select Categories for elements

You can select any number of Categories for each element without restrictions.

There are no automatic checks to consider; elements that reference each other do not need to be in the same Category. If you are using a Category Filter, the Categories included in the Category Filter are added to new elements when you create them. You can also manually select other Categories for elements and remove the automatically added Categories.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Click **Select** next to the **Category** field in the properties of a new or existing element.
   The **Category Selection** dialog box for the element opens.

2) Add or remove Categories as needed.

3) Click **OK**.

Related tasks

Create new Category elements on page 202
Activate Category Filters on page 204

Activate Category Filters

In most views, you can select a Category Filter to restrict which elements are displayed. You can also filter by more than one Category at a time.

The Category Filters are selected in the toolbar of the SMC Client. You can create a custom Category Filter containing any combination of Categories. For example, you can combine a Category for a particular geographic location and a Category for critical servers. This custom Category Filter combination only displays elements related to the critical servers at one location. The Category Filter is applied in all views.

Note

The selected Category Filter is applied *in all views* until you select a different Category Filter or **Category Filter Not Used**.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) If the Category Filter selection is not visible in the toolbar, select **Menu** > **View** > **Layout** > **Category Filter Toolbar**.

2) Select an existing Category.
   - If the Category you want to use is not listed, select **Select**, select the Category, then click **Select**.
   - To display the elements that do not belong to any Category, select the **Not Categorized** filter.
   - To display the predefined system elements, select the **System Elements** filter.

Result

Only the elements that belong to the selected Category are displayed. To display all elements again, select **Category Filter Not Used**.

Related tasks

Combine Category elements in custom filters on page 206

Activate the default Category Filters for Domain elements

When you create a Domain, you can set the default Category Filters that are automatically used when you log on to the Domain.

If you change the Category Filter, you can revert to the default Category Filter for the Domain.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) If the Category Filter selection is not visible in the toolbar, select **Menu** > **View** > **Layout** > **Category Filter Toolbar**.

   The Category Filter selection is displayed in the toolbar.

2) In the Category Filter toolbar, select **Default Category Filter for Domain**.

Related tasks

Create new Category elements on page 202
Activate Category Filters on page 204
Create Domain elements on page 436

Filter elements by Category Tag

Use Category Tags to filter categorized elements in different SMC Client views.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) If the Category Filter selection is not visible in the toolbar, select **Menu** > **View** > **Layout** > **Category Filter Toolbar**.

2) Select a Category Tag from the **Category Filter** menu to filter elements by Category Tag.



Note

If the Category Tag you need is not listed, select **Other** and navigate to the Category Tag.

Result

The elements in the view are filtered to only show elements in Categories that have the selected Category Tag.

Combine Category elements in custom filters

You can combine several Category Filters to display elements that are in any of the selected Categories.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the Category Filter toolbar, select **Define Custom Category Filter**.
2) Select the Categories you want to add, then click **Add**.
3) (Optional) If you want to view elements that do not have a Category (they belong to the Not Categorized Category), select **Show Not Categorized**.
4) (Optional) If you want to view elements that are predefined elements (they belong to the System Elements Category), select **Show System Elements**.
5) Click **OK**.
   Only the elements assigned to the selected Categories are displayed.

Related tasks

Create new Category elements on page 202
Activate Category Filters on page 204

Examples of Categories

The examples illustrate some common uses for Categories and general steps on how each scenario is configured.

Example: Creating category elements for a engine and an IPS configuration

This scenario shows an example of using categories to show only Engine or IPS engine configurations.

Company A is a large enterprise planning a new system. The system includes several Engine and IPS engines. Each Engine and IPS engine has its own policy. The company's administrators are only required to manage the Engines and their policies or the IPS engines and their policies at a time. To restrict which engines and policies are displayed, the following steps are taken:

1) The headquarters administrator creates two Categories: one for the elements that belong to the Engine configuration and another for the elements that belong to the IPS configuration.
2) The headquarters administrator creates the elements that represent the Engines, Engine policies, IPS engines, and IPS policies. The administrator then selects the appropriate Category for each element while defining its properties.

3) The administrators select the appropriate Category as the Category Filter so that only the elements in the Engine or IPS configuration are displayed.

Example: Combining category elements

This scenario shows an example of combining categories that are used at different sites.

Company B has sites in New York, Toronto, and Mexico City. The company's administrators have defined separate Categories for the elements that belong to each site, because administrators usually work with the elements of one site at a time. Today, however, Administrator A must apply the same configuration changes to the New York and Toronto sites. Administrator A does not want to create a Category for this temporary need. To be able to filter the elements belonging to both the New York and Toronto sites, Administrator A does the following:

1) Creates a custom Category Filter that contains the New York and Toronto Categories. The elements at both the New York and Toronto sites are displayed, and elements in the Mexico City Category are filtered out.

2) Makes the configuration changes to the elements in the New York and Toronto sites.

3) Selects the Category Filter Not Used filter to display all elements again.

Legacy elements and options

Elements and options related to features that are no longer supported might be visible in the SMC Client if you configured them using an earlier Forcepoint Network Security Platform software version, and are included only for backward compatibility.

For information about these features in earlier software versions, see the version-specific documentation.

Feature	No longer supported starting from
McAfee Endpoint Intelligence Agent (McAfee EIA)	Forcepoint Network Security Platform 6.3.0
McAfee Advanced Threat Defense	Forcepoint Network Security Platform 6.4.0

Part IV
Monitoring

Contents

- Monitoring Forcepoint Network Security Platform components on page 211
- Application Health Monitoring on page 255
- Monitoring third-party devices on page 261
- Viewing and exporting logged data on page 281
- Reports on page 311
- Filtering data on page 333

You can use the SMC to monitor system components and third-party devices. You can also view and filter logs, and create Reports from them.

Chapter 12

Monitoring Forcepoint Network Security Platform components

Contents

- Getting started with monitoring the system on page 211
- System monitoring tools in the SMC Client on page 212
- Overviews and how they work on page 227
- Monitoring users on the Dashboard on page 233
- Monitoring connections, block lists, VPN SAs, users, routing, SSL VPNs, and neighbors on page 236
- View and compare Element Snapshot elements on page 243
- Monitoring connections using Geolocation elements on page 244
- Monitoring configurations and policies on page 247
- Monitor administrator actions on page 248
- Monitor tasks on page 248
- Traffic captures and how they work on page 249
- Checking maintenance contract information on page 250
- Upcoming event notification on page 252

You can monitor Forcepoint Network Security Platform components and view system summaries in the SMC Client.

Getting started with monitoring the system

There are several ways to monitor the system in the SMC Client.

- Monitor the status of individual components and view a summary of the system status.
- Monitor the status of elements that belong to different administrative Domains.
- Create customizable overviews of the system.
- View user information and user alerts
- Monitor enforced Block Lists, open connections, active VPN SAs, active users, routing, and SSL VPN sessions.
- View, approve, and commit pending changes made to configurations and policies of engines.
- Check which configurations and policies are currently applied in the system.
- Check which actions administrators take.
- Check the status of Tasks that schedule commands to run automatically.
- Monitor the status of the maintenance contract.

Related concepts
Getting started with monitoring third-party devices on page 261
Getting started with the Logs view on page 281
Getting started with reports on page 311

System monitoring tools in the SMC Client

There are various tools and views that you can use to monitor the Security Engine system.

The Domain Overview

Domain elements allow you to group elements that belong to specific configurations. The Domain Overview allows you to see the status of all Domains and their elements.

If the configurations are divided into different administrative Domains, the **Domain Overview** is shown as the first view after logon. The Domain Overview is only available to administrators who have permissions for several domains. You can then select the Domain that you want to manage.

Note

In a HA environment with multiple Management Servers, when you log on to a Domain from the Domain Overview, the Domain is by default opened on the active Management Server.

Related concepts
Getting started with Domain elements on page 433

How the Dashboard is arranged

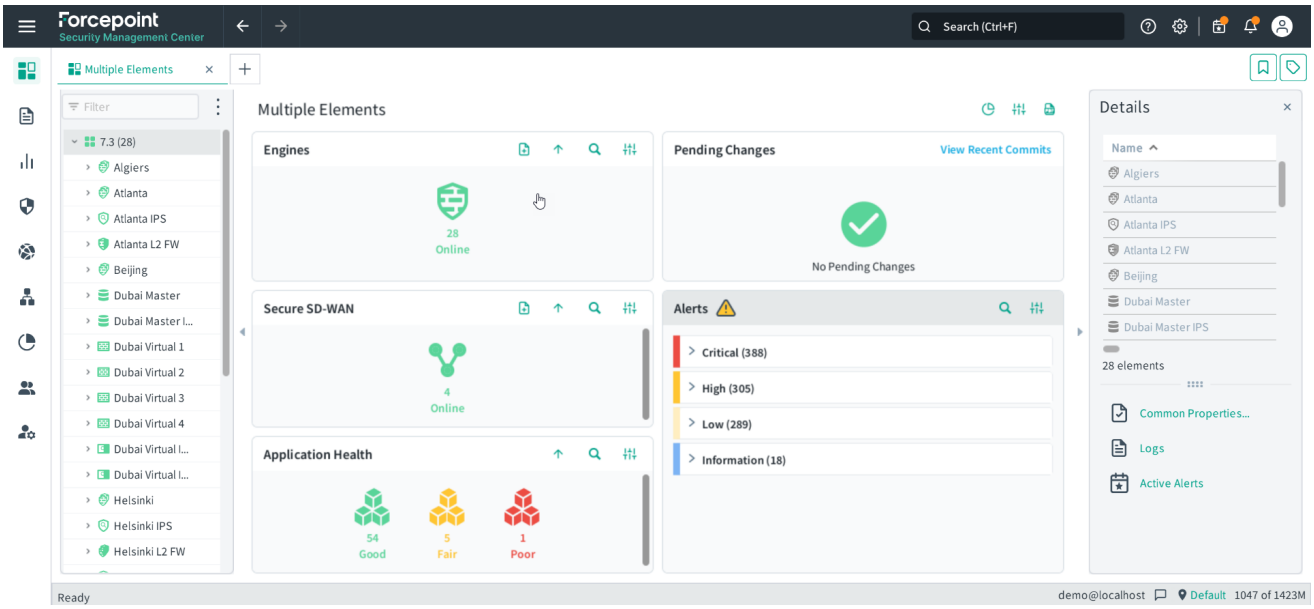
On the **Dashboard**, you can view the status of Security Engine components and monitored third-party devices.

By default, when you start the SMC Client, you see the **Dashboard**. This view provides the operating and connectivity status of SMC components and third-party components that are set up to be monitored through the SMC. The status information is stored on Log Servers. The Management Server compiles the Dashboard view based on data from all Log Servers.

There are several ways to open the **Dashboard**. For example:

- Select **Dashboard** > **Engines Dashboard**.
- Open a blank screen, then select **Monitoring** > **Engines Dashboard**.

Dashboard



1. The **Status** tree shows the status of monitored system elements. To see the status of a component in the Status tree, expand the tree, then place the cursor over any element. You can see the element's IP address and status in a tooltip.
2. Status cards for Security Engines, VPNs, and other monitored elements or the Home page for the selected element show detailed information about the status and configuration of monitored elements.

   Click the card for an element to open the element's home page. Click + **New** to add an element of the selected type.
3. The **Pending Changes** pane shows configuration and policy changes that have not yet been transferred. The **Recent Commits** pane opens in the same place and shows recent policy uploads.
4. The **Details** pane shows details of the selected element. Blank if no element is selected in the **Status** tree.
5. The **Drill-downs** pane contains shortcuts to more details of the selected element. Blank if no element is selected in the **Status** tree.
6. The **Application Health** pane shows network applications being monitored.
7. The ⋮ **More actions** menu allows you to organize alerts according to the severity or situation type.
8. Location for monitoring the system

Card size toolbar



1. **Auto** — Selects the card size automatically based on the space available.
2. **Sum** — Shows the number of engines for each status. This card size is useful for large installations.
3. Small cards show only the status and whether there are any active alerts for the element.
4. Medium cards show the status of the element, when the last policy upload was made, and whether there are any active alerts for the element.
5. Large cards show the status of the element, configuration information about the element, and whether there are any active alerts for the element. For Security Engines, the cards also show the current load.

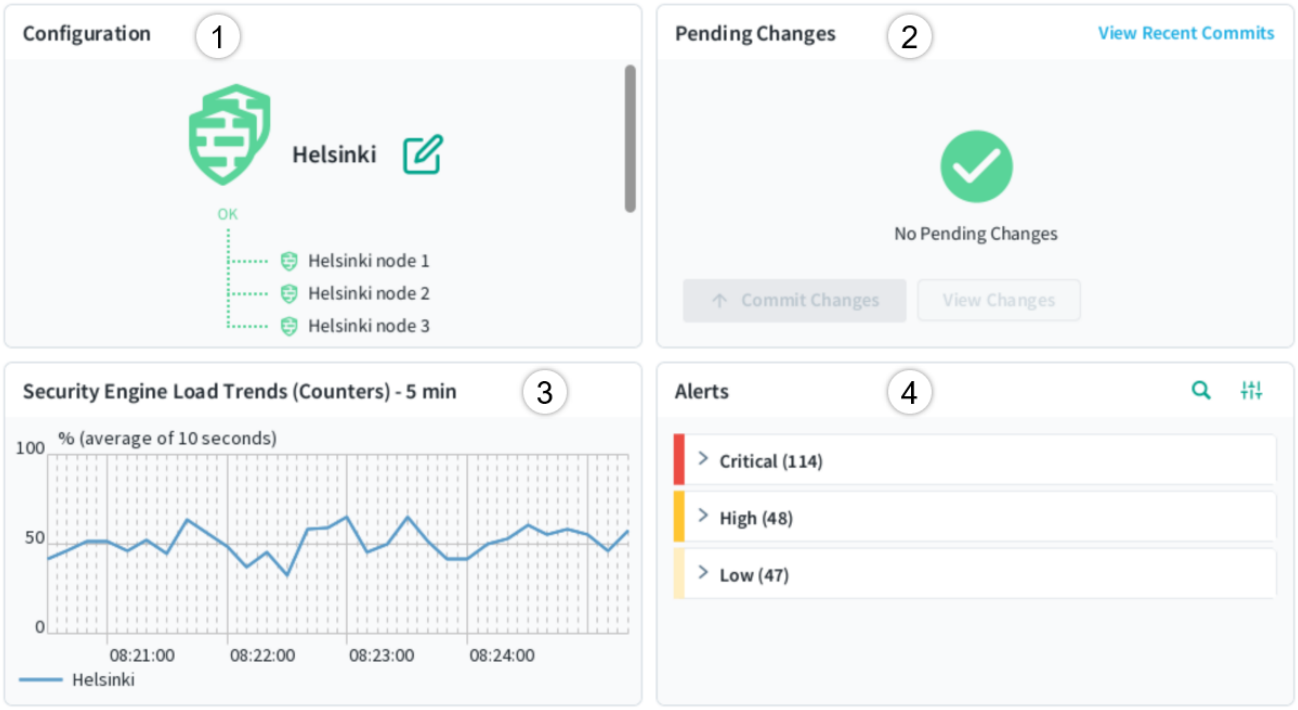
What the Security Engine home pages show

Home pages show different information depending on the type of element.

Click the status card of an element or select the element in the **Status** tree to open the home page for the element.

To edit the layout of the element home pages, click **Edit** at the top right corner of the view. You can re-size panes, or drag new panes into the view. To remove a pane, drag the pane to the top of the view.

Security Engine home page



1. Shows information about the configuration status and the VPN branch status of the Security Engine.

   To view detailed information about the VPN branch associated with the Security Engine, click the branch name. To preview the Security Engine properties or the policy, click the Security Engine or policy name. To edit the properties or the policy, click **Edit**.

   If the configuration of an Security Engine has not yet been completed, you can continue the configuration from the home page. For example, you can save the initial configuration or install a policy. The uncompleted configuration steps are shown on the home page.
2. The **Pending Changes** pane shows configuration and policy changes that have not yet been transferred to the Security Engines. Provides options for viewing, approving, and committing pending changes. The **Recent Commits** pane shows recent policy uploads. It opens in the same place as the **Pending Changes** pane.
3. Shows how the traffic load on the Security Engine has changed over time.
4. Shows active alerts for the Security Engine. The **More actions** menu allows you to organize alerts according to the severity or situation type.

When you select an individual Security Engine node in the **Status** tree, the hardware diagram page opens, showing details about the status of network ports. More detailed information is shown in the Info pane for network interfaces and hardware (appliance) status.

Engine Status History

The Security Engine the status history view by engine shows status changes over time, and engine specific traffic trends and top bandwidth network applications.

- On **Dashboard** view, right click any of the Security Engines, and select **Monitoring** > **Status History**.

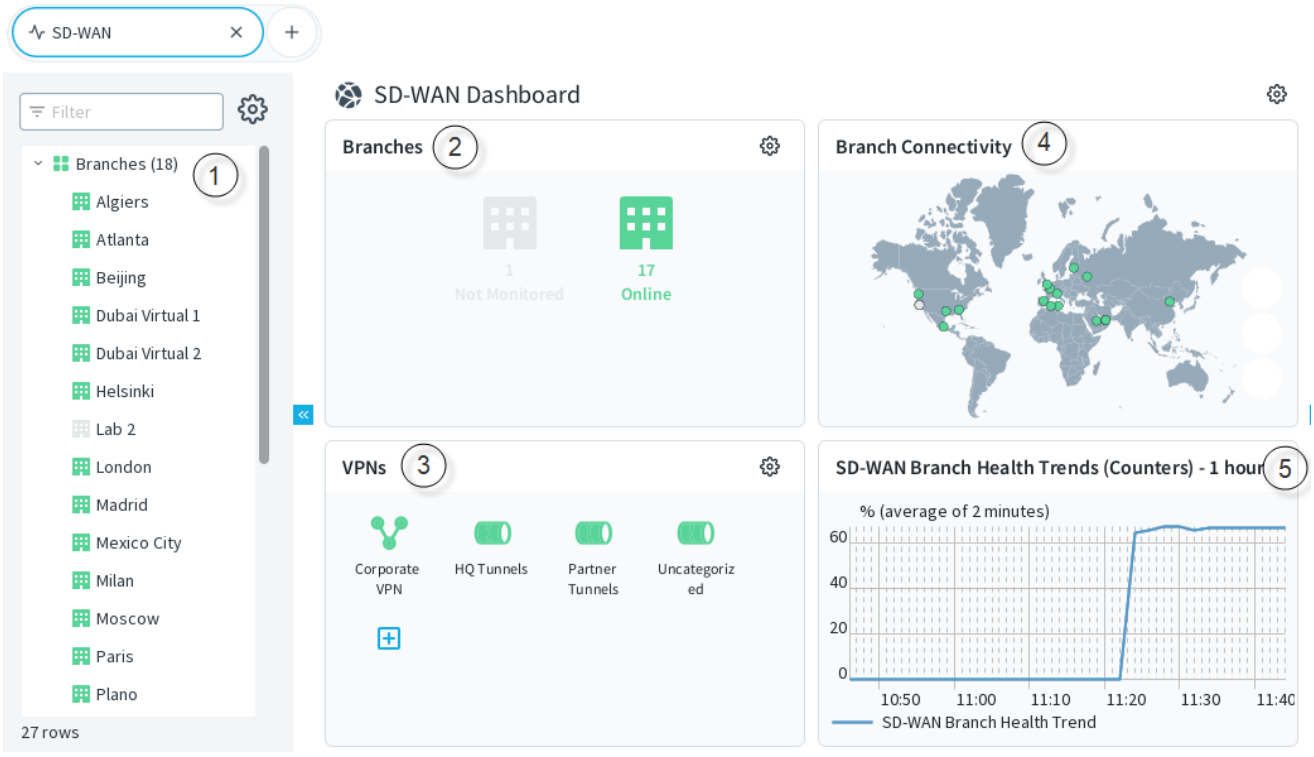
  By selecting, different time periods, view the Security Engine statuses over the selected period.

What the Secure SD-WAN dashboard shows

The Secure SD-WAN dashboard allows you to monitor Secure VPN features, such as Multi-Link and VPNs, and to view statistics and reports related to VPN features.

Branches represent the VPN gateways and NetLink elements associated with each Security Engine. The SD-WAN dashboard summarizes status information for all branches and VPNs. Various statistics items related to VPN monitoring are available when you customize the SD-WAN dashboard. You can also use these statistics items in Reports and Overviews.

Secure SD-WAN dashboard



1. The status tree shows the status of monitored branches, Route-based Tunnels, and policy-based VPNs. A branch is shown for an Security Engine if the Security Engine has VPN gateways that are used in a policy-based VPN or route-based tunnel, or if NetLinks associated with the Security Engine are used in an Outbound Multi-Link element.
2. Contains status cards that show information about the VPN health and NetLink status of each branch. When you select an individual branch, the Dashboard view shows detailed monitoring information about the NetLinks, VPN tunnels, and traffic associated with the branch.
3. Shows the overall status of VPNs and the status of gateways and tunnels in the VPNs. When you select an individual policy-based VPN or group of Route-based Tunnels, the Dashboard view shows detailed monitoring information about the policy-based VPN or Route-based Tunnels.
4. Shows the status of all monitored branches on a map.
5. Shows how the health has changed over time.

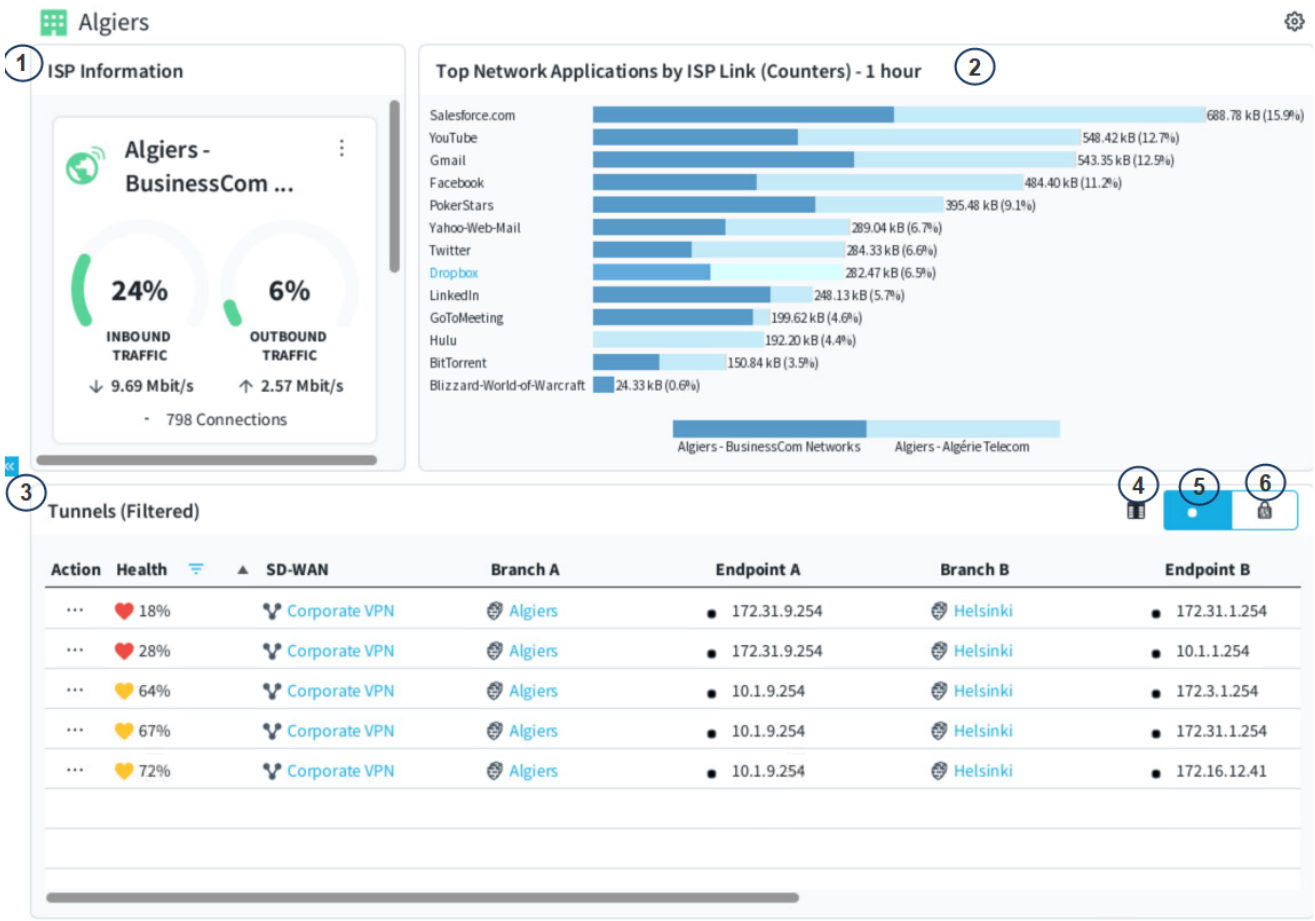
What the branch home pages show

Branch home pages show detailed monitoring information about the NetLinks, VPN tunnels, and traffic associated with the branch.

Click the status card of a branch or select the branch in the **Status** tree to open the home page for the branch.

To edit the layout of the home pages, click **Edit** at the top right corner of the view. You can re-size panes, or drag new panes into the view. To remove a pane, drag the pane to the top of the view.

Branch home page



1. Shows the information about the Netlinks that represent each ISP connection.

   Click ⋮ > **View NetLink history** to show ISP link status changes over time, sent and received traffic, traffic by network application, application and network latency, packet-loss, and jitter trends for the ISP links.

2. Shows the most used Network Applications according to the ISP connection that the traffic used.
3. Shows the status of the tunnels in the VPNs associated with this branch.

   On a specific row, click the **...** > **View tunnel history**. This shows the VPN status history, health, jitter, latency and packet-loss, and traffic trends on the gateway or the individual tunnel level.

4. Allows you to configure the column to view in the table.
5. Shows the status of endpoint-to-endpoint tunnels.
6. Shows the status of gateway-to-gateway tunnels.

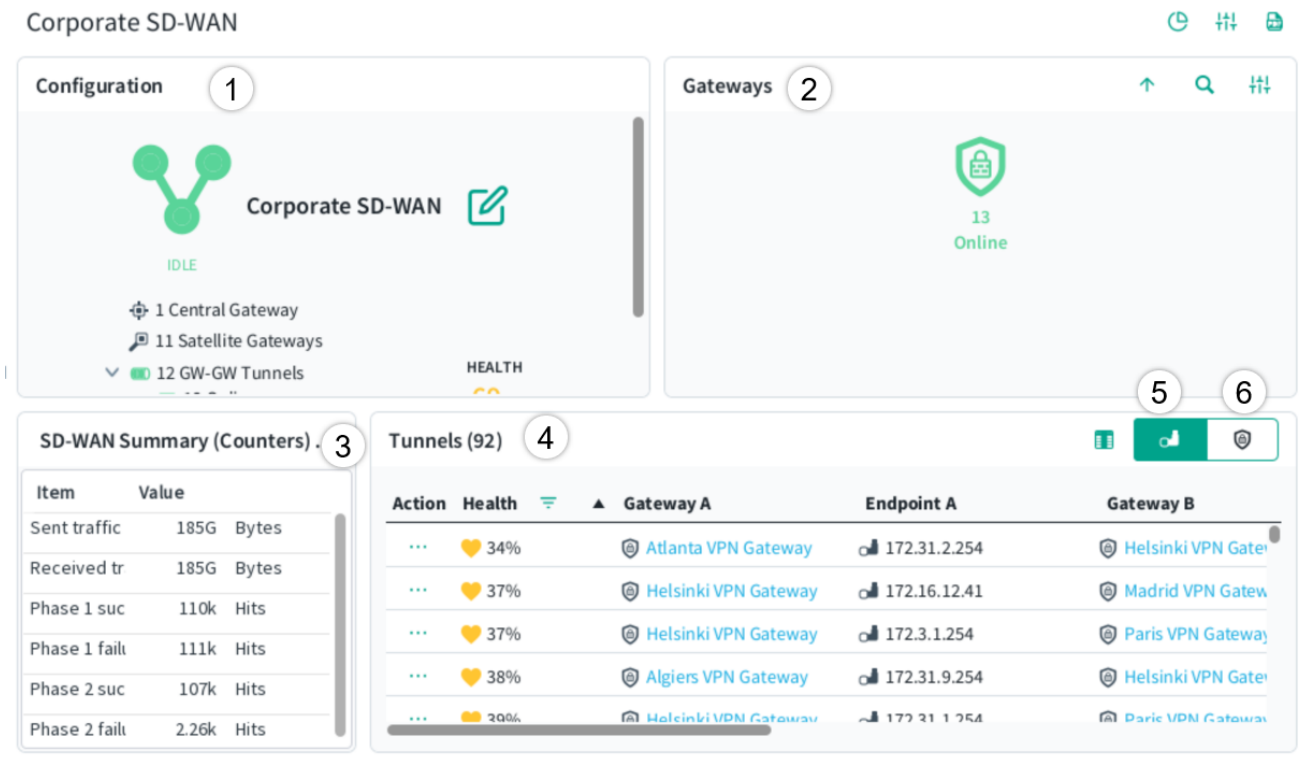
What the VPN (SD-WAN) home pages show

Home pages show different information depending on the type of VPN element.

Click the status card of an element or select the element in the **Status** tree to open the home page for the element.

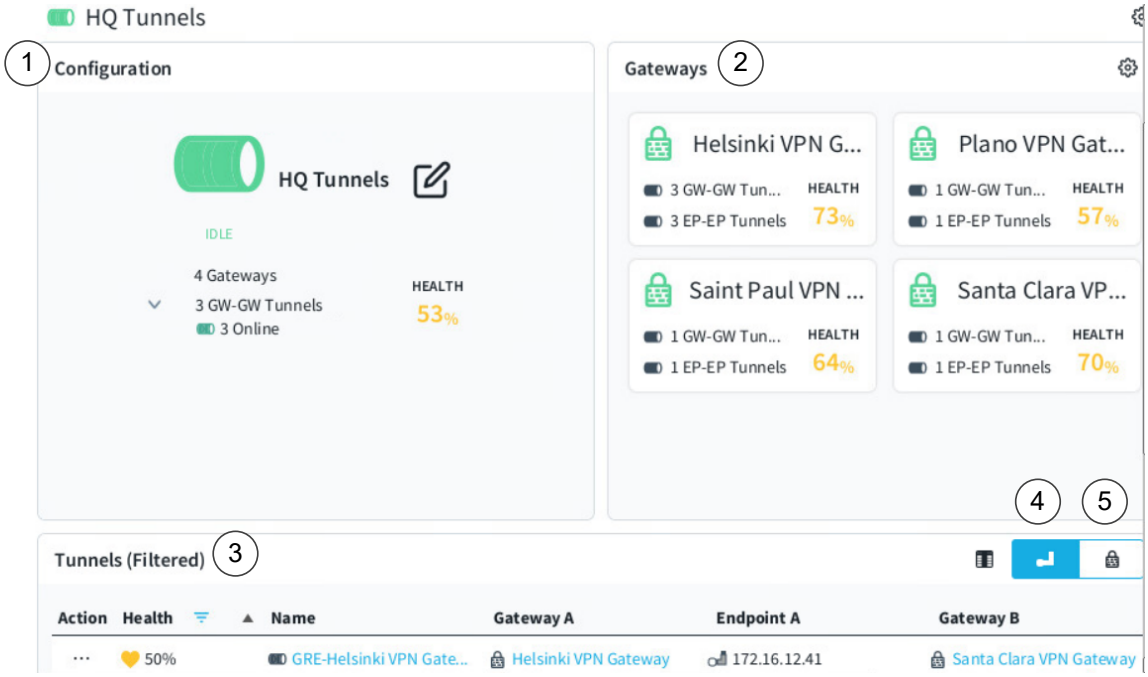
To edit the layout of the element home pages, click **Edit**. You can re-size panes, or drag new panes into the view. To remove a pane, drag the pane to the top of the view.

Policy-based VPN element home page



1. Shows information about the configuration status of the VPN. To preview the VPN, click the name of the VPN. To edit the VPN, click **Edit**.
2. Shows the status of the gateways in the VPN.
3. Shows a summary of statistics related to the VPN.
4. Shows the status of the tunnels in the VPN.
5. Shows the status of endpoint-to-endpoint tunnels.
6. Shows the status of gateway-to-gateway tunnels.

Home page for route-based tunnels

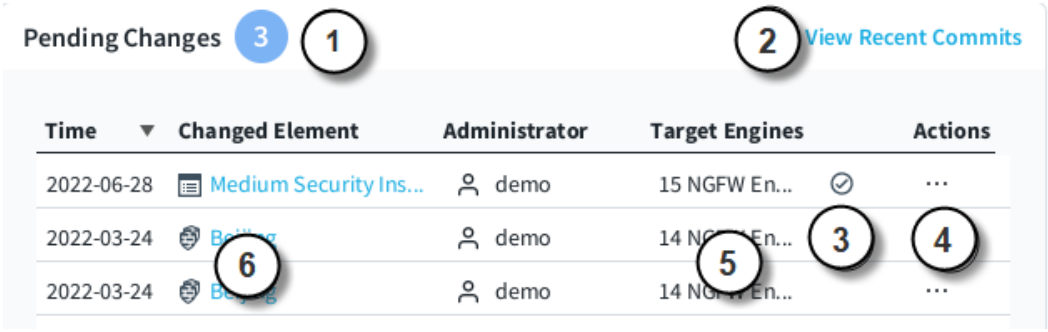


1. Shows an overview of the tunnels in the group.
2. Shows the status of the gateways in the tunnel group.
3. Shows the status of the tunnels.
4. Shows the status of endpoint-to-endpoint tunnels.
5. Shows the status of gateway-to-gateway tunnels.

What the Pending Changes pane shows

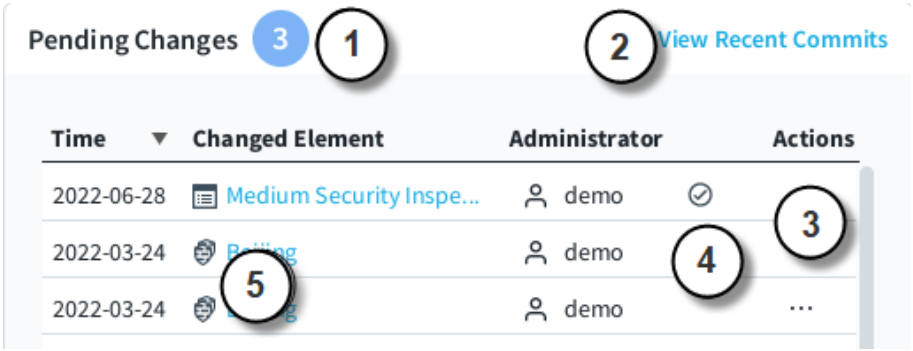
The **Pending Changes** pane shows configuration and policy changes that have not yet been transferred to the Security Engines. What the pane shows depends on whether an Security Engine's Dashboard page or the Dashboard view showing all Security Engines is open.

Pending Changes pane on the Dashboard page for all Security Engines



1. The total number of pending changes for all Security Engines.
2. Opens the **Recent Commits** pane.
3. The approval status of the pending change.
4. A menu with options to view audit data and to find out where the element is used.
5. The Security Engines affected by the pending change.
6. The element where the change was made.

Pending Changes pane on the Dashboard page for an Security Engine

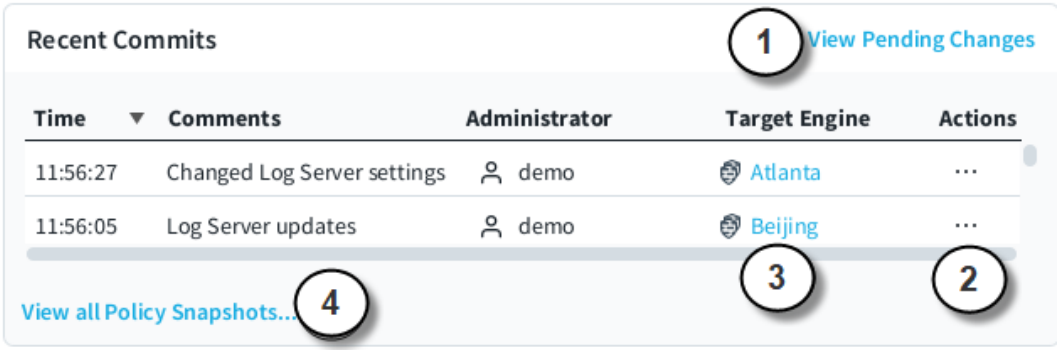


1. The number of pending changes for the Security Engine.
2. Opens the **Recent Commits** pane.
3. A menu with options to view audit data and to find out where the element is used.
4. The approval status of the pending change.
5. The element where the change was made.

What the Recent Commits pane shows

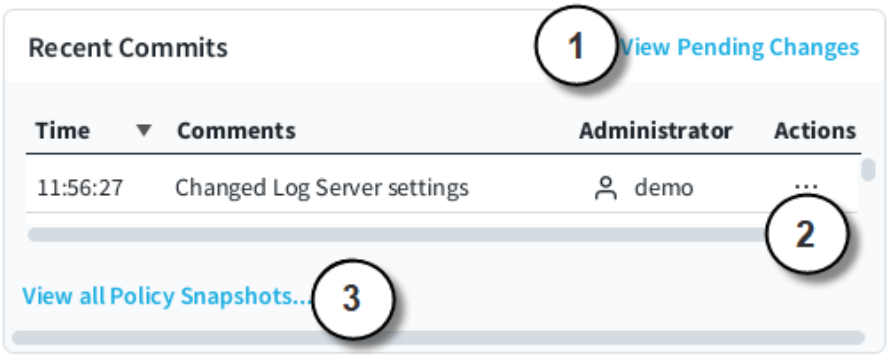
The Recent Commits pane shows the list of recent policy uploads on the selected Security Engine or all Security Engines, depending on whether an Security Engine's Dashboard page or the Dashboard view showing all Security Engines is open.

Recent Commits pane on the Dashboard page for all Security Engines



1. Opens the **Pending Changes** pane.
2. Options for viewing, comparing, and restoring Policy and Element Snapshots.
3. Shows the Security Engines to which the policy was uploaded.
4. Opens the Policy Snapshot view for the selected Security Engine type.

Recent Commits pane on the Dashboard page for an Security Engine



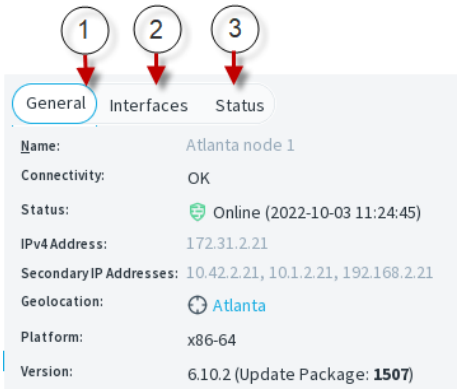
1. Opens the **Pending Changes** pane.
2. Options for viewing, comparing, and restoring Policy and Element Snapshots.
3. Opens the Policy Snapshot view for the selected Security Engine type.

What the Details pane shows

The **Details** pane is shown by default in most views. In addition to element details, the details pane shows the most important status information for components.

The type of element determines which tabs are shown.

Details pane for a node in an Security Engine

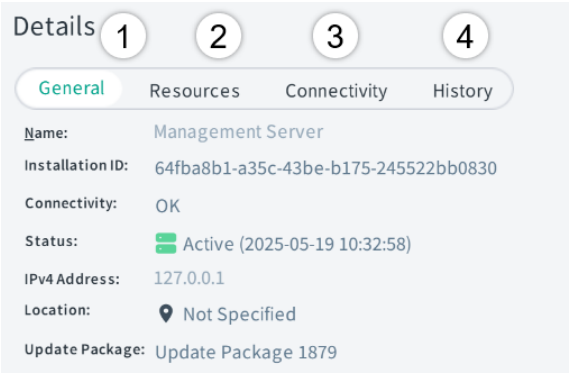


1. **General** tab — Shows general information about the element.
2. **Interfaces** tab — Shows information about the network ports of the selected engine node, such as speed and duplex.
3. **Status** tab — Shows the status of hardware and services.

   For Security Engine appliances, the Status tab shows the hardware status of the selected device.

   If the anti-malware feature is used, the status of the anti-malware signature database is shown.

Details pane for a Management Server



1. **General** tab — Shows general information about the element.
2. **Resources** tab — Shows information about resource usage on the computer where the Management Server, Log Server, or Web Access Server is installed for troubleshooting purposes.

   To refresh the information, click **Update**.

   If the memory usage gets too high, the Management Server, Log Server, or Web Access Server automatically restarts, and an alert and an audit entry are generated.
3. **Connectivity** tab — Shows information about the connectivity between the selected SMC server and other system components.
4. **History** tab — Shows audit information about when the element was created and modified.

Run diagnostics on SMC servers

You can run diagnostics to check the status of the communication between all the SMC servers (Management Servers, Log Servers, and Web Access Servers).

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Management Server, then select **SMC Servers Replication Diagnostics**.
   A dialog box that shows diagnostics about the status of all the SMC servers opens.

2) View the diagnostics and check if there are issues that you should resolve about any of the servers, then click **Close**.

Next steps

If the diagnostics indicate any issues with any SMC servers, such as issues with certificates or licenses, resolve the issue.

View the status of appliance configurations

When you configure a Forcepoint Network Security Platform appliance using the plug-and-play configuration method, you can view the status of the configuration process.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Settings** > **View Appliance Configuration Status**.

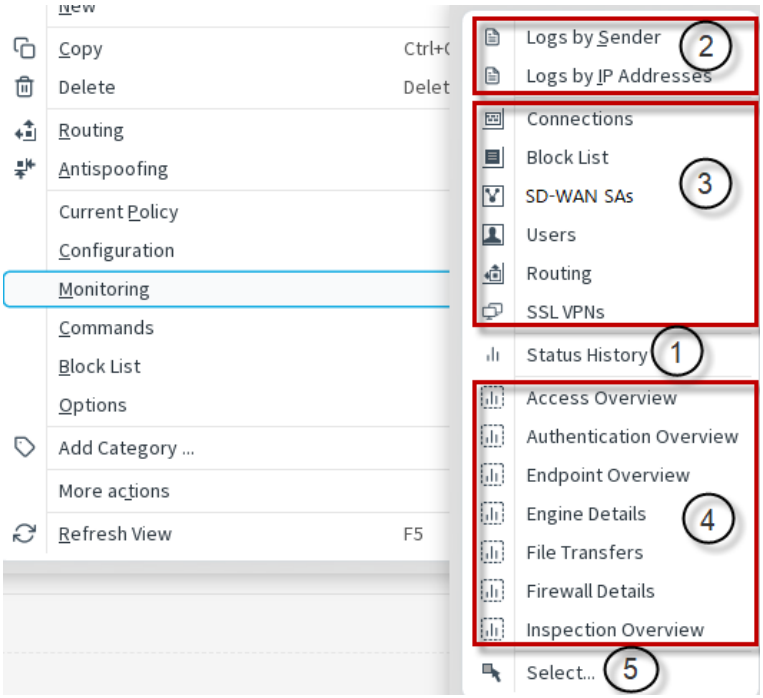
Related concepts

Management connections for Security Engines and how they work on page 629

Shortcuts for monitoring different elements

Actions for monitoring a component are available in the Monitoring branch of the element's right-click menu. The available actions depend on the component type.

Monitoring submenu in right-click menu for components — for Engines



1. View component status.
2. View log data for component.
3. View active connections, block list entries, VPN SAs, users, routes, and SSL VPNs.
4. View or edit a detailed Overview or third-party Overview.
5. Select an Overview that is not listed.

Related concepts

Getting started with the Logs view on page 281
Getting started with reports on page 311

How component statuses are indicated

The status of Security Engine components and monitored third-party components is indicated by colors in most views where the elements are displayed. The status of various system communications is indicated by colors in monitoring diagrams and in the Info pane.

In addition to element status colors, the following icons indicate status:

- Pending configuration or policy changes on an engine are indicated by a blue balloon showing the number of changes next to the Security Engine in the **Status** tree.

- A blue icon on an Security Engine node in the **Dashboard** view indicates that a policy upload is in progress.
- Hardware malfunctions are indicated with special icons in the **Status** tree. If any problems are indicated, always check the logs and alerts to find out what is causing the problems.

Icons that indicate malfunctions

Engine hardware malfunction is indicated with an icon on top of the affected engine's icon in the Status tree and on all top-level branches of the tree.

For more information about the hardware malfunction, select the engine in the Status tree, then click the **Status** tab in the **Info** pane.

Hardware malfunction icons

Icon	Hardware status	Description
	Warning	A predefined Warning level has been reached in hardware monitoring (for example, the remaining file system capacity is less than 15%). The system also generates a log entry.
	Alert	A predefined Alert level has been reached in hardware monitoring (for example, the remaining file system capacity is less than 5%). The system also generates an alert entry.

Replication malfunction icon

In an environment with more than one Management Server, the configuration data is replicated to all Management Servers that are online. If the replication of configuration data among the Management Servers fails, an exclamation point on a yellow triangle is shown. You can see this icon on top of the Management Server's icon in the Status tree for each Management Server that is not synchronized with the other Management Servers.

Element status colors

The element-level status gives an overview of the status of all engine nodes that are represented by the element (also shown for single-node components).

Element-level status

Color	Element status	Description
Green	All OK	All nodes have a normal status (online or standby).
Yellow	Warning	Some nodes have an abnormal status or have been commanded offline, but are still sending status updates normally.
Red	Alert	All nodes have an abnormal status, there are one or more nodes that have not sent expected status updates, or all nodes have been commanded offline.
Gray	Unknown status	No policy has been installed on any of the nodes.
White	Not monitored	An administrator has disabled monitoring for all nodes.

Node status colors

The node status gives more detailed information about individual engines.

Node-level status

Color	Node status	Description
Green	Node or server online	The node or server is online.
Green (with slot)	Locked online	The node is locked online to prevent automatic status transitions. The node does not change state unless commanded by an administrator.
Cyan	Standby mode	Used with clustered engines when the cluster is in Standby mode. The node is in standby mode. One of the standby nodes goes online when the previous online node goes to a state in which it does not process traffic.
Purple	Node offline	The node is offline and does not process traffic.
Purple (with slot)	Locked offline	The node is locked offline to prevent automatic status transitions. The node does not change state unless commanded by an administrator.
Gray	Timeout or unknown status	The Management Server does not know the status of the node.
White	Not monitored	An administrator has disabled monitoring for the node.
Gray outline	Under work, no contact to SMC yet	The node is under work and has not contacted the SMC yet.
Green outline	Under work, no first policy upload yet	The node is under work, has contacted the SMC, but the first policy upload has not yet been made.

NetLink status colors

NetLink status shows the status of the network links in a Multi-Link configuration.

Note

The NetLink elements are queried and the status is displayed only if probing settings are configured in the NetLink elements and the Outbound Multi-Link element is included in the engine configuration.

NetLink status icons

Color	NetLink status	Description
Green	OK	The NetLink is up.
Orange	Mixed	At least one Security Engine that uses this NetLink has reported an error status for the NetLink.
Gray	Unknown status	The Management Server does not know the status of the NetLink.
White	Not monitored	An administrator has disabled monitoring for the NetLink.

VPN status colors

The VPN status shows the health of the VPN tunnels.

VPN status

Color	Cluster status	Description
Green	Tunnels up	All tunnels have a normal status (online or standby) and there is traffic.
Yellow	Warning	An error was detected, at least for some traffic, but the tunnels are usable in principle, and some other traffic might be getting through.
Red	Error	Some or all tunnels are down.
Blue	Idle	The tunnels are valid, but there has not been recent traffic.
White	Not configured	The VPN has no valid tunnels, because the VPN configuration is not complete or does not allow any valid tunnels.

Connectivity status colors

Element diagrams and the **Connectivity** tab in the **Info** pane show the status of the connectivity between elements.

See the tooltip for the status color for more details.

Connectivity status

Color	Status	Explanation
Green	OK	The connection is active (there have been communications within the last 2 minutes) and no problems have been reported.
Red	Error	The Management Server received a report that connection attempts have failed.
Cyan	Idle	Connection between components is still open but there is a pause in communications.
Yellow	Warning	There are problems with heartbeat or state synchronization between nodes in a Engine Cluster. Only one of the interfaces used for heartbeat or state synchronization functions properly. This warning does not affect how the cluster functions.
Blue	Closed	The connection was closed by one of the components.
Gray	Timeout, Unknown	The Management Server does not know the status of the connection. The connection might or might not be working.

Overviews and how they work

Customizable Overviews contain information on the system's status, including shortcuts to frequently used views and statistical charts of the system's operation. An example of a frequently used view is a log filtered by criteria that you configured. Examples of statistical charts are engine load and traffic flow.

You can create new Overviews on an empty template or start your customization based on one of the default Overview templates.

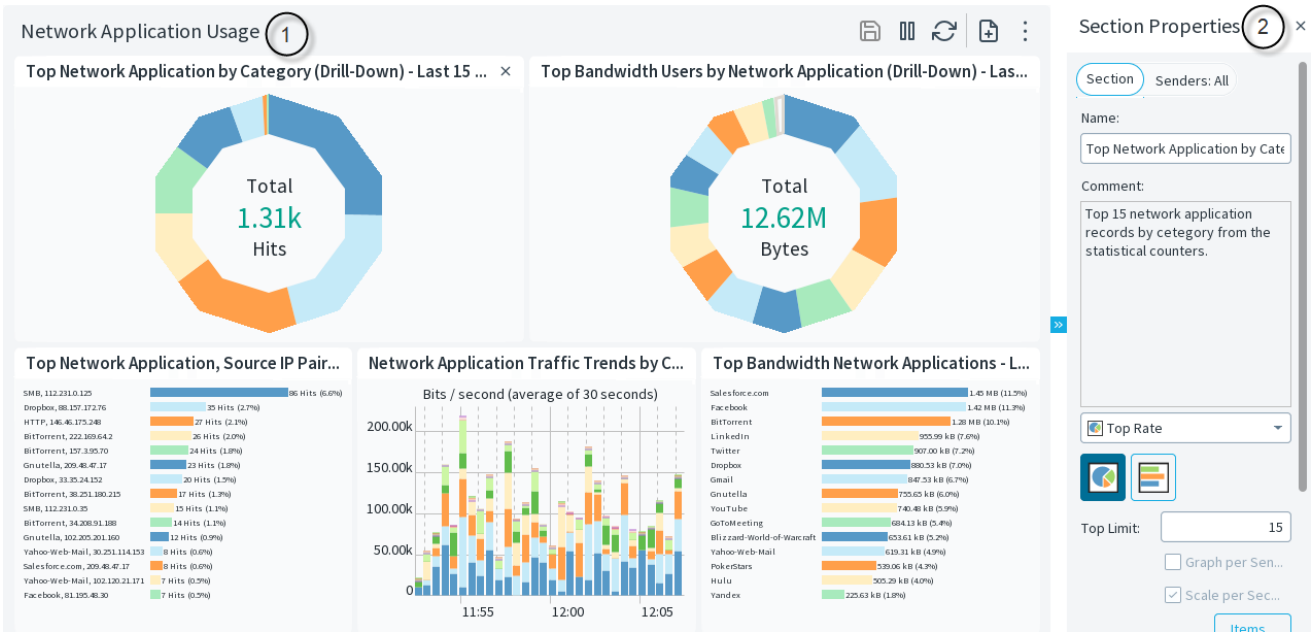
What Overview elements show

Overviews contain high-level status and statistical information.

You can use overviews to view:

- The system status
- Bookmarks, such as logs filtered with specific criteria
- Statistical charts on system operation (such as engine load) and the traffic flow

Example overview



1. Information is displayed as resizeable sections.
2. Use this pane to conveniently edit the content and appearance of the selected section.

You can create several different overviews for different purposes. Several default overviews are provided as templates.

In addition to status information, you can add various statistics related to the traffic and the operating state of components. You can display information in various ways, such as tables, maps, and different types of charts. Statistics can trigger an alert when the value of a monitored item reaches a limit you set.

Create Overview elements

After you create an Overview element, you can add a summary of the system status or statistics.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Overviews** > **New Overview**.

2) Select the Overview template:
   - To add your own Overview sections to an empty grid, select **Empty Overview**, then click **OK**.
   - To use one of the predefined Overview templates, select the template from the list, then click **OK**.

Modify Overview elements

Add System Summary sections and Statistics sections to customize your Overview.

Add System Summary Sections to Overview elements

The system summary is shown in the default start view, but you can also add it to your own Overviews.

It is possible to add more than one system summary to the same overview, but the information displayed is always the same.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Overviews** and an Overview.

2) Select **New** > **System Summary Section**.

3) (Optional) Adjust the placement and size of the new section by dragging and dropping the section or its edges. Resizing is based on preset grid positions. For resizing to work, drag the edge until it snaps to the next position on the screen.

4) Click **Save** or select **More actions** > **Save As**.

Add Statistics Sections to Overview elements

If you want to see statistical information in table or chart form, add a Statistics Section.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Overviews** and an Overview.

2) Select **New**, then select a section from the list.

> **Tip**
>
> If you cannot find an appropriate section in the list, select **Select**, then select a section in the **Select Section** dialog box.

> **Tip**
>
> If you want to add a section based on a statistical item, select **Create from Item**, then select an item from the **Select Item** dialog box.

3) Define the basic section properties in the Section Properties pane.

4) (Optional) Click the **Senders** tab, then select which elements are shown in the section.

5) (Optional) Adjust the placement and size of the new section by dragging and dropping the section or its edges.

> **Note**
>
> Resizing is based on preset grid positions. For resizing to work, drag the edge until it snaps to the next position on the screen.

6) Click **Save** or select **More actions** > **Save As**.

Create Statistics Sections

You can save a section you have customized in one Overview as a Statistics Section. Saving allows you to create the same type of section with the same settings in other Overviews.

After you create a Statistics Section, you can add the new section to other Overviews.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Overviews** and an Overview.

2) Right-click any section in the Overview, then select **Save As New Section**.

3) Define the basic section properties on the **General** tab of the **Section Properties** dialog box.

4) Select the diagram settings in the **Visualization** tab.

5) Select the diagram type.

Note

You can change this selection when you include the section in an Overview. The options available depend on the diagram settings you made in the previous step. If available, enter the number of items to be included in the **Top Limit** field. If you selected **Progress** as the diagram type, you can select **Graph per sender** or **Scale per Second**.

6) (Optional) Click the **Items** tab, then select or remove statistics items for the section.

7) (Optional) Click the **Senders** tab, then select which elements are shown in the section.

8) Click **OK**.

9) Click **Save** or select **More actions** > **Save As**.

Related concepts

Creating and editing local filters on page 337

How Statistical Items help you visualize data

Statistics process and visualize data. They help you focus on the most relevant information in the system (for example, trends in network traffic) and find changes and anomalies in network traffic.

You can use statistics in Overviews and Reports, and when you browse Logs or Connections. Filters are available to help you find information.

Statistical items count log entries, summaries of log fields included in those entries (like traffic volumes in log data containing accounting information), or specified statistical data (counter items). The items are organized based on the component types, as the runtime data they produce is different.

Note

Log entries are referred to as records in the item names.

Add Statistical Items to Overview elements

To see a representation of statistical data in an Overview, add a Statistical Item.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Overviews** and an Overview.

2) In the **Section Properties** pane, click **Items**.

3) Click **Add**.

4) Select the items.

5) In the **Properties** dialog box, click **OK**.
The items are added to the section and their data is displayed in the section.

Remove Statistical Items from Overview elements

If you no longer find a Statistical Item useful or relevant, you can remove it.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the **Section Properties** pane, click **Items**.

2) From the section, select the item that you want to remove, then click **Remove**.

3) In the **Properties** dialog box, click **OK**.
The item is removed from the section.

Set thresholds for monitored items in Overview elements

Thresholds activate automatic tracking of monitored items in Overviews. The values of the monitored items are checked once an hour.

For Progress items, the total of the values is compared to the threshold limit. The threshold is considered exceeded if the average level of the curve is above the threshold limit during the tracking period.

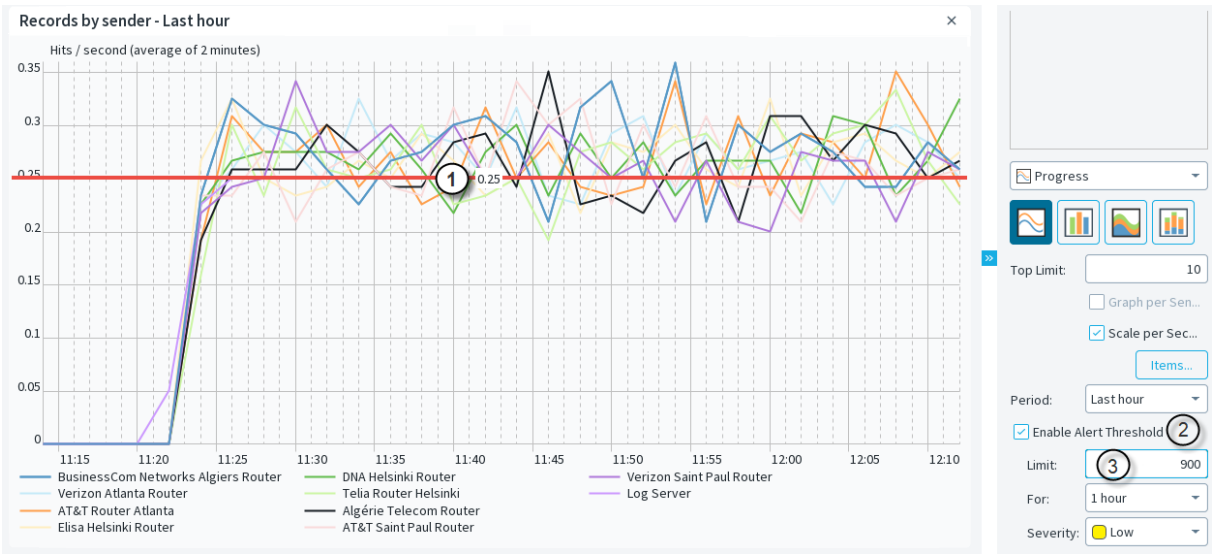
For Top Rate and Top Rate Curve items, the highest value is compared to the threshold limit. The threshold is considered exceeded if the highest value is above the threshold limit during the tracking period. If the threshold limit is exceeded, an alert is sent.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Overviews** and an Overview.

2) Select **Enable Alert Threshold**.



1	Drag the threshold line to set the alert limit
2	Alert threshold option enabled
3	Enter the alert limit manually

3) Specify the threshold limit in one of the following ways:
   - Drag and drop the threshold line in the overview section.
   - Enter the **Limit** as a number in the **Section Properties** pane.

4) (Optional) Select the tracking period during which monitored items are compared to the threshold limit from the **For** drop-down list. By default, one hour is selected.

5) (Optional) Select the **Severity** of the alert that is sent when the threshold limit is exceeded. By default, **Low** is selected.

6) Click **Save** or select **More actions** > **Save As**.

Related tasks

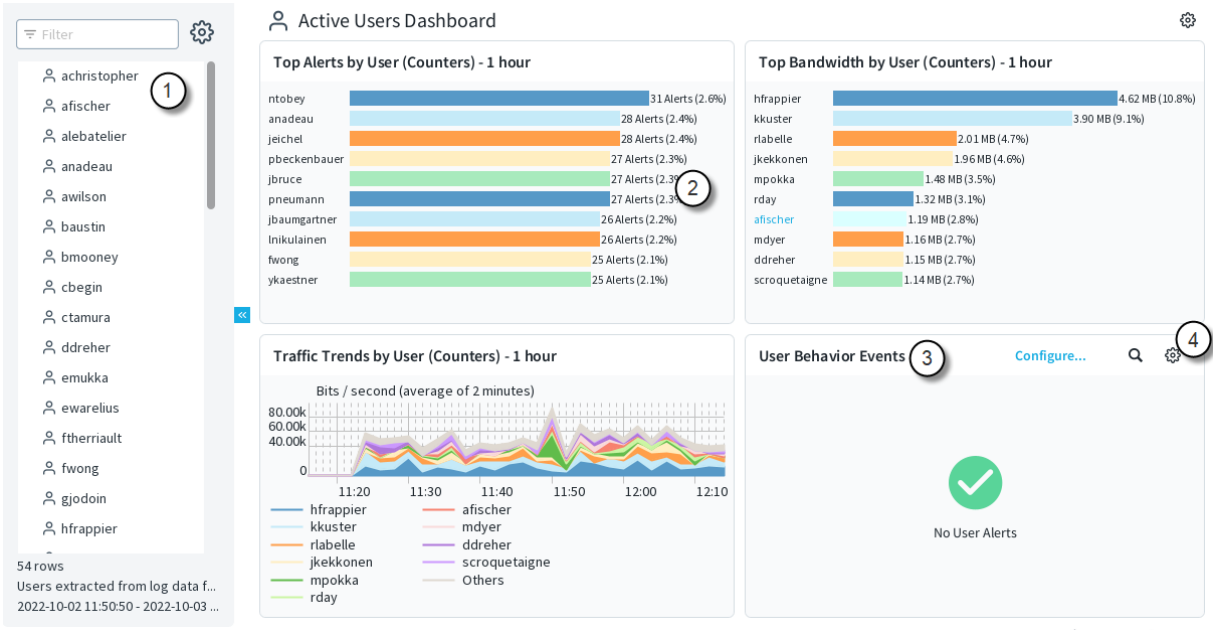
Create Overview elements on page 229

Monitoring users on the Dashboard

In the Dashboard view of the SMC Client, there are user dashboards where you can see an overview of user activity.

For example, you can see if there is any activity that indicates suspicious behavior, such as the use of certain network applications, attempts to access specific networks, or if a user has been associated with an attack Situation.

User dashboard



1 When users have been active and have caused log data to be generated, they are shown in the **Users** list. You can configure the time period within which a user must have been active. If there are no user names stored in log data, or in regions where privacy laws require that users must not be easily identified, you can show the IP addresses of users instead of their names.

> **Note**
>
> To be able to monitor users by name, you must enable the logging of user information in the Engine IPv4 and IPv6 Access rules.

2 The Statistics panes contain charts and general statistics of user activities, and if you select an individual user, you can see more detailed information about the user and their activities. If user information from Active Directory (AD) and the Endpoint Context Agent (ECA) service is available, the information is shown in separate panes in the Dashboard view.

3 The **User Behavior Events** pane shows alerts related to User Alert Checks. There are a set of system User Check Alerts, and you can add your own custom alerts. After configuring the rules, the generated alerts are shown here.

4 The **Tools** menu allows you to organize the information in the pane by Activity, User, User Alert Check Type, User Alert, and Severity.

Follow these general steps to configure showing users in the Dashboard view:

1) Enable the showing of user information in the Dashboard view.

2) (Optional) Create custom User Alerts.

3) Define rules that generate User Alerts.

Enable showing users on the Dashboard

To monitor users in the Dashboard view, you must enable the option in the global system properties.

The settings defined in the global system properties apply to all administrators in all Administrative Domains.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Settings** > **Global System Properties**.
2) On the **Global Options** tab, select **Show Users in the Dashboard View**.
3) Configure the other settings in the **User Information** section.
4) Click **OK**.

Create User Alerts for User Alert Checks

When users exceed a threshold defined in a User Alert Check, User Alerts are generated.

In the Custom Alert, use the Situation type User Behavior Check.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.
2) Right-click **Alert Configurations**, then select **New** > **Custom Alert**.
3) Configure the settings, then click **OK**.

Define rules that generate User Alerts

You can use different kinds of checks that generate User Alerts.

To see all the available User Alert Checks, select **Engine Configuration**, then browse to **Administration** > **Other Elements** > **User Alert Checks**.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Settings** > **Configure User Alerts**.
2) Click **Add**, then select **More actions** > **New** > **User Alert Check**.

3) Configure the settings, then click **OK**.

4) To disable a User Alert Check, deselect the checkbox to the left of the User Alert Check.
If you do not want to use a system User Alert Check, you must disable the check.

5) Click **OK**.

Monitoring connections, block lists, VPN SAs, users, routing, SSL VPNs, and neighbors

Engines track allowed connections, active VPN SAs, active users, routing, SSL VPN sessions, and directly connected neighbors in the network. Engine, Layer 2 Engine, and IPS engines also track combinations of IP addresses, ports, and protocols that are block listed.

You can monitor in the following ways:

- View currently open connections, enforced block list entries, active VPN SAs, active users, routing, SSL VPNs, and directly connected neighbors in the network.
- Save, view and compare snapshots of currently open connections, enforced block list entries, active VPN SAs, active users, routing, and SSL VPN sessions.



Note

To monitor LLDP neighbors, LLDP must be enabled for the Security Engine. If LLDP is not enabled, the Neighbor Monitoring view only shows ARP and IPv6 neighbor discovery protocol (NDP) entries.

To monitor users by name, you must enable the logging of user information in the Engine IPv4 and IPv6 Access rules. When monitoring users, you can only monitor the users connected to a particular Security Engine. To see a summary of the activity of all active users, enable showing users in the Dashboard view.

Related tasks

Define logging options for Access rules on page 901
Enable showing users on the Dashboard on page 235

View connections, block lists, VPN SAs, users, routing, SSL VPNs, and neighbors

There are several views in which you can monitor the status of the system.

The Block list view does not show whether connections matching the entries are blocked by the Engine, Layer 2 Engine, or IPS engine. Block list entries are added and removed according to their duration. The Block list view does not show any history of entries that have already expired. Use the Logs view to see information about actual connections that are allowed and denied. Depending on the logging options selected in the policy, you can also use the Logs view to see past block list entry creation.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click an Security Engine, then select **Monitoring** > **<view type>**.

   To open the Neighbor monitoring view for an Security Engine cluster, right-click an individual node in the cluster, then select **Monitoring** > **Neighbors**.

   Tip

   If the **Drill-downs** pane is open, you can also select **Monitoring** > **<view type>** from there.

2) To browse the data, select an option from the toolbar.

   You can adjust the view and the displayed data in a similar way as in the Logs view.

3) To filter data or open or compare snapshots, select the **Filter** tab or the **Snapshots** tab in the Query pane.

4) To combine entries having the same value for a given column type, right-click the heading of the corresponding column, then select **Aggregate by <column name>**.

5) Select one or more entries in the table, then right-click for a menu of actions you can take.
   - To view more information about related log events, select **Show Referenced Events**.
   - To block list connections manually, select **New Block list Entry** or **New Entry**.
   - To close a connection in the Connections view, select **Terminate**.
   - To remove a block list entry in the Block list view, select **Remove Entry**.
   - To force an SA to renegotiate in the VPN SAs view, select **Delete**.
   - To close the end user's session in the Users view, select **Delete**.

Related concepts

Getting started with the Logs view on page 281
Browsing log data on page 288
Block listing traffic and how it works on page 1093

Related tasks

Block list traffic manually on page 1098

Terminate connections manually

In the Connections view, you can manually terminate any current connection.

For example, you can remove an inactive connection that has not been properly closed. Terminating an open connection alone does not prevent any new connection from opening again.

You can terminate connections manually on the following types of engines and interfaces:

- Layer 3 physical interfaces on Engines
- Inline Interfaces on Engines, IPS engines, and Layer 2 Engines

Note

You cannot terminate connections manually on Capture Interfaces on Engines, IPS engines, or Layer 2 Engines.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click an Security Engine, then select **Monitoring** > **Connections**.
2) Select one or more connections in the table.
3) Right-click a selected row, then select **Terminate**.

Save monitoring snapshots

You can save snapshots of connections, block lists, VPN SAs, users, routing, SSL VPNs, and neighbors so that you can later view and compare them, or restore earlier versions of elements.

The saved snapshots are stored in the following directories.

Snapshot storage directories

Snapshot type	Server	Directory
Block list	Log Server	<installation directory>/data/storage/snapshots/BlackList/
Connection	Log Server Management Server	<installation directory>/data/storage/snapshots/connections/
VPN SA	Log Server	<installation directory>/data/storage/snapshots/VPN SAs/
User	Log Server	<installation directory>/data/storage/snapshots/users/
Routing	Log Server	<installation directory>/data/storage/snapshots/routing/
SSL VPN	Log Server	<installation directory>/data/storage/snapshots/SSL VPNs/
Neighbor	Log Server	<installation directory>/data/storage/snapshots/Neighbors/

Note

If you installed the Management Server in the `C:\Program Files\Forcepoint\SMC` directory in Windows, some program data might be stored in the `C:\ProgramData\Forcepoint\SMC` directory.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click an Security Engine, then select **Monitoring** > **<view type>**.

Tip

If the **Drill-downs** pane is open, you can also select **Monitoring** > **<view type>** from there.

2) To select the entries for the snapshot, click **Pause**.
   The system automatically creates a temporary snapshot of the currently displayed entries. The name of the temporary snapshot is displayed on the **Snapshots** tab in the **Query** pane. The temporary snapshot is automatically deleted.

3) Click **Save**.

4) Enter a name for the snapshot, then click **OK**.
   The name of the snapshot is displayed on the **Snapshots** tab in the **Query** pane.

Export monitoring snapshots

You can export snapshots from the Monitoring view to save stored snapshots elsewhere (for example, on your local workstation).

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Monitoring Configuration**.

2) Browse to **Other Elements** > **Monitoring Snapshots**, then browse to one of the following:
   - **Block list** > **Log Server**
   - **Connections** > **Log Server** or **Management Server**
   - **Logs** > **Management Server**
   - **Routing** > **Log Server**
   - **SSL VPNs** > **Log Server**
   - **Users** > **Log Server**
   - **VPN SAs** > **Log Server**
   - **Neighbors** > **Log Server**

3) Right-click a snapshot, then select **Export**.

4) Select the location to save the snapshot to, then click **Select**.

View monitoring snapshots

The snapshots are listed in the Monitoring view. You can also open snapshots saved on your local workstation.

View snapshots stored on SMC servers

You can view the snapshots you have saved.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Monitoring Configuration**.

2) Browse to **Other Elements** > **Monitoring Snapshots**, then browse to one of the following:
   - **Block list** > **Log Server**
   - **Connections** > **Log Server** or **Management Server**
   - **Logs** > **Management Server**
   - **Routing** > **Log Server**
   - **SSL VPNs** > **Log Server**
   - **Users** > **Log Server**
   - **VPN SAs** > **Log Server**
   - **Neighbors** > **Log Server**

3) Right-click a snapshot, then select **Open**.

View snapshots stored on local workstations

If you have exported a snapshot to your local workstation, you can later view the snapshot.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Monitoring Configuration**.

2) Browse to **Other Elements** > **Monitoring Snapshots**, right-click one of the items in the tree, then select **Open Local Snapshot**.

3) Select the snapshot, then click **Open**.

Compare monitoring snapshots

You can compare snapshots of connections, block lists, VPN SAs, users, routing, SSL VPNs, and neighbors with another snapshot of the same type. You can also compare a snapshot with the current block list, connections, VPN SAs, users, or routing entries.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Monitoring Configuration**.

2) Browse to **Other Elements** > **Monitoring Snapshots**, then browse to one of the following:
   - **Block list** > **Log Server**
   - **Connections** > **Log Server** or **Management Server**
   - **Logs** > **Management Server**
   - **Routing** > **Log Server**
   - **SSL VPNs** > **Log Server**
   - **Users** > **Log Server**
   - **VPN SAs** > **Log Server**
   - **Neighbors** > **Log Server**

3) Right-click the snapshot, then select **Open**.

4) In the **Query** pane, click the **Snapshots** tab.

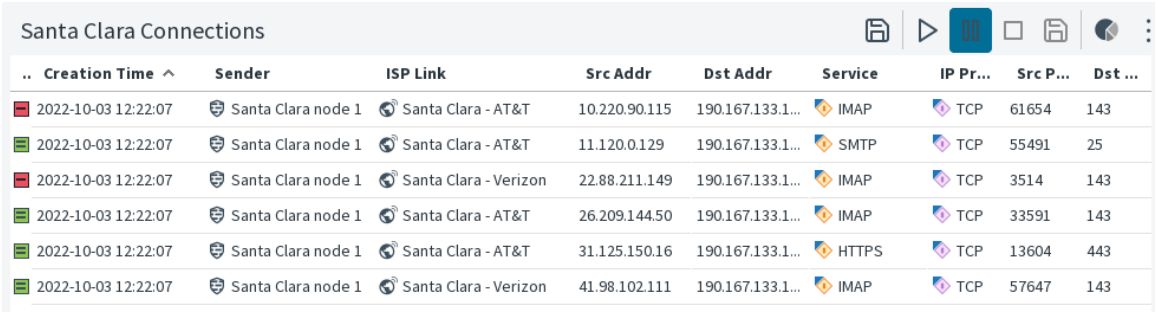
5) Using one of the following methods, select the first snapshot for comparison:
   - To compare a temporary snapshot of the current entries with a saved snapshot, click **Pause** to create the temporary snapshot. The name of the temporary snapshot is displayed in the first snapshot selection field.
   - Otherwise, click **Select** next to the first snapshot selection field, then select a snapshot.

6) Select the **Compare with** checkbox.

7) Click **Select** next to the second snapshot selection field, then select a second snapshot.

8) Click **Apply**.
   The results of the comparison are highlighted.

Snapshot comparison results

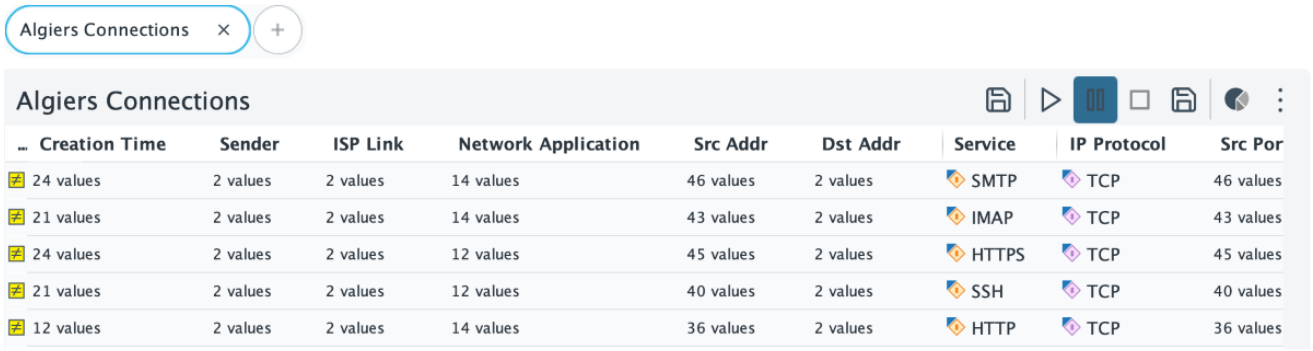


..	Creation Time	Sender	ISP Link	Src Addr	Dst Addr	Service	IP Pr...	Src P...	Dst ...
	2022-10-03 12:22:07	Santa Clara node 1	Santa Clara - AT&T	10.220.90.115	190.167.133.1...	IMAP	TCP	61654	143
	2022-10-03 12:22:07	Santa Clara node 1	Santa Clara - AT&T	11.120.0.129	190.167.133.1...	SMTP	TCP	55491	25
	2022-10-03 12:22:07	Santa Clara node 1	Santa Clara - Verizon	22.88.211.149	190.167.133.1...	IMAP	TCP	3514	143
	2022-10-03 12:22:07	Santa Clara node 1	Santa Clara - AT&T	26.209.144.50	190.167.133.1...	IMAP	TCP	33591	143
	2022-10-03 12:22:07	Santa Clara node 1	Santa Clara - AT&T	31.125.150.16	190.167.133.1...	HTTPS	TCP	13604	443
	2022-10-03 12:22:07	Santa Clara node 1	Santa Clara - Verizon	41.98.102.111	190.167.133.1...	IMAP	TCP	57647	143

The icons in the first column signify whether the entry has changed.

Icon	Description
	Entry only exists in snapshot 1
	Entry is included in both snapshots
	Entry only exists in snapshot 2
	Aggregated entries
No icon	There are more than 100 entries that match the **Aggregate by <column name>** selection. No further comparison can be done.

Example of snapshot comparison with entries aggregated by service



Tip

To open a snapshot for comparison directly in the Monitoring view, right-click the snapshot, then select **Compare to Current**.

Related concepts

Browsing log data on page 288

View and compare Element Snapshot elements

You can view earlier configurations of an element and compare them to the current configuration with Element Snapshots. You can also restore earlier configurations of an element.

Element Snapshots are automatically generated and saved in Audit logs each time element properties are saved. An Element Snapshot contains all properties of an element saved in the **Properties** dialog box.

Steps

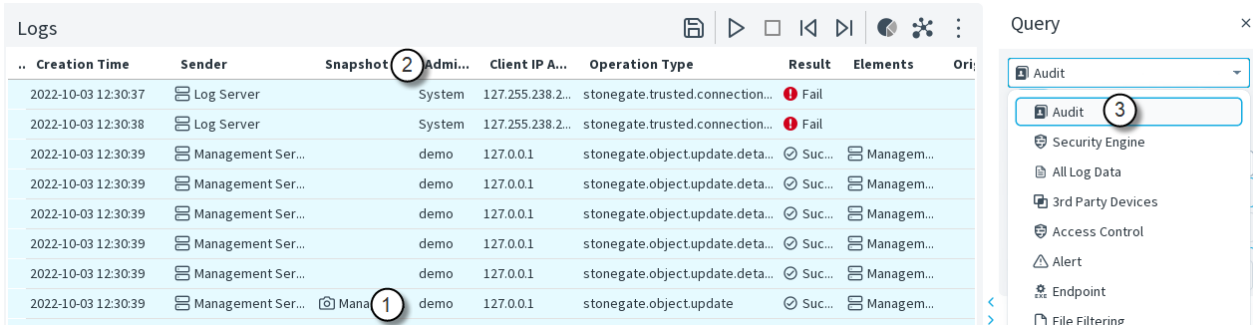
For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Logs**.

2) Select **Audit** from the drop-down list in the Query pane, then click **Apply**.

   All Element Snapshots generated during the defined time range are displayed in the Snapshot column of the log entry table.

> **Tip**
>
> If the column is not displayed, to add it, select **More actions** > **Columns** > **Column Selection**.



1	Element snapshot
2	Snapshot column
3	Audit selected

3) To view an Element Snapshot in more detail, right-click the Audit entry, then select **View Element Snapshot**.

4) To compare an Element Snapshot to the current configuration of the same element, right-click an Audit entry, then select **Compare to Current Element**.



Note

If the Element Snapshot properties differ from the current element properties, a red border is displayed around the Audit Log Version (snapshot) and Current Version of the element.



Tip

To display all values of the snapshot and the current element in XML format with differences indicated in red, select **Show: XML**.

5) To close the **Compare Elements** dialog box, click **OK**.

Related concepts

Benefits of exporting or importing elements on page 189

Related tasks

Select columns in the log entry table on page 300

Monitoring connections using Geolocation elements

Using Geolocation elements, you can define the physical locations of network elements, such as Hosts.

You can also see, for example, how much traffic the elements create.

This product includes GeoLite data created by MaxMind, available from https://www.maxmind.com. The location information for public IP addresses is based on the GeoLite data created by MaxMind. The IP addresses in the Geolocation database are updated when the Management Server is upgraded.

You can also create Geolocations manually in the SMC. Geolocations based on internal IP addresses must be configured manually, as these addresses cannot be found in the Geolocation database based on public IP addresses. Creating a Geolocation manually for a public IP address overrides location data found for the address in the Geolocation database.

Geolocation maps are available in Reports, Statistics, and Overviews.



Note

You cannot use Geolocation elements to filter network traffic. Use Country elements for that purpose.

Create Geolocation elements

A Geolocation represents the physical location of a network element. When you create a Geolocation, select the elements that belong to it.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.
2) Browse to **Other Elements** > **Geolocations**.
3) Right-click **Geolocations**, then select **New Geolocation**.
4) Enter the **Name** and **Address**.
5) To define the Coordinates, select an option:
   - To automatically resolve the Geolocation coordinates, click **Resolve from Address**.
   - Enter the **Latitude** and **Longitude** in Decimal Degrees format (for example, latitude 49.5000° and longitude -123.5000°).
6) Click the **Logo** tab, then select or upload a new custom logo image and then select it for the Geolocation.

   Note

   - A custom logo can only be used for a custom created Geolocation.
   - If a custom logo is defined for an IP Address and the IP address belongs to a custom Geolocation, the logo is displayed in the logs instead of the country flag.
7) Click the **Content** tab, then select the elements that belong to the Geolocation.
8) Click **OK**.

Set a Geolocation for an element

You can set many elements to use the same Geolocation.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click an element, then select **More actions** > **Set Geolocation**.
2) Select a Geolocation for the element, then click **Select**.

Geolocations and IP addresses in Google Maps

You can view the actual location of a Geolocation element or an IP address in more detail in Google Maps.

You can use the Show in Google Maps option in:

- Geolocation maps included in Overviews and Reports
- The Logs view
- The **Whois Information** dialog box

View Geolocation element locations in Overviews and Reports

Geolocation elements are displayed in Google Maps based on the location data that was entered for them.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) On a Geolocation map in an Overview or Report section, right-click a location, then select **Show in Google Maps**.
   The location is opened in Google Maps in your default web browser.

Related concepts

Overviews and how they work on page 227
Designing reports on page 315

View IP address locations in the Logs view

In the **Logs** view, you can see a location (for example, a city or street address) for an IP address in Google Maps.

Only IP addresses associated with a location can be displayed in Google Maps. In the **Logs** view, these IP addresses are indicated with a country flag icon next to the IP address.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Logs**.

2) Select a log entry with an IP address that has a flag icon associated with it.

3) Right-click the IP address in the **Fields** pane, then select **Show in Google Maps**.
   The location is opened in Google Maps in your default web browser.

View IP address locations from the Whois Information dialog box

When checking the Whois information for an IP address, you can see the location (for example, a city or street address) in Google Maps.

Only IP addresses associated with a location (for example, a city or street address) can be displayed in Google Maps.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Logs**.

2) Right-click the IP address of a log entry that has a flag icon associated with it, then select **Whois <IP address>**.

3) Right-click the **Whois Information** dialog box, then select **Show in Google Maps**.
   The location is opened in Google Maps in your default web browser.

Monitoring configurations and policies

The engines receive their configuration when a policy is installed.

You can monitor the policies and configurations installed on the engines in the following ways:

- You can check which policy is being enforced and when it was last installed.
- You can quickly view the most recent version of the installed policy.
- You can view the configurations that were transferred in each past policy installation and compare them to each other or the current policy stored on the Management Server.
- You can view, approve and commit the changes to the configuration and policies of engines that have not been transferred yet.

Related concepts

Getting started with policies on page 799

Related tasks

Check the currently installed policy on page 813
View, approve, and commit pending changes on page 107

Monitor administrator actions

A record of administrator actions is maintained in the SMC.

The records can only be viewed by administrators who are allowed to view Audit logs. They can be browsed like any other logs in the **Logs** view.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Settings** > **Audit**.
   The **Logs** view opens with the Audit logs selected for viewing.

2) Browse and filter the logs.

Related concepts
Getting started with the Logs view on page 281

Monitor tasks

You can check the status of running tasks and executed tasks (for example, upgrade tasks and system tasks).

Note

System tasks that run automatically are not visible in the **History** branch.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.

2) Browse to **Tasks** > **History**.
   You can see the following columns:
   - **Progress** — Shows the progress of a running Task.
   - **Info** — Shows additional details about the execution of a Task.
   - **State** — Shows the status of the Task.

3) To show tasks that have previously been run, select **More actions** > **Show Executed Tasks**.

Related concepts
Creating Task Definitions on page 1322

Related tasks

Schedule Tasks on page 1327
Start Tasks manually on page 1328
Stop running Tasks on page 1330

Traffic captures and how they work

You can capture network traffic data for network troubleshooting purposes. This data helps you to analyze network traffic to and from the engines.

It is also often useful to have this data available when contacting Forcepoint Customer Hub.

Traffic capture creates a .zip file that contains a tcpdump CAP file, which is compatible with standard “sniffer” tools such as tcpdump, WinDump, or Wireshark. You can select whether to include full packet information or only IP address headers in the tcpdump. You can also include a free-form description and information about your configuration and trace files in the traffic capture .zip file.

The data can be archived and analyzed later, as the traffic capture .zip file is saved on the Management Server or in a directory on your local workstation.

Traffic captures can only be taken on nodes that are online and have a policy uploaded.



Note

You must have permissions to send Advanced Commands to be able to take traffic captures.

You can stop or cancel a traffic capture at any point once it has been started.

- If you stop a traffic capture, all captured tcpdump data is compressed and sent to the Management Server or to your local workstation.
- If you cancel a traffic capture, all captured tcpdump data is deleted.

Related tasks

Create Administrator Role elements on page 378

Take traffic captures

If you want to analyze network traffic, capture the network traffic data.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click an Security Engine, then select **More actions** > **Capture Traffic**.
3) Select one or more interfaces whose traffic you want to capture, then click **Select**.

4) (Optional) Click **Add** to add more interfaces to the traffic capture.
You can also add interfaces from other types of engines.



Tip

You can create tcpdump files for several different interfaces in the same Traffic Capture task. The Traffic Capture .zip file contains a separate CAP file for each interface included in the capture.

5) (Optional) To limit the scope of the traffic capture, click the **Limit by** field, then enter an IPv4 or IPv6 address.
The IP address must match either the source or destination of the packets included in the capture.

6) Define the other traffic capture options.

7) Click **Start Capture**.

Related concepts
Alert log messages for troubleshooting on page 1375
Log messages for troubleshooting on page 1377
Error messages for troubleshooting on page 1382

Checking maintenance contract information

You can view maintenance contract and support level information for your licenses in the SMC Client by allowing your Management Server to contact Forcepoint servers.

This information is available for each license in the **Licenses** > **All Licenses** branch of the **Administration Configuration** view, if the Management Server can contact the servers.

To enable viewing maintenance contract and support level information permanently for your licenses, you must allow the Management Server to connect to the servers.

Related concepts
Getting started with automatic updates and upgrades on page 1291

Enable or disable automatic checking of maintenance contract information

You need to allow the Management Server to connect to Forcepoint servers.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Settings** > **Global System Properties**.
2) On the **Updates** tab, select **Allow Sending License and Installation Telemetry Data to Forcepoint Servers**.
3) Click **OK**.

View current maintenance contract information

When contract checking is enabled, you can view information on your support contract.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.
2) Browse to **Licenses**, then browse to **All Licenses** or a component-specific branch.
   The licenses are shown in the table in the right pane.
3) Select the license that you want to view.

Note

If information is not available, make sure that you have enabled the automatic checking of maintenance contract information or manually fetch the latest information.

Manually fetch latest maintenance contract information

If you do not allow contacting Forcepoint servers automatically, you can fetch the information manually.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.

2) Browse to **Licenses**, then browse to **All Licenses** or a component-specific branch.

3) Right-click the selected branch, then select **Check Maintenance Contract** or **More actions** > **Check Maintenance Contract**.
You are prompted to confirm that you want to send proof of license codes to Forcepoint.

4) Click **Yes**.
If the Management Server can connect to Forcepoint servers, the SMC Client displays the maintenance contract and support level information for the licenses. The information is available in the SMC Client until the Management Server is restarted.

Upcoming event notification

The upcoming events notification feature informs about events that are happening soon, such as expiration of licenses and certificates, and failures of scheduled tasks, that require administrator action.

Related tasks
Global settings on page 252
Profile settings on page 253
Events on page 253

Global settings

The settings defined in the **Global settings** apply to all administrators in all Administrative Domains.

Global settings dialog box is used to enable or disable reporting of event categories for all SMC administrators.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Click the **Events** icon.

2) Click the **Global settings** tab.

3) Select or deselect the check boxes under **On** to enable or stop reporting of certain event categories.

Global settings dialog box

Use the **Global settings** dialog box to set the event notification settings for all SMC administrators. SMC administrators with unrestricted permissions (superusers) can change the global settings.

Option	Definition
On	Select or deselect the checkbox to enable or stop reporting of event categories.

Option	Definition
Events	Event category description.
Days	Period left when event needs to be reported.

Profile settings

The settings defined in the **Profile settings** apply only to specific administrator accounts.

Settings under **Profile settings** tab:

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Click the **Events** icon on the top right corner.
2) Click the **Profile settings** tab.
3) Select or deselect the check boxes under **On** to enable or stop reporting of certain event categories.

Profile settings dialog box

Use the **Profile settings** dialog box to set the event notification settings for a specific administrator account.

Option	Definition
On	Select or deselect the checkbox to enable or stop reporting of event categories.
Events	Event category description.

Events

The Events dialog box, shows the list of upcoming events that require administrator action.

- Expiration of licenses.
- Expiration of VPN certificates.
- Expiration of HTTPS certificates for communication with external components and browser-based user authentication.
- Failure of scheduled tasks.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Click the **Events** icon on the top right.
2) Click the **Events** tab to view the list of upcoming events.

Events dialog box

The upcoming **Events** dialog box shows a list of upcoming events that require administrator action and the date by which the action must be completed.

Option	Definition
Due Date	Date by which the administrator should fix the issue.
Event	Event description
Element	Element that reported the event.
Related Element(s)	Impacted element like the certificates or the task itself.
Actions	**Where Used?**: To be able to retrieve references of the element **View Element**: To be able to drive the element on the corresponding folder **Hide events like this from me**: To be able to deselect from my profile, this kind of events.

Chapter 13

Application Health Monitoring

Contents

- Enable Application Health Monitoring on page 255
- View status of Network Applications on page 256
- Disable Application Health Monitoring on page 260

The **Application Health Monitoring** dashboard lets administrators monitor network and application layers connection quality.

The Engine measures the network application quality metrics, such as network latency, packet-loss/retransmissions, and application latency. These quality metrics are then collected by SMC and presented as application health on the **Application Health Monitoring** dashboard.

Enable Application Health Monitoring

To monitor application health and user experience on the **Application Health Monitoring** dashboard, you must enable this feature in the **Global System Properties** page.

Note

The settings defined in the **Global System Properties** apply to all administrators in all Administrative Domains.

Steps

1) Select **Settings** > **Global System Properties**.
   Alternatively, you can click the **Configuration** icon on the left navigation pane in the **Application Health Monitoring** dashboard, and then click **Configure Application Health Monitoring Options** to navigate to the **Global System Properties** page.

2) On the **Global Options** tab, configure the following settings in the **Network Application Health Monitoring** section.
   a) In **Health Monitoring**, select either **Network Application** or **Top Network Applications**.

b) The following options are available based on your selection in **Health Monitoring**:

- Top Network Applications– Selects top network applications automatically from selected application usage categories based on traffic statistics. Click the **Preview** option to see the list of selected top network applications. The **Top limit** option lets you to set the maximum limit of network applications being monitored and displayed on the dashboard. By default, the maximum limit is 10. However, this value is configurable.
  - Network application logging with accounting data is required for top network applications. For this, you need to select **Log Accounting Information** in the **Connection Closing** section of **Logging Select Rule Options** dialog box.
  - In the **Logging Select Rule Options** dialog box, select **Enforced** for Log Network Applications in the **Logging Enforcements** section. This populates the network application information in the logs.

  Note

  This selection is recommended for both network applications and top network applications.
- Network Application– Selects individual network applications to be added for monitoring. You can add or remove an application.

c) Click **OK**

3) Create a Engine Policy and do the following:
   - In the **Select Rule Action Options** dialog box, set the **Network Application Monitoring Latency** to **On**.

4) Save the Engine Policy.

Related reference
Global System Properties dialog box — Global Options tab
Select Rule Action Options dialog box (Allow)
Logging - Select Rule Options dialog box

View status of Network Applications

The **Application Health Monitoring** dashboard presents the application health of the network applications being monitored.

To view network application quality metrics.

Steps

1) Login to SMC using your login credentials.

2) Click **Dashboards**, and then select **Application Health Monitoring** dashboard to view the network applications being monitored.

Application Health Monitoring Dashboard

The **Application Health Monitoring** dashboard displays the health status of the network applications that are being monitored.

It includes the following widgets:

- **Application Heath:** Displays the overall health status of network applications being monitored.
- **Application Health Map:** Displays the health status of applications that are monitored for the site on the map. You can hover over a site on the map to view more details.
- **Application Latency Map:** Displays the application latencies grouped by the geographical location of the application server IP address on the map.
- **Network Latency Map:** Displays the network latencies grouped by the geographical location of the application server IP address on the map.
- **Application Health Tree:** Displays the statistics that are reported for each monitored application.
- **Top Application Latencies by Application:** Displays the chart with details of applications with top application latency.
- **Top Packet Loss by Application:** Displays the chart with details of applications with top packet loss.
- **Top Network Latencies by Application:** Displays the chart with details of applications with top network latency.

You can click the **settings** icon in the upper-right corner of the **Application Health Monitoring** dashboard to configure the widgets to display in the dashboard. By default, the following widgets are displayed in the dashboard:

- Application Health
- Application Latency Map
- Application Health Tree

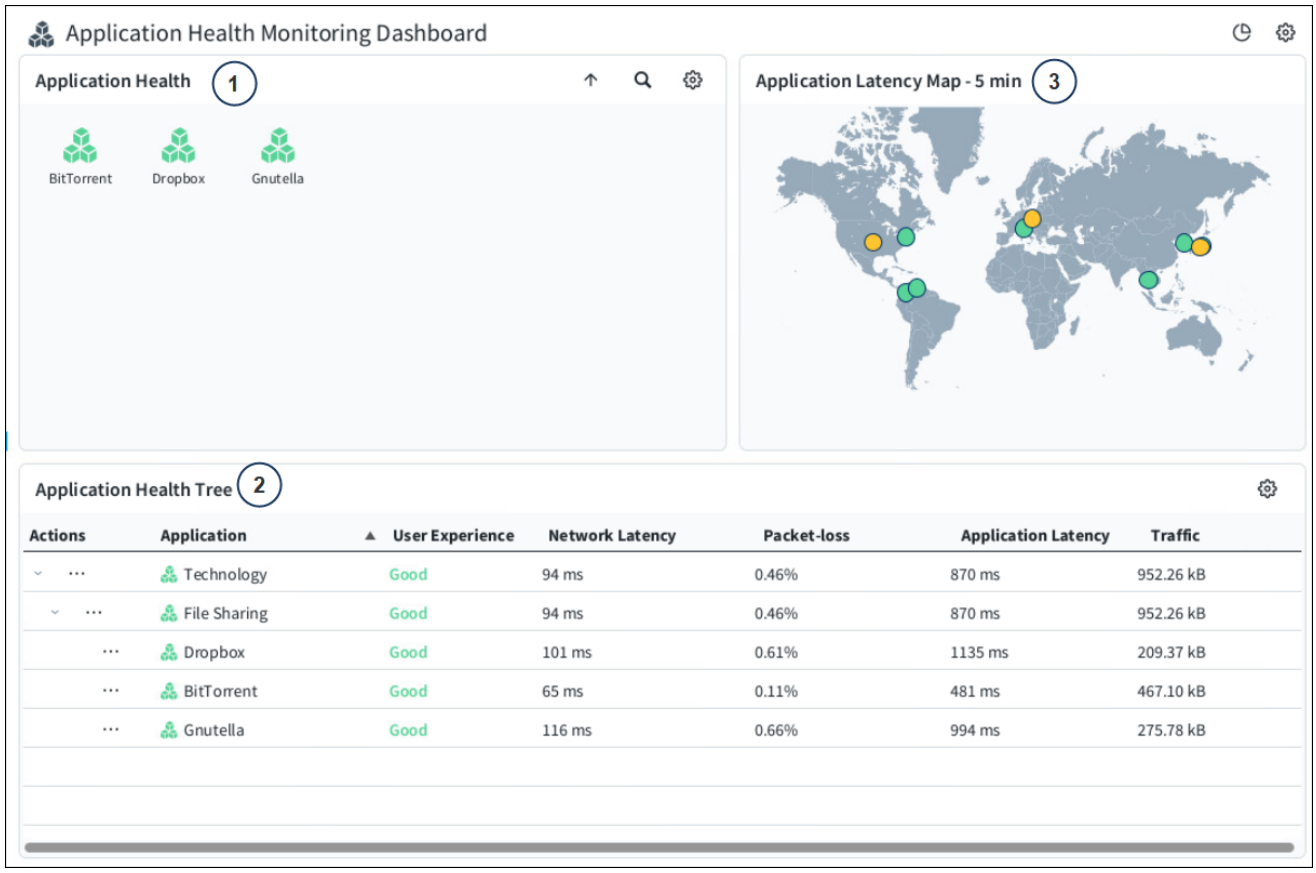
Also, you can click the **Statistics Time-Range** icon in the upper-right corner of the **Application Health Monitoring** dashboard to select the time-range between which you want to view the health status history of network applications in the dashboard. The details in the widgets are updated as per the selected time-range.

Note

1) To monitor application layer latency on TLS encrypted traffic, the TLS inspection must be enabled for the monitored traffic.

2) To accurately monitor health of the UDP traffic of an application that needs active probing, you must enable active monitoring for the applications in the access policy **Action** options. This allows the engine to use active probing as needed for the application traffic that match the rule of the access policy.

Application Health Monitoring Dashboard



1 Displays the overall health status of network applications being monitored and categorizes the applications as follows:

- **Good (Green)** – When all of the following is true:
  - Application latency value is less than 1200 milliseconds.
  - Network latency value is less than 150 milliseconds.
  - Packet loss value is less than 70 per-myriad.
- **Fair (Yellow)** – When one of the following is true:
  - Application latency value is more than 1200 milliseconds but less than 4800 milliseconds.
  - Network latency value is more than 150 milliseconds but less than 250 milliseconds.
  - Packet loss value is more than 70 per-myriad but less than 250 per-myriad.
- **Poor (Red)** – When one of the following is true:
  - Application latency value is more than 4800 milliseconds.
  - Network latency value is more than 250 milliseconds.
  - Packet loss value is more than 250 per-myriad.

Also, you can do the following from the **Application Health** widget:

- You can hover over an application to view more details related to the application health.
- You can hover over an application, and then click one of the following icons:
  - **Logs by situation icon:** Opens the **Logs** page with the log details filtered for the application.

- **Connections icon:** Opens the **Connections** page with the network application details for the application.
- **Status History icon:** Opens the **Application health status history** dashboard. Its displays the health status history trend, network latency, application latency trend, packet loss trend, traffic trend, and connection trend for the application. Also, you can select the time interval from the drop-down list in the upper-right corner of the dashboard to display the application health monitoring information for the selected time interval.

2 Displays the application latencies grouped by the geographical location of the application server IP address on the map..



Note

You can hover over a site on the map to view more details related to all applications for the site.

3 Displays the following statistics that are reported for each monitored application:

- **Action Menu:** Click the **…** menu under the **Action** column in a specific row to view the following action options for each application or application category:
  - **Logs by Situation:** Opens the **Logs** page with the log details filtered for the application.
  - **Connections:** Opens the Connections page with the network application details for the application.
  - **Status History:** Opens the Application health status history dashboard. Its displays the health status history trend, network latency, application latency trend, packet loss trend, traffic trend, and connection trend for the application. Also, you can select the time interval from the drop-down list in the upper-right corner of the dashboard to display the application health monitoring information for the selected time interval.
- **Application:** Displays the application category name or the application name.
- **User Experience:** Displays the overall application experience impacted by network and application latency, traffic transfer rate and any packet loss in transmission. The values that are displayed are Good (green), fair (yellow), and poor (red). Also, the time is displayed along with the user experience value to indicate when the last status change has occurred. For example, the **Fair (5 min)** value indicates that the status has changed from **Good** to **Fair** 5 minutes ago.
- **Network Latency:** Display the amount of time in milliseconds, the data packet takes to travel across the network for each application based on application health monitoring values received from engines during a periodic interval. By default, 5 minutes is used as the periodic interval to receive the application health monitoring values, which can be customized in SMC configuration file.
- **Packet-loss:** Displays the percentage of packets that are lost after being transmitted across the network for each application based on application health monitoring values received from engines during a periodic interval. By default, 5 minutes is used as the periodic interval to receive the application health monitoring values, which can be customized in SMC configuration file.
- **Application Latency:** Display the amount of time in milliseconds, the application takes to respond based on application health monitoring values received from engines during a periodic interval. By default, 5 minutes is used as the periodic interval to receive the application health monitoring values, which can be customized in SMC configuration file.
- **Traffic:** Displays the average transfer rate of traffic in the network for each application based on application health monitoring values received from engines during a periodic interval. By default, 5 minutes is used as the periodic interval to receive the application health monitoring values, which can be customized in SMC configuration file.

Disable Application Health Monitoring

You can disable the Application Health Monitoring feature with either of the following ways:

Steps

- Select **Disable Application Health Monitoring** option in the **Advanced Settings** screen to disable the feature for a specific Engine.
- On the **Global Options** tab, select **None** for **Health Monitoring** option to disable the feature for all Engines at once.

Related reference
Engine Editor > Advanced Settings

Chapter 14

Monitoring third-party devices

Contents

- Getting started with monitoring third-party devices on page 261
- Converting logs from third-party devices on page 262
- Methods for monitoring third-party device status on page 273
- Configuring Log Servers to monitor third-party devices on page 276
- Activate monitoring of third-party devices on page 277
- Configuring third-party devices for monitoring on page 278
- Changing the ports for third-party device monitoring on page 279
- Activate or deactivate third-party device monitoring alerts on page 280

The SMC can be configured to log and monitor other manufacturers' devices in much the same way as SMC components are monitored.

Getting started with monitoring third-party devices

As well as SMC components, you can monitor third-party devices, with some limitations.

What the third-party device monitoring feature does

You can configure Log Servers for a full range of monitoring features for third-party devices:

- Log Servers can receive a syslog stream and store the information in SMC log format. The stored logs can then be processed just like logs generated by SMC components. For example, the data can be included in reports you generate.
- Log Servers can receive SNMP statistics information and NetFlow (v5 and v9) and IPFIX data from third-party devices. You can view this information as part of your Overviews or create reports based on the received data.
- Log Servers can probe devices through several alternative methods. You can monitor the device status in the SMC Client in the same way as for the SMC components.

Limitations

Each Log Server can monitor a maximum of 200 third-party devices. This limit is enforced automatically.

SNMP statistics for third-party devices are limited to the amount of free and used memory, CPU load, and interface statistics.

Your Management Server license might limit the number of managed components. Each monitored third-party device is counted as a fifth of a managed unit.

Third-party device monitoring configuration overview

The steps you follow depend on the types of third-party devices you want to monitor.

Follow these general steps to configure the monitoring of third-party devices:

1) (Optional, for syslog data only) If you want to receive syslog data, define the syslog reception settings for each type of third-party device.

2) (Optional) If you want to monitor the status of third-party devices and receive statistics from them, define the status and statistics monitoring settings for each type of device.

3) Activate monitoring for the third-party device by adding monitoring settings in the properties of each element that represents a third-party device. (Third-party devices are: Router, Host, Active Directory Server, LDAP Server, RADIUS Authentication Server, and TACACS+ Authentication Server.)

4) Depending on features used, configure the third-party device for monitoring.

Related concepts

Converting logs from third-party devices on page 262
Methods for monitoring third-party device status on page 273
Configuring third-party devices for monitoring on page 278
Changing the ports for third-party device monitoring on page 279
Configuring Log Servers to monitor third-party devices on page 276

Related tasks

Activate monitoring of third-party devices on page 277
Activate or deactivate third-party device monitoring alerts on page 280

Converting logs from third-party devices

You can set up most external devices to send logs to the Log Server in syslog format.

The Log Server can convert incoming syslog entries to SMC log entries. You can use predefined Logging Profile elements or create new elements to determine how the field values are selected from a syslog entry and inserted into an SMC log entry. A Logging Profile must have at least one *logging pattern*. Logging patterns determine how the fields from syslog entry are parsed to the appropriate log fields in an SMC log entry.

You can create logging patterns in the following ways:

- *Ordered Fields* — Use when the fields in the syslog message are not separated by keywords and the type of field can only be deduced from its position. The received syslog entries are parsed in a sequence that you define in the Logging Profile. If the incoming logs vary in structure, you must define a different sequence for each type of structure. You can define several patterns in one Logging Profile.
- *Key-Value Pairs* — Use when the syslog message contains keywords that describe the type of field. The received syslog entries are parsed based on key values that you define in the Logging Profile. You can define the key values in any order. A single definition can be used even if logs vary in structure.

It is easier to configure a pattern using key-value pairs. We recommend that you use key-value pairs if a third-party device formats the relevant parts of the syslog packet as key-value pairs. Ordered fields can be used to process all syslog data regardless of its format, but it is more difficult to configure a pattern as ordered fields.

If a match is found, the system simply converts the matching syslog entry to an SMC log field. You can define Field Resolvers for more complex operations.

You can categorize incoming logs from third-party devices by selecting specific Log Data Tags for them. You can categorize logs based on the log type, or the feature or product that generates the logs. For example, you can associate the “Engine” Log Data Tag with third-party engine logs.

You can create categories by dividing the logging patterns in a Logging Profile in sections. Both ordered fields and key-value pairs can be divided into sections. You can select one or several Log Data Tags for each section. The selected Log Data Tags are shown for the matching log entries in the Fields pane of the Logs view if the Log Data Tags column is enabled. In addition to the Log Data Tags you define in the Logging Profile, the default “Third Party” and “Log Data” Log Data Tags are associated with all logs from third-party devices.

You can also use Log Data Tags as filtering criteria in the Logs view, in Reports, and in Local Filters for various elements. (Elements include: Administrator, Log Server, and Management Server elements and Correlation Situations).

Related concepts

Changing the ports for third-party device monitoring on page 279
Getting started with the Logs view on page 281
Getting started with filtering data on page 333

Related tasks

Create Logging Profile elements on page 263
Add Field Resolvers in Logging Profile elements on page 271

Create Logging Profile elements

A syslog packet consists of three parts: <PRI>, HEADER, and MSG. In a Logging Profile element, you define patterns for converting the MSG part of the syslog packet to a SMC log entry.

A Logging Profile parses the data in a syslog message to the corresponding SMC log fields when the syslog entry is converted to an SMC log entry. The parts of the syslog packet are explained in more detail in the following table.

Parts of the syslog packet

Section	Description
<PRI>	Contains facility and priority information. The Log Server automatically extracts the Facility value from the <PRI> part and converts it to the Syslog Facility field in SMC logs. You do not define patterns for mapping this section in the Logging Profile.
HEADER	Contains a time stamp and the host name or IP address of a device. The Log Server automatically extracts the data in the HEADER part. This section is optional in syslog packets, so not all devices send this data.

Section	Description
MSG	Contains the text of the syslog message. In the Logging Profile, you define the mapping for parsing this part of the syslog packet.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Monitoring Configuration**.
2) Browse to **Third-Party Devices** > **Logging Profiles**.
3) Right-click **Logging Profiles**, then select **New** > **Logging Profile**.
4) Enter a name for the Logging Profile, then click **OK**.

5) (Optional) To insert fields, drag and drop items to the **Header** field from the **Fields** branch in the Resource pane, or use type-ahead search.



Note

You can add fields that are the same for all logging patterns that you define in the **Patterns** pane. To omit a portion of data, add an **Ignore** field.

Important

Type or copy and paste from the syslog message any tokens that appear before and after the field values. If you do not insert the appropriate tokens, the data cannot be parsed.

In the illustration, the header of the syslog entry contains the following data common for all patterns:

```
<Cisco time><space><Ignore><space><Ignore><space><Cisco original time>
```

As a result, the header contains the following data:

```
<Sep 21 04:04:56> <cisco-example.stonesoft.com> <1815452:> <Sep 21 04:04:55> %
```

Because the **Ignore** field is used for `<cisco-example.stonesoft.com>` and `<1815452:>`, the values are not converted to SMC log entry format.

6) Select how Patterns are parsed:

- **Ordered Fields** — The syslog entries are parsed in the specified order. If the incoming logs vary in structure, you must define several patterns.
- **Key-Value Pairs** — The syslog entries are parsed based on key-value pairs that you define. You can add key-value pairs in any order. You can use one pattern for all logs even if the logs vary in structure.

7) Click **Save**.

Related concepts

Converting logs from third-party devices on page 262

Related tasks

Define logging patterns as ordered fields in Logging Profile elements on page 266
Define logging patterns as key-value pairs in Logging Profile elements on page 268

Define logging patterns in Logging Profile elements

You can define logging patterns in the Logging Profile using ordered fields or key-value pairs.

Define logging patterns as ordered fields in Logging Profile elements

The pattern that you define in a Logging Profile must be an exact match for the incoming syslog entry. If incoming logs vary in structure, you must define a different pattern for each type of entry.

If several patterns match, the system uses the pattern with the most matching entries.

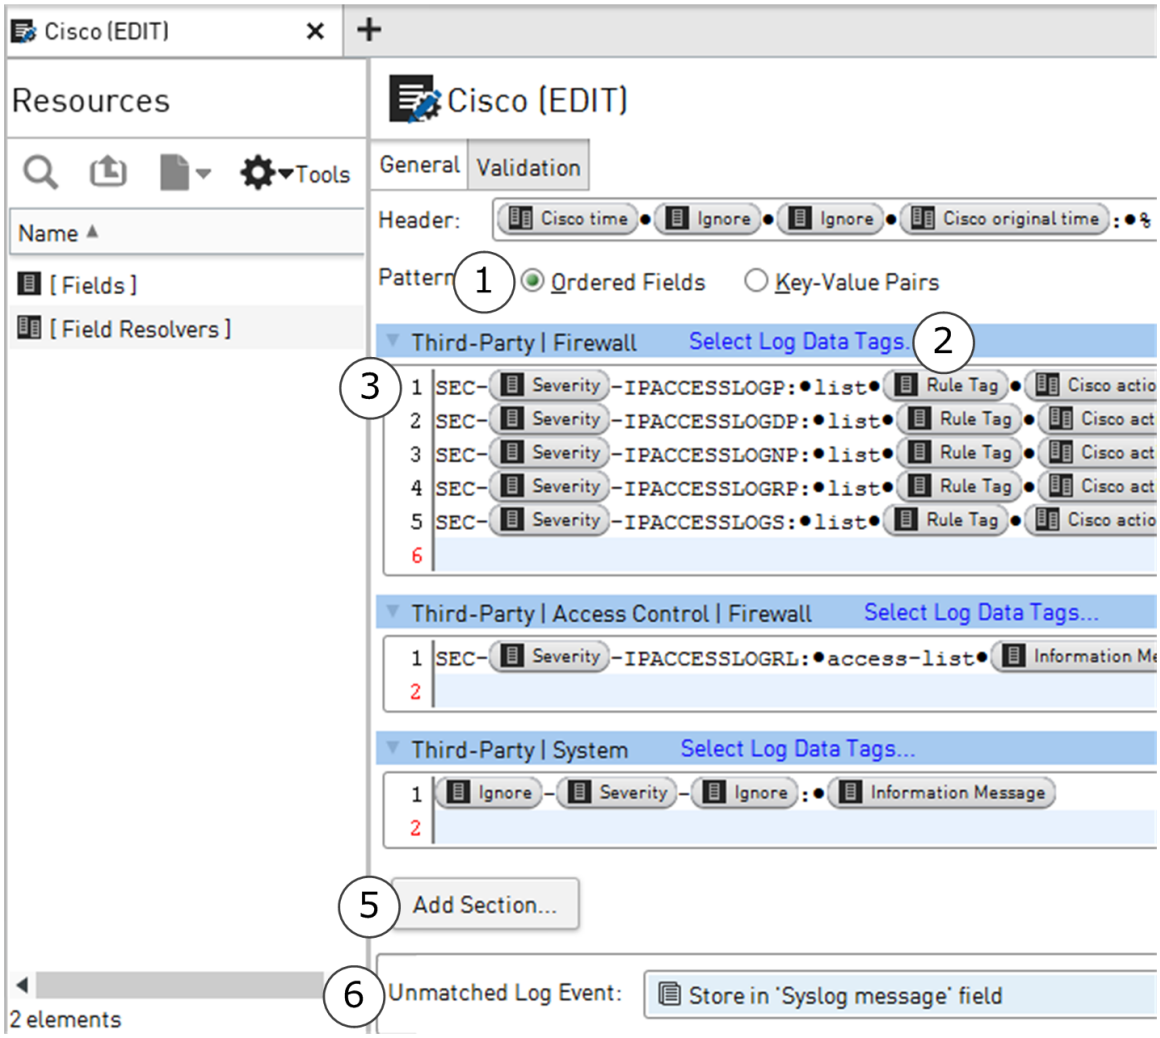
Each received syslog entry is converted to an SMC log entry. The field values that match a specified pattern are copied without further processing to an SMC log field. Also, you can create Field Resolvers to convert specific values in the syslog data to specific values in SMC logs.

You can use sections in the Logging Profile to organize the logging patterns. To create categories, you can associate one or several Log Data Tags with each section. The Log Data Tags improve the way log entries can be viewed and stored. However, they do not affect the way third-party log entries are converted into SMC log entries. If you do not select specific Log Data Tags for a section, only the default “Third Party” and “Log Data” Log Data Tags are shown for the matching log entries.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the Logging Profile, select **Ordered Fields** as the **Pattern**.



2) (Optional) In the header row of the Patterns section, click the **Select Log Data Tags** link.

Select the Log Data Tags according to the type of traffic that matches the ordered fields in the section, then click **Add**. The selected Log Data Tags are added to the Content list.

Note

Log Data Tags make the converted third-party log data records visible in the appropriate log data contexts. They also generate log data storage indexes, which speed up the filtering by data tags.

3) To insert the field values, drag and drop items from the **Fields** branch in the left pane to the empty space in the **Patterns** section. Or use type-ahead search.

   Alternatively, you can define a Field Resolver, then add it to the pattern instead of a log field. To omit a portion of data, add an **Ignore** field.

   Important

   Type or copy and paste from the syslog message any tokens that appear before and after the field values. If you do not insert the appropriate tokens, the data is not parsed.

4) (Optional) If some incoming log entries have a different structure, press **Enter** to add more rows to the Patterns section.

5) (Optional) To create another section in the same Logging Profile, click **Add Section**, then configure the new section.

6) In the **Unmatched Log Event** section, select the action for handling syslog data that does not match any defined logging patterns:
   - **Store in 'Syslog message' field** — A log entry is created and all data is inserted into the Syslog Message log field. The log entry is stored on the Log Server.
   - **Display in 'Syslog message' field (Current mode only)** — A log entry is created and all data is inserted into the Syslog Message log field. The log entry is displayed in the Current Events mode in the Logs view, but it is not stored.
   - **Ignore** — The data is discarded.

Related tasks

Add Field Resolvers in Logging Profile elements on page 271
Validate Logging Profile elements on page 272
Activate monitoring of third-party devices on page 277

Define logging patterns as key-value pairs in Logging Profile elements

When you define key-value pairs for converting syslog data, the Log Server parses each received syslog entry data based on the defined key-value pairs.

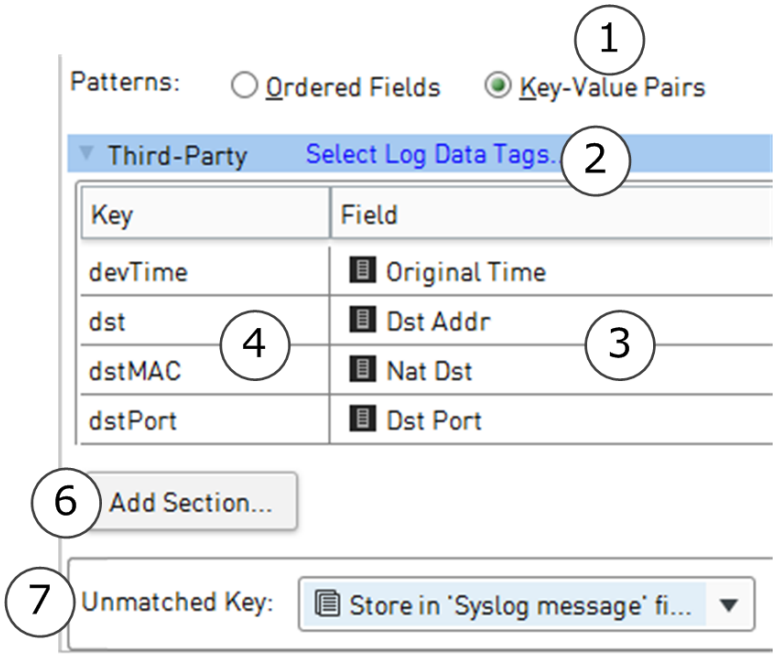
The data in the incoming syslog message must be formatted as key-value pairs.

You can use sections in the Logging Profile to organize the logging patterns. To create categories, you can associate one or several Log Data Tags with each section. The Log Data Tags improve the way log entries can be viewed and stored. However, they do not affect the way third-party log entries are converted into SMC log entries. If you do not select specific Log Data Tags for a section, only default “Third Party” and “Log Data” Log Data Tags are shown for matching log entries.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the Logging Profile, select **Key-Value Pairs** as the **Pattern**.



2) (Optional) Click the **Select Log Data Tags** link in the header row of the Patterns section.

Select the Log Data Tags according to the type of traffic that matches the key-value pairs in the section, then click **Add**. The selected Log Data Tags are added to the Content list.

Note

Log Data Tags make the converted third-party log data records visible in the appropriate log data contexts. They also generate log data storage indexes, which speed up the filtering by data tags.

3) Drag and drop SMC log fields from the **Fields** branch in the left pane to the **Field** column.

Alternatively, you can define a Field Resolver and add it to the pattern instead of a log field. To omit a portion of data, add an **Ignore** field. By default, the **Ignore** field is added to the **Field** column in the new section.

4) Double-click the **Key** column for the log field that you added, then type the corresponding key value as it appears in the syslog message (for example, `devTime`).

5) (Optional) To add more key-value pairs to a section, right-click a row, then select **Add Row**.

The key values can be added in any order. The key values are converted to SMC log entries based on the key values alone.

6) (Optional) To create another section in the same Logging Profile, click **Add Section**, then configure the new section.

7) In the **Unmatched Key** section, select the action for handling syslog data that does not match any defined logging patterns:

- **Store in 'Syslog message' field** — A log entry is created and all data is inserted into the Syslog Message log field. The created log entry is stored on the Log Server.
- **Ignore** — The data is discarded.

Related concepts
Converting logs from third-party devices on page 262

Related tasks
Define logging patterns as ordered fields in Logging Profile elements on page 266
Add Field Resolvers in Logging Profile elements on page 271
Validate Logging Profile elements on page 272
Activate monitoring of third-party devices on page 277

Benefits of adding Field Resolvers in Logging Profile elements

Field Resolvers convert values in incoming syslog fields to different values in SMC logs.

There are two types of Fields Resolvers: multi-value field resolvers and date field resolvers.

Multi-valued field resolvers

You can use multi-valued field resolvers in the following case:

To convert one value to several log fields — In some cases, a single value can have several corresponding log fields in SMC logs. A Field Resolver can parse a single value into multiple SMC log fields. For example, SMC components set an Action, a Situation, and an Event for traffic filtering decisions. If the external component notifies a “permitted” action, the Field Resolver can set the corresponding SMC log values for all 3 log fields.



Note

You can also use multi-value field resolvers if the log data has a pre-set range of values on the external devices and in the SMC. However, the possible values are different. For example, you can map a range of alert severities in the original data to similar alert severities in SMC logs (Info/Low/High/Critical).

Date field resolvers

You can use date field resolvers in the following case:

Converting time stamps — Different external devices use different date and time formats. A Field Resolver for each different incoming format maps the times and dates correctly to the SMC log format. The date and time syntax in Field Resolvers follows the Java standard.

Add Field Resolvers in Logging Profile elements

For converting syslog values, add a field resolver for either multiple values or for date and time.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the Logging Profile, click the **Field Resolvers** branch in the left pane.
2) Drag and drop a Field Resolver to the **Header** field or to the **Patterns** pane.
3) If the Field Resolver you want is not listed, right-click a Field Resolver, then select **New** > **Field Resolver**. Define the Field Resolver for either multiple values or for date and time.

Related tasks

Define Field Resolvers for multiple values in Logging Profile elements on page 271
Define Field Resolvers for date and time in Logging Profile elements on page 272
Validate Logging Profile elements on page 272

Define Field Resolvers for multiple values in Logging Profile elements

You can set the Field Resolver to convert various fields to the correct format, or to combine more than one field into one.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the properties of the Field Resolver, select **Multi-valued** as the **Field Type**.
2) In **Fields**, click **Add**, then select the appropriate SMC log fields.
   The incoming syslog data is inserted into the log fields you select.
3) To add a row to the table, click **Add**.
4) Enter the value that is used on the third-party device in the **Value** cell.
5) Enter or select the value you want to use for each selected SMC log field.
6) Click **OK**.

Related tasks
Define logging patterns as ordered fields in Logging Profile elements on page 266
Define logging patterns as key-value pairs in Logging Profile elements on page 268
Define Field Resolvers for date and time in Logging Profile elements on page 272
Validate Logging Profile elements on page 272

Define Field Resolvers for date and time in Logging Profile elements

You can set the Field Resolver to convert date and time values.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the properties of the Field Resolver, select **Date** as the **Field Type**.
2) Click **Select** next to **Time Field**, then select the log field for which you want to define a time stamp.
3) In the **Time Format** field, enter the format that the third-party device uses for the time stamp.
   Type the format according to Java standards. For the syntax, see:
   https://docs.oracle.com/javase/1.5.0/docs/api/java/text/SimpleDateFormat.html
   Example: Typing `MMM dd HH:mm:ss` maps the log time stamp as Jan 30 13:23:12.
4) Click **OK**.

Related tasks
Define logging patterns as key-value pairs in Logging Profile elements on page 268
Define logging patterns as ordered fields in Logging Profile elements on page 266
Define Field Resolvers for multiple values in Logging Profile elements on page 271
Validate Logging Profile elements on page 272

Validate Logging Profile elements

To verify that the syslog data is converted correctly to SMC log fields, you can test a Logging Profile that you created.

Steps

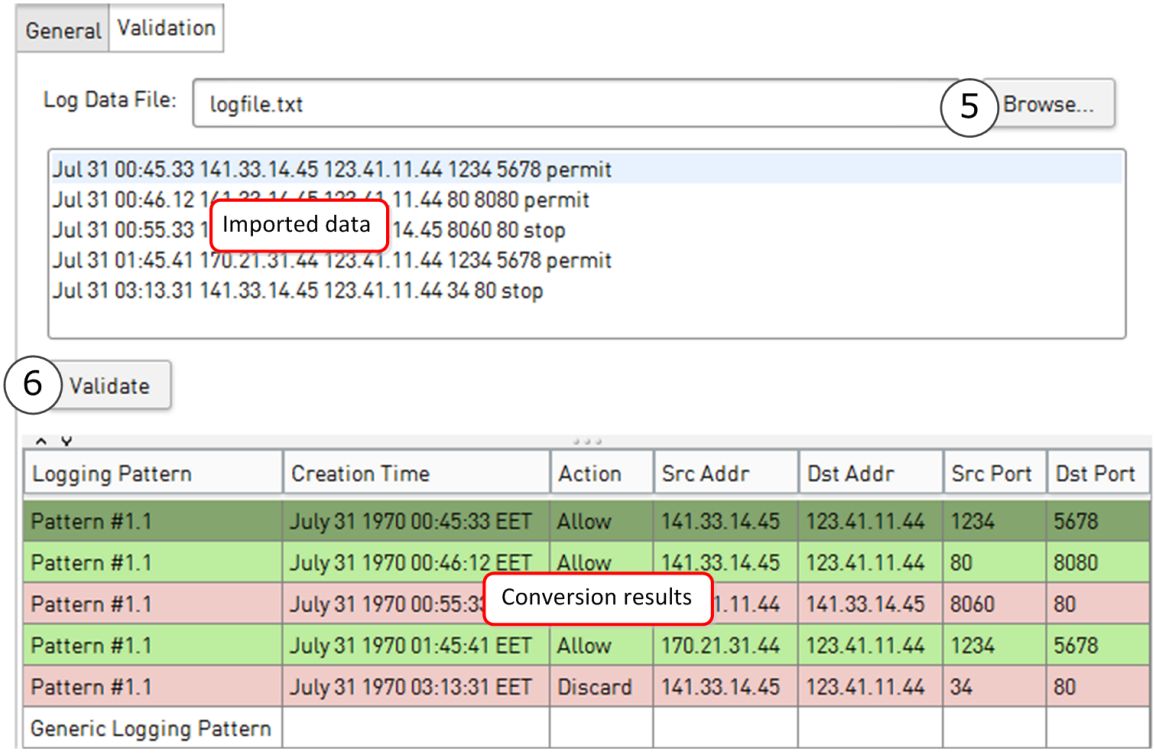
For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Monitoring Configuration**.
2) Browse to **Third-Party Devices** > **Logging Profiles**.

3) Right-click the Logging Profile that you want to validate, then select **Edit Logging Profile**.

4) Click the **Validation** tab.

5) To select a file with syslog data, click **Browse**.



6) Click **Validate**.
The imported data is displayed in the first pane. The validation results are displayed in the second pane. The first column of the results pane shows which logging pattern is used to convert each syslog entry.

Related tasks
Create Probing Profile elements for monitoring third-party devices on page 275
Activate monitoring of third-party devices on page 277

Methods for monitoring third-party device status

The Log Server can actively probe the status of third-party components using several alternative methods.

The supported methods are:

- SNMPv1
- SNMPv2c
- SNMPv3

- Pings
- TCP connection probes

When one of the SNMP status probing methods is used, you can also set up statistics reception for the device. Statistics reception relies on SNMPv1 traps sent by the third-party device.

The SMC supports statistical monitoring of the following details:

- Amount of free and used memory
- CPU load
- Interface statistics

Related concepts

Changing the ports for third-party device monitoring on page 279

Related tasks

Import MIBs for monitoring third-party devices on page 274
Create Probing Profile elements for monitoring third-party devices on page 275
Activate monitoring of third-party devices on page 277

Import MIBs for monitoring third-party devices

You can import third-party MIBs (management information bases) to support third-party SNMP monitoring.

Before you begin

You must have a MIB file for the device from the device vendor.

The OIDs (object identifiers) allow resolving the SNMP traps when they appear in log entries. If the OIDs are not resolved, they appear in the logs using a more difficult to read dotted notation. Only SNMPv1 Trap Reception is supported.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Monitoring Configuration**.
2) Browse to **Third-Party Devices** > **MIBs**.
3) Right-click **MIBs**, then select **Import MIBs**.
4) Browse for the MIB file to import, then click **Import**.
5) (Optional) When the import is finished, click **Close**.
6) In the navigation pane, browse to **MIBs** > **All MIBs** or **MIBs** > **By Vendor**.

7) To view the contents of a MIB, right-click it, then select **Properties**.

   The **General** tab shows the contents of a MIB as is. To view information about the objects that the MIB defines, click the **OID Tree** tab. The information includes the object identifiers, their OIDs in dot notation, and possibly a description of the object.

8) Click **OK**.

Related tasks

Create Probing Profile elements for monitoring third-party devices on page 275

Create Probing Profile elements for monitoring third-party devices

Probing Profiles define the monitoring of third-party device status (using SNMPv1/SNMPv2c/SNMPv3/Ping/TCP) and the settings for receiving statistics information from them (using SNMPv1).

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Monitoring Configuration**.

2) Browse to **Third-Party Devices** > **Probing Profiles**.

3) Right-click **Probing Profiles**, then select **New** > **Probing Profile**.

4) Define the probing profile settings on the **General** tab.

5) On the **Status** tab, define the probing profile status settings.

6) (SNMP/SNMPv2c/SNMPv3 probing only) On the **Statistics** tab, define the settings for statistics reception using SNMPv1 traps.

7) Click **OK**.

Related concepts

Converting logs from third-party devices on page 262

Related tasks

Activate monitoring of third-party devices on page 277

Configuring Log Servers to monitor third-party devices

Log Servers can be configured to monitor third-party devices.

You can select:

- A Probing Profile that defines how the Log Server monitors the device.
- A Logging Profile that defines how the received syslog data is converted into SMC log entries.

You can also configure Log Servers to receive SNMP traps or NetFlow (v5 or v9) or IPFIX data from third-party devices.

Define monitoring rules on a Log Server

Configure the rules that send monitoring data to a Log Server.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Dashboard** > **Servers / Devices**.
2) Browse to **Log Server**.
3) Right-click the Log Server to which you want to send monitoring data, then select **Properties**.
4) Switch to the **Monitoring** tab.
5) To create a rule, click **Add**.
6) Configure the monitoring rules.
7) To remove a rule, select the rule, then click **Remove**.
8) Click **OK**.

Create Access rules allowing third-party monitoring

If a third-party device and the Log Server are separated by a Engine or Layer 2 Engine, edit the Policy to allow traffic from the device to the Log Server.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Expand **Policies**, then browse to the type of policy you want to edit.
3) Right-click the Engine or Layer 2 Engine policy, then select **Edit Engine Policy** or **Edit Layer 2 Engine Policy**.
4) Switch to the **IPv4 Access** or **IPv6 Access** tab, then add an Access rule with the following values:
   - **Source** — the third-party element.
   - **Destination** — your Log Server.
   - **Service** — ICMP Ping, SNMP (UDP), or SNMP (TCP).

     The Service depends on the probing method that is used in the Probing Profile selected in the Monitoring rule.
   - **Action** — Allow.
   - **Logging** — None.
5) Save and install the policy to start using the new configuration.

Activate monitoring of third-party devices

You can activate status monitoring, log reception, or both for a third-party device.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Dashboard** > **Servers / Devices**.
2) Browse to **Third Parties**.
3) Right-click a third-party device, then select **Properties**.
4) On the **Monitoring** tab, select the Log Server to which the logs from the third-party device are sent.

5) (Optional) To receive status information, select **Status Monitoring**, then select the Probing Profile.

6) (Optional) To receive logging information, select **Log Reception**, then select the Logging Profile.

7) (Optional) To receive SNMP traps from the third-party device, select **SNMP Trap Reception**.

8) (Optional) To receive NetFlow/IPFIX data from the third-party device, select **NetFlow Reception**.

9) Click **OK**.

Related concepts

Getting started with monitoring the system on page 211
Converting logs from third-party devices on page 262
Configuring third-party devices for monitoring on page 278
Changing the ports for third-party device monitoring on page 279
Getting started with the Logs view on page 281
Configuring Log Servers to monitor third-party devices on page 276
Defining IP addresses as elements on page 919
Getting started with directory servers on page 1103
Getting started with user authentication on page 1127

Related tasks

Create Probing Profile elements for monitoring third-party devices on page 275
Activate or deactivate third-party device monitoring alerts on page 280
Create Access rules allowing third-party monitoring on page 277

Configuring third-party devices for monitoring

For any type of monitoring, confirm that the connections between the third-party device and the Log Server are allowed through any possible traffic filtering in your network. When configuring third-party devices for monitoring, be aware of some conditions.

- Ping and TCP status monitoring do not usually require additional configuration on the target device.
- SNMP-based polling usually requires that the target device is configured to respond to the Log Server's SNMP queries.
- Statistics sending (as SNMPv1 traps) must always be configured on the third-party device. For instructions on these tasks, see the documentation of the third-party device.
- NetFlow or IPFIX reception requires that the third-party device is configured to send NetFlow or IPFIX data. This configuration includes activating the service on the device and defining the reception port and IP address of the NetFlow or IPFIX collector (the Log Server).
- If necessary, you can change the ports that the Log Server uses to listen to syslog, SNMP, NetFlow, and IPFIX transmissions.

Related concepts
Changing the ports for third-party device monitoring on page 279
Configuring Log Servers to monitor third-party devices on page 276

Related tasks
Activate monitoring of third-party devices on page 277
Activate or deactivate third-party device monitoring alerts on page 280
Create Access rules allowing third-party monitoring on page 277

Changing the ports for third-party device monitoring

It is recommended to set your third-party devices to send data to the default ports that the Log Server listens to.

The default listening ports are:

- Windows — The Log Server listens to syslog on port 514 and SNMP traps on port 162.
- Linux — The Log Server listens to syslog on port 5514 and SNMP traps on port 5162.
- Windows and Linux — The Log Server listens to NetFlow and IPFIX data on port 2055.

If necessary, you can change the ports for syslog, SNMP, NetFlow, and IPFIX reception, but the port number in Linux must always be higher than 1024.

If it is not possible to reconfigure the third-party device to send syslog data, SNMP traps, NetFlow data, or IPFIX data to the correct port, you have other options. You can redirect traffic to a different port using an intermediate network device or on the Log Server, using iptables in Linux:

```
iptables -t nat -A PREROUTING -p udp -m udp --dport 514 -j REDIRECT --to-ports 5514
```

Related concepts
Configuring Log Servers to monitor third-party devices on page 276

Related tasks
Activate monitoring of third-party devices on page 277
Activate or deactivate third-party device monitoring alerts on page 280
Create Access rules allowing third-party monitoring on page 277
Edit Log Server configuration parameters on page 464

Activate or deactivate third-party device monitoring alerts

You can activate or deactivate the status surveillance feature, which generates an alert if a monitored component's status remains unknown for 15 minutes.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Dashboard** > **Servers / Devices Dashboard**.
2) Browse to **Third Parties**.
3) Right-click an element, then select or deselect **Options** > **Status Surveillance**.

Related concepts

Configuring Log Servers to monitor third-party devices on page 276

Related tasks

Activate monitoring of third-party devices on page 277
Create Access rules allowing third-party monitoring on page 277

Chapter 15

Viewing and exporting logged data

Contents

- Getting started with the Logs view on page 281
- Browsing log data on page 288
- Changing how log entries are displayed on page 299
- Exporting data from the Logs view on page 301
- Save data in PDF or HTML format on page 304
- Create rules from log entries on page 306
- Forwarding log data to an Elasticsearch cluster on page 307

You can view log, alert, and audit entries through the log browsing views. You can view data from SMC servers, all types of engines, and from third-party components that are configured to send data to the SMC.

Getting started with the Logs view

The **Logs** view displays all log, alert, and audit entries for the SMC.

You can view many types of entries from any number of components together or individually.

While you can view active alerts in the **Logs** view, you must acknowledge active alerts in the **Active Alerts** view.

Open the Logs view

There are several ways to access the Logs view.

Steps For more details about the product and how to configure features, click **Help** or press **F1**.

1) To access the Logs view, select an option:
   - Select **Logs**.
   - To view logs sent by a component, right-click an element that produces logs, then in the **Monitoring** submenu, select a log-related item.
   - To open the Logs view with different filtering criteria, create different bookmarks.

Related concepts

What the Logs view shows on page 282
Browsing log data on page 288
Changing how log entries are displayed on page 299
Exporting data from the Logs view on page 301

Related tasks

Create rules from log entries on page 306

What the Logs view shows

The **Logs** view can show entries generated by any SMC components and third-party components that send data to the SMC.

Depending on the permissions defined for your administrator account, the logged data can also include alert and audit entries. There are four different arrangements:

- Records
- Statistics
- Details
- Log Analysis

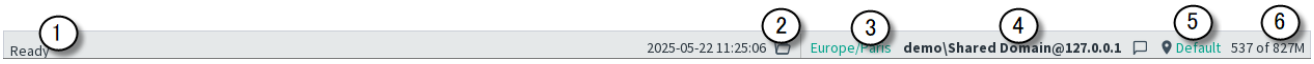
Panes in the Logs view

You can select and deselect panes from ☰ **Menu** > **View** > **Panels**.

The following panes are available in most arrangements:

- **Fields** pane — Provides quick access to categorized log entry details.
- **Query** pane — Allows you to filter the log entries so that you can find the information you need. You can filter the data by any combination of details that exist in the records.
- **Task Status** pane — Only available in the records arrangement. Shows the status of log-related tasks, such as a log export that you start from the **Logs** view.
- **Hex** pane — Shows traffic recordings generated by the Excerpt logging option of an Inspection rule (other recordings are viewed using an external viewer).
- **Summary** pane — Textual explanation of the event that generated the record.
- **Event Visualization** pane — A graphic showing important information about the event.
- **Deatils** pane — Shows detailed information about a selected log entry.

Status bar in the Logs view



1. Activity status
2. Server connectivity status
3. Set the time zone
4. The logged in user
5. The location — defines the Log Server contact address if there is NAT between the SMC Client and a Log Server
6. Real-time memory usage of the SMC Client

Note

If NAT is applied between your SMC Client and a Log Server, you must select the correct Location for your SMC Client to see the logs.

Related tasks

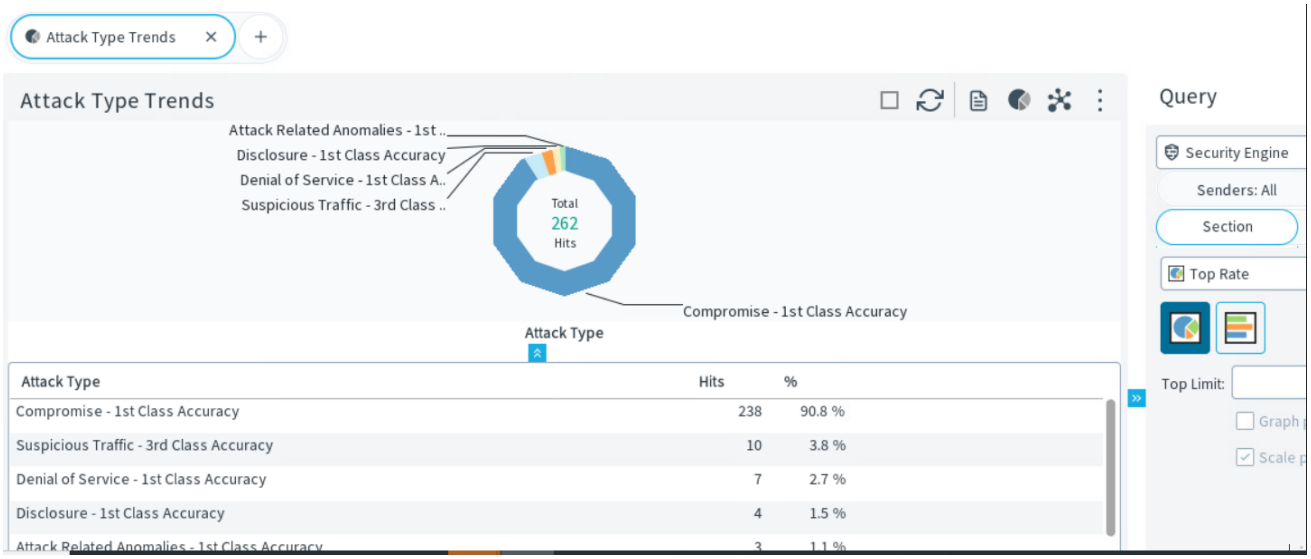
Use filters for browsing log entries on page 289
Browse log entries on a timeline on page 296
Check Whois records for IP addresses in log entries on page 297
Copy log entries in CSV format on page 301
Export log entries on page 302
Export IPS traffic recordings on page 303
Create rules from log entries on page 306

The Logs view in the Statistics arrangement

The Statistics arrangement allows you to generate basic summaries of the log data currently shown in the Logs view.

The basic summaries are similar to the charts in overviews, and include the ability to drill into the logs through individual chart items.

Statistics arrangement



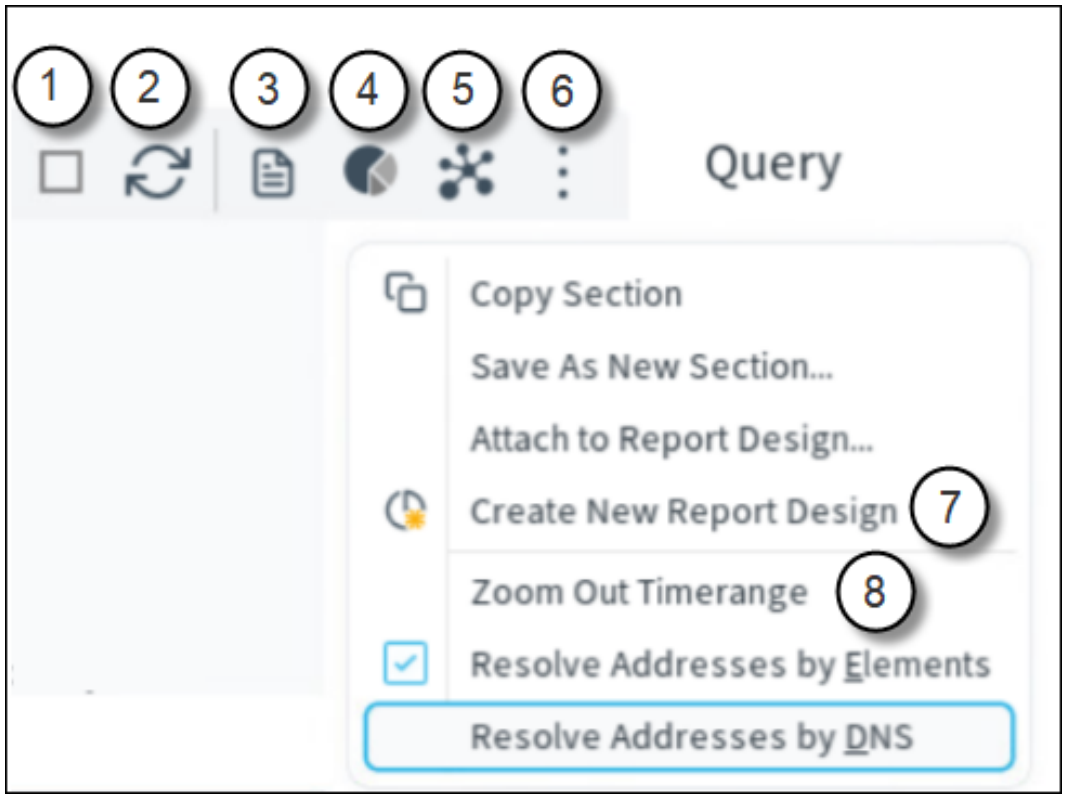
Tip

Right-click the chart sections for options that allow you to drill-down into the details.

In the Statistics arrangement of the Logs view, you can view charts of multiple events interactively. You can create a quick report of the log entries that match the active query. You can then further refine the query by viewing log entries that correspond to a chart segment.

The **Query** pane in the Statistics arrangement includes another **Section** tab. You can use the tabs to control the statistical display. The data can also be filtered in the same way as in other arrangements.

Toolbar in the Statistics arrangement



1. Stop the ongoing operation
2. Refresh
3. Show the **Records** arrangement
4. View graphical summaries based on the log entries
5. Opens the **Log Analysis** view
6. Opens the **More actions** menu
7. Options for creating new sections and reports
8. Generate a new chart with a wider time range

The chart area in the Statistics arrangement can contain a pie chart, a bar chart, a line chart, stacked line chart, or a map chart (based on an internal geolocation database). The available options depend on the chart type that is selected:

- **Top rate** charts can be displayed as a pie chart, bar chart, or a map. A top rate chart shows the total numbers of records that match the query.
- **Progress** charts can be displayed as a line chart, stacked line chart, bar chart, or stacked bar chart. A progress chart illustrates the numbers of records plotted over time (similar to the timeline, but in more detail).

When a chart is generated, you can right-click for a menu of actions related to the section and possibly the element that the section corresponds to. The actions available vary by section. Some of the most important actions are listed as follows:

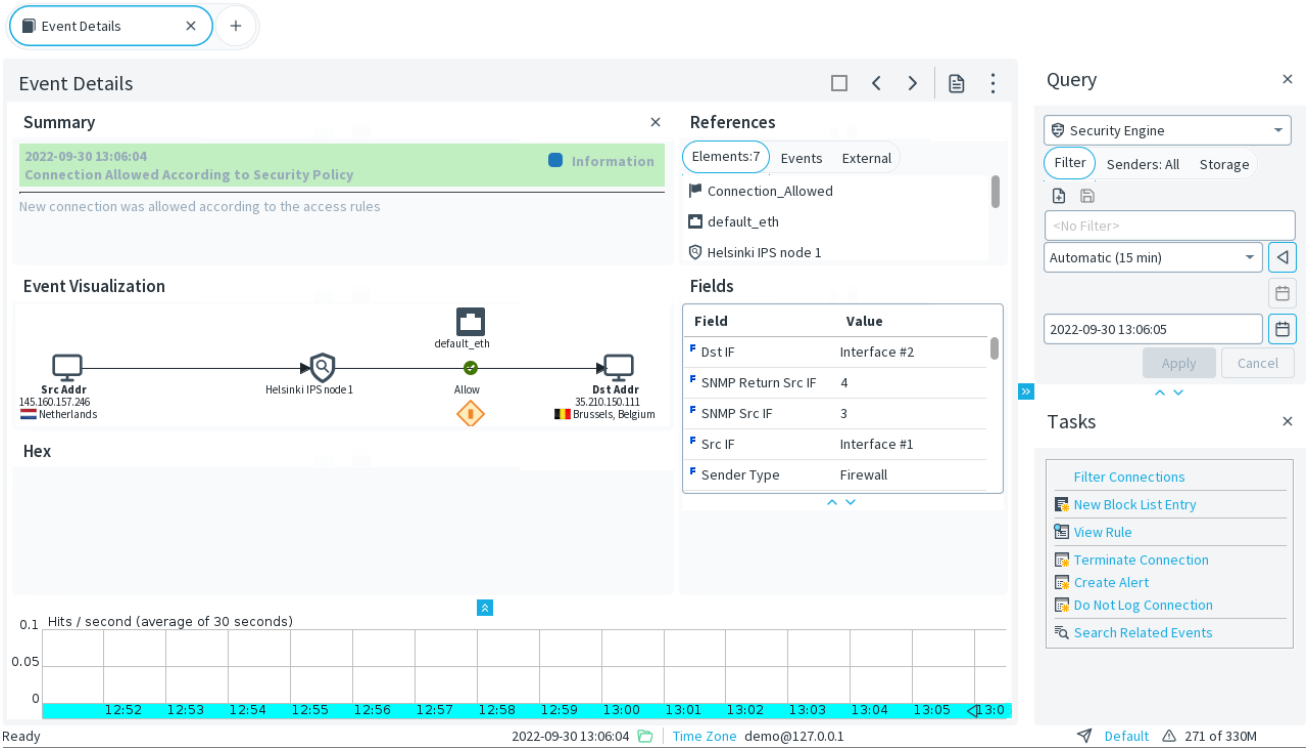
- **Show Records** — Opens the Records arrangement filtered to show the entries that comprise the chart section you right-clicked.
- **Add to Current Filter** — Allows you to use sections to filter data by adding the section in question to the **Filter** tab of the Query pane.

- **Statistics** item shortcuts — Drill down to create a chart from data that matches the previous chart section.

The Logs view in the Details arrangement

The Details arrangement shows an overview of an individual log entry.

Details arrangement



Toolbar in the Details arrangement



1. Stop the ongoing operation
2. Previous or next record
3. Show the **Records** arrangement
4. Opens the **More actions** menu

The Details arrangement also has the following panes:

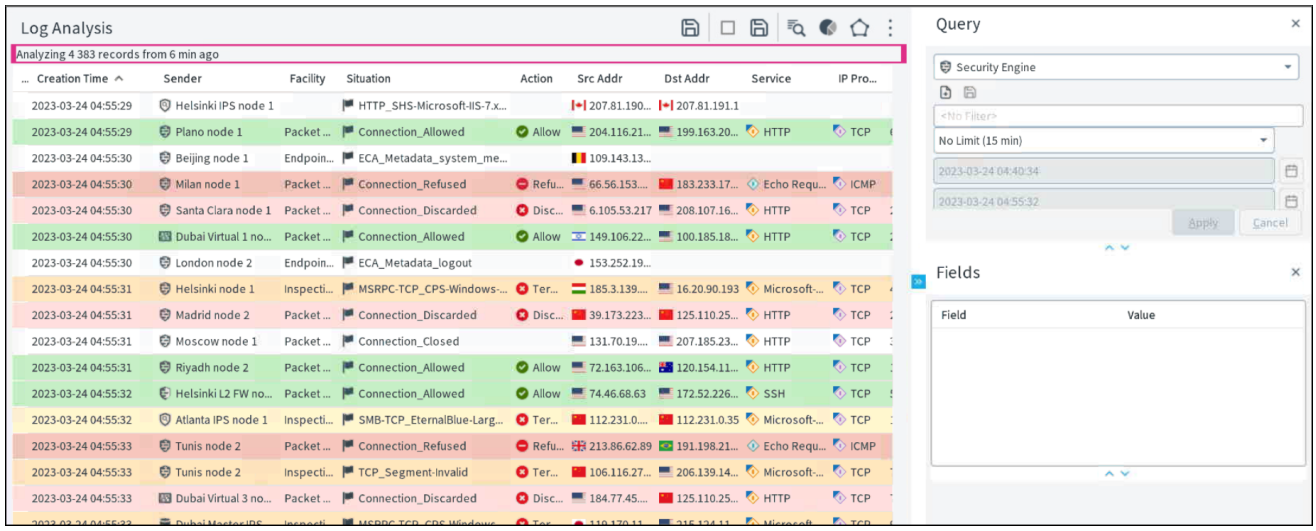
- **References** pane (shown by default) — Displays a list of elements corresponding to the details in the record and possibly more information about related records for some special records that are part of a larger event.
- **Tasks** pane — Shortcuts to configuration tasks that you can start based on the displayed entry (as in the Records arrangement in the right-click menu for entries).

The Logs view in the Log Analysis arrangement

The Log Analysis arrangement provides various tools to analyze and visualize log data.

For example, you can combine logs by service or situation, sort logs by column type, view the data as charts or diagrams. The various tools make it easier to notice patterns and anomalies in traffic.

Log Analysis arrangement



- To combine logs by Service or Situation, select **Aggregate** > **Aggregate by Service** or **Aggregate** > **Aggregate by Situation**.
- To sort logs by column type, select **Aggregate** > **Sort by Column**, then click the heading of the corresponding column.
- To view the data as charts, click **Statistics**, then select one of the predefined statistical items. Select **Select** to select an item from a complete list of statistical items.
- To view the data as a diagram, click **Visualizations**, then select one of the visualization options.

Visualization options

Option	Explanation
Attack Analysis	Displays information on Situations of the type Threat - 1st Class Accuracy or Suspicious Traffic - 1st Class Accuracy. Indicates allowed and disallowed connections between users and applications.
Audit Map	Displays information on how users manipulate elements.
Application and Executable Usage	Displays users and the applications that they use or access. Indicates allowed and disallowed connections between users and applications.
Service Map	Displays access to services in the network.

You can zoom in on the data presented in the visualization diagrams with the mouse wheel. Right-clicking elements in the diagrams opens a pop-up menu with various options to further analyze the elements and add them to filters. You can also, for example, drag and drop objects from the visualization diagram to the **Query** pane to create a filter.

Browsing log data

You can browse, filter, and search for log data in the **Logs** view.

Related concepts

Benefits of filtering log entries on page 289
Using Log Data Context elements on page 291
Viewing temporary log entries on page 296

Related tasks

View log entry details in the Fields pane on page 288
View log entries from specific components on page 290
Analyze log, alert, and audit entries on page 294
Sort log entries on page 294
Save snapshots of log, alert, and audit entries on page 295
View snapshots of log, alert, and audit entries on page 296
Browse log entries on a timeline on page 296
Check Whois records for IP addresses in log entries on page 297

View log entry details in the Fields pane

The Fields pane provides several alternative views to the log entry that is selected.

The pane is most useful in the Records arrangement. This arrangement shows a subset of fields and the information contained in the selected field. Using the Records arrangement, you can quickly browse the logs table for the exact details you are looking for without scrolling sideways or rearranging the columns.

The **Watchlist** item allows you to create a customized list of fields for your own use. The Watchlist is specific to each SMC Client installation.

You can look up a selected IP address in the online Whois database.

If an IP address in a log entry has a country flag icon next to it, a Geolocation (for example, a street, city, or country) has been associated with it. You can view the physical location of these IP addresses in Google Maps.

To change your personal Watchlist of log fields, follow these steps:

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Logs**.

2) If the Fields pane is not visible, select **Menu** > **View** > **Panels** > **Fields**.

3) In the list at the top of the pane, select **Watchlist**.

4) Change the selection to the fields you want to use:
   - To remove fields, right-click the field, then select **Remove** (to remove one field) or **Clear** (to remove all fields).

- To add more fields, drag and drop cells from the log entry table to the Fields pane. (The value of the field is irrelevant in this case.)
- To add a field to the Watchlist from other views in the pane, right-click the field, then select **Add to Watchlist**.

Related tasks

View IP address locations in the Logs view on page 246
Check Whois records for IP addresses in log entries on page 297

Benefits of filtering log entries

Efficient use of the logs requires that you filter the records displayed in the Logs view.

Filtering is done using the Query pane, which allows you to select the type of log data that it displayed. It also contains the follow tabs for filtering log data:

- **Filter** tab allows you to filter entries based on any information in the entries. The Log Data Context specifies the log data type.
- **Senders** tab allows you to filter entries according to the component that created the entry. Filtering by sender speeds up log browsing when there are many log sending components, but you are only interested in a limited set.
- **Storage** tab allows you to filter entries according to the servers on which the entries are stored.

Options on the three tabs allow you to set more filtering criteria.

Related concepts

Using Log Data Context elements on page 291

Related tasks

Use filters for browsing log entries on page 289

Use filters for browsing log entries

You can quickly create local filters by dragging and dropping. You can filter logs by time, use criteria stored in Filter elements, and save a Query as a permanent filter.



Note

The time selection refers to the entries' creation time stamp (not the reception time at the Log Server, which is also included in many entries). Internally, the SMC and engines always use universal time (UTC). The times are displayed according to the time zone selected in your SMC Client's status bar.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Logs**.

2) If the **Query** pane is not visible, select **Menu** > **View** > **Panels** > **Query**.

   You can drag and drop any field from the log entries to the **Filters** tab to create a Filter, select existing Filter elements, or add a filtering criterion. To add a criterion, use the toolbar icon and type in the detail. You can then further change and use the Filters you have created.

3) Select an option:
   - To change a detail manually, double-click the detail.
   - Right-click a field in the log entry table or in the **Fields** pane, then select **Add Filter: <field name>** to add the item and its value as a new filter row.
   - Right-click an item in the log entry table or in the **Fields** pane, then select **New Filter: <item name>** to define a value for the item and add it as a new filter row.
   - To add an empty row, right-click a filter row or empty space, then select **Row**.
   - To search based on a word or a string, right-click the **Query** pane, select **New** > **Filter: String**, then type your search string.
   - To remove a detail, right-click it, then select **Remove <detail description>**.
   - To remove a whole row, right-click something other than a detail on the row you want to remove, then select **Remove**.
   - Temporarily disable a filter row by right-clicking it, then selecting **Disable**.
   - To save the current filtering criteria as a permanent Filter element, click **Save** at the top of the **Filter** tab in the **Query** pane.

4) After you make changes to filters, click **Apply**.

Related concepts
Getting started with filtering data on page 333

Related tasks
View log entries from specific components on page 290

View log entries from specific components

You can filter logs based on the components that created the entries.

If the **Senders** tab is empty, data from all components is displayed in the **Logs** view. The **Senders** tab allows you to maintain the sender filtering independent of changes on the **Filter** tab. Restricting the included senders makes log browsing faster when there are many components in the SMC.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Logs**.

2) Switch to the **Senders** tab in the **Query** pane.

3) Click **Select** at the top of the **Senders** tab.

4) Select the elements you want to use as the senders, then click **Select**.

5) Click **Apply**.
The log data is refreshed, and only logs from the selected senders are displayed.

Related tasks
View log entries from specific servers and archive folders on page 291

View log entries from specific servers and archive folders

You can specify which servers and storage folders to include.

By default, the **Logs** view fetches data from the active log storage folder for all data storage servers, except those Log Servers that are excluded from log browsing.

You view logs from the active storage folder on specific Log Servers and Management Servers. You can also view logs from archives stored on Log Servers or archives stored locally on the computer where you are using the SMC Client. In an environment with multiple Management Servers, active alerts are automatically replicated between the Management Servers. Log Servers store all logs that other components have sent to it as well as audit data for the Log Server's own operation.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Logs**.

2) In the **Query** pane, switch to the **Storage** tab.

3) In the server list, select the servers and storage folders that you want to include.

4) Click **Apply**.
The log data is refreshed and filtered according to the selected servers and folders.

Related tasks
Edit Log Server configuration parameters on page 464

Using Log Data Context elements

You can use Log Data Contexts to select which type of log data is displayed in the Logs view and in the Reports view.

You can select a predefined Log Data Context or create a Log Data Context. You can also define the selection of columns for each Log Data Context.

Tip

To view log entries for the SMC Appliance, select the **SMC Appliance** log data context in the **Query** pane.

Related tasks

Create Log Data Context elements on page 292
Select log entry columns for Log Data Context elements on page 292

Create Log Data Context elements

The Log Data Context allows you to select which type of log data to display.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Logs**.
2) Click the Log Data Context drop-down list in the **Query** pane or in the **Log Type** section in the **Report Properties** pane, Report Section properties, or Report Item properties. Then select **New**.
3) Enter a **Name** for the new Log Data Context.
4) To add Filters to the Log Data Context, click **Select**.

   Note

   Log Data Tags index Log Data. We recommend that you select a Log Data Tag as a Filter.

5) Click **OK**.

Related tasks

Use filters for browsing log entries on page 289
Select log entry columns for Log Data Context elements on page 292

Select log entry columns for Log Data Context elements

You can edit the selection of columns that are displayed in a Log Data Context.

You can also save user-specific settings, save the updated column selection as the default settings, or reset the columns to the default settings for each Log Data Context.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Logs**.
2) Verify that the Log Data Context is selected in the **Query** pane or Log Type section in the **Report Properties** pane, Report Section properties, or Report Item properties.
3) Select ⋮ **More actions** > **Columns** > **Column Selection**.
4) Add the columns that you want to be displayed, then click **OK**.
5) Save the current column selection:
   - To save the column selection as your personal settings for the Log Data Context, select ⋮ **More actions** > **Columns** > **Save Your Local Settings**.
   - (Custom Log Data Contexts only) To save the column selection as the default settings for all administrators, select ⋮ **More actions** > **Columns** > **Save Default Settings**.
   - To discard changes to the column selection and revert to the previously saved default settings, select ⋮ **More actions** > **Columns** > **Reset to Default Settings**.

Related tasks

Select columns in the log entry table on page 300

Edit Log Data Context elements

You can change the name and edit the filters for a Log Data Context that you have created.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Logs**.
2) From the Log Data Context drop-down list, select **Select**.
3) Right-click the **Log Data Context** that you want to change, then select **Properties**.
4) Edit the properties, then click **OK**.

Related tasks

Use filters for browsing log entries on page 289

Add statistical items to a section of the Statistics arrangement of the Logs view

You can add statistical items to a section of the Statistics view.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the Query pane, click **Items**.
2) Click **Add**.
3) Select one or more items from the list, then click **Select**.
4) Click **OK**.
5) In the Query pane, click **Apply** to update the view.

Analyze log, alert, and audit entries

The Log Analysis view provides various tools to analyze logs, alerts, and audit entries.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Logs**.
2) Select **Analyze**.

Related concepts

The Logs view in the Log Analysis arrangement on page 287

Related tasks

Save snapshots of log, alert, and audit entries on page 295
View snapshots of log, alert, and audit entries on page 296

Sort log entries

By default, log entries are sorted according to their creation time. You can alternatively sort log entries according to any other column heading.

Large numbers of logs can require significant resources to be sorted. The **Log Analysis** view can shorten your selected time range if your current Query matches too many records to be efficiently sorted.

Note

The **Current Events** view is always sorted according to entry creation time. Sorting can only use stored data, so any temporary data visible in the view is permanently lost if you change the sorting.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Logs**.
2) Select **Analyze**.
3) Click the column heading by which you want to sort the logs.
   Depending on the column you click, the sort can take a while.

Save snapshots of log, alert, and audit entries

You can save snapshots of log, alert, and audit entries in the **Log Analysis** view.

The snapshots are saved on the Management Server, and are listed in the **Monitoring** view.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Logs**.
2) Select the entries for the snapshot.
   You can select a maximum of 100000 entries.
3) Select **Analyze**.
4) Click **Save**.
5) Enter a name for the snapshot, then click **OK**.

Related concepts

Browsing log data on page 288

Related tasks

View snapshots of log, alert, and audit entries on page 296

View snapshots of log, alert, and audit entries

The snapshots of log, alert, and audit entries are listed in the Monitoring view.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Monitoring Configuration**.
2) Browse to **Other Elements** > **Monitoring Snapshots** > **Logs** > **Management Server**.
3) Right-click a snapshot, then select **Open**.

Browse log entries on a timeline

You can skip around logs from different time periods using the timeline.

In the Records and Details view arrangements, part of the timeline is hidden by default. You can view the full timeline by dragging its upper edge.

Depending on your selection, the timeline allows you to browse freely (the **Automatic** option) or stops when the first or last entry within the specified time range is reached.

When you are browsing within a set time range, you cannot accidentally browse out of the time range set in the **Query** pane. Square brackets are shown at each end to show the limits of the range.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Drag the arrowhead to browse.
   - The arrow also indicates the selected start position (from the beginning or the end of the time range).
   - The chart plots the number of matching entries over time.

Viewing temporary log entries

The **Logs** view has two operating modes. One mode shows a fixed time frame, the other is a stream of current log entries, which also includes temporary entries.

In the normal mode, you can browse entries freely from any time period. When you activate the Current Events mode by clicking ▶ **Play**, the log entries update automatically to show the stream of log entries as they arrive at the Log Servers. Typically, you must filter out some entries to keep the pace of the Current Events mode slow enough that you can keep up with the entries.

The Current Events mode also displays temporary entries that are not stored on the Log Server (Transient log entries and log entries that are pruned out before permanent storing) so you might see more logs than in the normal mode. Temporary entries only exist within the current view and are permanently lost when the view is refreshed or closed. The updates in the Current Events mode are automatically deactivated when you select an entry or start browsing manually.

Note

Under some operating conditions, a small portion of log entries can arrive in mixed order. Because the Current Events mode attempts to maintain a logical flow of events, out-of-sequence entries are not shown. You might see a notification message if this happens.

Check Whois records for IP addresses in log entries

To get more information about the source of traffic that triggered a log entry, you can look up the Whois record of IP addresses in log entries.

The Whois record contains registration information and related contact details provided at the time of domain registration. The contents of the Whois record vary depending on the information provided by the owner of the domain or network segment. For IP addresses used by customers of an ISP, the information shown in the Whois record is usually the ISP's information.

The Whois information is queried from the relevant Regional Internet Registry (RIR). These registries include the ARIN (American Registry for Internet Numbers), the RIPE NCC (Réseaux IP Européens Network Coordination Centre), and the APNIC (Asia Pacific Network Information Centre). More information about the main RIRs can be found at the following links:

- ARIN at a glance: https://www.arin.net/about_us/overview.html
- RIPE Database: https://www.ripe.net/manage-ips-and-asns/db
- About APNIC: https://www.apnic.net/about-APNIC/organization

The computer running the SMC Client performs the Whois query. To be able to perform Whois queries, the security policy applied on the computer running the SMC Client must meet the following criteria:

- DNS queries must be allowed so that the SMC Client can resolve the relevant RIR server IP address.
- Opening TCP43 (Whois) connections must be allowed.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Logs**.

2) Right-click an IP address, then select **Whois**.

Tip

You can also view the physical location of an IP address in Google Maps.

The Whois record for the IP address is displayed.

Related concepts

Monitoring connections using Geolocation elements on page 244
Geolocations and IP addresses in Google Maps on page 246

Related tasks
View IP address locations from the Whois Information dialog box on page 247

Browse log and alert entries on the command line of Security Engines

If you have saved copies of the most recent log and alert entries locally on the Security Engine, you can browse the log and alert entries on the command line of the Security Engine.

Before you begin

Enable the storage of log entries on the Security Engine on the **Advanced Settings** > **Log Handling** branch of the Engine Editor.

Browsing log and alert entries locally on the Security Engine allows you to quickly troubleshoot problems that are specific to the location where the Security Engine is installed. You can browse log and alert entries even if the log and alert entries have already been sent to the Log Server, or if the connection to the SMC is not available.



Note

The root user and any other users who are allowed to access the Security Engine command line can view the saved log and alert entries.

The log and alert files are stored in the /spool/log/archive directory on the Security Engine.

You can use the following filtering when you browse log and alert entries on the command line of the Security Engine:

- Time range
- Facility
- IP address
- User name

Browsing log and alert entries on the command line of Security Engines has the following limitations:

- A limited number of log and alert entries are stored on the Security Engine for a limited time.
- In an environment with Master Engines and Virtual Engines, you can only browse log and alert entries, including log and alert entries for Virtual Engines, locally on the command line of Master Security Engines. You cannot browse log and alert entries locally on the command line of individual Virtual Security Engines.

Steps

1) Connect to the command line of the Security Engine.

2) To view log and alert entries, enter commands in the following format:

```
sg-log-view [options]
```

For details about the options, see the information about Forcepoint Network Security Platform commands. To show usage information on the command line of the Security Engine, enter the following command:

```
sg-log-view -h
```

Related tasks
Configure log handling settings on page 685
Access the Security Engine command line on page 364

Related reference
Facility field values

Changing how log entries are displayed

There are various ways in which you can customize how entries in the Log view are displayed.

Related concepts
Selecting the time zone for log browsing on page 300
Exporting data from the Logs view on page 301

Related tasks
Increase or decrease text size in log entries on page 299
Select columns in the log entry table on page 300

Increase or decrease text size in log entries

You can increase, decrease, and reset the text size in the **Logs** view.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Logs**.

2) Select ⋮ **More actions** > **Text Size**, then select an option.

Selecting the time zone for log browsing

The SMC and engines use universal time (UTC) internally. The times in the **Logs** view are changed to your selected time zone as they are displayed.

By default, this time zone is the local time zone of the computer you are using. Changing the time zone does not affect the actual time stamps of the log entries.

If the times in the log entries seem incorrect, verify that the time and time zone are set correctly in your operating system, on the Management Server, and on all Log Servers.

Select columns in the log entry table

You can select which columns are shown in the Logs view and customize how the columns are shown. You can add and remove columns and change the order and width of columns.

You can save the column selection and their settings for each Log Data Context. You can also view subsets of column information in the Fields pane.

You can arrange the columns in the following ways:

- To change the order of the columns — Drag the column header to a different location.
- To expand the column to the width of its contents — Double-click the column header.
- To view a menu of actions for adjusting the column widths — Right-click a column header.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Logs**.
2) Select ⋮ **More actions** > **Columns** > **Column Selection**.
3) Select **Add** and **Remove** to include and exclude selected fields.
   The Columns to Display list on the right shows your selections.
4) To organize selected fields on the Columns to Display list, select **Up** or **Down**.
   Fields at the top are shown at the left of the log record table.
5) Click **OK**.

Related concepts

Using Log Data Context elements on page 291

Log entry fields

Related tasks

View log entry details in the Fields pane on page 288

Tools for customizing the Logs view

To make logs easier to read, you can customize how the data is displayed in the Logs view.

The following options are available in the **More actions** menu:

- **Show Milliseconds** — Shows milliseconds in the log creation time.
- **Use Color Filters** — Enables log entry highlighting. Different colors highlight different types of logs.
- **Show Icons** — Shows the icons of the elements.

You can resolve IP addresses, protocols, and senders as DNS names or SMC elements. Resolving affects the view only and does not affect stored log data.

The following options for resolving are available in the **More actions** menu:

- **Resolve Addresses by DNS** — Enables IP address resolution using DNS.
- **Resolve Addresses by Elements** — Enables IP address resolution using element definitions.
- **Resolve Senders** — Enables the IP addresses of engines and SMC servers to be resolved using element definitions.



Note

IP address and port resolving works by comparing the information in the logs to internal and external information sources. If the information available is incorrect or ambiguous, the result might not reflect the actual hosts and services involved in the communications. For example, if a detail matches two elements, the first match is displayed even if the other element was used in the corresponding rule.

Exporting data from the Logs view

You can export log entries in various ways and formats.

Related tasks

Save elements, log data, reports, and statistics on page 304
Copy log entries in CSV format on page 301
Export log entries on page 302
Export IPS traffic recordings on page 303

Copy log entries in CSV format

Log, alert, and audit data can be copied directly from the **Logs** view, then pasted in comma-separated values (CSV) format.

For a limited number of entries, a simple copy and paste is the quickest export method.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Logs**.

2) Highlight the rows you want to copy, then copy and paste them to the other application, for example a spreadsheet application.

   The entries are copied with the column titles.

Export log entries

Log, alert, and audit data can be exported directly from the **Logs** view. Use the export command for large numbers of entries.

To schedule export tasks that are executed automatically, use the Log Data Tasks tool to export logs instead.

To export the data in a human-readable format, we recommend saving the entries in a .pdf or .html file instead. You can use this option when the exported data does not need further processing.

If you have defined an export banner, the text of the banner is added at the beginning of each exported HTML file to indicate that the export contains sensitive or classified data.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Logs**.

2) (Optional) To export only some of the entries that match your current Query, select some rows in the Records arrangement.

3) Right-click one of the entries, then select **Export** > **Export Log Events**.

4) Configure the settings, then click **OK**.

Result

The **Task Status** pane opens and shows the progress of the export.

Related concepts

Log data management and how it works on page 1307

Related tasks

Save elements, log data, reports, and statistics on page 304
Export IPS traffic recordings on page 303
Edit Log Server configuration parameters on page 464

Export IPS traffic recordings

You can set IPS Inspection rules to record network traffic as a logging option in both the Exceptions and the Rules tree.

These recordings are stored on the Log Servers. Recordings generated by the **Excerpt** option are shown directly in the **Logs** view. Longer recordings, however, are meant to be viewed in an external application and are not directly viewable.

Tip

To display the **Hex** pane, select **Menu** > **View** > **Panels** > **Hex**.

To view the recording, you can:

- Retrieve the recording through the log entry.
- Define a Task for exporting IPS recordings.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Logs**.

2) Highlight the rows that are associated with recordings.

Note

To browse for more entries that have a recording, change the selection in the **Fields** pane to **Full Capture**. (This selection is available when an entry that has an associated recording is selected.) The **Record ID** field is displayed with an identification number for entries that are associated with a recording.

3) Right-click a selected entry, then select **Export** > **Export IPS Recordings**.

4) From the **File Export Format** drop-down list, select the file format.

5) Select where to export the file.

6) Specify what happens when a previous file with the same name exists in the same folder.

7) Click **OK**.
The **Task Status** pane opens and shows the progress of the export.

Related concepts

Inspection Policy elements and how they work on page 867

Related tasks

Export log entries on page 302
Create an Export Log Task on page 1316

Save data in PDF or HTML format

You can save lists of elements, logged data, reports, statistics, and diagrams in PDF format or as HTML. You can customize the format of the PDF files.

Save elements, log data, reports, and statistics

You can save lists of elements, logged data, reports, statistics, and diagrams in PDF format or as HTML.

If you have defined an export banner, the text of the banner is added at the beginning of each exported HTML file to indicate that the export contains sensitive or classified data.

Note

Export banners are not added to log data that is exported or forwarded.

For exports in PDF, you can modify the style template to indicate that the export contains sensitive or classified data.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In most elements and views, select **Menu** > **File** > **Print**.

Use style templates in PDFs

You can use the background style templates when saving as PDF. You can use the templates in the SMC Client and the Web Portal as permitted by account permissions and Domain boundaries.

The style template is a standard PDF file you create with some or all of the following elements:

- One or more cover pages that are attached to the printed PDF before the content pages.
- A content page background with a header and footer. The height of the header and footer can be adjusted. The same single page is used as the background for all content pages.
- One or more trailing pages that are attached to the printed PDF after the content pages.
- Your PDF template file can contain pages that you do not want to use. These pages are ignored.
- A one-page PDF file is used as a content page. Your PDF template file must contain at least two pages if you want to add cover and trailing pages.

You can create the template, for example, by creating a suitable document in a word processor and saving it as a PDF. Design separate templates for the different page sizes (A3, A4, A5, or Letter) and orientations (portrait or landscape) you anticipate using.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Create a PDF file that contains a template page for the content and optionally one or more cover and trailing pages.

2) Open the **Print to** or **Print Elements to** dialog box, for example, by right-clicking a log entry and selecting **Print**.

3) Under **Page Setup**, select **New** from the **Style Template** list.

4) Enter a unique **Name** for the new Style Template.

5) Click **Browse** and select the PDF file you want to use as a template.

6) Select how the pages are used:
   - (Optional) The **Front Pages from** are inserted before the content pages without modifications. Fill in just the first field for a single-sheet front page.
   - The **Content Page** is used as the background for all pages that have system-generated content.
   - The **Header Height** and **Footer Height** define how much space (in millimeters) is left between the top and bottom of the page, and the first or last line of content. This setting prevents the generated content from overlapping text or graphics on your content page.
   - (Optional) The **Back Pages from** are inserted after the content pages without modifications. Fill in just the first field for a single-sheet trailing page.



Tip

You can use the same pages for different roles. For example, you can select the same page as a content page and a back page to add an empty page at the end of the PDF. The PDF template file must have at least two pages, even if you only use one of the pages.

7) Click **OK**.

Manage PDF style templates

You can change PDF style template settings and delete templates you no longer need.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Open the **Print to** dialog box by, for example, right-clicking a log entry, then selecting **Print**.

2) Under **Page Setup**, select **Select** from the **Style Template** list.

3) Right-click a Style Template, then select an action from the right-click menu. You can **Select** the selected Style Template or select **Properties** to adjust the template settings. You can also select **Copy** to copy the template name or **New Style Template** to create a new style template.

Create rules from log entries

You can use log entry details to generate new rules.

To convert a log entry to a rule, the log entry must be created based on a rule (the entry contains a rule tag). Creating a rule this way allows you to make quick exceptions to the current policy. You can create the following types of rules:

- A rule that allows a connection from an entry that logs stopped traffic
- A rule that stops a connection from an entry that logs allowed traffic
- A rule that changes the log level or stops the logging of matching connections

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Logs**.

2) Highlight the rows you want to include in the operation. You can select multiple log entries to create several rules in the same operation.

Note

Do not include incompatible entries in the selection:

- If you select multiple log entries, the **Sender** of all entries must be the same component.
- All selected entries must have a value in the **Rule Tag** field. (Entries must be created by rules in a policy.)

3) Right-click a selected log entry. Under **Create Rule**, select an option.

Note

The selection determines how the handling of matching connections is changed.

4) (Optional) Click **Select**, then change to the policy where the new rule is added. (For example, you can insert the rule in a subpolicy instead of the main policy.)

5) (Optional) Edit the **Comment**.

The comment is added to the rule's Comment cell.

6) Select the **Action**. All actions create the displayed rules at the beginning of the first insert point in the selected policy. You can also optionally install the policy with the new rule or open the policy for editing (with the new rule highlighted for you).

Note

You cannot edit the rule in this dialog box. To edit the rule, select **Add Rules and Edit the Policy**.

7) Click **OK**.

Related concepts
The different parts of the policy editing view on page 887

Related tasks
Install policies on page 813

Forwarding log data to an Elasticsearch cluster

Elasticsearch is an open-source search engine that runs on an external Elasticsearch server cluster. You can forward log data from Log Servers and Management Servers to an Elasticsearch cluster to improve the performance of browsing and searching for log entries, report generation, and other log-related features.

Important

Forwarding log data to an Elasticsearch cluster is an advanced feature that requires knowledge of how to configure Elasticsearch. You must already have an Elasticsearch cluster deployed and configured in your environment.

For more information about requirements for using Elasticsearch with the SMC, see Knowledge Base article 17583.

You can browse log entries that have been forwarded to an Elasticsearch cluster using the SMC Client in the same way as for other log entries. The Log Server automatically maps log fields to the corresponding Elasticsearch fields.

Elasticsearch indexes the following kinds of log fields:

- Log fields that can be used for filtering and browsing log entries.
- Log fields that can be used for reporting.

Elasticsearch indexes log data only when a log data file on the Log Server is complete. Typically, log data files are completed about once every hour. If a large number of log entries are received, the Log Server might create multiple log files each hour.

The configuration consists of these general steps:

1) Configure Elasticsearch in your environment.
   a) Deploy and configure an external Elasticsearch cluster.
   b) (Recommended) Configure TLS and client authentication in Elasticsearch.
2) In the SMC Client, create an Elasticsearch Cluster element.
3) (Optional) Override the settings for client authentication that are defined in the Elasticsearch Cluster element in the properties of the Log Server or the Management Server.

Create an Elasticsearch Cluster element

The Elasticsearch Cluster element defines the settings for contacting the Elasticsearch cluster. You can create one Elasticsearch Cluster element.

Before you begin

- You must already have an Elasticsearch cluster deployed and configured in your environment.
- You must create a TLS Profile element if you want to use an imported certificate to secure the connection between the Log Server or Management Server and the Elasticsearch cluster.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the SMC Client, select **Engine Configuration**, then browse to **Network Elements**.
2) Browse to **Servers**.
3) Right-click **Servers**, then select **New** > **Elasticsearch Cluster**.
4) Configure the settings, then click **OK**.

Override the settings for client authentication defined in the Elasticsearch Cluster element

In the properties of the Log Server or the Management Server, you can optionally override the settings for client authentication that are defined in the Elasticsearch Cluster element.

Before you begin

- Create an Elasticsearch Cluster element.
- You must create a TLS Credentials element if you want to use an external certificate to secure the connection between the Log Server and the Elasticsearch cluster.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Network Elements**.
2) Browse to **Servers**.
3) Right-click the Log Server or the Management Server, then select **Properties**.

4) (Optional) On the **Elasticsearch** tab, select **Override**, then select the TLS certificate that is used to secure the connection between the Log Server or the Management Server and the Elasticsearch cluster.



Note

The Elasticsearch tab is only visible after you have created an Elasticsearch Cluster element.

5) Click **OK**.

Browse log entries that have been forwarded to an Elasticsearch cluster

You can browse log entries that have been forwarded to an Elasticsearch cluster using the SMC Client in the same way as for other log entries.

Before you begin

Configure the Log Server or the Management Server to forward log data to the Elasticsearch cluster.

Some statistics items are not supported for log entries that are stored on an Elasticsearch cluster. Statistics items that are not supported are indicated with an icon in the SMC Client.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Logs**.

2) On the **Storage** tab of the **Query** pane, select **Elasticsearch**.

Result

Log entries that have been forwarded to Elasticsearch clusters are shown.

Create reports using log entries that have been forwarded to an Elasticsearch cluster

You can generate some types of reports using log entries that have been forwarded to an Elasticsearch cluster.

Before you begin

Configure the Log Server or the Management Server to forward log data to the Elasticsearch cluster.

Report designs that are supported for log entries that are stored on an Elasticsearch cluster are grouped together in the SMC Client. Some statistics items are not supported for log entries that are stored on an Elasticsearch cluster. Statistics items that are not supported are indicated with an icon in the SMC Client.

Note

If the report includes report items that are not supported, empty Report Sections might be generated in the report.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Monitoring Configuration**.
2) Browse to **Reports** > **Design**.
3) Right-click a Report Design, then select **Start**.
4) On the **Storage** tab of the **Report Operation Properties** dialog box, select **Elasticsearch**.
5) Click **OK**.

Related concepts

Getting started with reports on page 311

Chapter 16

Reports

Contents

- Getting started with reports on page 311
- Restricting a report's scope on page 313
- Designing reports on page 315
- Generate and view reports on page 322
- Exporting reports on page 328
- Status reporting on page 332
- Example of reports on page 332

Reports are summaries of logs and statistics that allow you to combine large amounts of data into an easily viewable form.

You can process data from logs and engine statistics and generate easy-to-read diagrams, charts, and tables. The reports you generate are based on a Report Design. A Report Design can be a predefined design, a predefined design that you modify, or a custom design that you create.

Getting started with reports

The SMC Client provides extensive reporting tools for generating reports on information stored in the SMC. The summaries that make up the reports can be illustrated with different types of charts and tables.

Reports allow you to gather and visualize data in an easy-to-read format that provides an overview of what is happening in the network and that you can customize. Reports are configured and generated in the Monitoring view. You can view reports as graphs, charts, tables, and geolocation maps.

You can generate reports based on two types of runtime data:

- **Log data** — Consists of distinct events (for example, a connection opening or closing). It contains all details about each event including the exact time when the event occurred. Log data can be filtered granularly, but running statistics from the raw logs can be slow, especially when using a long data period.
- **Counter data** — Consists of pre-processed summaries of statistics that are based on sums or averages of events or traffic units within a certain period. Counter data that is older than an hour is consolidated by the hour. Counter items produce statistics quickly, even for long periods of data, but they can only be filtered by sender.

You can create reports on log, alert, and audit entries and statistical monitoring information.

You can generate reports based on predefined Report Designs and Report Sections or on Report Designs that you have created yourself. You can use your own Style Template for PDF creation to give the reports a unified corporate look.

You can view the reports in the SMC Client and in the Web Portal.

You can export reports in PDF, HTML, or plain text format, so that the files can be printed and shared. You can also directly email reports as they are generated.

Various ready-made Report Designs are provided. You can customize the existing templates or design new reports to meet your needs.

In addition to creating and generating reports based on Report Designs, you can also quickly create reports in the **Logs** view, in the Log Analysis arrangement, for example.

To provide auditing information in compliance with regulatory standards, you can generate a purpose-built System Summary report that summarizes elements, administrators, policies, and other details about system configuration and events.

Related concepts

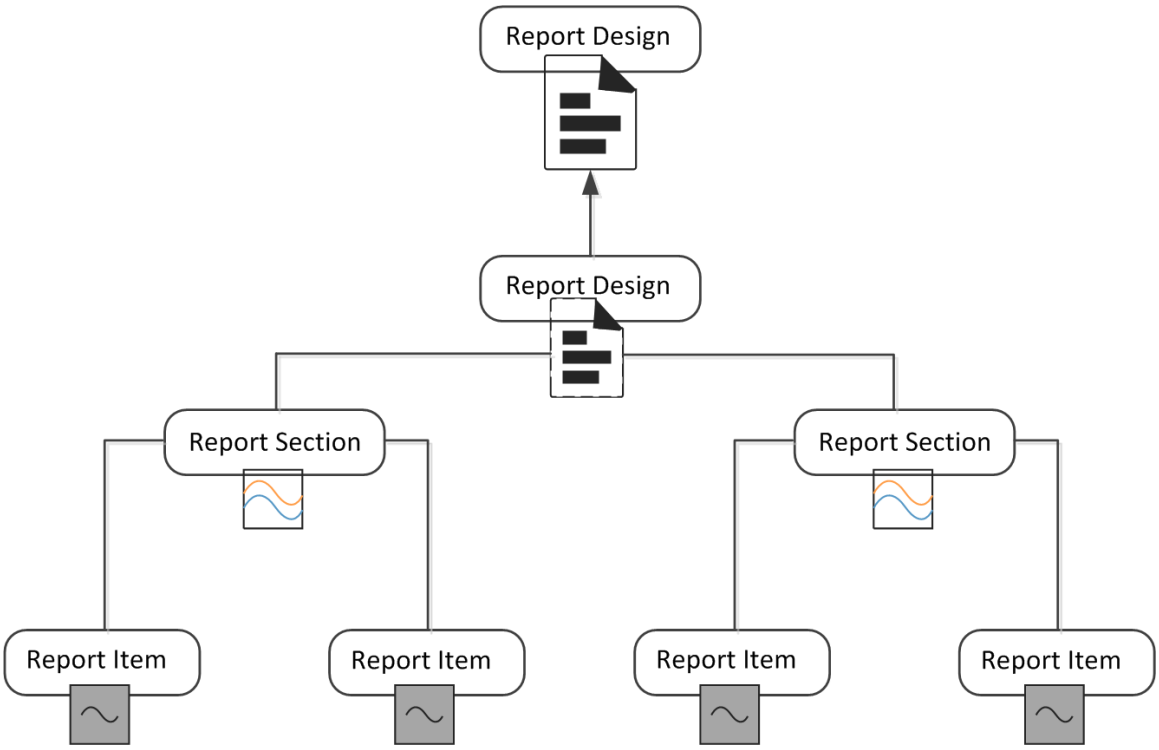
The Logs view in the Statistics arrangement on page 283

How reports are constructed

Reports are summaries of log data and statistical monitoring information. Reports consist of *Report Items*, *Report Sections*, and *Report Designs*.

The following illustration shows their relationships.

Reporting objects and elements



The Report Design is the main container for a particular type of report. The Report Designs are used as the basis for Report Tasks that generate the reports that you can view.

The Report Design consists of one or more Report Sections. A Report Section defines parameters for all items within it. It mainly defines how the information is displayed, such as the type of chart and the number of top items shown. Each Report Section in the Report Design creates a separate chart and table in the generated report.

Each Report Section contains one or more Report Items. Each Report Item represents a way to summarize the data. For example, a Report Item might summarize the total number of connections counted between the start and end times defined for the task that generates the report.

Each item adds specific information to the chart or table that the Report Section generates. For example, the items, total traffic volume, sent traffic volume, and received traffic volume, are shown as separate lines on the curve chart that the Report Section generates.

Report configuration overview

Reports consist of Report Designs that include relevant sections and items.

Follow these general steps to create and generate reports:

1) Customize an existing Report Design or create a new Report Design.

2) Customize or add Report Sections or Items.

3) Generate the report.

Related concepts

Designing reports on page 315
Defining what Report Sections to use in a Report on page 317

Related tasks

Generate and view reports on page 322

Restricting a report's scope

There are various tools you can use to focus the scope of your Reports.

Data filters for reports

Filters are the main tool for increasing the granularity of reports.

The Report Items define some general criteria for selecting data. For example, you can produce a report of connections by source IP address. Defining general criteria is a good way to get an overview, but in many cases, you want more specific information.

For example, a general item such as the total number of logged connections can be made much more specific. To make this item more specific, apply a filter that matches a single pair of source and destination IP addresses. The filter only counts connections between these two hosts.

Only log-based raw data is suitable for log filters. Counter Statistics Data items use pre-counted statistical data instead of logs, so most log data filters cannot be used with Counter Statistics Report Items. You can only filter Counter Statistics Report Items by sender (the component that generated the statistics).

You can define a filter at one or more of the following levels:

1) **Report Task** — Applies to the single report produced by a particular task or to all reports produced by a task that is scheduled to run regularly.

2) **Report Design** — Applies to all reports produced using the Report Design.

3) **Report Section** — Applies to all items included in the section.

4) **Report Item** — Applies to that specific item only.

If filters are applied at several of these levels, all filters are applied and the log entries are filtered top-down in the order listed. Each filtering stage completely excludes non-matching log entries for the stage in which the filter is applied and all further stages. For example, select the filter, TCP destination port 80, for a Report Section. Then, all items in that section only process information in log entries that mention TCP destination port 80.

Reports within administrative Domains

If there are administrative Domains configured, the reports are Domain-specific.

While logged on to a Domain, you can only create reports concerning the components that belong to that Domain. If you are allowed to log on to the Shared Domain and have unrestricted permissions (superuser), you can create Reports concerning components in any Domain.

The Reports created in the Shared Domain are visible to administrators in all other Domains. If the reports contain sensitive data that must not be displayed to all administrators in all Domains, create the reports in a specific Domain.

Related concepts
Getting started with Domain elements on page 433

Reporting on the configuration database and audit logs

While other reports are based on logs and statistics, the System Report is based on information collected from the Management Server's configuration database and audit logs.

The report includes details like administrator and Web Portal user activity, account settings, configuration of and changes to the Engine and IPS engines, and configuration of the Management Server.

The report can help you provide the required data for auditing in compliance with regulations, such as the Payment Card Industry (PCI) Data Security Standard.

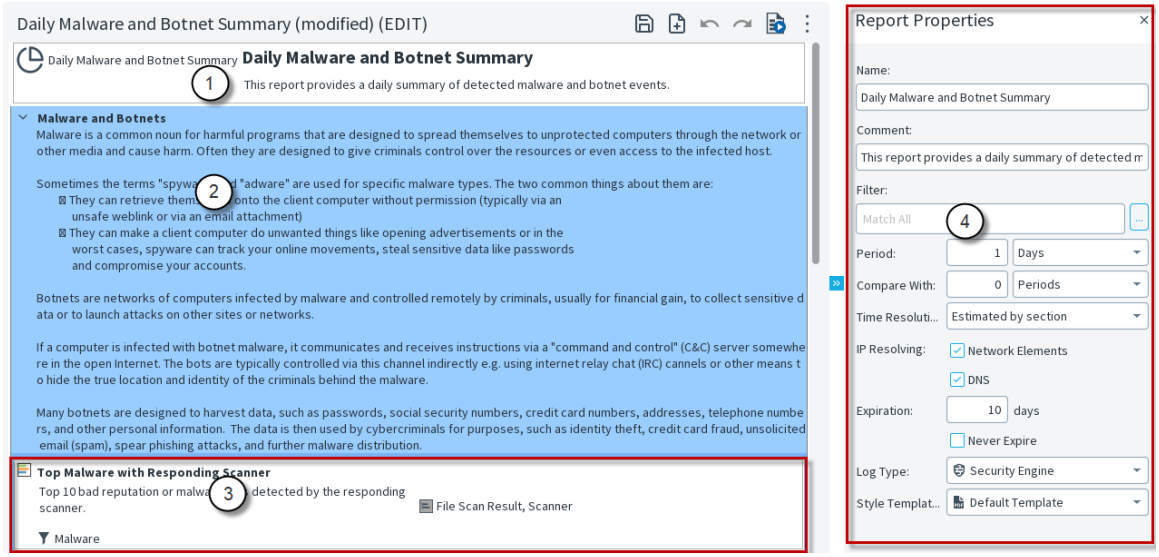
The report is generated, exported, and edited in the same way as other types of reports. The only difference is the content of the report.

Designing reports

Report Designs determine how to process the data and how the results are displayed.

They can also determine which template is used for PDF exports and which charts appear on them. Ready-made Report Designs serve as a useful guide for constructing your own Report Designs. You can also create custom Report Designs.

Report Design



1. Report overview
2. Heading Section
3. Report Sections
4. Properties of the selected element

There are several ways to create Report Designs. Although you can start by defining a new empty Report Design, it is often easier to use one of the predefined Report Designs as a template. The comments in the properties of predefined Report Designs and Sections explain their general purpose.

The Report Design properties specify, for example, the time period for the report. You can define a period comparison for the Report Design. This feature allows you to compare values between two identical time periods. For example, if the time period is one week, you can compare the results for this week to the results from the previous week.

Related tasks

Modify Report Design elements on page 316
Create Report Design elements on page 316

Modify Report Design elements

If you have an existing Report Design, you can add or remove sections or items.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Monitoring Configuration**.
2) Browse to **Reports** > **Design**.
3) Right-click the Report Design that you want to edit, then select **Edit Report Design**.
4) After making your changes, click **Save** or select **More actions** > **Save As**.

Related tasks

Create Report Design elements on page 316

Create Report Design elements

First you create a Report Design, then you add items to it.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Monitoring Configuration**.
2) Browse to **Reports** > **Design**.
3) Right-click **Design**, then select **New Report Design**.
4) Select a template for the new Report Design, then click **OK**.
5) Enter a name for the new Report Design.
6) (Optional) Select a **Filter**.



Note

When you generate a report, all filters defined in the report task properties, in the Report Design, Report Sections, and individual Report Items are used to filter the data. If the filters do not intersect, empty Report Sections might be generated in the report.

7) Adjust the other properties as needed.
8) Click **Save** or select **More actions** > **Save As**.

Related concepts
Using Log Data Context elements on page 291
Defining what Report Sections to use in a Report on page 317
Creating and editing local filters on page 337

Defining what Report Sections to use in a Report

A Report Design consists of one or more Report Sections, which define parameters for all Report Items. You can modify and create Report Sections in a Report Design.

Each Report Section in the Report Design creates a separate chart or table (or both) in the generated report.

Tip

To browse and edit predefined Report Sections, select **Engine Configuration**, then browse to **Monitoring** > **Reports** > **Sections**.

Tip

You can create customized report sections from the Statistics arrangement of the **Logs** view.

Customize Report Sections and Items

You can add predefined Report Sections to your Report Design and then modify their contents and properties according to your needs.

Tip

Use Heading Sections that contain a description to group the different Report Sections.

A Report Section represents a collection of Report Items in reports. Each Report Section adds a separate summary (chart or table) to the report. Depending on the summary type, the summary can be presented in one or more of the following ways:

- Bar chart
- Stacked bar chart
- Curve chart
- Stacked curve chart
- Pie chart
- Geolocation map
- Table

The available types of Report Section summaries are explained in the following table.

Summary types

Summary type	Description	Visualization
Progress	Illustrates how events are spread out within the reporting period. This summary type is useful for finding trends in the data. Example: A line chart showing the volume of traffic during a 24-hour period.	A bar chart, stacked bar chart, curve chart, or stacked curve chart.
Top Rate	Illustrates events with the highest occurrences. This summary type is useful for highlighting the most common values in the data. Example: A bar chart showing the number of connections to the five IP addresses that have received the most connections yesterday. The first Report Item in a top rate summary section must have a sorting criteria “by X” (for example, allowed connections by source IP address). The sorting criteria is applied to all items in the section for ranking the top rates.	A bar chart, a pie chart, or a geolocation map. A bar chart is more suitable for displaying many top rates, whereas a pie chart is better at illustrating the relative proportions. A geolocation map shows the distribution of events according to physical location.
Summary Table	A simple table for displaying the exact event counts. This summary type is useful for providing data for further processing, for example, in a spreadsheet application.	A table.
System Information	Summarizes current configuration information in the Management Server’s internal database. Example: A listing of all engines with the software versions, names of the currently installed policies, and the latest policy upload times.	A table.

A Report Item represents a value that you want counted in log data or statistical monitoring information. (Allowed traffic in bits or the number of allowed connections are examples of values that can be counted.)

The data for the Report Items is generated in the following ways:

- A simple count of how many log entries have a certain value within the reporting period. For example, the Allowed Connections Report Item counts the log entries that have the value Allowed in the Action field. A simple count is how the results for most Report Items are summed.
- A count of how many log entries have a certain value within the reporting period grouped “by X” criteria. For example, Allowed connections by source IP address presents a chart for an adjustable number of IP addresses that have the most allowed connections within the reporting period.
- Sums or averages of traffic volumes in log entries for Report Items of the “traffic” type (for example, Allowed traffic). Access rules that have the accounting option enabled in the Engine Policy generate the data for “traffic” items. Interface statistics often provide more accurate total volumes because accounting (and logging in general) is not active for all rules.
- Values stored in the Management Server’s database for System Information items. The statistical data is pre-summarized. It is not as detailed as the monitoring statistics displayed in the **Dashboard** view and cannot be filtered in detail like the log data.

Related concepts

Data filters for reports on page 313

Related tasks
Modify Report Sections on page 319
Create Report Sections on page 320

Modify Report Sections

You can add predefined Report Sections to your Report Design, then edit their contents and properties according to your needs.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Monitoring Configuration**.
2) Browse to **Reports** > **Design**.
3) Right-click the Report Design that you want to edit, then select **Edit Report Design**.
4) Select the section in the Report Design.
   The section properties are displayed in the **Section Properties** pane.
5) Edit the section properties.
6) (Optional) Group Sections under a Heading Section.
   a) Right-click a Section that you want to add to a group, then select **Move to New Heading Section**.
   b) In the **Section Properties** pane, enter a name and optionally a description for the group of Sections.
7) (Optional) To change the order of the Sections or to add more Sections under a Heading Section, drag them to the order you want.
8) Click **Save** or select **More actions** > **Save As**.

Related concepts
Using Log Data Context elements on page 291
Creating filters on page 334

Related tasks
Create Report Sections on page 320

Create Report Sections

You can create new Report Sections in a Report Design.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Monitoring Configuration**.
2) Browse to **Reports** > **Design**.
3) Right-click the Report Design that you want to edit, then select **Edit Report Design**.
4) Right-click in the Report Design, then select **Add New Section**.
5) Select a section from the list.

   Note

   If you cannot find an appropriate section in the list, select **Select**, then select a section in the **Select Section** dialog box. If you want to add a section based on a statistical item, select **Create from Item**, then select an item from the **Select Item** dialog box.

   The new section is added to the Report Design and the background of the new section is highlighted.
6) Modify the section properties.
7) Click **Save** or select **More actions** > **Save As**.

Related tasks

Modify Report Sections on page 319

Create and modify Report Items

You can add and edit statistical items in Report Sections.

Statistical items count the following types of data:

- Log entries (referred to as *records* in the item names)
- Summaries of some log fields included in those entries (such as traffic volumes in log data that contains accounting information)
- Gathered statistical data (counter items)
- System information to summarize current configuration information in the Management Server's internal database

Add Report Items

You can add Items to a Report Section.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Monitoring Configuration**.
2) Browse to **Reports** > **Design**.
3) Right-click the Report Design that you want to edit, then select **Edit Report Design**.
4) Select the Report Section to which you want to add Items.
5) In the **Section Properties** pane, click **Items**.
6) Select **Add** > **Add Item**, then select the Items to add.



Tip

You can add one nested Item under the first Item. Select **Add** > **Add Secondary Item**, then select the Item to add. The Item is nested in the generated Report.

7) Click **OK**.

Related tasks

Modify Report Items on page 321

Modify Report Items

You can modify Items that are included in a Report Section.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Monitoring Configuration**.
2) Browse to **Reports** > **Design**.
3) Right-click the Report Design that you want to edit, then select **Edit Report Design**.
4) Select the Report Section in which you want to modify the Item selection.
5) In the **Report Properties** pane, click **Items**.
6) Right-click the item that you want to modify, then select **Properties**.

7) Modify the Item properties.

8) To save the changes to the Item properties, click **OK**.

9) (Optional) To change the order of the added Items, drag and drop them.

10) Click **OK**.

11) Click **Save** or select ⋮ **More actions** > **Save As**.

Related concepts
Using Log Data Context elements on page 291

Generate and view reports

Reports are generated from the Report Designs that are under the Reports tree of the **Monitoring Configuration** view.

When you generate a report, the Management Server sends the task to all Log Servers that are not excluded from processing. The task's progress and possible errors are shown next to the task under the selected Report Design.

Each Log Server processes the task and returns the summary data for each Report Section. The Management Server merges the data from the Log Servers into one report. If one of the selected Log Servers cannot be contacted for any reason, the execution of the task is delayed until the Log Server becomes available.

Related concepts
Designing reports on page 315

Related tasks
Generate reports on page 322
View reports on page 326

Generate reports

After you have created your Report Design, you can set the time period that you want details about, then generate your report.

Note

When you generate a report, all filters defined in the report task properties, in the Report Design, Report Sections, and individual Report Items are used to filter the data. If the filters do not match any data, empty Report Sections might be generated in the report.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Monitoring Configuration**.
2) Browse to **Reports** > **Design**.
3) Right-click a Report Design, then select **Start**.
4) Select options to define how to generate the report.
5) Click **OK**.

Related tasks

Define report tasks on page 323
Cancel ongoing report tasks on page 326
View reports on page 326

Define report tasks

You can set the task to be repeated daily, and also set the output format for the task.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Monitoring Configuration**.
2) Browse to **Reports** > **Design**.
3) Right-click a Report Design, then select **Start**.
4) Click the **Task** tab.
5) Select how often you want to **Repeat** the report generation.
   The time period selected in the Report Design determines the available choices.
6) (Optional) Select a **Category**.
7) Select one or more outputs to be produced directly.



Note

To view the report before deciding if you want to process it further, leave only **Store Report** selected.

- **Text Export/PDF Export/HTML Export** — The report is stored as a text, PDF, or HTML file on the Management Server in the <installation directory>/data/reports/files/<Report Design name>/ directory. The report is named according to the time range chosen.

  Note

  If you installed the Management Server in the `C:\Program Files\Forcepoint\SMC` directory in Windows, some program data might be stored in the `C:\ProgramData\Forcepoint\SMC` directory.

- **Post Process** — The report is generated according to options chosen and then a script is started. By default, the script is SgPostProcessReport.bat, which is in the <installation directory>/data/reports/bin directory on the Management Server.

8) (Optional) Enter the **E-mail Address** to which the completed report is sent directly as email. Separate addresses with commas.

   Note

   The SMTP server for sending reports must be defined in the Management Server's properties.

9) (PDF exports only) Select a **Style Template**. If you use the default Style Template, you can select whether to create a portrait or landscape PDF. If you use a customized template, the orientation is defined in the template.

Related concepts

Tab-delimited text report files on page 330

Related tasks

Use style templates in PDFs on page 304
Select data sources for reports on page 325
View reports on page 326
Modify Management Server elements on page 481

Post-processing report files

You can customize reports by post processing them as part of the Report Task.

Post-processing runs the <installation directory>/data/reports/bin/sgPostProcessReport script on the Management Server and passes command arguments to the script. The following table explains the possible command arguments.

Command arguments for post-processing reports

Command argument	Explanation
`-creation_time YYYY/MM/DD hh:mm:ss`	The report creation time.
`-filter_categ name`	The category assigned to the task filter.
`-filter_name filter`	The name of the filter assigned to the task.
`-html_file filename`	The file name of a report exported as HTML.

Command argument	Explanation
`-pdf_file filename`	The file name of a report exported as PDF.
`-period_begin YYYY/MM/DD hh:mm:ss`	The begin time of the reporting period.
`-period_end YYYY/MM/DD hh:mm:ss`	The end time of the reporting period.
`-report_categ name`	The category assigned to the report (and to the report file).
`-report_file_title title`	The title of the report file.
`-report_name name`	The name of the report.
`-text_file filename`	The file name of a report exported as plain text.

The script parses the values from the command arguments to use the values for post-processing. Only parameters that have a defined value are forwarded to the post-processing script.

When a parameter has multiple values, each of the values is forwarded as a separate command argument.

Example of using multiple values

When the report has the two categories “Example Corporation” and “weekly report”, these values are forwarded to the script as `-report_categ Example Corporation` and `-report_categ weekly report`.

Select data sources for reports

By default, the Management and Log Servers are selected as data sources, but you can select the source of your choosing.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Monitoring Configuration**.
2) Browse to **Reports** > **Design**.
3) Right-click a Report Design, then select **Start**.
4) Click the **Storage** tab.
5) Select the data storage type:
   - **Default** — The Management Servers and Log Servers are used as the data sources.
   - **Primary archive** — Archived data is used as the data source.
   - **Custom** — A combination of archived data and data provided by the Management and Log Servers is used as the data source.
6) Select the Management Servers and Log Servers from which you want to include data in the report.
7) To start generating the report, click **OK**.

Related tasks
View reports on page 326

Cancel ongoing report tasks

If the report you are generating includes large amounts of data, generating the report can take a long time.

This situation might result when, for example, the time frame is wide and the data filter in use does not restrict the data sources.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Monitoring Configuration**.
2) Browse to **Reports** > **History**.
3) Right-click the task you want to cancel, then select **Abort** or **Delete**.
   - Aborted tasks can be edited and restarted by double-clicking them.
   - Deleted tasks are permanently removed.

View reports

After you generate a report, it is available for viewing in the Reports view.

A report might be automatically deleted according to the expiration setting defined in the Report Design. If you want to keep a generated report permanently, we recommend exporting it.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Monitoring Configuration**.
2) Browse to **Reports** > **History**.
   The created reports are automatically grouped according to their creation date.
3) To group reports by date or design, select ⋮ **More actions** > **Organize By** > **By Date** or ⋮ **More actions** > **Organize By** > **By Design**.

4) Double-click the report you want to view.



Tip

If you want to see the data of a report section in table format, right-click the section, then select **Show Table**. The table is added to the section.

The contents of the report are shown.

Related concepts
Exporting reports on page 328

Change the properties of generated reports

You can change the Category and expiration time of a report.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Monitoring Configuration**.
2) Browse to **Reports** > **History**.
3) Right-click the title of the report you want to change, then select **Properties**.
4) (Optional) Select a **Category**.
5) (Optional) To change the time when the report is automatically deleted, change the value in the **Expiration** field. To delete the report yourself, select **Never**.
6) Click **OK**.

Related concepts
How Categories help you view only certain elements on page 201
Exporting reports on page 328

Related tasks
Generate and view reports on page 322

Exporting reports

You can export PDF, HTML, and text reports manually for previously generated reports or automatically when generating the report.

Automatically exported files can be automatically sent out as email or saved in the <installation directory>/data/reports/files/report_design/ directory on the Management Server.

The report files are named according to the report's time range as follows: startdate_starttime_enddate_endtime_N, where N is a sequential number (starting from 1) that identifies files from the same time range.

Example report file name

20150423_100000_20150424_180000_1.txt is the first text report generated for the time range from 23 April 2015 10:00:00 to 24 April 2015 18:00:00.

Related tasks

Generate and view reports on page 322

Export reports as PDF files

When you export a report file as a PDF file, a default template is automatically used for the report.

You can also import a Style Template for the report you are about to export.

Note

After exporting, check the result. If the report text or charts are placed on top of your template background, you might need to adjust the headers and footers in your template.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Monitoring Configuration**.

2) Expand the **Reports** branch, then select **History**.

3) Double-click the report you want to export.
   The report opens.

4) Select what part of the report to export:
   - To export the whole report — Select **Menu** > **File** > **Print**.
   - To export a section of the report — Right-click the section, then select **Print section**.

5) Select **PDF** as the **Format**.

6) Select a printing option:
   - To open the PDF in your default PDF reader — Select **Print to PDF reader**.

- To save the PDF — Select **Print to File**, then browse to the location where you want to save the file.

7) (Optional) Select the **Style Template**.
   If you use the default Style Template, you can select whether to create a portrait or landscape PDF. If you use a customized template, the orientation is defined in the template.

8) Click **OK**.
   The PDF is generated.

Related tasks
Use style templates in PDFs on page 304
Generate and view reports on page 322

Export reports as HTML files

When you export a report as an HTML file, a default template is automatically used as the background for the report.

If you have defined an export banner, the text of the banner is added at the beginning of the HTML file to indicate that the export contains sensitive or classified data.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Monitoring Configuration**.

2) Expand the **Reports** branch, then select **History**.

3) Double-click the report you want to export.
   The report opens.

4) Select what part of the report to export:
   - To export the whole report — Select **Menu** > **File** > **Print**.
   - To export a section of the report — Right-click the section, then select **Print section**.

5) Select **HTML** as the **Format**.

6) Click **Browse**, then select where you want to save the HTML files.

7) Click **OK**.
   The HTML files are saved in the defined location. The HTML report opens in your default web browser.

Related tasks
Generate and view reports on page 322

Export reports as tab-delimited text files

You can export a plain text file, without any formatting.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Monitoring Configuration**.
2) Expand the **Reports** branch, then select **History**.
3) Double-click the report you want to export.
4) Select what part of the report to export:
   - To export the whole report — Select **Menu** > **File** > **Print**.
   - To export a section of the report — Right-click the section, then select **Print section**.
5) Select **TXT** as the **Format**.
6) Click **Browse**, then select where you want to save the file.
7) Click **OK**.

Tab-delimited text report files

Tab-delimited text files can be used for further processing.

The tab-delimited text files contain the statistics in tabulated tables. The tab characters and the operating environment-specific line endings delimit the text.

Structure of a tab-delimited text report file

Line no.	File content	Description
1	<report title>, <start time> - <end time>	Start and end time define the reporting period in format YYYY/MM/DD hh:mm:ss.
2	<empty line>	
3	<section content for each section>	Each section of the report follows the format described in the following table.

Line no.	Section content	Description
1	<section name>[; <section comment>]	Optional section comment with a leading semicolon (;) might follow the section name.
2	<empty line>	
3	<section data> | “No data”	Section data follows the format described in the following table. If the section contains no data, the text “No data” is displayed instead.

Line no.	Section content	Description
Section's last line	<empty line>	

Line no.	Section data content	Description
1	<table heading>	Tab delimited column labels. Some columns might not have a label, and labels might be padded with trailing spaces.
2	<empty line>	
3	<table rows>	Tab delimited values. Value in any given column can be empty. The column values are not padded.
Section data's last line	<empty line>	

Email reports

Reports can be emailed as soon as they are generated.

Before you begin

Define the SMTP server settings in the Management Server's properties.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Monitoring Configuration**.
2) Browse to **Reports** > **Design**.
3) Right-click a Report Design, then select **Start**.
4) Click the **Task** tab.
5) Enter an email address in the **E-mail Address**.
   To enter multiple email addresses, separate them with commas.
6) Click **OK**.

Related tasks

Modify Management Server elements on page 481

Status reporting

Create status reports on Security Engines or VPNs.

When creating a report design through the report editor, either create a section of type Status then select a status item or just create a section from a status item directly using **Section Create from Item**, then select an appropriate related element.

Example of reports

The example illustrates a common use for reports and general steps on how the scenario in question is configured.

Example: Identifying a disruptive internal user

An example of administrators looking for a specific trend in network activity to identify a specific user.

Administrators at Company A notice that downloads have gone up dramatically over the past week. They suspect that there might be an individual user that is excessively downloading files from the Internet. To confirm their suspicions, the Administrators decide to run a report that shows them who has used the most bandwidth in the network.

The administrators take the following steps:

1) Activate Log Accounting Information for each rule that allows connections from internal hosts to the Internet and install the policy. (Incoming connections to internal workstations are not allowed.)

2) Wait for a full workday for the logs with accounting information to be generated.

3) Create a filter that matches the IP address space of regular workstations as the source address and any external IP addresses as the destination address.

4) Create a Report Design based on the Engine Daily summary and attach the filter created in the previous step to the Report Design.

5) Increase the “Top Limit” value for the section “Traffic by src. IP” to see more results.

6) Generate a report for the previous day to check the traffic volumes for the top hosts.

Chapter 17

Filtering data

Contents

- Getting started with filtering data on page 333
- Creating filters on page 334
- Organizing Filter elements on page 344
- Examples of filters on page 346

Filters allow you to select data based on values that it contains. Most frequently, you use filters when viewing logs, but filters can also be used for other tasks, such as exporting logs and selecting data for reports.

Getting started with filtering data

Network traffic can generate a large amount of log data. You can use filters to select data for many operations such as viewing log entries in the **Logs** view or generating statistical reports.

Related concepts

Getting started with the Logs view on page 281
Creating filters on page 334
Organizing Filter elements on page 344

Related tasks

Use filters for browsing log entries on page 289

What filters do

Filters allow you to efficiently manage the large amounts of data that the system generates. Filters select entries by comparing values defined in the Filter to each data entry included in the filtering operation. The operation can use the filter to either include or exclude matching data.

You can use filters for selecting data in the following tasks:

- Browsing logs, alerts, and audit data.
- Browsing in all session monitoring views.
- Creating reports.
- Selecting which logs administrators who have restricted accounts or Web Portal User accounts are allowed to view.
- Defining how logs are highlighted in the Logs view.
- Forwarding logs to external third-party devices.

- Forwarding audit data to external third-party devices.
- Browsing which IP addresses, ports, and protocols are on the engines' block lists.
- Pruning log data.
- Exporting and deleting log data and alerts.
- Creating Correlation Situations to analyze engine and Log Server events.

Filter types

There are two types of filters in the SMC Client: local and permanent filters.

- Local filters are specific to a view or an element. You cannot use a local filter anywhere else in the SMC Client.
- Permanent Filter elements, which you can use anywhere in the SMC Client. There are predefined permanent Filters. You can also create new permanent Filters.

How filters are created

You can create filters in four basic ways.

- Based on criteria you define — You can create a local filter or permanent Filter element, and define filtering criteria in the Filter properties.
- Based on other Filters — You can create a copy of a Filter element or copy-and-paste parts of filters to other filters.
- Based on existing log entries — You can create local filters in Monitoring views where you can view logs, then save them as permanent Filter elements.
- Based on element configuration — Some local filters are created automatically by your selections in specific views or elements.

You cannot edit the predefined Filters, but you can create editable copies. Filter elements can be imported and updated when you activate new dynamic update packages, so the selection and names of predefined filters can change.

The default Filter elements are in the System Elements Category and have the Tags *System* or *Correlation* (for filters used in Correlation Situations).

Creating filters

You can create local filters that are specific to an element or view, then save them as permanent Filter elements to use globally.

Related concepts

Creating and editing local filters on page 337

Related tasks

Create or edit Filter elements on page 340

Basics of constructing filters

You construct filters using fields, values of the fields, and operations.

Filters are constructed from the following parts:

- The *fields* that you want to match in the data (for example, there are separate fields for source IP address and port in logs). You can filter data according to any field.
- The *values* in those fields that you want to match (for example, the exact port number or IP address you are interested in).
- *operations* define how the fields and values are matched to data entries (especially if there are several fields included as the filtering criteria).

A filter can have one or several fields. The more fields you have in a filter, the more specific the selection of log data becomes. For example, you can use the **IP source** field in the filter and get a selection of log data that matches the source IP address you specify. To limit the selection of log data even further, you could add a field for the destination port used.

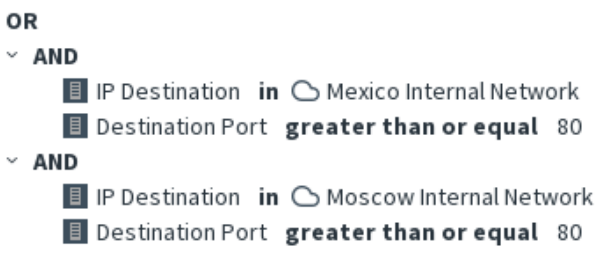
Different types of data entries contain different types of information, so the fields you add also restrict the general type of data that your filter matches. It is possible to create a filter that can never match any data if the combination of fields is not found in any single entry. However, everything depends on the general structure of the filter. It is possible to create filters that match related data in different types of entries using different fields as criteria. Depending on the field, you can define one to several values that you want to look for in the data. There are some operations (for example, Defined) for which a field value is not needed.

Operations define how field values in log data are compared to the field values defined in the filter. You can have as many operations in a filter as necessary, and you can also nest operations inside other operations. When you add two fields, you must always combine the fields with an operation. Each field in a filter is attached to one of these operations:

- Calculations (*BITWISE* and *SUM OF*)
- Comparisons (for example, *EQUAL TO*, *GREATER THAN*, *SMALLER THAN*)
- Logicals (*AND*, *NOT*, *OR*)

Filter example

Matching events with a filter



The illustration shows a Filter with several fields and operations. This Filter matches if the destination IP address is in the 192.168.11.0/24 network AND the destination port is 80 or greater OR if the destination IP address is in the 192.168.12.0/24 network AND the destination port is 80 or greater.

A data entry of a connection to host 192.168.11.10 on port 80 matches the first AND operation in the example filter. The same connection does not match the second AND operation in the Filter. Because the two AND operations are combined with OR, the Filter as a whole is considered a match and the data is selected for the task that is being carried out.

Filters that match a single value

A filter that matches a single source IP address:

```
Src Addr EQUAL TO 192.168.1.101
```

Where `Src Addr` is a field, `EQUAL TO` is the operation, and the IP address is a value.

Filters that match several values

A filter that matches any non-empty value for destination port:

```
Dst Port IS DEFINED
```

A filter that matches all destination ports between 1024 and 49152:

```
Dst Port BETWEEN 1024 AND 49152
```

A filter that matches any of three alternative destination ports:

```
Dst Port IN 51111, 52222, 53333
```

Complex filters that use logical operations

You can add the logical operations NOT, AND, and OR. The NOT operation negates the criteria you set.

A filter that matches all destination ports except ports between 1024 and 49152:

```
NOT
```

```
Dst Port BETWEEN 1024 AND 49152
```

When you add more than one field to a filter, you must define how the fields are used in relation to each other. You must use either AND (all field values must be present in the same entry) or OR (a data entry matches the filter if any one field value is found).

A filter that matches if the destination port is lower than 1024 and the source is a particular IP address:

```
AND
```

```
Src Addr EQUAL TO 192.168.1.101
```

```
Dst Port SMALLER THAN 1024
```

A filter that matches either of two destination ports:

```
OR
```

```
Dst Port EQUAL TO 80
```

```
Dst Port EQUAL TO 8080
```

You can apply the AND and OR operations to other AND and OR statements to create more complex statements. You can also negate whole AND and OR sections with NOT

Related concepts

Creating and editing local filters on page 337

Related tasks

Create or edit Filter elements on page 340

Creating and editing local filters

Local filters are valid only in the view or in the element in which they are created.

Local filters that you create in the following Monitoring views are temporary. They are only available until you close the view:

- Logs view.
- Connections view.
- Block list view.
- Users view.
- VPN SAs view.
- Routing Monitoring view.
- SSL VPNs Monitoring view.

Local filters created or edited for the following elements are saved with the element:

- Report Designs.
- Administrator elements.
- Log Server elements.
- Management Server elements.
- Correlation Situations.

You can save local filters as permanent Filter elements in all views. These Filter elements can then be used anywhere in the SMC Client.

Create local filters

You can create local filters in various views in the SMC Client.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Open a view or element for editing.

2) Depending on the view or dialog box, create the local filter in one of the following ways:

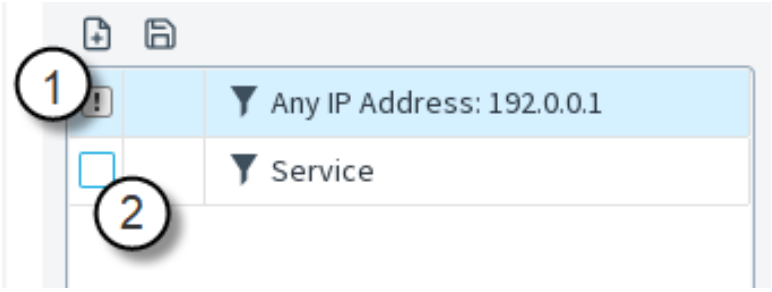
View/Dialog box		Configuration steps
Monitoring views (Logs, Connections, block list, VPN SAs, Users, Routing, or SSL VPNs)		On the **Filter** tab in the **Query** pane, click **New**, then select **New** > **Filter** to create a filter from a list of available fields, or select **New** > **Filter:<field name>** to create a filter based on a preselected field.  **Tip** You can also drag and drop log data rows onto the **Filter** tab
Reports view	Report Design	1) Select the name of the Report Design. 2) In the **Report Properties** pane, click the **Select Element** button next to the **Filter** field.
	Report Section	1) Select the Report Section. 2) In the **Section Properties** pane, click the **Select Element** button next to the **Filter** field.
	Report Item	Double-click the Report Item.
Administrator Properties dialog box		1) Switch to the **Permissions** tab. 2) Click **Select** next to **Filter**. The **Local Filter Properties** dialog box opens.
Log Server Properties or **Management Server Properties** dialog box		1) For the Log Server, switch to the **Log Forwarding** tab, then click **Add** to add a Log Forwarding rule. 2) For the Management Server, switch to the **Audit Forwarding** tab, then click **Add** to add an Audit Forwarding rule. 3) Double-click the **Filter** cell. The **Local Filter Properties** dialog box opens.
Correlation Situation Properties dialog box		1) Switch to the **Context** tab. 2) Click **Select** next to the **Context** field, then select **Compress**, **Group**, **Match**, or **Sequence**. 3) For Compress or Match, click **Select** next to the **Compress filter** field. The **Local Filter Properties** dialog box opens. 4) For Group or Sequence, double-click the **Event Match** cell. The **Local Filter Properties** dialog box opens.

3) (If not in a Monitoring view) Select **Add** > **New** > **Filter** to create a filter from a list of available fields or **Add** > **New** > **Filter:<field name>** to create a filter based on a preselected field.

4) Edit the filter in the **Filter Properties** dialog box.
   - If you are creating a filter from a list of available fields, enter an optional name in the **Name** field. Select the setting for **Undefined Value Policy**, then define the filtering criteria.
   - If you are creating a filter based on a preselected field, edit the filter properties. The available options depend on the field type.

5) Click **Apply**.
   The filter is added to the **Local Filter Properties** dialog box.
   In a Monitoring view, the filter is added to the **Filter** tab in the **Query** pane.

6) (Optional) Edit the local filter:
   - To negate a filter row, click the corresponding checkbox. This option filters out entries that match the filter.
   - If a row contains more than one item, click the operator cell to toggle between the **and** and **or** operators.



1	Negate a filter.
2	Toggle between operators.

7) (Optional) To save a local filter as a permanent Filter element, click **Save**. In a Monitoring view, click **Save**. You can use Filter elements in any view.

8) Click **OK**. In a Monitoring view, click **Apply**.

Related tasks
Save local filters as permanent Filter elements on page 340

Save local filters as permanent Filter elements

Local filters are specific to the view or element for which they have been created. Saving local filters as permanent Filter elements allows you to apply these filters anywhere in the SMC Client.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Create a local filter.
2) In the **Local Filter Properties** dialog box, click **Save**. In a Monitoring view, click **Save**.
3) Enter a name for the filter in the **Name** field, then click **OK**.

Related tasks

Create local filters on page 337

Create or edit Filter elements

All the permanent Filter elements are stored and can be edited from one view.

Tip

An easy way to create a permanent Filter element is to create a local filter, then save it as a Filter element.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Monitoring Configuration**.
2) Expand the **Other Elements** branch.
3) To create a new filter:
   a) Right-click **Filters**.
   b) Select **New** > **Filter**.
4) To edit an existing filter:
   a) Browse to **Filters** > **All Filters**.
   b) Right-click the filter, then select **Properties**.
5) Edit the filter properties, then define the filtering criteria.

6) Click **OK**.

Related concepts
Benefits of filtering log entries on page 289

Related tasks
Save local filters as permanent Filter elements on page 340
Add or edit criteria in Filter elements on page 341

Add or edit criteria in Filter elements

You can add new criteria or edit the criteria of Filters that you have created.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Monitoring Configuration**.

2) Browse to **Other Elements** > **Filters** > **All Filters**.

3) Right-click a filter you have created, then select **Properties**.



Tip

To edit a filter based on an existing system filter, right-click the filter, then select **New** > **Duplicate**.

4) Select the setting for **Undefined Value Policy**.

This setting defines how data entries are matched when a field is included in the filter, but is not present in the entry. For example, a filter defines a range of destination ports, but the operation encounters a log entry from a protocol that does not use ports, such as ICMP.

5) If there is no logical operation (AND or OR) at the correct level, add one using the shortcut buttons above the editing pane.

You can nest logical operations to create more complex filters. For example, you can create two AND sections under an OR condition to match either of the two sets of criteria.

6) To change a logical operation, right-click the operation, select **Change To**, then select the new operation.

7) To add a field:

a) Right-click the logical operation to which you want to add a field, then select **New** > **Select**.

b) Click **Fields**, then browse to the field group that contains the field you want to add, or browse to **All Fields** for a list of all available fields.

c) Select the field, then click **Select**.

8) To edit the field:

a) Right-click the field that was added, then select **Edit**.

b) Select the **Comparison**.

The available comparison selection depends on the selected field and whether the field already contains one or more values.

The most common comparisons are:

- **Any Value** — Allows you to match any non-empty value in the field.
- **Between** — Allows you to match a range (for example, a range of TCP/UDP ports).
- **Contains** — Allows you to match any of several alternative values (for example, both an IPv4 address and an IPv6 address).
- **In** — Allows you to match a single value (for example, an IP address or Network element).

c) Depending on the comparison and type of field, define the values that you want the filter to match in one of the following ways:

- Enter one or more values. For the **In** or **Contains** comparison, click **Add** to add the entered value to the value list.
- Double-click the empty space in the value list, then select an element.
- To edit a value that has been added to the value list, double-click the value.
- To remove a value, right-click the value, then select **Remove**.

d) Click **Apply**.

9) To remove criteria, right-click the criteria, then select **Remove Row** or **Remove**.

If you select **Remove Row**, all criteria nested under the row is moved up one level.

If you select **Remove**, all criteria nested under the row is also removed.

10) Click **OK**.

Related concepts
Log entry fields

How missing values are handled in Filter elements

You can adjust what happens when the Filter element is matched to data that does not contain any value for a field that the filter defines.

By default, log data matches the filter only if all fields in the filter are also found in log data.

Because there are different types of data entries, some entries might not contain any value for some field that a filter contains. For example, an Alert entry warning you that the monitoring connection from a Engine has been lost does not contain any source or destination IP address information. The reason for this is that the entry is not related to traffic processing. If you apply a filter that matches an IP address in the **Logs** view, the Alert is filtered out of the view. Missing values that cannot be verified as matching or non-matching are called *undefined values* in the configuration.

To define in more detail how missing fields are handled, you have two options:

- The *Undefined value policy* setting defines whether log data matches the filter if there are missing fields.
- The *Any Value Comparison* operation allows you to define specific fields in the filter that the log data must always have. The value that the field contains is not taken into account. Data entries that do not have these fields do not match the filter.

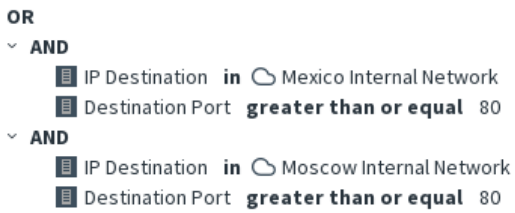
You can use one of the four Undefined value policy settings to define how missing values are handled. The setting works differently depending on the structure of the filter. The results of logical operations (AND, OR, NOT) in the filter depend on the Undefined value policy setting. A logical operation is typically either *true* or *false*. However, if a field in the filter does not exist in a data entry, the logical operation is left *undefined*.

Undefined value policy settings

Setting	Description
False by comparison	A Comparison operation is *false* if log data does not have all fields used in the filter. Depending on the structure of the filter, the log data does or does not match the Filter. For example, if the outermost operation in the filter is AND, the log data does not match the filter if any of the inner operations are *false*.
False by filter	Log data does not match the filter if the outermost operation in the filter is *undefined* because log data does not have all fields used in the filter. The filter is *false*.
True by filter	Log data matches the filter if the outermost operation in the filter is *undefined* because log data does not have all fields used in the filter. The filter is *true*.
Undefined	If the outermost operation is *undefined* because log data does not have all fields used in the filter, the *undefined* result is passed to the component that uses the filter. The handling of the *undefined* result varies according to the component that uses the filter. In most cases, this setting works in the same way as “False by filter”. If the outermost operation is *undefined* because log data does not have all fields in the filter, the data does not usually match the filter.

Undefined value policy settings

Undefined values when matching an event



A filter has the **IP destination** and **Destination port** fields. ICMP traffic, for example, does not have the **Destination port** field. If ICMP traffic is matched with the example filter, the filtering results vary according to the selected Undefined value policy:

- **False by comparison** — The AND operations are *false*. As a result, the OR operation is also *false*. The event does not match the filter.
- **False by filter** — The AND operations are *undefined* (neither *true* nor *false*). As a result, the OR operation is also *undefined*. The setting interprets the *undefined* result as *false*. The event does not match the filter.
- **True by filter** — The AND operations are *undefined* (neither *true* nor *false*). As a result, the OR operation is also *undefined*. The setting interprets the *undefined* result as *true*. The event matches the filter.
- **Undefined** — The AND operations are *undefined* (neither *true* nor *false*). As a result, the OR operation is also *undefined*. The Undefined setting passes the *undefined* value to the component that uses the log data. The handling of the data varies according to the component. Most components handle the data in the same way as False by filter, so that the event does not match this filter.

Organizing Filter elements

Adding Filter Tags to permanent Filter elements that you create makes it easier to find filters when you want to use them. You can select several Filter Tags for each Filter element that you create.

Create Filter Tag elements

Filter Tags help you organize Filter elements.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Monitoring Configuration**.
2) Browse to **Other Elements** > **Filters**.
3) Right-click **Filters**, then select **New** > **Filter Tag**.
4) Enter a name in the **Name** field.
5) Click **OK**.

Related tasks

Change the Filter Tags of Filter elements on page 345

Change the Filter Tags of Filter elements

You can add Filter Tags to Filter elements, to make it easier to find a filter.

Tip

These steps show how to add or remove Filter Tags to multiple Filters, but you can also edit Filter Tags on the **Tags** tab in the properties of an individual Filter element.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Monitoring Configuration**.

2) Browse to **Other Elements** > **Filters**, then select one or more Filter elements.

   Tip

   If you want to remove all references to a Filter Tag so that you can delete the Filter Tag, select all the filters in the **All Filters** branch. The selection includes filters that do not reference the Filter Tag that you want to remove.

3) To add a tag:

   a) Right-click a Filter, then select **Add Tag** > **Filter Tag**.

   b) Select the Filter Tag you want to add.

      Note

      We do not recommend that you assign the **System** or **Correlations** Filter Tags.

   c) Click **Select**.

4) To remove a tag:

   a) Right-click a Filter, then select **Remove Tag** > **Remove**.

   b) Select the Filter Tag you want to remove.

      Note

      You cannot remove the **System** or **Correlations** Filter Tags from predefined system elements.

   c) Click **Select**.

Related tasks

Create Filter Tag elements on page 344

Examples of filters

These examples illustrate some common uses for filters and general steps on how the scenarios are configured.

Example: Create a filter for log entries associated with specific users

This scenario shows an example of using filters to include logs in a report.

Company A wants a report of users who have authenticated themselves within a certain time frame. To create the report, the company's administrator needs a filter to select the logs concerning the authenticated users. The administrator:

1) Creates a Filter element.

2) Selects the **Auth. User** field to filter the user names of authenticated users.

3) Selects the **in** operation.

4) Adds the wildcard * as the value to the **Auth. User** field to match all authenticated users in log data.

Example: Create a filter for pings in a network

This scenario shows an example of using filters to exclude logs from a report.

Company B's administrator has noticed that the number of ping attempts (ICMP echo requests) in the internal network has increased. The administrator wants a report of all recent pings in the local network to make sure an outsider has not taken over the servers in the internal network. The administrator frequently pings from the HOST 2 workstation in the internal network. The administrator knows that pings coming from HOST 2 are legitimate, and wants to exclude pings from HOST 2 from the report.

The administrator needs a new filter for generating the report. The administrator:

1) Creates a Filter element in which the source IP address field in log data is compared to the internal network's addresses, and the ICMP type is compared to Echo.

2) Adds a condition that the IP address in the log data must not belong to the HOST 2 workstation.

Part V
Controlling Security Engines

Contents

- Controlling Security Engine operation on page 349
- Working on the Security Engine command line on page 363

You can command and set options for engines through the SMC Client or on the engine command line. You can also stop traffic manually.

Chapter 18

Controlling Security Engine operation

Contents

- Commanding Security Engines remotely on page 349
- Set Security Engine options on page 356
- Change a NetLink state manually on page 359
- Disable cluster nodes temporarily on page 360
- Re-enable disabled cluster nodes on page 361
- Editing Security Engine configurations on page 362

You can command and set options for Engines, Layer 2 Engines, IPS engines, Master Engines, Virtual Engines, Virtual IPS engines, and Virtual Layer 2 Engines through the SMC Client.

Commanding Security Engines remotely

You can send commands to Security Engines remotely through the SMC Client.

You can control Engines, Layer 2 Engines, IPS engines, Master Engines, Virtual Engines, Virtual IPS engines, and Virtual Layer 2 Engines through each engine element's right-click menu. The commands available depend on the type of component. In a cluster, the commands that affect the operating state of the engines can only be given to the individual nodes, not to the whole cluster.

You can also give commands and set options for more than one engine at a time by Shift-selecting or Ctrl-selecting the elements.



Note

For abnormal situations, there are limited tools for giving some basic commands (such as go online or offline) through the engine's command line interface. Under normal circumstances, you should control the engines remotely through the SMC Client.

Turn Security Engines online

You can turn engines in the offline state online through the right-click menu.

When engines are in the offline state, the status icon is blue and the status text reads offline. You can turn engines online if there are no configuration issues that would prevent the node or cluster from operating normally. Typical issues that can prevent a node from going online include policy issues, automatic tests failing, or heartbeat connection problems between nodes in clusters.

Note

You might also be able to give commands to nodes in the unknown state (gray icon), but you might not see a change of status. Because the actual operating status is not available, the node might already be online. In this case, you might receive an error if you try to command the node online.

Note

If the cluster is set to standby mode, only one node at a time can be online. Commanding a standby node online switches the current online node to standby.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Dashboard** > **Engines Dashboard**.
2) Expand the nodes of the engine that you want to turn online.
3) Right-click the engine node, then select **Commands** > **Go Online** or **Commands** > **Lock Online**.
4) (Optional) In the **Confirmation** dialog box that opens, enter an **Audit Comment**.
   The comment is included in the audit log entry that is generated.
5) Click **Yes**.
   The engine is turned online shortly.

Related tasks

Adjust general Engine clustering options on page 677
Adjust IPS clustering options on page 679
Adjust Layer 2 Engine clustering options on page 680
Adjust general Master Engine clustering options on page 681
Troubleshoot Security Engines that do not go or stay online on page 1385

Turn Security Engines offline

In the offline state, engines stop processing traffic, but remain otherwise operational and ready to be turned online again.

Engines in the offline state can be turned on either automatically or by an administrator's command, depending on the configuration.

CAUTION

When you turn a node offline, it stops processing traffic. On Engines, Layer 2 Engines, and Master Engines, traffic is stopped unless other cluster nodes can take over. On Virtual Engines, Virtual IPS engines, and Virtual Layer 2 Engines, traffic is always stopped. On IPS engines, the behavior depends on the Failure Mode of the interfaces.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Dashboard** > **Engines Dashboard**.
2) Expand the nodes of the engine that you want to turn offline.
3) Right-click the node, then select **Commands** > **Go Offline** or **Commands** > **Lock Offline**.
4) (Optional) In the **Confirmation** dialog box that opens, enter an **Audit Comment**.
   The comment is included in the audit log entry that is generated.
5) Click **Yes**.
   The engine is turned offline shortly.

Related concepts

Configuring interfaces for IPS engines on page 572

Related tasks

Turn Security Engines online on page 349

Set nodes to standby mode

When a cluster runs in standby mode, only one node at a time processes traffic. The other running nodes are on standby.

Standby nodes monitor the traffic so that they can take over if the active node fails. Only one node at a time is in the online state, and the rest are either in the standby or offline state. When you command an online node to standby, a standby node in the cluster (if there is one) automatically goes online to take over the traffic.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Dashboard** > **Engines Dashboard**.
2) Expand the nodes of the engine that you want to set to standby mode.
3) Right-click a node, then select **Commands** > **Standby**.
4) (Optional) In the **Confirmation** dialog box that opens, enter an **Audit Comment**.
   The comment is included in the audit log entry that is generated.
5) Click **Yes**.
   The node is set to standby mode shortly.

Related concepts
Getting started with advanced Security Engine settings on page 675

Switch Active Node

If needed, you can switch the engine active node for high availability backup unit. For more information on high availability backup unit, refer to the *Configure a backup unit or a connection synchronization for external high availability* topic.

Steps

1) Select **Dashboard** > **Engines Dashboard**.
2) Right-click the engine, then select **Commands** > **Switch Active Node**.
3) Click **Yes**.
   The active node is switched shortly.

Note

You can navigate to the **General** tab, in the **Info** section of the SMC UI to view information regarding the **Active Node** and the **Backup Node Version** for the engine.

Related tasks
Configure a backup unit or a connection synchronization for external high availability on page 524

Reboot nodes

In rare cases, you might need to reboot Forcepoint Network Security Platform nodes.

CAUTION

If you are rebooting a cluster, reboot the nodes one by one to avoid breaks in service. If you command all nodes to reboot, all nodes reboot at the same time.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Dashboard** > **Engines Dashboard**.
2) Expand the nodes of the engine that you want to reboot.
3) Right-click the node, then select **Commands** > **Reboot**.

Next steps

Monitor the rebooting process by following the changes in the status of the element.

Power off the Security Engine

If needed, you can power off the Security Engine.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Dashboard** > **Engines Dashboard**.
2) Expand the nodes of the Security Engine that you want to power off.
3) Right-click the node, then select **Commands** > **Power Off**.

Reset the Security Engine appliance to factory settings

If needed, you can reset the Security Engine appliance to factory settings. To make sure that confidential information is removed, the stored data can be overwritten multiple times.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Dashboard** > **Engines Dashboard**.
2) Expand the nodes of the Security Engine that you want to reset.
3) Right-click the node, then select **Commands** > **Reset to Factory Settings**.
4) Enter how many times you want the stored data on the file system to be overwritten.
5) Click **OK**.

Refresh the currently installed policy

You can reinstall the currently installed policy of one or more components to transfer configuration changes since the last policy installation.

Each type of Security Engine has its own type of policy. Inspection Policies are used by all types of Security Engines.

Note

In clusters, all nodes must be either operational or explicitly disabled for the policy installation to succeed.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click the Security Engine, then select **Current Policy** > **Refresh**.
3) Click **OK**.

Related concepts

Getting started with policies on page 799

Related tasks

Disable cluster nodes temporarily on page 360

Remove Virtual Engines from Master Engines

You can remove a Virtual Engine from a Master Engine if the Virtual Engine is no longer needed.

When you remove a Virtual Engine from a Master Engine, the Virtual Engine goes offline and stops processing traffic. The Virtual Engine element is kept in the Security Management Center. You can associate the Virtual Engine with a different Virtual Resource to activate the Virtual Engine on a different Master Engine. The Master Engine must host Virtual Engines in the same role as the Virtual Engine you want to activate.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click the Virtual Engine, then select **Commands** > **Remove Virtual Engine from Master Engine**.
   A new tab opens to show the progress of the operation. The Virtual Engine is removed from the Master Engine.

Move Virtual Engines from one Master Engine node to another

You can optionally move a Virtual Engine from one node to another on the same Master Engine for load balancing.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click the Virtual Engine, then select **Commands** > **Move to Master Engine Node**.
3) Select the Master Engine node on which to run the Virtual Engine, then click **Select**.

Note

The node selection might change later due to automatic load balancing.

Join an Security Engine to a Domain

You can join an Engine to a Domain to allow Integrated Windows Authentication (IWA) of the users. For more details, refer to the *Configuring the Integrated Windows Authentication* topic.

Steps

1) Select **Dashboard** > **Engines Dashboard**.
2) Right-click the engine node, then select **Commands** > **Join to an AD domain**. The **Join to an Active Directory domain** dialog box is displayed.
3) Configure the settings.
4) Click the **OK** button.

Un-join an Security Engine from a Domain

Follow these steps to remove an engine from the domain.

Steps

1) Select **Dashboard** > **Engines**.

2) Right-click the engine node, then select **Commands** > **Unjoin from AD domain**. The **Unjoin Active Directory domain** dialog box is displayed.

3) Configure the settings.

4) Click the **OK** button.

Set Security Engine options

You can set options for Engines, Layer 2 Engines, IPS engines, Master Engines, Virtual Engines, Virtual IPS engines, and Virtual Layer 2 Engines through the SMC Client.

Enable or disable status monitoring

By default, monitoring is automatically activated for all engines, but can be turned off as necessary.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click the Security Engine, then select **Options**.

3) Deselect or select **Monitored**.
Shortly, the status changes to **Not Monitored** and the icons associated with the element turn white.

Related concepts
Getting started with monitoring the system on page 211

Enable or disable diagnostics

Diagnostics mode provides more detailed log data for troubleshooting purposes.

Note
Disable the diagnostics after troubleshooting to avoid overloading the Log Server with log data.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click the Security Engine, then select **Options** > **Diagnostics**.

3) Select or deselect the **Enable diagnostic logs for the selected features** option at the top of the dialog box.

4) Select the features for which you want diagnostic log data on the engine when diagnostics are enabled.

5) Click **OK**.
The changes are applied immediately.

Related concepts
Getting started with the Logs view on page 281

Enable or disable user database replication

You can enable or disable the replication of the Management Server's internal LDAP database to a Engine or a Master Engine.

The Management Server's internal LDAP database stores accounts for end users for authentication purposes. Engines have a local replica of the Management Server's internal LDAP database. By default, all changes are immediately replicated from the Management Server's internal LDAP database to the local replicas on Engines.

Master Engines have one combined local replica of the Management Server's internal LDAP database for each Domain in which a Virtual Engine has users in the internal LDAP database. By default, changes are replicated from the Management Server's database to the local replicas on the Master Engines. The information that is replicated to Master Engines depends on the User Authentication configuration of the Virtual Engines.

Note

Changing the replication of the Management Server's internal LDAP database for a Master Engine also changes the replication of the Management Server's internal LDAP database for all Virtual Engines hosted by the Master Engine.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click the Security Engine, then select **Options**.

3) Select or deselect **User DB Replication**.

Related concepts
Getting started with directory servers on page 1103

Enable or disable status surveillance

Status surveillance generates an alert when engines change to an unknown state.

By default, there is no warning to administrators if the status of the engines changes to an unknown state. You can optionally activate the status surveillance feature. The status surveillance feature generates an alert if a single engine or none of the engines in a cluster do not send a status update for 15 minutes.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click the Security Engine, then select **Options**.
3) Select or deselect **Status Surveillance**.

Next steps

Make sure that System Alerts are escalated so that the notification is sent if status surveillance detects a failure.

Related concepts

Alert escalation and how it works on page 411

Enable or disable SSH access

Secure remote access to the engines is provided by the SSH daemon process. This process can be started and stopped remotely.

For maximum security, we recommend disabling SSH access whenever it is not used.

Alternatively, you can enable and disable SSH access when logged on to the node.

SSH uses TCP port 22. Make sure that the connections are allowed in the policies of any Engines or Layer 2 Engines involved in the communications (including the Engine or Layer 2 Engine that you are trying to contact).

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Dashboard** > **Engines Dashboard**.
2) Expand the nodes of the engine for which you want to enable or disable SSH access.
3) Right-click a node, then select **Commands** > **Enable SSH** or **Commands** > **Disable SSH**.

Result

The SSH process is started or stopped on the engine.

Related tasks
Configuring SSH access to the SMC Appliance on page 146
Reconfigure Security Engine settings on page 365

Change the Security Engine root password

The password for access to the Security Engine command line can be changed remotely through the SMC Client.

The user account for accessing the command line is always root. Alternatively, if you remember the old password, you can change the password when logged on to the node.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Dashboard** > **Engines Dashboard**.
2) Expand the nodes of the engine for which you want to change the root password.
3) Right-click a node, then select **Commands** > **Change Password**.
4) Enter the new password in both fields, then click **OK**.
   The new password is effective immediately.

Related tasks
Reconfigure Security Engine settings on page 365
Change your own local Security Engine password on page 389

Change a NetLink state manually

You can manually command NetLinks to the Enable or Disable state.

To change a NetLink's state, it must be operational. You cannot send commands to NetLinks that have the unknown (gray) status.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Dashboard** > **Secure SD-WAN Dashboard**.
2) Browse to **Branches**, and then select a branch.

3) In the **ISP Information** section:

   a) To enable the NetLink, select ⋮ **Tools** > **Force NetLink Enable**.

   b) To disable the NetLink, select ⋮ **Tools** > **Force NetLink Disable**.

Related tasks
Create NetLink elements for Multi-Link configuration on page 737

Disable cluster nodes temporarily

Disabling a cluster node allows continued management of the other cluster members if one node goes out of operation.

Before you begin

Turn off the cluster node that you want to disable and disconnect the network cables.

You can disable nodes in Engine Clusters, IPS Clusters, Layer 2 Engine Clusters, or Master Engines. When you disable a node, you can physically remove it from the cluster without removing its definition from the system.

Disabling a node indicates to the other nodes and the Management Server that it is not necessary to try to contact it. Disabling a node prevents unnecessary communication attempts, alerts, and test failures. Disabling a node also allows policy installations on the other nodes when one node is shut down or malfunctions. No commands can be sent to a disabled node and no monitoring information is available for it.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the SMC Client, select **Engine Configuration**.

2) Right-click the Security Engine, then select **Edit <element type>**.

3) In the navigation pane on the left, browse to **General** > **Clustering**.

4) In the **Nodes** table, select **Disabled** for the nodes you want to disable.

5) Click **Save and Refresh**.

Related tasks
Re-enable disabled cluster nodes on page 361

Re-enable disabled cluster nodes

You can re-enable nodes in a cluster that you have temporarily disabled.

Before you begin

Before connecting network cables to the disabled node or to a replacement for it, set the node to the initial configuration state using the Security Engine Configuration Wizard (sg-reconfigure) on the engine command line.

Note

If you reintroduce a disabled node that has a working configuration, the node must receive the heartbeat traffic from other nodes and accept it based on certificates. Otherwise, the node considers itself the only available cluster member and goes online. Cluster nodes that do not communicate with each other can prevent the whole cluster from processing traffic.

When a Engine Cluster, IPS Cluster, Layer 2 Engine Cluster, or Master Engine node has been disabled, its configuration is typically made obsolete by policy installations done on the other cluster nodes. Having an obsolete configuration prevents the node from operating normally and might in some cases disturb the operation of the whole cluster.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the SMC Client, select **Engine Configuration**.
2) Right-click the Security Engine, then select **Edit <element type>**.
3) In the navigation pane on the left, browse to **General** > **Clustering**.
4) Deselect the **Disabled** option in the **Nodes** table for the nodes you want to re-enable, then click **OK**.
5) Click **Save and Refresh** to ensure that all nodes have the same configuration.

   Note

   If the policy installation is unsuccessful, return the previously disabled node to the initial configuration state.

6) (Optional) In the **Engines Dashboard** view, right-click the node, then select **Commands** > **Go Online** or **Commands** > **Standby** to return the node to operation.

   The node is set to online or standby mode shortly.

Related tasks

Reconfigure Security Engine settings on page 365

Editing Security Engine configurations

The Security Engines can be configured in various ways.

- The network card drivers, mapping of physical interfaces on the network cards to Interface IDs, and speed/duplex settings are defined using the Security Engine Configuration Wizard. You can use the Security Engine Configuration Wizard on the engine command line or in a web browser.
- Other engine-specific settings are defined in the Engine Editor in the SMC Client.

Related concepts

Editing existing Security Engines on page 509

Related tasks

Disable cluster nodes temporarily on page 360
Re-enable disabled cluster nodes on page 361
Reconfigure Security Engine settings on page 365
Convert a Single Engine to a Engine Cluster on page 514
Convert a Single IPS engine to an IPS Cluster on page 516
Convert a Single Layer 2 Engine to a Layer 2 Engine Cluster on page 517

Chapter 19

Working on the Security Engine command line

Contents

- Considerations for working on the Security Engine command line on page 363
- Access the Security Engine command line on page 364
- Reconfigure Security Engine settings on page 365
- Create Security Engine scripts on page 367
- Send commands to Virtual Security Engines on page 368

Although the engines are managed remotely, some operations on the Linux command line on the engines are useful for troubleshooting and local maintenance operations.

Considerations for working on the Security Engine command line

Nearly all engine configuration is done through the SMC Client, but some engine settings and options must be defined and configured on the command line.

What you can do on the Security Engine command line

- Reconfigure the engine's keyboard layout, time zone, network card settings, and network card to Interface ID mapping.
- Create scripts that run when the engine changes its state.
- Establish contact between the engine and the Management Server.
- Manually revert to the previous configuration.
- Run various troubleshooting tools, both general and specific to Forcepoint Network Security Platform.

Limitations of the Security Engine command line

Changes made on the engine command line apply only to the node on which they were made. If you want to change settings for other engines, such as all nodes in a cluster, you must make the same changes separately on the command line of each engine.

Some engine configuration options, such as network interface settings, cannot be changed through an SSH console. To be able to change these settings, you must connect using a serial cable or connect a display and keyboard directly to the engine hardware.

The Management Server contact settings that are displayed in the Security Engine Configuration Wizard (sg-reconfigure) do not show the engine's actual working configuration (transferred whenever the engine's policy is installed or refreshed). The Security Engine Configuration Wizard displays the values that were set when the node was initialized.

If you are not a root user on the engine, your permissions to execute commands might be limited. This might be the case if your SMC account has been replicated on the engine and your permissions have been limited in the local sudo configuration file.

What do I need to know before I begin?

All command-line tools that are available for single Security Engines are also available for Virtual Security Engines that have the same role. However, there is no direct access to the command line of Virtual Security Engines. Commands to Virtual Security Engines must be sent from the command line of the Master Engine that hosts the Virtual Security Engines.

Related tasks

Send commands to Virtual Security Engines on page 368
Access the Security Engine command line on page 364
Reconfigure Security Engine settings on page 365
Create Security Engine scripts on page 367

Access the Security Engine command line

There are several ways you can connect to the Security Engine and access the Security Engine command line.

Steps

1) Connect to the Security Engine in one of the following ways:
   - Physically using a serial cable and a terminal console program with these settings:
     - **Bits per second** — 115,200
     - **Data bits** — 8
     - **Parity** — None
     - **Stop bits** — 1.

     Note

     The serial console port speed is 115,200 bps in most Security Engine appliances. The speed is 9600 bps in older Security Engine appliance models. See the hardware guide for your Security Engine appliance model for more information.
   - Physically using a monitor and keyboard connected directly to the appliance.
   - Remotely using an SSH client. SSH access to the Security Engine can be enabled and disabled through the SMC Client.

2) If you have root administrator permissions, log on and enter the Security Engine password.
   If you forget the password, you can change it in the SMC Client.

3) If you do not have root permissions and your administrator account has been replicated on the Security Engine, log on using your administrator user name and personal password for the Security Engine.
   If you forget the Security Engine-specific password, ask an SMC administrator with Manage Administrators rights to create you a new one in the SMC Client.

Related tasks
Enable or disable SSH access on page 358
Change the Security Engine root password on page 359
Reconfigure Security Engine settings on page 365
Create Security Engine scripts on page 367

Reconfigure Security Engine settings

On the command line of the Security Engine, you can use the Security Engine Configuration Wizard to change settings that were defined during the installation of the Security Engine.

The Security Engine Configuration Wizard also allows you to re-establish a trust relationship between the Security Engine and the Management Server if the trust is lost.

Note

On Security Engines that are fully configured, you can change each setting individually without changing the other settings. All steps are optional.

Steps

1) Start the Security Engine Configuration Wizard using one of the following commands:
   - `sg-reconfigure --no-shutdown` — The Security Engine Configuration Wizard starts without shutting down the Security Engine. You cannot change network interface settings in this mode.
   - `sg-reconfigure` — The Security Engine shuts down and the Security Engine Configuration Wizard starts. All options are available if you have a local connection. If you have a remote SSH connection, you cannot change network interface settings.

2) Change the general settings.
   - Change the keyboard layout for command-line use.
   - Change the time zone for command-line use.
   - Change the host name of the engine.
   - Enable or disable SSH access to the engine command line.

   Note

   Unless you have a specific reason to enable SSH access to the engine command line, we recommend leaving it disabled.

3) Change the password for the root user account.
   a) Highlight **Change**, then press **Enter**.
   b) Enter and confirm the new password for the root user account.
   c) Highlight **OK**, then press **Enter**.

4) Change the bootloader password.
   The bootloader password prevents unauthorized editing of parameters in the second-level grub menu on the Security Engine.
   a) Highlight **Change**, then press **Enter**.
   b) Enter and confirm the new bootloader password.
   c) Highlight **OK**, then press **Enter**.

5) Change the network card settings and the mapping of network cards to Interface IDs.

6) Change the settings on the **Prepare for Management Contact** screen.



Note

The Management Server contact details are not used by the Security Engine after a policy has been installed from the Management Server. They are shown for your reference only.

- To re-establish the trust relationship between the Security Engine and the Management Server, select **Contact Management Server**, then enter a new one-time password.
  Select this option when you want to replace a missing or expired certificate, or if the trust relationship with the Management Server is lost for any other reason, such as changing the Management Server's IP address.



CAUTION

If there is a Engine or Layer 2 Engine between a remote Security Engine and the Management Server, you must allow the connection in the Engine or Layer 2 Engine Access rules. If there is a NAT device between a remote Security Engine and the Management Server, you must also configure NAT rules for the connection in the Engine Policy. Otherwise, the Security Engine cannot contact the Management Server.

- To reset the Security Engine to the post-installation state, select **Switch to Initial Configuration**.



CAUTION

Selecting this option removes all configuration and policy information that has been transferred to the Security Engine. The post-installation state uses a policy that allows communication only between the Security Engine and the Management Server. You must install a policy on the Security Engine before it can be operational again.

Related concepts

Connect Security Engines to the SMC on page 631

Related tasks
Access the Security Engine command line on page 364

Related reference
Security Engine commands on page 1445

Create Security Engine scripts

Security Engine scripts run when the Security Engine changes its state.

The script names and locations cannot be changed. If the scripts are not found, engine operation continues as normal. If a script is found, it is executed and a log entry is created. To stop scripts from running, you must delete or move the script.

Note

If you want to use a script in a cluster, create or copy the script on all nodes in the cluster. Then all nodes function in the same way when their state changes.

Steps

1) Create a text file with the commands you want the engine to execute (the first line of the script must be `#!/bin/sh`) in one of the following ways:
   - Create and edit the script on the engine's command line using the vi text editor.
   - Create and edit the script on a different host and transfer the file to the engine, for example, using SSH.

2) Save the script in the correct folder on the engine.

   Possible scripts on the engines

Triggering event	Script location and name
Engine operating system boots	/data/run-at-boot
Administrator refreshes or installs the policy	/data/run-at-policy-apply
Engine enters the Online state	/data/run-at-online
Administrator issued the 'Lock Online' command	/data/run-at-locked-online
Engine enters the Offline state	/data/run-at-offline
Administrator issued the 'Lock Offline' command	/data/run-at-locked-offline
Engine enters the Standby state	/data/run-at-standby

3) Make the file executable by typing the following command:
   `chmod a+x /data/<script name>`

Result

The script is executed whenever the engine encounters the triggering event for running the script.

Related tasks
Access the Security Engine command line on page 364

Related reference
Security Engine commands on page 1445

Send commands to Virtual Security Engines

Commands to Virtual Security Engines are sent from the command line of the Master Engine that hosts the Virtual Security Engines.

All command-line tools that are available for single Security Engines are also available for Virtual Security Engines that have the same role.

Steps

1) Connect to the command line on the Master Engine.

2) Enter commands in the following format:

`se-virtual-engine [options]`

Options for `se-virtual-engine` Command

Option	Description
`-h | --help`	Shows the help message for the se-virtual-engine command.
`-l | --list`	Lists the active Virtual Security Engines.
`-v <ID> |` `--virtual-engine=<ID>`	Specifies the ID of the Virtual Engine on which to execute the command.
`-e | --enter`	Enters the command shell for the Virtual Engine specified with the `-v` or `--virtual-engine` option. To exit the command shell, type exit. Using the command shell is recommended if you want to send multiple commands to the Virtual Engine.
`-E "<command [options]>" |` `--execute="<command [options]>"`	Executes the specified command on the Virtual Engine specified with the `-v` or `--virtual-engine` option. Executing individual commands is recommended if you only want to send a few commands to the Virtual Engine. You can also execute individual commands to send the same command to multiple Virtual Security Engines.

Related tasks

Access the Security Engine command line on page 364

Related reference

Security Engine commands on page 1445

Part VI

SMC configuration

Contents

- Administrator accounts on page 373
- Alert escalation on page 411
- Domain elements on page 433
- Getting Started with the Web Portal on page 447
- Using the SMC Client in a web browser on page 451
- SMC Client downloads from the Management Server on page 455
- Configuring the Log Server on page 457
- Configuring SMC servers for high availability on page 467
- Reconfiguring the SMC and Security Engines on page 481

SMC Manager configuration allows you to customize how the SMC components work.

Chapter 20

Administrator accounts

Contents

- Getting started with administrator accounts on page 373
- How administrator accounts work on page 374
- Default administrator account elements on page 375
- Administrator account configuration overview on page 377
- Creating Administrator Role and Access Control List elements on page 377
- Add administrator accounts on page 379
- Enforce an approval workflow on page 383
- Restrict the log data an administrator can view on page 383
- Customize log colors for administrators on page 384
- Replicate administrator accounts on page 385
- Enable and define password policy settings on page 386
- Change administrator passwords on page 387
- Authenticate administrators using OpenID authentication method on page 389
- Authenticate administrators using RADIUS or TACACS+ methods on page 392
- Authenticate administrators on engines using the Radius authentication method on page 394
- Authenticate administrators using SAML v2 authentication method on page 396
- Using LDAP authentication for administrators on page 398
- Authenticate administrators using certificate-based authentication on page 399
- Disable administrator accounts on page 405
- Delete administrator accounts on page 406
- API client accounts and how they work on page 407
- Configure SMC API on page 407

Administrator accounts define administrator rights and permissions in the SMC.

Getting started with administrator accounts

An administrator account specifies the actions for which the administrator has permissions, such as creating elements and browsing logs.

You can define administrator rights for each administrator. You can give different permissions to each administrator globally, for specific administrative Domains, for specific groups of elements, and even for individual elements. Depending on the element, there are different levels of access that you can grant.

The Management Server contains information about all elements to make sure that administrator actions are limited by the rights defined in the administrator account. Administrators can edit an element only if they are allowed to edit all configurations where the element is used. The Management Server also prevents administrators from deleting elements that are still used in some other configuration, from editing the same Policy element simultaneously, and from making conflicting changes to the same element.

How administrator accounts can be configured

- An unrestricted (superuser) administrator is created during the installation of an SMC Appliance.
- You can configure administrators in the SMC Client with these steps:
  1) Sets of administrator permissions are defined as reusable lists.
  2) Each list of permissions is applied to a specific group of elements.
  3) Define the administrator permissions.
     Several different pairs of permissions and elements can be applied to a single administrator account. These permissions can include, for example, viewing access to some elements and editing access to other elements. You can also create unrestricted accounts for “superusers” that have permissions for any action on any element. Some maintenance tasks require an unrestricted account.

Command-line administrator rights are available for engines and for the all-in-one SMC Appliance. To log on to the SMC Appliance command line, Administrators must have SMC Appliance superuser administrator permissions. Administrators with unrestricted permissions (superusers) are allowed to log on to the SMC Appliance command line only if there are no administrators with SMC Appliance Superuser permissions. All administrator accounts with SMC Appliance Superuser permissions are automatically replicated to the SMC Appliance and can execute root-level commands using the sudo tool.

In the SMC Client, administrator accounts can be configured to replicate to engines. If needed, administrator accounts can also be granted sudo permission to engines.

How administrator accounts work

Many elements are used in the configuration of administrator accounts, including Security Engine, Policy, Access Control List, Administrator Role, and Filter elements.

Two types of elements represent administrator accounts in the SMC:

- *Administrator* elements define accounts for administrators who are allowed to manage elements through the SMC Client or on the engine and view information in the Web Portal.
- *Web Portal User* elements define accounts for users who are allowed to view information in the Web Portal.

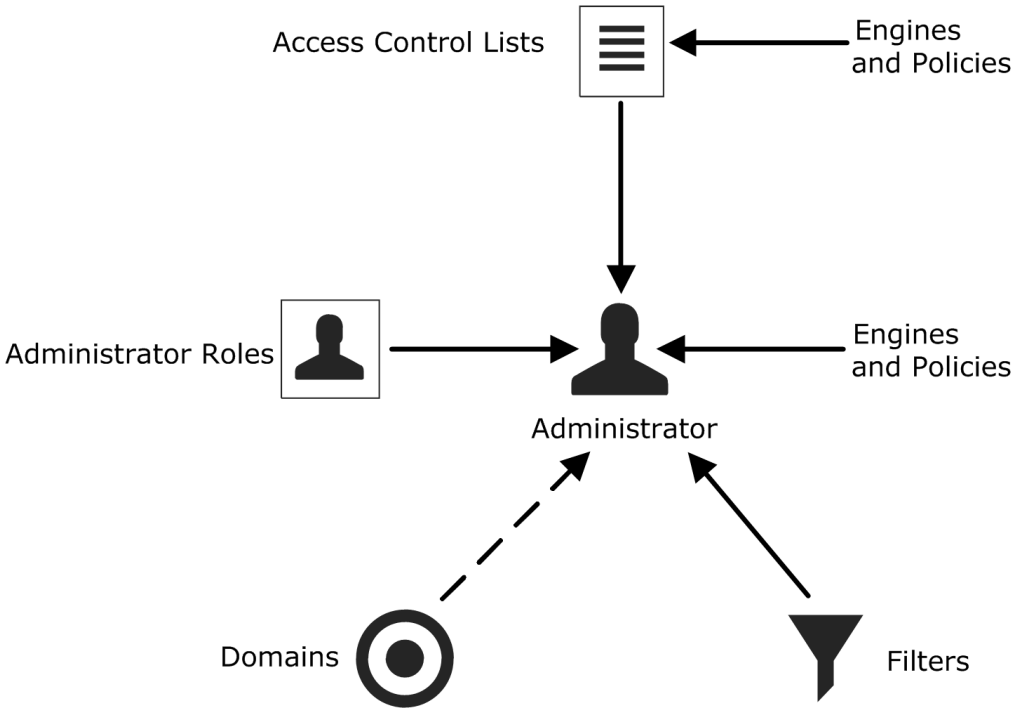


Note

Accounts that are used to log on to the SMC Client can also be used to log on to the Web Portal. Web Portal User accounts that can only be used to log on to the Web Portal are created separately.

The elements used with Administrator elements:

Elements for administrator account definitions



- *Administrator Roles* define sets of allowed actions.
- *Access Control Lists* contain elements and allow you to more easily apply the Administrator Roles to several engines and policies. There are some default Access Control List elements that are automatically populated and can represent additional element types.
- If an administrator is allowed to view logs, you can use *Filters* to select which logs are displayed to the administrator.
- If you use administrative Domains, you can give administrators access to any number of Domains.

Default administrator account elements

There are several predefined Administrator Roles and Access Control Lists that help you configure Administrator permissions. You cannot edit the predefined elements.

The following table describes the predefined Administrator Roles that you can optionally use instead of or in addition to customized Administrator Roles you create. All permissions listed here are always applied to a specific set of elements that you define.

Predefined administrator roles

Administrator role	Permissions given
Editor	Editors can: ■ View the properties of elements. ■ Send commands to engines, refresh policies, upload policies, and browse logs and alerts (if applied to components that send logs). ■ Create, edit, and delete elements.

Administrator role	Permissions given
Operator	Operators can: ■ View the properties of elements. ■ Send commands to engines, refresh policies, upload policies, and browse logs and alerts (if applied to components that send logs).
Owner	When an administrator creates an element, the administrator is automatically set as an owner of that element. Owners can: ■ View the properties of elements. ■ Create, edit, and delete elements.
Viewer	View the properties of elements.

All elements automatically belong to one or several predefined Access Control List elements in addition to the Access Control Lists you create yourself.

Predefined Access Control List elements

Access Control List	Description
All Elements	All elements that are defined in the system.
All Domains	All Domain elements in the system. Can be used with Administrator elements only if Domain elements have been configured.
All Administrators	All elements of the type mentioned in the name of the Access Control List.
All API Clients	
All Cloud Elements	
All Engine Policies	
All Engines	
All Inspection Policies	
All IPS Engines	
All IPS Policies	
All Layer 2 Engine Policies	
All Layer 2 Engines	
All Layer 2 Interface Policies	
All Third Party Devices	
All Web Portal Users	
All SSL VPN Gateways	Legacy SSL VPN Gateway elements.
All Simple Elements	All elements except elements that have a dedicated system Access Control List.

The contents of the Access Control Lists are Domain-specific if Domain elements have been configured in the system. For example, in the Shared Domain, **ALL IPS Policies** refers to all IPS Policies that belong to the Shared Domain.

Administrator account configuration overview

You must configure an administrator account for each administrator. You can create customized task and element lists that can be used to define permissions for administrators.

Follow these general steps to configure administrator accounts:

1) (Optional) Define customized reusable lists of allowed tasks for accounts with restricted permissions.

2) (Optional) Define customized reusable lists of elements for defining access rights for restricted accounts.

3) Create an administrator account for each administrator.

4) (Optional) Configure the password policy requirements for administrator passwords.

CAUTION

Do not use shared accounts. Using shared accounts makes auditing difficult and can make it difficult to discover security breaches.

Related tasks

Create Administrator Role elements on page 378
Create Access Control List elements on page 379
Add administrator accounts on page 379
Enable and define password policy settings on page 386

Creating Administrator Role and Access Control List elements

You can use Administrator Role and Access Control List elements in accounts that define restricted administrator permissions.

You can either use the predefined Administrator Roles and Access Control Lists or create custom ones.

Related tasks

Create Administrator Role elements on page 378
Create Access Control List elements on page 379

Create Administrator Role elements

Administrator Role elements specify a restricted set of permissions that include the right to create, edit, and delete elements.

Each administrator can have several different Administrator Roles applied to different sets of elements. There are some default Administrator Roles, but if you want to customize the permissions in any way, you must create custom Administrator Role elements. The Administrator Role contains a fixed list of permissions that you can activate.

Important

Select only the minimum necessary permissions for each role. Administrators who are allowed to edit administrator accounts can freely give themselves any permissions.

CAUTION

Changes made to an Administrator Role are applied immediately to every administrator account that uses the role (possibly including the account you are currently logged on with). Make sure that the permissions are correct before you apply changes to existing Administrator Roles.

If you change the permissions for existing administrator accounts, the administrators are notified that their permissions have changed the next time that they log on to the SMC Client.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.
2) Right-click **Access Rights** and select **New** > **Administrator Role** or right-click an existing Administrator Role to edit and select **Properties**.
3) (New Administrator Role only) In the **Name** field, enter a unique.
4) Select the permissions that are applied to the elements selected for the role.
5) Click **OK**.

Related concepts

Traffic captures and how they work on page 249

Related tasks

Create Access Control List elements on page 379
Add administrator accounts on page 379

Create Access Control List elements

An Access Control List defines a group of granted elements for which an administrator has rights.

If an Administrator Role gives the rights to install policies and browse logs and alerts, you must apply the Administrator Role to Security Engines in the Administrator element. The Access Control Lists that you create can include engines and policies.

The predefined Access Control Lists (in **Administration** > **Access Rights** > **Access Control Lists**) allow you to give access to all elements of a certain type. When you create an element, it is automatically added to the relevant default Access Control List. For example, a new Engine element is automatically included both in the **ALL Elements** and **ALL Engines** Access Control Lists.

Note

You must create custom Access Control Lists if you want to give access to a limited number of elements within one type.

If you change the permissions for existing administrator accounts, the administrators are notified that their permissions have changed the next time that they log on to the SMC Client.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.
2) Right-click **Access Rights** and select **New** > **Access Control List** or right-click an existing Access Control List and select **Properties**.
3) (New Access Control List only) In the **Name** field, enter a unique name.
4) Select the elements you want to add to the Access Control List from **Resources** and click **Add**.
   The selected elements are added to the **Granted Elements** list in the right pane.
5) Click **OK**.

Related tasks

Add administrator accounts on page 379
Define administrator permissions for Security Engines on page 648

Add administrator accounts

Administrator elements represent administrator accounts in the SMC. Administrators configure and monitor the SMC and the Security Engines.

An account with unrestricted permissions (superuser) is automatically created during installation to guarantee that a superuser account is available in the SMC. With this first account, you can create the necessary administrator accounts for daily management tasks. For the SMC Appliance, the account created during installation is also a user for the appliance.

The administrator accounts for the users of the optional Web Portal are defined with Web Portal User elements. All other administrator accounts are defined with Administrator elements.

There are several ways to authenticate administrator logons:

- You can authenticate administrators using a password stored in the internal database of the SMC.
- You can use a RADIUS or TACACS+ authentication method provided by an external authentication server.
- You can authenticate administrators using simple password authentication against integrated external LDAP databases.
- You can authenticate administrators by using a SAML based identity provider.
- You can authenticate administrators by using an OpenID provider.
- You can authenticate administrators using an X.509 certificate stored in the Windows certificate store or on a smart card, such as a Common Access Card (CAC).

Note

Certificate-based authentication is not supported for Web Portal Users.

Note

We highly recommend that you define a unique administrator account for each administrator. Using shared accounts makes auditing difficult and can make it difficult to discover security breaches.

There are two general permission levels for the administrators:

- *Unrestricted permissions* give the administrators the right to manage all elements without restriction, and the right to run scripts that require the administrators to authenticate themselves.
  Administrators with unrestricted permissions (superusers) can optionally also have SMC Appliance Superuser permissions that allow the administrators to log on to the SMC Appliance command line.
- *Restricted permissions* allow you to define the administrator's rights in detail using the Administrator Roles with individual elements and Access Control Lists.

If you change the permissions for existing administrator accounts, the administrators are notified that their permissions have changed the next time that they log on to the SMC Client.

If you use administrative Domains, there are some more considerations:

- You must create administrator accounts with unrestricted permissions in the Shared Domain.
- You must select Domains for each administrator role.
- Restricted accounts in the Shared Domain cannot access elements from any other Domains.
- Restricted accounts in other Domains can be granted elements that belong to the Shared Domain. However, the granted elements must belong to a Domain that is allowed for the administrator role selected for the account. For example, an administrator account in another Domain has the operator role in the Shared Domain. The administrator can be granted a policy template from the Shared Domain. The administrator can view the full contents of the policy.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.

2) Right-click **Access Rights** and select **New** > **Administrator**.

3) From the **Type** drop-down list, select where the administrator account is stored.

4) If you selected **Linked to LDAP**, select the user and user group in the integrated external LDAP directory to which the administrator account is linked.

   a) Click **Select** next to the **User** field, then select a User element.

   b) (Optional) Click **Select** next to the **Group** field, then select the User Group to which the User element must belong for SMC access to be allowed.

5) (When **Local** is selected) In the **Name** field, enter a unique name.

   The administrator uses this user name to log on to the SMC Client. For administrator accounts that are linked to user accounts in an integrated external LDAP directory, the name is filled in automatically.

6) To authenticate administrator logons using a user name and password on the Management Server, configure these options.

   a) From the **Authentication** drop-down list, select **User Password**.

   b) In the **Password** fields, enter and confirm the password, or click **Generate Password** to generate a temporary random password.

   Generated passwords are one-time passwords. The administrator is automatically prompted to enter a new password at the first logon.

   Note

   If you replicate administrator accounts as local accounts on engines, you must define a separate password for the local engine accounts.

   c) (Optional, manually entered passwords) To require the administrator to enter a new password at the first logon, select the **Require Administrator to Change Password at First Logon** checkbox.

7) To authenticate administrator logons using RADIUS or TACACS+ authentication method, from the **Authentication** drop-down list, select **RADIUS** or **TACACS+**.

8) To authenticate administrator logons using simple password authentication against an integrated external LDAP database, select **LDAP**.

9) To authenticate administrator logons by using an Identity Provider, from the **Authentication** drop-down list, select a **SAML** element.

   Note

   SAML username must contain the same username that is configured in the Identity Provider and must match with the current SMC username.

10) To authenticate administrator logons by using an OpenID Provider, from the **Authentication** drop-down list, select an **OpenID** element.

   Note

   OpenID username must contain the same username that is configured in the OpenID Provider and must match with the current SMC username.

11) To authenticate administrator logons using certificate-based authentication, from the **Authentication** drop-down list, select **Client Certificate**.

12) On the **Permissions** tab, define the administrator permissions.



CAUTION

Select only the minimum necessary permissions for each Administrator account.

13) For administrator accounts with restricted permissions, define the rights and granted elements.

a) Click **Add Role**.
A new Administrator Role appears in the list.

b) Click the **Role** cell and select the administrator role that defines the rights you want to set.

c) Right-click the **Granted Elements** cell for the role and select **Edit Granted Elements**.

d) Select the elements to which the rights granted by the administrator role apply.
The **Set to ALL** action depends on the type of elements. For example, if you browse to **Engines** and click **Set to ALL**, the item **All Engines** is added. You can also select one or more predefined or user-created Access Control Lists. **Simple elements** includes all elements except elements that have a dedicated system Access Control List. For example, there are dedicated Access Control Lists for different types of Security Engines and their policies.

e) (Optional) If Domain elements have been configured, click the **Domains** cell to select the Domains in which the rights granted by the administrator role and the selected elements apply.

f) (Optional) If Domain elements have been configured, leave **Allow Administrators to Log On to the Shared Domain** selected to allow the administrator to log on to the Shared Domain. Otherwise, the administrator is only allowed to log on to the specified Domains.

14) Click **OK**.

Related concepts

The Domain Overview on page 212
Creating Administrator Role and Access Control List elements on page 377
Getting started with Domain elements on page 433

Related tasks

Enable and define password policy settings on page 386
Authenticate administrators using RADIUS or TACACS+ methods on page 392

Enforce an approval workflow

You can optionally enable an approval workflow in which an administrator must approve changes before they are committed and transferred to the engines.

Administrators with the following permissions can view the changes, approve the changes, and transfer the configurations to the engines:

- Administrators that have the Approve Changes permission
- Administrators with unrestricted permissions (superusers)

By default, the same administrator who made the changes cannot approve the changes. You can optionally allow administrators to approve their own changes.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Settings** > **Global System Properties**.
2) On the **Change Management** tab, select **Require Approval for Changes in Engine Configuration**.
3) Click **OK**.

Related tasks

View, approve, and commit pending changes on page 107

Restrict the log data an administrator can view

If an administrator is allowed to view logs and alerts, you apply local filters to the log data before it is displayed to the administrator.

The filters that you create here are specific only to the Administrator element in question, unless you save them as permanent Filter elements.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.
2) Expand the **Access Rights** branch and click **Administrators**.
3) Right-click the Administrator and select **Properties**.
4) Click the **Permissions** tab.

5) Under **Log Filters**, click **Select**.

6) Define the Local Filter's properties.

7) Click **OK**.

Related concepts
Getting started with filtering data on page 333
Creating and editing local filters on page 337

Related tasks
Save local filters as permanent Filter elements on page 340
Add administrator accounts on page 379
Customize log colors for administrators on page 384

Customize log colors for administrators

By default, certain logs are shown with a colored background in the **Logs** view. Using administrator-specific log colors makes it easier to draw the administrator's attention to particular logs.

You can customize the default log colors used by default in all administrator accounts or define administrator-specific log colors in the Administrator element's properties. To use customized colors for logs, you must also create one or more filters that match those logs. Only administrators with the right to manage administrator accounts can customize log colors.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.

2) Expand the **Access Rights** branch and click **Administrators**.

3) Right-click the Administrator and select **Properties**.

4) In the **Administrator Properties** dialog box, click the **Color Filters** tab.

5) Select the log type for which you want to change color filters.

6) (Optional) To add a new filter for a log type, click **Add** and double-click the **Filter** cell in the new color filter row.

7) Select the color.
   - Double-click the **Color** cell of the filter for which you want to change the color and select a color from the palette.

- Click **More Colors** to select a custom color.

The selected colors are assigned to the filters and they are used whenever logs match the filter.

8) Click **OK**.

Related concepts
Creating and editing local filters on page 337

Related tasks
Disable administrator accounts on page 405
Monitor administrator actions on page 248

Replicate administrator accounts

You can replicate SMC administrator accounts as local administrator accounts on selected engines. This enables several administrators to access an engine locally with the security privileges of the root user.

Before you begin

Before replication, each administrator must have an existing SMC administrator account. However, they must not have existing accounts on the engine.

Several administrators might need to access a single engine for troubleshooting or for configuring features that are not yet available through the SMC Client. It is a good security practice to create each of them a separate account with a personal password and permissions. This practice enables more granular and accurate auditing as well.

The root administrator can limit and configure the engine administrators' permissions individually in the local engine *sudo* security policy. When an administrator is allowed to use sudo commands to execute root-level commands on the engine, by default all commands are allowed on the engine. You can limit the commands allowed for an administrator by editing the configuration for the sudo package. Engine configuration files for sudo are in the /data/config/sudoers.d/ directory on the engine.



Note

Administrator Permissions and Roles or other configurations done in the SMC Client are not replicated on the engine.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.

2) Expand the **Access Rights** branch and click **Administrators**.

3) Right-click the Administrator and select **Properties**.

4) Click the **Account Replication** tab.

5) Select **Replicate Account on Selected Engines**.

6) (Optional) To allow the use of sudo commands to execute root-level commands, select **Allow executing root-level commands with the sudo tool**.

7) In the **Password** field, enter the password and confirm it in the **Confirm** field. You can also click **Generate Password** to generate a random 7-digit alphanumeric password.

8) To select the Security Engines elements where the accounts are replicated to, click **Add**.

9) Select **Access Control Lists**, **Domains**, or **Security Engines** and then select the element by clicking **Select**.

10) Click **OK**.

Result

The administrator account is replicated on the engines if the engines are online and have a connection to the Management Server.

Enable and define password policy settings

If you authenticate administrators or Web Portal users with internal authentication, you can enforce a password policy.

Before you begin

You must be logged on using an administrator account with sufficient permissions to change the password policy settings. Permissions to manage Administrator elements or unrestricted permissions (superuser) are required. If administrative Domains are configured, you must be logged on to the Shared Domain.

The settings in the password policy are applied to:

- Administrator and Web Portal user accounts defined using Administrator and Web Portal User elements.
- SMC administrator accounts that are replicated as local administrator accounts on Security Engines.
- The root account on Security Engines.
- The Management Server database password.
  For the Management Server database password, only requirements for length, uppercase characters, lowercase characters, and numbers are applied. Special characters are not allowed in the Management Server database password.

You can define the following settings in the password policy:

- Session limits and idle timeouts
- Restrictions on failed logon attempts
- Automatic disabling of inactive accounts
- Requirements for password age and expiration
- Requirements for password strength

Note

If you have previously changed the default password policy settings in the SGConfiguration.txt file, the settings are automatically applied on the **Password Policy** tab. Any further modifications you make to the SGConfiguration.txt file have no effect.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Settings** > **Global System Properties**.

2) Click the **Password Policy** tab.

3) Select **Enforce Password Settings for All the Administrators and Web Portal Users**.

Note

The password policy is enforced by default.

4) Select the password policy settings.

   For information about the options that you must select in a Common Criteria certification environment, see the *Common Criteria Certification User's Guide*.

5) Click **OK**.

Change administrator passwords

If you have not configured administrator passwords to automatically expire, we recommend that you change administrator passwords regularly.

Related tasks

Enable and define password policy settings on page 386

Change passwords of other administrators

An administrator who has the right to manage administrator accounts can change any other administrator's password.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.
2) Select **Access Rights** > **Administrators** or **Access Rights** > **Web Portal Users**.
3) Right-click the Administrator or Web Portal User element and select **Properties**.
4) In the **Password** field, enter and confirm the password.
5) Click **OK**.

Change your own administrator password

All administrators can change their own passwords in the SMC Client or the Web Portal.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Settings** > **Password** > **Change Password**.
2) In the **Old Password** field, enter your current password.
3) In the **New Password** field, enter a new password.
4) In the **Confirm New Password** field, confirm the new password.
5) Click **OK**.

Change your own local Security Engine password

If administrator accounts have been replicated on engines, administrators can change their own local engine passwords using the SMC Client.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the SMC Client, select **Settings** > **Password** > **Change Engine Password**.
2) In the **New Password** field, enter a new password.
3) In the **Confirm New Password** field, confirm the new password.
4) Click **OK**.

Related tasks

Change the Security Engine root password on page 359

Authenticate administrators using OpenID authentication method

You can authenticate administrators by using the OpenID connect authentication method to facilitate single sign-on to the SMC Web Access portal. An OpenID Provider is used to authenticate administrators to grant access to the SMC Web Access portal.

Before you begin

You must have the following:

1) An OpenID Provider is configured. Please contact your OpenID Provider support team for details.
2) OpenID Discovery URL.
3) The Client ID.
4) The Client Secrets.

Note

1) The OpenID connect authentication method can only be used with the SMC Web Access portal. Also, this authentication method can only be configured for an administrator.
2) You must restart the SMC Web Access portal for the changes made to come into effect.

Steps

1) Create a SAML authentication method element:
   a) Browse to **User Authentication** > **Authentication Methods**.
   b) Right-Click and select **New Authentication Method**.
   c) Enter a unique name for the authentication method element in the **Name** field.
   d) From the **Type** drop-down list, select **OpenID**.

   Note

   The fields below the **Type** drop-down list changes as per the options selected from the **Type** drop-down list.

   e) Enter the URL from where SMC will fetch the details about the OpenID connect authorization server in the **OpenID Discovery URL** field.
   f) Enter the public identifier for SMC Web Access application in the **Client ID** field.
   g) Enter the client secret that is used by SMC Web Access to authenticate itself in the **Client Secret** field.
   h) Enter the ID token that is claimed to be used as the username in the **Username Attribute Name** field.
   i) Select a certificate authority that is used to connect to the OpenID server. Click **Select** to select the element.
   j) Optionally, add a comment in the **Comment** field for your future reference.
   k) Click **OK** to save the changes.

2) Configure the OpenID authentication in the properties of the administrator:
   a) Browse to **Administration** > **Access Rights** > **Administrators**.
   b) Right-click an Administrator element, then select **Properties**.
   c) From the **Authentication** drop-down list, select the OpenID authentication element.
   d) Click **OK**.

3) Configure the SMC Web Access. For more details on how to enable or configure the SMC Web Access, refer to the following sections in the *Forcepoint Network Security Platform Online Help* documentation:
   - Enable SMC Web Access
   - Management Server Properties dialog box

Authenticate administrators using RADIUS or TACACS+ methods

You can authenticate administrators and Web Portal users using RADIUS or TACACS+ authentication methods.

Before you begin

You must have an external authentication server that provides RADIUS or TACACS+ authentication methods.

The Management Server's internal user database does not allow external authentication servers to query the administrator account information. To use external authentication, you must manually create an account both in the SMC for defining the permissions and in the external directory for logon authentication. The administrator's user name for the Management Server and for the directory that the external authentication server uses must match exactly.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Add one of the following types of server elements to integrate the external server, then define the shared secret used in the communications in the server element.
   - Add a RADIUS Authentication Server element, then add a RADIUS Authentication Method.
   - Add a TACACS+ Authentication Server element, then add a TACACS+ Authentication Method.
   - Add an Active Directory Server element, then add a RADIUS Authentication Method.

   Note

   To use a RADIUS or TACACS+ Authentication Server that has an IPv6 address, the Management Server must also have an IPv6 address.

2) Add an Access rule that allows traffic from your Management Server to the external authentication server.

3) Select **Network Elements**.

4) Browse to **Servers**.

5) Right-click the Management Server, then select **Properties**.

6) From the **RADIUS Method** or **TACACS+ Method** drop-down list, select the authentication protocol for authenticating the Management Server's communications with the external authentication server.

   CAUTION

   To guarantee the security of the SMC, communications between the Management Server and the external authentication server must remain confidential. We recommend transferring these connections over secure networks only.

7) (RADIUS Authentication Servers only) Set up the external server for use with the Management Server.
   a) Define the Management Server as a RADIUS client on your server.
   b) Define the same authentication method on your server as you selected in the Management Server properties in the previous step.

8) In the SMC Client, configure RADIUS or TACACS+ authentication in the properties of each Administrator or Web Portal User account.
   a) Select **Administration**.
   b) Select **Access Rights** > **Administrators**.
   c) Right-click an Administrator element, then select **Properties**.
   d) From the **Authentication** drop-down list, select **RADIUS** or **TACACS+**.
   e) From the **Authentication Method** drop-down list, select an Authentication Method element, or click **Select** to select a different Authentication Method element.
   f) Click **OK**.

Related tasks

Add administrator accounts on page 379
Create Active Directory Server elements on page 1108
Create RADIUS or TACACS+ Authentication Server elements on page 1134
Define Authentication Method elements for external servers on page 1135

Authenticate administrators on engines using the Radius authentication method

You can authenticate administrators on an engine by using the RADIUS authentication method.

Before you begin

- This feature is not supported for virtual engines.
- This feature is only supported on engine versions 7.1.3 and later.
- You must replicate administrator accounts on the engine to authenticate administrators on that engine using the RADIUS authentication method. For more details on replicating administrator accounts, refer to the *Replicate administrator accounts* topic.
- If a user disables the root account and downgrades the engine to a version that does not support radius authentication for engine, then the root account will remain unavailable even if the password is reset from SMC. Hence, it is recommended that the root account is enabled before downgrading the engine version.

Steps

1) Configure the RADIUS authentication settings for the administrator. For more information, refer to the **Authenticate administrators using RADIUS or TACACS+ methods** topic.

2) Select **Engine Configuration**.

3) Right-click an engine, then select **Edit <element type>**.

4) Navigate to **Advanced Settings** > **Authentication**.

5) In the **Root and Administrator Authentication** section:

   a) From the **Root Password Login** drop-down list, select one of the following options:

   - **Login Allowed via SSH and Console:** The root password login to an engine is allowed via SSH and console.

     Note

     By default, this option is selected if the engine is upgraded.

   - **Login Allowed via Console Only:** The root password login to an engine by using SSH is not allowed. But root password login by using console is allowed.

     Note

     By default, this option is selected when we create a new engine.

   - **Root Account Disabled (Super User Privileges through sudo):** The root password login to an engine is disabled.

   b) From the **Authentication Method** drop-down list, select an authentication method element from the below options:

   - **Local Password:** Allows authentication using the local password.
   - **[Select…]:** Select this option to view the available radius authentication method elements.

     Note

     The authentication method options are displayed as per the radius authentication server elements that are configured. For more details on how to create a radius authentication server element, refer to the **Define Authentication Method elements for external servers** topic.

   c) From the **SSH Passwordless Login** drop-down list, select one of the following options:

   - **Allow:** The SSH password less login is allowed.
   - **Deny:** The SSH password less login is denied.

     Note

     This applies only to administrators replicated on the engine. For more details on administrator account replication, refer to the **Add administrator accounts** topic.

6) Click the **Save and Refresh** icon.

Authenticate administrators using SAML v2 authentication method

You can authenticate administrators by using a Security Assertion Markup Language (SAML) based identity providers to facilitate Single Sign-on (SSO) to the SMC Web Access portal.

Before you begin

You must have the following:

1) A SAML IdP configured. Please contact your SAML IdP support team for details.
2) The identity provider metadata URL.
3) Service Provider Entity ID.

The authentication is done by transferring identity data between two parties, that is an Identity Provider (IdP) and a Service Provider (SP).

Identity Provider: It performs the authentication and passes the identity data of the administrator and authorization level to the service provider.

Service Provider: It trusts the identity provider and in turn authorizes the user to access the requested resource.

Note

1) The SAML authentication method can only be used with the SMC Web Access portal. Also, this authentication methods can only be configured for an administrator.
2) You must restart the SMC Web Access portal for the changes made to come into effect.

Steps

1) Create a SAML authentication method element:
   a) Browse to **User Authentication** > **Authentication Methods**.
   b) Right-click and select **New Authentication Method**.
   c) Enter a unique name for the authentication method element in the **Name** field.
   d) From the **Type** drop-down list, select **SAML**.

   Note

   The fields below the **Type** drop-down list changes as per the options selected from the **Type** drop-down list.

   e) Enter the URL from where SMC will fetch the details about the SAML configuration in the **Identity Provider Metadata URL** field.
   f) Enter the unique identifier for SMC (Service Provider) in the **Service Provider Entity ID** field.
   g) From the **Name ID Policy Format** drop-down list, select a policy format. The following policy formats are supported:
      - **Persistent:** Use this policy format if you want a user to sign-in to the identity provider as one user, but sign-in to the service provider as a different user.

        Note

        Before you can use this policy format, you must link the user at the identity provider with the user at the service provider. Also, you can choose to have the user linked during the single sign-on or by using the alias service.

      - **Transient:** Use this policy format if you want a user to sign-in as a shared anonymous user irrespective of which user they use to sign-in at the identity provider.
      - **Email Address:** Use this policy format if you want a user to sign-in at the service provider as the same user that they use to sign-in at the identity provider.
      - **Unspecified:** Use this option if you do not want to specify a policy format.
   h) Enter the name of SAML2 attribute that defines the username in the **Username Attribute Name** field.
   i) Select a TLS profile to use to connect to an Identity Provider. Click **Select** to select the element.
   j) Select the TLS credentials to use to sign in SAML requests, and decryptSAML responses. Click **Select** to select the element.
   k) Optionally, add a comment in the **Comment** field for your future reference.
   l) Click **OK** to save the changes.

2) Configure the SAML authentication in the properties of the administrator:
   a) Browse to **Administration** > **Access Rights** > **Administrators**.
   b) Right-click an Administrator element, then select **Properties**.
   c) From the **Authentication** drop-down list, select the SAML authentication element.
   d) Click **OK**.

3) Configure the SMC Web Access. For more details on how to enable or configure the SMC Web Access, refer to the following sections in the *Forcepoint Network Security Platform Online Help* documentation:
   - Enable SMC Web Access
   - Management Server Properties dialog box

Using LDAP authentication for administrators

When you use LDAP authentication for administrators, administrator accounts are linked to user accounts in an integrated external directory server. The external directory server where the user accounts are stored verifies the user credentials.

Note

To use LDAP authentication for administrators, you must have an integrated external directory server where the administrator accounts are stored.

When administrators authenticate to the Management Server, the Management Server sends the user name and password to the external directory server for authentication. The external directory server checks the user name and password against the user's credentials in the directory. If a user group is defined for the administrator, the external directory server also checks whether the linked user account is still a member of the specified group. The external directory server responds to the Management Server whether authentication succeeds or fails.

Note

Because the user name and password are sent through the LDAP connection, we recommend using LDAPS or Start TLS when you use LDAP Authentication.

Authenticate administrators using certificate-based authentication

You can authenticate administrators using an X.509 certificate stored in the Windows certificate store, on a smart card like a Common Access Card (CAC), or via SMC Web Access.

One of the following ways can be used to authenticate administrators using certificate-based authentication to log on to SMC:

- **Certificate-based authentication for Installed Clients:** The smart card reader or certificate files can be used to authenticate administrators. For more details:
  - On how to log on to SMC using certificates, refer to the *Log on to the SMC using certificate-based authentication* topic.
  - On how to configure certificate-based authentication for installed clients, refer to the *Configuring certificate-based authentication for installed clients* topic.

  Note

  - If the smart card reader is used to authenticate administrators, you must have the smart card reader and corresponding software installed.
  - If the certificate files are used to authenticate administrators, you must save the certificates in the Windows certificate store.

- **Certificate based authentication for SMC Web Access:** Administrators can be authenticated by using certificates in the browser when using SMC Web Access. For more details on how to configure certificate-based authentication for SMC Web Access, refer to the *Configuring certificate-based authentication for SMC Web Access* topic.

  Note

  You must save the certificates in the Windows certificate store.

- **Certificate based authentication for both Installed Client and SMC Web Access:** You can enable both the certificate-based authentication for Installed Client and the certificate-based authentication for SMC Web Access at the same time to authenticate administrators. For more details on how to configure certificate-based authentication for SMC Web Access, refer to the *Configuring certificate-based authentication for both Installed Client and SMC Web Access* topic.

A client certificate in the Windows certificate store is used for client authentication. There is also a trusted certificate authority (CA) for the client certificate in the Windows certificate store. There are two ways to store the private key for the client certificate:

- The private key can be stored on a smart card, from which the client certificate can be populated to the Windows certificate store.
- A Windows software provider can be used for key storage.

Note

Certificate-based authentication is only supported for SMC Clients installed in Windows 10. Certificate-based authentication is not supported for Web Portal Users.

Configuring certificate-based authentication for installed clients

You can use smart card readers or certificate files to authenticate administrators.

Steps

1) Create a TLS Credentials element to define the certificate that is used to authenticate the Management Server in communications for certificate-based authentication. For more details, refer to the *Configure TLS inspection for server protection* topic.



Note

- You can generate and sign a new certificate request or import an existing certificate.
- The certificate defined in the TLS Credentials element is used for server authentication. The SMC Client validates the server certificate path using the trusted CA certificates in the Windows certificate store.
- If SMC is configured to use externally signed certificates for internal management communication, the **TLS Credentials** field can be left empty. In this case, the same management server certificate that is used for other management communications is used in SMC Client communication for client certificate authentication.

2) Create a TLS profile element to define the trusted CAs for the Management Server and the client certificates. For more details, refer to the *Create TLS Profile elements* topic.



Note

Make sure that the TLS Profile element includes the trusted CAs for both the Management Server's certificate and for the client certificates. The trusted CA can be the same for the certificate of the Management Server and for the client certificates.

3) Configure the Management Server for certificate-based authentication.
   a) Select **Network Elements**.
   b) Browser to **Servers**.
   c) Right-click the Management Server, then select **Properties**.
   d) From the **General** tab, next to the **TLS Credentials** field, click **Select**, then select a TLS Credentials element.
   e) Next to the **TLS Profile** field, click **Select**, then select a TLS Profile element.
   f) Click **OK**.

4) In the properties of each Administrator, configure certificate-based authentication.
   a) Select **Administration**.

b) Select **Access Rights** > **Administrators**.

c) Right-click an Administrator element, then select **Properties**.

d) From the **Authentication** drop-down list, select **Client Certificate**.

e) From the **Client Identity Type** drop-down list, select the certificate attribute that is used to identify the administrator.

f) Specify the value of the certificate attribute in one of the following ways:

- In the **Identity Value** field, enter the value of the certificate attribute.
- Click **Fetch From Certificate**, then import the certificate to get the value from the certificate.

g) Click **OK**.

5) If the certificate for the Management Server was not signed using a CA that is already trusted by the administrators' client operating systems, add the CA that signed the certificate as a trusted CA on each administrator's computer.

a) Export the CA certificate from the CA that signed the certificate for the Management Server.

b) Import the CA certificate on each administrator's computer.

c) Configure the operating system to trust the CA certificate.

Related concepts

Creating certificates on page 158

Related tasks

Add administrator accounts on page 379
Log on to the SMC using certificate-based authentication on page 101

Configuring certificate-based authentication for SMC Web Access

You can use certificates in the browser to authenticate administrators via SMC Web Access.

Steps

1) Create a TLS Credentials element to define the certificate that is used to authenticate the Management Server in communications for certificate-based authentication. For more details, refer to the *Configure TLS inspection for server protection* topic.



Note

- You can generate and sign a new certificate request or import an existing certificate.
- The certificate defined in the TLS Credentials element is used for server authentication. The SMC Client validates the server certificate path using the trusted CA certificates in the Windows certificate store.
- If SMC is configured to use externally signed certificates for internal management communication, the **TLS Credentials** field can be left empty. In this case, the same management server certificate that is used for other management communications is used in SMC Client communication for client certificate authentication.

2) Create a TLS profile element to define the trusted CAs for the Management Server and the client certificates. For more details, refer to the *Create TLS Profile elements* topic.



Note

Make sure that the TLS Profile element includes the trusted CAs for both the Management Server's certificate and for the client certificates. The trusted CA can be the same for the certificate of the Management Server and for the client certificates.

3) Configure the Management Server for certificate-based authentication.

a) Select **Network Elements**.

b) Browse to **Servers**.

c) Right-click the Management Server, then select **Properties**.

d) From the **General** tab, next to the **TLS Profile** field, click **Select**, then select a TLS Profile element.

e) From the SMC Web Access tab, select the **Client Certificate Authentication** checkbox.



Note

You must select the **Client Certificate** option as the authentication method in the administrator properties.

f) Click **OK**.

g) Restart the SMC Web Access:

i) Select **Dashboards** > **Servers / Devices Dashboard**.

ii) Right-click the management server, and then select the **More actions** > **Restart Web Access** option.

4) In the properties of each Administrator, configure certificate-based authentication.
   a) Select **Administration**.
   b) Select **Access Rights** > **Administrators**.
   c) Right-click an Administrator element, then select **Properties**.
   d) From the **Authentication** drop-down list, select **Client Certificate**.
   e) From the **Client Identity Type** drop-down list, select the certificate attribute that is used to identify the administrator.
   f) Specify the value of the certificate attribute in one of the following ways:
      - In the **Identity Value** field, enter the value of the certificate attribute.
      - Click **Fetch From Certificate**, then import the certificate to get the value from the certificate.
   g) Click **OK**.
5) If the certificate for the Management Server was not signed using a CA that is already trusted by the administrators' client operating systems, add the CA that signed the certificate as a trusted CA on each administrator's computer.
   a) Export the CA certificate from the CA that signed the certificate for the Management Server.
   b) Import the CA certificate on each administrator's computer.
   c) Configure the operating system to trust the CA certificate.

Configuring certificate-based authentication for both Installed Clients and SMC Web Access

You can use certificates stored in the Windows certificate store, on smart cards, or via SMC Web Access to authenticate administrators.

Steps

1) Create a TLS Credentials element to define the certificate that is used to authenticate the Management Server in communications for certificate-based authentication. For more details, refer to the *Configure TLS inspection for server protection* topic.



Note

- You can generate and sign a new certificate request or import an existing certificate.
- The certificate defined in the TLS Credentials element is used for server authentication. The SMC Client validates the server certificate path using the trusted CA certificates in the Windows certificate store.
- If SMC is configured to use externally signed certificates for internal management communication, the **TLS Credentials** field can be left empty. In this case, the same management server certificate that is used for other management communications is used in SMC Client communication for client certificate authentication.

2) Create a TLS profile element to define the trusted CAs for the Management Server and the client certificates. For more details, refer to the *Create TLS Profile elements* topic.



Note

Make sure that the TLS Profile element includes the trusted CAs for both the Management Server's certificate and for the client certificates. The trusted CA can be the same for the certificate of the Management Server and for the client certificates.

3) Configure the Management Server for certificate-based authentication.

a) Select **Network Elements**.

b) Browse to **Servers**.

c) Right-click the Management Server, then select **Properties**.

d) From the **General** tab, next to the **TLS Credentials** field, click **Select**, then select a TLS Credentials element.

e) Next to the **TLS Profile** field, click **Select**, then select a TLS Profile element.

f) From the SMC Web Access tab, select the **Client Certificate Authentication** checkbox.



Note

You must select the **Client Certificate** option as the authentication method in the administrator properties.

g) Click **OK**.

4) In the properties of each Administrator, configure certificate-based authentication.

a) Select **Administration**.

b) Select **Access Rights** > **Administrators**.

c) Right-click an Administrator element, then select **Properties**.

d) From the **Authentication** drop-down list, select **Client Certificate**.

e) From the **Client Identity Type** drop-down list, select the certificate attribute that is used to identify the administrator.

f) Specify the value of the certificate attribute in one of the following ways:

- In the **Identity Value** field, enter the value of the certificate attribute.
- Click **Fetch From Certificate**, then import the certificate to get the value from the certificate.

g) Click **OK**.

5) If the certificate for the Management Server was not signed using a CA that is already trusted by the administrators' client operating systems, add the CA that signed the certificate as a trusted CA on each administrator's computer.

a) Export the CA certificate from the CA that signed the certificate for the Management Server.

b) Import the CA certificate on each administrator's computer.

c) Configure the operating system to trust the CA certificate.

Disable administrator accounts

If an administrator account is no longer needed, you can disable the administrator account to remove access for the administrator.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.

2) Expand the **Access Rights** branch and click **Administrators**.

3) Right-click the Administrator and select **Disable Administrator**.
A **Confirmation** dialog box opens.

4) Click **Yes** to confirm that you want to disable the Administrator.
The Administrator element is marked as **Obsolete** and all scheduled Tasks created by the Administrator are ignored.

Related tasks
Add administrator accounts on page 379
Delete administrator accounts on page 406

Delete administrator accounts

When you permanently delete an administrator account, all history of changes the administrator has made to elements is also deleted permanently.

Before you begin

You must disable the administrator account before you can delete it.

A history of the changes that an administrator makes to elements is saved in the SMC. If you delete the administrator account, all history information about the changes the administrator has made is lost. Audit entries that reference the administrator are preserved.

Note

There must be at least one account with unrestricted permissions (superuser) in the SMC. It is not possible to delete the last remaining unrestricted account.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.

2) Expand the **Access Rights** branch and click **Administrators**.

3) Right-click the Administrator element, then select **Delete**.
A **Confirmation** dialog box opens.

CAUTION

Deletion is permanent. There is no undo. To recover a deleted administrator account, you must either recreate it or restore it from a previously created backup that contains the Administrator element.

4) Click **Yes**.
The Administrator element is permanently deleted.

Related tasks
Disable administrator accounts on page 405
Restore elements from the Trash on page 199

API client accounts and how they work

You can use the application programming interface (API) of the SMC to run certain actions in the SMC remotely using an external application or script.

You must grant all API clients access to the SMC in the SMC Client. You can define the API clients and their permissions using API Client elements. For more information on how to configure API Client elements in the SMC Client, see the *Forcepoint Security Management Center API User Guide*.

Configure SMC API

The Application Programming Interface (API) of SMC allows external applications to connect with the SMC.

Note

If there is a engine between SMC and the other applications, make sure that there is an Access rule to allow communication.

The SMC API can be used to run actions remotely using an external application or script. For more information about using SMC API, see the *Forcepoint Security Management Center API User Guide*.

Create TLS credentials for SMC API Clients

If you want to use encrypted connections, the SMC API Client needs TLS credentials to connect with the Management Server.

Note

You can import the existing private key and certificate if they are available.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the SMC Client, select **Administration**.
2) Browse to **Certificates** > **TLS Credentials**.
3) Right-click **TLS Credentials**, then select **New TLS Credentials**.
4) Complete the certificate request details.
   a) In the **Name** field, enter the IP address or domain name of SMC.
   b) Complete the remaining fields as needed.
   c) Click **Next**.

5) Select **Self Sign**.

6) Click **Finish**.

Result

The TLS Credentials element is added to **Administration** > **Certificates** > **TLS Credentials**. The **State** column shows that the certificate has been signed.

Enable SMC API

To allow other applications to connect using the SMC API, enable SMC API on the Management Server.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the SMC Client, select **Dashboard** > **Servers / Devices Dashboard**.

2) Browse to **Management Server**.

3) Right-click the Management Server, then select **Properties**.

4) Click the **SMC Client API** tab, then select **Enable**.

5) (Optional) In the **Host Name** field, enter the name that the SMC API service uses.



Note

API requests are served only if the API request is made to this host name. To allow API requests to any host name, leave this field blank.

6) Make sure that the listening port is set to the default of 8082 on the Management Server.

7) If the Management Server has several IP addresses and you want to restrict access to one, enter the IP address in the **Listen Only on Address** field.

8) Select the TLS Credentials element that is used for HTTPS connections. Click the **Select** button against the **Server Credentials** field, then select an element.

9) If you want to use encrypted connections, click the **Select** button against the **Server TLS Cryptography Suite Set** field, then select the TLS Credentials element and the Cryptography Suite Set element.

10) Click **OK**.

Create an API Client element

External applications use API clients to connect to SMC.

Before you begin

SMC API must be enabled for the Management Server.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.
2) Browse to **Access Rights**.
3) Right-click **Access Rights** and select **New** > **API Client**.
4) In the **Name** field, enter a unique name for the API Client.
5) Use the initial authentication key or click **Generate Authentication Key** to generate a new one.

   Important

   This key appears only once, so be sure to record it. The API Client uses the authentication key to log on to SMC API.

6) Click the **Permissions** tab.
7) Select the permissions for actions in the SMC API. As a minimum, set the **Viewer** permission to **All Simple Elements**.
8) Click **OK**.

Chapter 21

Alert escalation

Contents

- Alert escalation and how it works on page 411
- Creating Alert elements on page 414
- Configure notifications for alerts on page 416
- Alert Chain elements and how they work on page 417
- Creating Alert Policy elements on page 419
- Install Alert Policy elements on page 420
- Acknowledge active alerts on page 421
- How custom scripts for alert escalation work on page 422
- Create SMTP Server elements on page 423
- Use a script for SMS notification on page 424
- SNMP for the SMC Appliance on page 425
- Test alerts on page 428
- Examples of alert escalation on page 429

The SMC can escalate the alerts generated so that notifications are sent to the administrators through multiple channels.

Alert escalation and how it works

Alerts notify you if something unexpected or suspicious happens. It is important for administrators to respond to alerts to maintain the health of the SMC.

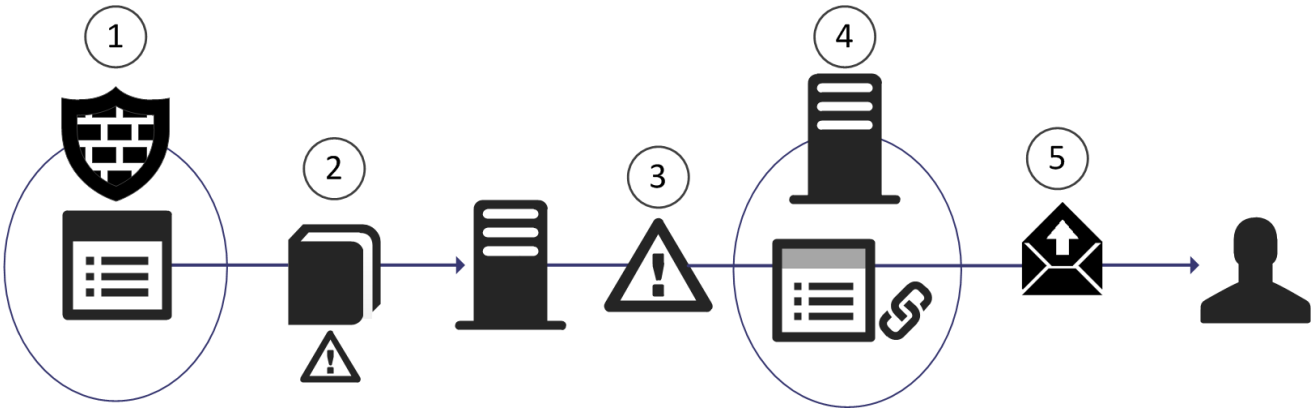
Alert entries inform the administrators when an event in the system requires their attention. For example, alerts are sent when there is a problem with the system, when a test or task fails, or when a rule that is configured to trigger an alert matches. Alerts can also be sent when a threshold for a user alert check is exceeded.

Active alerts are stored on the Management Server until the alerts are acknowledged. In an environment with multiple Management Servers, each active alert is stored on each Management Server. Alert entries are displayed in the **Active Alerts** view and in the **Logs** view with other types of log entries.

If you have configured the SMC Client to show users in the Dashboard view, you can see a summary of user alerts. Select a user to see the user alerts that the user has generated.

The Management Server can send out different types of notifications to administrators. Alert escalation stops when one of the administrators acknowledges the alert or when all configured alert notifications have been sent. When an alert entry is acknowledged, it is removed from the **Active Alerts** view and from the Management Server, and an audit entry is created.

Alert escalation



1. An event on a system component triggers an alert entry.
2. The alert entry is sent to the Log Server, which stores it.
3. The Log Server forwards the alert entry to the Management Server, where it is handled as an active alert.
4. The Management Server matches the alert entry to the *Alert Policy* to select the correct *Alert Chain*.
5. The Alert Chain triggers a series of notifications that are sent to administrators.

For example, an Alert Chain can first notify one of the administrators by email and wait for acknowledgment for 10 minutes. If the alert is not acknowledged in time, the Management Server can send another notification as an SMS text message.

Limitations

- The SMC does not support authentication or TLS encryption for SMTP.
- Only one email recipient can be configured for each notification. To send an email to several people at the same time, you must configure an email group on the mail server or configure several notifications consecutively without delays.
- Only SMC servers can send Test Alerts. Test Alerts always have default Severity and Situation information.
- By default, the maximum number of active alerts is 3000 per Domain. You can change the default number of active alerts per Domain by adjusting the `MAX_ACTIVE_ALERTS` and `CRITICAL_ACTIVE_ALERTS_EXTRA_SPACE` parameters in the `SGConfiguration.txt` file that is stored on the Management Server. The maximum number of alerts of any Severity is 2000. After 2000 alerts of any Severity have been sent, only Critical alerts are still sent until the total number of active alerts is 3000.

Default alert escalation elements

The alert system includes some predefined elements for configuring alert escalation.

The following default alert escalation elements are included:

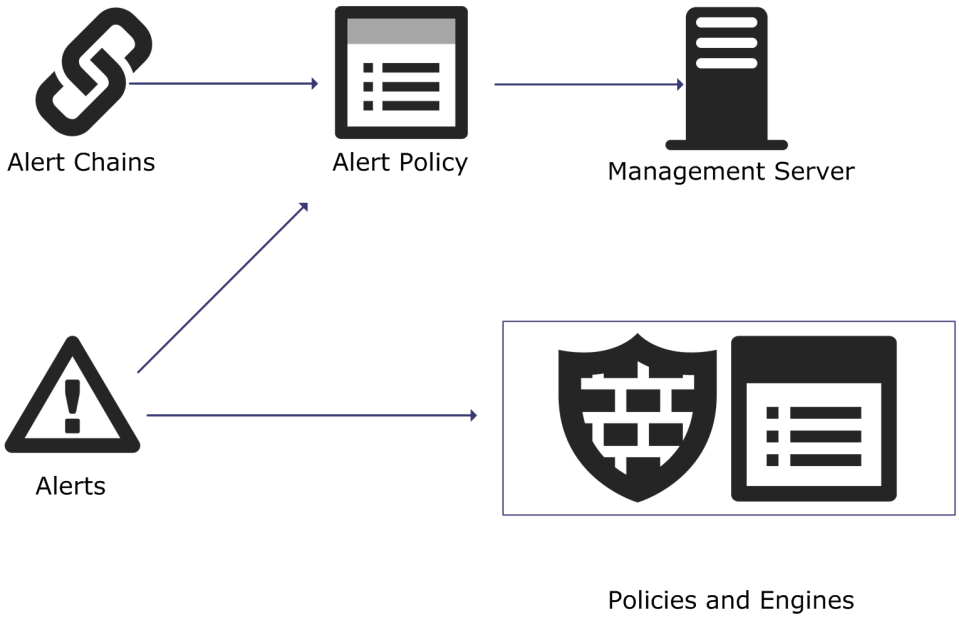
- **System Situations** — System Situations contain definitions for events in the system that trigger a System Alert. There are no configurable parameters for System Situations and you cannot adjust when these Situations are triggered.

- **System Alert** — This Alert element is used for alerts triggered by critical events in the system's internal operation (System Situations). System Alerts always require administrator action. Make sure that your Alert Policies escalate System Alerts.
- **Default Alert** — This Alert element defines the alert that is triggered if no other alert is specified. The Default Alert is used in the default Inspection Policy. You can also use it in your own custom configurations.
- **Test Alert** — This Alert element is used when you test alert handing.
- **Default Alert Chain** — This Alert Chain escalates all alerts to all administrators through user notification in the SMC Client.
- **Default Alert Policy** — This Alert Policy contains a rule that escalates all alerts using the Default Alert Chain.

Alert escalation configuration overview

Configuring how alerts are escalated involves several general steps.

Elements in the configuration



The configuration consists of the following general steps:

1) (Optional) Create Custom Alert elements for more precise matching in the Alert Policy.

2) Define which events trigger alerts.

3) Define how alert notifications are sent to administrators.

4) Create lists of actions to take when an alert is escalated.

5) Define which alerts are matched against which Alert Chain.

Information included in alert notifications

The amount of information the alert notification includes depends on the Alert Notification used.

Contents of escalated alerts

Alert notification	Information included in the notification
Custom script	Depends on the script you create.
Email	Includes the full details of the alert, the full situation description, and the contents of all hex viewable fields as a hexadecimal dump and ASCII.
SMS (text message)	SMS messages sent using an external script are limited to one line in length, with no character limit for that one line. SMS messages sent by SMTP and HTTP do not have this limit. The standard character limit for an SMS message is 160 characters. Only log fields that fit into the notification message are selected, in the following order: Situation name, Severity, source IP address, destination IP address, destination port, Sender, Logical Interface, and the creation time of the alert.
SNMP	Only log fields that fit into the notification message are selected, in the following order: Situation name, Severity, source IP address, destination IP address, destination port, Sender, Logical Interface, alert creation time, traffic recording excerpt, and the application protocol.

Rule order in alert policies and alert chains

The system processes Alert Policies and Alert Chains from top down, so the order of the rules is important.

- In Alert Policies, rules must proceed from rules with the most limited scope to rules that are the most general.
- In Alert Chains, the order of the rules determines the order in which the alert notifications are sent.

Custom alert scripts for alert escalation

You can write a script that executes custom commands for alert escalation.

To send alert notifications using Custom Alert Scripts, you must define the Root Path on the Management Server where custom alert scripts are executed. The default location is `<installation directory>/data/notification`. All custom scripts must be stored in the same root path that is defined in the properties of the Management Server that controls the Shared Domain.

The example notification script `notify.bat` in Windows and `notify.sh` in Linux can be edited for your own use. In Linux, the sgadmin user needs read, write, and execute permissions in the script's directory.

Creating Alert elements

You can use one of the predefined Alert elements or define a Custom Alert element.

There are three predefined Alert elements in the SMC:

- The *System Alert* is reserved for alerts about the system operation.
- The *Default Alert* is a ready-made element that defines the alert that is triggered if no specific alert is triggered.
- The *Test Alert* is used in Alert Policies for testing Situations.

You can also define Custom Alert elements, which are useful if you want to configure different alert notifications for different types of events. You can create Custom Alerts and create specialized handling rules for them in your Alert Policy. System Events are automatically associated with the System Alert element.

Related tasks
Test alerts on page 428

Create Custom Alert elements

You can define your own custom Alert elements for more precise matching in the Alert Policy.

Defining custom Alert elements allows you to configure different alert notifications for different types of events and write more specific alert messages. If you set up the Management Server to send SNMP traps when alerts are triggered, you can define the SNMP code for the associated alert entries.

Alert entries are always triggered by an event, and information regarding the event is automatically included in the alert entry. How much information is included depends on the type of Alert Notification.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.
2) Right-click **Alert Configurations**, then select **New** > **Custom Alert**.
3) Configure the settings, then click **OK**.

Next steps

If you edited a Custom Alert that is already used in an Alert Policy, refresh the Alert Policy on the Domain or Domains.

Defining what triggers an alert

There are several different types of events that can trigger an alert.

The following events can trigger alerts:

- A warning or error in the operation of the SMC
- A test failure
- A match to a rule
- A threshold in a user alert check is exceeded
- A match to a pattern defined in a Situation element.

System Alerts and custom alerts are always triggered by an event in the system. In addition to the System Alerts triggered by internal events in the SMC, you can configure the following events to trigger alerts:

- You can configure a rule in your Engine, Layer 2 Engine, Layer 2 Interface, or IPS Policy to trigger an alert. .
- You can activate Status Surveillance on engines to trigger an alert when the Management Server does not receive status updates for a while.
- You can configure the engine tester to issue an alert whenever a test fails (for example, when a network link goes down). Some tests that run on the engine by default might already be configured to issue alerts.
- Server Pool Monitoring Agents can trigger alerts when they detect problems with the servers.
- You can set thresholds for user alert checks to trigger alerts when the threshold is reached.
- You can set thresholds for monitored items in Overviews to trigger alerts when the threshold is reached.

Related concepts

Getting started with the Security Engine tester on page 637
Getting started with inbound traffic management on page 749

Related tasks

Set thresholds for monitored items in Overview elements on page 232
Enable or disable status monitoring on page 356
Define logging options for Access rules on page 901

Configure notifications for alerts

Alert Notifications are ways to send notifications to administrators.

By default, alert notifications are only sent to administrators through user notification in the SMC Client. You can also send alerts in the following ways:

- **E-Mail** — Alert notifications are sent as email using an SMTP server.
- **SMS** — Alert notifications are sent as an SMS text message over HTTP, using an SMTP server, or with a script that forwards the message to a third-party tool. You can add multiple SMS Channel Types. If the first SMS channel fails, the subsequent SMS channels are used in the order in which they are listed.
- **SNMP** — The SNMP Trap Code specified in the custom alert is sent using an SNMP server.
- **Custom Alert Scripts** — Alerts are sent for processing to a script you create.

SMS messages sent by script are limited to one line in length. The maximum length of the SMS messages sent by script depends on the third-party tool that is used in sending the messages. The standard character limit for SMS messages is 160 characters. To use a script, install a third-party tool that forwards the SMS messages, such as gnokii, and the drivers for the tool on the same host. If the tool is not installed on the Management Server host, configure the script for sending the alert notifications to access the tool remotely.



Note

If you have installed a third-party tool, make sure that your Engine Policy or Layer 2 Engine Policy allows the traffic from the Management Server to the host.

If you want to send alerts in one or more of these ways, you must integrate external components, such as a GSM modem or an SMTP server. This integration is done in the properties of the Management Server element. In an environment with multiple Management Servers, you must define alert notifications for all Management Servers, even if a Management Server does not currently control any Domains.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Dashboard** > **Servers / Devices Dashboard**.
2) Browse to **Management Server**.
3) Right-click the active Management Server, then select **Properties**.
4) Click the **Notifications** tab, then configure one or more notification methods.
5) Click **OK**.

Alert Chain elements and how they work

Alert Chain elements define which notification channels are used to send alert notifications to administrators.

Alert Chains contain rows that are read from top to bottom. Alert Chains are used in Alert Policies. The Alert Policy defines which Alerts trigger Alert Chains. You can also add delays between the notifications to give the administrators time to respond.

The Final Action row determines what happens when all Alert Channels in the Alert Chain have been tried, but none of the administrators have acknowledged the alert:

- The alert escalation can stop.
- The alert can be automatically acknowledged.
- The alert can be redirected to some other Alert Chain.
- The alert processing can return to the Alert Policy for further matching.

Create Alert Chain elements

Alert Chain elements specify how alert notifications are sent to administrators.

You can create multiple Alert Chains.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.
2) Browse to **Alert Configurations** > **Alert Chains**.
3) Right-click **Alert Chains**, then select **New Alert Chain**.
4) Configure the settings, then click **OK**.

Result

The Alert Chain opens for editing.

Edit Alert Chain elements

Alert Chain elements are composed of rows ordered from top to bottom. Each row specifies a notification method and a recipient.

You can only enter one recipient and one notification method per row. You must add more rows in the following cases:

- You want to use the same notification method for more than one recipient.
- You want to use more than one notification method for the same recipient.

The Final Action row determines what happens when all Alert Channels in the Alert Chain have been tried, but none of the Administrators have acknowledged the alert.



Tip

It is not mandatory to add any rows to an Alert Chain. For example, you can use only the Final Action to automatically acknowledge or stop the escalation of alert entries that the Alert Policy directs to the chain.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.

2) Browse to **Alert Configurations** > **Alert Chains**.

3) Right-click an Alert Chain, then select **Edit <name>**.

4) Add a rule:
   - In an empty Alert Chain, right-click the **Final Action** row, then select **Rule** > **Add Rule**.
   - In an Alert Chain with existing rules, right-click a rule, then select **Rule** > **Add Rule Before** or **Rule** > **Add Rule After**.

5) Select the Alert Channel.

6) Specify the Destination of the alert notification.
   The destination information varies according to the selected alert channel.

7) (Recommended) Double-click the **Threshold to Block** cell, then set a limit for how many alerts the designated recipient is sent.

8) (Mandatory for the Delay channel, optional for other channels) In the **Delay** cell, enter the number of minutes before the next row of the alert chain is processed.
   - The purpose of the delay is to give to the recipient of the notification enough time to acknowledge the alert before the next notification is sent.
   - If sending the notification through the selected channel fails, the delay entered here is ignored. If you want to add delays that are always valid, add a row with **Delay** as the alert channel, then set the delay on that row.

9) Select the Final Action that the SMC takes if the last row of the Alert Chain is reached.

10) Click **Save**.

Creating Alert Policy elements

Alert Policies determine the criteria for selecting which alerts generated by various sources are escalated to which Alert Chains.

Engines, Layer 2 Engines, IPS engines, and SMC servers are possible sources for alerts. If Domain elements have been configured, you can select a Domain as a Sender in an Alert Policy in the Shared Domain.

An Alert Policy contains rules for matching incoming alert entries. Alert entries that match an Alert Policy rule are escalated to the Alert Chain defined in the rule. Make sure that your Alert Policies also escalate System Alerts. If an alert entry does not match any rule in the Alert Policy, the alert entry is not escalated.

Create Alert Policy elements

Alert Policies specify Alert Chains and the Alerts that trigger them.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.

2) Browse to **Alert Configurations** > **Alert Policies**.

3) Right-click **Alert Policies**, then select **New Alert Policy**.

4) Configure the settings, then click **OK**.

Result

The Alert Policy opens for editing.

Edit Alert Policy rules

Alert Policy rule settings include the Alert Sender, the Alert and Situation, Time, and Severity.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.
2) Browse to **Alert Configurations** > **Alert Policies**.
3) Right-click an Alert Policy, then select **Edit <name>**.
4) Add a rule:
   - In an empty Alert Policy, right-click the rule table, then select **Rule** > **Add Rule**.
   - In an Alert Policy with existing rules, right-click a rule ID, then select **Rule** > **Add Rule Before** or **Rule** > **Add Rule After**.
5) Specify the rule settings.
6) Select which **Alert Chain** is processed when an alert event matches this rule.
7) Click **Save**.

Install Alert Policy elements

Changes made to Alert Policies or the Alert Chains used by the Alert Policy take effect when you install the Alert Policy on a Domain.

You can install the same Alert Policy on multiple Domains. In this case, you must install the Alert Policy from the Shared Domain or from the Domain to which the Alert Policy belongs.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.
2) Browse to **Alert Configurations** > **Alert Policies**.
3) Right-click the Alert Policy, then select **Install Policy**.
4) Select the Domains on which you want to install the Alert Policy, then click **Add**.
5) Click **OK**.

Acknowledge active alerts

New alerts are handled as active alerts until they are acknowledged. To stop alert escalation, you acknowledge active alerts.

When an SMC component generates an alert, it sends the alert to the Log Server. The Log Server stores the alert entry. A new alert entry is handled as an active alert by the Management Server. A Domain's active alerts are visible when you are logged on to the Domain. Active alerts are stored on the Management Server until the alerts are acknowledged. In an environment with multiple Management Servers, active alerts are automatically replicated between the Management Servers.

Alert entries are displayed in the **Active Alerts** view and in the **Logs** view with other types of log entries. You can also view alert entries in the Web Portal.

You can acknowledge alert entries in the **Active Alerts** view. When an alert entry is acknowledged, it is removed from the **Active Alerts** view and from the Management Server. An audit entry is created when an alert is acknowledged. All Alert Chain processing for that alert entry is stopped. You can acknowledge alerts one by one. You can alternatively aggregate similar types of alerts as a group and acknowledge the whole group of alerts at the same time.

Note

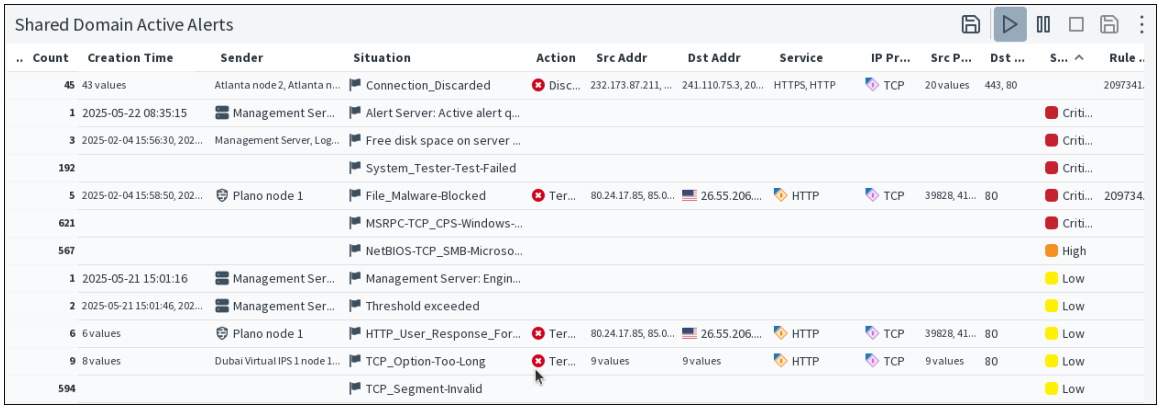
When you acknowledge an alert entry, alert escalation stops for that alert entry and no new notifications are sent out from the Management Server to administrators.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Click **Active Alerts** in the bottom right corner of the SMC Client.
   By default, the **Active Alerts** view opens with all active alerts aggregated by Situation.

Active Alerts view



2) Select one or more alert entries.
   - To aggregate the alert entries by time or sender, select ⋮ **More actions** > **Aggregate** > **Sort by Time** or ⋮ **More actions** > **Aggregate** > **Aggregate by Sender**.
   - To view the individual alerts, select ⋮ **More actions** > **Details**.
   - The **Query** pane allows you to filter the active alert entries so that you can find the information you need.

3) Right-click the selected alerts, then select **Acknowledge**.

Related concepts
Benefits of filtering log entries on page 289

How custom scripts for alert escalation work

Before writing a custom alert script, review the following information about the command arguments and script file location.

All custom scripts must be stored in the same root path that is defined in the properties of the Management Server that controls all Domains.

The example notification script notify.bat in Windows and notify.sh in Linux can be edited for your own use. In Linux, the sgadmin user needs read, write, and execute permissions in the script's directory.

The alert information is given to the script as command arguments as described in the following table.

Arguments passed to the custom scripts

Argument Number	Content	Description
1	Alert ID	The unique identifier for the alert.
2	Alert Name	The name defined in the alert properties.
3	Alert Originator	The IP address of the component that generated this alert.
4	Alert Date	The date when the alert was originally generated.
5	Alert Message	A short alert description.
6	Alert Severity	The Severity value of the alert from 1–10, where 1 is the least severe and 10 is the most severe. The numeric Severity value corresponds to the following Severity value in the generated alert: 1= Info, 2–4=Low, 5–7=High, and 8–10=Critical.
7	Alert Short Description	The contents of the Comment field in the alert properties.
8	Event ID	IPS only: reference to the event ID that triggered the alert.
9	Situation Description	Long description of the Situation that triggered the alert.

Alert scripts stored in the directory defined in the Management Server element's properties can be called from Alert Chains by their name.

When the alert script is executed, the output (`stdout`) is appended to the `notify.out` file in the script's directory. The error output (`stderr`) is appended to the `notify.err` file in the script's directory. The Linux script in the following illustration is an example of how to create an operating system log entry using the custom script alert notification.

Example custom alert script

```
#!/bin/sh
# This script uses the ‘logger’ utility to create an operating system
# log entry of the alert notification.
PATH=/bin:/usr/bin

# Create log entry: “SMC Alert (<ALERT_ID>): Severity <SEVERITY>
#    : <ALERT_NAME> : <ALERT_DESCRIPTION>”

/usr/bin/logger “SMC Alert ($1): Severity $6 : $2 : $5”

exit 0
```

Create SMTP Server elements

SMTP Server elements define the properties of SMTP servers that you can use, for example, to send email and SMS notifications.

You can also use SMTP Server elements as the Source or Destination in a policy. SMTP Server elements can be used in the properties Management Server elements for emailing Reports and sending alerts as email and SMS.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Network Elements**.
2) Browse to **Servers**.
3) Right-click **Servers**, then select **New** > **SMTP Server**.
4) Configure the settings, then click **OK**.

Related tasks

Email reports on page 331

Use a script for SMS notification

You can use a script to send alert notification SMS messages through a third-party tool that forwards the SMS messages to the administrators.

Before you begin

If there is a Engine or a Layer 2 Engine between the Management Server and the host on which the tool for forwarding the SMS messages is installed, you must allow the traffic from the Management Server to the host. Do this in the Engine or Layer 2 Engine Policy. You must install the tool that forwards the SMS messages from the Management Server.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Install a third-party tool (for example, gnokii) to forward SMS messages and the drivers for the tool on the host.
   See the tool-specific instructions for information about configuring the tool.
2) In the SMC Client, select **Dashboard** > **Servers / Devices Dashboard**.
3) Browse to **Management Server**.
4) Right-click the Management Server, select **Properties**, then click the **Notifications** tab.
5) Select **Add** > **Script**.
6) Define the following script properties:
   - **Name** — The file name of the script.
   - **Script Path** — The full path to the script that sends the SMS, or the relative path from the execution directory.
   - **Execution Path** — The directory in which the script is executed. The execution log is stored in this directory.
7) (Optional) Click **Test** if you want to test that the SMS messages are sent correctly.
8) Click **OK**.
9) Click **OK** in the Management Server properties to save the changes to the notifications.

SNMP for the SMC Appliance

SNMP traps and monitoring access can be configured to receive information about the health of the SMC Appliance.

Using SNMP monitoring software or an SNMP station, the appliance can be polled for information about available disk space, memory utilization, and running processes. Polling the appliance regularly helps you to identify appliance events early, such as partitions becoming full or processes running hot. SNMP Traps can be configured so that the appliance alerts the specified user or community.

Both SNMP 2c and 3 are supported. Configuring SNMP access to the SMC Appliance is done from the SMC Client. The list of available Net-SNMP common MIBs can be found in the /usr/share/snmp/mibs directory. Visit https://www.net-snmp.org for MIB descriptions.

Note

SNMP users and communities have read-only access to the appliance.

Related reference

SNMP traps and MIBs on page 658

Create an SNMP Agent for SNMP version 1 or 2c

Configure an SNMP Agent for SNMP version 1 or 2c so that Security Engines can share network management information using the SNMP protocol, or for SNMP version 2c so that the SMC Appliance can share network management information using the SNMP protocol.

Note

The SMC Appliance does not support SNMP v1.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Browse to **Other Elements** > **Engine Properties** > **SNMP Agents**.

3) Right-click **SNMP Agents**, then select **New SNMP Agent**.

4) In the **Name** field, enter a unique for the SNMP Agent.

5) From the **Version** drop-down list, select **v1** or **v2c**.
   For the SMC Appliance, you must select **v2c**.

6) (Optional) In the **Monitoring** section, click **Add**, then enter the community string.
   The community string is used for authentication in monitoring.

7) (Optional) In the **Listening Port** field, enter the UDP port number that the SNMP agent listens to.

8) In the **Contact** field, enter the contact information for the person responsible for the Security Engines or the SMC Appliance.

9) Click **OK**.

Create an SNMP Agent for SNMP version 3

Configure an SNMP Agent for SNMP version 3 so that Security Engines or the SMC Appliance can share network management information using the SNMP protocol.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Browse to **Other Elements** > **Engine Properties** > **SNMP Agents**.

3) Right-click **SNMP Agents**, then select **New SNMP Agent**.

4) In the **Name** field, enter a unique for the SNMP Agent.

5) From the **Version** drop-down list, select **v3**.

6) In the **User Names** section, add one or more users names.
   a) Click **Add**.
   b) In the **User Name** field, enter the user name.
   c) From the **Protocol** options, select the authentication protocol, then enter a password in the **Password** field.
   d) From the **Privacy** options, select the privacy protocol, then enter a password in the **Privacy Password** field.

7) (Optional) In the **Monitoring** section, click **Add**, then select the user for monitoring.

8) (Optional) In the **Listening Port** field, enter the UDP port number that the SNMP agent listens to.

9) In the **Contact** field, enter the contact information for the person responsible for the Security Engines or the SMC Appliance.

10) Click **OK**.

Configure what triggers SNMP traps

The trap parameters define where and how SNMP traps are sent from Security Engines and the SMC Appliance.

The same SNMP Agent element can be used for Security Engines and the SMC Appliance. Some settings only apply to Security Engines. Settings that are not supported for the SMC Appliance are ignored when the SNMP Agent is used for the SMC Appliance.

In addition to the general events, the tester on each Security Engine can send SNMP traps when a test fails.

Note

If the **Destinations** field is left empty, no traps are sent, and the other trap parameters are ignored. If the **Destinations** field has a value, the rest of the trap parameters must also have a value.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Browse to **Other Elements** > **Engine Properties** > **SNMP Agents**.

3) Open the SNMP Agent properties in one of the following ways:
   - Right-click an existing SNMP Agent element, then select **Properties**.
   - To create an SNMP Agent element, right-click **SNMP Agents**, then select **New SNMP Agent**.

4) In the **Traps** section, specify the sender of the SNMP trap.
   - SNMPv1 — (Security Engines only) In the **Community** field, enter a community string.
   - SNMPv2c — In the **Community** field, enter a community string.
   - SNMPv3 — From the **User Name** drop-down list, select a user name.

5) Click **Add**, then enter the IP address and UDP port where the traps are sent.

6) (Security Engines only) In the **Active Traps** section, select the events for which you want to set a trap.

7) Click **OK**.

Activate the SNMP agent for the SMC Appliance

You can configure access for SNMP software or stations to gather data about the SMC Appliance.

Before you begin

Create and configure an SNMP Agent element.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Settings** > **Global System Properties**
2) Click the **SNMP** tab.
3) From the **SNMP Agent** drop-down list, select the SNMP Agent that you want to activate.
4) (SNMPv3 only) In the **SNMP Engine ID** field, enter a unique identifier for the SMC Appliance.
5) In the **SNMP Location** field, enter the string that is returned on queries to the SNMPv2-MIB or SNMPv2-MIB-sysLocation object.
6) Click **OK**.

Related tasks

Activate the SNMP agent on Security Engines on page 662

Test alerts

You can test the correct functioning of the alerts by sending a Test Alert.

Test Alerts have a severity value of **Info** and have their own specific Situation. Engine, Layer 2 Engine, and IPS engines cannot send Test Alerts. The sender is always an SMC server, so it is not possible to test how alerts from other components are handled using a Test Alert.

Note

A Test Alert is escalated only if the Alert Policy rules match Test Alerts.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.

2) Browse to **Alert Configurations** > **Alerts**.

3) Right-click the alert to test, then select **Test Alert**.

4) Select the server that sends the alert, then click **Select**.

5) Verify that the alert entry appears in the **Active Alerts** view.

Related tasks
Acknowledge active alerts on page 421

Examples of alert escalation

These examples show some common uses for alert escalation, and general steps on how each example is configured.

Example: disabling alert escalations for a specific situation

You might want to disable alert escalations for a specified Situation, as shown in the following example.

The administrators at company A notice that the system issues an alert every time someone mistypes their password when logging on using the SMC Client. They have set up their IPS system to detect if there are several failed logons within a short period (which could indicate malicious activity). For this reason, they decide that they do not want to receive alert notifications about failed logons.

The administrators:

1) Create an Alert Chain and name it Auto-acknowledge.

2) Set the final action for the Auto-acknowledge Alert Chain to Acknowledge without adding any new rows.

3) Add the following new rule at the top of the Alert Policy:

Sender	Alert and Situation	Chain
ANY	"Management Server: Login Failed" Situation Element	"Auto-acknowledge" Alert Chain

4) Refresh the Alert Policy on the Shared Domain.

Alerts for failed logons are still generated and stored, but they do not trigger any alert notification and they are never shown in the **Active Alerts** view. For example, reports can still include information about failed logon attempts to highlight excessive logon failures.

Example: escalating alerts to specific administrators and sites

You can escalate alerts to specified locations and administrators, as shown in the following example.

Company B has two sites, a branch office (BO) and a headquarters (HQ) site, which both have their own administrators. Both sites have a Engine and a Log Server, and the shared Management Server is at the HQ site. Domains are not used, so all elements are in the Shared Domain. The administrators decide to set up alert escalation.

For the most severe alert entries, alerts are sent as an SMS text message to the shared mobile phone each site has for the administrator on duty. If the administrator at one site does not acknowledge the alert entry within 15 minutes, the alert notification is sent to the administrator at the other site.

For less severe alert entries, the alerts are only escalated to the site where the alert entry is created. At first, the less severe notifications are sent only through a User Notification in the SMC Client. After an hour, the alert notification is sent as an SMS text message to the shared mobile phone of the site where the alert entry is created.

The administrators:

1) Create new Alert Chains for high-severity and low-severity alert entries for both the HQ and the BO sites. There are four Alert Chains in total.
"HQ Important Alerts" contains the following rules:

Channel	Destination	Delay
SMS	[Phone number for HQ shared mobile phone]	15 min
SMS	[Phone number for BO shared mobile phone]	

"HQ Minor Alerts" contains the following rules:

Channel	Destination	Delay
User Notification	HQ Administrator A HQ Administrator B [other administrators]	60 min
SMS	[Phone number for HQ shared mobile phone]	

The "BO Important Alerts" and "BO Minor Alerts" Alert Chains are the same as the HQ Alert Chains, but with the BO Administrators and a different phone number.

2) Create an Alert Policy with the following rules:

Sender	Alert and Situation	Severity	Chain
HQ Engine HQ Log Server Management Server	ANY	High... Critical	HQ Important Alerts
HQ Engine HQ Log Server Management Server	ANY	Info...Low	HQ Minor Alerts

Sender	Alert and Situation	Severity	Chain
BO Engine BO Log Server	ANY	High...Critical	BO Important Alerts
BO Engine BO Log Server	ANY	Info...Low	BO Minor Alerts

3) Configure SMS Notification in the Management Server's properties.

4) Install the new Alert Policy on the Shared Domain.

Chapter 22

Domain elements

Contents

- Getting started with Domain elements on page 433
- How Domain elements work on page 434
- Create and modify Domain elements on page 436
- Log on to a Domain on page 438
- Log off from all Domains on page 441
- Move elements between Domains on page 442
- View Domain status on page 443
- Delete a Domain on page 443
- Examples of Domain elements on page 444

Domain elements allow you to restrict which elements are displayed to the administrators in the SMC Client and in the optional Web Portal. They also allow you to define in which administrative Domains an administrator has permissions. Configuring Domains requires a special license.

Getting started with Domain elements

Domain elements help you manage large networks and define administrator permissions.

In a large system, there can be different geographical sites that are managed by different administrators. Typically, most of the administrators only manage SMC components at their own site. Only a few main administrators are responsible for the overall system health across all sites. Domain elements allow you to group elements that belong to specific configurations (for example, elements that belong to a particular site or customer). The elements in different Domains are kept separate from each other.

The administrators' rights within a Domain depend on the permissions defined in the administrator accounts. You can grant access for an administrator to one or more Domains and define the permissions for each Domain in fine detail.

How Domains can be configured

- Domain elements allow you to group elements that belong to specific configurations (for example, elements that belong to a specific customer or site).
- You can use Domains to divide responsibilities between administrators, so that administrators only have access to elements in specific Domains.
- You must have a special license to be able to configure Domain elements. The number of Domains that you can create depends on the license.
- The **ALL Domains** Access Control List is a default Access Control List that you can use in administrator accounts to grant access to all defined Domains.

- The predefined *Shared Domain* is meant for all elements that do not belong to a particular customer or site. All predefined system elements belong to the Shared Domain. If there is no Domain license in the SMC or no Domains have yet been configured, all elements belong to the Shared Domain.

Shared Domain

- The elements in the Shared Domain are displayed to all administrators when they are logged on to any Domain in the SMC Client.
- Domains, Management Servers, Log Pruning Filters, and Administrator accounts with unrestricted permissions (superusers) are elements that automatically belong to the Shared Domain. You can only create these elements in the Shared Domain, and you cannot move them to any other Domain.
- Licenses and update packages always belong to the Shared Domain.
- If you have Master Engine and Virtual Engine elements, the Master Engine must either belong to the Shared Domain or to the same Domain as the Virtual Security Engines.

Related concepts
Configuration of Master Engines and Virtual Engines on page 527

How Domain elements work

Elements belong to one or more Domains. Administrator permissions in the Domain determine which elements administrators can manage.

When Domains are used, each element automatically belongs to a Domain. An element can only belong to one Domain at a time. By default, all elements belong to the Domain in which they are created. The Shared Domain is meant for elements that are used in several Domains, for example, high-level policy templates. All predefined system elements also automatically belong to the Shared Domain.

When administrators log on to a Domain, they can manage the elements in the Domain according to the permissions granted for that specific Domain. Administrators can also view most elements that belong to the Shared Domain even when they are not allowed to log on to the Shared Domain. However, the contents of the elements are only displayed to administrators who have permission to view those elements' contents. Elements in the Shared Domain can only be edited from within the Shared Domain.

If there are existing elements when you first start using Domains, all existing elements belong to the Shared Domain. You can move the elements to other Domains as necessary. In an environment with more than one Management Server, you can also change the active Management Server that controls all Domains.

Using elements with administrative Domains

Each element automatically belongs to either the Domain in which it was created or to the Shared Domain. When you create elements, first log on to the correct Domain and then create the elements so that the elements belong to the right Domain.

You can freely decide to which Domain most elements belong, except for the following elements:

- Domains, Management Servers, Log Pruning Filters, and Administrator accounts with unrestricted permissions (superusers) are elements that automatically belong to the Shared Domain. You can only create these elements in the Shared Domain, and you cannot move them to any other Domain.

- Licenses and update packages always belong to the Shared Domain.
- The Management Server's internal LDAP user database (the LDAP Domain element called InternalDomain). Configure external LDAP servers in the Domains to create Domain-specific accounts for end-user authentication.
- If you have Master Engine and Virtual Engine elements, the Master Security Engine must either belong to the Shared Domain or to the same Domain as the Virtual Security Engines.

In addition, there are limitations for selecting the Domain for some elements that are closely associated with other elements:

- A Log Server that is selected as the Log Server for a Management Server must belong to the Shared Domain.
- If a Log Server has a backup Log Server, both Log Servers must belong to the same Domain.
- A Log Server and the Security Engines that send their event data to the Log Server must be in the same Domain.
- A Task and the target of the Task (for example, an Export Log Task and the target Log Servers) must be in the same Domain. Otherwise, the Task cannot be run.
- By default, all elements used in a VPN must belong to the same Domain. You can also use some elements that belong to the Shared Domain when you configure a VPN in another Domain. These elements include the VPN Client gateway, Certificate Authorities, Gateway Certificates, Gateway Profiles, Gateway Settings, and VPN Profiles.



Note

The elements in the Shared Domain are displayed to all administrators when they are logged on to any Domain in the SMC Client.

When an administrator modifies a simple element, Management Server checks whether the administrator has permissions to all the granted elements that the edited element refers to. This is done to avoid a situation where one administrator modifies an element used in several sub-domains without the administrators of other sub-domains getting notified about this change.

The Management Server reference check is enabled on all the SMC versions. To allow administrators to make changes to elements shared in several sub-domains, you need perform the following steps to disable the reference check.

- On the Management Server host, open the `<smc_installation_folder>/data/SGConfiguration.txt` file for editing.
- Add the `CHECK_REFERENCES_DURING_EDITION=false` definition.
- Save the file.
- Restart Management Server service.



Note

The other sub-domain administrators can notice the change from **Pending Changes**.

Related concepts

How Categories help you view only certain elements on page 201

Create and modify Domain elements

Create Domain elements, for example, to group elements in different configurations and restrict administrator permissions.

You can import and set a logo for each Domain.

Create Domain elements

Only administrators that have unrestricted permissions (superusers) can create Domains.

Before you begin

To create or manage Domains, you must log on to the active Management Server.

You can create as many Domains as you need. Your Domain license defines how many Domains you can create.

A service break for the SMC is highly recommended when introducing Domains into the system, assigning the existing elements to the correct Domains, and changing the administrator accounts.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.
2) Browse to **Other Elements** > **Domains**.
3) Right-click **Domains**, then select **New Domain**.
4) Give the Domain a unique **Name.**
5) (Optional) Enter a **Comment** for your own reference.
6) (Optional) Fill in the **E-mail Address** and **Phone Number** fields with information that you want to be displayed in the Domain Overview.
   You can enter, for example, the information for the contact person at your company or the administrator responsible for the Domain.
7) (Optional) Click **Add** and select the Categories to use in the **Default Category Filter**.
8) Click **OK**.

Related concepts

How Categories help you view only certain elements on page 201

Related tasks
Log on to a Domain on page 438
Move elements between Domains on page 442

Set a Domain logo

You can define a logo for each Domain. The Domain logo is displayed in the upper right corner of the SMC Client window and next to the Domain's name in the Domain Overview. The Domain logo is also displayed in the Web Portal.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.
2) Expand the **Other Elements** branch and click **Domains**.
3) Right-click the Domain and select **Properties**.
4) Click the **Logo** tab in the **Domain Properties** dialog box.
5) Select the logo option from the list:
   - Select **None** to remove a previously selected logo so that the Domain has no logo.
   - Select a logo from the list or select **Select** and select a logo in the dialog box that opens.
   - Select **New** to import a new logo.
6) Click **OK**.
   The **Logo** tab shows a preview of the selected Logo.
7) Click **OK**.

Related tasks
Write announcements to Web Portal users

Import a Domain logo

Before you can associate a logo with a domain, you must import it.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.
2) Expand the **Other Elements** branch and click **Domains**.
3) Right-click the Domain and select **Properties**.
4) In the **Logo Properties** dialog box, enter a unique **Name** for the logo.
5) Click **Browse** and select the image (a .jpg, .gif, .png, .tif, .bmp, or .pnm file).
6) Click **OK**.
   The **Logo** tab shows a preview of the selected Logo.
7) Click **OK**.

Log on to a Domain

If you only have permissions in a single Domain, you automatically log on to the Domain when you log on to the SMC Client. If you have permissions in more than one Domain, you must log on to the correct Domain before managing elements that belong to the Domain.

Note

You can be logged on to more than one Domain at the same time.

Related tasks

Create Domain elements on page 436

Log on to a Domain when logging on to the SMC Client

If you have permissions in more than one Domain, you can log on to a specific Domain when you log on to the SMC Client.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Start the SMC Client.
2) In the SMC Client **Logon** dialog box, select or enter the server's IP address in the **Server Address** field, then enter a forward (/) or backward (\) slash and the name of the Domain.

For example, `192.168.200.31/Example Domain` or `192.168.200.31\Example Domain`.



Tip

Select **Remember Server Address** to avoid entering the information when logging on in future.



Note

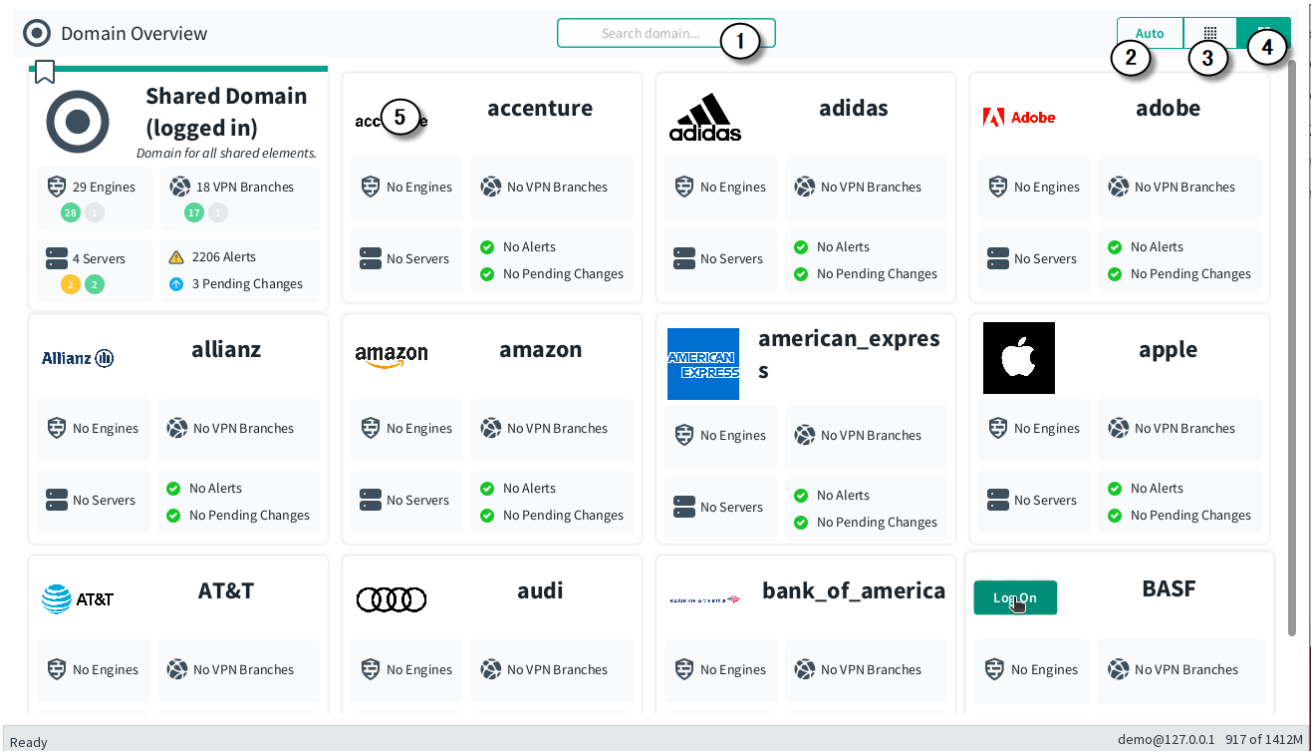
In an environment with multiple Management Servers, if you log on to a standby Management Server and you also log on to a Domain, the Domain opens on the standby Management Server.

Log on to a Domain from the Domain Overview

If you have permissions in multiple Domains, the **Domain Overview** opens after you have logged on to the SMC Client. You must log on to the correct Domain before managing elements that belong to the Domain.

For more details about the product and how to configure features, click **Help** or press **F1**.

Domain Overview



1 **Search Field:** You can use the **Search** field to filter domain cards as per the domain name.

2 **Auto:** Selects the domain card size (small or large) to display automatically based on the space available in the **Domain Overview** window.

3 **Small Card:** Displays small size domain cards for the domains that are available to the administrator. The card only contains the domain logo and name.

4 **Large Card:** Displays large size domain cards for the domains that are available to the administrator. The card contains the domain logo, domain name, summary of elements within the domain, and their current status along with the number of pending changes and active alerts.

5 **Domain Card:** Contains information related to a specific domain. You can hover over the domain card, and then click **Log On** to logon to the domain to manage elements that belong to that domain.

Note

By default, the **Shared Domain** card is bookmarked and is always placed as the first item in the **Domain Overview** window.

Steps

1) If the **Domain Overview** is not open, select **Menu** > **File** > **New Tab**, then select **Domain Overview**. The **Domain Overview** window opens.

2) Hover over the Domain card, then select **Log On**.

Note

In a HA environment with multiple Management Servers, when you log on to a Domain from the **Domain Overview**, the Domain is by default opened on the active Management Server.

If you want to log on to a Domain from a standby Management Server, select **File** > **SMC Client HA Administration** in the **Domain Overview**, right-click the standby Management Server in the **SMC Client HA Administration** dialog box that opens, then select **Log On**. A new **Domain Overview** window opens for the standby Management Server. Hover over the Domain card that you want to open on the standby Management Server, then select **Log in (read only) <Domain name>**.

Log on to a Domain from the Configuration view

If you have permissions in more than one Domain, you must log on to the correct Domain before managing elements that belong to the Domain. You can log on to a Domain from the **Configuration** view of the SMC Client.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.

2) Browse to **Other Elements** > **Domains**.

3) Right-click the Domain, then select **Log On to <Domain name>**.

Log off from all Domains

If you have permissions in more than one Domain, you can log off from all Domains through the Domain Overview.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Menu** > **File** > **New Tab**, then select **Domain Overview**.

2) Select **File** > **Log out from All**.

Move elements between Domains

You can move most elements from one Domain to another.

You can move existing elements from one Domain to another. Only administrators with unrestricted permissions (superusers) can move elements between Domains. When you start moving elements from one Domain to another, the Management Server automatically searches for element references. You can then either remove the references between the elements or move the referred or referencing elements. In addition to individual elements, you can also move all elements associated with a Category. Using Categories can make moving elements easier if you need to move many elements.

The following elements always belong to the Shared Domain and cannot be moved:

- Predefined system elements
- Domains
- Management Servers
- Licenses
- Update packages
- Log Pruning Filters.
- Administrator accounts with unrestricted permissions (superusers)

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Log on to the Domain from which you want to move elements.

2) Select the elements that you want to move to another Domain.



Tip

To hide elements that belong to the Shared Domain, select ⋮ **More actions** > **Show Only Current Domain**.

3) Right-click one of the selected elements, then select **More actions** > **Move to Domain**.

4) Click **Select** next to the **Target Domain** field, then select a Domain.

5) Click **Move**.

The Management Server automatically searches for references to or from the selected elements in the current Domain.

A new view opens and shows the reference status.

- If the element you are moving references another element, you must either remove the references or move the referenced element as well.

6) If the **Element References** pane shows **Referring** or **Referenced Elements**, expand the branches to see the referring or referenced elements and resolve the element references:
   - If more detailed information is available for an element, double-click **...** in the **Details** column next to the element.
   - If you also want to move the referring or referenced element, select the referring or referenced element, then click **Add**.

- If you do not want to move the referring or referenced element, right-click the element, then edit it to remove the references to the element that you are moving.

If you want to move a Host element that is used in a policy, but not the policy itself, remove the Host from the policy before moving the Host to a different Domain.

7) If you have resolved a reference, click **Refresh View** to update the status of element references.
   - You must resolve all element references before moving the selected elements.
   - If the **Element References** pane is empty, there are no element references.

8) Click **Continue** to move the elements to the selected Domain.

View Domain status

The Domain Overview allows you to see at a glance the status of the Domains and their elements. You do not need to log on to each Domain to monitor its status.

The Domain Overview is available only to administrators who have permissions in more than one Domain. The Domain Overview only shows information from the Domains in which the administrator has permissions. The information in the Domain Overview depends on the administrator's rights. The Domain Overview shows the statuses of the elements and the number of alerts in each Domain. The Domain Overview also shows any other information (for example, email addresses and telephone numbers) defined in the Domain properties.



Note

The Domain Overview automatically opens when an administrator who has permissions in more than one Domain logs on to the SMC Client without specifying a Domain.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Menu** > **File** > **New Tab**, then select **Domain Overview**.

Related tasks

Log on to a Domain on page 438
Log off from all Domains on page 441

Delete a Domain

If you delete a Domain, all elements that belong to the Domain are also deleted.

If there are elements that you do not want to delete, move them to another Domain before deleting the Domain. You cannot delete the predefined Shared Domain. Only administrators who have unrestricted permissions (superusers) can edit and delete Domains.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.

2) Expand the **Other Elements** branch and click **Domains**.

3) Right-click the Domain and select **Delete.**
   A **Confirmation** dialog box opens.

4) If the Domain you are deleting is used in any configuration, click **Open References** in the **Confirmation** dialog box to view the references. Right-click each element, then select **Edit** to remove the references.

5) Click **Yes**.
   A **Confirmation** dialog box opens.

6) Type `YES` to confirm that you want to permanently delete the Domain.

Related tasks

Create Domain elements on page 436
Move elements between Domains on page 442

Examples of Domain elements

You might want to create a different Domain for each of your organization's customers or sites.

Example: Create Domain elements for different customers

You can create a different Domain for each of your organization's customers, as shown in the following example.

Company A is a Managed Security Service Provider (MSSP) with many customers. The networks of different customers must be kept separate. The administrators who manage the customer networks must only be allowed to see the networks for which they are responsible. Most of the administrators only manage a single customer's network, but some of the administrators are responsible for several customers' networks.

The administrator decides to use Domain elements to group the elements belonging to each customer and to make it easier to manage the different customer networks. The administrator also decides to use Category elements to tag the existing elements that are included in each Domain. As the user database information must not be available across Domains, the administrator decides to use an external LDAP server in each Domain for user authentication.

Company A's administrator:

1) Arranges a service break with the customers before introducing Domains into the system.

2) Logs on to the Shared Domain and creates the following elements:

- A separate Domain element for each customer.
- The Administrator elements (the administrator accounts) for the administrators who manage several customers' networks in several Domains.
- A Category element for each customer's elements.

3) Defines a default Category Filter that includes the customer-specific Category for each customer's elements.

4) Logs on to each customer's Domain and creates the Administrator elements (the administrator accounts) for the administrators who manage only that particular customer's network.

5) While logged on to each Domain, configures the elements for using an external LDAP server for authenticating the users in the Domain and for storing the Domain's user database.

6) While logged on to the Shared Domain, moves all customer-specific elements from the Shared Domain to the correct customer-specific Domain.
   - To make it easier to move the elements, the administrator first selects the customer-specific Category and then all elements that belong to the Category.

7) When all customers' Domains and their elements have been configured and the service break is over, the administrators for each customer company log on to the SMC Client.
   - The administrators who are responsible for a single customer's networks automatically log on to the Domain assigned to them when they log on to the SMC Client. They only see the elements that belong to their own configuration and the elements in the Shared Domain.
   - The administrators who have permissions in several Domains must select the Domain when they have logged on to the SMC Client.

Example: Create Domain elements for different sites

You can create a different Domain for each of your organization's sites, as shown in the following example.

Company B is a large enterprise planning a new system. The system includes 12 different sites, each of which contains 10 networks. The administrators at each site only need to be able to see the networks at their own sites. The headquarters administrator decides to use the Management Server's internal LDAP user database for user authentication in all Domains. This means that all administrators in each Domain are able to view the user database information.

The headquarters administrator:

1) Logs on to the Shared Domain and creates Domains to represent each of the 12 sites.

2) Configures the user database and user authentication using the internal LDAP directory of the SMC while logged on to the Shared Domain.

3) Logs on to each Domain that represents a site's configuration and creates the elements for the Domain:
   - The Administrator elements (the administrator accounts) for the administrators of each site.
   - All other elements that belong to each Domain.

When the administrators at each site log on to the SMC Client, they also automatically log on to the Domain assigned to them. They only see the elements that belong to their own site's configuration and also the elements in the Shared Domain.

Chapter 23

Getting Started with the Web Portal

Contents

- Enabling the Web Portal Client on page 448
- Create Web Portal User accounts on page 449

Using the Web Portal, customers of managed service providers can access information about their systems.



Note

You must have a license for running a Web Portal and creating Web Portal users.

What the Web Portal Does

The Web Portal provides restricted client-less access to Logs, Reports, and Policy Snapshots. It is useful for managed service providers for providing customers with information about their systems. There is no software for end users to install. They can access the information using a web browser.

The process of setting up a Web Portal consists of the following steps:

1) Enable the Web Portal Client. For more details, refer to the Enabling the Web Portal Client topic.

2) Create Web Portal User accounts for the end users. For more details, refer to the Create Web Portal User accounts topic.



Note

SMC Client administrator accounts are also valid in the Web Portal.

Enabling the Web Portal Client

The Web Portal Client provides browser-based access to Reports, Policy Snapshots, and Logs for specific authorized users.

Before you begin

- You must have the Client API setting configured and enabled before you enable the Web Portal Client feature. For more details on SMC API settings, refer to the *Configure SMC API* topic.
- When you upgrade from SMC version 7.2.X or earlier to version 7.3 and the SMC API is not configured and enabled. The web portal will not function properly after the upgrade. To ensure the web portal works, you must first manually configure and enable the SMC API.
- It is strongly recommended that you configure the SMC API in HTTPS mode.

Steps

1) Select **Dashboard** > **Servers / Devices Dashboard**.
2) Right-click the Management Server, then select **Properties**.
3) Click the **Client API** tab.
4) In the **Web Portal Client** section, select the **Enable** checkbox.
5) The URL next to the **Enable** checkbox can be used to access the Web Portal in a web browser.

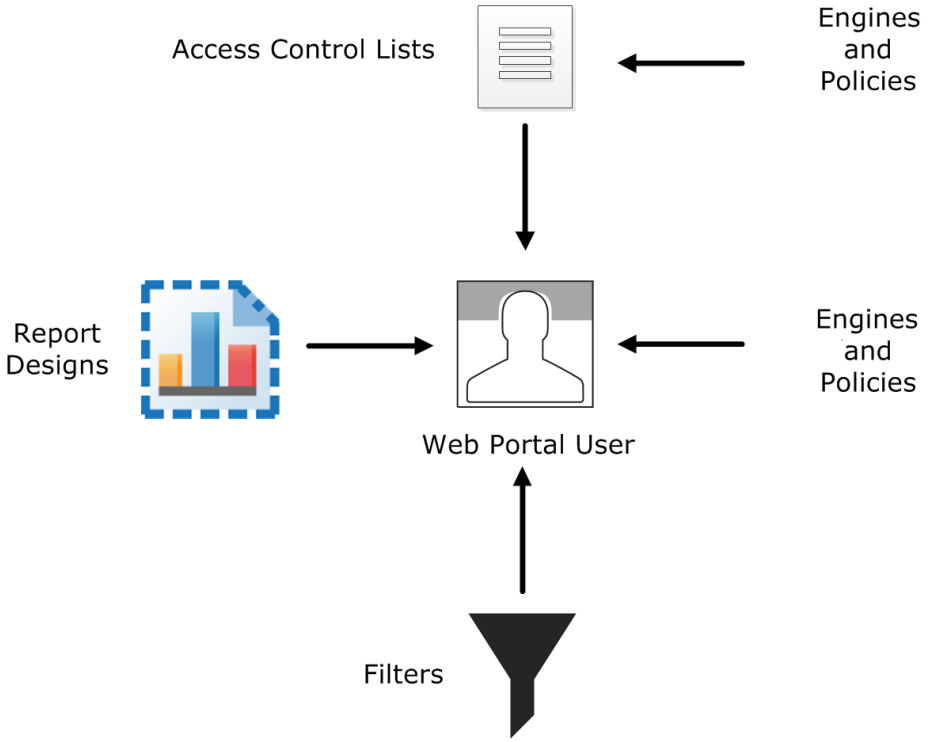
Note

The URL for the web portal includes the same host name and port number used for the SMC API. For example, in the URL http://example.com:8082/webclient/, the *example.com* represents the host name, and the *8082* is the port number used in the SMC API configuration.

Create Web Portal User accounts

The accounts for the optional Web Portal are defined with Web Portal User elements. It is highly recommended to create a unique Web Portal User account for each Web Portal User.

Elements for Web Portal User accounts



- Engine elements define which logs, reports, or policy snapshots are displayed.
- Policies, sub-policies, and template policies define which parts of the Policy Snapshots are displayed.
- *Report Designs* define which reports are displayed. The Web Portal user is allowed to view all generated reports that are based on the granted Report Designs.
- *Filters* define which logs are displayed. You can also add Filters that the Web Portal User can choose to apply when browsing logs.

Web Portal Users can also use internal authentication or external RADIUS authentication.

If administrative Domains are used, there are some more considerations:

- Each Web Portal User account is limited to a single Domain.
- The Web Portal User is allowed to see all information in the Policy Snapshots from the granted engines. If a policy's template is in the Shared Domain, the Web Portal User can also see the rules inherited from the template in the Policy Snapshot.
- The Web Portal Users might be allowed to view reports generated in the Shared Domain depending on their granted elements.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.

2) Right-click **Access Rights** and select **New** > **Web Portal User**.

3) In the **Name** field, enter the user name that the Web Portal user uses to log on to the Web Portal.

4) From the **Authentication** drop-down list, select an authentication method for Web Portal Users.

5) Select engines from which the Web Portal User is allowed to view logs and Policy Snapshots.
   a) Click the **Engines** tab.
   b) To grant all engines to the Web Portal User, select **Allow ANY**.



Note

If Domain elements have been configured, only the engines in the current Domain are granted to the Web Portal user.

   c) To select individual engines, click **Add**, then select the engines and click **Select**.

6) On the **Policies** tab, select the policies from which the Web Portal User is allowed to view Policy Snapshots. You can only select policies that are installed on engines granted to the Web Portal User. You can define in detail which parts of the policies are shown in the Policy Snapshots.

7) On the **Logs** tab, select log browsing permissions for the Web Portal User.

8) On the **Reports** tab, select the kinds of reports that the Web Portal User can access.

9) Click **OK**.

Chapter 24

Using the SMC Client in a web browser

Contents

- Configuration overview on page 452
- Start the SMC Client in a web browser on page 453

To avoid installing the full Java-based SMC Client on each workstation that an administrator uses, you can run the SMC Client in a web browser.

You can enable the Web Access feature on the Management Server or Web Access Server. Administrators log on to the SMC Client on a web page, and the SMC Client runs as an HTML5 application in the web browser. The web browser is the only requirement on the workstation.

You can also connect to and manage multiple versions of the SMC. This removes the requirement to have the SMC and the locally-installed SMC Client be the same version.

You can configure Web Access in the properties of the Management Server or Web Access Server. You can also enable the Web Access feature during the installation of the Management Server.

Limitations and recommendations

- By default, the SMC allows a maximum of five sessions using Web Access at the same time. To change the maximum number of concurrent sessions, see Knowledge Base article 17248.
- Web browser support is limited to Google Chrome and Mozilla Firefox.
- It is not possible to log on using certificate-based authentication.
- Interacting with the local file system is limited. Each user has a folder located at `%installation%\data\%servertype%\webswing\users\admin%id%` where `%servertype%` is `datamgtserver` for the Management Server and `datawebserver` for the Web Access Server and `%id%` is the ID of the administrator in the database. Users can import or export elements to this folder, for example.
- When you copy text using Ctrl+C, you must manually allow the copy operation in the bottom-right corner of the screen.



Note

Using right-click actions to paste text from the system clipboard is not supported. You must use Ctrl+V to paste text into the SMC Client.

- Web Access can consume resources. Especially if many administrators will be using the feature, we recommend that you enable the feature on the Web Access Server.
- If the Management Server or Web Access Server is installed on a Linux platform, xvfb-run must be installed.
- If the Management Server and Web Access Server are installed on the same computer, we recommend that you do not enable Web Access on both servers.

Configuration overview

Follow these main steps to use the SMC Client in a web browser.

1) Enable the Web Access feature in the Management Server or Web Access Server properties.
If you enabled the feature during the installation of the Management Server, you can configure additional options in the Management Server properties.

2) Use a web browser to start the SMC Client.

Enable Web Access

You can enable and configure the feature in the properties of the Management Server or Web Access Server.

Note

To use Web Access for a Management Server or Web Access Server installed on a Linux platform, xvfb-run must be installed in /usr/bin.

Note

When Web Access is enabled, the Agentless ZTNA Application setup lets you to access SMC UI through FONE portal.

Note

Administrator must enable the ZTNA Connector to access the SMC, and add regular access policy rules to allow the traffic to the services behind the engine. For more information, see *Engine Editor* > *Add-Ons* > *ZTNA Connector*.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Network Elements**.

2) Browse to **Servers**.

3) Right-click the Management Server or Web Access Server, then select **Properties**.

4) On the **Web Access** tab, select **Enable**.

5) Configure the settings, then click **OK**.

Management Server and Web Access Server Properties	
Option	**Definition**
Host Name (Optional)	Enter the host name that the service uses. Leave the field blank to allow requests to any of the server's host names.
Port Number	Enter the TCP port number that the service listens to. By default, port 8085 is used when Web Access is enabled on the Management Server and port 8083 when enabled on the Web Access Server. **Note** Make sure that the listening port is not in use on the server.
Listen Only on Address (Optional)	If the server has several addresses and you want to restrict access to one address, specify the IP address to use.
Session Timeout	Enter the timeout in seconds after which the session expires. While the session is still active, the administrator does not need to log on again if they close the web browser.
Server Credentials	Select the TLS Credentials element that is used for HTTPS connections. Click **Select** to select an element.
Server TLS Cryptography Suite Set	Select the TLS Cryptography Suite Set element that defines the allowed algorithms for HTTPS connections. Click **Select** to select an element.
Use SSL for session ID (Optional)	Track sessions in your web application. Do not select this option if your network requires you to use cookies or URIs for session tracking.

Start the SMC Client in a web browser

Administrators can log on to the SMC Client and perform their duties using a web browser.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In a web browser, browse to the URL of the server that you configured the Web Access feature on.
   The URL can be the IP address of the server or the host name that you defined in the properties of the server. Make sure that you include the port number at the end of the URL.
   Example where Web Access is enabled on the default port 8085 on the Management Server:
   `https://127.0.0.1:8085`

2) Enter your user name and password, then click **Log On**.

3) Use the SMC Client as you normally would.

Chapter 25

SMC Client downloads from the Management Server

Contents

- Enable and configure SMC Client downloads on the Management Server on page 455
- Download the SMC Client from the Management Server on page 456

When the Management Server provides the SMC Client for download, administrators can download and install the SMC Client from the SMC Downloads page.

Enable and configure SMC Client downloads on the Management Server

To allow administrators to download and install the SMC Client from the SMC Downloads page, enable SMC Client downloads on the Management Server.

Steps

1) Select **Network Elements**.
2) Browse to **Servers**.
3) Right-click the Management Server, then select **Properties**.
4) If it is not already selected, select **Enable** on the **Client Downloads** tab.
5) Select **Management Client Download**.
6) Configure the settings, then click **OK**.

Management Server Properties

Option	Definition
Host Name (Optional)	Enter the host name that the service uses. Leave the field blank to allow requests to any of the server's host names.

Option	Definition
Port Number	Enter the TCP port number that the service listens to. By default, port 8080 is used for new SMC installations, and port 8084 is used when you upgrade the SMC. **Note** Make sure that the listening port is not in use on the server.
Listen Only on Address (Optional)	If the server has several addresses and you want to restrict access to one address, specify the IP address to use.
Server Credentials	Select the TLS Credentials element that is used for HTTPS connections. Click **Select** to select an element.
Generate Server Logs (Optional)	Select if you want to log all file load events for further analysis with external web statistics software.

Download the SMC Client from the Management Server

Download and install the SMC Client from the SMC Downloads page.

Steps

1) In a web browser, browse to the URL of the Management Server.
   The URL can be the IP address of the Management Server or the host name that you defined in the properties of the Management Server. Make sure that you include the port number at the end of the URL.

2) To download the SMC Client installation package, click **Download Management Client for <operating system>**.

   Tip
   The operating system is automatically detected. To download the SMC Client installation package for another operating system, click **Click here for other platforms**.

3) Install the SMC Client with administrator privileges.

4) Log on and use the SMC Client as you normally would.

Chapter 26

Configuring the Log Server

Contents

- Modify Log Server elements on page 457
- Select backup Log Servers for high availability on page 458
- Forwarding log data from Log Servers to external hosts on page 459
- Edit Log Server configuration parameters on page 464
- Certify Log Servers on page 465

You can modify a Log Server element, configure settings for Log Servers, and recertify Log Servers.

Modify Log Server elements

One Log Server element is automatically created during SMC installation. You can change the settings as necessary.

You can:

- Rename the Log Server element.
- Change the Log Server's IP address.
- Change the platform on which the Log Server runs.
- Define other Log Servers that you can use as backup Log Servers.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Dashboard** > **Servers / Devices Dashboard**.

2) Browse to **Log Server**.

3) Right-click the Log Server, then select **Properties**.

4) Change the Log Server properties.



Note

We recommend that you always use the default port 3020 if possible. To use a non-standard port, manually add Access rules to allow communications using the new port from the Security Engines to the Log Server.



Note

Be careful when excluding Log Servers from reporting. If you select this setting for a Log Server that is in use, there is no warning that generated reports are missing data.

5) Click **OK**.

Related concepts

Define contact IP addresses on page 127
Getting started with reports on page 311
Element-based NAT and how it works on page 633

Related tasks

Create Location elements on page 127

Select backup Log Servers for high availability

You can select one or more backup Log Servers for each Log Server for high availability.

Before you begin

You must already have more than one Log Server.

The same Log Server can simultaneously be the main Log Server for some components and a backup Log Server for components that primarily use another Log Server. You can also set Log Servers to be backup Log Servers for each other so that whenever one goes down, the other Log Server is used.



Note

The SMC Appliance does not support high availability for the Management Server or the Log Server.

If Domain elements have been configured, a Log Server and its backup Log Server or Log Servers must belong to the same Domain.

CAUTION

If the log volumes are high, make sure that the backup Log Server can handle the traffic load in failover situations.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Dashboard** > **Servers / Devices Dashboard**.
2) Browse to **Log Server**.
3) Right-click the Log Server for which you want to select a backup Log Server, then select **Properties**.
4) Click the **High Availability** tab.
5) Click **Add**.
6) Select one or more Log Servers, then click **Select**.
7) Click **OK**.

Related concepts

Log Servers for HA on page 468

Forwarding log data from Log Servers to external hosts

You can forward log data from Log Servers to external hosts to back up the data or to process the data in an external system.

You can define which type of log data you want to forward and in which format. You can also use Local Filter elements to specify in detail which log data is forwarded.

Log data does not need to be stored on the Log Server to be sent to the external host. If log pruning is applied, any log data that the Immediate Discard log pruning filters delete is not forwarded to the external host.

Add log forwarding rules to Log Servers

Add log forwarding rules to the Log Server to enable log forwarding.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Dashboard** > **Servers / Devices Dashboard**.
2) Browse to **Log Server**.
3) Right-click the Log Server from which you want to forward log data, then select **Properties**.
4) Click the **Log Forwarding** tab.
5) To create a rule, click **Add**.



Tip

To remove a rule, select the rule, then click **Remove**.

6) Configure the log forwarding rules.
7) Click **OK**.

Enable TLS protection for log or audit data forwarding

You can optionally enable TLS protection for log or audit data forwarding to an external syslog server.

Before you begin

Because there is a connection to an external system, public key infrastructure (PKI) integration, including certificate revocation list (CRL) checking, must already be configured.

You can optionally configure TLS server identity to verify the identity of the syslog server to which log data is forwarded from the Management Server or the Log Server.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Dashboard** > **Servers / Devices Dashboard**.
2) Browse to **Log Server** or **Management Server**.

3) Right-click the Log Server or Management Server from which you want to forward log or audit data, then select **Properties**.

4) Click the **Log Forwarding** or **Audit Forwarding** tab.

5) Add rules for log or audit data forwarding.
   a) To add a rule, click **Add**.
   b) To select the external host to which the log or audit data is forwarded, double-click the **Target Host** cell, select a Host element, then click **Select**.
   c) In the **Service** cell, select a **network protocol with TLS** option from the drop-down list.
   d) In the **Port**, **Format**, **Data Type** (Log Server only), and **Filter** cells, select the settings according to your needs.
   e) Double-click the **Kafka Topic** cell to specify a name for the kafka topic, if you have selected the **Kafka with TLS** option from the **Service** drop-down list.
   f) To select the TLS profile for TLS-protected log data forwarding, double-click the **TLS Profile** cell, select a TLS Profile element, then click **Select**.

6) (Optional) Configure the TLS Server Identity.
   a) Double-click the **TLS Server Identity** cell.
   b) From the **TLS Server Identity** drop-down list, select the server identity type field to be used.
   c) (Optional) Click **Fetch from Certificate** to fetch the value of the server identity type field from a certificate.

   Note

   You can fetch the value of the server identity field from a certificate only if the server identity field is **Distinguished Name**, **SHA-1**, **SHA-256**, **SHA-512**, or **MD5**).

   d) In the **Identity Value** field, enter the value of the server identity field.

7) Define the Log Server or Management Server TLS certificate options.
   This certificate is used as the client certificate when connecting to the external syslog server.
   - To use the server's internal certificate, select **Use Internal Certificate**.
   - To use the certificate contained in a TLS Credentials element, select **Use Imported Certificate**, then click **Select**.
     Select a TLS Credentials element.
   - To leave the server's certificate unauthenticated, select **No Client Authentication**.

8) Click **OK**.

Enable logging for monitored traffic

To generate log data that can be forwarded to an external host, you must enable logging for the traffic that you want to monitor.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Expand **Policies**, then browse to the type of policy you want to edit.
3) Right-click the policy and select **Edit <policy type>**.
4) Click the **IPv4 Access** or **IPv6 Access** tab, then edit the rule that allows the traffic that you want to monitor.
   If there is no Access rule for the traffic that you want to monitor, create the Access rule.
5) Double-click the **Logging** cell.
6) Select **Override Collected Values Set With "Continue" Rules**.
7) From the **Log Level** drop-down list, select **Stored** or **Essential**.
8) (Optional) If you want to forward logs using the NetFlow or IPFIX format, select **Log Accounting Information** in the **Connection Closing** drop-down list.
9) Click **OK**.
10) Save and install the policy to start using the new configuration.

Related concepts
Getting started with Access rules on page 831

Define general syslog settings

You can adjust general log forwarding settings by editing the LogServerConfiguration.txt. Adjusting these settings is optional.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Stop the Log Server:
   - If you run the Log Server as a service in Windows, you can stop it in the Windows Control Panel's **Services** list.

- In Linux, run the script <installation directory>/bin/sgStopLogSrv.sh.

2) Create a text file on the Log Server that lists the fields to forward in the correct order.
See Knowledge Base article 10010 for more information.



Tip

The <installation directory>/data/fields/syslog_templates/ directory contains example configuration files.

3) Change the parameters in LogServerConfiguration.txt.
The file is located in <installation directory>/data/.

Log Server configuration

Parameter	Value	Description
`SYSLOG_CONF_FILE`	`<File name>`	Path to the file you created in Step 2, which defines the fields that are forwarded and their order.
`SYSLOG_MESSAGE_PRIORITY`	`0-191 a`	The priority of the syslog message is included at the beginning of each UDP packet (the default is 6). a) As defined in RFC 3164 (https://www.ietf.org/rfc.html).

`SYSLOG_USE_DELIMITER`	`ALWAYS_EXCEPT_NULL`
	`NEVER`
	`ALWAYS`
		Defines whether to use double quotes (“) in syslog messages to delimit the field values.
		The default setting “ALWAYS_EXCEPT_NULL” uses double quotes only for non-empty fields. “NEVER” does not use delimiters. “ALWAYS” uses double quotes as delimiters for all empty and non-empty field values.

4) Save the file and restart the Log Server.

Related concepts
Log entry fields

Related reference
Syslog entries

Add Access rules allowing traffic from Log Servers to external hosts

If the external host and Log Server are separated by a Engine or Layer 2 Engine, you must add rules to allow traffic from the Log Server to the host.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Expand **Policies**, then browse to the type of policy you want to edit.
3) Right-click the policy, then select **Edit Engine Policy** or **Edit Layer 2 Engine Policy**.
4) Click the **IPv4 Access** or **IPv6 Access** tab, then add an Access rule with the following values:
   - **Source** — Log Server
   - **Destination** — Host element
   - **Service** — Syslog (UDP), Syslog (TCP), or NetFlow (UDP), depending on the protocol used. For TLS-protected traffic, select TCP with TLS.

     The same Service and Port that was selected in the Log Forwarding rule must be selected here.
   - **Action** — Allow
   - **Logging** — **None** (recommended in most cases)

   Note

   Logging the log forwarding can create a loop where the log forwarding creates a log entry each time. If you want to log the log forwarding, create a local filter in the Log Forwarding rule to exclude logs related to forwarding.
5) Save and install the policy to start using the new configuration.

Edit Log Server configuration parameters

To configure the Log Server in detail, you can edit LogServerConfiguration.txt. Normally, it is not necessary to configure the Log Server outside of the SMC Client. However, under special circumstances, you might want more control over the way the Log Server behaves.

For more information, see Knowledge Base article 19177.

Certify Log Servers

If the Log Server was not certified during the installation or if it needs a new certificate, certify the Log Server manually.

Steps

1) Stop the Log Server:
   - If the Log Server is installed as a service in Windows, stop it in the Windows Control Panel's **Services** list.
   - In Linux, run the script <installation directory>/bin/sgStopLogSrv.sh.

2) Request the certificate:
   - In Windows, run the script <installation directory>/bin/sgCertifyLogSrv.bat.
   - In Linux, run the script <installation directory>/bin/sgCertifyLogSrv.sh.

3) Enter the credentials for an administrator account with unrestricted permissions (superuser).

4) If administrative Domains are configured and the Log Server does not belong to the Shared Domain, enter the name of the Domain.

5) Wait for the certification to finish and start the Log Server again.
   - In Windows, start it in the Windows Control Panel's **Services** list.
   - In Linux, run the script <installation directory>/bin/sgStartLogSrv.sh.

Chapter 27

Configuring SMC servers for high availability

Contents

- Using additional SMC servers for high availability on page 467
- Management Server HA configuration overview on page 468
- Log Server HA configuration overview on page 472
- Manage HA Management Servers and Log Servers on page 476

You can install several Management Servers and Log Servers to provide high availability for the SMC.

Note

The SMC Appliance does not support high availability for the Management Server or the Log Server.

Using additional SMC servers for high availability

You can install additional Management Servers and Log Servers for high availability.

The high availability (HA) solution includes automatic incremental replication of the configuration data stored on the Management Server. This way, manual intervention is minimized, and the SMC can be fully managed and monitored without manually reinstalling and restoring a backup.

A Management Server and a Log Server are required for configuration changes to and system monitoring of engines. Although engines work independently without the SMC according to their installed configuration, configuration changes and system monitoring are not possible without a Management Server and Log Server. The Management Server in particular is a critical component, as it is the only place where the full configuration information is stored.

Management Servers for HA

When a Management Server is active, the Management Server has control of all Domains and can be used for configuring and managing the SMC. Only one Management Server at a time can be the active Management Server.

The changes made on the active Management Server are replicated incrementally to the other Management Servers: only the changed parts of the management database are replicated, and the replication takes place in real time.

To use additional Management Servers, you must have a special Management Server license that includes the high availability feature. The license is a combined license for all Management Servers and it must list the IP addresses of all Management Servers.

Additional Management Servers automatically replicate the configuration data on the active Management Server. If the active Management Server becomes unusable, you must manually activate another Management Server, which allows you to work normally.

Note

The additional Management Servers are meant for backup and disaster recovery purposes. Only one Management Server at a time can be used for configuring and managing the SMC.

Log Servers for HA

Installing additional Log Servers ensures that system monitoring continues if one Log Server fails.

CAUTION

The logs are not copied between the Log Servers. To back up log data, set up an Archive Log Task to copy log data from the Log Server to some other location.

Alert escalation proceeds normally, new logs can be browsed, and the engine status and statistics can be examined. However, the log data is not automatically replicated between the Log Servers, so some log data is always unavailable during outages. We recommend that you select the same Log Server for all engines in the same location. If different engines send events to different Log Servers, it is not possible to correlate events, as none of the Log Servers see all events.

Any Log Server can be used both as the main Log Server for some components and as a backup Log Server for one or more other Log Servers. However, consider the load of the Log Servers before you set up a Log Server as a backup Log Server, to avoid overloading when failover occurs.

You can set up additional Log Servers with normal Log Server licenses. A separate license for each Log Server is required, even if the Log Server is used only as a backup.

You can install a new Log Server that works as a backup for another Log Server. You can also set existing Log Servers to be used as backups for another Log Server.

Management Server HA configuration overview

Configuring additional Management Servers requires the following overall steps.

1) Create elements for the additional Management Servers.
2) License the additional Management Servers.
3) Allow communications to the new Management Servers through engines as necessary.
4) Install the Management Server software on the target servers.

Create additional Management Server elements

Any new Management Server elements you add to an existing SMC are considered to be additional Management Servers. You can set up several additional Management Servers.

The number of additional Management Servers you can create depends on the Management Server license.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Dashboard** > **Servers / Devices Dashboard**.

2) Select **More actions** > **New** > **Server** > **Management Server**.

3) In the **Name** field, enter a unique name.

4) In the **IP Address** field, enter the IP address.

5) (Optional) If NAT is used to translate addresses in communications between this server and other SMC components, define the **Location** and **Contact Address**.

   Note

   Management Server replication does not use contact addresses. Make sure that other Management Servers can communicate with this Management Server using its untranslated IP address.

6) From the **Log Server** drop-down list, select the Log Server to which the Management Server sends its logs.

7) If you use a RADIUS server or TACACS+ server for authenticating administrators, select the authentication method from the **RADIUS Method** or **TACACS Method** drop-down list.

8) (Optional) If you want to send alert notifications or reports from the Management Server, configure the options on the **Notifications** tab.
   - If you want to send alerts by email, as SMS text messages, through a custom script, or as SNMP traps, define the alert channels.
   - You can use an SMTP server to send generated reports directly from the Management Server. Select the SMTP Server or create an SMTP Server element.

     In the **Sender Address** field, enter the email address that is shown as the email sender. Remember to allow these connections in the Engine Policy if necessary.

9) Click **OK**.

Related concepts

Define contact IP addresses on page 127

Related tasks

Create Location elements on page 127
Email reports on page 331
Authenticate administrators using RADIUS or TACACS+ methods on page 392
Configure notifications for alerts on page 416
Create SMTP Server elements on page 423

Install licenses for additional Management Servers

Using additional Management Servers requires a special combined license that lists the IP addresses of all Management Servers within the same SMC.

After receiving the proof-of-license (POL) code, generate the license at https://stonesoftlicenses.forcepoint.com.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Settings** > **Install Licenses**.
2) Browse to the license file on your local workstation and click **Install**.
3) Select **Administration**.
4) Browse to **Licenses** > **Unlicensed Components** > **Servers** and make sure that the additional Management Server is not listed there.

   If the server is listed, check that the IP address information is the same in the license and the additional Management Server element.

Related tasks

Generate licenses on page 1333

Access rules for additional Management Servers

The Firewall Template policy contains rules that allow connections between the Engine and all SMC servers. If other components connect through a Engine to additional Management Servers, add rules that allow this traffic.

You might also have to add NAT rules.

The rules in the Firewall Template policy are IPv4 Access rules that allow the system communications.

Related concepts
Getting started with Access rules on page 831

Related reference
Forcepoint Security Management Center ports on page 1457
Security Engine ports on page 1460

Install additional Management Server software

Install additional Management Servers using the SMC Installation Wizard.

Before you begin

To install the software, you need the SMC installation files.

Steps

1) Start the installation in one of the following ways:
   - From a .zip file: unzip the file and run setup.exe on Windows or setup.sh on Linux.
   - From a DVD: insert the installation DVD and run the setup executable from the DVD:

Operating System	Path to Executable
Windows 64-bit	\Forcepoint_SMC_Installer\Windows-x64\setup.exe
Linux 32-bit	/Forcepoint_SMC_Installer/Linux/setup.sh
Linux 64-bit	/Forcepoint_SMC_Installer/Linux-x64/setup.sh



Note

If the DVD is not automatically mounted in Linux, mount the DVD with the following command: `mount /dev/cdrom /mnt/cdrom`

2) Proceed according to instructions in the Installation Wizard until you are prompted to select which components you want to install.



Note

If you install the SMC in `C:\Program Files\Forcepoint\SMC`, the installation creates an extra `C:\ProgramData\Forcepoint\SMC` folder, which duplicates some of the folders in the installation directory and also contains some of the program data.

3) If you also want to install a Log Server and a local SMC Client on this computer, leave **Typical** selected and click **Next**. Otherwise, select **Custom**, select the components you want to install and click **Next.**

4) Select the IP address of the Management Server from the list or type it in.

Note

This IP address must be the IP address defined for the corresponding Management Server element.

5) Type in the IP address of the Log Server for sending alerts.

6) Select **Install as an Additional Management Server for High Availability**.

7) Click **Next** and follow the instructions to start the installation.
A logon prompt for replication opens.

8) Log on using an unrestricted administrator account.
The **Management Server Selection** dialog box opens.

9) Select the correct Management Server from the list and click **OK**.
The databases are synchronized.

Note

If the synchronization fails for some reason (such as a network connection problem), run the sgOnlineReplication script on the additional Management Server when connectivity is restored.

Tip

You can view replication information in the **Info** pane when you select the Management Server.

Related tasks

Obtain SMC installation files on page 1347

Log Server HA configuration overview

Configuring additional Log Servers requires several high-level steps.

1) Create an element for the additional Log Server.

2) License the additional Log Server.

3) Define the Log Server as a backup Log Server for another Log Server.

4) Make sure the communications between SMC components and the additional Log Server are allowed.

5) Install the Log Server software on the target server.

CAUTION

The logs are not copied between the Log Servers. To back up log data, set up an Archive Log Task to copy log data from the Log Server to some other location.

Create additional Log Server elements

You can set up several backup Log Servers for each Log Server.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Dashboard** > **Servers / Devices Dashboard**.
2) Select **More actions** > **New** > **Server** > **Log Server**.
3) In the **Name** field, enter a unique name.
4) In the **IP Address** field, enter the IP address.
5) (Optional) If NAT is used to translate addresses in communications between this server and other SMC components, define the **Location** and **Contact Address**.
6) (Optional) Change the Log Server's TCP Port number if necessary. We recommend always using the default port 3020 if possible.

   Note

   If you use a non-standard port, add rules to allow communications from the engines to the Log Server ports, even when using the Firewall Template policy.

7) (Optional) If you do not want the Log Server to gather statistical information for monitoring and reports, select **Exclude from Log Browsing, Statistics and Reporting**.

   Tip

   You might want to select this option if the Log Server is not used for daily operations.

Related concepts

Define contact IP addresses on page 127

Related tasks

Create Location elements on page 127

Install licenses for additional Log Servers

Each Log Server requires a separate license, even if it is only used as a backup for some other Log Server.

After receiving the proof-of-license (POL) code, generate the license at https://stonesoftlicenses.forcepoint.com.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Settings** > **Install Licenses**.
2) Browse to the license file on your local workstation and click **Install**.
3) Select **Administration**.
4) Browse to **Licenses** > **Unlicensed Components** > **Servers** and make sure that the additional Log Server is not listed there.

   If the server is listed, check that the IP address information is the same in the license and the additional Log Server element.

Related tasks

Generate licenses on page 1333

Access rules for additional Log Servers

The Engine Template policy contains rules that allow connections between the Engine and all SMC servers. If other components connect through a Engine to additional Log Servers, add rules that allow this traffic.

You might also have to add NAT rules.

The rules in the Firewall Template are IPv4 Access rules that allow the system communications.

Related concepts

Getting started with Access rules on page 831

Related reference

Forcepoint Security Management Center ports on page 1457
Security Engine ports on page 1460

Install additional Log Server software

Install additional Log Servers using the SMC Installation Wizard.

Before you begin

To install the software, you need the SMC installation files.

Steps

1) Start the installation in one of the following ways:
   - From a .zip file: unzip the file and run setup.exe on Windows or setup.sh on Linux.
   - From a DVD: insert the installation DVD and run the setup executable from the DVD:

Operating System	Path to Executable
Windows 64-bit	\Forcepoint_SMC_Installer\Windows-x64\setup.exe
Linux 32-bit	/Forcepoint_SMC_Installer/Linux/setup.sh
Linux 64-bit	/Forcepoint_SMC_Installer/Linux-x64/setup.sh



Note

If the DVD is not automatically mounted in Linux, mount the DVD with the following command:

```
mount /dev/cdrom /mnt/cdrom
```

2) Proceed according to instructions in the Installation Wizard until you are prompted to select which components you want to install.

3) Select **Custom**, then select **Log Server**.

4) Finish the installation according to instructions in the Installation Wizard.
   You must certify the Log Server before the Log Server can connect to the Management Server.

Result

The additional Log Server is now installed and shown in green (Online) in the **Dashboard** view.

Related tasks

Obtain SMC installation files on page 1347

Manage HA Management Servers and Log Servers

When multiple Management Servers are configured, one Management Server is defined as active. You can manually change the active Management Server.

You can also disable database replication and synchronization for selected Management Servers.

You can use the SMC Client, the command line of the Management Server, or the SMC API to manage HA Management Servers.

Set the active Management Server

If the currently active Management Server is unreachable, set a different Management Server to active. If is also recommended to set a different Management Server active before you change the platform of the Management Server that is currently active.



Note

Changing the active Management Server is primarily meant for disaster recovery. We do not recommend changing the active Management Server unless it is necessary.



CAUTION

If the currently active Management Server is unreachable when you change the active Management Server, but might possibly recover, disconnect the active Management Server from the network. Disconnecting the currently active Management Server prevents having two active Management Servers online at the same time.

When a Management Server is the active Management Server, it has full control of all administrative Domains. Setting a Management Server to standby releases full control of all administrative Domains. Another Management Server can take full control of all administrative Domains and become the new active Management Server.

The engines continue to operate normally even when no Management Server is reachable, so there is no interruption to any network services.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Dashboard** > **Servers / Devices Dashboard**.

2) Browse to **Management Server**.

3) Right-click the management server, and then select **SMC Client HA Administration**.

4) Right-click an additional Management Server, then select **Set Active**.
   You are prompted to confirm setting the Management Server to active.

5) Click **Yes**.

6) Set the previous active Management Server to standby, then manually synchronize the databases to return to normal operation.

Related reference
Forcepoint Security Management Center commands on page 1429

Disable or enable automatic database replication from active Management Server to additional Management Servers

By default, changes in the configuration information stored on the currently active Management Server are automatically copied to all other Management Servers.

Before you begin

There must be a route between the SMC Client and the Management Servers. If there is no route between the SMC Client and the Management Servers, you cannot send a command through the SMC Client HA Administration dialog box.

Automatic management database replication is incremental and continuous; only the changed parts of the database are replicated and the replication is done in real time.



Note

The management database is not synchronized automatically between the Management Servers after a Management Server upgrade. You must synchronize the database manually after the upgrade.

Disabling automatic database replication is not recommended unless you have a specific need to do so. For example, you might need to prevent excessive failed replication attempts when you know that the Management Server will be unreachable for a long time. An example scenario would be changing the Management Server's hardware or the Management Server's IP address.



CAUTION

Re-enable automatic database replication in the Management Server's properties when the Management Server is reachable. The management database is automatically synchronized after you have re-enabled automatic database replication and the Management Server is working normally. After successful database replication, the management database status is shown as OK.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Dashboard** > **Servers / Devices Dashboard**.

2) Browse to **Management Server**.

3) Right-click the management server, and then select **SMC Client HA Administration**.

4) Right-click the Management Server, then select **Replication** > **Include Server and Full Database Replication** or **Replication** > **Exclude Server from Replication**.
You are prompted to confirm that you want to include or exclude the server.

5) Click **Yes**.
The Management Server is included or excluded until the next full database replication.

Synchronize databases between the active Management Server and additional Management Servers

You must synchronize the configuration information manually through the SMC Client after upgrading the Management Servers or after restoring a backup.

Before you begin

There must be a route between the SMC Client and the Management Servers. If there is no route between the SMC Client and the Management Servers, you cannot send a command through the SMC Client HA Administration dialog box.

Manual management database synchronization is primarily meant for resynchronizing the databases after upgrading the SMC. We do not recommend using manual database synchronization unless you have a specific need to do so.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Connect to the active Management Server using the SMC Client.

2) Select **Dashboard** > **Servers / Devices Dashboard**.

3) Browse to **Management Server**.

4) Right-click the management server, and then select **SMC Client HA Administration**.

5) If the SMC Client is in a different network than the additional Management Server, select the **Location** from which to send the command.

6) For each additional Management Server that you want to synchronize:

   a) Right-click the Management Server, then select **Replication** > **Full Database Replication**.

b) When prompted to confirm the replication, click **Yes**.
All existing configurations on the additional Management Server are overwritten.

c) Click **OK** to acknowledge the completion of the synchronization, then wait for the Management Server to restart.

Result

After the Management Server has restarted, its Replication Status is updated in the **SMC Client HA Administration** dialog box. Click **Close** to close the **SMC Client HA Administration** dialog box.

Chapter 28

Reconfiguring the SMC and Security Engines

Contents

- Modify Management Server elements on page 481
- Configure SMC API on page 483
- Forward audit data from Management Servers to external hosts on page 486
- Change the Management Server database password on page 489
- Change the Management Server or Log Server platform on page 490
- Change Management Server and Log Server IP addresses on page 491
- Troubleshooting connecting to Management Servers on page 495
- Things to consider when changing the Security Engine role on page 496

You can modify settings for Management Servers, change hardware platforms or the IP addresses used in system communications, change the type of certificate authority, and change the role of Security Engines.

Modify Management Server elements

One Management Server element is automatically created during SMC installation. You can change the settings as necessary.



Note

You can rename the Management Server element freely, but you cannot delete the Management Server that the SMC is connected to.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) (Multiple Management Servers only) Select **Dashboard** > **Servers / Devices Dashboard**, Browse to **Management Server**, Right-click the management server, and select **SMC Client HA Administration**, then temporarily exclude the Management Servers you are about to edit from database replication.

   Temporarily exclude the Management Servers from database replication, for example, for troubleshooting purposes, while changing the Management Server's hardware, or changing the Management Server IP address.

2) Select **Dashboard** > **Servers / Devices Dashboard**.

3) Browse to **Management Server**.

4) Right-click the Management Server, then select **Properties**.

5) (Optional) Change the **Name** of the Management Server.

6) (Optional) If NAT is used to translate addresses in communications between this server and other SMC components, define the **Location** and **Contact Address**.

7) Select the **Log Server** to which the Management Server sends its logs.

8) If you use a RADIUS server for authenticating administrators, select the **RADIUS Method** that is used in the authentication.

9) (Optional) If you want to send alert notifications or reports from the Management Server, select the **Notifications** tab.
   - If you want to send alerts by email, as SMS text messages, through a custom script, or as SNMP traps, define the alert channels.
   - You can use an SMTP server to send generated reports directly from the Management Server. Select the SMTP Server or create an SMTP Server element. Enter the **Sender Address** that is shown as the email sender. Remember to allow these connections in the Engine Policy if necessary.

10) (Optional) If the connection from the Management Server to the Forcepoint servers requires a proxy server, select the **Connection** tab and configure the **Proxy Settings**.
    a) Select **Use Proxy Server for HTTPS Connection**, then enter the **Proxy Address** and **Proxy Port**.
    b) If the proxy server requires user authentication, select **Authenticate to the Proxy Server**, then enter the **Proxy User Name** and **Proxy User Password**.

    By default, passwords and keys are not shown in plain text. To show the password or key, deselect the **Hide** option.

    Note

    The **Proxy Address** field must contain only the proxy hostname, without http:// or https://.

11) If the replication status in the **Info** pane indicates a problem with database replication, synchronize the management databases manually.

12) Click **OK**.

Related concepts

Define contact IP addresses on page 127

Related tasks

Create Location elements on page 127
Email reports on page 331
Authenticate administrators using RADIUS or TACACS+ methods on page 392
Configure notifications for alerts on page 416
Create SMTP Server elements on page 423
Synchronize databases between the active Management Server and additional Management Servers on page 478

Configure SMC API

The Application Programming Interface (API) of SMC allows external applications to connect with the SMC.

Note

If there is a engine between SMC and the other applications, make sure that there is an Access rule to allow communication.

The SMC API can be used to run actions remotely using an external application or script. For more information about using SMC API, see the *Forcepoint Security Management Center API User Guide*.

Create TLS credentials for SMC API Clients

If you want to use encrypted connections, the SMC API Client needs TLS credentials to connect with the Management Server.

Note

You can import the existing private key and certificate if they are available.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the SMC Client, select **Administration**.
2) Browse to **Certificates** > **TLS Credentials**.
3) Right-click **TLS Credentials**, then select **New TLS Credentials**.
4) Complete the certificate request details.
   a) In the **Name** field, enter the IP address or domain name of SMC.
   b) Complete the remaining fields as needed.

c) Click **Next**.

5) Select **Self Sign**.

6) Click **Finish**.

Result

The TLS Credentials element is added to **Administration** > **Certificates** > **TLS Credentials**. The **State** column shows that the certificate has been signed.

Enable SMC API

To allow other applications to connect using the SMC API, enable SMC API on the Management Server.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the SMC Client, select **Dashboard** > **Servers / Devices Dashboard**.

2) Browse to **Management Server**.

3) Right-click the Management Server, then select **Properties**.

4) Click the **SMC Client API** tab, then select **Enable**.

5) (Optional) In the **Host Name** field, enter the name that the SMC API service uses.

Note

API requests are served only if the API request is made to this host name. To allow API requests to any host name, leave this field blank.

6) Make sure that the listening port is set to the default of 8082 on the Management Server.

7) If the Management Server has several IP addresses and you want to restrict access to one, enter the IP address in the **Listen Only on Address** field.

8) Select the TLS Credentials element that is used for HTTPS connections. Click the **Select** button against the **Server Credentials** field, then select an element.

9) If you want to use encrypted connections, click the **Select** button against the **Server TLS Cryptography Suite Set** field, then select the TLS Credentials element and the Cryptography Suite Set element.

10) Click **OK**.

Create an API Client element

External applications use API clients to connect to SMC.

Before you begin

SMC API must be enabled for the Management Server.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.
2) Browse to **Access Rights**.
3) Right-click **Access Rights** and select **New** > **API Client**.
4) In the **Name** field, enter a unique name for the API Client.
5) Use the initial authentication key or click **Generate Authentication Key** to generate a new one.

   Important

   This key appears only once, so be sure to record it. The API Client uses the authentication key to log on to SMC API.

6) Click the **Permissions** tab.
7) Select the permissions for actions in the SMC API. As a minimum, set the **Viewer** permission to **All Simple Elements**.
8) Click **OK**.

Forward audit data from Management Servers to external hosts

You can configure Management Servers to forward audit data to external hosts in CSV, XML, JSON, CEF, LEEF, or McAfee ESM format. Forwarding audit data enables you to handle and process audit data with external devices.

Add rules for forwarding audit data from Management Servers

Configure the rules that forward audit data from a Management Server.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Dashboard** > **Servers / Devices Dashboard**.

2) Browse to **Management Server**.

3) Right-click the Management Server from which you want to forward audit data, then select **Properties**.

4) Click the **Audit Forwarding** tab.

5) To create a rule, click **Add**.



Tip

To remove a rule, select the rule and click **Remove**.

6) In the select **Target Host** cell, select the external host to which the audit data is forwarded.
   a) Double-click the **Target Host** cell.
   b) Select a Host element.
   c) Click **Select**.

7) In the **Service** cell, select the network protocol for forwarding the audit data.

8) (Optional) Double-click the **Port** cell, then enter the port that is used for audit data forwarding.



Note

You might have to define an Access rule that allows traffic to the target host. In this case, make sure that the Port you select is also used as the Port in the Access rule.

9) In the **Format** cell, select the audit data forwarding format.

10) (Optional) Double-click the **Filter** cell, then define a Local Filter element that defines which audit data is forwarded.

Note

The Local Filter is only applied to the audit data that matches the Audit Forwarding rule.

11) Click **OK**.

Related tasks

Enable TLS protection for log or audit data forwarding on page 487

Enable TLS protection for log or audit data forwarding

You can optionally enable TLS protection for log or audit data forwarding to an external syslog server.

Before you begin

Because there is a connection to an external system, public key infrastructure (PKI) integration, including certificate revocation list (CRL) checking, must already be configured.

You can optionally configure TLS server identity to verify the identity of the syslog server to which log data is forwarded from the Management Server or the Log Server.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Dashboard** > **Servers / Devices Dashboard**.

2) Browse to **Log Server** or **Management Server**.

3) Right-click the Log Server or Management Server from which you want to forward log or audit data, then select **Properties**.

4) Click the **Log Forwarding** or **Audit Forwarding** tab.

5) Add rules for log or audit data forwarding.

   a) To add a rule, click **Add**.

   b) To select the external host to which the log or audit data is forwarded, double-click the **Target Host** cell, select a Host element, then click **Select**.

c) In the **Service** cell, select a **network protocol with TLS** option from the drop-down list.

d) In the **Port**, **Format**, **Data Type** (Log Server only), and **Filter** cells, select the settings according to your needs.

e) Double-click the **Kafka Topic** cell to specify a name for the kafka topic, if you have selected the **Kafka with TLS** option from the **Service** drop-down list.

f) To select the TLS profile for TLS-protected log data forwarding, double-click the **TLS Profile** cell, select a TLS Profile element, then click **Select**.

6) (Optional) Configure the TLS Server Identity.

a) Double-click the **TLS Server Identity** cell.

b) From the **TLS Server Identity** drop-down list, select the server identity type field to be used.

c) (Optional) Click **Fetch from Certificate** to fetch the value of the server identity type field from a certificate.



Note

You can fetch the value of the server identity field from a certificate only if the server identity field is **Distinguished Name**, **SHA-1**, **SHA-256**, **SHA-512**, or **MD5**).

d) In the **Identity Value** field, enter the value of the server identity field.

7) Define the Log Server or Management Server TLS certificate options.

This certificate is used as the client certificate when connecting to the external syslog server.

- To use the server's internal certificate, select **Use Internal Certificate**.
- To use the certificate contained in a TLS Credentials element, select **Use Imported Certificate**, then click **Select**.
  Select a TLS Credentials element.
- To leave the server's certificate unauthenticated, select **No Client Authentication**.

8) Click **OK**.

Related concepts

Creating certificates on page 158

Related tasks

Create TLS Profile elements on page 157

Create Access rules allowing traffic from Management Servers to external hosts

If the external host and Management Server are separated by a Engine or Layer 2 Engine, edit the Policy to allow traffic from the Management Server to the host.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Browse to **Policies**
3) Browse to **Engine Policies** or **Layer 2 Engine Policies**, depending on the type of policy that you want to edit.
4) Right-click the Engine or Layer 2 Engine policy, then select **Edit Engine Policy** or **Edit Layer 2 Engine Policy**.
5) Click the **IPv4 Access** or **IPv6 Access** tab, then add an Access rule with the following values:
   - **Source**: your Management Server
   - **Destination**: the target Host element
   - **Service**: Syslog (UDP) or Syslog (TCP), depending on the protocol used. For TLS-protected traffic, select TCP with TLS. The same **Service** and **Port** that was selected in the Audit Forwarding rule must be selected here.
   - **Action**: Allow
6) If you have finished editing the policy, click **Save and Install**.

Change the Management Server database password

The Management Server contains a database for storing the configuration information. The password for the database is automatically created during the installation of the Management Server. In some rare cases, you might need this password to complete an operation. In these cases, you can change the database password.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Settings** > **Password** > **Change Database Password**.
2) Enter the password in the **New Password** and the **Confirm New Password** fields.

3) Click **OK** to confirm the change.

Change the Management Server or Log Server platform

You can change the hardware or the operating system that the Management Server or Log Server components run on.

Note

For any Web Access Server, install the new Web Access Server as a fresh installation as instructed in the *Forcepoint Network Security Platform Installation Guide*.

Note

If you want to change the platform of the active Management Server in an environment with multiple Management Servers, it is recommended to set another Management Server as the active Management Server before you start changing the platform of the Management Server that used to be the active Management Server.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Take backups of your Management Servers and your Log Servers.

   Note

   The Web Access Server has no information to back up.

2) (Multiple Management Servers only) Open the **SMC Client HA Administration** dialog box and temporarily exclude the additional Management Servers from database replication.

3) Install new Management Servers and Log Servers on the new platforms.

   See the *Forcepoint Network Security Platform Installation Guide*.

4) (Single Management Server only) Restore the Management Server and Log Server backups from the old environment to the new installation.

   Note

   If you change the platform of a standby Management Server in an environment with multiple Management Servers, do not restore a backup on the standby Management Server. To restore the management database on the standby Management Server, synchronize the management databases manually between the active Management Server and the standby Management Server when the standby Management Server is back online.

5) (Hardware change) Shut down and disconnect the old Management Server or Log Server hardware.

6) (Hardware change) Connect the new Management Server or Log Server to the network environment.

7) (Multiple Management Servers only) If database replication between Management Servers fails, run SMC server diagnostics to troubleshoot the situation.

Related concepts
Restoring backups on page 1299

Related tasks
Synchronize databases between the active Management Server and additional Management Servers on page 478
Back up system configurations on page 1297

Change Management Server and Log Server IP addresses

You can change the IP addresses of the Management Server and Log Server without losing management connectivity. The steps to change the Management Server's IP address depend on whether the Management Server and Log Server are on separate or on the same appliance.

When you change IP addresses, other connections between the different components might be temporarily lost. Make sure that the connections return to normal after the IP address changes.

Before changing the IP address, we recommend making a backup of the Management Server and Log Server.

Related tasks
Change the IP address of combined Management Servers and Log Servers on page 494
Back up system configurations on page 1297

Change the Management Server IP address

When the Management Server and Log Server are installed on different appliances, you can change the IP address of the Management Server by following these steps.

Before changing the IP address, we recommend making a backup of the Management Server and the Log Server.

Note

If any Engines between the Management Server and other components do not use a policy based on the Firewall Template, check that they allow all necessary connections.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Request an IP address change for the Management Server license at https://stonesoftlicenses.forcepoint.com.

2) (Multiple Management Servers only) Open the **SMC Client HA Administration** dialog box and temporarily exclude the Management Servers for which you are changing the IP address from database replication.
   For more information, see the topic that explains how to synchronize Management databases manually.

3) Add Engine IPv4 Access rules (and possibly NAT rules) that allow policy upload connections from the new IP addresses to the Engine.
   The services needed for the communications between the different components are explained in the topic that lists SMC ports.

4) (Security Engines with Node-Initiated contact to Management Server only) Open the Management Server Properties and add the new Management Server IP address as a Contact Address.
   The Security Engine must be able to contact the Management Server at both the current Management Server IP address and the new Management Server IP address.

5) Refresh the Engine Policies.

6) Stop the Management Server and Log Server services.

7) Change the IP address of the host server in the operating system.

8) On the Management Server, run the command
   ```
   sgChangeMgtIPOnMgtSrv <new Management Server IP address>
   ```

9) On all Log Servers, run the command
   ```
   sgChangeMgtIPOnLogSrv <new Management Server IP address>
   ```

10) Start the Management Server and Log Server services and log on using the SMC Client.

11) Install the new Management Server license when prompted.

12) Remove the Engine IPv4 Access rules that you created in Step 3 and refresh the Engine Policies.
    After running the IP address change scripts, the Alias elements in the inherited rules translate to the right IP addresses.

13) If the replication status in the **Info** pane indicates a problem with database replication, synchronize the management databases manually.

Related tasks

Define Management Server or Log Server contact addresses on page 128
Synchronize databases between the active Management Server and additional Management Servers on page 478
Change the IP address of combined Management Servers and Log Servers on page 494
Back up system configurations on page 1297

Related reference

Forcepoint Security Management Center ports on page 1457

Change the Log Server IP address

When the Management Server and Log Server are installed on different appliances, you can change the IP address of the Log Server by following these steps.

When you change the Log Server's IP address, the traffic between the Log Server and the engines is interrupted and the logs are spooled on the engines. Changing the IP address might also mean that the transfer of engine status and statistics information is temporarily interrupted.

Before changing the IP address, we recommend making a backup of the Management Server and the Log Server.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Request a license binding change for the Log Server if the license is bound to the Log Server's IP address at https://stonesoftlicenses.forcepoint.com and install the new license.

2) Edit the <installation directory>/data/LogServerConfiguration.txt file on the Log Server and update the Log Server IP address.

3) In the SMC Client, open the Log Server properties and update the Log Server IP address.

4) Stop and restart the Log Server service.

5) Refresh the policies of all engines that send data to the Log Server.

Related tasks

Edit Log Server configuration parameters on page 464
Back up system configurations on page 1297

Change the IP address of combined Management Servers and Log Servers

When the Management Server and the Log Server are installed on the same system, you can change the IP address for both Servers in one procedure.

When you change the Log Server's IP address, the traffic between the Log Server and the engines is interrupted and the logs are spooled on the engines. Changing the IP address might also mean that the transfer of engine status and statistics information is temporarily interrupted.

Before changing the IP addresses, we recommend making a backup of the Management Server and the Log Server.

Note

If any Engines between the Management Server and other components do not use a policy based on the Firewall Template, check that they allow all necessary connections.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Request a license binding change to the new IP address for the Management Server license, and also the Log Server if the license is bound to an IP address.

   Make the request at https://stonesoftlicenses.forcepoint.com.

2) (Multiple Management Servers only) Open the **SMC Client HA Administration** dialog box and temporarily exclude the Management Servers for which you are changing the IP address from database replication.

   For more information, see the topic that explains how to synchronize Management databases manually.

3) Add Engine IPv4 Access rules (and possibly NAT rules) that allow policy upload connections from the new IP addresses to the Engine.

   The services needed for the communications between the different components are explained in the topic that lists SMC ports.

4) (Security Engines with Node-Initiated contact to Management Server only) Open the Management Server Properties and add the new Management Server IP address as a Contact Address.

   The Security Engines must be able to contact the Management Server at both the current Management Server IP address and the new Management Server IP address.

5) Refresh the Engine Policies.

6) Stop the Management Server and Log Server services.

7) Change the IP address of the host server in the operating system.

8) Run the sgChangeMgtIPOnMgtSrv script on the Management Server.

   For more information, see the topic that explains SMC commands.

9) Run the sgChangeMgtIPOnLogSrv script on the Log Server.

10) Edit the <installation directory>/data/LogServerConfiguration.txt file on the Log Server and update the Log Server IP address.
For more information, see the topic that explains changing Log Server configuration parameters.

11) Start the Management Server service and log on using the SMC Client.

12) Install the new licenses when prompted.

13) Open the **Log Server** properties and update the IP address.

14) Start the Log Server service.

15) Remove the Engine IPv4 Access rules that you created in Step 3 and refresh the Engine Policies.
After running the IP address change scripts, the Alias elements in the inherited rules translate to the right IP addresses.

16) If the replication status in the **Info** pane indicates a problem with database replication, synchronize the management databases manually.

Related tasks
Edit Log Server configuration parameters on page 464
Synchronize databases between the active Management Server and additional Management Servers on page 478
Change the Management Server IP address on page 491
Change the Log Server IP address on page 493
Start Tasks manually on page 1328
Back up system configurations on page 1297

Related reference
Forcepoint Security Management Center commands on page 1429
Forcepoint Security Management Center ports on page 1457

Troubleshooting connecting to Management Servers

If an engine that has been changed cannot connect to the Management Server, you have several troubleshooting options.

If an engine cannot connect to the Management Server because of changes in the configuration, you can restore the contact. Generate a new initial configuration for the engine (through the engine's right-click menu), then run the Security Engine Configuration Wizard on the command line. See the *Forcepoint Network Security Platform Installation Guide*.

If you suspect that the engine's configuration is out of date, select the option to return the engine to the initial configuration state. For more information, see the topic that explains reconfiguring engine settings.

CAUTION

Returning the engine to the initial configuration state clears the engine's configuration. Only management communications are allowed, and no traffic is allowed to pass through the engine.

Related tasks

Reconfigure Security Engine settings on page 365

Things to consider when changing the Security Engine role

You can change the role of an Security Engine, converting one type of Security Engine to another, if you have a specific need to do so.

Consider these things when changing the role:

- You can only change the Security Engine role for engines that currently have Forcepoint Network Security Platform software installed. To change the role of engines that currently have specific engine software (for example, Engine/VPN role) installed, you must reinstall the engine software. See the *Forcepoint Network Security Platform Installation Guide*.
- Changing the engine role is only supported on modular appliances, for engines installed on a virtualization platform, or for engines installed on your own hardware. You cannot change the engine role on small appliances.
- You must have an Security Engine license that is valid for all engine roles. You cannot change the role of engines that have a license for a specific type of engine.
- If using the Security Engine Configuration Wizard on the engine command line, you must connect through a serial console or VGA console. It is not possible to change the engine's role using an SSH connection to the engine.

Prepare for changing the Security Engine role

Before you change the role of an Security Engine from one type to another, perform these high-level steps to create an engine element for the new role.

Note

You cannot use the same primary control IP address in multiple elements. You must either change the primary control IP address in the engine's current interface configuration or delete the existing engine element before creating the new engine element.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Create the correct type of engine element for the new role and define the basic properties.

2) Configure the engine's interfaces.

3) Configure the routing.
   - Routes to directly connected networks are automatically added for all Security Engine roles.
   - For Security Engines, add a default route and any routes through next-hop gateways to networks that are not directly connected to the Engines.
   - You might need to define a default route for IPS engines and Layer 2 Engines if the SMC components are not on a directly connected network.

4) Generate the initial configuration for the engine and save the configuration on a USB drive.
   For more information, see the topic that explains management contact procedure for engines.

Related concepts
Getting started with Security Engines on page 503
Network interfaces for Security Engine on page 543
Connect Security Engines to the SMC on page 631

Clear the existing engine configuration before changing the Security Engine role

Before you change the role of an Security Engine from one type to another, clear the existing engine configuration.

Steps

1) Connect to the engine's command line using a monitor and keyboard or a serial cable, and log on as root.
   If you use a serial cable, use a terminal console program to connect to the Security Engine appliance with these settings:
   - **Bits per second** — 115,200
   - **Data bits** — 8
   - **Parity** — None
   - **Stop bits** — 1.

   Note
   The serial console port speed is 115,200 bps in most Security Engine appliances. The speed is 9600 bps in older Security Engine appliance models. See the hardware guide for your Security Engine appliance model for more information.

2) Enter the command `sg-clear-all`.
   The engine restarts and you are prompted to select the system restore options.

   CAUTION
   Using the system restore options clears all configuration information from the engine. It is not possible to recover the engine's previous configuration.

3) Type `2` and press **Enter**.
You are prompted to select the number of overwrites for the data partition.

4) Type `0` and press **Enter**.
You are prompted to select the number of overwrites for the swap partition.

5) Type `0` and press **Enter**.
You are prompted to select whether to remove the spool partition.

6) Type `2` and press **Enter**.
You are prompted to select the number of overwrites for the spool partition.

7) Type `0` and press **Enter**.
You are prompted to confirm that you want to clear the configuration.

8) Type `YES` and press **Enter**.
The system starts clearing the engine's configuration.



CAUTION

Do not turn off or restart the engine while the configuration is being cleared. Doing so can cause a serious error that prevents the engine from starting.

9) Press **Enter** to restart the engine when prompted.
The engine restarts and the Security Engine Configuration Wizard starts on the command line.

Reconfigure the Security Engine after clearing the existing configuration

After you create an engine element and clear the existing engine configuration, you can change the role of an Security Engine from one type to another.

For details about how to make initial contact using the Security Engine Configuration Wizard in a web browser, see the *Forcepoint Network Security Platform Installation Guide*.

Steps

1) Select **Role** and press **Enter**.

2) Select the new role for the Security Engine and press **Enter**.
The role-specific Security Engine Configuration Wizard starts.

3) Select one of the following configuration methods:
   - Select **Import** and press **Enter** to import a saved configuration.
   - Select **Next** and press **Enter** to manually configure the engine's settings.

4) (Manual configuration only) Configure the **Operating System Settings** and **Network Interfaces**.
   See the *Forcepoint Network Security Platform Installation Guide*.

5) Select **Switch Engine node to initial configuration** and fill in the Management Server information.

6) Select **Contact** and enter the **Management Server IP address** and the **one-time password**.



CAUTION

Select 256-bit Security Strength only if the engine is not able to communicate with the Management Server after you start using a new internal ECDSA certificate authority.

7) (Optional) Select **Edit Fingerprint** and press **Enter**. Fill in the Management Server's certificate fingerprint (also shown when you saved the initial configuration).
   Filling in the certificate fingerprint increases the security of the communications.

8) Select **Finish** and press **Enter**.
   The engine makes initial contact with the Management Server.

9) Install a policy on the engine.

Related tasks

Install policies on page 813

Part VII

Security Engine configuration

Contents

- Creating and modifying Security Engines on page 503
- Creating and modifying Master Engine and Virtual Engine elements on page 527
- Network interface configuration on page 543
- Connecting Security Engine to the SMC on page 629
- Element-based network address translation (NAT) on page 633
- Configuring the Security Engine tester on page 637
- Engine permissions on page 647
- DNS Relay on page 651
- Setting up SNMP for Security Engines on page 657
- Setting up LLDP for Security Engines on page 665
- Alias element translations for Security Engines on page 669
- Add-on features for Security Engines on page 673
- Advanced Security Engine settings on page 675

You can create and modify Engines, IPS engines, Layer 2 Engines, Master Engines and Virtual Security Engines. You can configure the Security Engine properties, activate optional features, and configure advanced Security Engine settings.

Chapter 29

Creating and modifying Security Engines

Contents

- Getting started with Security Engines on page 503
- Creating Security Engine elements on page 504
- Editing existing Security Engines on page 509
- Configure a backup unit or a connection synchronization for external high availability on page 524
- Configure global contact policy settings for node-initiated contact to the Management Server on page 524
- Synchronizing the time on Security Engines on page 526

Security Engine elements contain the configuration information that is directly related to the Engines, IPS engines, and Layer 2 Engines. The configuration information includes interface definitions, cluster mode selection, tester settings, and other options specific to the Security Engine.

Getting started with Security Engines

Security Engine elements are used for the configuration and management of your Engines, IPS engines, and Layer 2 Engines.

Security Engine elements contain settings that cannot be reused in the configuration of other components, such as the network interface configuration. Security Engine elements also determine which reusable elements are included in the configuration of a particular component. For example, a component can send its log data to the Log Server, which is a reusable element. All Security Engines are centrally managed through the Management Server.

If you are creating Security Engine elements for the first time, follow the instructions in the *Forcepoint Network Security Platform Installation Guide*.

Security Engine configuration overview

Configuring Engines, IPS engines, and Layer 2 Engines consists of several general procedures.

1) Install a license for the Security Engine.

2) Create an Security Engine element and define the basic properties.

3) Configure the Security Engine interfaces.

4) Configure the routing.

5) Generate the initial configuration for the Security Engine and use it to establish a connection between the engine and the Management Server.

6) Install a policy on the Security Engine.

Related concepts

Creating Security Engine elements on page 504
Network interfaces for Security Engine on page 543
Configuration of Master Engines and Virtual Engines on page 527
Connect Security Engines to the SMC on page 631
Getting started with licenses on page 1331

Related tasks

Install policies on page 813

The Engine Editor and how it works

The Engine Editor combines all Security Engine configuration information into one view that is easy to access and navigate. You can quickly save and validate all changes using the Engine Editor toolbar.

The Engine Editor has the following branches:

- **General** — General Security Engine settings
- **Interfaces** — Interface configuration
- **Routing** — Routing and antispoofing configuration
- **Add-Ons** — Settings related to Security Engine add-ons and additional services.
- **Policies** — Information about the Security Policy used on the Security Engine and settings related to element-based NAT and Aliases. Also contains settings for Automatic rules
- **VPN** (Engines only) — Settings related to the VPN configuration and SSL VPN configuration
- **Advanced Settings** — Advanced Security Engine settings (Traffic Handling, SYN Rate Limits, Log Handling, Scan Detection, DoS Protection, Idle Timeouts, and Tunneling)

The available branches can vary depending on the type of Security Engine.

You can use the Engine Editor to edit or view an Security Engine configuration:

- To edit an Security Engine, right-click the Security Engine, then select **Edit**.
- To view an Security Engine, right-click the Security Engine, then select **Preview**. You can also double-click the Security Engine.

Creating Security Engine elements

You can either create an element or copy and edit an existing element.

Before you create an Security Engine element, make sure that you have a license for it. The element can be configured without a license, but you must have a license to make the Security Engine operational.

You can create Security Engine elements in the following ways:

- Single Engine and Engine Cluster elements one by one
- Several new Single Engine and Engine Cluster elements at the same time with a wizard
- IPS and Layer 2 Engine elements one by one

Create Security Engines

You can create the elements that represent Security Engines in the Security Engines view.

You can create Single Engines, IPS Engines, or Layer 2 Engines. You can also create clustered Engines, IPS Engines or Layer 2 Engines. In a cluster, 2-16 appliances can be used. You can later convert a single Security Engine into a clustered Security Engine.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Select **New** > **Engine** > **Single Engine** or **Engine Cluster**.
3) Configure the settings on the branches of the Engine Editor.
   The branches shown can vary depending on the type of Security Engine.
4) To save your changes, click **Save**.
5) When you have finished making all your changes and want to transfer the configuration to the Security Engine, click **Save and Refresh**.

Related tasks

Create additional Log Server elements on page 473

Creating multiple Engines at the same time using a wizard

You can create multiple Single Engines or Engine Clusters at the same time using a wizard. To simplify the configuration, you can create the Engines based on an existing Engine.

Using the wizard offers several benefits:

- It is easy to create several Engine elements at the same time.
- You can select IP addresses as Endpoints for the VPN Gateway elements that represent the Engines in VPNs.
- You can define a policy that is automatically installed on the Engines when they make initial contact to the Management Server.

The Engine properties you define in the wizard are common to all the Engines you create. Consider which properties all Engines can share and which properties must be defined separately for each Engine. After you have created the Engines, you can change the properties of each individual Engine.

Interfaces

You must define at least one layer 3 physical interface and one IPv4 address for the Engines. Make sure that the IP addresses that are assigned to the Engines are not used by any other components.

To use a Layer 3 Physical Interface for communication with the Management Server, define a Layer 3 Physical Interface with an IP address (dynamic IP address for Single Engines). The Layer 3 Physical Interface is assigned Interface ID 0. When connecting the cables to the appliance, connect the cable for the control connection to Ethernet port 0. See the relevant *Hardware Guide* for detailed information about mapping the Interface IDs with specific ports on the appliances.

Considerations for Single Engines

- You can optionally use the Proof-of-Serial (POS) codes that are delivered with the Forcepoint Network Security Platform appliances to create the Single Engine elements.
- When you use POS codes in the wizard, all appliances must be the same model. If you have POS codes for different types of appliances, you must run the wizard separately for each appliance model to create the elements. Before you create Engine elements, note the serial numbers and geographical locations of the appliances.
- In plug-and-play configuration, the appliances configure themselves automatically after they are plugged in and connected to the network.

Note

Only specific Forcepoint Network Security Platform appliances can use plug-and-play configuration.

Note

There are special considerations when using plug-and-play configuration. For example, both the SMC and the Security Engines must be registered for plug-and-play configuration before you configure the engines. See Knowledge Base article 9662.

Add interfaces in the following order:

1) Layer 3 Physical Interfaces

2) Integrated ADSL modems

Note

ADSL Interfaces are only supported on specific legacy appliances that have an integrated ADSL network interface card.

3) Integrated mobile broadband modems

4) Tunnel Interfaces

5) Integrated wireless modems

6) SSID Interfaces

7) Integrated switches and Port Group Interfaces

Considerations for Engine Clusters

- You cannot use the Proof-of-Serial (POS) codes that are delivered with the Forcepoint Network Security Platform appliances.
- The new Engine Clusters must be based on an existing Engine Cluster.

Add interfaces in the following order:

1) Layer 3 Physical Interfaces

2) Tunnel Interfaces

Start the wizard

Start the wizard to create multiple Engines at the same time. Define the general settings for the new Engines.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Select **New** > **Engine** > **Multiple Single Engines** or **Multiple Engine Clusters**.

3) (Single Engines only) To create multiple Single Engine elements based on POS codes, enter the POS codes in the **Proof-of-Serial Codes** Codes field.

4) If you do not have POS codes, enter the number of Engines to create.
   You can create up to 1000 Engines.

5) Select the Engine on which you want to base the configuration from the **Base Configuration On** drop-down list.
   This step is optional when creating Single Engines.

6) To progress through the Wizard, click **Next**. To go back a page, click **Previous**.
   After most configuration pages there is a page where you can review and edit the configuration. Double-click the fields to directly edit the details.

7) When you have completed the wizard and reviewed the summary page, click **Finish**.

Create a Connection Synchronization Group element

This topic provides information about how to create a connection synchronization group element.

Limitations:

- You can add a maximum of 16 engine nodes that have the **Connection Synchronization for External High Availability** feature enabled in a group.
- This is only supported for SMC managed appliances.

Steps

1) Select **Engine Configuration**.
2) Browse to **Other Elements** > **Engine Properties** > **Connection Synchronization Groups**.
3) Right-click, then select **New Connection Synchronization Group**.
4) Configure the settings, then click **OK**.

Duplicate existing Security Engine elements

If you have similar configurations at several sites, you can duplicate an existing Security Engine element to reduce the need for manual configuration.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click the existing engine element, then select **New** > **Duplicate**.
3) Give the engine element a unique **Name**.
4) Adjust the rest of the properties as needed.
   Change at least the IP address used for communication with the Management Server. (Or configure the duplicated element to use a different control IP address.) The control IP address must be unique for each engine.
5) Click **OK**.
   The new engine element is added to the element tree.

Related concepts

Network interfaces for Security Engine on page 543

Related tasks
Edit individual Security Engines on page 509

Editing existing Security Engines

You can change the properties of an individual Security Engine element or the common properties shared by several Security Engine elements.

Existing Security Engine elements are shown in the **Engine Dashboard** view under specific branches. Engine, IPS, and Layer 2 Engine elements are shown in the **Engine Configuration** view under the **Engines** branch. Changing the common properties of several engine elements restricts you to actions and values that are applicable to all engine elements at the same time. You can also convert a Single Engine, a Single IPS engine, or a Single Layer 2 Engine into a cluster.

Related concepts
Changing control IP addresses on page 520
Configuration of Master Engines and Virtual Engines on page 527

Related tasks
Edit common properties of several Security Engines at the same time on page 510
Convert a Single IPS engine to an IPS Cluster on page 516
Convert a Single Layer 2 Engine to a Layer 2 Engine Cluster on page 517
Add nodes to clusters on page 519

Edit individual Security Engines

You can change the properties specific to one individual engine element, such as IP address definitions, by editing the engine element.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click an engine, then select **Edit <element type>**.
3) Use the navigation pane on the left to find the settings that you want to edit.
4) Edit the configuration information as required.
5) Select from the following:
   - To validate the changes, select **More actions** > **Validate**.
   - To validate and save the changes, click **Save**.

- To validate and save the changes and refresh the security policy on the engine, click **Save and Refresh**.

Note

Validation issues are displayed in the **Issues** pane. Double-click an issue to return to the section in which the issue can be fixed.

Edit common properties of several Security Engines at the same time

You can select several Security Engines and change the properties that are common to all of them.

The following limitations apply when you change the properties of several engines at once:

- Properties specific to one individual Security Engines element, such as IP address definitions, are never available in the common properties.
- If you select both single and clustered Security Engines elements, the cluster-specific options are not available.
- If you select Security Engines of different types, you can only change properties that are supported for all of the selected types of elements.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) **Shift**-select or **Ctrl**-select the Security Engines that you want to edit.

Tip

To have the maximum number of properties available, select Security Engines of the same type.

3) Right-click one of the selected Security Engines, then select **Common Properties**.

4) Configure the settings, then click **OK**.
   The properties that you can configure depend on how similar in type the selected Security Engines are.

Converting Single Engines to Engine Clusters

You can use a conversion tool to change an existing Single Engine into a Engine Cluster.

The conversion tool:

- Maintains the relationship of the Security Engine element with other configurations in the system, such as VPNs
- Allows you to maintain some existing interface configurations, such as VLANs defined on interfaces
- Minimizes service interruptions.

The following limitations apply when you convert Single Engine elements to Engine Cluster elements:

- It is not possible to combine two Single Engine elements into a Engine Cluster.
- A Single Engine can only be converted to a two-node Engine Cluster. If you want to add more nodes to the cluster, you must add the nodes separately after the conversion.

Due to differences in the supported configurations, the following configurations prevent you from converting from a Single Engine to a Engine Cluster:

Unsupported configurations on Engine Clusters

Configuration	Notes
Wireless interfaces	Engine clusters do not support wireless interfaces.
Dynamic IP addresses	Engine clusters can only have static IP addresses. Clusters cannot use a dynamically assigned (DHCP or PPPoE) IP address.
Modem interfaces	Engine clusters do not support integrated mobile broadband modems, such as LTE modems. You must change to a configuration that uses an external mobile broadband modem through an Ethernet connection to convert to a cluster.
Integrated switch	Engine clusters do not support integrated switches. To convert to a cluster, you must change to a configuration that uses an external switch that the engines access through an Ethernet connection.



Note

If you change the control IP address of the existing node after you start the conversion tool, the connection between the engine and the SMC is lost.

The configuration consists of these general steps:

1) Prepare your environment for converting the Single Engine to a Engine Cluster.
2) Prepare interfaces and IP addresses for converting the Single Engine to a Engine Cluster.
3) Convert the Single Engine element to a Engine Cluster element.
4) Activate the new Security Engine configuration on the Engine Cluster.

Prepare the environment for converting a Single Engine to a Engine Cluster

Make sure your environment meets these requirements before you convert a Single Engine to a Engine Cluster.

Steps

1) Select one Single Engine element to convert to a cluster.
2) Make sure that you have a license for each node in the Engine Cluster.

3) If you are not using identical hardware, make sure that the performance levels match your needs.

   You can use hardware with different performance levels for load-balanced clustering. For standby clustering, the performance level of each node must be high enough to handle all traffic.

4) Make sure that enough IP addresses are available in the network, especially if the Single Engine is managed remotely.

   Each node in the Engine Cluster needs at least one dedicated IP address for its management communications. Also, the traffic that the nodes inspect requires at least one dedicated IP address per cluster.

5) Make sure that the Security Engine hardware is running software versions that are compatible with the Security Management Center, and that both Security Engines are running the same version.

6) If the Security Engine hardware that you are adding to the cluster already has a working configuration from previous use, return it to the initial configuration state using the Security Engine Configuration Wizard (sg-reconfigure) on the command line.

   Do not establish a connection with the Management Server before the Engine Cluster element is ready.

   CAUTION

   If the Engine has a working configuration, it goes online and processes traffic when you turn it on to configure it for the Engine Cluster.

7) Connect the network cables to the new Security Engine hardware and turn it on.

Related concepts

Load balancing on page 34
Standby operation on page 34
Getting started with licenses on page 1331

Related tasks

Access the Security Engine command line on page 364
Reconfigure Security Engine settings on page 365
Upgrade Security Engines remotely on page 1354

Prepare interfaces and IP addresses for converting a Single Engine to a Engine Cluster

Engine Clusters have different IP addressing requirements than Single Engines. You must change the IP address that is used for a particular role if the new interface type is not compatible with that role.

Engine Clusters must have two types of IP addresses:

- NDI (Node Dedicated IP Address) — An IP address that is used for traffic to or from an individual node in a cluster. Each node in the cluster has a specific IP address that is used as the NDI.
- CVI (Cluster Virtual IP Address) — An IP address that is used to handle traffic routed through the cluster for inspection. All nodes in a cluster share this IP address. Allows other devices to communicate with the Engine Cluster as a single entity. If other network devices, such as a default gateway or VPN endpoint, select the

engine's IP address, converting the IP address to a CVI allows those external configurations to remain the same.

Interface type requirements by role on Engine Clusters

Role	Type Required	Notes
Control interface (Management connections)	NDI	Each node requires its own NDI address. Often, the same IP address on a Single Engine is used for both the engine's own communications and the traffic that the engine processes. In these cases, you can convert the IP address that processes the traffic to a CVI. With the conversion, you can avoid reconfiguring external equipment and you can add new NDI addresses for the nodes. Make sure that enough IP addresses are available in the network, especially if the Single Engine is managed remotely.
DHCP relay	CVI	Configured in the physical interface properties.
DHCP relay for VPN clients	NDI	Configured in the VPN settings in the Engine Editor.
Heartbeat interface	NDI	Heartbeat and state synchronization communications between clustered engines. We recommend using a dedicated interface for the heartbeat, as reliable transmissions are critical to the operation of the cluster. If the heartbeat traffic passes through a switch, make sure that the switch does not throttle or block multicast traffic between the clustered engines.
Routing	CVI	Traffic that is sent to an NDI address is not routed to any other destinations. Surrounding network devices that use the engine as a default gateway must use a CVI address. If the internal DHCP server is used and configured to assign the engine as the default gateway for clients, the default gateway IP address must be a CVI. (Configure the CVI in the physical interface properties.)
VPN endpoints	CVI	Configured in the VPN settings in the Engine Editor.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) If you plan to convert the Single Engine's existing IP address for management connections to a CVI that processes the traffic in the Engine Cluster, configure a new IP address for management connections.
   a) In the properties of the Single Engine element, add a new IP address to the interface that you want to use for management connections.
   b) In the interface options, select the new control IP address from the **Backup** control IP address drop-down list.
      Management communication is not yet allowed to the new IP address. Adding the new IP address as a backup control IP address prevents management communication from being interrupted.
   c) Edit the Access and NAT rules of any engines on the communications path so that both current and new control IP addresses are allowed, then refresh the policies of these engines.
   d) Refresh the policy of the Single Engine that you plan to convert.

e) Deselect the new control IP address from the **Backup** control IP address drop-down list.

f) Select the new control from the **Primary** control IP address drop-down list.

2) Add any new IP addresses that are required for the selected interface roles and configure the settings to use those IP addresses.
Make sure to add IP addresses to all physical interfaces or VLAN interfaces.

3) If configured, remove dynamic IP addresses, modem interfaces, ADSL interfaces, integrated switches, and port group interfaces. These configurations are not supported on clusters.

Related tasks
Add IPv4 addresses to Single Engine interfaces on page 562
Select system communication roles for engine interfaces on page 570
Add Access rules on page 899
Add NAT rules on page 903

Convert a Single Engine to a Engine Cluster

Use the conversion tool to change an existing Single Engine into a Engine Cluster.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click the Single Engine element, then select **Configuration** > **Upgrade to Cluster**.
An interface mapping dialog box opens.

2) Click the **Upgrade to** cell for each interface and select the IP address types for the interfaces.
You can create a CVI and an NDI for the same physical interface. This configuration is recommended for all interfaces. More IP addresses are generated automatically to create the CVIs and NDIs.

Note
Each selection is validated and you might not be able to select a type if it is incompatible with the selected role of the interface.

3) Click **OK**.
The properties dialog box for the new Engine Cluster element opens.

4) On the Interfaces tab, add the interfaces and addresses needed for the cluster.
Make sure that the IP addresses on all interfaces are unique and unassigned, and change them if necessary.

5) Select **Packet Dispatch** as the CVI mode and enter the related unicast MAC address in the properties of all physical interfaces.

6) Click **Options**, then define which IP addresses are used in particular roles in system communications.

7) If the internal DHCP server is configured to assign the engine as the default gateway for clients, verify that the default gateway IP address is a CVI on the **DHCP** tab of the **Physical Interface Properties** dialog box.

8) Click **OK**.
   The Single Engine element is converted to a Engine Cluster.

Activate the Security Engine configuration after converting a Single Engine to a Engine Cluster

You must activate the new configuration to finish converting a Single Engine element to a Engine Cluster element.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) If any external device uses the engine as a default gateway or VPN endpoint and the previously used IP address is converted to an NDI, reconfigure the external equipment to reference a CVI address.

2) Run the Security Engine Configuration Wizard on the command line (sg-reconfigure) or in a web browser.

3) Make sure the interface IDs are mapped to the correct network ports on the hardware.

4) Make initial contact between the Security Engine nodes and the Management Server.
   Install and configure any new Security Engine nodes as part of the cluster in the same way as in a new installation. See the *Forcepoint Network Security Platform Installation Guide*.

5) Install the policy on the cluster.
   If any new nodes have not yet been initialized, set the inactive nodes to disabled before you refresh the policy of the existing node. Otherwise, the policy installation fails due to a lack of connectivity to all nodes.

Next steps

If there are problems with the clustered configuration, you can return to single-node operation. To do so, command one node offline through the right-click menu or through the command line.

Related concepts
Connect Security Engines to the SMC on page 631

Related tasks
Disable cluster nodes temporarily on page 360

Convert a Single IPS engine to an IPS Cluster

You can convert an existing Single IPS element to an IPS Cluster element.

Converting a Single IPS element to an IPS Cluster element maintains the relationship of the Security Engine element with other configurations in the system. The conversion requires you to select one Single IPS element to convert to an IPS Cluster.

The following limitations apply when you convert a Single IPS to an IPS Cluster:

- It is not possible to combine two Single IPS elements into an IPS Cluster element.
- A Single IPS engine can only be converted to a two-node IPS Cluster. If you want to add more nodes to the cluster, you must add the nodes separately after the conversion.

CAUTION

If you change the control IP address of the existing node in this process, the connection between the Security Engine and the SMC is lost.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Make sure that both Security Engines are licensed.

   The licensing of clustered Security Engine nodes is done in the same way as the licensing of two Single IPS engines. All current IPS engine licenses allow clustering the nodes, so no license changes are required to activate the feature.

2) Make sure that the Security Engines are running software versions that are compatible with the Security Management Center, and preferably that both Security Engines are running the same version.

   Although the cluster can be installed with the Security Engines running different software versions (unless otherwise stated in the Release Notes), long-term use with mismatched versions is not supported.

3) If the new IPS engine you want to add to the IPS Cluster already has a working configuration from previous use, return it to the initial configuration state.

   You can do so in the Security Engine Configuration Wizard (sg-reconfigure) on the command line.

   Note

   Do not establish a connection with the Management Server before the IPS Cluster element is ready.

4) Connect the network cables to the new node and power it on.

5) Right-click the Single IPS element that you want to upgrade to an IPS Cluster, then select **Configuration** > **Upgrade to Cluster**.

6) Browse to **Interfaces** > **Interface Options**.

7) Define which IP addresses are used in particular roles in system communications.

8) Click **Save**.

Note

You can still close the Engine Editor without saving the changes to return to the previous configuration and undo the conversion.

9) Make initial contact between each node and the Management Server.

Install and configure any new Security Engine nodes as part of the cluster as in a new installation.

10) Install the policy on the IPS Cluster.

To refresh the policy of the existing node before the new nodes are initialized, disable the inactive nodes on the **Clustering** pane in the Engine Editor. Otherwise, the policy installation fails due to a lack of connectivity to all nodes.

Related concepts

Considerations for working on the Security Engine command line on page 363
Changing control IP addresses on page 520
Connect Security Engines to the SMC on page 631
Getting started with licenses on page 1331
Getting started with upgrading Security Engines on page 1351

Related tasks

Add nodes to clusters on page 519
Select system communication roles for IPS interfaces on page 584

Convert a Single Layer 2 Engine to a Layer 2 Engine Cluster

You can use a conversion tool to convert an existing Single Layer 2 Engine to a Layer 2 Engine Cluster.

Using the conversion tool maintains the relationship of the Single Layer 2 Engine element and other configurations in the system. The conversion requires you to select one Single Layer 2 Engine element to convert to a Layer 2 Engine Cluster.

The following limitations apply when you convert a Single Layer 2 Engine to a Layer 2 Engine Cluster:

- It is not possible to combine two Single Layer 2 Engine elements into a Layer 2 Engine Cluster element.
- A Single Layer 2 Engine can only be converted to a two-node Layer 2 Engine Cluster. To add more nodes to the cluster, add the nodes separately after the conversion.

CAUTION

If you change the control IP address of the existing node in this process, the connection between the Security Engine and the SMC is lost.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Make sure that both Security Engines are licensed.

   The licensing of clustered Security Engine nodes is done in the same way as the licensing of two Single Layer 2 Engines. All current Layer 2 Engine licenses allow clustering the nodes, so no license changes are required to activate the feature.

2) Make sure that the Security Engines are running software versions that are compatible with the SMC, and preferably that both Security Engines are running the same version.

   Although the cluster can be installed with the Security Engines running different software versions (unless otherwise stated in the Release Notes), long-term use with mismatched versions is not supported.

3) If the new Layer 2 Engine has a working configuration from previous use, return it to the initial configuration state.

   You can do so in the Security Engine Configuration Wizard (sg-reconfigure) on the command line.

   Note

   Do not establish a connection with the Management Server before the Layer 2 Engine Cluster element is ready.

4) Connect the network cables to the new node and power it on.

5) Right-click the Single Layer 2 Engine element that you want to upgrade to a Layer 2 Engine Cluster, then select **Configuration** > **Upgrade to Cluster**.

6) Browse to **Interfaces** > **Interface Options**.

7) Define which IP addresses are used in particular roles in system communications.

8) Click **Save**.

   Note

   You can still close the Engine Editor without saving the changes to return to the previous configuration and undo the conversion.

9) Make initial contact between each node and the Management Server.

   Install and configure any new Security Engine nodes as part of the cluster as in a new installation.

10) Install the policy on the Layer 2 Engine Cluster.

    To refresh the policy of the existing node before the new nodes are initialized, disable the inactive nodes on the **Clustering** pane in the Engine Editor. Otherwise, the policy installation fails due to a lack of connectivity to all nodes.

Related concepts

Considerations for working on the Security Engine command line on page 363
Changing control IP addresses on page 520
Connect Security Engines to the SMC on page 631
Getting started with licenses on page 1331
Getting started with upgrading Security Engines on page 1351

Related tasks

Select system communication roles for Layer 2 Engine interfaces on page 596

Add nodes to clusters

By default, the Engine Editor displays two nodes in the **Clustering** pane. You can add new nodes to the cluster.

Before you begin

Before adding a node to the configuration, install the physical Security Engine device and connect the cables. At a minimum, connect the cables for the interfaces that enable communication with the Management Server and between the clustered Security Engines.

If the device has a working configuration from previous use, use the Security Engine Configuration Wizard (sg-reconfigure) on the command line to return the device to the initial configuration state. Set up the initial configuration state before connecting the device to the network. Do not make initial contact with the Management Server.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click a Security Engine, then select **Edit <element type>**.

3) Browse to **General** > **Clustering**.

4) Click **Add Node**.

5) Give the node a unique name.

6) In the **Nodes** table, review and update the information in the **IP Address**, **Contact IP Address**, and **Comment** columns for each NDI.

   Double-click the value you want to change.

   Engine CVI details are identical between the nodes, so adding a node to a Engine Cluster does not require changing the CVI configuration in any way.

7) Click **OK**.

8) Click **Save**.

9) Save the initial configuration for the new Security Engine node to create a one-time password.

10) Make initial contact between the new Security Engine and the Management Server.

11) Refresh the policy to transfer the changes to the Security Engines.

To refresh the policy of the existing node before the new nodes are initialized, disable the inactive nodes on the **Clustering** pane in the Engine Editor. Otherwise, the policy installation fails due to a lack of connectivity to all nodes.

Related concepts

Connect Security Engines to the SMC on page 631

Changing control IP addresses

You can change the control IP address of an Security Engine without losing management connectivity.

When you change IP addressing, other connections between the different components might be temporarily lost. You must make sure that the connections return to normal after the IP address changes.

Note

We recommend that you do not use the IP address of an Aggregated Link interface as the primary or secondary control IP address of the Engine.

Related tasks

Change the control IP addresses within the same network on page 520
Change control IP addresses to a different network on page 521

Change the control IP addresses within the same network

You can change the control IP address of an Security Engine to a new address that belongs to the same network as the old address.

The new control IP addresses of IPS engines and Layer 2 Engines must always belong to the same network as the existing control IP addresses. If management connectivity is no longer needed, change the control IP address in the SMC and reinitialize the Security Engine through the command line using a new one-time password.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) If you have an IP-address-bound license for the Security Engine, request a new Management Server POL code bound license at https://stonesoftlicenses.forcepoint.com.

This change is required, because IP-address-bound licenses are no longer supported.

2) Install and bind the new license to the Security Engine.

3) In the Engine Editor, create an interface for the new IP address and set the address as the backup control IP address.

4) Install the policy on the Security Engine.
From this point on, you can start using the new address in the network.

5) In the Engine Editor, set the old and new control IP addresses as the backup and primary control IP addresses, respectively.

Note

If your Security Engine cannot use the old and new control IP addresses simultaneously, remove the old control IP address from the **Interfaces** pane in the Engine Editor. Also remove the corresponding network from the **Routing** pane in the Engine Editor.

6) Click **Save and Refresh**.

7) Remove the old control IP address from the **Interfaces** pane and the **Routing** pane in the Engine Editor.

8) Click **Save and Refresh** again.

Note

If the connection with the Management Server is lost while you try to change IP addressing, run the Security Engine Configuration Wizard (sg-reconfigure) on the Security Engine command line. This action returns the Security Engine to the initial configuration state and re-establishes initial contact between the Security Engine and the Management Server.

Related concepts

Network interfaces for Security Engine on page 543
Management connections for Security Engines and how they work on page 629
Connect Security Engines to the SMC on page 631
Getting started with licenses on page 1331

Change control IP addresses to a different network

You can change the control IP address of an Security Engine to a new IP address in a different network than the old one.

Because these steps require the configuration of Outbound Multi-Link, you can only change the control IP address of Engines to a different network. For all other Security Engine roles, you must change the IP address within the same network.

If management connectivity is no longer needed, change the control IP address in the SMC and reinitialize the Security Engine through the command line using a new one-time password.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) If you have an IP-address-bound license for the Security Engine, request a new Management Server POL code bound license at https://stonesoftlicenses.forcepoint.com.
   This change is required, because IP-address-bound licenses are no longer supported.

2) Install and bind the new license to the Security Engine.

3) Edit the Single Engine or Engine Cluster element in the Engine Editor and add an interface.
   - Define the new primary control address as the backup control IP address.
   - If your engine is a cluster and you do not want to lose any connections, also define a new CVI for the cluster.

4) Configure Outbound Multi-Link.
   Create two NetLinks: one for the old control IP address and one for the new control IP address.

5) Install the policy on the Security Engine.
   From this point on, you can start using the new address in the network.

6) To set the new and old control IP addresses as the primary and backup IP addresses, respectively, edit the Single Engine or Engine Cluster element in the Engine Editor.

   Note

   If your Security Engine cannot use the old and new control IP addresses simultaneously, remove the interface with the old control IP address from the **Interfaces** pane in the Engine Editor. Also remove the elements and rules you created for the Multi-Link configuration.

7) Click **Save and Refresh**.

8) Remove the interface with the old control IP address from the **Interfaces** pane in the Engine Editor.

9) Remove the elements and rules you created for the Multi-Link configuration.

10) Click **Save and Refresh** again.

   Note

   If the connection with the Management Server is lost while you try to change IP addressing, run the Security Engine Configuration Wizard (sg-reconfigure) on the Security Engine command line. This command returns the Security Engine to the initial configuration state and re-establishes initial contact between the Security Engine and the Management Server.

Related concepts
Editing existing Security Engines on page 509
Network interfaces for Security Engine on page 543
Management connections for Security Engines and how they work on page 629
Connect Security Engines to the SMC on page 631
Getting started with outbound traffic management on page 731
Getting started with licenses on page 1331

Configure Layer 2 Settings for Security Engines in the Engine/VPN role

Layer 2 Settings for Security Engines in the Engine/VPN role define the Layer 2 Interface Policy for the Security Engine, and advanced settings for layer 2 physical interfaces on the Security Engine.

Layer 2 Interface Policies contain rules for traffic detected by layer 2 physical interfaces on Security Engine in the Engine/VPN role. To use layer 2 physical interfaces, you must select the Layer 2 Interface Policy for the Security Engine. All layer 2 physical interfaces on the Security Engine use the same Layer 2 Interface Policy.

Note

When an Engine is configured with Layer 2 inline or capture interfaces, the Layer 2 Interface Policy does not have an option to select an Inspection Policy or File Filtering Policy. Instead, the Inspection Policy and File Filtering Policy options that you select in the Inspection tab on the Engine Policy page is used to process the traffic on the Layer 2 interfaces. For more details, refer to the Which Inspection and File Filtering Policy Is Used for Layer 2 Interfaces Knowledge Base Article.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click a Security Engine, then select **Edit <element type>**.
3) Browse to **General** > **Layer 2 Settings**.
4) From the **Policy for Layer 2 Interfaces** drop-down list, select the Layer 2 Interface Policy.
5) (Optional) Configure one or more of these advanced settings for layer 2 physical interfaces:
   - Connection tracking options
   - Bypass options for Capture Interfaces and Inline IPS Interfaces
6) Click **Save and Refresh** to transfer the changes.

Configure a backup unit or a connection synchronization for external high availability

You can now for an engine, configure a backup unit or configure a connection synchronization for external high availability.

Limitations:

- You cannot enable both the **Backup Unit** and the **Connection Synchronization for External High Availability** feature at the same time.
- This is only supported for SMC managed appliances.

The **Backup Unit** feature is a high availability functionality that is used to manage small remote sites where engine clustering is not possible or only a single IP address can be used. This functionality allows synchronization of the connection between two engine nodes and only one node is active at a time. Hence, if one of the engine nodes is unavailable the engine operations are not interrupted as the second engine node becomes active.

The **Connection Synchronization for External High Availability** feature is used to enable high availability in Public Clouds by allowing synchronization of connections between 2 to 16 single engine configurations.

Steps

1) Select **Engine Configuration**.
2) Browse to **Engine** > **Engines**.
3) Right-click an engine, then select Edit <element type>.
4) In the navigation pane on the left, browse to **Advanced Settings** > **High Availability**.
5) Configure the settings.
6) Click **Save**.

Configure global contact policy settings for node-initiated contact to the Management Server

Communication between the Management Server and an Security Engine, Master Engine, or Virtual Security Engine can be reversed. In this case, the Security Engine opens a connection to the Management Server and keeps it open to wait for any commands.

Before you begin

An Security Engine, Master Engine, or Virtual Security Engine with a dynamic Control IP Address has been configured.

Reversing communication might be necessary in the following cases:

- The Security Engine does not have a static IP address that the Management Server can contact. For example, instead of a static IP address, the Security Engine has a dynamic IP address on the control interface or there is intermediate dynamic NAT.
- The Management Server's connections are blocked because of a traffic filtering device between the components.

The settings for communication between the Management Server and the engines are set in the SGConfiguration.txt file stored on the Management Server. You can either use the default values for each setting or change the settings by adding parameters and values to the SGConfiguration.txt file.

Steps

1) On the Management Server computer, browse to the <installation directory>/data directory.



Note

If you installed the Management Server in the `C:\Program Files\Forcepoint\SMC` directory in Windows, some program data might be stored in the `C:\ProgramData\Forcepoint\SMC` directory.

2) Edit the SGConfiguration.txt file and add the following parameters as needed.

SGConfiguration parameters

Parameter name	Description
`DCP_INITIAL_DELAY`	Time (in seconds) to wait after initialization before the first connection attempt to the Management Server. The default value is 5 seconds.
`DCP_CONNECTION_INTERVAL`	Time (in seconds) to wait before connecting again to the Management Server after a successful connection. The default value is 25 seconds.
`DCP_RETRY_INTERVAL`	Time (in seconds) to wait before connecting again to the Management Server after a failed connection attempt. The default value is 25 seconds.
`DCP_IDLE_TIMEOUT`	Time (in seconds) before an idle connection is closed. The default value is 1800 seconds (30 minutes).

3) Save and close the file.

4) Refresh the policies of the engines to transfer the changes.

Related tasks

Select system communication roles for engine interfaces on page 570
Select system communication roles for IPS interfaces on page 584
Select system communication roles for Layer 2 Engine interfaces on page 596

Synchronizing the time on Security Engines

The time on Security Engines is automatically synchronized to match the time of the Management Server. You can optionally use an NTP server to synchronize the time.

By default, Security Engines get time setting commands from the Management Server. If an Security Engine is configured to use NTP and it can successfully get the time from an external NTP server, the Security Engine ignores time setting commands from the Management Server.

In environments where there are Master Engines and Virtual Engines, you can use NTP servers to synchronize engine times only for Master Engines. Virtual Engines do not communicate directly with NTP servers.

If the Management Server, Log Server, and the engines do not have the same time, there might be problems with logging and monitoring. Also make sure that the computer you are using for SMC Client access has the time and time zone set correctly. Correct settings prevent time synchronization problems when you view statistics or logs, generate reports, or schedule automatic maintenance tasks.

Related tasks

Enable NTP time synchronization for Security Engines on page 146

Chapter 30

Creating and modifying Master Engine and Virtual Engine elements

Contents

- Configuration of Master Engines and Virtual Engines on page 527
- Create Master Engines or Virtual Engines on page 530
- Create Virtual Resource elements on page 531
- Add additional nodes to Master Engines on page 532
- Moving a Virtual Engine to a different Master Engine on page 533
- Convert Engines to Master Engines and Virtual Engine elements on page 533
- Example: deploying Virtual Engines for MSSP customers on page 541

Virtual Engines are logically separate Security Engines that run as virtual instances on a physical Security Engine appliance. A Master Engine is a physical appliance that provides resources for Virtual Engines.

Configuration of Master Engines and Virtual Engines

Master Engines are physical devices that provide resources for multiple Virtual Engines.

Using Virtual Engines allows the same physical engine device to support multiple policies or routing tables, or policies that involve overlapping IP addresses. This is especially useful in a Managed Security Service Provider (MSSP) environment, or in a network environment that requires strict isolation between networks.

A *Virtual Resource* element defines the set of resources on the Master Engine that are allocated to a Virtual Engine. Virtual Resource elements associate Virtual Engines with Physical Interfaces or VLAN Interfaces on the Master Engine.

Virtual Engines associated with the same Master Engine can belong to different administrative Domains. However, the Master Engine must either belong to the Shared Domain or to the same Domain as the associated Virtual Engines. For example, the Master Engine can belong to the Shared Domain, while each associated Virtual Engine belongs to a different Domain.

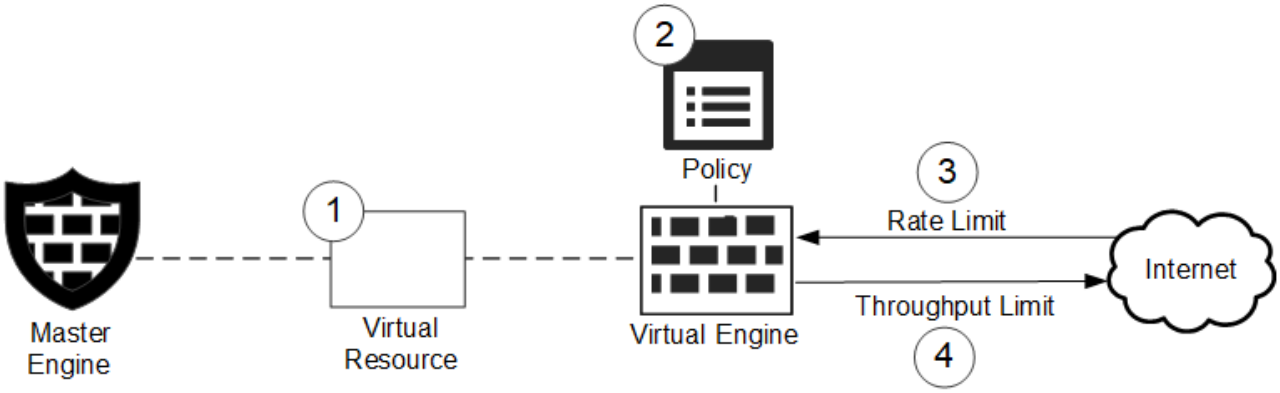
Any Security Engine that has a license that allows the creation of Virtual Resources can be used as a Master Security Engine.

Before you define a new Master Engine element, make sure that you have an Security Engine license for each Master Engine node. Virtual Engines do not require individual licenses. Instead, the Security Engine license for the Master Engine defines how many Virtual Resources can be created. The number of Virtual Resources limits the number of Virtual Engines: one Virtual Engine at a time can be associated with each Virtual Resource.

Protecting Virtual Engines

In the Virtual Resource, you can set the rate limit and throughput limit for the Virtual Engine. Setting the rate limit helps protect the other Virtual Engines by ensuring that a single Virtual Engine does not consume all the resources of a Master Engine.

Example of using rate limit and throughput limit for a Virtual Engine



1. In the properties of the Master Engine, open the Virtual Resource, then set the limits for the rate limit or throughput limit, or for both.
2. Refresh the policy on the Virtual Engine.
3. When incoming network traffic exceeds the rate limit, the packets are dropped. If a rate limit is defined, the limit must be much higher than the throughput limit.
4. When outgoing network traffic exceeds the throughput limit, the packets are queued. If there is a QoS Policy set for the Virtual Engine, the policy handles the prioritization as normal.

Limitations

The following limitations apply to Master Engines and Virtual Engines:

- To use more than one Virtual Engine role, you must create a separate Master Engine for each Virtual Engine role. Each Master Engine must be on a separate physical Master Engine device.
- Virtual Engines do not support dynamic IP addresses or Wireless Interfaces.
- If there are multiple administrative Domains, the Master Engine must either belong to the Shared Domain or to the same Domain as the Virtual Engines.
- Virtual Engines handle only the traffic routed through the Virtual Engine for inspection. All other traffic, including communication between the Virtual Engines and the SMC components, is proxied by the Master Engine. Virtual Engines do not communicate directly with other Virtual Engines.

Master Engine and Virtual Engine configuration overview

Master Engine and Virtual Engine configuration consists of creating Master Engines and associating Virtual Engines with the Master Engines.

By default, a Master Engine element has placeholders for two nodes when the element is created. A Master Engine can have 1–16 nodes. If you do not need to use clustering on the Master Engine, you can remove one of the automatically created nodes.

Note

All Virtual Engines on the same Master Engine must have the same Virtual Engine role (Engine/VPN role, IPS, or Layer 2 Engine). To use more than one Virtual Security Engine role, you must create a separate Master Engine for each Virtual Engine role. Each Master Engine must be on a separate physical Master Engine appliance.

The configuration consists of the following general steps:

1) Generate and install Security Engine licenses for the Master Engine.
2) Create a Master Engine element.
3) Create a Virtual Resource element.
4) Configure Physical or VLAN Interfaces for the Master Engine and assign Virtual Resources to the interfaces.
5) Create Virtual Engine elements.
6) Configure Physical, VLAN, or Tunnel Interfaces for the Virtual Engines.
7) Configure routing for the Master Engine and for Virtual Engines.

   Note

   You cannot configure routing for Virtual IPS engines or Virtual Layer 2 Engines.

8) Install or refresh the policy on the Master Engine to transfer changes to the Master Engine's Physical/VLAN Interfaces and the mapping of Virtual Engines to Master Engine Interfaces.
9) Install or refresh the policy on the Virtual Engines.

Related concepts

Getting started with licenses on page 1331
Adding routes for Master Engines and Virtual Engines on page 695

The Engine Editor and how it works

The Engine Editor combines all Security Engine configuration information into one view that is easy to access and navigate. You can quickly save and validate all changes using the Engine Editor toolbar.

The Engine Editor has the following branches:

- **General** — General Security Engine settings
- **Interfaces** — Interface configuration
- **Routing** — Routing and antispoofing configuration
- **Add-Ons** — Settings related to Security Engine add-ons and additional services.
- **Policies** — Information about the Security Policy used on the Security Engine and settings related to element-based NAT and Aliases. Also contains settings for Automatic rules
- **VPN** (Engines only) — Settings related to the VPN configuration and SSL VPN configuration
- **Advanced Settings** — Advanced Security Engine settings (Traffic Handling, SYN Rate Limits, Log Handling, Scan Detection, DoS Protection, Idle Timeouts, and Tunneling)

The available branches can vary depending on the type of Security Engine.

You can use the Engine Editor to edit or view an Security Engine configuration:

- To edit an Security Engine, right-click the Security Engine, then select **Edit**.
- To view an Security Engine, right-click the Security Engine, then select **Preview**. You can also double-click the Security Engine.

Create Master Engines or Virtual Engines

A Master Engine is a physical engine device that provides the resources for Virtual Engines. One physical Master Engine can support multiple Virtual Engines.

Selecting a Virtual Resource for the Virtual Engine automatically adds the Virtual Engine to the Master Engine where the Virtual Resource is used.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Select from the following:
   - **New** > **Other** > **Master Engine**.
   - **New** > **Engine** > **Virtual Engine**.

3) If creating a Master Engine, select the role for the Virtual Engines that this Master Engine will host, then click **OK**.

Note

You cannot change the role after you create the Master Engine element.

4) Configure the settings on the branches of the Engine Editor.
   The branches shown can vary depending on the type of Security Engine.

5) To save your changes, click **Save**.

6) When you have finished making all your changes and want to transfer the configuration to the Security Engine, click **Save and Refresh**.

Related concepts
Configuring interfaces for Virtual Engines on page 609

Create Virtual Resource elements

Virtual Resources associate Virtual Engines with Physical Interfaces or VLAN Interfaces on the Master Security Engine.

When you select the same Virtual Resource for a Physical Interface or VLAN Interface on the Master Engine and for a Virtual Security Engine, the Virtual Engine is automatically associated with the Master Security Engine. Create one Virtual Resource for each Virtual Engine that you plan to add.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click the Master Engine element, then select **Edit Master Engine**.

3) Browse to **Interfaces** > **Virtual Resources**.

4) Click **Add**.

5) Configure the settings.

6) Click **OK**.

7) Click **Save**.

Next steps

Continue the configuration in one of the following ways:

- Configure Master Engine interfaces.
- Associate the Virtual Resource with a Master Engine interface and with a Virtual Engine.

Related tasks
Add physical interfaces for Master Engines on page 603

Add additional nodes to Master Engines

By default, the Master Engine element has two nodes on the **Clustering** pane of the Engine Editor. You can add new nodes to the Master Engine element. Each Master Engine supports up to 16 nodes.

Before you begin

Install the additional physical engine device. Connect the cables for at least the interface for communication with the Management Server and the interface for communications between the clustered engines. If the device already has a working configuration from previous use, return it to the initial configuration state in the Security Engine Configuration Wizard on the command line before connecting it to the network. Do not make initial contact with the Management Server before you add the node to the Master Security Engine configuration in the SMC Client.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click a Security Engine, then select **Edit <element type>**.
3) Browse to **General** > **Clustering**.
4) Click **Add Node**.
5) Give the node a unique **Name**.
6) In the **Nodes** table, double-click the **IP Address**, **Contact IP Address**, or **Comment** fields if you want to change the values for each NDI.
7) Click **OK**.
8) Click **Save**.
9) Save the initial configuration for the new Master Engine node to create a one-time password.
10) Make initial contact between the new Master Engine and the Management Server.
11) Refresh the policy of the Master Engine to transfer the changes to the Security Engines.

Related concepts

Connect Security Engines to the SMC on page 631

Related tasks

Reconfigure Security Engine settings on page 365

Moving a Virtual Engine to a different Master Engine

The Virtual Resource selected in the properties of a Virtual Engine determines the Master Engine to which the Virtual Engine belongs.

To move a Virtual Engine to a different Master Engine, you select a Virtual Resource that is associated with a different Master Engine in the Virtual Engine's properties. The move becomes effective when you refresh the policy on the Master Engine.

Convert Engines to Master Engines and Virtual Engine elements

You can use a conversion tool to change a Single Engine or Engine Cluster element to a Master Engine and Virtual Engine elements.

The Master Engine must always be based on a Engine. The Virtual Engines can only be in the Virtual Engine role. Only Single Engines or Engine Clusters that are running on 64-bit hardware can be converted to Master Engines and Virtual Engines.

Single Engines and Engine Clusters with the following configurations cannot be converted to Master Engines and Virtual Engines:

- Engine Clusters that have a CVI on the Control Interface
- Engine Clusters that have a CVI on the Heartbeat Interface
- Single Engines that have ADSL Interfaces
- Single Engines that have Wireless Interfaces
- Single Engines that have a dynamic IP address
- Single Engines or Engine Clusters that have a dynamic Contact Address
- Single Engines or Engine Clusters that have DHCP settings on the interface for communication with the Management Server

Related concepts

Defining Host elements on page 922

Start the Convert Engine to Master Engine and Virtual Engines wizard

Start the conversion tool and define general properties for the Master Engine and Virtual Engines.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click a Single Engine or Engine Cluster and select **Configuration** > **Convert to Master Engine and Virtual Engine(s)**.
   The **Convert Engine to Master Engine and Virtual Engines** wizard starts.

3) (Optional) Select a Engine on which to base the configuration from the **Base Configuration On** list.

4) Enter the **Number of Virtual Engines** to create.
   The specified number of Virtual Resources are added to the table, and a Virtual Engine is associated with each Virtual Resource.

5) (Recommended) Double-click the **Virtual Resource Name** field and edit the automatically generated Virtual Resource Name for each Virtual Engine.

6) (Recommended) Double-click the **Virtual Engine Name** field and edit the automatically generated Virtual Engine Name for each Virtual Engine.

7) Click **Next**.
   The **Define Basic Information for the Master Engine** page opens.

8) Enter a **Name** for the Master Engine.
   The name is also used to automatically generate the names of the nodes.

9) Select the **Log Server** to which the Master Engine sends its log data.

10) (Optional) In **DNS IP Addresses** field, add one or more DNS IP addresses.
    DNS IP addresses are IP addresses of external DNS servers. Master Engines use these DNS servers to resolve Domain names to IP addresses. Master Engines need DNS resolution to contact services that are defined using URLs or domain names, and to resolve fully qualified domain names (FQDNs) used in policies.
    - To enter a single IP address manually, click **Add** and select **IP Address**. Enter the IP address.
    - To define an IP address using a network element, click **Add** and select **Network Element**.

11) Select the **Location** for this Master Engine if there is a NAT device between this Master Engine and other SMC components.

12) Define other settings according to your environment:
    - To include the Master Engine in predefined categories, select the appropriate **Categories**.

- To add custom commands to the Master Engine's right-click menu, add a **Tools Profile**.
- Add, edit, or remove nodes.

13) Click **Next**.
The **Define Interfaces for the Master Engine** page opens.

Define interfaces for converted Master Engine elements

All Master Engine interfaces must either have Interface Options defined or be associated with a Virtual Resource.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) If you want to use a Physical Interface to host a Virtual Engine, right-click a Physical Interface and select **Edit Physical Interface**.

CAUTION

The Control Interface is used for Master Engine communications. Do not add a Virtual Resource to the Physical Interface that is used as the Control Interface.

2) Select the **Virtual Resource** to associate with the interface.

3) (Optional) Select **Allow VLAN Definition in Virtual Engine** to allow VLAN Interfaces to be added to the automatically created Physical Interfaces in the Virtual Engine that is associated with this interface.

4) Select the **Virtual Engine Interface ID**.
This is the Physical Interface ID of the Virtual Engine that is associated with this interface.

5) Click **OK** to close the **Physical Interface Properties** dialog box.

6) (Optional) Click **Options** to select which IP addresses are used in particular roles in system communications.

7) (Optional) Add **ARP Entries**.

8) (Optional) Click **Virtual Resources** if you want to edit the name of the Virtual Resources.

9) Click **Next** and continue the configuration in one of the following ways:
   - If Tunnel Interfaces are defined for the Single Engine or Engine Cluster on which the configuration is based, the **Distribute Tunnel Interfaces to Virtual Engines** page opens. Move the Tunnel Interfaces into the interface configuration for the Virtual Engines.
   - If the Single Engine or Engine Cluster on which the configuration is based is associated with a VPN Gateway element, the **Review Distribution of Internal Gateways to Virtual Engines** page opens. Move the VPN Gateway elements into the configuration for the Virtual Engines.

- Otherwise, the **Define Routing for the Master Engine** page opens. Define routing for the converted Master Engines.

Distribute Tunnel Interfaces to converted Virtual Engine elements

If Tunnel Interfaces are defined for the Single Engine or Engine Cluster element on which the configuration is based, you must move the Tunnel Interfaces into the interface configuration for the Virtual Engine elements.

Master Engines cannot use Tunnel Interfaces.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Review the distribution of Tunnel Interfaces to Virtual Engines.

2) (Optional) If you want to change how the Tunnel Interfaces are distributed to Virtual Engines, click the **Virtual Engine** field for a Tunnel Interface and select the correct Virtual Engine from the list.

3) Click **Next** and continue the configuration in one of the following ways:
   - If the Single Engine or Engine Cluster on which the configuration is based is associated with a VPN Gateway element, the **Review Distribution of Internal Gateways to Virtual Engines** page opens. Move the VPN Gateway elements into the configuration for the Virtual Engines.
   - Otherwise, the **Define Routing for the Master Engine** page opens. Define routing for the converted Master Engines.

Distribute VPN gateways to converted Virtual Engine elements

If the Single Engine or Engine Cluster element on which the configuration is based is associated with a VPN Gateway element, you must move the VPN Gateway configuration to the Virtual Engines.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Review the distribution of VPN Gateways to Virtual Engines.

2) (Optional) If you want to change how the VPN Gateways are distributed to Virtual Engines, click the **Virtual Engine** field for a VPN Gateway. Then select the correct Virtual Engine from the list.

3) Click **Next**.

Define routing for converted Master Engine elements

Routes to directly connected networks are automatically added. You must add a default route and any routes through next-hop gateways to networks that are not directly connected to the Master Engine.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Review and edit the routing.

2) Click **Next**.
The **Select Additional Configuration Options** page opens.

Related concepts
Adding routes for Master Engines and Virtual Engines on page 695

Select additional configuration options for converted Master Engine elements

You can define additional properties for the Master Engine on the **Select Additional Configuration Options** page.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Define Additional Master Engine Properties** and click **Next**.
The **Define Tester Settings for the Master Engine** page opens.

2) (Optional) Define tester settings.

3) Click **Next**.
The **Define Permissions for the Master Engine** page opens.

4) (Optional) Define permissions.

5) Click **Next**.
The **Define Advanced Settings for the Master Engine** page opens.

6) (Optional) Define advanced settings.

7) Click **Next**.
The **Review Basic Information for Virtual Engines** page opens.

Define NTP settings for converted Master Engines

You can configure Master Engines to use external NTP servers.

Note

You can select NTP servers only for Master Engines. Virtual Engines do not communicate directly with NTP servers.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Enable time synchronization from NTP server**.
2) To add a row to the table, click **Add**.
3) To add an NTP server, right-click the **NTP Server** cell, select **Select**, then select an NTP Server element.
4) (Optional) If there is more than one NTP server, select the preferred NTP server.
5) Click **Next**.

Change the basic properties of converted Virtual Engine elements

Most of the basic settings for the Virtual Engine elements are inherited from the Master Engine. Some settings can be adjusted.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select a Virtual Engine to review or edit its properties.
2) In the **Name** field, edit the name for the Virtual Engine.
3) (Optional) In **DNS IP Addresses** field, add one or more DNS IP addresses.

   These addresses are the IP addresses of the DNS servers that the Master Engine with which the Virtual Engine is associated uses to resolve domain names. Virtual Engines need DNS resolution to contact services that are defined using URLs or domain names, and to resolve fully qualified domain names (FQDNs) used in policies. When DNS relay is configured, these DNS servers are used unless domain-specific DNS servers are specified in a DNS Relay Profile element.

   Note

   If you have defined NetLink-specific DNS IP addresses, adding DNS IP addresses overrides the NetLink-specific DNS IP addresses.

- To enter a single IP address manually, click **Add** and select **IP Address**. Enter the IP address.
- To define an IP address using a network element, click **Add** and select **Network Element**.

4) Select the appropriate **Categories**.

5) Click **Next**.
The **Review Interfaces for Virtual Engines** page opens.

Change the interfaces and routing for converted Virtual Engine elements

The Virtual Engine interface configuration is automatically generated based on the Master Engine interface configuration. Some settings can be adjusted.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Click **Options** to select which IP addresses are used in particular roles in system communications.

2) (Optional) Add ARP Entries.

3) (Optional) Click **Multicast Routing** to define static multicast or IGMP-based multicast forwarding (IGMP proxying).

4) Click **Next**.
The **Review and Edit Routing for Virtual Security Engines** page opens.

5) Review and edit the routing.

6) Click **Next**.
The **Review NAT Definitions for Virtual Security Engines** page opens.

7) Continue the configuration in one of the following ways:
   - If you want to add NAT definitions, review the NAT definitions for the converted Virtual Engine elements.
   - Otherwise, click **Next**. The **Select Additional Configuration Options** page opens.

Add NAT definitions for converted Virtual Engine elements

On the **Define NAT Definitions for Virtual Security Engines** page, you can define how Virtual Engines translate network addresses.

NAT rules are automatically generated and organized in the Engine Policy based on the NAT definitions created in the Engine properties.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) To use the default NAT address as the Public IP Address if there is not a more specific NAT definition that matches the traffic, select **Use Default NAT Address for Traffic from Internal Networks**.

2) To view the Default NAT Address properties, click **Show Details**.

3) In the **NAT Definitions for** list, select the Virtual Engine for which you want to add NAT definitions .

4) Add or edit the NAT definitions.

5) Click **Next**.
   The **Select Additional Configuration Options** page opens.

6) Continue the configuration in one of the following ways:
   - If you want to define additional options, select additional configuration options for the converted Virtual Engine elements.
   - Otherwise, Click **Next**. The **Summary** page opens.

Select additional configuration options for converted Virtual Engine elements

You can define additional properties for the Virtual Engine elements on the **Select Additional Configuration Options** page.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) (Optional) Select **Define Additional Properties for Virtual Security Engines** and click **Next**.
   The **Review Permissions for Virtual Security Engines** page opens.

2) (Optional) Select the Virtual Engine for which you want to define permissions and configure the permissions.

3) Click **Next**.
   The **Review Add-Ons for Virtual Security Engines** page opens.

4) (Optional) Select the Virtual Engine for which you want to define add-ons and configure the add-ons.

5) Click **Next**.
The **Review Advanced Settings for Virtual Security Engines** page opens.

6) (Optional) Select the Virtual Engine for which you want to define advanced engine properties and define the advanced properties.

7) Click **Next**.
The **Summary** page opens.

Finish converting Engines to Master Engines and Virtual Engine elements

When you have finished configuring the Master Engine and Virtual Engine elements, you must finalize the conversion.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Click **Finish**.
A new tab opens to show the progress of the conversion.

2) Click **Close** to close the tab when the conversion is complete.

3) Install or refresh the policy on the Master Engine.

4) Install a policy on the Virtual Engines.

Example: deploying Virtual Engines for MSSP customers

An example of configuring Master Engines and Virtual Engines in an MSSP environment.

Company A is an MSSP (Managed Security Services Provider). Customer 1 and Customer 2 are customers of Company A. The customers each want one Virtual Engine with two Physical Interfaces. The administrators at Company A decide to use their existing Security Engine appliance as a Master Engine to host Virtual Engines for Customer 1 and Customer 2. Separate administrative Domains have already been configured for each customer. The engine already has a license that allows the creation of Virtual Resources.

The administrators at Company A:

1) Create a Master Engine element in the Shared Domain.

2) Create one Virtual Resource element for each customer's Virtual Engine and select the appropriate Domain for each Virtual Resource:

Virtual resources details

Virtual resource name	Domain
Customer 1 Virtual Resource	Customer 1 Domain
Customer 2 Virtual Resource	Customer 2 Domain

3) Create the following Physical Interfaces on the Master Engine:

Physical interfaces details

Interface ID	Description
0	Physical Interface for the Master Engine's own traffic
1	Physical Interface for hosted Virtual Engine traffic

4) Add an IPv4 address for each Master Engine node to Physical Interface 0.

5) Add the following VLAN Interfaces to Physical Interface 1 and select the appropriate Virtual Resource for each VLAN Interface:

VLAN interfaces details

Interface ID	Virtual resource	Description
VLAN 1.1	Customer 1 Virtual Resource	VLAN Interface for the first Physical Interface on the Virtual Engine for Customer 1
VLAN 1.2	Customer 1 Virtual Resource	VLAN Interface for the second Physical Interface on the Virtual Engine for Customer 1
VLAN 1.3	Customer 2 Virtual Resource	VLAN Interface for the first Physical Interface on the Virtual Engine for Customer
VLAN 1.4	Customer 2 Virtual Resource	VLAN Interface for the second Physical Interface on the Virtual Engine for Customer 2

6) Create a Virtual Engine element for each customer and select the appropriate Virtual Resource for each Virtual Engine:

Virtual engine details

Virtual engine	Virtual resource
Customer 1 Virtual Engine	Customer 1 Virtual Resource
Customer 2 Virtual Engine	Customer 2 Virtual Resource

7) Add IP addresses to the Physical Interfaces on the Virtual Engines.

8) Refresh the policy on the Master Engine.

9) Refresh the policy on the Virtual Engines.

Chapter 31

Network interface configuration

Contents

- Network interfaces for Security Engine on page 543
- Configuring interfaces for Engines on page 547
- Configuring interfaces for IPS engines on page 572
- Configuring interfaces for Layer 2 Engines on page 585
- Configuring interfaces for Master Engines on page 598
- Configuring interfaces for Virtual Engines on page 609
- Add manual ARP entries to Security Engine on page 615
- Examples of interface configurations on page 615

The network interface configuration for Security Engines is stored on the Management Server in the properties of Single Engine, Engine Cluster, Single IPS, IPS Cluster, Single Layer 2 Engine, Layer 2 Engine Cluster, Master Security Engine, and Virtual Security Engine elements.

Network interfaces for Security Engine

The network interface configuration stored on the Management Server contains most of the settings for network interfaces.

The interface configuration is done using the SMC Client. The interface configuration stored on the Management Server does not contain:

- The network card driver selection
- The mapping of the operating system port numbers to the numbers used in the SMC Client
- The network card speed/duplex settings



Note

If you use automatic configuration for appliances, the Interface IDs that you select in the SMC Client must match the port numbers on the physical appliances. Check the port numbers in the relevant hardware guide.



Note

The configuration transferred from the Management Server overwrites the settings that can be defined through the engine command line. The settings contain the details for initial contact with the Management Server to establish a trusted communications channel.

Types of interfaces for Security Engines with Layer 3 Interfaces

You can configure several types of interfaces for Security Engines with Layer 3 Interfaces.

Types of interfaces for Security Engines with Layer 3 Interfaces

Interface type	Purpose of interface	Limitations
Layer 3 physical	System communications and traffic inspection.	You cannot add both VLAN Interfaces and IP addresses to a Physical Interface. If an IP address is already configured for a Physical Interface, adding a VLAN Interface removes the IP address. If you plan to use VLAN Interfaces, configure the VLAN Interfaces first and then add IP addresses to the VLAN Interfaces.
Layer 2 physical	Traffic inspection. Layer 2 interfaces on Security Engines with Layer 3 Interfaces allow the engine to provide the same kind of traffic inspection that is available for Security Engine in the IPS and Layer 2 Engine roles.	You cannot add layer 2 physical interfaces of the Inline Layer 2 Engine type to Engine Clusters in Load Balancing mode. Only Standby mode is supported. You cannot add IP addresses to layer 2 physical interfaces on Security Engines with Layer 3 Interfaces. VLAN retagging is not supported on layer 2 physical interfaces of the inline IPS type.
VLAN	Divides a single physical interface into several virtual interfaces.	■ You cannot add VLAN interfaces on top of other VLAN Interfaces (nested VLANs). ■ You cannot create valid VLAN Interfaces in a Virtual Engine if the Master Engine interface that hosts the Virtual Engine is a VLAN Interface.
Modem (Single Engines only)	Represents a mobile broadband modem connected to a USB port on a purpose-built Security Engine appliance.	■ A Modem Interface is only supported on Single Engines that run on specific Security Engine appliances. ■ Modem Interfaces do not support VLAN tagging.
Tunnel	A logical interface that is used as an endpoint for tunnels in route-based VPNs.	■ Tunnel Interfaces can only have static IP addresses. ■ Tunnel Interfaces do not support VLAN tagging.
VPN Broker	A specialized interface for use with the VPN Broker. For more information about VPN Broker, see the *Forcepoint Security Engine Manager and VPN Broker Product Guide*.	This type of interface is only supported for use with the VPN Broker.
Wireless (Single Engines only)	Represents a wireless network interface card of a purpose-built Security Engine appliance.	A Wireless Interface is only supported on Single Engines that run on specific Security Engine appliances that have a wireless network interface card.

Interface type	Purpose of interface	Limitations
Switch (Single Engines only)	Represents the switch functionality on a purpose-built Security Engine appliance.	■ The switch functionality is only supported on Single Engines that run on specific Security Engine appliances that have an integrated switch. ■ The ports in the integrated switch do not support VLAN tagging or PPPoE. ■ You cannot use ports on the integrated switch as the control interface.

Types of interfaces for Security Engines in the IPS and Layer 2 Engine roles

Interface definitions are an important part of IPS and Layer 2 Engine elements.

Types of interfaces for Security Engines in the IPS and Layer 2 Engine roles

Interface type	Purpose of interface	Limitations
Physical (Normal type)	System communications. These interfaces are used when the engine is the source or the final destination of the communications. An example is control communications between the engine and the Management Server. Define at least one interface that is dedicated to system communications for each IPS engine or Layer 2 Engine.	
Physical (Capture Interface or Inline Interface type)	Traffic inspection. Define one or more traffic inspection interfaces for each IPS engine or Layer 2 Engine.	
VLAN	Divides a single physical interface into several virtual interfaces.	■ You cannot add VLAN Interfaces on top of other VLAN Interfaces (nested VLANs). ■ You cannot create valid VLAN Interfaces in a Virtual Engine if the Master Engine interface that hosts the Virtual Engine is a VLAN Interface.

Interface numbering

The interfaces have their own numbering in the SMC called the interface ID. The interface IDs are mapped to the corresponding network interfaces on the Security Engine when you configure the Forcepoint Network Security Platform software.

Interface numbering for Security Engines

Interface type	Interface numbering in the SMC
Layer 3 physical (Securtiy Engine with layer 3 interfaces)	Each physical interface has a unique interface ID number.
Layer 2 physical (Securtiy Engine with layer 3 interfaces)	
Physical (IPS and Layer 2 Engine roles)	
VLAN	Each VLAN interface has a VLAN number. The defined VLAN interfaces are displayed, for example, as “5.202” for network interface 5 with VLAN 202.
Wireless	The wireless interface has a unique interface ID number. An SSID (service set identifier) interface represents an 802.11 wireless LAN. You can add several SSID interfaces to the wireless interface.
Modem	Modem Interfaces are identified with modem numbers. The modem number is mapped to the modem’s IMEI (international mobile equipment identity) number. Each modem is assigned a unique ID when you connect the modem to the engine. You can change the mapping between the modem’s IMEI number and the modem ID through the engine command line, if necessary.
Tunnel	Tunnel interfaces are numbered with tunnel interface ID numbers. The mapping of Tunnel Interfaces to physical network interfaces on the engine is done automatically by the engine operating system based on the routing configuration.
Integrated switch	Integrated switches are identified with switch IDs. Integrated switches have predefined switch IDs. For example, the switch ID is 0 on Forcepoint Network Security Platform 110 appliances. You can add port group interfaces to switches. Port group interfaces are identified by port group IDs. The defined switches and port group interfaces are displayed, for example, as *0.1* for switch ID 0 with port group 1.

Network interface configuration overview

The network interface configuration process depends on the Forcepoint Network Security Platform role.

The interface configuration proceeds as follows:

1) Define the interfaces and IP addresses according to the engine role.

2) Configure additional related settings depending on the features you want to use.

Configuring interfaces for Engines

The interface configuration process consists of several general steps.

The configuration proceeds as follows:

1) Add the required number of network connections:
   - Add physical interfaces.
   - (Optional, Single Engines only) Add a wireless interface for the integrated wireless router.
   - (Optional, Single Engines only) Add modem interfaces for mobile broadband modem connections.
   - (Optional, Single Engines only) Add a switch and port group interfaces for the integrated switch.
2) (Optional, physical interfaces only) Add VLAN interfaces to physical interfaces.
3) (Optional) Add tunnel interfaces for the route-based tunnels.
4) (Not applicable to modem interfaces) Configure the IP address settings.
5) (Optional) Add loopback IP addresses to assign IP addresses that do not belong to any directly connected networks to the engine.
6) Select the interfaces that are used for system communications.

Add layer 3 physical interfaces for Security Engine in the Engine/VPN role

A layer 3 physical interface represents an actual network port on the engine.

Layer 3 physical interfaces correspond to network ports on the Security Engine. The number of defined layer 3 physical interfaces can be lower than the number of network ports on the hardware.

By default, the numbering of the layer 3 physical interfaces in the SMC Client corresponds to the operating system interface numbering on the engine. For example, Interface ID 0 is mapped to eth0 and ID 1 to eth1. However, the mapping is not fixed and you can change it through the engine command line. See the relevant *Hardware Guide* for details about which Interface IDs to map with which network ports. On Engine Clusters, this mapping can be done differently from node to node. You must take care that the interface that represents the same network interface on each node is correctly cabled to the same network.

An interface of the type **None** represents a single network port on the engine. An *Aggregated Link* represents two or more network ports on the engine. An Aggregated Link in high-availability mode provides protection against hardware failure. You can also use an Aggregated Link in load-balancing mode to increase throughput. You can combine two network ports in an Aggregated Link in high-availability mode. You can combine up to eight network ports in an Aggregated Link in load-balancing mode.

Note

We recommend that you do not use the IP address of an Aggregated Link interface as the primary or secondary control IP address of the Engine.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Single Engine or Engine Cluster and select **Edit <element type>**.
2) In the navigation pane on the left, select **Interfaces**.
3) Right-click the empty space and select **New** > **Layer 3 Physical Interface**.
4) Define the interface properties.
5) Click **OK**.
6) Continue the configuration in one of the following ways:
   - If you want to use VLANs, add the VLANs before adding IP addresses.
   - Otherwise, add IP addresses directly to the layer 3 physical interfaces.

Related concepts

IP addresses for Single Engine interfaces on page 560
Defining Zone elements on page 923
Quality of Service (QoS) and how it works on page 973

Related tasks

Activate the internal DHCP server on a engine interface on page 572

Adding layer 2 physical interfaces for Security Engines in the Engine/VPN role

Layer 2 physical interfaces on Security Engines in the Engine/VPN role pick up traffic for inspection.

You can add one or more capture interfaces, inline IPS interfaces, and inline Layer 2 Engine interfaces to Security Engines in the Engine/VPN role.

Types of layer 2 physical interfaces for Security Engines in the Engine/VPN role

Interface Type	Description
Capture interface	Capture interfaces listen to traffic that is not routed through the Security Engine. Connections picked up through capture interfaces can be reset through reset interfaces.

Interface Type	Description
Inline IPS interface	The interface is directly on the traffic path so that traffic passes through the interface to reach its destination. The Security Engine can inspect the traffic coming from one interface and either stop the traffic or send it out through the other interface. The default action for network traffic in Access rules is **Allow**. When Bypass mode is used, if the interface is unable to process traffic, all traffic is allowed without inspection.
Inline Layer 2 Engine interface	The interface is directly on the traffic path so that traffic passes through the interface to reach its destination. The Security Engine can inspect the traffic coming from one interface and either stop the traffic or send it out through the other interface. The default action for network traffic in Access rules is **Discard**. Bypass mode cannot be used. If the interface is unable to process traffic, all traffic is blocked.

Configure layer 2 physical interfaces for engines in the following order:

1) (Optional) Add Logical Interfaces.

2) (Optional) Add Reset Interfaces for Capture Interfaces.

3) Add Capture Interfaces, Inline IPS Interfaces, or Inline Layer 2 Engine Interfaces.

Note

When you use layer 2 interfaces on Security Engines in the Engine/VPN role, follow the same cable connection guidelines as for IPS and Layer 2 Engines.

Related concepts

Cable connection guidelines for IPS and Layer 2 Engines on page 90

Add logical interfaces

Logical interface elements allow you to group interfaces together according to network segment and interface type.

Logical interfaces are used in the configuration of the following types of interfaces to represent one or more network interfaces:

- Capture interfaces on Engines, IPS Engines, and Layer 2 Engines
- Inline interfaces on IPS engines and Layer 2 Engines
- Inline IPS interfaces on Engines
- Inline Layer 2 Engine interfaces on Engines

You cannot use the same logical interface to represent both capture interfaces and inline interfaces on the same engine. On Engines, you cannot use the same logical interface to represent both inline IPS interfaces and inline Layer 2 Engine interfaces. Otherwise, a logical interface can represent any number or combination of physical interfaces or VLAN Interfaces.

There is one predefined logical interface element called **default_eth**. If you want to create both capture interfaces and inline interfaces on the same Security Engine, you must add at least one more logical interface.

On IPS engines and Layer 2 Engines, a logical interface element called **System Communications** is automatically assigned to interfaces that have an IP address that is used as the primary or backup Control IP address. You can use the **System Communications** logical interface to represent all Control IP addresses in IPS and Layer 2 Engine Policies.

Note

You cannot use the **System Communications** logical interface on engines for Capture interfaces, Inline IPS interfaces, or Inline Layer 2 Engine interfaces.

You can use logical interfaces in IPS Policies, Layer 2 Engine Policies, and Layer 2 Interface Policies to limit the scope of your rules. You can use logical interfaces to create rules that match based on which interface the traffic was picked up from. For example, you can create a different logical interface for each VLAN and use them to create rules that apply only to traffic from a specific VLAN.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Browse to **Other Elements**.
3) Right-click **Logical Interfaces**, then select **New Logical Interface**.
4) Configure the settings, then click **OK**.

Add reset interfaces

Reset interfaces interrupt communications picked up through capture interfaces when the traffic matches a rule that terminates connections.

Reset interfaces can deliver TCP resets and ICMP “destination unreachable” messages to interrupt communications picked up through capture interfaces when the traffic matches a rule that terminates connections.

The resets are sent using the source and destination addresses and MAC addresses of the communicating hosts, so an IP address is not mandatory for a reset interface. You can optionally add an IP address if you also want to use this interface for system communications.

VLANs are supported for sending resets, but the correct VLAN is selected automatically. The interface you want to use as the reset interface must not have any manually added VLAN configuration.

You can use an existing system communications interface for sending resets if the reset interface connects to the same networks as the capture interface, and there are no VLANs on the system communications interface.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click an engine element, then select **Edit <element type>**.
2) In the navigation pane on the left, select **Interfaces**.
3) Create a new Physical Interface.
   - For IPS engines and Layer 2 Engines, right-click the empty space and select **New Physical Interface**.

- For Engines, right-click the empty space and select **New** > **Layer 3 Physical Interface**.

4) Select an **Interface ID**.

5) Select the interface type according to the engine role:
   - For IPS engines and Layer 2 Engines, select **Normal Interface**.
   - For Engines, select **None**.

6) Click **OK**.

Result

The Physical Interface is added to the interface list.

Note

When you set up the physical network, make sure that the reset interface connects to the same networks as the capture interfaces.

Next steps

Set up the capture interfaces that use this reset interface.

Add layer 2 physical interfaces for Security Engine in the Engine/VPN role

Layer 2 physical interfaces on Security Engine in the Engine/VPN role provide traffic inspection.

Layer 2 physical interfaces have definitions for the corresponding logical interface that the interface belongs to. The logical interface represents one or more network interfaces that capture the traffic for inspection.

On inline IPS interfaces and inline Layer 2 Engine interfaces, you must select a logical interface for each inline interface.

Note

You cannot use the same logical interface to represent both inline IPS interfaces and inline Layer 2 Engine interfaces.

For capture interfaces, the configuration depends on the deployment.

- When a capture interface is connected to a switch SPAN port, you must select a logical interface for each capture interface. You can optionally select the same logical interface for more than one capture interface.
- When a network TAP device is used, you must select the same logical interface for two capture interfaces. The monitored traffic going to different directions is captured through these two related network interfaces and is then combined into a complete traffic flow on the logical interface.

To define a capture interface, you must select a reset interface for it. You can use Layer 3 physical interfaces that meet the following requirements as reset interfaces:

- The Layer 3 physical interface connects to the same networks as the capture interfaces.
- There are no VLANs on the Layer 3 physical interface.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Single Engine or Engine Cluster and select **Edit <element type>**.
2) In the navigation pane on the left, select **Interfaces**.
3) Right-click the empty space and select **New** > **Layer 2 Physical Interface**.
4) From the **Type** drop-down list, select the interface type.
5) Define the physical interface properties.

Note

You cannot select the same logical interface for a capture interface and an inline interface on the same Security Engine.

CAUTION

Using **Bypass** as the **Failure Mode** requires a fail-open network interface card. If the ports that represent the interfaces cannot fail open, policy installation fails on the Security Engine. Bypass mode is not compatible with VLAN retagging. In network environments where VLAN retagging is used, normal mode is automatically enforced.

6) Click **OK**.

Related concepts

Cable connection guidelines for IPS and Layer 2 Engines on page 90

Add VLAN interfaces for engines

VLANs divide a single physical network link into several virtual links. You can define VLANs for both Single Engines and Engine Clusters.

A Virtual Local Area Network (VLAN) is a logical grouping of hosts and network devices that allows creating several separated networks on a single physical link. To allow this separation, the Engine supports VLAN tagging as defined in the IEEE 802.1q standard.

VLANs also make it easier to deploy geographically distributed Engine Clusters (for example, a cluster whose nodes are located in different buildings). Fewer physical interfaces and less cabling is needed. When you create a VLAN interface, the CVI mode and MAC address are defined commonly for all virtual interfaces configured for the same Physical interface definition.

One network interface can support up to 4094 VLANs. The defined VLAN interfaces are displayed, for example, as “5.202” for network interface 5 with VLAN 202. The VLANs must also be defined in the configuration of the switch or router to which the interface is connected.

Note

If an IP address is already configured for a Engine physical interface, adding a VLAN interface removes the IP address. If you plan to use VLAN interfaces, configure the VLAN interfaces first and then add IP addresses to the VLAN interfaces.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Single Engine or Engine Cluster, then select **Edit <element type>**.
2) In the navigation pane on the left, browse to **Interfaces**.
3) Right-click a physical interface, then select **New** > **VLAN Interface**.
4) Configure the settings, then click **OK**.

Result

The specified VLAN ID is added to the physical interface.

Next steps

Continue the configuration in one of the following ways:

- Add other types of interfaces.
- Add IP addresses to the physical interfaces or VLAN Interfaces.

Related concepts

IP addresses for Single Engine interfaces on page 560
Defining Zone elements on page 923
Quality of Service (QoS) and how it works on page 973

Related tasks

Activate the internal DHCP server on a engine interface on page 572

Add ADSL Interfaces for Single Engines

You can configure ADSL Interfaces on some legacy Forcepoint Network Security Platform appliances.

Note

ADSL Interfaces are only supported on specific legacy appliances that have an integrated ADSL network interface card.

You can only configure one ADSL Interface for each Single Engine. ADSL Interfaces are not supported on Engine Clusters.

The supported ADSL standards are ANSI T1.413 issue 2n, G.dmt, G.lite, ADSL2 DMT, ADSL2 G.lite, Annex A, and Annex B.

Use the number of the ADSL port on the appliance as the Interface ID of the ADSL Interface. For information about mapping the ADSL port on the appliance, see the relevant *Hardware Guide*.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Single Engine and select **Edit Single Engine**.
2) In the navigation pane on the left, select **Interfaces**.
   The **Interfaces** pane opens on the right.
3) Right-click the empty space and select **New** > **ADSL Interface**.
4) Define the ADSL Interface properties.
5) Click **OK**.
   The ADSL Interface is added to the interface list.

Next steps

You are now ready to add IP addresses for the Single Engine.

Add Wireless Interfaces for Single Engines

You can define a Wireless Interface to use a Single Engine as a wireless access point.

You can configure one Wireless Interface on a Single Engine. Wireless Interfaces are only supported on specific Forcepoint Network Security Platform appliances that have an integrated wireless network interface card. Wireless Interfaces are not supported on Engine Clusters.

You can define several wireless LANs for the Wireless Interface. A wireless LAN is defined by adding an SSID (service set identifier) interface for the Wireless Interface.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Single Engine and select **Edit Single Engine**.
   The Engine Editor opens.
2) In the navigation pane on the left, select **Interfaces**.
   The **Interfaces** pane opens on the right.
3) Right-click the empty space and select **New** > **Wireless Interface**.

4) Define the Wireless Interface properties.



Note

Use the number of the wireless port on the appliance as the **Interface ID** of the Wireless Interface.

5) Click **OK**.
The Wireless Interface is added to the interface list.

Next steps

Define SSID interfaces for the Single Engine.

Related tasks
Add SSID Interfaces for Single Engines on page 555

Add SSID Interfaces for Single Engines

You can define several SSID Interfaces for the Wireless Interface.

An SSID (service set identifier) interface represents an 802.11 wireless LAN. You can define several SSID Interfaces for the Wireless Interface.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Single Engine and select **Edit Single Engine**.
The Engine Editor opens.

2) In the navigation pane on the left, select **Interfaces**.
The **Interfaces** pane opens on the right.

3) Right-click the Wireless Interface, then select **New SSID Interface**.

4) Define the SSID Interface properties.

5) If you want to configure additional settings for the SSID Interface, continue the configuration in one of the following ways:
   - Configure security settings for the wireless connections.
   - Filter wireless connections based on the clients' MAC addresses.
   - Activate the internal DHCP server on the SSID Interface.
   - Add an IPv4 address to the interface.

6) Otherwise, save the changes to the SSID Interface.
   a) Click **OK** to close the **SSID Interface Properties** dialog box.

b) Click **Save and Refresh**.

Related concepts
Defining Zone elements on page 923
Quality of Service (QoS) and how it works on page 973

Related tasks
Activate the internal DHCP server on a engine interface on page 572
Add IPv4 addresses to Single Engine interfaces on page 562

Add Modem Interfaces for Single Engines

A Modem Interface represents the settings of a mobile broadband modem that provides a wireless link from a Single Engine to the Internet or to the Management Server. You can optionally define one or more Modem Interfaces for each Single Engine.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Single Engine, then select **Edit Single Engine**.
2) In the navigation pane on the left, select **Interfaces**.
3) Right-click the empty space, then select **New** > **Modem Interface**.
4) Define the basic Modem Interface properties.
5) If necessary, define the contact address information.
   If components from some Locations cannot use the Default contact address, click **Exceptions** to define Location-specific contact addresses.
6) Click **OK**.
7) If you are configuring a new Single Engine, or if you want to change the roles the different interfaces have in the configuration, select system communication roles for engine interfaces.
8) Otherwise, update the routing configuration and transfer the changes to the engine.
   a) Select **Routing** in the navigation pane on the left.
   b) Adjust the routing configuration as necessary.
   c) Click **Save and Refresh**.

Related concepts
Defining Zone elements on page 923

Related tasks
Define contact addresses for a single Security Engine or a Cluster Virtual IP Address on page 138
Change or remove the PIN codes of Modem Interfaces on page 557
Select system communication roles for engine interfaces on page 570
Add manual ARP entries to Security Engine on page 615

Change or remove the PIN codes of Modem Interfaces

In some cases, you might need to change or remove your PIN code.

If you change the SIM card of a mobile broadband modem that is connected to a Single Engine and the PIN code is enabled on the new SIM card, you must change the PIN code information about the related Modem Interface. You must remove the PIN code if a PIN code was enabled on the old SIM card but a PIN code is not enabled on the new SIM card.

You must also change the PIN code if you have received the message “PIN code differs from initial contact settings” at policy installation after you have finished configuring the Engine. This message means that the PIN code in the Modem Interface properties does not match the PIN code that you have entered for the mobile broadband modem in the command-line interface. If there are other problems with the PIN code, you can check through the engine’s command line interface which PIN code was used in the initial configuration. You can then change the PIN code information if necessary.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) On the Security Engine command line, enter the sg-reconfigure command to start the Security Engine Configuration Wizard, then change the mobile broadband modem’s PIN code information.
2) In the SMC Client, right-click the Security Engine, then select **Edit <element type>**.
3) Browse to **Interfaces**.
4) Right-click the Modem Interface, then select **Properties**.
5) Change the PIN code information, then click **OK**.
6) Click **Save and Refresh**.

Related tasks
Add Modem Interfaces for Single Engines on page 556

Add tunnel interfaces for engines

Tunnel Interfaces allow routing information to be used to determine the correct VPN tunnel to use in route-based VPNs.

Any traffic that is routed to a tunnel interface and allowed by Access rules is automatically sent through the tunnel to the peer endpoint defined in the Route-based Tunnels element. Tunnel interfaces are only used in route-based VPNs.

You can optionally add IPv4 or IPv6 addresses to a tunnel interface. Tunnel interfaces can only have static IP addresses. Any IP address can be added to a tunnel interface, even if the same IP address is used on another interface or as a loopback IP address. Adding an IP address to a tunnel interface allows you to define the source IP address of traffic sent from the Security Engine itself. For example, an IP address is recommended to provide a source IP address for dynamic routing daemons, for IGMP proxy, and for Protocol Independent Multicast - Sparse-Mode (PIM-SM) configuration. If no IP address is added to the tunnel interface, the source IP address for traffic sent from the Security Engine is automatically selected. The selection is done according to the **Bypass Default IP Address** setting in the loopback interface configuration for the Security Engine.

The mapping of tunnel interfaces to physical network interfaces on the Security Engine is done automatically based on the routing configuration.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click an Security Engine, then select **Edit <element type>**.
3) Browse to **Interfaces**.
4) Select **Add** > **Tunnel Interface**.
5) Configure the settings.
6) Click **OK**.
7) If you want to add a source IP address for traffic sent from the engine node, add IPv4 addresses or IPv6 addresses to the tunnels.
8) If you do not want to add IP addresses, select system communication roles for engine interfaces to define how the source IP address for traffic sent from the engine node is selected.
9) Click **Save and Refresh**.

Related concepts

Quality of Service (QoS) and how it works on page 973
Types of VPNs in Forcepoint Network Security Platform on page 1157

Related tasks

Add tunnel interfaces for Virtual Engines on page 611
Select system communication roles for engine interfaces on page 570
Add IPv4 addresses to Single Engine interfaces on page 562
Add IPv6 addresses to Single Engine interfaces on page 565
Add IPv4 and IPv6 addresses to Engine Cluster interfaces on page 567

Add integrated switches for Single Engines

An integrated switch represents the switch functionality on purpose-built Forcepoint Network Security Platform appliances. Integrated switches eliminate the need for an external switch device and reduce costs and clutter.

The switch functionality is only supported on specific Forcepoint Network Security Platform appliances that have a hardware or software integrated switch. For more information, see the model-specific *Forcepoint Network Security Platform Hardware Guide* for your Forcepoint Network Security Platform appliance.

- On Forcepoint Network Security Platform appliances that have hardware integrated switches, you can configure one integrated switch and one or more port group interfaces.
- On Forcepoint Network Security Platform appliances that have software integrated switches, you can configure one or more integrated switches, and one port group interface on each integrated switch.

You can only use the integrated switch if the appliance has been configured as a Single Engine.

Note

The ports in the integrated switch do not support VLAN tagging or PPPoE. You cannot use ports on the integrated switch as the control interface.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Single Engine and select **Edit Single Engine**.
   The Engine Editor opens.
2) In the navigation pane on the left, select **Interfaces**.
3) In the right pane, right-click the empty space and select **New Switch**.
4) In the **Switch Properties** dialog box, define the switch properties.
5) Click **OK**.
   The switch is added to the interface list.

Next steps

Add port group interfaces to the switch.

Add Port Group Interfaces for Single Engines

Port groups simplify port and network segment configuration. Traffic inside a port group is not inspected. The traffic between port groups is inspected by the engine in the same way as other traffic.

Before you begin

You must add the integrated switch before you can add port group interfaces.

Depending on the Forcepoint Network Security Platform appliance model, you can define one or more port group interfaces and add different types of interfaces to the port group:

- On Forcepoint Network Security Platform appliances that have hardware integrated switches, you can define one or more port group interfaces on the integrated switch.
  You can add physical interfaces to the port group interface.
- On Forcepoint Network Security Platform appliances that have software integrated switches, you can define one port group interface on each integrated switch.
  You can add physical and SSID interfaces to the port group interface. You must first add the physical and SSID interfaces in the engine editor without any IP address configuration, before adding these interfaces to the port group.

For more information about the type of integrated switch that your appliance has, see the model-specific *Forcepoint Network Security Platform Hardware Guide* for your Forcepoint Network Security Platform appliance.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Single Engine, then select **Edit Single Engine**.
2) In the navigation pane on the left, select **Interfaces**.
3) In the right pane, right-click the switch and select **New Port Group Interface**.
4) In the **Port Group Interface Properties** dialog box, define the port group interface properties.
5) Click **OK**.
   The port group interface is added to the interface list. The defined switches and port group interfaces are displayed, for example, as *0.1* for switch ID 0 with port group 1.

Next steps

Add IP addresses to the port group interface.

IP addresses for Single Engine interfaces

You can define several types of IP addresses for the Single Engine's interfaces.

The interfaces of a Single Engine can have the following types of IP addresses:

- Each layer 3 physical interface can have one or more static or dynamic IP addresses. A layer 3 physical interface can have multiple dynamic IP addresses only if you add VLAN interfaces and the VLAN interfaces

each have a dynamic IP address. Otherwise, a layer 3 physical interface can only have a single dynamic IP address. The same interface can have both IPv4 and IPv6 addresses.

- Each Aggregated Link interface can have one or more static IP addresses. The same interface can have both IPv4 and IPv6 addresses.
- Each VLAN interface can have one or more static IP addresses or a single dynamic IP address. The same interface can have both IPv4 and IPv6 addresses.
- Each SSID interface defined for the optional Wireless interface can have a single IPv4 or IPv6 address.
- A Tunnel interface can have one or more static IP addresses.
- A Port Group interface of an integrated Switch can have one or more static IP addresses or a single dynamic IP address. The same interface can have both IPv4 and IPv6 addresses.

Note

You cannot add IP addresses to layer 2 physical interfaces.

You might need to define a contact address if you enter a private static address and NAT is used to translate it to a different external IP address. The external IP address must be configured as the contact address if other SMC components must use the external IP address to contact this Engine or if the IP address is used as a VPN endpoint.

To use the engine as an IGMP proxy for multicast routing, the IP addresses for the downstream interfaces must be the lowest IP addresses among all IGMP queries in the local networks.

Related concepts

Define contact IP addresses on page 127

Dynamic IP addresses for Single Engine interfaces and how they work

Single Engines support the use of DHCP, PPPoA, PPPoE, and SLAAC to assign dynamic IPv4 or IPv6 addresses on the engine's network interfaces. PPP is only supported for IPv4 addresses.

Typically, a dynamic IP address is used in smaller sites with xDSL connections that have no static IP address available for the engine. Modem Interfaces always have dynamic IP addresses that are provided through PPP.

Instead of an IP address, each interface with a dynamic IP address is identified in the SMC by a DHCP Index. This is a number that is used to distinguish different interfaces with dynamic IP addresses from one another.

When a engine has a fixed IP address, the Management Server can contact the engine whenever there is need. When the engine has a dynamic IP address, the Management Server does not have a fixed IP address to contact, so the engine contacts the Management Server instead. You can also define that a engine that has a fixed IP address contacts the Management Server. The frequency of these contacts can be adjusted as necessary. If the contact is lost (for example, the Internet connection goes down), the Management Server queues the commands you have made to the engine and executes them when contact is re-established.

If the address the engine uses for system communications is dynamic, the engine opens a communications channel to the Management Server and the Management Server never attempts to contact the engine outside this connection. If the management traffic flows through an interface that has a dynamic address, you can adjust timeouts and other settings related to these communications.

Dynamic IP addresses also affect policy-based VPNs: other VPN gateways cannot open VPN connections to the gateway if the address is dynamic. Instead, the gateway with the dynamic endpoint must always initiate the VPN. After the VPN tunnel is established, connections can be made in both directions as usual. VPN client

connections can be forwarded through a site-to-site VPN from some gateway that has a static IP address (VPN hub configuration).

Because the dynamic IP address assignment includes the next-hop gateway address, the routing for interfaces with a dynamic address is defined using special dynamic Router and NetLink elements.

There are default Alias elements that can be used in the Engine's policy to represent its own dynamic addresses. For each dynamic address interface there are four Alias elements distinguished by the DHCP index number:

- **$$ DHCP Interface X.ip**: The current dynamic IP address allocated to this interface
- **$$ DHCP Interface X.gateways**: The received IP address for the default router
- **$$ DHCP Interface X.dns**: The received IP address of the DNS server
- **$$ DHCP Interface X.net**: The network behind the interface with a dynamic IP address

Note

These Aliases are meant for use in the policies of the Engine that has the dynamic IP address. They are translated to the values of the Engine the policy is installed on, so they cannot be used in the policies of other components.

Related tasks

Configure global contact policy settings for node-initiated contact to the Management Server on page 524
Add Modem Interfaces for Single Engines on page 556

Add IPv4 addresses to Single Engine interfaces

You can add IPv4 addresses to layer 3 physical interfaces, VLAN interfaces, SSID interfaces, tunnel interfaces, and port group interfaces of an integrated switch.

Dynamic IPv4 addresses are supported on layer 3 physical interfaces, VLAN interfaces, and port group interfaces. Dynamic IP addresses are not supported on Aggregated Link interfaces.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Single Engine, then select **Edit Single Engine**.

2) Browse to **Interfaces**.

3) Right-click a layer 3 physical interface, VLAN interface, SSID interface, tunnel interface, or port group interface, then select **New > IPv4 Address**, or right-click an ADSL interface, then select **New IPv4 Address**.

4) Configure the IP address information in one of the following ways:
   - Select **Static**, then enter the **IPv4 Address**. The Network Settings are automatically entered.

Tip

To resolve the IP address from a DNS name, right-click the field, then select **Resolve From DNS Name**.

- (Physical, VLAN, ADSL, and port group interfaces only) Select **Dynamic**, then select the DHCP index from the **Dynamic Index** drop-down list. The index is used for identification in other parts of the configuration (such as Engine Policies) to represent the possibly changing IP address.

5) (Physical, ADSL, SSID, VLAN, and port group interfaces only) If necessary, define the contact address information.
   - Enter the **Default** contact address or select **Dynamic** (next to the **Default** field) if the interface has a dynamic contact address. The Default contact address is used by default whenever a component that belongs to another Location connects to this interface.
   - If components from some Locations cannot use the Default contact address, click **Exceptions** to define Location-specific contact addresses.
   - Dynamic contact addresses are not supported on SSID interfaces.

6) (Static IPv4 addresses only) Check the automatically filled-in **Netmask**, then adjust it as necessary.

7) (Optional) Configure additional features for this interface:
   - (Static IPv4 addresses only) If you want to use VRRP on the physical interface, VLAN interface, or port group interface of an integrated switch, add virtual routers to Single Engine interfaces.
   - (Dynamic IPv4 addresses only) If the interface requires PPPoE or PPPoA support, add point-to-point protocol clients to Single Engine interfaces.
   - (Dynamic IPv4 addresses only) If you do not want a default route to be created through the interface, deselect **Automatic Default Route**.

8) Click **OK**.

9) Click **Save**.

10) Continue the configuration in one of the following ways:
   - Add IPv6 addresses to a physical interface, VLAN interface, SSID interface, tunnel interface, or port group interface.
   - Add modem interfaces.
   - If you are creating a Engine, or if you want to change the roles the different interfaces have in the configuration, select system communication roles for engine interfaces.
   - If you added IP addresses to tunnel interfaces, define routing for route-based VPNs.

Related concepts

Dynamic IP addresses for Single Engine interfaces and how they work on page 561

Related tasks

Add loopback IP addresses to engines on page 571
Add manual ARP entries to Security Engine on page 615

Add virtual routers to Single Engine interfaces

Virtual Router Redundancy Protocol (VRRP) allows high-availability router configurations.

VRRP support is only available on Physical Interfaces, VLAN Interfaces, or Port Group Interfaces of Single Engines. One virtual router can be configured for each Physical Interface or VLAN Interface. The virtual router can have either backup or active status. The virtual router configured for one interface does not take into account the status of possible virtual routers configured on other interfaces. VRRP is only supported for static IPv4 addresses.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the IP Address properties, click **VRRP Settings**.
2) Select **Enable VRRP**.
3) Enter the **ID**, **Priority**, and **IPv4 Address** according to the configuration of the virtual router.
4) Click **OK** to close the **VRRP Settings** dialog box.
5) Click **OK** to close the **IP Address Properties** dialog box.
6) If you are creating a Engine, or if you want to change the roles the different interfaces have in the configuration, select system communication roles for engine interfaces.
7) Otherwise, click **Save and Refresh**.

Related tasks

Add Modem Interfaces for Single Engines on page 556
Add IPv4 addresses to Single Engine interfaces on page 562
Add IPv6 addresses to Single Engine interfaces on page 565
Select system communication roles for engine interfaces on page 570

Add point-to-point protocol clients to Single Engine interfaces

You can use PPPoE or PPPoA clients to connect the Engine to an external ADSL modem without configuring routing and NAT settings on the ADSL modem.

Point-to-point protocol over Ethernet (PPPoE) and point-to-point protocol over ATM (PPPoA) clients on the Engine simplify the installation of an appliance when PPPoE or PPPoA is used by the network link connected to the interface. Activating the PPPoE or PPPoA client on the Engine allows you to connect the Engine to an external ADSL modem in transparent bridge mode. You must also activate the PPPoE or PPPoA client on the Engine if you have a specific appliance that has an ADSL port and the ISP for the ADSL connection requires PPPoE or PPPoA.

Note

- The PPPoE client is supported for both dynamic IPv4 addresses and dynamic IPv6 addresses on single engines.
- For a SMC managed single engine, the same PPP settings are shared if the interface has both IPv4 and IPv6 PPPoE configured.
- The PPPoE interface supports working in conjunction with single IP HA, route probing, and netlinks.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the **IP Address properties** dialog box, click **PPP Settings**. The **PPP Settings** dialog box opens.
2) Configure the settings, then click **OK**.
   If you do not have the PPP details, contact your service provider.
3) Click **OK** to close the **IP Address Properties** dialog box.
4) If you are creating a Engine, or if you want to change the roles the different interfaces have in the configuration, select system communication roles for engine interfaces.
5) Otherwise, click **Save and Refresh**.

Related tasks

Add Modem Interfaces for Single Engines on page 556
Add IPv4 addresses to Single Engine interfaces on page 562
Add IPv6 addresses to Single Engine interfaces on page 565
Select system communication roles for engine interfaces on page 570

Add IPv6 addresses to Single Engine interfaces

You can add IPv6 addresses to layer 3 physical interfaces, VLAN interfaces, SSID interfaces, tunnel interfaces, and port group interfaces of an integrated switch on Single Engines.

Dynamic IPv6 addresses are supported on layer 3 physical interfaces, VLAN interfaces, and port group interfaces. Dynamic IP addresses are not supported on Aggregated Link interfaces.

Note

If you have added VLAN interfaces to physical interfaces, you must add the IPv6 Addresses to the VLAN interfaces.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Single Engine, then select **Edit Single Engine**.

2) Browse to **Interfaces**.

3) Right-click a layer 3 physical interface, VLAN interface, SSID interface, tunnel interface, or port group interface, then select **New > IPv6 Address**.

4) Configure the IP address information in one of the following ways:
   - Select **Static**, then enter the IPv6 Address. The Network Settings are automatically entered.

     Tip

     To resolve the IP address from a DNS name, right-click the field, then select **Resolve From DNS Name**.
   - (Physical, VLAN, and port group interfaces only) Select **Dynamic**, then select the DHCP index from the **Dynamic Index** drop-down list. The index is used for identification in other parts of the configuration (such as Engine Policies) to represent the possibly changing IP address.

5) (Static IPv6 addresses only) Check the automatically filled-in **Prefix Length**, then adjust it if necessary.

6) (Optional) Configure additional features for this interface:
   - (Dynamic IPv6 addresses only) If the interface requires PPPoE or PPPoA support, add point-to-point protocol clients to Single Engine interfaces.
   - If you do not want a default route to be created through the interface, deselect **Automatic Default Route**.

     Note

     This step is applicable to dynamic IPv6 addresses only.

7) Click **OK**.

8) Click **Save**.

9) Continue the configuration in one of the following ways:
   - Define Modem interfaces.
   - If you are creating a Engine, or if you want to change the roles the different interfaces have in the configuration, select system communication roles for engine interfaces.
   - If you added IP addresses to tunnel interfaces, define routing for the route-based tunnels.

Related tasks

Add Modem Interfaces for Single Engines on page 556
Select system communication roles for engine interfaces on page 570
Add loopback IP addresses to engines on page 571

Add IPv4 and IPv6 addresses to Engine Cluster interfaces

You can add IPv4 and IPv6 addresses to layer 3 physical interfaces, VLAN interfaces, and tunnel interfaces on Engine Clusters.

IPv6 addresses are supported on Engine Clusters with dispatch clustering mode. IPv6 and IPv4 addresses can be used together on the same Engine Cluster.

Engine Clusters can have two types of IP addresses.

Types of IP addresses for Engine Clusters

IP address type	Description	When to use it
Cluster Virtual IP address (CVI)	An IP address that is used to handle traffic routed through the cluster for inspection. All nodes in a cluster share this IP address. Allows other devices to communicate with the Engine Cluster as a single entity. Each CVI inherits the MAC address defined for the physical interface. The MAC/IP address pair always remains the same as only the location of the MAC address changes to the current dispatcher node (packet dispatch). This configuration makes the external network equipment forward traffic to the correct node for dispatching. The CVIs on different physical interfaces cannot have duplicate MAC addresses.	Define a CVI for the interface if traffic that the engine inspects is routed to or from the interface.
Node Dedicated IP address (NDI)	An IP address that is used for traffic to or from an individual node in a cluster. Each node in the cluster has a specific IP address that is used as the NDI. NDIs are used for the following purposes: ■ Node-to-node communications, such as heartbeat connections between the engines in a cluster and other traffic to or from individual nodes. ■ Traffic between each individual node and the Management Server and Log Server ■ Communications with external components (such as authentication servers, or hosts that are probed in network connectivity tests) When you define NDIs, you must define both node-specific properties (such as the node's IP address) and properties that all nodes in the cluster share. All nodes must have the same netmask value for their NDI.	Define at least 2 NDIs: one for management connections and one for the heartbeat traffic between the nodes. We recommend that you define an NDI for each interface that has a CVI, if practical. Some features might not work reliably without an NDI. If there is a CVI without a corresponding NDI from the same network segment, communications that require an NDI 'borrow' an IP address. The address can be borrowed from another NDI on the same physical interface, VLAN interface, or aggregated link interface. If there is no NDI on the same physical interface, VLAN interface, or aggregated link interface, the default IP address for outgoing traffic is used. The 'borrowed' IP address can be used without issues with routers that strictly follow the ARP standard. You might need to create a static ARP entry if some routers do not strictly follow the ARP standard.

You can define one or more CVI or NDI for the same physical interface or VLAN interface. You can also define only a CVI or only an NDI for a physical interface or VLAN interface. If the physical interface is an aggregated link, all interfaces that belong to the aggregated link share the IP address definitions.

You might also need to define a contact address if the CVI or NDI is private and NAT is used to translate the IP address to a different external IP address. The external IP address must be configured as the contact address in the following cases:

- Other SMC components must use the external IP address to contact this Engine (NDI).
- This IP address is a VPN endpoint (CVI).

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Engine Cluster, then select **Edit Engine Cluster**.

2) Browse to **Interfaces**.

3) Right-click a layer 3 physical interface, VLAN interface, or tunnel interface, then add the IP address in one of the following ways:
   - To add an IPv4 address, select **New** > **IPv4 Address**
   - To add an IPv6 address, select **New** > **IPv6 Address**

   Note

   If you have added VLAN interfaces to a physical interface, add the IP addresses to the VLAN interfaces.

4) (Optional) If this interface does not receive or send IPv6 traffic that the Engine examines, deselect **Cluster Virtual IP Address**.

   Note

   By default, both **Cluster Virtual IP Address** and **Node Dedicated IP Address** are selected.

5) To add a CVI address, enter the IP address in the **IPv4 Address** or **IPv6 Address** field.

   Tip

   To resolve the IP address from a DNS name, right-click the field, then select **Resolve From DNS Name**.

6) (IPv4 addresses only) If necessary, define the contact address for the Engine Cluster.
   - In the **Default** field, enter the default contact address. The default contact address is used by default whenever a component that belongs to another Location connects to this interface.
   - If components from some Locations cannot use the default contact address, click **Exceptions** to define Location-specific contact addresses.

7) To add NDI addresses for the nodes, click the **IPv4 Address** or **IPv6 Address** cell in the table, then enter the IP address for each node.



Tip

To resolve the IP address from a DNS name, right-click the field, then select **Resolve From DNS Name**.

8) (IPv4 addresses only) If necessary, double-click the **Contact Address** cell in the table, then define the contact address for each node.
   - In the **Default** field at the top of the dialog box, enter the default contact address.
   - If components from some Locations cannot use the default contact address, click **Add** to define Location-specific contact addresses.

9) (IPv4 addresses only) Check the automatically filled-in **Netmask**, then adjust it as necessary.

10) (IPv6 addresses only) Check the automatically filled-in **Prefix Length**, then adjust it as necessary.

11) Click **OK**.

12) Click **Save**.

13) Continue the configuration in one of the following ways:
   - If you are creating a Engine Cluster, or if you want to change the roles the different interfaces have in the configuration, select system communication roles for engine interfaces.
   - If you added IP addresses to tunnel interfaces, define routing for route-based tunnels.
   - Otherwise, refresh the policy to transfer the changes.

Related concepts

IP addresses for Single Engine interfaces on page 560

Related tasks

Define contact addresses for a single Security Engine or a Cluster Virtual IP Address on page 138
Define contact addresses for Node Dedicated IP Addresses on page 139
Select system communication roles for engine interfaces on page 570
Add manual ARP entries to Security Engine on page 615

Select system communication roles for engine interfaces

You can select which IP addresses are used for particular roles in system communications (for example, in communications between the Engine and the Management Server).

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Single Engine or Engine Cluster and select **Edit <element type>**.

2) In the navigation pane on the left, select **Interfaces** > **Interface Options**.

3) In the **Interface Options** pane on the right:

   a) From the **Primary** control IP address drop-down list, select the primary control IP address for communications with the Management Server.

   Note

   We recommend that you do not use the IP address of an Aggregated Link interface as the primary or secondary control IP address of the Engine.

   b) (Optional, recommended) From the **Backup** control IP address drop-down list, select a backup control IP address for Management Server contact (used if the primary fails).

   c) (Optional) If the Security Engine is behind a device that applies dynamic NAT to the inbound management connections or blocks them, select **Node-Initiated contact to Management Server**.

   When this option is selected, the engine opens a connection to the Management Server and maintains connectivity.

   d) (Engine Clusters only) Select the primary Heartbeat Interface for communications between the nodes.

   We recommend that you use a Physical Interface, not a VLAN Interface. We strongly recommend that you do not direct any other traffic through this interface. A dedicated network helps ensure reliable and secure operation.

   CAUTION

   Primary and Backup Heartbeat networks exchange confidential information. If dedicated networks are not possible, configure the cluster to encrypt the exchanged information.

   e) (Engine Clusters only) Select a backup Heartbeat Interface that is used if the Primary Heartbeat Interface is unavailable.

   It is not mandatory to configure a backup Heartbeat Interface, but we strongly recommend it. If heartbeat traffic is not delivered, the cluster cannot operate and traffic is disturbed. We strongly recommend that you use a dedicated interface for the backup heartbeat as well.

   f) In the **Default IP Address for Outgoing Traffic** field, select the IP address that the nodes use if they have to initiate connections through an interface that has no Node Dedicated IP Address.

4) Click **OK**.

5) Click **Save and Refresh**.

Next steps

You are now ready to configure routing.

- If this is a new Engine, add routes for Engines.
- If you are configuring Tunnel Interfaces for route-based tunnels, define the routing for the route-based tunnels.

Related tasks

Adjust general Engine clustering options on page 677

Add loopback IP addresses to engines

You can use Loopback IP addresses to assign IP addresses that do not belong to any directly connected networks to the Single Engine or Engine Cluster.

Loopback IP addresses are not connected to any physical interface and they do not create connectivity to any network.

- Any IP address that is not used to route traffic on another interface can be used as a loopback IP address, and you can add several loopback IP addresses to each Engine.
- Any IP address that is not already used as a Cluster Virtual IP Addresses (CVI) or Node Dedicated IP Addresses (NDI) on another interface can be used as a loopback IP address.
  - A CVI loopback IP address is used for loopback traffic that is sent to the whole cluster. All the nodes in the cluster share it.
  - An NDI loopback IP address is used for loopback traffic that is sent to a specific node in the cluster. NDI loopback IP addresses must be unique for each node. You must define an NDI loopback IP address for all nodes.
- The same IP address can be used as a loopback IP address and as the IP address of a Tunnel Interface.
- Loopback IP addresses can be used as the IPv4 Identity for Authentication Requests or IPv6 Identity for Authentication Requests, the IPv4 Source for Authentication Requests or IPv6 Source for Authentication Requests, and the Default IP Address for Outgoing Traffic.
- Loopback IP addresses cannot be used as Control IP addresses for communication with the Management Server or as Heartbeat Interfaces.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Security Engine, then select **Edit <element type>**.

2) Browse to **Interfaces** > **Loopback**.

3) Configure the settings.

4) Click **Save and Refresh**.

Activate the internal DHCP server on a engine interface

You can use the internal DHCP server in a Single Engine or Engine Cluster to assign IPv4 addresses to hosts in the protected network.

This solution is meant for small installations. It might be more practical to assign the IP addresses using the engine rather than relay the DHCP requests from a separately maintained local DHCP server or from some other site's DHCP server through a VPN.

The internal DHCP server can be set up independently on several Physical Interfaces, VLAN Interfaces, and Port Group Interfaces of an integrated Switch. When VLAN or Port Group Interfaces are configured, the DHCP server must be set up separately for each VLAN or Port Group. Only IPv4 addresses are supported. To use this feature, the Engine interface must have at least one IPv4 address.

Note

You can use the internal DHCP server to provide IP addresses to the VPN client Virtual Adapter only if you use Single Engines as VPN gateways.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Security Engine, then select **Edit <element type>**.
2) Browse to **Interfaces**.
3) Right-click a Physical Interface, VLAN Interface, SSID Interface, or Port Group Interface, then select **Edit <interface type>**.
4) On the **DHCPv4** tab, select **DHCPv4 Server** from the **DHCP Mode** drop-down list.
5) Configure the settings, then click **OK**.
6) Click **Save and Refresh** to transfer the new configuration to the Security Engine.

Related concepts

Getting started with Access rules on page 831

Configuring interfaces for IPS engines

The interface configuration for Single IPS engines and IPS Clusters consists of several main steps.

Follow these general steps to configure IPS interfaces:

1) Define the required number of Physical Interfaces.
2) Add the required number of VLANs.

3) Add IP addresses or a traffic inspection role to the interface.
   - IP addresses are required for interfaces that are used for system communications.
   - Interfaces that have a traffic inspection role work transparently in the network and do not have IP addresses.

4) Select the interfaces that are used for system communications.

Heartbeat network for IPS Clusters

The nodes in an IPS cluster exchange status information through a heartbeat network using multicast transmissions.

If an IPS node becomes unavailable, the other nodes of the cluster immediately notice this, and connections are reallocated to the available nodes. A dedicated network is recommended for at least the primary heartbeat communications.

Add system communication interfaces for IPS engines

Normal Interfaces are used for communication between the IPS engine and the Management Server.

Physical Interfaces correspond to network ports on the IPS engine. In an IPS Cluster, each physical interface definition represents a network interface on all nodes of the cluster. By default, the numbering of the Physical Interfaces in the SMC Client corresponds to the operating system interface numbering on the engine. For example, Interface ID 0 is mapped to eth0, and Interface ID 1 is mapped to eth1. However, the mapping is not fixed and you can change it through the engine command line. This mapping can be done differently from node to node as long as you take care that the same interface on each node is correctly cabled to the same network.

In a Single IPS, Normal Interfaces are used for:

- Communication between the IPS engine and the Management Server.
- Sending event information and traffic recordings to Log Servers.
- As the Reset Interface for sending TCP Reset responses.

In an IPS Cluster, Normal Interfaces handle all traffic for which the end-point of the communication is a node itself. Normal Interfaces are used for:

- Heartbeat communication between the nodes.
- Communication between each individual node and the Management Server.
- Sending event information and traffic recordings to Log Servers.
- For any other traffic between the node itself and some other host.

Normal Interfaces in an IPS Cluster are also used as the Reset Interface for sending TCP Reset responses.

Each Single IPS engine or node in an IPS Cluster needs at least one Normal interface for communicating with other SMC components. You can define more than one system communication interface if it is necessary in your network environment. It is recommended to create a separate Normal Interface that is used for communication with the Management Server rather than using the same Normal Interface for sending event information and traffic recordings to Log Servers, and for communication with the Management Server.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Single IPS or IPS Cluster and select **Edit <element type>**.
   The Engine Editor opens.

2) In the navigation pane on the left, select **Interfaces**.
   The **Interfaces** pane opens on the right.

3) Right-click the empty space and select **New Physical Interface**.

4) Define the Physical Interface properties.
   Select **Normal Interface** as the **Type**.

   CAUTION

   Make sure that you set the interface speed correctly. When the bandwidth is set, the IPS engine always scales the total amount of traffic on this interface to the bandwidth you defined. This scaling happens even if there are no bandwidth limits or guarantees defined for any traffic.

5) Click **OK**. The Physical Interface is added to the interface list.

6) Continue the configuration in one of the following ways:
   - Add VLAN interfaces to the Physical Interface.
   - Add IP addresses to the system communication interfaces.

Related concepts

IP addresses for IPS engines on page 576
Defining Zone elements on page 923
Quality of Service (QoS) and how it works on page 973

Add VLAN Interfaces for IPS engines

VLANs divide one physical network link into several virtual links.

A *Virtual Local Area Network* (VLAN) is a logical grouping of hosts and network devices that appear as a single network segment regardless of the physical topology. IPS engines support VLAN tagging as defined in the IEEE 802.1q standard. One physical interface can support up to 4094 VLANs.

VLAN tagging can be used:

- To inspect VLAN tagged traffic (no VLAN Interface configuration required on the IPS engine).
- To define different inspection rules for different VLANs (requires defining VLAN Interfaces for the IPS engine).
- For sending the IPS engine's management and logging connections through a directly connected VLAN segment.

Traffic picked up from a VLAN tagged interface can be inspected without configuring VLAN tagging on the IPS engine. However, you must configure the VLANs on the IPS engine if you want to create different traffic inspection rules for different VLANs. Even then, not all VLANs necessarily have to be specified on the IPS engine. The traffic inspection is customized for the VLANs by defining different Logical Interfaces for the different

VLAN Capture Interfaces. The Logical Interface elements are then used in the IPS Policy rules to define which rules are used for which VLANs.

When you use VLAN Inline Interfaces, the interface numbers must be different and the VLAN identifier must be identical in both of the Inline Interfaces. For example, 3.101 and 4.101 would be a valid pair of VLAN Inline Interfaces. Also, when a VLAN Interface is used for an Inline Interface, it cannot be simultaneously used for any other types of interfaces. The VLAN identifiers you configure on the IPS engine must match the switch or router configuration.

When you use VLAN with Capture Interfaces, the network interface used as the *Reset Interface* for sending TCP Reset responses must be defined in the Capture Interface's properties. The reset is automatically tagged for the same VLAN that triggers a reset. The Reset Interface must be connected to the same VLAN/Broadcast domain as the Capture Interface to reach the communicating hosts.

If the IPS engine encounters unknown VLANs, it might or might not inspect the traffic. This is controlled by the **Inspect Unspecified VLANs** option in the Inline and Capture Interface definitions (by default, the option is set so that all traffic is inspected).

CAUTION

Do not add any manual VLAN definitions to an interface you want to use for sending resets. Adding VLANs prevents selecting the interface as a Reset Interface and also removes the Reset Interface from any existing selections. The IPS engine automatically uses the correct VLAN when sending resets.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Single IPS or IPS Cluster and select **Edit <element type>**.

2) In the navigation pane on the left, select **Interfaces**.

3) Right-click a Physical Interface and select **New** > **VLAN Interface**.

4) Define the VLAN Interface properties.

CAUTION

The throughput for each VLAN Interface must not be higher than the throughput for the Physical Interface to which the VLAN Interface belongs.

CAUTION

Make sure that you set the interface speed correctly. When the bandwidth is set, the IPS engine always scales the total amount of traffic on this interface to the bandwidth you defined. This scaling happens even if there are no bandwidth limits or guarantees defined for any traffic.

CAUTION

The MTU for each VLAN Interface must not be higher than the MTU for the Physical Interface to which the VLAN Interface belongs.

5) Click **OK**.

6) Click **Save and Refresh**.

Result

The VLAN Interface is now ready. The VLAN Interface is identified as Interface-ID.VLAN-ID, for example 2.100 for Interface ID 2 and VLAN ID 100.

Related concepts

IP addresses for IPS engines on page 576
Defining Zone elements on page 923
Quality of Service (QoS) and how it works on page 973

IP addresses for IPS engines

You can add IP addresses to system communication interfaces on IPS engines.

An IPS engine's system communication interfaces (Normal interfaces) can have the following types of IP addresses:

- A Physical Interface can have one or more static or dynamic IP addresses. A Physical Interface can have multiple dynamic IP addresses only if you add VLAN Interfaces on the Physical Interface and the VLAN Interfaces each have a dynamic IP address. Otherwise, a Physical Interface can only have a single dynamic IP address.
- A VLAN Interface can have one or more static IP addresses or a single dynamic IP address.

When a Normal Interface is used for communication with the Management Server, as the Heartbeat Interface in an IPS Cluster, or for communication with the Log Server, an IP Address is needed. When the same Normal Interface that is used for communication with the Management Server and Log Server is also used as a Reset Interface for sending TCP Reset responses, it can have an IP address. When a Normal Interface is used only as a Reset Interface, it must not have an IP address.

All nodes in an IPS Cluster must have the same netmask value for the IP address of their respective Normal Interfaces. The IP addresses specified for each node are used whenever the nodes need to be contacted individually.

You might need to define a contact address if you enter a private static address and NAT is used to translate it to a different external IP address. The external IP address must be configured as the contact address if other SMC components need to use the external IP address to contact the engine.

Related concepts

Define contact IP addresses on page 127

Related tasks

Add IPv4 and IPv6 addresses to IPS Cluster interfaces on page 578

Add IPv4 and IPv6 addresses to Single IPS interfaces

You can add IPv4 and IPv6 addresses to system communication interfaces on Single IPS engines.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Single IPS element and select **Edit Single IPS**.
   The Engine Editor opens.

2) In the navigation pane on the left, select **Interfaces**.
   The **Interfaces** pane opens on the right.

3) Right-click a Physical Interface or VLAN Interface and select **New** > **IPv4 Address** or **New** > **IPv6 Address**.

4) Configure the IPv4 or IPv6 address settings in one of the following ways:
   - Select **Static** and enter the IPv4 Address or IPv6 Address. The Network Settings are automatically entered.

     Tip

     To resolve the IP address from a DNS name, right-click the field, then select **Resolve From DNS Name**.
   - Select **Dynamic** and select the DHCP index from the **Dynamic Index** drop-down list if the interface gets its IP address from a DHCP server. The DHCP Index is a number for your own reference to identify the DHCP interface.

5) (IPv4 addresses and dynamic IPv6 addresses only) If necessary, define the contact address information.
   - Enter the **Default** contact address or select **Dynamic** (next to the **Default** field) if the interface has a dynamic contact address. The Default contact address is used by default whenever a component that belongs to another Location connects to this interface.
   - If components from some Locations cannot use the Default contact address, click **Exceptions** to define Location-specific contact addresses.

6) (Static IPv4 addresses only) Check the automatically filled-in **Netmask** and adjust it as necessary.

7) (Static IPv6 addresses only) Check the automatically filled-in **Prefix Length** and adjust it as necessary.

8) (Optional) Configure additional features for this interface:
   - (Dynamic IPv4 or IPv6 addresses only) If you do not want a default route to be created through the interface, deselect **Automatic Default Route**.
   - (Dynamic IPv6 addresses only) If you want to use DHCPv6 to get the IPv6 address, select **Use DHCPv6 to get IPv6 Address**.

9) Click **OK**.

10) Continue the configuration in one of the following ways:

- If you are creating a new IPS element, add traffic inspection interfaces.
- If you are editing an existing IPS element, select system communication roles for IPS interfaces.

Related concepts
Types of traffic inspection interfaces for IPS engines on page 579

Related tasks
Define contact addresses for a single Security Engine or a Cluster Virtual IP Address on page 138
Select system communication roles for IPS interfaces on page 584

Add IPv4 and IPv6 addresses to IPS Cluster interfaces

You can add IPv4 and IPv6 addresses to system communication interfaces on IPS engines.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click an IPS Cluster and select **Edit IPS Cluster**.
The Engine Editor opens.

2) In the navigation pane on the left, select **Interfaces**.

3) Right-click a physical interface or a VLAN interface and add the IP address in one of the following ways:
   - To add an IPv4 address, select **New** > **IPv4 Address**
   - To add an IPv6 address, select **New** > **IPv6 Address**

4) Click the **IPv4 Address** or **IPv6 Address** cell in the table and enter the IP address for each node.



Tip

To resolve the IP address from a DNS name, right-click the field, then select **Resolve From DNS Name**.

5) (IPv4 addresses only) If necessary, double-click the **Contact Address** cell in the table and define the contact address for each node.
   - In the **Default** field at the top of the dialog box, enter the default contact address. The default contact address is used by default whenever a component that belongs to another Location connects to this interface.
   - If components from some Locations cannot use the default contact address, click **Add** to define Location-specific contact addresses.

6) (IPv4 addresses only) Check the automatically filled-in **Netmask** and adjust it as necessary.

7) (IPv6 addresses only) Check the automatically filled-in **Prefix Length** and adjust it as necessary.

8) Click **OK**.

9) Click **Save**.

10) Continue the configuration in one of the following ways:
- If you are creating a new IPS element, add traffic inspection interfaces.
- If you are editing an existing IPS element, select system communication roles for IPS interfaces.

Related concepts
Types of traffic inspection interfaces for IPS engines on page 579

Related tasks
Define contact addresses for an IPS Cluster or a Layer 2 Engine Cluster on page 140
Select system communication roles for IPS interfaces on page 584

Types of traffic inspection interfaces for IPS engines

Capture Interfaces and Inline Interfaces on IPS engines pick up traffic for inspection.

IPS engines pick up traffic from the network for inspection. There are two ways to install the IPS engines:

- In an IDS-like configuration in which the traffic is only captured for inspection.
- In a full IPS configuration where the IPS engine is installed inline, directly on the traffic path so that traffic must always pass through the IPS engine to reach its destination. Only traffic that attempts to pass through Inline Interfaces can be actively filtered.

Connections picked up through *Capture Interfaces* can be reset through specially set up *Reset Interfaces*. To define a Capture Interface, you must define a Reset Interface for it. Capture Interfaces and *Inline Interfaces* can be defined on the same IPS engine and used simultaneously.

Logical Interface elements allow you to group interfaces together according to network segment and interface type. You can then use the Logical Interface elements as matching criteria when you edit the rules in your IPS policies. There is one predefined Logical Interface element (called **default_eth**) that can be used in interface configurations. If you want to create both Capture and Inline Interfaces on the same IPS engine, you must add at least one more Logical Interface.

Related tasks
Add reset interfaces on page 593
Add Capture Interfaces for IPS engines on page 582
Add Inline Interfaces for IPS engines on page 583

Add logical interfaces

Logical interface elements allow you to group interfaces together according to network segment and interface type.

Logical interfaces are used in the configuration of the following types of interfaces to represent one or more network interfaces:

- Capture interfaces on Engines, IPS Engines, and Layer 2 Engines
- Inline interfaces on IPS engines and Layer 2 Engines
- Inline IPS interfaces on Engines
- Inline Layer 2 Engine interfaces on Engines

You cannot use the same logical interface to represent both capture interfaces and inline interfaces on the same engine. On Engines, you cannot use the same logical interface to represent both inline IPS interfaces and inline Layer 2 Engine interfaces. Otherwise, a logical interface can represent any number or combination of physical interfaces or VLAN Interfaces.

There is one predefined logical interface element called **default_eth**. If you want to create both capture interfaces and inline interfaces on the same Security Engine, you must add at least one more logical interface.

On IPS engines and Layer 2 Engines, a logical interface element called **System Communications** is automatically assigned to interfaces that have an IP address that is used as the primary or backup Control IP address. You can use the **System Communications** logical interface to represent all Control IP addresses in IPS and Layer 2 Engine Policies.



Note

You cannot use the **System Communications** logical interface on engines for Capture interfaces, Inline IPS interfaces, or Inline Layer 2 Engine interfaces.

You can use logical interfaces in IPS Policies, Layer 2 Engine Policies, and Layer 2 Interface Policies to limit the scope of your rules. You can use logical interfaces to create rules that match based on which interface the traffic was picked up from. For example, you can create a different logical interface for each VLAN and use them to create rules that apply only to traffic from a specific VLAN.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Browse to **Other Elements**.
3) Right-click **Logical Interfaces**, then select **New Logical Interface**.
4) Configure the settings, then click **OK**.

Add reset interfaces

Reset interfaces interrupt communications picked up through capture interfaces when the traffic matches a rule that terminates connections.

Reset interfaces can deliver TCP resets and ICMP “destination unreachable” messages to interrupt communications picked up through capture interfaces when the traffic matches a rule that terminates connections.

The resets are sent using the source and destination addresses and MAC addresses of the communicating hosts, so an IP address is not mandatory for a reset interface. You can optionally add an IP address if you also want to use this interface for system communications.

VLANs are supported for sending resets, but the correct VLAN is selected automatically. The interface you want to use as the reset interface must not have any manually added VLAN configuration.

You can use an existing system communications interface for sending resets if the reset interface connects to the same networks as the capture interface, and there are no VLANs on the system communications interface.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click an engine element, then select **Edit <element type>**.

2) In the navigation pane on the left, select **Interfaces**.

3) Create a new Physical Interface.
   - For IPS engines and Layer 2 Engines, right-click the empty space and select **New Physical Interface**.
   - For Engines, right-click the empty space and select **New** > **Layer 3 Physical Interface**.

4) Select an **Interface ID**.

5) Select the interface type according to the engine role:
   - For IPS engines and Layer 2 Engines, select **Normal Interface**.
   - For Engines, select **None**.

6) Click **OK**.

Result

The Physical Interface is added to the interface list.



Note

When you set up the physical network, make sure that the reset interface connects to the same networks as the capture interfaces.

Next steps

Set up the capture interfaces that use this reset interface.

Add Capture Interfaces for IPS engines

Capture Interfaces listen to traffic that is not routed through the IPS engine.

You must define Capture Interfaces if you want to use the IPS engine to inspect traffic that does not flow through the IPS engine. You can have as many Capture Interfaces as there are available physical ports on the IPS engine. External equipment must be set up to mirror traffic to the Capture Interface.

Capture Interfaces have definitions for the corresponding Logical Interface that the interface belongs to. The Logical Interface represents one or more network interfaces that capture the traffic for inspection:

- When a Capture Interface is connected to a switch SPAN port, each Capture Interface is bound to one Logical Interface. More than one Capture Interface can optionally be bound to the same Logical Interface.
- When a network TAP device is used, two Capture Interfaces are bound to the same Logical Interface. The monitored traffic going to different directions is captured through these two related network interfaces and is then combined into a complete traffic flow on the Logical Interface.

You cannot select the same Logical Interface for a Capture Interface and an Inline Interface on the same IPS engine.

A Reset Interface can be selected for a Capture Interface to send TCP Reset responses for the traffic captured from the interface. The Reset Interface is a Normal Interface that can reach the communicating components with the TCP Reset (for example, a Normal Interface connected to the monitored network).

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Single IPS or IPS Cluster and select **Edit <element type>**.
   The Engine Editor opens.

2) In the navigation pane on the left, select **Interfaces**.
   The **Interfaces** pane opens on the right.

3) Right-click the empty space, then select **New Physical Interface**.

4) Define the interface properties.

   Note

   Select **Capture Interface** as the interface **Type**.

5) Click **OK** to save the interface settings.

6) Continue the configuration in one of the following ways:
   - If you are creating a IPS element, select system communication roles for IPS interfaces.
   - Otherwise, click **Save and Refresh** to activate the new interface configuration.

Related concepts

Defining Zone elements on page 923

Add Inline Interfaces for IPS engines

Define Inline Interfaces if you want to position a Single IPS or Inline IPS Cluster directly in the traffic path so that any traffic that is to be inspected goes through the IPS engine.

Inline Interfaces allow traffic to flow through the IPS engine, so that traffic that is deemed harmful can be actively filtered out. An Inline Interface is configured with two Interface IDs, representing two physical interfaces or two VLANs. The IPS engine can inspect the traffic coming from one interface and either stop the traffic or send it out through the other interface. The two interfaces are equal in the configuration. Traffic that is allowed through is always forwarded from one interface to the other (there is no routing decision involved). Inline Interfaces do not have an IP address or a MAC address visible to the network.

Note

Some Forcepoint Network Security Platform appliances use a fail-open network card. Fail-open network cards have fixed pairs of ports. Take particular care to map these ports correctly. Otherwise, the network cards do not correctly fail open when the IPS engine is offline.

In addition to the Interface IDs, Inline Interfaces also have definitions for the corresponding Logical Interface that the interface belongs to. A single Logical Interface can represent one or more pairs of Inline Interfaces. The Logical Interface element can be used to represent the interfaces in the IPS Policy. You cannot select the same Logical Interface for a Capture Interface and an Inline Interface on the same IPS engine.

The number of Inline Interfaces you can configure is limited by the IPS engine's license.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Single IPS or IPS Cluster and select **Edit <element type>**.
   The Engine Editor opens.

2) In the navigation pane on the left, select **Interfaces**.
   The **Interfaces** pane opens on the right.

3) Right-click the empty space, then select **New Physical Interface**.

4) Define the interface properties.
   Select **Inline Interface** as the interface **Type**.

   CAUTION

   Using **Bypass** as the **Failure Mode** requires a fail-open network interface card. If the ports that represent the interfaces cannot fail open, policy installation fails on the Security Engine. Bypass mode is not compatible with VLAN retagging. In network environments where VLAN retagging is used, normal mode is automatically enforced.

5) Click **OK** to save the interface settings.

6) Continue the configuration in one of the following ways:
   - If you are creating an new IPS element, select system communication roles for IPS interfaces.
   - Otherwise, click **Save and Refresh** to activate the new interface configuration.

Related concepts
Defining Zone elements on page 923
Quality of Service (QoS) and how it works on page 973

Select system communication roles for IPS interfaces

You can optionally change which interfaces are used for which types of system communications.

The Interface Options allow you to define:

- Which IP addresses are used as the primary and backup Control IP address
- Which interfaces are used as the primary and backup Heartbeat Interface (IPS Clusters only)
- The default IP address for outgoing traffic

By default, the first IP address you add to a Normal Interface is automatically selected as:

- The primary Control IP address
- The primary Heartbeat Interface (IPS Clusters only)
- The default IP address for outgoing traffic

You can optionally change which interface is used for each of these purposes, and define a backup Control IP address and backup Heartbeat Interface (IPS Clusters only).

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Single IPS or IPS Cluster and select **Edit <element type>**.
   The Engine Editor opens.

2) In the navigation pane on the left, select **Interfaces** > **Interface Options**.

3) In the **Interface Options** pane on the right:
   a) Select the Primary Control IP address for Management Server contact.

   b) (Optional) Select a Backup Control IP address that is used if the Primary Control IP address is not available.

   c) (Optional, single IPS only) If the Security Engine is behind a device that applies dynamic NAT to the inbound management connections or blocks them, select **Node-Initiated contact to Management Server**.
      When this option is selected, the engine opens a connection to the Management Server and maintains connectivity.

d) (IPS Clusters only) Select the **Primary** Heartbeat Interface for communications between the nodes.

We recommend that you use a Physical Interface, not a VLAN Interface. We strongly recommend that you do not direct any other traffic through this interface. A dedicated network helps ensure reliable and secure operation.

CAUTION

Primary and Backup Heartbeat networks exchange confidential information. If dedicated networks are not possible, configure the cluster to encrypt the exchanged information.

e) (IPS Clusters only) Select a Backup Heartbeat Interface that is used if the Primary Heartbeat Interface is unavailable.

It is not mandatory to configure a backup Heartbeat Interface, but we strongly recommend it. If heartbeat traffic is not delivered, the cluster cannot operate and traffic is disturbed. We strongly recommend that you use a dedicated interface for the backup heartbeat as well.

f) In the **Default IP Address for Outgoing Traffic** field, select the IP address that the nodes use if they have to initiate connections through an interface that has no Node Dedicated IP Address.

4) Continue the configuration in one of the following ways:
   - If you are creating a new IPS element, configure the routing.
   - Otherwise, click **Save and Refresh** to activate the new interface configuration.

Related concepts

Connect Security Engines to the SMC on page 631

Related tasks

Adjust IPS clustering options on page 679
Add routes for IPS engines or Layer 2 Engines

Configuring interfaces for Layer 2 Engines

The interface configuration for Single Layer 2 Engines and Layer 2 Engine Clusters consists of several main steps.

1) Add the required number of Physical Interfaces.

2) Add the required number of VLAN Interfaces.

3) Add IP addresses or a traffic inspection role to the interfaces.
   - IP addresses are required for interfaces that are used for system communications.

- Interfaces that have a traffic inspection role work transparently in the network and do not have IP addresses.

4) Select the interfaces that are used for system communications.

Network interfaces for using Layer 2 Engines in Passive Engine mode

You can configure Layer 2 Engines in Passive Engine mode. In Passive Engine mode, the engine captures network traffic for inspection but does not actively filter traffic.

The most common way to configure a Layer 2 Engine in Passive Engine mode is to define Capture Interfaces for listening to network traffic that does not flow through the Layer 2 Engine. If you configure only Capture Interfaces, the engine always functions in Passive Engine mode. You can also use a Layer 2 Engine that has Inline Interfaces in Passive Engine mode. To do this, you configure the engine to only log connections.

Add system communication interfaces for Layer 2 Engines

Normal Interfaces are used for communication between the Layer 2 Engine and the Management Server.

A Physical Interface element corresponds to a network port on the Layer 2 Engine. By default, the numbering of the Physical Interfaces in the SMC Client corresponds to the operating system interface numbering on the engine. For example, Interface ID 0 is mapped to eth0, and Interface ID 1 is mapped to eth1. However, the mapping is not fixed and you can change it through the engine command line.

In a Single Layer 2 Engine, Normal Interfaces are used for the following types of communication:

- Communication between the Layer 2 Engine and the Management Server.
- For sending event information and traffic recordings to Log Servers.

In a Layer 2 Engine Cluster, Normal Interfaces handle all traffic for which the endpoint of the communication is a node itself. Normal Interfaces are used for the following types of communication:

- The Heartbeat communication between the nodes.
- Communication between each individual node and the Management Server.
- For sending event information and traffic recordings to Log Servers.
- For any other traffic between the node itself and some other host.

Each Layer 2 Engine needs at least one interface for communicating with other SMC components. You can define more than one system communication interface if it is necessary in your network environment. Multiple Normal Interfaces can be configured for the same physical network interface. It is recommended to create a separate Normal Interface for each of the following uses:

- For communication with the Management Server.
- For sending event information and traffic recordings to Log Servers.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Single Layer 2 Engine or Layer 2 Engine Cluster and select **Edit <element type>**.
   The Engine Editor opens.

2) In the navigation pane on the left, select **Interfaces**.

3) In the **Interfaces** pane on the right, right-click the empty space and select **New Physical Interface**.

4) Define the Physical Interface properties.
   Select **Normal Interface** as the **Type**.

   CAUTION

   Make sure that you set the Interface speed correctly if you enter a Throughput value for QoS. When the bandwidth is set, the Layer 2 Engine always scales the total amount of traffic on this interface to the bandwidth you defined. The bandwidth is scaled even if there are no bandwidth limits or guarantees defined for any traffic.

5) Click **OK**. The Physical Interface is added to the interface list.

6) Continue the configuration in one of the following ways:
   - If the system communications must use a particular VLAN on the directly connected network segment, add VLAN interfaces.
   - Otherwise, add IP addresses to the system communication interfaces.

Related concepts

IP addresses for Layer 2 Engines on page 589
Defining Zone elements on page 923
Quality of Service (QoS) and how it works on page 973

Add VLAN interfaces for Layer 2 Engines

VLANs divide a single physical network link into several virtual links.

VLANs divide a single physical network link into several virtual links. VLANs can be defined for both single and clustered Layer 2 Engines. The maximum number of VLANs for a single Physical Interface or Inline Interface is 4094. The VLANs must also be defined in the configuration of the switch or router to which the interface is connected.

Traffic picked up from a VLAN tagged interface can be inspected without configuring VLAN tagging on the Layer 2 Engine. However, you must configure VLANs on the Layer 2 Engine if you want to create different traffic inspection rules for different VLANs. Even then, not all VLANs necessarily have to be specified on the Layer 2 Engine. VLANs can optionally also be used for sending the Layer 2 Engine's management and logging connections through a directly connected VLAN segment.

By default, all VLAN traffic is inspected in the same way as non-VLAN traffic. Configure VLAN Interfaces for the physical interfaces if you want to customize traffic inspection for the different VLANs. The traffic inspection is customized for the VLANs by defining different Logical Interfaces for the different VLAN Interfaces. The Logical

Interface elements are then used in the Layer 2 Engine Policy rules to define which rules are used for which VLANs.

If the Layer 2 Engine encounters unknown VLANs, it might or might not inspect the traffic. The **Inspect Unspecified VLANs** option in the Inline Interface definitions defines whether the Layer 2 Engine inspects the traffic. By default, the option is set so that all traffic is inspected.

When you use VLANs with Inline Interfaces, the interface numbers must be different and the VLAN identifier must be identical in both of the Inline Interfaces. For example, 3.101 and 4.101 would be a valid pair of VLAN Inline Interfaces. Also, when a VLAN Interface is used for an Inline Interface, it cannot be simultaneously used for any other types of interfaces.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Single Layer 2 Engine or Layer 2 Engine Cluster and select **Edit <element type>**.
   The Engine Editor opens.

2) In the navigation pane on the left, select **Interfaces**.
   The **Interfaces** pane opens on the right.

3) Right-click a Physical Interface and select **New** > **VLAN Interface**.

4) Define the VLAN Interface properties.

   CAUTION

   The throughput for each VLAN Interface must not be higher than the throughput for the Physical Interface to which the VLAN Interface belongs.

   CAUTION

   Make sure that you set the interface speed correctly. When the bandwidth is set, the Layer 2 Engine always scales the total amount of traffic on this interface to the bandwidth you defined. The bandwidth is scaled even if there are no bandwidth limits or guarantees defined for any traffic.

   CAUTION

   The MTU for each VLAN Interface must not be higher than the MTU for the Physical Interface to which the VLAN Interface belongs.

5) Click **OK**.
   The specified VLAN ID is added to the Physical Interface.

6) Continue the configuration in one of the following ways:
   - If you added a VLAN Interface to a Normal interface, add IP addresses to the VLAN Interface.
   - If you added a VLAN Interface to an Inline Interface, the Inline Interface is ready to use. Click **Save and Install** to activate the new interface configuration.

Related concepts
IP addresses for Layer 2 Engines on page 589
Types of traffic inspection interfaces for Layer 2 Engines on page 592
Defining Zone elements on page 923
Quality of Service (QoS) and how it works on page 973

Related tasks
Add system communication interfaces for Layer 2 Engines on page 586

IP addresses for Layer 2 Engines

You can add IP addresses to system communication interfaces on Layer 2 Engines.

A Layer 2 Engine's system communication interfaces (Normal interfaces) can have the following types of IP addresses:

- A Physical Interface can have one or more static or dynamic IP addresses. A Physical Interface can have multiple dynamic IP addresses only if you add VLAN Interfaces on the Physical Interface and the VLAN Interfaces have dynamic IP addresses. Otherwise, a Physical Interface can only have a single dynamic IP address.
- A VLAN Interface can have one or more static IP addresses or a single dynamic IP address.

You might need to define a contact address if you enter a private static address and NAT is used to translate it to a different external IP address. The external IP address must be configured as the contact address if other SMC components need to use the external IP address to contact the engine.

Related concepts
Define contact IP addresses on page 127

Related tasks
Add IPv4 and IPv6 addresses to Single Layer 2 Engine interfaces on page 589
Add IPv4 and IPv6 addresses to Layer 2 Engine Cluster interfaces on page 591

Add IPv4 and IPv6 addresses to Single Layer 2 Engine interfaces

You can add IPv4 and IPv6 addresses to system communication interfaces on Single Layer 2 Engines.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Single Layer 2 Engine and select **Edit Single Layer 2 Engine**.
The Engine Editor opens.

2) In the navigation pane on the left, select **Interfaces**.
The **Interfaces** pane opens on the right.

3) Right-click a Physical Interface or a VLAN Interface and select **New** > **IPv4 Address** or **New** > **IPv6 Address**.

4) Configure the IPv4 or IPv6 address settings in one of the following ways:
- Select **Static** and enter the IPv4 Address or IPv6 Address. The Network Settings are automatically entered.



Tip

To resolve the IP address from a DNS name, right-click the field, then select **Resolve From DNS Name**.

- Select **Dynamic** and select the DHCP index from the **Dynamic Index** drop-down list if the interface gets its IP address from a DHCP server. The DHCP Index is a number for your own reference to identify the DHCP interface.

5) (IPv4 addresses and dynamic IPv6 addresses only) If necessary, define the contact address information.
- Enter the default contact address or select **Dynamic** (next to the **Default** field) if the interface has a dynamic contact address. The default contact address is used by default whenever a component that belongs to another Location connects to this interface.
- If components from some Locations cannot use the default contact address, click **Exceptions** to define Location-specific contact addresses.

6) (Static IPv4 addresses only) The **Netmask** field is automatically populated. Adjust it as necessary.

7) (Static IPv6 addresses only) The **Prefix Length** field is automatically populated. Adjust it as necessary.

8) (Optional) Configure additional features for this interface:
- (Dynamic IPv4 or IPv6 addresses only) If you do not want a default route to be created through the interface, deselect **Automatic Default Route**.
- (Dynamic IPv6 addresses only) If you want to use DHCPv6 to get the IPv6 address, select **Use DHCPv6 to get IPv6 Address**.

9) Click **OK**.

10) Continue the configuration in one of the following ways:
- If you are defining a new Layer 2 Engine element, add traffic inspection interfaces.
- If you are editing an existing Layer 2 Engine element, select system communication roles for Layer 2 Engine interfaces.

Related concepts

Types of traffic inspection interfaces for Layer 2 Engines on page 592

Related tasks

Define contact addresses for a single Security Engine or a Cluster Virtual IP Address on page 138
Select system communication roles for Layer 2 Engine interfaces on page 596

Add IPv4 and IPv6 addresses to Layer 2 Engine Cluster interfaces

You can add IPv4 and IPv6 addresses to system communication interfaces on Layer 2 Engine Clusters.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Layer 2 Engine Cluster and select **Edit Layer 2 Engine Cluster**.
The Engine Editor opens.

2) In the navigation pane on the left, select **Interfaces**.

3) Right-click a physical interface or a VLAN interface and add the IP address in one of the following ways:
   - To add an IPv4 address, select **New** > **IPv4 Address**
   - To add an IPv6 address, select **New** > **IPv6 Address**

4) Click the **IPv4 Address** or **IPv6 Address** cell in the table and enter the IP address for each node.

   Tip

   To resolve the IP address from a DNS name, right-click the field, then select **Resolve From DNS Name**.

5) (IPv4 addresses only) If necessary, double-click the **Contact Address** cell in the table and define the contact address for each node.
   - In the **Default** field at the top of the dialog box, enter the default contact address. The default contact address is used by default whenever a component that belongs to another Location connects to this interface.
   - If components from some Locations cannot use the default contact address, click **Add** to define Location-specific contact addresses.

6) (IPv4 addresses only) Check the automatically filled-in **Netmask** and adjust it as necessary.

7) (IPv6 addresses only) Check the automatically filled-in **Prefix Length** and adjust it as necessary.

8) Click **OK**.

9) Click **Save**.

10) Continue the configuration in one of the following ways:
   - If you are creating a new Layer 2 Engine Cluster element, add traffic inspection interfaces.

- If you are editing an existing Layer 2 Engine Cluster element, select system communication roles for IPS interfaces.

Related tasks

Define contact addresses for an IPS Cluster or a Layer 2 Engine Cluster on page 140
Select system communication roles for Layer 2 Engine interfaces on page 596

Types of traffic inspection interfaces for Layer 2 Engines

Capture Interfaces and Inline Interfaces on Layer 2 Engines pick up traffic for inspection.

Layer 2 Engines inspect network traffic. Layer 2 Engines are typically installed inline, directly on the traffic path so that traffic must always pass through the Layer 2 Engine to reach its destination. Only traffic that attempts to pass through *Inline Interfaces* can be actively filtered.

You can also configure a Layer 2 Engine in Passive Engine mode. In Passive Engine mode, a Layer 2 Engine has *Capture Interfaces* defined for inspections that listen to and log network traffic.

Connections picked up through *Capture Interfaces* can be reset through specially set-up Reset Interfaces. Capture Interfaces and Inline Interfaces can be defined on the same Layer 2 Engine and used simultaneously.

Logical Interface elements allow you to group interfaces together according to network segment. You can then use the Logical Interface elements as matching criteria when you edit the rules in your Layer 2 Engine policies.

Related tasks

Add logical interfaces on page 592
Add Capture Interfaces for IPS engines on page 582
Add Inline Interfaces for IPS engines on page 583

Add logical interfaces

Logical interface elements allow you to group interfaces together according to network segment and interface type.

Logical interfaces are used in the configuration of the following types of interfaces to represent one or more network interfaces:

- Capture interfaces on Engines, IPS Engines, and Layer 2 Engines
- Inline interfaces on IPS engines and Layer 2 Engines
- Inline IPS interfaces on Engines
- Inline Layer 2 Engine interfaces on Engines

You cannot use the same logical interface to represent both capture interfaces and inline interfaces on the same engine. On Engines, you cannot use the same logical interface to represent both inline IPS interfaces and inline Layer 2 Engine interfaces. Otherwise, a logical interface can represent any number or combination of physical interfaces or VLAN Interfaces.

There is one predefined logical interface element called **default_eth**. If you want to create both capture interfaces and inline interfaces on the same Security Engine, you must add at least one more logical interface.

On IPS engines and Layer 2 Engines, a logical interface element called **System Communications** is automatically assigned to interfaces that have an IP address that is used as the primary or backup Control IP address. You can use the **System Communications** logical interface to represent all Control IP addresses in IPS and Layer 2 Engine Policies.

Note

You cannot use the **System Communications** logical interface on engines for Capture interfaces, Inline IPS interfaces, or Inline Layer 2 Engine interfaces.

You can use logical interfaces in IPS Policies, Layer 2 Engine Policies, and Layer 2 Interface Policies to limit the scope of your rules. You can use logical interfaces to create rules that match based on which interface the traffic was picked up from. For example, you can create a different logical interface for each VLAN and use them to create rules that apply only to traffic from a specific VLAN.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Browse to **Other Elements**.
3) Right-click **Logical Interfaces**, then select **New Logical Interface**.
4) Configure the settings, then click **OK**.

Related tasks

Add Capture Interfaces for IPS engines on page 582
Add Inline Interfaces for IPS engines on page 583
Add Capture Interfaces for Layer 2 Engines on page 594
Add Inline Interfaces for Layer 2 Engines on page 595

Add reset interfaces

Reset interfaces interrupt communications picked up through capture interfaces when the traffic matches a rule that terminates connections.

Reset interfaces can deliver TCP resets and ICMP “destination unreachable” messages to interrupt communications picked up through capture interfaces when the traffic matches a rule that terminates connections.

The resets are sent using the source and destination addresses and MAC addresses of the communicating hosts, so an IP address is not mandatory for a reset interface. You can optionally add an IP address if you also want to use this interface for system communications.

VLANs are supported for sending resets, but the correct VLAN is selected automatically. The interface you want to use as the reset interface must not have any manually added VLAN configuration.

You can use an existing system communications interface for sending resets if the reset interface connects to the same networks as the capture interface, and there are no VLANs on the system communications interface.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click an engine element, then select **Edit <element type>**.
2) In the navigation pane on the left, select **Interfaces**.
3) Create a new Physical Interface.
   - For IPS engines and Layer 2 Engines, right-click the empty space and select **New Physical Interface**.
   - For Engines, right-click the empty space and select **New** > **Layer 3 Physical Interface**.
4) Select an **Interface ID**.
5) Select the interface type according to the engine role:
   - For IPS engines and Layer 2 Engines, select **Normal Interface**.
   - For Engines, select **None**.
6) Click **OK**.

Result

The Physical Interface is added to the interface list.

Note

When you set up the physical network, make sure that the reset interface connects to the same networks as the capture interfaces.

Next steps

Set up the capture interfaces that use this reset interface.

Related tasks

Add system communication interfaces for IPS engines on page 573
Add system communication interfaces for Layer 2 Engines on page 586
Add Capture Interfaces for Layer 2 Engines on page 594

Add Capture Interfaces for Layer 2 Engines

Capture Interfaces listen to traffic that is not routed through the Layer 2 Engine.

You can have as many Capture Interfaces as there are available physical ports on the Layer 2 Engine. External equipment must be set up to mirror traffic to the Capture Interface.

Capture Interfaces are used to configure a Layer 2 Engine in Passive Engine mode. In Passive Engine mode, the Layer 2 Engine only listens to traffic that does not flow through it.

Capture Interfaces have definitions for the corresponding Logical Interface that the interface belongs to. The Logical Interface represents one or more network interfaces that capture the traffic for inspection. Each Capture Interface is bound to one Logical Interface. More than one Capture Interface can optionally be bound to the same Logical Interface.

You cannot select the same Logical Interface for a Capture Interface and an Inline Interface on the same Layer 2 Engine.

A Reset Interface can be selected for a Capture Interface to send TCP Reset responses for the traffic captured from the interface. The Reset Interface is a Normal Interface that can reach the communicating components with a TCP Reset (for example, a Normal Interface connected to the monitored network).

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Single Layer 2 Engine or Layer 2 Engine Cluster and select **Edit <element type>**.
   The Engine Editor opens.

2) In the navigation pane on the left, select **Interfaces**.
   The **Interfaces** pane opens on the right.

3) Right-click the empty space and select **New Physical Interface**.

4) Define the interface properties.

   Note

   Select **Capture Interface** as the interface **Type**.

5) Click **OK** to save the interface settings.

6) Continue the configuration in one of the following ways:
   - If you are creating a new Layer 2 Engine element, select system communication roles for Layer 2 Engine interfaces.
   - If you added an interface to an existing Layer 2 Engine element, click **Save and Refresh** to activate the new interface configuration.

Related concepts

Defining Zone elements on page 923

Add Inline Interfaces for Layer 2 Engines

Inline Interfaces allow traffic to flow through the Layer 2 Engine, so that traffic that is harmful can be actively filtered out.

An Inline Interface consists of two physical interfaces or two VLANs. This way, the Layer 2 Engine can inspect the traffic coming from one interface and either stop the traffic or send it out through the other interface. Traffic that is allowed through is always forwarded from one interface to the other (there is no routing decision involved in the interface selection).

In addition to the Interface IDs, Inline Interfaces also have definitions for the corresponding Logical Interface that this interface belongs to. A single Logical Interface can represent one or more pairs of Inline Interfaces. The Logical Interface element can be used to represent the interfaces in the Layer 2 Engine Policy. You cannot select the same Logical Interface for a Capture Interface and an Inline Interface on the same Layer 2 Engine.

Some Forcepoint Network Security Platform appliances have a fail-open network card, so the Inline Interfaces must be configured for those specific ports. Inline Interfaces do not have an IP address or a MAC address visible to the network.

The number of Inline Interfaces you can configure is limited by the Layer 2 Engine's license.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Single Layer 2 Engine or Layer 2 Engine Cluster and select **Edit <element type>**.
   The Engine Editor opens.

2) In the navigation pane on the left, select **Interfaces**.
   The **Interfaces** pane opens on the right.

3) Right-click the empty space and select **New Physical Interface**.

4) Define the interface properties.

   Note

   Select **Inline Interface** as the interface **Type**.

5) Click **OK** to save the interface settings.

6) Continue the configuration in one of the following ways:
   - If you are creating a new Layer 2 Engine element, select system communication roles for Layer 2 Engine interfaces.
   - If you added an interface to an existing Layer 2 Engine element, click **Save and Refresh** to activate the new interface configuration.

Related concepts

Defining Zone elements on page 923
Quality of Service (QoS) and how it works on page 973

Select system communication roles for Layer 2 Engine interfaces

Interface options allow you to select which interfaces are used for which types of system communications.

You can define the following settings for system communication:

- Which IP addresses are used as the primary and backup Control IP address
- Which interfaces are used as the primary and backup Heartbeat Interface (Layer 2 Engine Clusters only)
- The default IP address for outgoing traffic

By default, the first IP address you add to a Normal Interface is automatically selected for the following roles:

- As the primary Control IP address

- As the primary Heartbeat Interface (Layer 2 Engine Clusters only)
- As the default IP address for outgoing traffic

You can optionally change which physical interface is used for each of these purposes. You can also define a backup Control IP address and backup Heartbeat Interface (Layer 2 Engine Clusters only).

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Single Layer 2 Engine or Layer 2 Engine Cluster and select **Edit <element type>**.
   The Engine Editor opens.

2) In the navigation pane on the left, select **Interfaces** > **Interface Options**.

3) In the **Interface Options** pane on the right:
   a) Select the Primary Control IP address for Management Server contact.

   b) (Optional) Select a Backup Control IP address that is used if the Primary Control IP address is not available.

   c) (Optional, single Layer 2 Engine only) If the Security Engine is behind a device that applies dynamic NAT to the inbound management connections or blocks them, select **Node-Initiated contact to Management Server**.
   When this option is selected, the engine opens a connection to the Management Server and maintains connectivity.

   d) (Layer 2 Engine Clusters only) Select the primary Heartbeat Interface for communications between the nodes.
   We recommend that you use a Physical Interface, not a VLAN Interface. We strongly recommend that you do not direct any other traffic through this interface. A dedicated network helps guarantee reliable and secure operation.

   CAUTION

   Primary and Backup Heartbeat networks exchange confidential information. If dedicated networks are not possible, configure the cluster to encrypt the exchanged information.

   e) (Layer 2 Engine Clusters only) Select a backup Heartbeat Interface that is used if the Primary Heartbeat Interface is unavailable.
   It is not mandatory to configure a backup Heartbeat Interface, but we strongly recommend it. If heartbeat traffic is not delivered, the cluster cannot operate and traffic is disturbed. We strongly recommend that you use a dedicated interface for the backup heartbeat as well.

   f) In the **Default IP Address for Outgoing Traffic** field, select the IP address that the nodes use if they initiate connections through an interface that has no Node Dedicated IP Address.

4) Click **OK**.

5) Continue the configuration in one of the following ways:
   - If you are creating a new Layer 2 Engine element, configure the routing.

- Otherwise, click **Save and Refresh** to transfer the configuration changes.

Related concepts
Connect Security Engines to the SMC on page 631

Related tasks
Adjust Layer 2 Engine clustering options on page 680
Add routes for IPS engines or Layer 2 Engines

Configuring interfaces for Master Engines

Master Engines can have two types of interfaces: interfaces for the Master Engine's own traffic, and interfaces that are used by the Virtual Engines hosted on the Master Engine.

You can add Physical Interfaces and VLAN Interfaces to a Master Engine. If you want to use a Physical Interface or VLAN Interface to host a Virtual Engine, you must select a Virtual Resource for the interface. The same Virtual Resource can be used on more than one Master Engine interface to allocate multiple interfaces to the same Virtual Engine. If you want the Virtual Engine to have multiple interfaces, you must use the same Virtual Resource on more than one Master Engine interface.

If you want to use a Physical Interface or VLAN Interface for the Master Engine’s system communications, you can add IP addresses to either:

- An interface that does not have a Virtual Resource assigned to it
- A shared interface that has Virtual Resources assigned to it

By default, the Physical Interface definitions for the Master Engine are mapped to the actual network interfaces on the Master Engine hardware in numerical order. If necessary, you can change the mapping using command-line tools on the Master Engine. This mapping can be done differently from one Master Engine node to another. Make sure that the interface that represents the same network interface on each Master Engine node is correctly cabled to the same network.

Shared interfaces

A shared interface is a single layer 3 physical interface on the Master Engine in the Engine/VPN role. You can assign multiple Virtual Resources to the interface, so the interface can be shared by multiple Virtual Engines. The shared interface can also have shared VLANs underneath it.

Aggregated interfaces

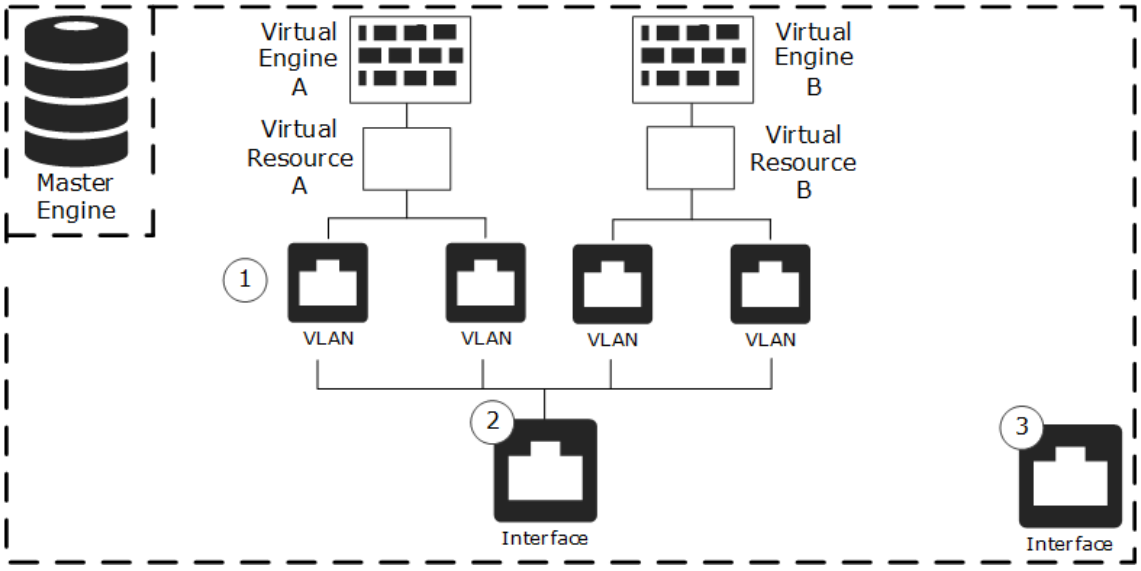
To use an aggregated interface as an interface for a Virtual Engine, you must do one of the following:

- Make the aggregate interface a shared interface.
- Make the aggregate interface a shared interface, add shared VLAN interfaces to the interface, then assign the Virtual Resources to the shared VLAN interfaces.

- Add VLAN interfaces to a regular aggregate interface, then assign the Virtual Resources to the VLAN interfaces.

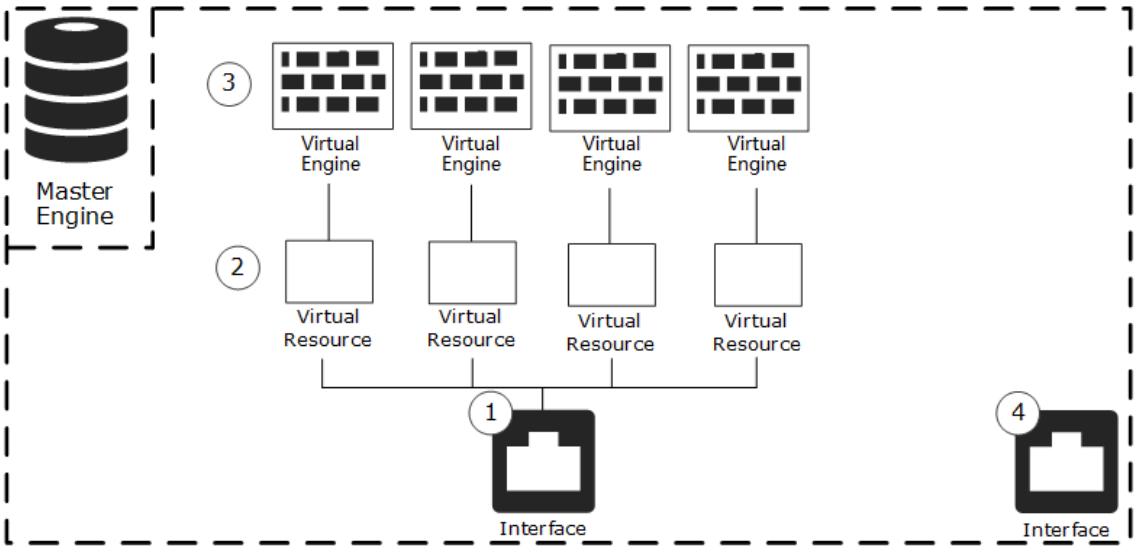
Interface examples

Example of Master Engine and Virtual Engines



1. VLAN Interfaces for hosted Virtual Engine traffic.
2. Physical Interface for hosted Virtual Engine traffic.
3. Physical Interface for the Master Engine system communications.

Example of shared interface on a Master Engine and Virtual Engines



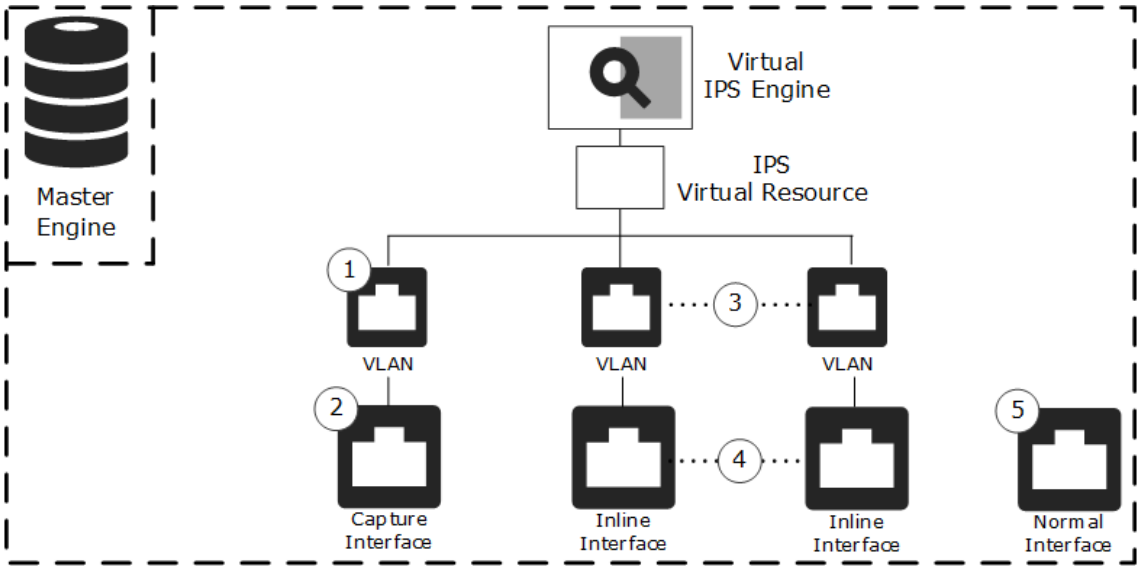
1 Shared physical interface for hosted Virtual Engine traffic.

2 Virtual Resources selected for the shared physical interface.

2 Virtual Engines associated with the Virtual Resources.

4 Physical Interface for the Master Engine system communications.

Example of Master Engine and Virtual IPS engines



1 VLAN Interface for hosted Virtual IPS engine traffic.

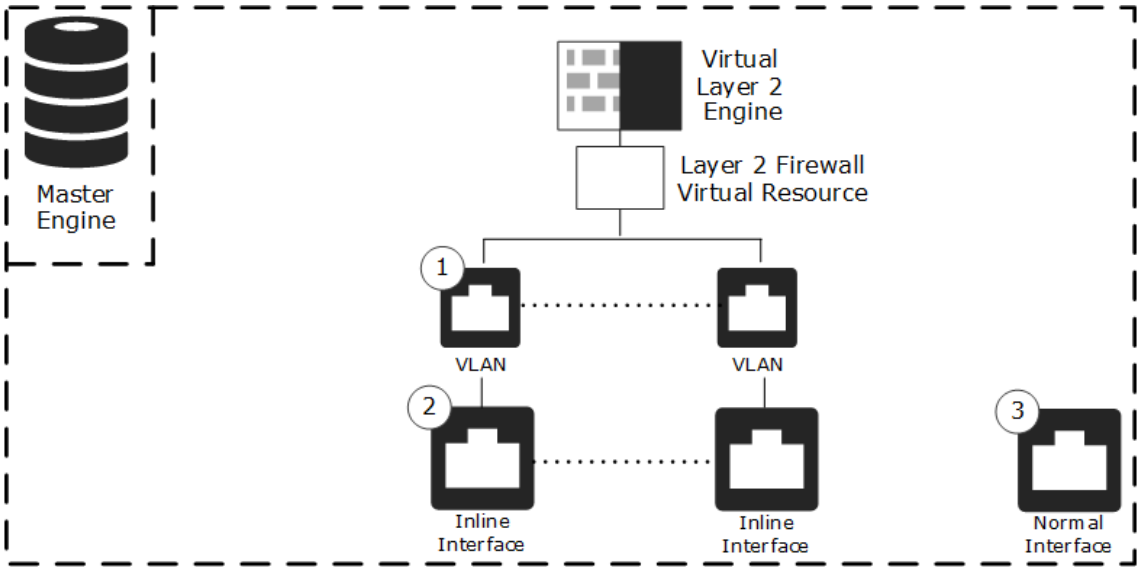
2 Capture Interface for hosted Virtual IPS engine traffic.

3 Inline VLAN Interface pair for hosted Virtual IPS engine traffic.

4 Inline Interface pair for hosted Virtual IPS engine traffic.

5 Physical Interface for the Master Engine system communications.

Example of Master Engine and Virtual Layer 2 Engine



1 Inline VLAN Interface for hosted Virtual Layer 2 Engine traffic.

2 Inline Interface for hosted Virtual Layer 2 Engine traffic.

3 Physical Interface for the Master Engine system communications.

Master Engine interface configuration overview

The interface configuration for Master Engines consists of several main steps.

1) Define Physical Interfaces for both the Master Engine system communications, and for hosted Virtual Engines.

2) (Optional) Add VLAN Interfaces.

3) Add IP addresses to the physical interfaces used for system communications.
   - IP addresses are required for interfaces that are used for system communications.
   - You cannot add both an IP address and a Virtual Resource to the same Physical Interface or VLAN Interface.

4) Select which interfaces are used for different types of system communications.

Related concepts

Configuring interfaces for Virtual Engines on page 609

System communication interfaces for Master Engines

Physical Interfaces correspond to network ports on the Master Engine. By default, the numbering of the Physical Interfaces in the SMC Client corresponds to the operating system interface numbering on the engine. For example, Interface ID 0 is mapped to eth0, and Interface ID 1 is mapped to eth1. However, the mapping is not fixed and you can change it through the Security Engine command line.

The types of Physical Interfaces that you can define for the Master Engine system communications depend on the role of the hosted Virtual Engines:

Master Engine layer 3 interface types for system communications

Role	Interface Type	Explanation
Virtual Engine	None	Corresponds to a single network interface on the Master Engine appliance.
	Aggregated Link in High Availability Mode	Represents two interfaces on the Master Engine appliance. Only the first interface in the aggregated link is actively used. The second interface becomes active only if the first interface fails. If you configure an Aggregated Link in High Availability mode, connect the first interface to one switch and the second interface to another switch.
	Aggregated Link in Load Balancing Mode	Represents up to eight interfaces on the Master Engine appliance. All interfaces in the aggregated link are actively used and connections are automatically balanced between the interfaces. Link aggregation in the Load Balancing Mode is implemented based on the IEEE 802.3ad Link Aggregation standard. If you configure an Aggregated Link in Load Balancing Mode, connect all interfaces to a single switch. Make sure that the switch supports the Link Aggregation Control Protocol (LACP) and that LACP is configured on the switch.
Virtual IPS	Normal Interface	Corresponds to a single network interface on the Master Engine appliance. Only Normal Interfaces can be used for Master Security Engine system communications when the hosted Virtual Engines are in the Virtual IPS role.
Virtual Layer 2 Engine	Normal Interface	Corresponds to a single network interface on the Master Engine appliance. Only Normal Interfaces can be used for Master Security Engine system communications when the hosted Virtual Engines are in the Virtual Layer 2 Engine role.

If the Master Engine is a cluster, it is recommended to add at least two layer 3 Physical Interfaces for the Master Engine:

- An interface for communications between the Management Server and the Master Engine.

  Note

  We recommend that you do not use the IP address of an Aggregated Link interface as the primary or secondary control IP address of the Engine.

- An interface for the heartbeat communications between the Master Engine nodes. The heartbeat traffic is critical to the functioning of the cluster, so it is highly recommended to have a dedicated physical interface as the heartbeat interface.
  You cannot use a shared interface as a heartbeat interface.

Add physical interfaces for Master Engines

In addition to physical interfaces that are used for hosted Virtual Engine communications, you must add at least one physical interface for system communications.

The types of interfaces that you can configure for the Master Engine depend on the role of the Virtual Engines that the Master Engine hosts. For Master Engines that host Virtual Engines, you can add Layer 3 physical interfaces and Layer 2 physical interfaces.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Security Engine, then select **Edit <element type>**.
2) Browse to **Interfaces**.
3) Depending on the role of the hosted Virtual Engine, select **Add** > **Layer 3 Physical Interface** or **Layer 2 Physical Interface**.
4) Configure the settings, then click **OK**.

Related tasks

Add IPv4 and IPv6 addresses to Master Engine interfaces on page 605
Select the same Virtual Resource for multiple interfaces on page 608

Select system communication roles for Master Engine interfaces

Select which Master Engine interfaces are used for particular roles in system communications.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Security Engine, then select **Edit <element type>**.
2) Browse to **Interfaces** > **Interface Options**.
3) In the **Interface Options** pane:
   a) From the **Primary** control IP address drop-down list, select the primary control IP address that the Master Engine uses for communications with the Management Server.



Note

We recommend that you do not use the IP address of an Aggregated Link interface as the primary or secondary control IP address of the Engine.

b) (Optional, recommended) From the **Backup** control IP address drop-down list, select a backup control IP address that the Master Engine uses for communications with the Management Server if the primary control IP address fails.

c) (Optional, Security Engine with Layer 3 Interfaces only) If the Master Engine is behind a device that applies dynamic NAT to outbound connections or in some other way blocks incoming connections, select **Node-Initiated contact to Management Server**.

When this option is selected, the engine opens a connection to the Management Server and maintains connectivity.

d) (Master Engine Cluster Only) From the **Primary** heartbeat drop-down list, select the primary interface for communications between the nodes.

We recommend using a physical interface, not a VLAN interface. We strongly recommend that you do not direct any other traffic through this interface. A dedicated network helps guarantee reliable and secure operation.

CAUTION

Primary and backup heartbeat networks exchange confidential information. If dedicated networks are not possible, configure the cluster to encrypt the exchanged information.

e) (Master Engine Cluster Only) From the **Backup** heartbeat drop-down list, select the backup heartbeat interface that is used if the primary heartbeat interface is unavailable.

It is not mandatory to configure a backup heartbeat interface, but we strongly recommend it. If heartbeat traffic is not delivered, the cluster cannot operate and traffic is disturbed. We strongly recommend that you use a dedicated interface for the backup heartbeat as well.

f) In the **Default IP Address for Outgoing Traffic** field, select the IP address that the nodes use if they have to initiate connections through an interface that has no Node Dedicated IP address.

4) Click **Save and Refresh**.

Next steps

Bind licenses to Master Engine elements.

Related concepts

Adjust Master Engine clustering options on page 681

Add VLAN interfaces to Master Engines

Master Engines can have two types of VLAN interfaces: VLAN interfaces for the Master Engine's own traffic, and VLAN interfaces that are used by the Virtual Engines hosted on the Master Engine.

The maximum number of VLANs for a single physical interface is 4094. The VLANs must also be defined in the configuration of the external switch or router to which the interface is connected.

On Master Engines that host Virtual IPS engines or Virtual Layer 2 Engines, the Virtual Engines can inspect traffic from VLAN interfaces without configuring VLAN tagging.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click a Security Engine, then select **Edit <element type>**.
3) Browse to **Interfaces**.
4) Right-click a physical interface, then select **New** > **VLAN Interface**.
5) Configure the settings, then click **OK**.
   The specified VLAN ID is added to the physical interface.
6) Click **Save**.

Next steps

Add IP addresses to the physical interfaces or VLAN interfaces for Master Engine system communications.

Related concepts

Defining Zone elements on page 923
Quality of Service (QoS) and how it works on page 973

Related tasks

Add IPv4 and IPv6 addresses to Master Engine interfaces on page 605
Select the same Virtual Resource for multiple interfaces on page 608

Add IPv4 and IPv6 addresses to Master Engine interfaces

You can add several IPv4 addresses to each Physical Interface or VLAN Interface that does not have a Virtual Resource associated with it.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click a Master Engine, then select **Edit <element type>**.
3) Browse to **Interfaces**.

4) Right-click a physical interface or a VLAN interface, then select **New** > **IPv4 Address** or **New** > **IPv6 Address**.

5) Configure the settings, then click **OK**.

6) Click **Save**.

Next steps

Continue the configuration in one of the following ways:

- If you are configuring a new Master Engine, or if you want to change the roles the different interfaces have in the configuration, select system communication roles for Master Engine interfaces.
- Otherwise, refresh the policy to transfer the configuration changes.

Related concepts

Define contact IP addresses on page 127

Related tasks

Select system communication roles for Master Engine interfaces on page 603

Sharing interfaces on Master Engines

As an alternative to using an external, physical switch, you can add a single layer 3 physical interface on a Master Engine that can be shared by up to 250 Virtual Engines. In addition, VLAN interfaces under the physical interface can be shared.

An example of where this could be beneficial is that a managed security services provider (MSSP) can have a single layer 3 physical interface that is shared by multiple Virtual Engines, where each Virtual Engine is dedicated to a different customer.

In addition to sharing a regular physical interface, the Virtual Engines can share aggregated link interfaces.

The Virtual Engines are identified by a unique unicast MAC address. The shared physical interface has a MAC address prefix (the first five octets of a MAC address) which groups the Virtual Engines together. The final octet of the MAC address, automatically taken from the Virtual Engine ID, identifies the individual Virtual Engine.

Underneath shared interfaces, you can also add shared VLAN interfaces that can be shared by multiple Virtual Engines.

The Virtual Engines that share an interface can communicate with each other if needed, but you must manually configure the routing and Access rules.

Limitations

Shared interfaces cannot be created when using the Convert Security Engine to Master Engine and Virtual Engines wizard. You must manually add the interfaces later.

Create shared physical interfaces for Master Engines

In the properties of the physical interface, you must add multiple Virtual Resources.

Before you begin

You have created Virtual Resources.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click a Security Engine, then select **Edit <element type>**.
3) Browse to **Interfaces**.
4) Select **Add** > **Layer 3 Physical Interface**.
5) From the **Virtual Resource** drop-down list, select **Multiple Virtual Resources**.
6) Click **Add**, then select the Virtual Resources that you want to use for the Virtual Engines.
7) Configure the additional settings, then click **OK**.

Next steps

Create Virtual Engines that use the Virtual Resources that you have created.

Create shared VLAN interfaces for Master Engines

You can optionally create shared VLANs on the shared Master Engine physical interface that can be used by multiple Virtual Engines.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Security Engine, then select **Edit <element type>**.
2) Browse to **Interfaces**.
3) Select **Add** > **Layer 3 Physical Interface**.

4) From the **Virtual Resource** drop-down list, select **Multiple Virtual Resources**.

Note

Do not add any Virtual Resources to the table.

5) Configure the additional settings, then click **OK**.

6) Right-click the shared physical interface that you created, then select **New** > **VLAN Interface**.

7) In the **Virtual Resources** section, click **Add**, then select the Virtual Resources that you want to use for the Virtual Engines.

8) Configure the additional settings, then click **OK**.

Next steps

Create Virtual Engines that use the Virtual Resources that you have created.

Select the same Virtual Resource for multiple interfaces

You can easily use the same Virtual Resource for multiple physical or VLAN interfaces.

Note

This procedure is not valid for shared interfaces.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Security Engine, then select **Edit <element type>**.

2) Browse to **Interfaces**.

3) Ctrl-select or shift-select multiple physical or VLAN interfaces that do not have an IP address.

4) Right-click the selected interfaces, then select **Select Virtual Resource**.

5) Select the Virtual Resource to associate with the interfaces, then click **OK**.

6) Click **Save and Refresh** to transfer the configuration changes.

Configuring interfaces for Virtual Engines

Physical interfaces in the properties of a Virtual Engine represent interfaces allocated to the Virtual Engine in the Master Engine.

All communication between Virtual Engine and the SMC is proxied by the Master Engine.

Physical interfaces for Virtual Engines are automatically created based on the interface configuration in the Master Engine properties. The number of physical interfaces depends on the number of interfaces allocated to the Virtual Engine in the Master Engine. You can optionally edit the automatically created physical interfaces.

You can add VLAN interfaces if the creation of VLAN interfaces for Virtual Engines is enabled in the properties of the physical interface on the Master Engine.

Both IPv4 and IPv6 addresses are supported on Virtual Engines. You can define one or more static IP addresses for Virtual Engine interfaces. On Virtual Engines, you can also optionally add tunnel interfaces for route-based VPNs.

You can optionally add loopback IP addresses to the Virtual Engine. Loopback IP addresses allow you to assign IP addresses that do not belong to any directly connected networks to the Virtual Engine. Loopback IP addresses are not connected to any physical interface and they do not create connectivity to any network. Any IP address that is not already used on another physical or VLAN interface in the same Virtual Engine can be used as a loopback IP address. The same IP address can be used as a loopback IP address and as the IP address of a tunnel interface. Loopback IP addresses can be used as the IPv4 Identity for Authentication Requests or IPv6 Identity for Authentication Requests, the IPv4 Source for Authentication Requests or IPv6 Source for Authentication Requests, and the Default IP Address for Outgoing Traffic.

By default, the interface definitions for the Virtual Engine are mapped to interfaces on the Master Engine in the order in which the interfaces are created on the Master Engine.

The interface configuration for Virtual Engines consists of the following main steps:

1) Edit the automatically created physical interfaces.
2) (Optional) Add the required number of VLANs.
3) (Optional, Virtual Engines only) Define tunnel interfaces for route-based VPNs.
4) (Virtual Engines only) Configure the IP address settings.
5) (Optional, Virtual Engines only) Define Loopback IP addresses to assign IP addresses that do not belong to any directly connected networks to the virtual engine.
6) (Virtual Engines only) Select the interfaces that are used in particular roles.

Change the properties of physical interfaces for Virtual Engines

Physical interfaces for Virtual Engines are automatically created based on the interface configuration of the Master Engine.

The number of physical interfaces depends on the number of interfaces allocated to the Virtual Engine in the Master Engine. It is not recommended to add physical interfaces to Virtual Engines, as the interfaces might not be valid. You can change the automatically created physical interfaces.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Security Engine, then select **Edit <element type>**.
2) Browse to **Interfaces**.
3) Right-click a physical interface, then select **Edit Physical Interface**.
4) Configure the settings, then click **OK**.
5) Continue the configuration in one of the following ways:
   - If you want to use VLANs, add the VLANs before adding IP addresses.
   - (Virtual Engine only) If you want to create route-based VPNs, add tunnel interfaces.
   - (Virtual Engine only) Otherwise, add IP addresses directly to the physical interfaces.
   - Otherwise, click **Save and Refresh** to transfer the configuration changes.

Related concepts

Adding routes for Master Engines and Virtual Engines on page 695
Defining Zone elements on page 923
Quality of Service (QoS) and how it works on page 973

Related tasks

Add logical interfaces on page 592
Add VLAN interfaces to Virtual Engine interfaces on page 610
Add tunnel interfaces for Virtual Engines on page 611

Add VLAN interfaces to Virtual Engine interfaces

VLANs divide a single physical network link into several virtual links.

VLAN interfaces can only be added for Virtual Engines if the creation of VLAN interfaces is enabled in the Master Engine Properties. The maximum number of VLANs for a single physical interface is 4094. The VLANs must also be defined in the configuration of the external switch or router to which the interface is connected.

Note

You cannot add VLAN interfaces on top of other VLAN interfaces. Depending on the configuration of the Master Engine, you might not be able to create valid VLAN interfaces for the Virtual Engine. Contact the administrator who configured the Master Engine.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Security Engine, then select **Edit <element type>**.
2) Browse to **Interfaces**.
3) Right-click a physical interface, then select **New** > **VLAN Interface**.
4) Configure the settings, then click **OK**.
   The specified VLAN ID is added to the physical interface.
5) Continue the configuration in one of the following ways:
   - (Virtual Engine only) If you do not want to add tunnel interfaces for route-based VPNs, add IP addresses directly to the physical interfaces.
   - Otherwise, click **Save and Refresh** to transfer the configuration changes.

Related concepts

Adding routes for Master Engines and Virtual Engines on page 695
Defining Zone elements on page 923
Quality of Service (QoS) and how it works on page 973

Related tasks

Add logical interfaces on page 592
Add tunnel interfaces for Virtual Engines on page 611

Add tunnel interfaces for Virtual Engines

Tunnel Interfaces allow routing information to be used to determine the correct VPN tunnel to use in route-based VPNs.

Any traffic that is routed to a tunnel interface and allowed by Access rules is automatically sent through the tunnel to the peer endpoint defined in the Route-based Tunnels element. Tunnel interfaces are only used in route-based VPNs.

You can optionally add IPv4 or IPv6 addresses to a tunnel interface. Tunnel interfaces can only have static IP addresses. Any IP address can be added to a tunnel interface, even if the same IP address is used on another interface or as a loopback IP address. Adding an IP address to a tunnel interface allows you to define the source IP address of traffic sent from the Security Engine itself. For example, an IP address is recommended to provide a source IP address for dynamic routing daemons, for IGMP proxy, and for Protocol Independent Multicast - Sparse-Mode (PIM-SM) configuration. If no IP address is added to the tunnel interface, the source IP address

for traffic sent from the Security Engine is automatically selected. The selection is done according to the **Bypass Default IP Address** setting in the loopback interface configuration for the Security Engine.

The mapping of tunnel interfaces to physical network interfaces on the Security Engine is done automatically based on the routing configuration.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click an Security Engine, then select **Edit <element type>**.
3) Browse to **Interfaces**.
4) Select **Add** > **Tunnel Interface**.
5) Configure the settings.
6) Click **OK**.
7) If you want to add a source IP address for traffic sent from the engine node, add IPv4 addresses or IPv6 addresses to the tunnels.
8) If you do not want to add IP addresses, select system communication roles for engine interfaces to define how the source IP address for traffic sent from the engine node is selected.
9) Click **Save and Refresh**.

Related concepts

Adding routes for Master Engines and Virtual Engines on page 695
Defining Zone elements on page 923
Quality of Service (QoS) and how it works on page 973
Types of VPNs in Forcepoint Network Security Platform on page 1157

Related tasks

Add tunnel interfaces for engines on page 558
Change the properties of physical interfaces for Virtual Engines on page 610
Select additional options for Virtual Engine interfaces on page 613

Add IP addresses to Virtual Engine interfaces

You can add one or more static IPv4 or IPv6 addresses for each Physical Interface, VLAN Interface, or Tunnel Interface on a Virtual Engine.

You can optionally add loopback IP addresses to the Virtual Engine. Loopback IP addresses allow you to assign IP addresses that do not belong to any directly connected networks to the Virtual Engine. Loopback IP addresses are not connected to any physical interface and they do not create connectivity to any network. Any IP address

that is not already used on another physical or VLAN interface in the same Virtual Engine can be used as a loopback IP address. The same IP address can be used as a loopback IP address and as the IP address of a tunnel interface. Loopback IP addresses can be used as the IPv4 Identity for Authentication Requests or IPv6 Identity for Authentication Requests, the IPv4 Source for Authentication Requests or IPv6 Source for Authentication Requests, and the Default IP Address for Outgoing Traffic.

You might need to define a contact address if you enter a private static address and NAT is used to translate it to a different external IP address. The external IP address must be configured as the contact address if the IP address is used as a VPN endpoint.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Security Engine, then select **Edit <element type>**.
2) Browse to **Interfaces**.
3) Right-click a Physical Interface, VLAN Interface, or Tunnel Interface, then select **New** > **IPv4 Address** or **New** > **IPv6 Address**.

   Note

   If you have added VLAN Interfaces to Physical Interfaces, add the IP Addresses to the VLAN Interfaces.

4) Configure the settings, then click **OK**.
5) Continue the configuration in one of the following ways:
   - If you are creating a new Virtual Engine, or if you want to change the roles the different interfaces have in the configuration, select interface options for Virtual Engine interfaces.
   - Otherwise, click **Save and Refresh** to transfer the configuration changes.

Related concepts

Define contact IP addresses on page 127

Select additional options for Virtual Engine interfaces

In the Virtual Engine's interface options, you can select which IP addresses are used in particular roles.

Interface Options can only be configured for Virtual Engines. All communication between Virtual Engines and the SMC is proxied by the Master Engine. Virtual Engines do not have any interfaces for system communication.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Security Engine, then select **Edit <element type>**.

2) Browse to **Interfaces** > **Interface Options**.

3) Configure the settings.

Next steps

Continue the configuration in one of the following ways:

- Add loopback IP addresses for the Virtual Engine.
- If you are configuring a new Virtual Engine, click **Save**, close the Engine Editor, then add routes for the Master Engine.
- Otherwise, click **Save and Refresh** to transfer the configuration changes.

Related concepts

Adding routes for Master Engines and Virtual Engines on page 695

Add loopback IP addresses to Virtual Engines

You can use Loopback IP addresses to assign IP addresses that do not belong to any directly connected networks to the Virtual Engine.

Loopback IP addresses are not connected to any physical interface and they do not create connectivity to any network.

- Any IP address that is not used to route traffic on another interface can be used as a loopback IP address, and you can add several loopback IP addresses to each Engine.
- The same IP address can be used as a loopback IP address and as the IP address of a Tunnel Interface.
- Loopback IP addresses can be used as the IPv4 Identity for Authentication Requests or IPv6 Identity for Authentication Requests, the IPv4 Source for Authentication Requests or IPv6 Source for Authentication Requests, and the Default IP Address for Outgoing Traffic.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Security Engine, then select **Edit <element type>**.

2) Browse to **Interfaces** > **Loopback**.

3) Configure the settings.

4) Continue the configuration in one of the following ways:
   - If you are configuring a new Virtual Engine, click **Save**, close the Engine Editor, then add routes for the Master Engine.
   - Otherwise, click **Save and Refresh**.

Related concepts

Adding routes for Master Engines and Virtual Engines on page 695

Add manual ARP entries to Security Engine

You can add manual ARP entries for IPv4 and neighbor discover entries for IPv6 in the Engine Editor.

ARP (Address Resolution Protocol) entries and neighbor discovery entries are normally managed automatically based on the routing configuration. It is not necessary to add manual ARP entries or neighbor discovery entries unless there are problems with the automatic entries, such as devices that do not respond to gratuitous ARP requests, or that impose a significant delay on such operations. The manual ARP entries and neighbor discovery entries are generated by the Security Engine regardless of the installed policy.

Engines support both static and proxy ARP entries defined with IPv4 and IPv6 addresses. IPS engines, Layer 2 Engines, Master Security Engines, and Virtual Engines support only static ARP entries defined with IPv4 addresses.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Security Engine, then select **Edit <element type>**.
2) Browse to **Interfaces** > **ARP Entries**.
3) Click **Add ARP Entry**.
4) Configure the settings.
5) Click **Save and Refresh** to transfer the configuration to the Security Engine.

Examples of interface configurations

Examples of interface configurations for Security Engine in different roles are described in the following sections.

Examples of Single Engine interface configuration

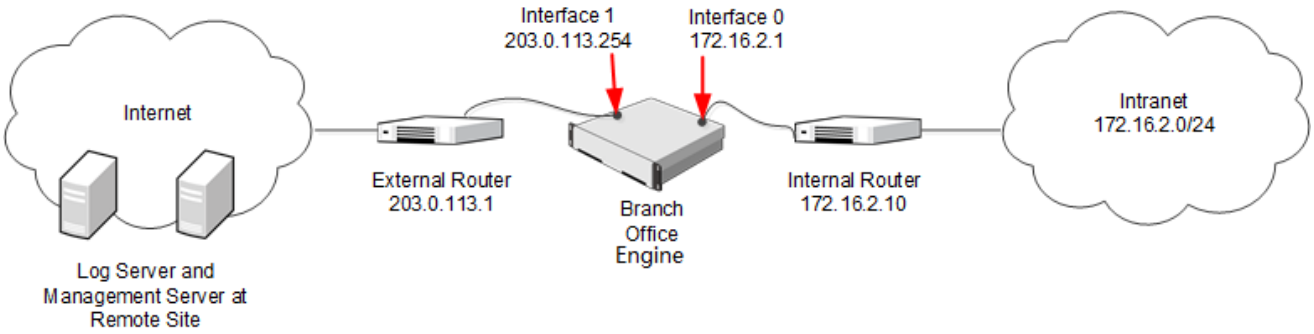
These examples illustrate common configurations for a Single Engine and general overviews of the steps for how each example is configured.

Example: setting up a Single Engine

An example of creating a new Single Engine and configuring the interfaces.

Company A has opened a new branch office. The administrator at the branch office is setting up a Single Engine in the branch office network.

Branch office network



The Branch Office Engine has two interfaces with internal and external routers:

- The internal router is connected to Interface ID 0.
- The external router is connected to Interface ID 1.

The SMC has already been installed at the remote site, and the branch office administrator is now ready to install and configure the Single Engine. The administrator:

1) Creates a Single Engine element (Branch Office Engine) and defines the Log Server at the remote site as its Log Server.

2) Creates an interface for connecting to the internal router and gives it the following properties:
   - Interface ID: 0.
   - IP Address: 172.16.2.1.

3) Creates an interface for connecting to the external router and gives it the following properties:
   - Interface ID: 1.
   - IP Address: 203.0.113.254.

4) Saves the initial configuration of the Branch Office Engine on a USB drive.

5) Installs the engine in the server room.

6) Inserts the USB drive in the engine, turns it on, and waits until the SMC Client shows that contact is established between the engine and the Management Server.

7) Checks the routing configuration and adds the first few rules for allowing traffic through the engine.

8) Installs a Engine Policy using the SMC Client to transfer the first working configuration to the engine.

Example: adding an additional interface to a Single Engine element

An example of adding an interface to a Single Engine.

In the previous example, the administrator initially configured the engine at the company's new branch office with just two interfaces. Now the administrator decides to add a physically separated DMZ network for access to/from that office's mail server to properly control both internal and external traffic with this publicly exposed server. The administrator:

1) Creates an interface for the DMZ and gives it the following properties:
   - Interface ID: 2
   - IP Address: 192.168.2.1.
2) Creates new rules in the engine's policy to allow traffic to/from the DMZ and NAT rules to translate between the private and public IP address of the mail server.
3) Connects the new DMZ router to the engine.
4) Installs a Engine Policy using the SMC Client to transfer the new working configuration to the Engine.

Examples of Engine Cluster interface configuration

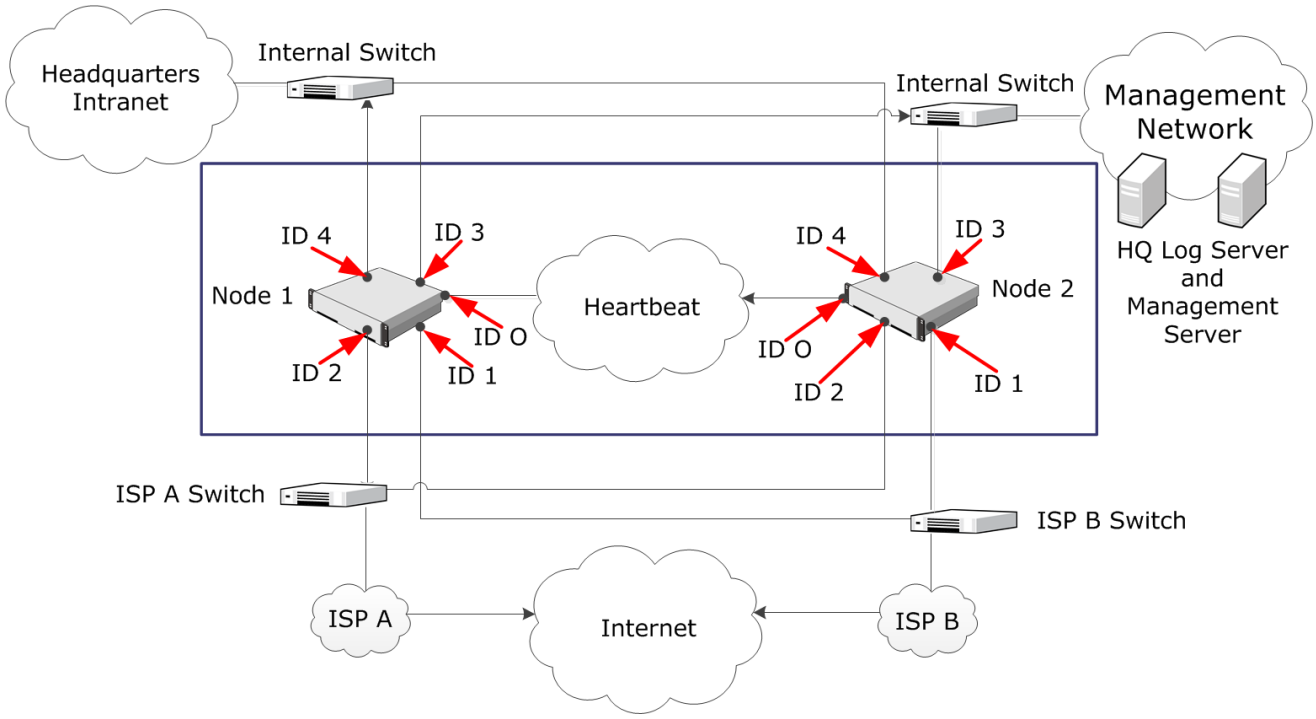
These examples illustrate the configuration of a Engine Cluster with general steps for how each example is configured.

Example: setting up a Engine Cluster element

An example of creating a Engine Cluster element and configuring the interfaces.

The administrators at the headquarters of Company A want to set up a Engine Cluster. The cluster consists of two cluster nodes: Node 1 and Node 2. The HQ Cluster Engine has a dedicated heartbeat network (10.42.1.0/24), and it is connected to two internal networks: Headquarters Intranet (172.16.1.0/24) and Management Network (192.168.10.0/24). It uses Multi-Link to ISP A and ISP B for its connection to the Internet.

Headquarters Network



The administrators:

1) Create a Engine Cluster element (HQ Cluster) and define HQ Log as its Log Server.

2) Define the physical interfaces 0–4.

3) Define the CVIs and NDIs for the physical interfaces. Except for the IP addresses, the node-specific properties for Node 1 and Node 2 are the same.

Cluster Interfaces

Interface ID	Type	IP Address	Comment
0	NDI for Node1	10.42.1.1	Heartbeat
0	NDI for Node2	10.42.1.2	Heartbeat
1	CVI	198.51.100.254	ISP B
1	NDI for Node1	198.51.100.21	ISP B
1	NDI for Node2	198.51.100.22	ISP B
2	CVI	203.0.113.254	ISP A
2	NDI for Node1	203.0.113.21	ISP A
2	NDI for Node2	203.0.113.22	ISP A
3	CVI	192.168.10.1	Management Network
3	NDI for Node1	192.168.10.21	Management Network
3	NDI for Node2	192.168.10.22	Management Network

Interface ID	Type	IP Address	Comment
4	CVI	172.16.1.1	Headquarters Intranet
4	NDI for Node1	172.16.1.21	Headquarters Intranet
4	NDI for Node2	172.16.1.22	Headquarters Intranet

4) Save the initial configuration of the engines in the SMC Client.

5) Map the interface identifiers in the configuration to the physical interfaces on each engine's command line and establish contact between each engine and the Management Server.

6) Install a Engine Policy on the Engine Cluster in the SMC Client to transfer the working configuration to the engines. The nodes exchange authentication information and begin to work as a cluster.

Example: adding a node to a Engine Cluster element

An example of adding a node to an existing Engine Cluster element and configuring the interfaces.

Company A's Engine currently consists of two nodes. However, the load on the Engine is exceptionally high, so the administrator has decided to add another node to ensure continuity of network services even when one of the nodes is offline. The administrator does the following:

1) Adds a third node in the Engine Cluster element's properties.

2) Defines the node-specific IP addresses for the NDI interfaces of the new node.
   - The cluster-level interface configuration does not need adjustments because it is shared by all nodes.

3) Installs the new engine and performs the initial configuration.

4) Refreshes the security policy of the Engine Cluster.

Examples of IPS engine interface configuration

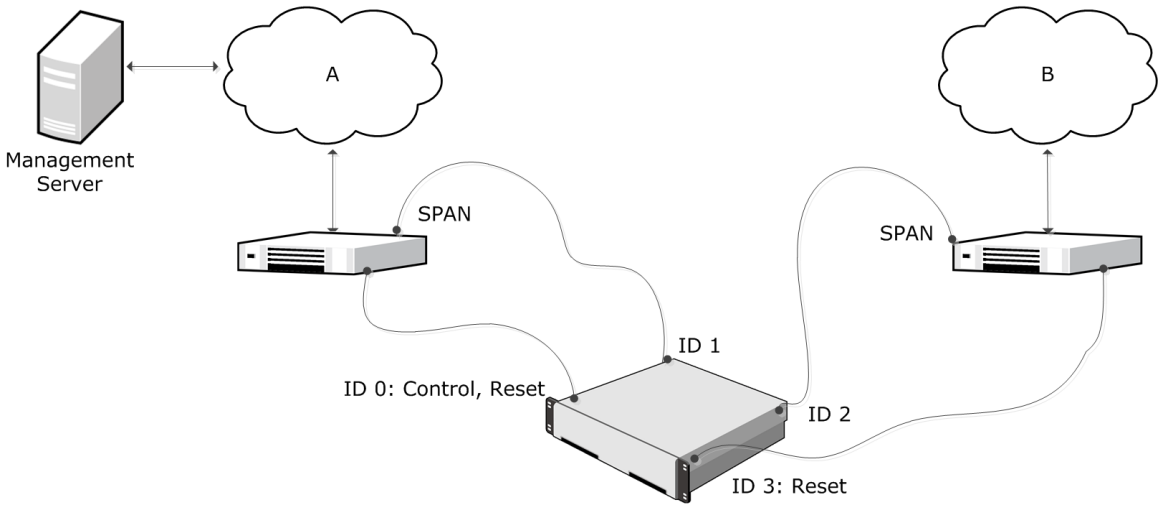
These examples illustrate some common uses for IPS engines and general steps for how each example is configured.

Example: IPS Capture Interface configuration with SPAN

An example of using SPAN ports on switches to duplicate packets for inspection.

The administrator at company A wants to set up a Single IPS engine and deploy it in IDS configuration using SPAN ports on the switches to duplicate packets for inspection. The following illustration shows the interfaces of the IPS engine in IDS configuration.

Capture Interfaces with SPAN



In this example, Interface ID 0 is a Normal Interface used for management connections, and sending TCP Reset responses for network segment A. Interface ID 1 is a Capture Interface for capturing network traffic from the network segment A switch for inspection. Interface ID 2 is a Capture Interface for capturing network traffic from the network segment B switch for inspection. Interface ID 3 is a Normal Interface used for sending TCP Reset responses for network segment B.

The administrator does the following:

1) Creates a Single IPS element and selects the Log Server to which it sends log data and the traffic recordings.

2) Defines Interface ID 0 as a Normal Interface and adds an IP address to it.
   - The IP address on Interface ID 0 is automatically selected as the Primary Control IP address because Interface ID 0 is the first Normal Interface with an IP address.

3) Defines Interface ID 3 as a Normal Interface without an IP address.
   - Because Interface ID 3 is used only as a Reset Interface, it must not have an IP address.

4) Defines Interface ID 1 as a Capture Interface and selects Interface ID 0 as the Reset Interface.

5) Defines Interface ID 2 as a Capture Interface and selects Interface ID 3 as the Reset Interface.

6) Saves the initial configuration of the engine in the SMC Client.

7) Maps the interface IDs to the physical interfaces in the Security Engine Configuration Wizard and makes initial contact with the Management Server.

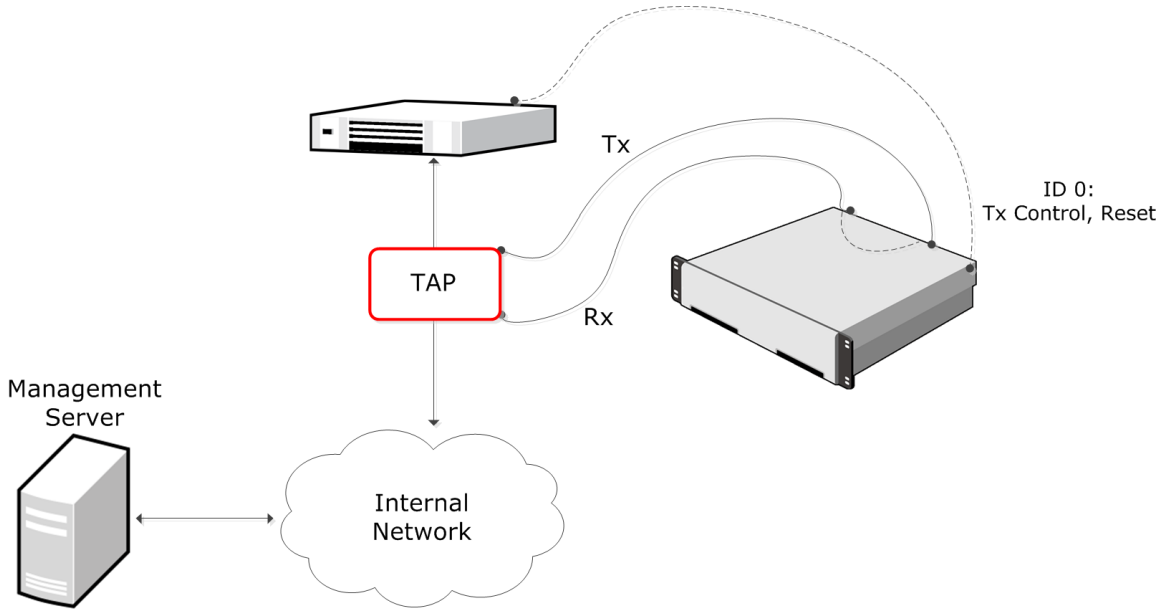
8) Installs an IPS Policy in the SMC Client to transfer the configuration to the engine.

Example: IPS Capture Interface configuration with TAP

An example of using a network TAP device to forward packets for inspection.

The administrator at company B wants to set up a Single IPS engine and deploy it in IDS configuration using a network TAP device. WireTAP copies transmitted (Tx) and received (Rx) packets from the monitored cable and forwards them to separate links for further analysis in the Single IPS engine. The following illustration shows the interfaces of the Single IPS engine in IDS configuration.

Capture Interfaces with TAP



In this example, Interface ID 0 is a Normal Interface used for management connections, and sending TCP Reset responses. Interface ID 1 is a Capture Interface that listens to the received (Rx) packets from the network TAP. Interface ID 2 is a Capture Interface that listens to transmitted (Tx) packets from the network TAP. Interface IDs 1 and 2 share the Logical Interface, which combines the traffic from both physical interfaces so that it can be inspected as a complete traffic flow.

The administrator does the following:

1) Creates a Single IPS element, and selects the Log Server to which it sends log data and the traffic recordings.

2) Creates a Logical Interface called Capture for the two Capture Interfaces.

3) Defines Interface ID 0 as a Normal Interface and adds an IP address to it.

4) Defines Interface ID 1 and Interface ID 2 as Capture Interfaces, selects Interface ID 0 as the Reset Interface, and selects the Logical Interface called Capture for both.

5) Saves the initial configuration of the engine in the SMC Client.

6) Connects the network cables to the appropriate NICs.

7) Maps the interface IDs to the physical interfaces in the Security Engine Configuration Wizard and makes initial contact with the Management Server.

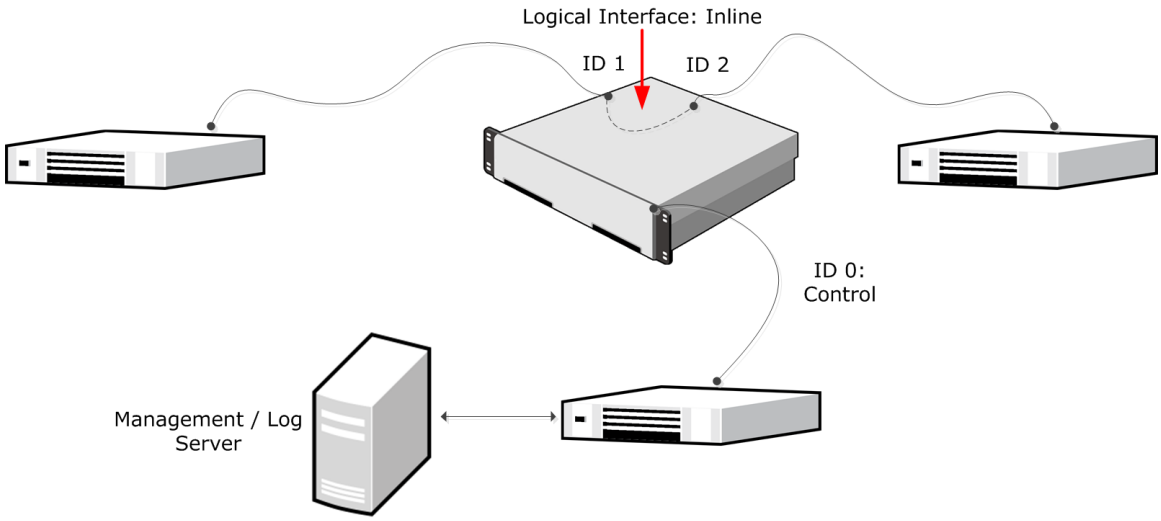
8) Installs an IPS Policy in the SMC Client to transfer the configuration to the engine.

Example: IPS Inline Interface configuration

An example of deploying a Single IPS in the traffic path.

The administrator at Company C wants to set up a Single IPS engine and deploy it in the traffic path. The following illustration shows the interfaces of the inline Single IPS engine.

Inline IPS engine



In this example, the IP address on Interface ID 0 is configured as the Control IP address for management connections. Interface ID 1 and Interface ID 2 are an Inline Interface pair that share the Logical Interface, called Inline. Traffic comes in through Interface ID 1. Any traffic that is allowed by the IPS engine leaves through Interface ID 2.

The administrator does the following:

1) Creates a Single IPS element and selects the Log Server to which it sends log data and the traffic recordings.

2) Creates a Logical Interface called Inline for the Inline Interface pair.

3) Defines Interface ID 0 as a normal interface and adds an IP address to it.

4) Defines Interface IDs 1 and 2 as an Inline Interface pair and selects the Logical Interface called Inline for the pair.

5) Saves the initial configuration of the engine in the SMC Client.

6) Connects the network cables to the appropriate NICs.

7) Maps the interface IDs to the physical interfaces in the Security Engine Configuration Wizard and makes initial contact with the Management Server.

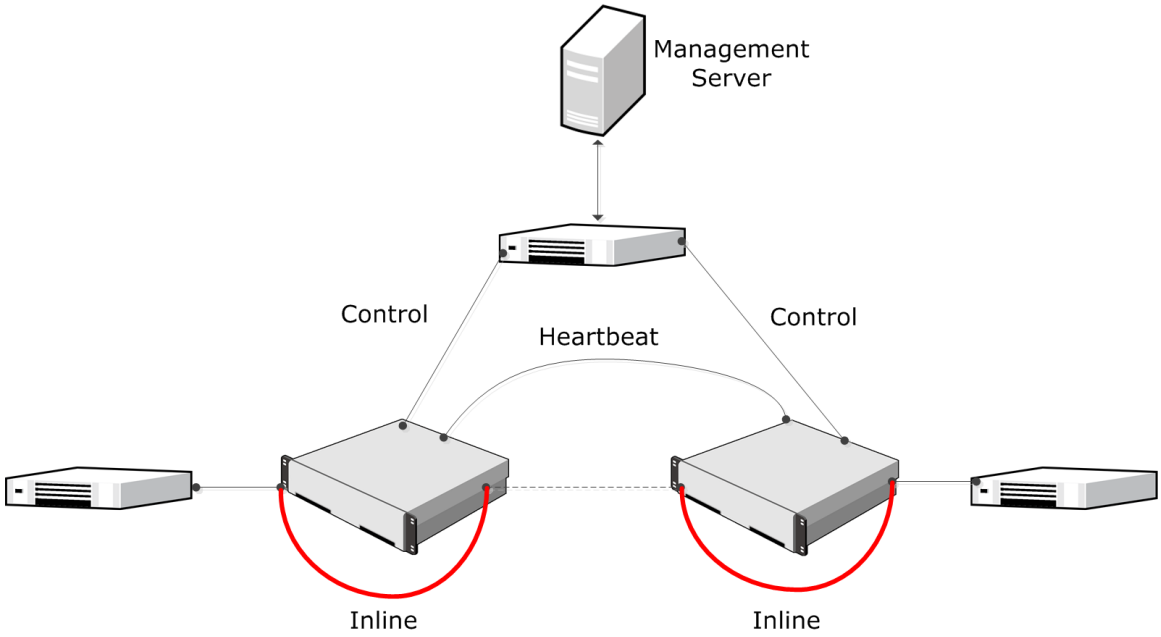
8) Installs an IPS Policy in the SMC Client to transfer the configuration to the engine.

Example: IPS Cluster serial inline deployment

This scenario shows an example of an IPS Cluster in a serial inline deployment.

The following illustration shows the interfaces of the inline IPS Cluster.

Inline serial IPS cluster



In this example, the IPS Cluster consists of two nodes. Interface ID 0 is a Normal Interface used for the heartbeat communication between the nodes. Interface ID 1 is a Normal Interface used for communication with the Management Server. Interface ID 2 and Interface ID 3 are an Inline Interface pair that share one Logical Interface, called Inline. Traffic enters each IPS node through Interface ID 2 and leaves through Interface ID 3.

The administrators:

1) Create an IPS Cluster element and select the Log Server to which the IPS Cluster sends event data and traffic recordings.

2) Define Interface ID 0 as a Normal Interface and add IP addresses for each of the nodes. The IP address on Interface ID 0 is automatically selected as the Primary Control IP address, the Primary Heartbeat Interface, and the Log communication source IP Address.

3) Define Interface ID 1 as a Normal Interface and add IP addresses for each of the nodes.

4) Define Interface IDs 2 and 3 as an Inline Interface pair and select the Logical Interface called Inline for the pair.

5) Select Interface ID 0 as the Primary Heartbeat Interface and select the IP address on Interface ID 1 as the Primary Control IP address in the Interface Options.

6) Save the initial configuration of the engine in the SMC Client.

7) Connect the Heartbeat and Inline Interfaces between the nodes with crossover cables, and the rest of the interfaces with straight cables.

8) Map the interface IDs to the physical interfaces in the Security Engine Configuration Wizard and make initial contact with the Management Server.

9) Install an IPS Policy on each of the nodes in the SMC Client to transfer the configuration to the IPS Cluster.

Examples of Layer 2 Engine interface configuration

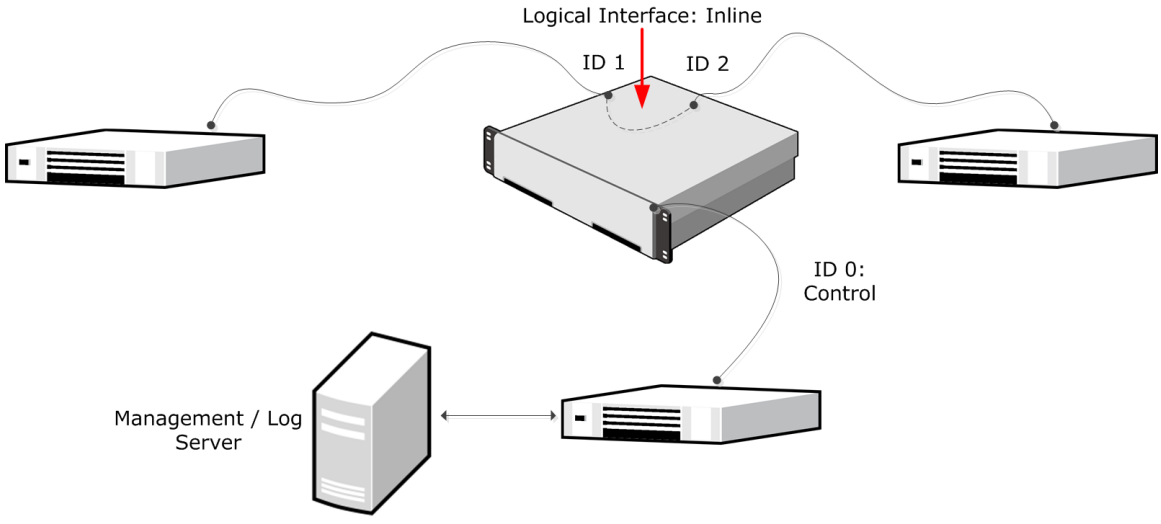
These examples illustrate some common uses for Layer 2 Engines and general steps for how each example is configured.

Example: Layer 2 Engine Inline Interfaces in inline mode

An example of deploying a Layer 2 Engine in the traffic path in inline mode.

The following illustration shows the interfaces of the inline Layer 2 Engine.

Inline Layer 2 Engine



In this example, the IP address on Interface ID 0 is configured as the Control IP address for management connections. Interface ID 1 and Interface ID 2 are an inline interface pair that share the Logical Interface, called Inline. Traffic comes in through Interface ID 1. Any traffic that is the Layer 2 Engine allows leaves through Interface ID 2.

The administrator does the following:

1) Creates a Single Layer 2 Engine element and selects the Log Server to which the Layer 2 Engine sends its log data.

2) Creates a Logical Interface called Inline for the Inline Interface pair.

3) Defines Interface ID 0 as a normal interface and adds an IP address to it.

4) Defines Interface IDs 1 and 2 as an inline interface pair and selects the Logical Interface called Inline for the pair.

5) Saves the initial configuration of the engine in the SMC Client.

6) Connects the network cables to the appropriate physical interfaces on the engine.

7) Maps the interface IDs to the physical interfaces in the Security Engine Configuration Wizard and makes initial contact with the Management Server.

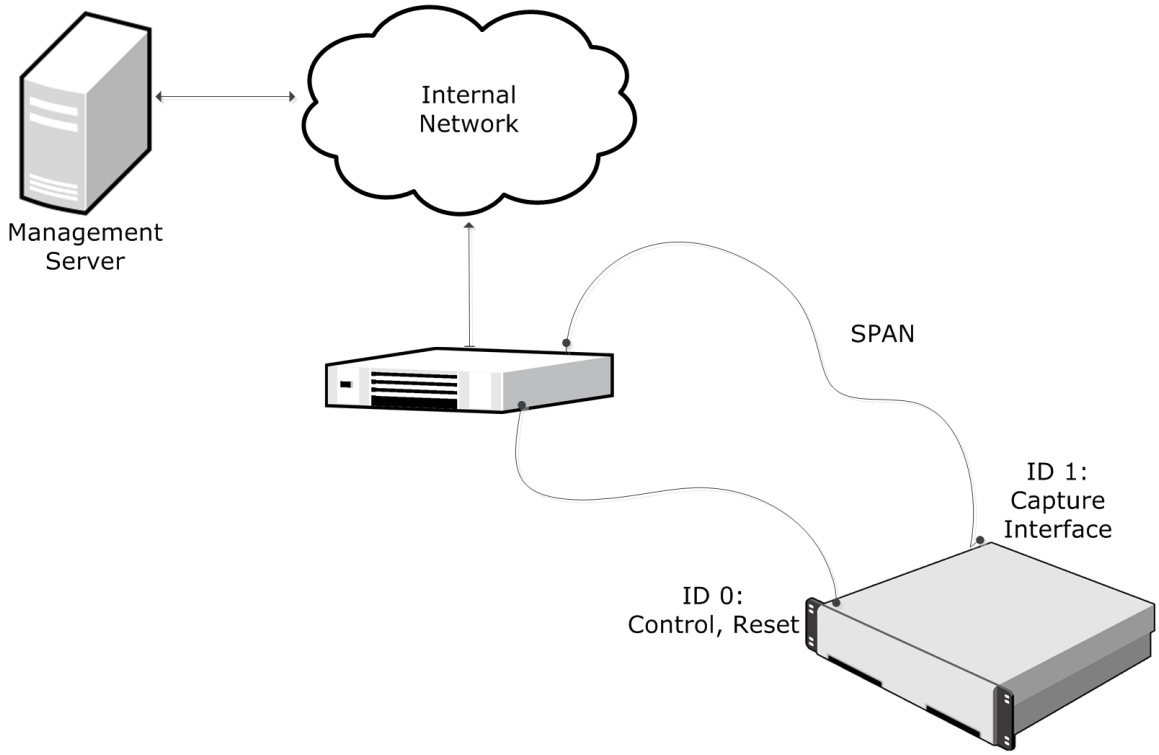
8) Installs a Layer 2 Engine Policy in the SMC Client to transfer the configuration to the engine.

Example: Layer 2 Engine Capture Interfaces in Passive Engine mode

An example of deploying a Layer 2 Engine in Passive Engine mode.

The administrator at company B wants to set up a Single Layer 2 Engine and deploy it in Passive Engine mode using SPAN ports on the switch to duplicate packets for inspection. The following illustration shows the interfaces of the Layer 2 Engine in Passive Engine mode with Capture Interfaces.

Capture Interfaces with SPAN



In this example, Interface ID 0 is a Normal Interface used for management connections and sending TCP Reset responses. Interface ID 1 is a Capture Interface used for capturing network traffic from the network switch for inspection.

The administrator does the following:

1) Creates a Single Layer 2 Engine element and selects the Log Server to which the Layer 2 Engine sends its log data.

2) Defines Interface ID 0 as a Normal Interface and adds an IP address to it.
   - The IP address on Interface ID 0 is automatically selected as the Primary Control IP address because Interface ID 0 is the first Normal Interface with an IP address.

3) Defines Interface ID 1 as a Capture Interface and selects Interface ID 0 as the Reset Interface.

4) Saves the initial configuration of the engine in the SMC Client.

5) Maps the interface IDs to the physical interfaces in the Security Engine Configuration Wizard and makes initial contact with the Management Server.

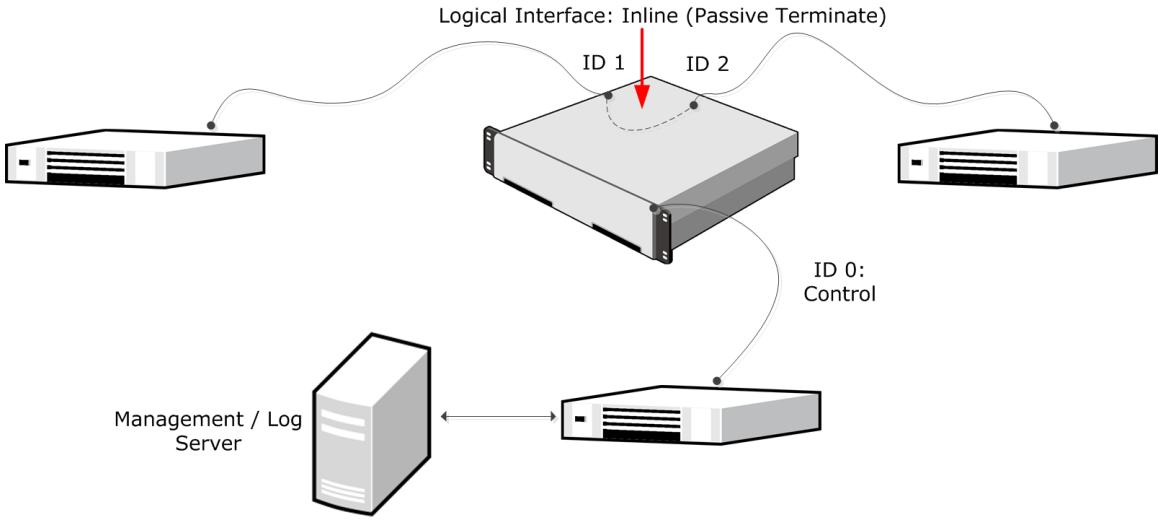
6) Installs a Layer 2 Engine Policy in the SMC Client to transfer the configuration to the engine.

Example: Layer 2 Engine Inline Interfaces in Passive Engine mode

An example of deploying a Layer 2 Engine in Passive Engine mode in the traffic path.

The administrator at company C wants to set up a Single Layer 2 Engine and deploy it in Passive Engine mode in an inline configuration. The following illustration shows the interfaces of the Single Layer 2 Engine in Passive Engine mode with Inline Interfaces.

Inline Interfaces in Passive Engine Mode



In this example, the IP address on Interface ID 0 is configured as the Control IP address for management connections. Interface ID 1 and Interface ID 2 are an inline interface pair that share the Logical Interface, called Inline (Passive Terminate). Traffic comes in through Interface ID 1 and leaves through Interface ID 2.

The administrator does the following:

1) Creates a Single Layer 2 Engine element and selects the Log Server to which the Layer 2 Engine sends its log data.

2) Creates a Logical Interface called Inline (Passive Terminate) for the Inline Interface pair.

3) Defines Interface ID 0 as a Normal Interface and adds an IP address to it.

4) Defines Interface IDs 1 and 2 as an inline interface pair and selects the Logical Interface called Inline for the pair.

5) Configures the Layer 2 Engine to only create Terminate (passive) log entries:
   - For all connections that match the Access rules with the Discard action in the Layer 2 Engine Policy.
   - All Inspection rules with the Terminate action in the Inspection Policy.

6) Saves the initial configuration of the engine in the SMC Client.

7) Connects the network cables to the appropriate physical interfaces on the engine.

8) Maps the interface IDs to the physical interfaces in the Security Engine Configuration Wizard and makes initial contact with the Management Server.

9) Installs a Layer 2 Engine Policy in the SMC Client to transfer the configuration to the engine.

Chapter 32

Connecting Security Engine to the SMC

Contents

- Management connections for Security Engines and how they work on page 629
- Configuration overview of connecting Security Engines to the SMC on page 630
- Connect Security Engines to the SMC on page 631

To maintain the security of your system, the Security Engine establish an authenticated and encrypted connection with Log Servers and Management Servers.

Management connections for Security Engines and how they work

When you connect the Security Engine to the SMC, the Security Engine makes initial contact with the Management Server and receives a certificate.

The certificate allows the Security Engine to authenticate itself to other components in all further communications. When components contact each other, they check if the other component's certificate is signed by the same internal certificate authority as their own certificate. The certificate authority runs on the Management Server, but is separate from the Management Server itself. The initial contact procedure is secured using a one-time password.

If using Forcepoint Network Security Platform appliances, you can connect them to the SMC using the plug-and-play configuration method. In plug-and-play configuration, you upload the initial configuration to the Installation Server. When the appliance is turned on with all cables connected, it downloads the initial configuration from the Installation Server. After this, the Security Engine automatically installs the initial configuration and makes initial contact with the Management Server. You can also specify a policy to be installed on the Security Engine when it makes initial contact with the Management Server.



Note

There are special considerations when using plug-and-play configuration. For example, both the SMC and the Security Engines must be registered for plug-and-play configuration before you configure the engines. See Knowledge Base article 9662.

Saving the initial configuration details on a USB drive allows automatic configuration by turning on the appliance with the USB drive inserted. Alternatively, you can import the configuration details from a USB drive in the Security Engine Configuration Wizard.

You can also save the initial configuration details in some other suitable location or on the clipboard. You can then copy and paste or enter them manually in the Security Engine Configuration Wizard.

CAUTION

The information must be handled securely when saving the initial configuration details on a USB drive or in some other location. The initial configuration files include the one-time password for establishing the trust relationship between the Management Server and the engine.

Limitations

- The plug-and-play configuration method is only available for Forcepoint Network Security Platform appliances. You must have a valid proof-of-serial (POS) code for each appliance you want to configure using the plug-and-play configuration method.
- Virtual Engines do not communicate directly with the SMC. All communication between Virtual Engines and the SMC is proxied by the Master Security Engine.

What should I know before I begin?

- Security Engine certificates expire three years after they are issued. If the automatic certificate renewal option is active, the certificate is renewed automatically before it expires.
- If the certificate of the Security Engine is lost or expires, the initial contact procedure must be repeated to reconnect the Security Engine to the other components.
- The internal certificate authority that signs the Security Engine certificates is valid for ten years. The internal certificate authority is automatically renewed six months before the expiration date and new certificates signed by the new internal certificate authority are automatically created for the Security Engines. If the automatic certificate renewal fails, you must again make initial contact with the Management Server so that the Security Engine receives a new certificate.
- When a new internal certificate authority is created, its initial status is **Ready to Use** and it is not yet **Active**. A new internal certificate authority in a **Ready to Use** state only signs Management Server certificates. Certificates for other SMC components are signed by the internal certificate authority that is used by the Management Server. In an environment with multiple Management Servers, the new internal certificate authority reaches **Active** status when all the Management Servers are using the new internal certificate authority.

Related tasks

View the status of appliance configurations on page 223

Configuration overview of connecting Security Engines to the SMC

Configuring management connections for Security Engines consists of these main steps.

1) Configure the Security Engine in the SMC Client.

2) Save an initial configuration on the Management Server. This triggers the creation of a one-time password that the Security Engine uses to connect to the Management Server.

3) Install the Security Engine to make the Security Engine contact the Management Server.

4) Install a policy on the Security Engine to transfer the full working configuration from the Management Server to the Security Engine.

Related concepts
Connect Security Engines to the SMC on page 631

Related tasks
Reconfigure Security Engine settings on page 365
Install policies on page 813

Connect Security Engines to the SMC

Save the initial configuration to enable the Security Engines to connect to the SMC.

Saving an initial configuration allows you to establish a management connection for Security Engines for the first time. If you are installing a new Security Engine or want to replace a previous working configuration, you can save relevant parts of the configuration on a USB drive and import it during the Security Engine installation.

Saving an initial configuration also allows you to reconnect previously configured Security Engines that have lost the connection. This might be because of a missing or expired certificate or because the internal certificate authority that signs the Security Engine certificates has been renewed and the Security Engines have not yet received a new certificate signed by the new internal certificate authority.

When you save the initial configuration, a one-time password is created. This password is required if you use the Security Engine Configuration Wizard to configure Security Engines.

By default, one-time passwords expire after 30 days if they are not used. You can optionally configure the expiration time in the **Global System Properties** dialog box.

The one-time password that is created is specific to each Security Engine. Keep track of the passwords. If you mix them up or lose them, you can repeat the procedure and create new initial configurations.

If there is a Engine between the Security Engine and the Management Server, allow the connection in the Engine's Access rules. If there is a NAT device between the Security Engine and the Management Server, also configure NAT rules for the connection.

Related concepts
Centralized management of global system settings on page 106

Save the initial configuration and generate the one-time password

Save the initial configuration to establish a management connection for Security Engines.

Saving the initial configuration generates the one-time password required for manual configuration using the Security Engine Configuration Wizard. You can alternatively upload the configuration details to the Installation Server or save them, for example, on a USB drive.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Dashboard** > **Engines Dashboard**.

2) Save the initial configuration information:
   - For an individual Security Engine node, right-click the node, then select **Save Initial Configuration**.
   - For all the Security Engines in a cluster, right-click the top-level cluster element, then select **Configuration** > **Save Initial Configuration**.

3) To manually enter details in the Security Engine Configuration Wizard or if the engine already has the correct configuration, select **View Details**, then write down the one-time password.

4) Configure the settings.

Next steps

- If you selected to upload the initial configuration to the Installation Server, connect the cables, then turn on the appliance. The appliance contacts the Installation Server and downloads the initial configuration.
- To configure the appliance automatically using a USB drive, turn on the appliance with a USB drive inserted.
- If you manually saved the details, turn on the appliance and import the configuration to the Security Engine Configuration Wizard.

For more information, see the *Forcepoint Network Security Platform Installation Guide*.

Installing policies on Security Engines

Install a policy using the SMC Client to transfer the complete configuration from the Management Server to an Security Engine.

After installation, the Security Engine is ready to start processing traffic. The Security Engines do not receive any clustering settings until the first time you install a policy on them and the working configuration is received from the Management Server. If you do not include a policy in the initial configuration, only the interface used for the control connection with the Management Server is configured after the Security Engine establishes contact with the Management Server.

Chapter 33

Element-based network address translation (NAT)

Contents

- Element-based NAT and how it works on page 633
- Add NAT definitions for element-based NAT on page 634
- Edit or remove NAT definitions for element-based NAT on page 635

Element-based NAT allows you to define NAT addresses in the properties of an element. The NAT definitions define how engines translate network IP addresses.

Element-based NAT and how it works

With element-based NAT, you select which elements have their own NAT address and define the NAT addresses for those elements.

Note

Element-based NAT is not intended for complex network environments. In more complex network environments, we recommend adding NAT rules to the Engine policy.

What element-based NAT does

You can add NAT definitions to the following types of elements:

- Security Engines — Single Engines, Engine Clusters, Master Engines, and Virtual Engines
- Servers — Active Directory Servers, DHCP Servers, Elasticsearch Clusters, External DNS Servers, ICAP Servers, LDAP Servers, Log Servers, Management Servers, NTP Servers, Proxy Servers, RADIUS Authentication Servers, SMTP Servers, TACACS+ Authentication Servers, and Web Access Servers
- Some elements in the **Network Elements** branch of the **Configuration** view, for example, Address Ranges, Groups, Hosts, Networks, and Routers.

In addition to using element-based NAT, you can manually add NAT rules to the Engine Policy if you want to configure NAT in more detail. Remember, however, that a more specific manually created NAT rule might prevent traffic from matching NAT rules that are automatically generated from NAT definitions. For more information about manually adding NAT rules, see the topic that explains how engine NAT rules work.

You can also use a default NAT address for all internal networks to automatically translate traffic from internal networks to the public IP address of the external interface. This can be useful, for example, in simple network environments. Default NAT can only be selected in the engine properties.

Note

When several IP addresses from the same network are available, the SMC automatically selects the smallest IPv4 address as the default NAT address.

An internal element that has a static NAT definition can be used in the **Destination** cell of an Access rule. Traffic also matches the destination IP address that corresponds to the element's public IP address in the NAT definition.

What do I need to know before I begin?

- NAT rules are automatically generated and organized in the Engine Policy based on the NAT definitions created in the element properties.
- NAT rules generated from NAT definitions are not visible in the Engine Policy, and are applied after the NAT rules that you have added manually to the policy.
- The SMC automatically generates both the source and destination NAT rules from the NAT definitions.

Related concepts

Getting started with NAT rules on page 851

Add NAT definitions for element-based NAT

NAT definitions define the NAT addresses for elements.

When you add a NAT definition to an engine, the NAT definition is also added to the elements that are included in the engine's NAT configuration. You primarily configure NAT definitions in the Engine Editor. It is also possible to configure NAT definitions in a network element's properties, depending on your permissions in the Domain to which the elements belong.

NAT definitions are automatically processed in the following order from the most specific to the least specific:

- Manually added NAT rules in the Engine Policy
- NAT definitions for element-based NAT
- Default NAT

If there is not a more specific match after the NAT rules and the NAT definitions are processed, default NAT is used. You can use NAT rules in the Engine Policy to create exceptions to NAT definitions and default NAT.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click an engine element and select **Edit <element type>**.

3) In the navigation pane on the left, browse to **Policies** > **Element-based NAT**.

4) Configure the settings, then click **OK**.

5) Click **Save and Refresh**.

Next steps

Refresh the policy on the Security Engine after you have edited the NAT definition of any element to transfer the changes.

Related concepts
Getting started with NAT rules on page 851

Related tasks
Edit or remove NAT definitions for element-based NAT on page 635

Edit or remove NAT definitions for element-based NAT

Edit outdated NAT definitions or remove unnecessary NAT definitions.

Note

You must refresh the Engine policy on the engine after you have edited the NAT definition of any element to transfer the changes.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click an engine element and select **Edit <element type>**.
The Engine Editor opens.

3) In the navigation pane on the left, browse to **Policies** > **Element-based NAT**

4) Select the NAT definition and click **Edit NAT Definition** or **Remove NAT Definition**.
   - If you selected **Edit NAT Definition**, edit the NAT definition settings and click **OK**.
   - If you selected **Remove NAT Definition**, click **Yes** in the confirmation dialog box that opens.

5) Click **Save and Refresh**.

Related tasks
Add NAT definitions for element-based NAT on page 634

Chapter 34

Configuring the Security Engine tester

Contents

- Getting started with the Security Engine tester on page 637
- Specify global tester settings on page 638
- Add Security Engine tests on page 638
- Configuring additional test-specific settings for Security Engine tests on page 639
- View configured Security Engine tests on page 643
- Remove Security Engine tests on page 643
- Disable or enable Security Engine tests on page 644

The Security Engine tester runs various checks on the Security Engine and initiates responses based on the success or failure of these tests.

Getting started with the Security Engine tester

The Security Engines can be configured with periodic self-tests to make sure that each Engine, IPS engine, Layer 2 Engine, and Master Security Engine is functioning correctly.

The tester runs checks at certain intervals, depending on the state of the engine (online, offline, or standby). Depending on the result, the tester can turn the engine online or offline, send alerts, and send SNMP traps.

The tester has the following limitations:

- The tester also runs internal system tests that cannot be edited or disabled. These tests are meant for recognizing certain configuration problems and internal error conditions, such as nodes in the cluster having different policies.
- Tests cannot be run on Virtual Engines.

Engine tester configuration overview

You can define the settings to run tests on Security Engines.

1) Configure the global tester settings that are common to all tests.
2) Add Security Engine tests.

3) Configure the settings for running individual tests.

Specify global tester settings

The global settings of the tester have default values that you can override to meet your needs.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click a Security Engine, then select **Edit <element type>**.
3) Browse to **General** > **Tester**.
4) Configure the global settings.

Add Security Engine tests

As well as test-specific settings, some tests share common settings.

You can receive notification of test failures as Alerts or as SNMP traps. A test can switch also nodes offline or online based on the result.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click a Engine, IPS engine, Layer 2 Engine, or Master Engine element, and select **Edit <element type>**.
3) Browse to **General** > **Tester**.
4) Under the test entry table, click **Add**, then select the test type.
5) Configure the common settings.
6) Click **OK**.
7) Click **Save and Refresh** to transfer the new configuration.

Related concepts
Getting started with SNMP configuration for Security Engines on page 657

Configuring additional test-specific settings for Security Engine tests

The settings available to you vary according to the test you select. Some of the tests have additional, test-specific settings.

Configure settings for the External test

For the tester to run external tests, configure the available settings.

CAUTION

This test allows administrators who have permissions to edit the properties of Security Engines to run arbitrary commands in the Security Engine operating system.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the **Retry Count** field, enter the number of times the tester tries to execute the test
   We recommend always setting the retry count to more than 1 to avoid creating overly sensitive tests that burden the system unnecessarily.

2) In the **Test Timeout** field, enter the timeout in milliseconds.
   - If the test being run does not return a response in the specified time, the test is considered to have failed.
   - We recommend a timeout of 500–1000 milliseconds, depending on the external test script.

3) In the **Command Line** field, enter the command or script path.
   Example:

```
/data/home/root/connectivity.sh
```

Tip

You can use custom properties profiles to upload the script to the Security Engine.

4) Click **OK**.

Note

The command or script must return an exit code of 0 (zero) if it succeeds. Any non-zero return value is a failure.

5) Click **Save and Refresh**.

Related concepts
Using custom properties profiles to upload custom scripts on page 687

Configure settings for the File System Space test

For the tester to run File System Space tests, configure the available settings.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Specify the partition in the **Partition** field.
2) Enter the minimum amount of **Free Space** in kilobytes.
   When the amount of free space drops below this amount, the engine executes the chosen action.
3) Click **OK**.
4) Click **Save and Refresh**.

Configure settings for the Free Swap Space test

For the tester to run Free Swap Space tests, configure the available settings.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Enter the minimum amount of **Free Space** in kilobytes.
   When the amount of free space drops below this amount, the engine executes the chosen action.
2) Click **OK**.
3) Click **Save and Refresh**.

Configure settings for the Inline Pair Link Speed test

For the tester to run Inline Pair Link Speed tests, configure the available settings.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Enter the **Test Timeout** in milliseconds.
   If the test being run does not return a response in the specified time, the test is considered to have failed. Avoid overly short timeout values.
2) Click **OK**.
3) Click **Save and Refresh**.

Configure settings for the Link Status test

For the tester to run Link Status tests, configure the available settings.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) From the **Interface** drop-down list, select the interface on which the test is run.

   Note

   On Engines, only the first interface that belongs to an Aggregated Link is shown in the list of interfaces. However, the Link Status test checks the status of all interfaces in the Aggregated Link.

2) (Aggregated link in load-balancing mode only) Select the percentage of aggregated links that must be down for the test to be considered failed.
3) Click **OK**.
4) Click **Save and Refresh**.

Configure settings for the Multiping test

For the tester to run Multiping tests, configure the available settings.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the **Retry Count** field, enter the number of times the tester tries to execute the test.

   We recommend always setting the retry count to more than 1 to avoid creating overly sensitive tests that burden the system unnecessarily.

2) In the **Test Timeout** field, enter the timeout value in milliseconds.

   If the test being run does not return a response in the specified time, the test is considered to have failed. Avoid timeout values that are too short.

3) In the **Source Address** drop-down list, select the source address for the test:

   - **DEFAULT** — The source IP address is selected automatically based on the routing configuration.

     In a cluster, if the physical interface that routes the ping packet out has an NDI (Node Dedicated IP address), this address is used as the source address. Otherwise, the NDI selected as **Default IP for Outgoing Traffic** is used.
   - A single physical interface, a VLAN interface, a modem interface (Single Engines only), an SSID interface (Single Engines only), or a port group interface (Single Engines only). If the **Node ID** selection is **ALL**, each node uses the IP address of the selected interface as the source IP address for the test. If a single node is selected in **Node ID** in a cluster, the source address and the multiping test itself apply to that node only.

4) Specify the target addresses of ICMP echo requests.

   Note

   The test is considered to have failed if none of the target addresses respond.

   a) To add a row to the **Target Addresses** table, click **Add**.

   b) Enter an IPv4 or IPv6 address.

5) Click **OK**.

6) Click **Save and Refresh**.

Configure settings for the Policy test

The Policy test is no longer supported in Security Engine version 6.1.0 and higher.

View configured Security Engine tests

The test entries that you have configured are displayed in the **Tester** pane of the Engine Editor.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click a Engine, IPS engine, Layer 2 Engine, or Master Engine element and select **Edit <element type>**.
   The Engine Editor opens.

3) In the navigation pane on the left, browse to **General** > **Tester**.
   The **Tester** pane opens on the right. All configured test entries are displayed in a table. The selected states, actions, and parameters for the tests are shown.

Remove Security Engine tests

Removing a test permanently removes the test settings from the selected Security Engine.

Alternatively, you can also stop a test from running without removing it

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click a Engine, IPS engine, Layer 2 Engine, or Master Engine element and select **Edit <element type>**.
   The Engine Editor opens.

3) In the navigation pane on the left, browse to **General** > **Tester**.
   The **Tester** pane opens on the right.

4) Select the test entry that you want to remove.

5) Click **Remove**.
   A **Confirmation** dialog box opens.

6) Click **Yes** to confirm that you want to remove the test.
   The test is permanently removed.

7) Click **Save and Refresh** to transfer the new configuration to the engines.

Disable or enable Security Engine tests

You can disable or enable any individual test or all tests that are run on a specified node.

You do not have to enable test entries that you configure to activate them. All test entries that you configure are automatically activated and available to the tester after a policy refresh or install.

Disable or enable individual tests

Temporarily disable individual unused engine tests, or enable the individual tests you want to run.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click a Engine, IPS engine, Layer 2 Engine, or Master Engine element and select **Edit <element type>**.
   The Engine Editor opens.
3) In the navigation pane on the left, browse to **General** > **Tester**.
   The **Tester** pane opens on the right.
4) Deselect or select the option in the **Active** column of the test.
5) Click **Save and Refresh** to transfer the new configuration to the engines.

Disable or enable all custom tests

Temporarily disable all custom tests on a node, or enable all custom tests you want to run on a node.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click a Engine, IPS engine, Layer 2 Engine, or Master Engine element and select **Edit <element type>**.
   The Engine Editor opens.
3) In the navigation pane on the left, browse to **General** > **Tester**.
   The **Tester** pane opens on the right.
4) In the **Global Settings** section, deselect or select the option in the **Active** column of one or more nodes.

5) Click **Save and Refresh** to transfer the new configuration to the engines.

Chapter 35

Engine permissions

Contents

- Getting started with permissions on page 647
- Define administrator permissions for Security Engines on page 648
- Select the allowed policies on page 649

You can set permissions to control the administration of the engines.

Getting started with permissions

You can define the permissions that enable or restrict administrators to edit and view an engine's properties.

What permission control does

You can user engine permissions control in two ways:

- To prevent some administrators from editing and viewing an engine’s properties to prevent unauthorized modifications and protect confidential details.
- To restrict the policies that can be installed on the engine to prevent service outages caused by the wrong policy being installed on the engine by accident.



Note

Engine permission control does not affect the local permissions to execute command-line commands on the engine. You can replicate administrator accounts on engines and then configure permissions to execute commands using the sudo tool. The Local Administrators section in the Administrator Permissions pane shows the local administrators defined, if any.

What do I need to know before I begin?

- Your administrator account must have editing permissions for the engine element.
- Permissions for Master Engines and Virtual Engines are configured separately. Otherwise, engine permissions are configured for Master Engines and Virtual Engines in the same way as for other types of engines.

Related concepts

Getting started with administrator accounts on page 373

Engine permissions configuration overview

You can define the permissions that permit users to edit and view engine elements, and allow policies to install on the engine.

Configuring engine permission involves the following general steps:

1) Define which administrators are allowed to edit and view the element.

2) Define which policies can be installed on the engine.

Define administrator permissions for Security Engines

Define the administrator permissions that permit users to access and view engine options.

Before you begin

Your administrator account must have editing permissions to the engine element.

You can either add an Access Control List or an individual Administrator-Administrator Role pair as permitted on the engine. The rights that the Access Control List grants to the administrators are defined in the properties of the administrator accounts (defined with Administrator elements).

Administrators with restricted permissions can refresh or upload an engine's policy only if the administrator is a permitted administrator for both the engine and the policy. The engines might not accept all policies.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Open the element for editing in one of the following ways depending on the element type:

Element type	Description
Engine element	1) Right-click an engine element, then select **Edit <element type>**. The Engine Editor opens. 2) In the navigation pane on the left, browse to **General** > **Permissions**.
Policy element	1) Right-click the policy and select **Properties**. 2) Click the **Permissions** tab.

2) Click **Add** for the **Access Control Lists** and select one or more Access Control Lists to which you want the engine to belong.
   To create an Access Control List, select ⋮ **More actions** > **New** > **Access Control List**.

3) To add a permission, click **Add Permission** under **Permissions**.
   A new row appears on the administrator list.

4) Click the **Administrator** cell and select the Administrator.

5) Right-click the **Administrator Role** cell and select **Edit Administrator Role**.

6) Select a role and click **Add**.

7) Click **OK** to close the **Select Elements** dialog box.

8) Save the changes in one of the following ways depending on the element type:

Element type	Description
Engine element	Click **Save**.
Policy element	Click **OK** to close the dialog box.

Note

Changes to administrator permissions are immediately distributed and taken into account in all related elements.

Related concepts
Creating Administrator Role and Access Control List elements on page 377

Related tasks
Create Access Control List elements on page 379

Select the allowed policies

Assign Policy or Template Policy permissions for engines.

Before you begin

Your administrator account must have editing permissions to the engine element.

By default, any policy can be installed on any engine as long as the policy is of appropriate type for the type of engine. To prevent accidental installations of the wrong policy, you can select the allowed policy for each engine in the engine element's properties. The policy selection is enforced regardless of the administrator permissions that the installing administrator has.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**

2) Right-click an engine element and select **Edit <element type>**.
   The Engine Editor opens.

3) In the navigation pane on the left, browse to **General** > **Permissions**.
   The **Permissions** pane opens on the right.

4) In the **Policies** section at the bottom, select the **Policy**:
   - To allow the installation of any policy, click **Set to Any**.
   - Otherwise, click **Add**.

5) Select the correct Policy or Template Policy.
   If you select a Template Policy, any policy based on the template can be installed.

6) Click **Save and Refresh**.

Chapter 36

DNS Relay

Contents

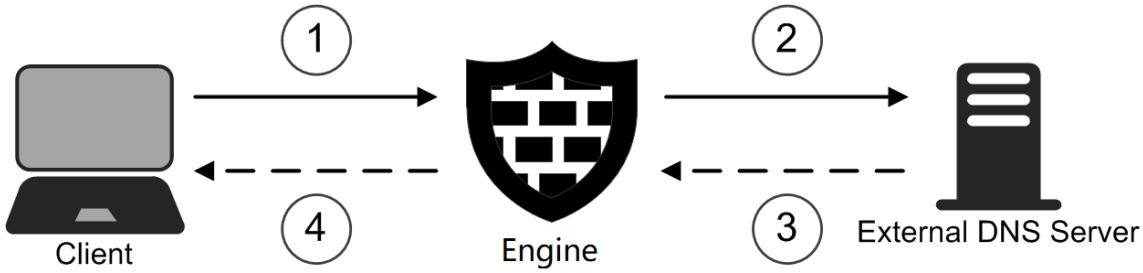
- Getting started with DNS relay on page 651
- Enable DNS relay on page 654

DNS relay allows the engine to provide DNS services for clients in internal networks.

Getting started with DNS relay

In DNS relay, clients send DNS requests to a DNS resolver, which forwards the requests to a remote DNS server. In Forcepoint Network Security Platform, the engine can act as a local DNS resolver for clients in the internal network.

How DNS relay works



1. Clients in the internal network send DNS requests to the engine.
2. The engine forwards the DNS requests to remote DNS servers.
3. Remote DNS servers send DNS responses to the engine.
4. The engine provides the responses to the clients in the internal network.

The engine temporarily stores the results of DNS requests in its cache until the time limit specified in the time to live (TTL) value for the DNS entry is reached. When a client makes a DNS request for a domain that has recently been requested, the engine provides the IP address from the cache. Caching reduces the load on upstream DNS servers and improves performance.

In addition to providing DNS services for clients in the internal network, the engine can also optionally do the following:

- Return fixed DNS results for specific hosts or domains.
- Forward DNS requests to different DNS servers depending on the domain in the DNS request.
- Translate IPv4 addresses resolved by external DNS servers to IPv4 addresses in the internal network.

You can configure DNS relay on Single Engines, Engine Clusters, and Virtual Engines.

DNS relay configuration overview

To use DNS relay, you must configure settings for the engine in the SMC Client, and configure clients in your network environment.

Follow these general steps to set up DNS relay:

1) (Optional) Create a DNS Relay Profile element.
2) Define DNS IP addresses for the engine in the Engine Editor.
3) Configure DNS Relay options for the engine in the Engine Editor.
4) In your internal network environment, configure devices to use the selected listening IP addresses on the engine as a DNS resolver.

Fixed DNS results

You can optionally configure the engine to return fixed DNS results for specific hosts or domains without relaying the request to any DNS server.

You can define fixed DNS results in two ways.

Ways to define fixed DNS results

Option for fixed DNS results	Description
Host name mappings	You statically map host names and aliases for host names to IPv4 or IPv6 addresses. When a client requests DNS resolution for a host name that is included in the fixed mappings, the engine resolves the IP address based on the mappings. Host name mappings simplify the configuration when you only need to resolve a small number of host names in internal networks to static IP addresses. You define host name mappings as pairs of IP addresses and host names in the **Host Name Mappings** section of DNS Relay Profile elements.
Fixed domain answers	The engine replies to requests for specific domain names with IPv4 addresses, IPv6 addresses, domain names, or empty DNS replies. When the engine provides an empty DNS reply, the client receives the same response as for domains for which no DNS record is found. Fixed domain answers are useful if you always want to direct requests for specific domains to specific destinations. For example, you can reply to all requests for the domain of an advertising network with an empty reply to block unwanted ads on web pages. You define fixed domain answers as pairs of domains names and values that the engine returns in the **Fixed Domain Answers** section of DNS Relay Profile elements.

Domain-specific DNS servers

The engine can use domain-specific DNS servers to forward DNS requests to different DNS servers depending on the requested domain.

For example, the engine can forward queries for internal domains to remote internal DNS servers, and forward other queries to a public DNS server, such as a DNS server maintained by your Internet Service Provider (ISP).

When you forward queries for external domains to public DNS servers, users get the DNS result that is geographically closest to them. Using the geographically closest IP address improves the quality of services that use DNS load balancing, such as cloud services that have regionally distributed data centers.

You define domain-specific DNS servers in DNS Relay Profile elements as pairs of domain names and DNS IP addresses in the **Domain-Specific DNS Servers** section of DNS Relay Profile elements.

In the Engine Editor, you specify the IP addresses that are used as source IP addresses when the engine makes domain-specific DNS queries. If you send DNS queries through a VPN tunnel, you must select source IP addresses. In other configurations, selecting source IP addresses is optional. If you do not select a source IP address, the source IP address is automatically selected based on the route to the external DNS server.

An automatic rule allows traffic from the engine to domain-specific DNS Servers. If you send the DNS queries through a VPN tunnel, you must disable the automatic rule.

Related concepts
Rules for DNS relay on page 653

DNS answer translations

DNS answer translations guarantee that DNS replies for users in internal networks always contain internal IP addresses even if a public DNS server returns an external IP address for the same domain.

DNS answer translations map IPv4 addresses resolved by external DNS servers to IPv4 addresses in the internal network. When an IPv4 address in a DNS result matches one of the specified DNS answer translations, the engine modifies the DNS result so that users in internal networks receive the internal IP address.

You define DNS answer translations in DNS Relay Profile elements as pairs of original IPv4 addresses and translated IPv4 addresses in the **DNS Answer Translations** section of DNS Relay Profile elements.

Rules for DNS relay

Traffic for DNS Relay is allowed by automatic rules by default. We recommend using automatic rules.

The automatic rules allow the following traffic for DNS relay:

Automatic rules for DNS relay

Automatic rule	Traffic allowed
Allow Traffic from Listening IP Addresses to DNS Relay Port	Traffic from the listening IP addresses of the engine to port 53/TCP and port 53/UDP for DNS relay.

Automatic rule	Traffic allowed
Allow Connections to Domain-Specific DNS Servers	Traffic from the engine to domain-specific DNS Servers. If you want to send the DNS traffic through a policy-based VPN, you must disable this automatic rule. If you disable this automatic rule, you must add IPv4 or IPv6 Access rules to allow traffic from the engine to the DNS servers. You must also add IPv4 or IPv6 NAT rules if you want to apply NAT or port translation to the DNS traffic.

If you create Access or NAT rules to match specific DNS traffic, use one or more of the following elements:

- **DNS** Service Group — Matches both TCP and UDP traffic on port 53
- **DNS (TCP)** Service — Matches TCP traffic on port 53
- **DNS (UDP)** — Matches UDP traffic on port 53

Enable DNS relay

To enable DNS relay, you must configure DNS Relay settings for the engine. You can optionally create custom DNS Relay Profile elements.

Create DNS Relay Profile elements

DNS Relay Profile elements contain the host name mappings, domain-specific DNS servers, fixed domain answers, and DNS answer translations that the engine uses when it provides DNS services to the internal network.

If you do not want to define custom settings, you can use the predefined **Cache Only** DNS Relay Profile element.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Browse to **Other Elements** > **Engine Properties** > **DNS Relay Profiles**.
3) Right-click **DNS Relay Profiles**, then select **New DNS Relay Profile**.
4) In the **Name** field, enter a unique name.
5) Configure options in one or more of the following sections:
   - **Host Name Mappings** — Statically map host names, aliases for host names, and unqualified names (a host name without the domain suffix) to IPv4 or IPv6 addresses.
   - **Domain-Specific DNS Servers** — Forward DNS requests to different DNS servers depending on the requested domain.
   - **Fixed Domain Answers** — Direct requests for specific domains to IPv4 addresses, IPv6 addresses, fully qualified domain names (FQDNs), or empty DNS replies.

- **DNS Answer Translations** — Map IPv4 addresses resolved by external DNS servers to IPv4 addresses in the internal network.

Note

You can add a maximum of 250 rows to the DNS Relay Profile element.

6) Click **OK**.

Configure DNS Relay settings in the Engine Editor

To enable DNS relay, configure DNS Relay settings for the engine in the Engine Editor.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click a Single Engine, Engine Cluster, or Virtual Engine element and select **Edit <element type>**.

3) In the **DNS IP Addresses** field, add the IP addresses of one or more external DNS servers to which the engine forwards DNS requests from clients in the internal network.
   - To enter a single IP address manually, click **Add** and select **IP Address**. Enter the IP address in the dialog box that opens.
   - To define an IP address using a network element, click **Add** and select **Network Element**.

4) In the navigation pane on the left, browse to **General** > **DNS Relay**.

5) From the **DNS Relay Profile** drop-down list, select a DNS Relay Profile element.
   - If you created a custom DNS Relay Profile element, select your custom element.
   - To enable DNS relay without defining custom settings, select the predefined **Cache Only** DNS Relay Profile element.

6) In the **Listening IP Addresses** field, add one or more IP addresses.
   Clients in the internal network can use all of the specified IP addresses to send DNS requests to the engine.

7) (Optional) From the **Source for Domain-Specific DNS Queries** drop-down list, select the interface that is used as the source IP address when the engine forwards DNS requests to domain-specific DNS servers.
   When **According to Routing** is selected, the source IP address is automatically selected based on the route to the external DNS server.

Next steps

Configure devices in the internal network to use the selected listening IP addresses on the engine as a DNS resolver.

Chapter 37

Setting up SNMP for Security Engines

Contents

- Getting started with SNMP configuration for Security Engines on page 657
- Create an SNMP Agent for SNMP version 1 or 2c on page 659
- Create an SNMP Agent for SNMP version 3 on page 660
- Configure what triggers SNMP traps on page 661
- Activate the SNMP agent on Security Engines on page 662
- Configure Engine access rules to allow SNMP queries from trusted SNMP managers on page 663

SNMP is a standard protocol that different equipment can use to send network management-related information to each other. You can configure Security Engines to send SNMP traps to external equipment.

Getting started with SNMP configuration for Security Engines

Security Engines can send SNMP traps to a central network monitoring system when system events occur, such as when a test fails.

SNMP Agent elements define the settings according to which the Security Engines send SNMP trap messages to compatible external software. The same SNMP Agent can be used by multiple Security Engines and by the SMC Appliance.

The SNMP Agent supports SNMPv1 (RFC1157), SNMPv2c (RFCs 1901 and 3416), and SNMPv3 (RFC 3414). See the documentation of the receiving software for information about which version to use.

The MIB files are included in the Mibs folder in the SMC installation package for all platforms (smc_<version>_<build>.zip). For descriptions of the MIBs, see the MIB files.

Configuring SNMP for Security Engines consists of these general steps:

1) Create an SNMP Agent element.

2) Configure the SNMP Agent element for SNMPv1, SNMPv2c, or SNMPv3.

3) Define what triggers SNMP traps.

4) Activate the SNMP Agent on the Security Engines.

5) Configure Engine access rules to allow SNMP queries from trusted SNMP managers.

SNMP traps and MIBs

You can use these SNMP traps and MIB objects with Security Engines.

SNMP traps for Security Engines

Trap name	Objects included	Description
fwPolicyInstall	fwSecurityPolicy	Policy was installed on the Engine.
ipsPolicyInstall	ipsSecurityPolicy	Policy was installed on the IPS engine.
nodeBoot	-	Node bootup complete.
nodeHwmon	nodeHwmonEvent	Hardware monitoring system has detected problems.
nodeOffline	nodeOperState	Node changed to offline or standby state.
nodeOnline	nodeOperState	Node changed to online state.
nodeShutdown	-	Node is shutting down.
nodeTestFailure	nodeTestIdentity	Test subsystem reported a test failure on the node.
nodeFailedUserLogin	-	Logon failed on the engine's console or through SSH.
nodeUserLogin	nodeLastLogin	Log on initiated on the engine's console or through SSH.
nodeUserLogout	-	Log off on the engine's console or through SSH.

Forcepoint Network Security Platform-specific MIBs

MIB	Role	Engine Version	Description
STONESOFT-SMI-MIB	-	6.10 or earlier	High-level registrations MIB. This MIB defines the top-level enterprise registrations for the Forcepoint Network Security Platform products in the .iso.org.dod.internet.private.enterprises.stonesoft branch (OID `.1.3.6.1.4.1.1369`). The other MIB modules depend on STONESOFT-SMI-MIB, so it is necessary to load STONESOFT-SMI-MIB first.
STONESOFT-NETNODE-MIB	FW, IPS, L2FW	6.10 or earlier	Generic network node MIB.
STONESOFT-FIREWALL-MIB	FW, L2FW	6.10 or earlier	Engine application MIB.
STONESOFT-IPS-MIB	IPS	6.10 or earlier	IPS application MIB.
FORCEPOINT-NGFW-ENGINE-MIB	FW, IPS, L2FW	6.11 or later	Engine application MIB.

Note

Traps and other Forcepoint NGFW or Forepoint Network Security Platform specific SNMP OIDs are defined in FORCEPOINT-NGFW-ENGINE-MIB file. For more details, refer to the *Enhanced SNMP Agent default in NGFW Engine version 7.0* Knowledge Base article.

Create an SNMP Agent for SNMP version 1 or 2c

Configure an SNMP Agent for SNMP version 1 or 2c so that Security Engines can share network management information using the SNMP protocol, or for SNMP version 2c so that the SMC Appliance can share network management information using the SNMP protocol.



Note

The SMC Appliance does not support SNMP v1.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Browse to **Other Elements** > **Engine Properties** > **SNMP Agents**.
3) Right-click **SNMP Agents**, then select **New SNMP Agent**.
4) In the **Name** field, enter a unique for the SNMP Agent.
5) From the **Version** drop-down list, select **v1** or **v2c**.
   For the SMC Appliance, you must select **v2c**.
6) (Optional) In the **Monitoring** section, click **Add**, then enter the community string.
   The community string is used for authentication in monitoring.
7) (Optional) In the **Listening Port** field, enter the UDP port number that the SNMP agent listens to.
8) In the **Contact** field, enter the contact information for the person responsible for the Security Engines or the SMC Appliance.
9) Click **OK**.

Create an SNMP Agent for SNMP version 3

Configure an SNMP Agent for SNMP version 3 so that Security Engines or the SMC Appliance can share network management information using the SNMP protocol.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Browse to **Other Elements** > **Engine Properties** > **SNMP Agents**.
3) Right-click **SNMP Agents**, then select **New SNMP Agent**.
4) In the **Name** field, enter a unique for the SNMP Agent.
5) From the **Version** drop-down list, select **v3**.
6) In the **User Names** section, add one or more users names.
   a) Click **Add**.
   b) In the **User Name** field, enter the user name.
   c) From the **Protocol** options, select the authentication protocol, then enter a password in the **Password** field.
   d) From the **Privacy** options, select the privacy protocol, then enter a password in the **Privacy Password** field.
7) (Optional) In the **Monitoring** section, click **Add**, then select the user for monitoring.
8) (Optional) In the **Listening Port** field, enter the UDP port number that the SNMP agent listens to.
9) In the **Contact** field, enter the contact information for the person responsible for the Security Engines or the SMC Appliance.
10) Click **OK**.

Configure what triggers SNMP traps

The trap parameters define where and how SNMP traps are sent from Security Engines and the SMC Appliance.

The same SNMP Agent element can be used for Security Engines and the SMC Appliance. Some settings only apply to Security Engines. Settings that are not supported for the SMC Appliance are ignored when the SNMP Agent is used for the SMC Appliance.

In addition to the general events, the tester on each Security Engine can send SNMP traps when a test fails.

Note

If the **Destinations** field is left empty, no traps are sent, and the other trap parameters are ignored. If the **Destinations** field has a value, the rest of the trap parameters must also have a value.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Browse to **Other Elements** > **Engine Properties** > **SNMP Agents**.

3) Open the SNMP Agent properties in one of the following ways:
   - Right-click an existing SNMP Agent element, then select **Properties**.
   - To create an SNMP Agent element, right-click **SNMP Agents**, then select **New SNMP Agent**.

4) In the **Traps** section, specify the sender of the SNMP trap.
   - SNMPv1 — (Security Engines only) In the **Community** field, enter a community string.
   - SNMPv2c — In the **Community** field, enter a community string.
   - SNMPv3 — From the **User Name** drop-down list, select a user name.

5) Click **Add**, then enter the IP address and UDP port where the traps are sent.

6) (Security Engines only) In the **Active Traps** section, select the events for which you want to set a trap.

7) Click **OK**.

Activate the SNMP agent on Security Engines

The SNMP Agent is responsible for SNMP-related tasks on the Security Engines.

Before you begin

If you use SNMPv3, there must be one or more user names defined in the properties of the SNMP Agent element.

When you use SNMPv3, you can specify the SNMP engine ID for each single Security Engine and each node of Security Engine clusters. The SNMP engine ID is a unique identifier for the Security Engine that is used by the SNMP agent. The engine ID is used with a hash function to generate keys for authentication and encryption of SNMPv3 messages. If you do not specify the SNMP engine ID, an SNMP engine ID is automatically generated.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**

2) Right-click an engine element, then select **Edit <element type>**.

3) In the navigation pane on the left, browse to **General** > **SNMP and LLDP**.

4) From the **SNMP Agent** drop-down list, select the SNMP Agent that you want to activate.

5) In the **SNMP Location** field, enter the string that is returned on queries to the SNMPv2-MIB or SNMPv2-MIB-sysLocation object.

6) In the **Listening IP Addresses** field, add one or more IPv4 or IPv6 addresses.

7) (Optional, SNMPv3 only) Specify the value of the **SNMP Engine ID** option.
   - Single Security Engines — In the **SNMP Engine ID** field, enter a unique identifier for the Security Engine.
   - Security Engine clusters — Browse to **General** > **Clustering**, then enter a unique identifier for each node in the **SNMP Engine ID** cell.

8) Click **Save and Refresh** to transfer the changes.

Configure Engine access rules to allow SNMP queries from trusted SNMP managers

Automatic rules do not allow SNMP traffic to and from the engine.

The administrator must do the following:

1) Create access rule that allow SNMP queries from trusted SNMP manager hosts.

2) Create access rule that allow engines to send SNMP traps to SNMP trap receivers.

Note

It is recommended to allow SNMP probing of the engine only from trusted sources.

For more details on how to create access rules, refer to the **Access rules** section.

For example, access rule to allow SNMP queries from trusted SNMP managers:

Source	Destination	Service	Action	Comment
SNMP-manager-host	$$ Local Cluster (NDI addresses only)	SNMP (UDP)	Allow	Rule to allow SNMP queries
$$ Local Cluster (NDI addresses only	SNMP-manager-host	SNMP Trap (UDP)	Allow	Rule to allow SNMP queries

Note

The SNMP-manager-host is the network element for the SNMP manager used in the environment.

Chapter 38

Setting up LLDP for Security Engines

Contents

- Getting started with LLDP for Security Engines on page 665
- LLDP configuration overview on page 666
- Create custom LLDP Profile elements on page 666
- Enable LLDP on Security Engines on page 667

Network devices can use the Link Layer Discovery Protocol (LLDP) to advertise their identity, capabilities, and neighbors on a local area network.

Getting started with LLDP for Security Engines

Security Engines can use LLDP to send information about themselves to directly connected devices on the network, and receive information that other devices on the network send.

LLDP makes it easier to deploy a large number of Security Engines. LLDP announcements from Security Engines allow other directly connected devices on the network to assign the correct VLAN IDs to ports on network switches to which the Security Engine is connected. LLDP announcements from directly connected devices on the network provide information about switch topology to Security Engines, such as which network switch and port the Security Engine is connected to, and which VLANs it can reach.

When LLDP is enabled for a Layer 3 Physical Interface on an Security Engine, the Security Engine always announces the following type-length-values (TLVs):

- Chassis ID — The MAC address of the first Ethernet port
- Port ID — The name of the interface in the format 'ifname <name>'
- Port Description — The name of the interface
- Time to Live — The period of time for which LLDP advertisements should be stored in the cache of neighboring LLDP-compliant devices. This value is automatically calculated based on the transmit delay and the hold time multiplier defined in the LLDP Profile element that the Security Engine uses.

The Security Engine can optionally announce the following TLVs:

- System Name — The name of the Security Engine or the node in the Security Engine cluster in the SMC.
- System Description — Operating system details about the Security Engine, such as operating system name, operating system version, and architecture.
- System Capabilities — A bit-map of the enabled capabilities of the interface as router, repeater, and other.
- Management Address — The IP addresses of the control interfaces

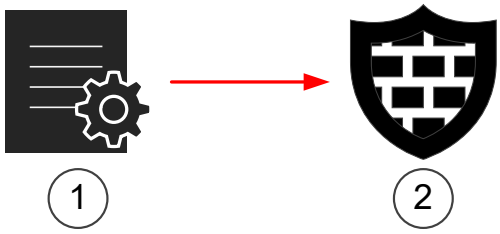
LLDP for Security Engines has the following limitations:

- LLDP is not supported on Virtual Engines.
- LLDP is supported only on Layer 3 Physical Interfaces.

LLDP configuration overview

LLDP configuration consists of several general steps.

Elements in the configuration



1. LLDP Profile elements define settings for LLDP announcements and the type-length-values (TLVs) that the Security Engine announces.

   You can create custom LLDP Profile elements or use the default LLDP Profile element.
2. In the properties of each Security Engine, you select the LLDP Profile that the Security Engine uses, and enable LLDP for individual Layer 3 Physical interfaces.

Create custom LLDP Profile elements

LLDP Profile elements define settings for LLDP announcements and the type-length-values (TLVs) that the Security Engine announces.

There is one default LLDP Profile element. If the default element meets your needs, it is not necessary to create a custom LLDP Profile element.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Browse to **Other Elements** > **Engine Properties** > **LLDP Profiles**.
3) Right-click **LLDP Profiles**, then select **New** > **LLDP Profile**.
4) Configure the settings, then click **OK**.

Enable LLDP on Security Engines

To enable LLDP, select an LLDP Profile for the Security Engine, then enable LLDP for individual Layer 3 Physical interfaces.

Before you begin

If you do not want to use the default LLDP Profile element, you must create a custom LLDP Profile element.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Security Engine, then select **Edit <element type>**.

2) Select an LLDP Profile for the Security Engine.

Note

The same LLDP Profile is used for all interfaces on which you enable LLDP.

   a) Browse to **General** > **SNMP and LLDP**.

   b) Next to the **LLDP Profile** field, click **Select**, then select an LLDP Profile element.

3) Enable LLDP for individual Layer 3 Physical interfaces.

   a) Browse to **Interfaces**.

   b) Right-click a Layer 3 Physical Interface, then select **Edit Physical Interface**.

   c) From the **LLDP Mode** drop-down list, select one of the options.

   d) Click **OK**.

4) Click **Save and Refresh** to transfer the changes.

Chapter 39

Alias element translations for Security Engines

Contents

- Getting started with Alias element translations on page 669
- Add Alias element translation values on page 670
- Remove Alias element translation values on page 671

Alias elements can be used to represent other network elements in configurations. The value an Alias takes in a configuration can be different on each Security Engine where the Alias is used.

Getting started with Alias element translations

Alias element translation values are set for each engine. You can change the translation values through the engine properties or in the properties of the Alias element itself.

What Alias elements do

Use Alias elements in your policies to represent IP addresses. Aliases differ from Group elements in that they do not represent all elements at once. The Alias elements do not contain any IP address information themselves. The values that the Aliases receive in a policy depend on the translation value you set for each Alias in the engine elements' properties. This way, the same policy can be used on several engines. The IP address information is filled in correctly according to the translation values for each engine. Alias elements are especially useful in policy templates and sub-policies.

What do I need to know before I begin?

Aliases are configured for Security Engines. The Security Engines use the translated values in the data they send to the Log Server.

If you use Master Engines and Virtual Security Engines, you must configure Aliases separately for the Master Engines and for the Virtual Security Engines.

There are some predefined Aliases in the SMC. Reserved Aliases cannot be edited. The names of reserved Aliases start with two `$$` symbols. Reserved Aliases receive their translation values automatically based on the engine's configuration. You can also create Aliases. User-created Aliases and predefined Aliases that you can edit start with one `$` symbol. Predefined Aliases that you can edit do not receive their translation values automatically.

Related concepts
Creating Security Engine elements on page 504
Editing existing Security Engines on page 509

Add Alias element translation values

To use Alias elements to represent other network elements in a policy, define the translation values of the Alias elements for each engine that uses the policy. The same Alias element can have different translation values on different engines.

Add Alias element translation values through the engine properties or in the properties of the Alias element itself.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Click **Security Engines**.

3) Right-click an engine element, then select **Edit <element type>**.

4) In the navigation pane on the left, browse to **Policies** > **Aliases**.

5) Select the Alias element whose value you want to edit.

6) Define the value for the Alias:
   - If you want the Alias to match any IP address on this engine, right-click the **Value** cell for the Alias and select **Set to Any**. Skip to Step 8.
   - Otherwise, right-click the **Value** cell and select **Edit Value**.

7) Browse in the **Resources** pane for an element you want to use as the translation value on this engine.
   If the element does not exist, you can create one through the right-click menu for the correct element type.

8) Select the elements and click **Add**.
   The elements are added to the **Alias Value** pane.

9) Click **OK** to close the **Alias Value Properties** dialog box.

10) Click **Save and Refresh**.

Remove Alias element translation values

When you no longer need Alias element translation values, remove them from the engine element properties.

If there are no translation values for a particular Alias element for an engine, the Translation Value is **None**. If Aliases with the value **None** are used alone as matching criteria in a rule in a policy, the rule never matches any traffic.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**
2) Right-click an engine element and select **Edit <element type>**.
3) In the navigation pane on the left, browse to **Policies** > **Aliases**.
4) Right-click the **Value** cell of the Alias and select **Edit Value**.
5) Remove the elements.
   - To remove selected elements, select the elements in the **Alias Value** pane and click **Remove**.
   - To remove all elements, click **Remove All**.
6) Click **OK** to close the **Alias Value Properties** dialog box.
7) Click **Save and Refresh**.

Chapter 40

Add-on features for Security Engines

Contents

- Getting started with add-on features on page 673
- Edit add-on settings for Security Engines on page 674

There are several add-on features that you can use on Engines, IPS engines, Layer 2 Engines, Virtual Engines, Virtual IPS engines, and Virtual Layer 2 Engines.

Getting started with add-on features

In the **Add-Ons** pane in the Engine Editor, you can enable several optional add-ons for engines, such as TLS inspection, QUIC inspection, browser-based user authentication, anti-malware, file reputation checks, and ZTNA connector.

Some add-ons are separately licensed features.

Supported roles for add-on features

Add-on	Supported on
TLS Inspection	Engines, IPS engines, Layer 2 Engines, Virtual Engines, Virtual IPS engines, Virtual Layer 2 Engines
ECA Settings	Engines, IPS engines, Layer 2 Engines, Virtual Engines
User Authentication	Engines, Virtual Engines
User Identification	Engines, IPS engines, Layer 2 Engines
Anti-Malware	Engines, IPS engines, Layer 2 Engines, Master Engines
File Reputation	Engines, IPS engines, Layer 2 Engines, Master Engines
IPv6 Transition	Single Engines, Engine Clusters, Virtual Engines
Sidewinder Proxy	Engines
Snort	Engines, IPS engines, Layer 2 Engines
ThreatSeeker	Engines, IPS engines, Layer 2 Engines, Virtual Engines, Virtual IPS engines, Virtual Layer 2 Engines
OPC UA Inspection	Engines, IPS engines, Layer 2 Engines

Related concepts
TLS inspection and how it works on page 1063
Enabling access control by user on page 1113
Getting started with anti-malware scanning on page 989
Creating Security Engine elements on page 504
Editing existing Security Engines on page 509

Related tasks
Enable browser-based user authentication on page 1137

Edit add-on settings for Security Engines

You can edit add-on settings in the Engine Editor.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click an engine element, then select **Edit** <element type>.
2) In the navigation pane on the left, expand the **Add-Ons** branch.
3) Browse to the add-on settings that you want to edit, then adjust the settings.
4) Click **Save and Refresh** to save the changes to the configuration and refresh the policy on the engine.

Related concepts
TLS inspection and how it works on page 1063
Enabling access control by user on page 1113
Getting started with anti-malware scanning on page 989

Chapter 41

Advanced Security Engine settings

Contents

- Getting started with advanced Security Engine settings on page 675
- Open the advanced settings on page 676
- Adjusting Engine clustering options on page 676
- Adjust IPS clustering options on page 679
- Adjust Layer 2 Engine clustering options on page 680
- Adjust Master Engine clustering options on page 681
- Configure inspection of tunneled traffic on page 683
- Set connection timeouts on page 684
- Configure SYN rate limits on page 684
- Configure log handling settings on page 685
- Configure DoS protection settings on page 686
- Configure scan detection settings on page 687
- Using custom properties profiles to upload custom scripts on page 687

Advanced settings cover various system parameters related to different features.

Getting started with advanced Security Engine settings

Advanced settings system parameters generally have default values that are appropriate for most installations without any need for adjustments. Some values can be overridden in policies rule-by-rule when exceptions are needed.



CAUTION

Improper adjustments to some of the advanced settings can seriously degrade the performance of the system. Do not adjust the advanced settings unless you have a specific need to do so.

Open the advanced settings

To adjust advanced settings for an Security Engine, you must open the Engine Editor.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**
2) Right-click an Security Engine, then select **Edit <element type>**.
3) Browse to **Advanced Settings**.
4) Adjust the settings.
5) Click **Save**.

Adjusting Engine clustering options

You can check and adjust some of the settings related to how the Engine Cluster behaves and exchanges information from node to node.

If you are an advanced user, it is also possible to tune settings related to how the traffic load is balanced between several online nodes. However, we strongly discourage you from changing the load-balancing settings unless Forcepoint Customer Hub instructs you to make particular changes.

Node state synchronization

The nodes of a Engine Cluster periodically exchange synchronization messages to synchronize state data.

State synchronization is essential for the following features:

- Dynamic load balancing
- Transparent switchover of nodes in case of failure or maintenance
- Handling of related connections when a service (for example, FTP) opens multiple connections

Regular, timer-launched synchronization events are needed to synchronize state data and to avoid cutting connections in case of node failure. Timed synchronization events are divided into full and incremental sync messages.

Sync messages

Type	Explanation
Full Sync Messages	Contain all connection data about the traffic handled by a node at the time when the message was sent. When new data is received, it replaces the existing data. Full sync requires more bandwidth and processing time.

Type	Explanation
Incremental Sync Messages	Contain only data on connections that were created or changed since the last full or incremental sync message. Incremental sync needs less bandwidth and processing time. Because the incremental changes are sent only once, the system might lose connections if the data is lost. While able to produce accurate data with frequent updates, incremental sync requires full sync to provide reliable synchronization data.

By default, a combination of full and incremental sync messages is exchanged between nodes. This way, frequent updates on incremental changes and recurrent reports on existing connections are combined.

In cases where synchronization of connection information between nodes is causing a disturbance to specific traffic, you can disable synchronization for the traffic using rule options in the Policy. Disabling synchronization reduces the traffic volume on the active heartbeat interface, but it also prevents transparent failover of connections to other nodes.

Security level for state synchronization

Because synchronization controls the inter-node traffic within a heartbeat network, you must ensure the security of the heartbeat and synchronization data.

Traffic between nodes can be authenticated, or both authenticated and encrypted. Traffic between nodes can also optionally be sent without authentication or encryption. However, this makes it possible to both sniff synchronization data and send fraudulent messages to open connections.

Note

Independent of the security level, all critical information such as passwords and encryption keys are protected. They are never sent in plain text.

Adjust general Engine clustering options

To tune the settings that relate to how the traffic load is balanced between several online nodes, adjust the Engine clustering options.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**

2) Right-click a Engine Cluster element and select **Edit Engine Cluster**.
The Engine Editor opens.

3) In the navigation pane on the left, browse to **General** > **Clustering**.

4) Configure the settings.

Note

Do not adjust the settings in the **Advanced Cluster Settings** dialog box unless you are certain it is necessary.

5) Click **Save and Refresh** to transfer the changes.

Manually tune the Engine load-balancing filter

The Engine Cluster's load-balancing filter can be manually edited if there is a specific need for modifications.

CAUTION

Do not manually tune the load-balancing filter unless you are certain it is necessary. Normally, there is no need to tune the load-balancing filter, because the configuration generates all required entries automatically. Unnecessary tuning can adversely affect the operation of the filter.

Any edited load-balancing parameters are combined with the automatically created filtering entries. However, editing the load-balancing parameters of the Engine Cluster without careful consideration can cause conflicts in filtering decisions.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**

2) Right-click a Engine Cluster element and select **Edit Engine Cluster**.
The Engine Editor opens.

3) In the navigation pane on the left, browse to **General** > **Clustering**.

4) In the **Clustering Mode** section, click **Clustering**.

5) On the **Manual LB Filters** tab of the **Advanced Cluster Settings** dialog box , select an option from the **Filter Mode** drop-down list to define how traffic is balanced between the nodes.

6) (Optional) Select **Load-Balancing Filter Uses Ports** to include a port value for selecting between all nodes.
This setting decreases the granularity of VPN load balancing, and increases the granularity of other traffic load balancing. In typical networks, traffic is balanced based on IP address information only. If there is a dominating pair of communication IP addresses, apply the **Use Ports** option in the load-balancing filter entry only to their traffic.

CAUTION

Enabling the **Load-Balancing Filter Uses Ports** option is not compatible with some features, such as mobile VPNs.

7) Click **OK**.

8) Click **Save and Refresh** to transfer the changes.

Add Engine load-balancing filter entries

You can manually add IP addresses to the Engine Cluster's load-balancing filter.

CAUTION

Do not manually tune the load-balancing filter unless you are certain it is necessary. Normally, there is no need to tune the load-balancing filter, because the configuration generates all required entries automatically. Unnecessary tuning can adversely affect the operation of the filter.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**
2) Right-click a Engine Cluster element and select **Edit Engine Cluster**.
3) In the navigation pane on the left, browse to **General** > **Clustering**.
4) In the **Clustering Mode** section, click **Clustering**.
5) Click the **Manual LB Filters** tab.
6) Click **Add** to generate a new filter entry row.
7) Double-click the **IP Address** field.
8) Select whether you want to filter an **IPv4 Network**, an **IPv6 Network**, or a **Range** of IP addresses.
9) Enter the IPv4 address and netmask, the IPv6 address and prefix, or the address range, and click **OK**.
10) Click the **Action** cell and select an action.
11) If you selected **Replace** by as the action, click the **Replacement IP** field and enter the replacement IP address.
12) (Optional) Select any additional options.
13) Click **OK**.
14) Click **Save and Refresh** to transfer the changes.

Adjust IPS clustering options

IPS Clusters operate by default in load-balancing mode. This means that all configured nodes in an IPS Cluster are online simultaneously and the traffic is distributed among the operational nodes. The load balancing aims to keep the traffic load as evenly distributed as possible.

Alternatively, the IPS Cluster can run in *standby* mode. In that case, only one IPS node at a time is online and processing traffic, while the others are in standby mode. Only if the online node fails, one of the standby nodes goes online to take over the connections being handled by the failed node.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**
2) Right-click an IPS Cluster element, then select **Edit IPS Cluster**.
   The Engine Editor opens.
3) In the navigation pane on the left, browse to **General** > **Clustering**.
   The **Clustering** pane opens on the right.
4) Configure the settings.
5) Click **Save and Refresh** to transfer the changes.

Adjust Layer 2 Engine clustering options

By default, Layer 2 Engine Clusters operate in active-standby mode.

Only one Layer 2 Engine node at a time is online and processing traffic, while the others are standby. Only if the online node fails, one of the standby nodes goes online to take over the connections being handled by the failed node.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**
2) Right-click a Layer 2 Engine Cluster element and select **Edit Layer 2 Engine Cluster**.
   The Engine Editor opens.
3) In the navigation pane on the left, browse to **General** > **Clustering**.
   The **Clustering** pane opens on the right.
4) Configure the settings.
5) Click **Save and Refresh** to transfer the changes.

Adjust Master Engine clustering options

You can check and adjust some of the settings related to how the Master Engine behaves and exchanges information from node to node.

If you are an advanced user, it is also possible to tune settings related to how the traffic load is balanced between several online nodes. However, we strongly discourage you from changing the load-balancing settings unless Forcepoint Customer Hub instructs you to make particular changes.

Adjust general Master Engine clustering options

To tune the settings that relate to how the traffic load is balanced between several online nodes, use the Master Engine Clustering options.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**

2) Right-click a Master Engine element and select **Edit Master Engine**.
The Engine Editor opens.

3) In the navigation pane on the left, browse to **General** > **Clustering**.

4) Configure the settings.

Note

Do not adjust the settings in the **Advanced Cluster Settings** dialog box unless you are certain it is necessary.

5) Click **Save and Refresh** to transfer the changes.

Manually tune the Master Engine load-balancing filter

The Master Engine's load-balancing filter can be manually edited if there is a specific need for modifications.

CAUTION

Do not manually tune the load-balancing filter unless you are certain it is necessary. Normally, there is no need to tune the load-balancing filter, because the configuration generates all required entries automatically. Unnecessary tuning can adversely affect the operation of the filter.

Any edited load-balancing parameters are combined with the automatically created filtering entries. However, editing the load-balancing parameters of the Master Engine without careful consideration can cause conflicts in filtering decisions.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click a Master Engine and select **Edit Master Engine**.
   The Engine Editor opens.
3) In the navigation pane on the left, browse to **General** > **Clustering**.
4) In the Clustering Mode section, click **Clustering**.
5) On the **Manual LB Filters** tab, select an option from the **Filter Mode** drop-down list to define how traffic is balanced between the nodes.
6) (Optional) Select **Load-Balancing Filter Uses Ports** to include a port value for selecting between all nodes.
   This setting decreases the granularity of VPN load balancing, and increases the granularity of other traffic load balancing. In typical networks, traffic is balanced based on IP address information only. If there is a dominating pair of communication IP addresses, apply the **Use Ports** option in the load-balancing filter entry only to their traffic.

   CAUTION

   Enabling the **Load-Balancing Filter Uses Ports** option is not compatible with some features, such as mobile VPNs.
7) Click **OK**.
8) Click **Save and Refresh** to transfer the changes.

Add Master Engine load-balancing filter entries

You can manually add IP addresses to the Master Engine's load-balancing filter.

CAUTION

Do not manually tune the load-balancing filter unless you are certain it is necessary. Normally, there is no need to tune the load-balancing filter, because the configuration generates all required entries automatically. Unnecessary tuning can adversely affect the operation of the filter.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click a Master Engine element and select **Edit Master Engine**.
The Engine Editor opens.

3) In the navigation pane on the left, browse to **General** > **Clustering**.
The **Clustering** pane opens on the right.

4) In the **Clustering Mode** section, click **Clustering**.

5) Click the **Manual LB Filters** tab.

6) Click **Add** to generate a new filter entry row.

7) Double-click the **IP Address** field.

8) Select whether you want to filter an **IPv4 Network**, an **IPv6 Network**, or a **Range** of IP addresses.

9) Enter the IPv4 address and netmask, the IPv6 address and prefix, or the address range, and click **OK**.

10) Click the **Action** cell and select an action.

11) If you selected **Replace** by as the action, click the **Replacement IP** field and enter the replacement IP address.

12) (Optional) Select any additional options.

13) Click **OK**.

14) Click **Save and Refresh** to transfer the changes.

Configure inspection of tunneled traffic

You can define in the Advanced Settings how the IPS engine or Layer 2 Engine inspects tunneled traffic.

If traffic is tunneled using IP-in-IP or Generic Routing Encapsulation (GRE), the payload of the tunneling packet can be checked against the Access rules several times.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click an IPS or Layer 2 Engine element, then select **Edit <element type>**.

3) In the navigation pane on the left, browse to **Advanced Settings** > **Tunneling**.

4) Define the settings.

5) Click **Save and Refresh** to transfer the configuration changes.

Set connection timeouts

You can define general timeouts for removing idle connections from the state table, including non-TCP communications that are handled like connections.

The timeout prevents wasting engine resources on storing information about abandoned connections. Timeouts are a normal way to clear traffic information with protocols that have no closing mechanism. The communicating client and server also have timeouts for closing inactive connections.

You can set timeouts by protocol and by TCP connection state. Idle timeouts set in Access rules override these global settings.

Timeouts do not affect active connections. The connections are kept in the state table as long as the interval of packets within a connection is shorter than the timeouts set.

CAUTION

Setting excessive timeouts for many connections consumes resources excessively and can disturb the operation of the engine.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click a Engine, IPS, or Layer 2 Engine element, then select **Edit <element type>**.

3) In the navigation pane on the left, browse to **Advanced Settings** > **Idle Timeouts**.

4) Click the **Timeout(s)** column and enter the timeout value for the protocol in seconds.

5) (Optional) Click **Add** to add a protocol to the list and enter the timeout for the protocol.

6) Click **Save and Refresh** to transfer the configuration changes.

Configure SYN rate limits

You can configure SYN rate limits to reduce the risk of SYN flood attacks against the Engine, IPS engine, Layer 2 Engine, Master Engine, or Virtual Engine.

SYN rate limits are applied to TCP connections. Each TCP connection starts with a SYN packet. If the SYN rate limits defined for the Security Engine are reached, the Security Engine drops new TCP connections.

The global SYN rate limits that you define in the Security Engine properties are applied as default settings on all interfaces. You can also define SYN rate limits that override the global settings in each interface's properties.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click a Security Engine, then select **Edit <element type>**.
3) Browse to **Advanced Settings** > **SYN Rate Limits**.
4) Configure the settings.
5) Click **Save and Refresh** to transfer the configuration changes.

Configure log handling settings

Log Handling settings allow you to adjust logging when the log spool on the Engine, IPS, Layer 2 Engine, or Master Engine fills up.

Logs are spooled locally when the Log Server is not available. The Master Engine spools its own logs and the logs sent by the Virtual Engines that the Master Engine hosts.

You can also configure Log Compression to save resources on the engine. By default, each generated Antispoofing and Discard log entry is logged separately and displayed as a separate entry in the **Logs** view. Log Compression allows you to define the maximum number of separately logged entries. When the defined limit is reached, a single Antispoofing log entry or Discard log entry is logged. The single entry contains information on the total number of the generated Antispoofing log entries or Discard log entries. After this, logging returns to normal and all generated entries are once more logged and displayed separately.

The general Log Compression settings you define in the Engine Editor are applied as default settings on all interfaces. You can also define Log Compression and override the global settings in each interface's properties.

You can optionally save copies of the most recent log entries locally on the Security Engine. You can browse the saved log entries on the command line of the Security Engine even if the log entries have already been sent to the Log Server.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click a Engine, IPS, Layer 2 Engine, or Master Engine element and select **Edit <element type>**.
3) In the navigation pane on the left, browse to **Advanced Settings** > **Log Handling**.
4) Configure the options according to your environment.

   Do not enable Log Compression if you want all Antispoofing and Discard entries to be logged as separate log entries (for example, for reporting or statistics).
5) Click **Save and Refresh** to transfer the configuration changes.

Configure DoS protection settings

You can configure three forms of protection that can help prevent Denial of Service (DoS) attacks: SYN flood protection and slow HTTP request protection against rate-based DoS attacks, and TCP reset protection against TCP reset attacks.

The following settings can be configured when the Rate-Based DoS Protection Mode is enabled:

- **SYN flood protection** — In a SYN flood attack, an attacker sends many TCP SYN packets to a server without any intention of completing the TCP handshake. The SYN packets are often sent with forged source IP addresses. If the rate of unanswered SYN-ACK packets exceeds the threshold set in the DoS protection options, the SYN flood protection is activated, and log data is generated.

  SYN flood protection can also be activated by the detection of too many half-open TCP connections. An attacker can create a large number of half-open TCP connections to use up resources on the Security Engine. To guard against this, you can set a limit for the number of half-open TCP connections per destination IP address. When the limit is exceeded, the SYN flood protection is activated, and log data is generated.

  When the SYN flood protection is activated, the Security Engine acts as a SYN proxy. The Security Engine completes the TCP handshake with the client, and only initiates the connection with the server after the client has completed the TCP handshake.
- **Slow HTTP request protection** — When the Security Engine receives an HTTP request, it analyzes the data transfer rate and length of time it takes to read the header fields of the HTTP request. If the sender of the request tries to keep the connection open for an unreasonable length of time, consuming excessive resources, the Security Engine block lists the sender's IP address for a specified length of time.

In addition, you can configure protection against DoS attacks that are based on TCP resets:

- **TCP reset protection** — In a TCP reset attack, an attacker sends forged TCP segments with an RST flag in an attempt to make the Security Engine drop TCP connections. The Security Engine detects the sequence numbers of the TCP RST segments to determine whether it is under a TCP Reset attack. If the segment's sequence number is in the current receive window but does not exactly match the expected sequence number, the Security Engine might send back a challenge ACK message. The connection is dropped only if the original sender responds to the challenge ACK with a new TCP reset that contains the correct sequence number.

You can enable all forms of DoS protection on all Security Engine and Virtual Engine types.

Note

If **Rate-Based DoS Protection Mode** is set to **On** or **Off** in the Engine Editor, you can override the setting in Access rules. If **Rate-Based DoS Protection Mode** is set to **Disabled**, you cannot enable rate-based DoS protection in an Access rule. You cannot override the TCP reset protection setting in Access rules.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click an Security Engine, then select **Edit <element type>**.
3) Browse to **Advanced Settings** > **DoS Protection**.
4) Configure the settings.

5) Click **Save and Refresh** to transfer the configuration changes.

Related tasks
Define Action options in Access rules on page 900

Configure scan detection settings

Before an attack, attackers might scan the network for open ports. When you enable scan detection on an engine, the number of connections or connection attempts within a time window is counted. If the number of events reaches the threshold set in the scan detection options, an alert is generated.

Note

If scan detection is enabled or set to Off in the Engine Editor, you can override the scan detection mode in Access rules. If scan detection is set to Disabled in the Engine Editor, you cannot enable scan detection in an Access rule.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click an engine element, then select **Edit <element type>**.
3) In the navigation pane on the left, browse to **Advanced Settings** > **Scan Detection**.
4) Configure the settings.
5) Click **Save and Refresh** to transfer the configuration changes.

Related tasks
Define Action options in Access rules on page 900

Using custom properties profiles to upload custom scripts

If you use custom scripts that you manually upload to the Security Engine, you can instead add the scripts to Custom Properties Profile elements.

If the custom properties profile is referenced in the configuration of the Security Engine, the script is automatically uploaded to all the Security Engine nodes when the policy is installed.

For example, if you use a custom script for the External Test for the Security Engine, you can use a custom properties profile to upload the script to the Security Engine. If the script is uploaded to the default location, you can refer to `/data/config/policy/latest/scripts/[script_name]` in the properties of the External Test.

This feature is supported on the Security Engine in the Engine, IPS, and Layer 2 Engine roles. For Virtual Engines, add the custom properties profile to the Master Engine.

You can upload custom scripts to the following paths:

```
/data
/data/config/base
/data/config/hooks/online
/data/config/hooks/offline
/data/config/hooks/standby
/data/config/hooks/policy-applied
/data/config/hooks/ve-active
/data/config/hooks/ve-deactive
```

Note

The scripts are not encrypted, even if the Security Engine configuration is otherwise encrypted.

In the custom properties profile, you can define additional attributes that your script can use. Additional attributes and their values are saved to the same location as your custom script in a file named `<script_name>_allow`. One attribute per line is stored in the file in the following format:

```
<attribute name>:<attribute value>
```

In this example, `/data/my_script.sh` has the additional attributes `test_attribute1` with the value `1` and `test_attribute2` with the value `2`. In the `/data` directory, there are two files:

```
my_script.sh
my_script.sh_allow
```

The file `my_script.sh_allow` contains the following:

```
test_attribute1:1
test_attribute2:2
```

For script examples, see Knowledge Base article 18290.

Note

Custom scripts for the Security Engine and custom scripts for Alert Chains in the SMC are configured separately and are separate scripts.

Create a Custom Properties Profile element

A custom properties profile contains a custom script that can run on the Security Engine.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Browse to **Other Elements** > **Engine Properties** > **Custom Properties Profiles**.

3) Select **New** > **Custom Properties Profile**.

4) Configure the settings, then click **OK**.

Enable a custom properties profile for an Security Engine

After you have created a custom properties profile, you must enable it in the properties of the Security Engine.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click a Security Engine, then select **Edit <element type>**.

3) Browse to **Advanced Settings** > **Custom Properties Profiles**.

4) Click **Add**, then select a Custom Properties Profile element.

5) Click **Save and Install**.

Part VIII
Routing

Contents

- Configuring routing and antispoofing on page 693
- Configuring dynamic routing on page 717
- Outbound traffic management on page 731
- Inbound traffic management on page 749
- Dynamic link selection on page 783

Use the SMC Client to configure static or dynamic routing, and use a Multi-Link configuration to manage and distribute inbound and outbound connections.

Chapter 42

Configuring routing and antispoofing

Contents

- Getting started with routing on page 693
- Routing configuration overview on page 694
- Add routers on page 696
- Add or view the default route on page 697
- Add static routes on page 697
- Using metrics or ECMP on multiple routes to the same destination on page 698
- Configure Forced Next Hop routing on page 701
- Configure policy routing on page 703
- Configuring multicast routing on page 703
- Configure DHCP message routing on page 707
- Check routes using the Route Query tool on page 709
- Remove static routes on page 709
- Modifying antispoofing on page 710
- Examples of routing configuration on page 713

Routing defines through which next hop router the Security Engine forwards traffic from a source address to a destination address. Antispoofing defines which addresses are considered valid source addresses for the networks connected to each interface.

Getting started with routing

To understand how the Security Engine reads routing definitions, look at the network interfaces in the routing or antispoofing configurations.

The SMC automates most of the routing and antispoofing configuration. Much of the configuration is generated automatically based on the IP addresses of the network interfaces.

The **Routing** pane in the Engine Editor shows the interfaces and a Network element for each network that is directly connected to the Security Engine. The routing information is stored on the Management Server. The Network is created based on the IP addresses that you define for each interface.

Use the **Display Mode** menu at the top of the pane to switch between the traditional tree view and a simple table view where IPv4 and IPv6 routes are shown in separate tables.

In the **Routing Tools** pane, you can view and create default routes on the **Default Route** tab, add simple routes on the **Add Route** tab, or check where packets with a certain IP address are routed on the **Query Route** tab. To add routes, we recommend you use the right-click menu in the tree view.

Routing decisions are made for each packet by matching from the most specific route definition to the most general. For packets subject to address translation or VPN tunneling, routing is always done after NAT or tunneling is applied using the translated IP addresses.

When the Security Engine reads routing definitions, it selects the most specific route and antispoofing definition it finds for each packet. The Security Engine:

1) Checks if there is a route defined for the specific destination IP address of the packet (Host elements).

2) Checks routes to the defined networks (Network elements).

3) Uses the default route (the Any network element) if no other route matches the packet's destination address. The default route typically leads to the Internet if the site has Internet access.

If there are overlapping definitions, the more specific one is considered first.

Routing configuration overview

Routing configuration involves adding a default route, adding routes to networks that are not directly connected, and adding routes to networks that can be reached through route-based VPNs.

Follow these general steps to configure routing:

1) Add the default route.

2) Add routes to networks that are not directly connected, but require a next hop gateway.

3) Add routes to networks that are reachable through the Tunnel Interfaces used in route-based VPNs.

Limitations

- Routing and antispoofing can only be configured for interfaces that have IP addresses. It is not possible to define routing or antispoofing for the following types of interfaces because they do not have IP addresses:
  - Capture Interfaces and Inline Interfaces on IPS engines or Layer 2 Engines, or on Master Engines that host Virtual IPS engines or Virtual Layer 2 Engines.
  - Capture Interfaces, Inline IPS Interfaces, and Inline Layer 2 Engine Interfaces on Security Engines in the Engine/VPN role.
  - All interfaces on Virtual IPS engines and Virtual Layer 2 Engines.
- Layer 2 physical interfaces on engines are not included in the routing and antispoofing configuration.
- The basic routing configuration does not determine which traffic is routed through policy-based VPNs. Routing is checked after policy-based VPN traffic is encapsulated inside encrypted packets with different source and destination IP address information.

Multi-Link

For information on using NetLinks to configure routing for Multi-Link, see the section about defining Multi-Link routes.

Related concepts
Defining Multi-Link routes on page 735

Default elements for routing and antispoofing

Networks that correspond to the IP addresses of each interface are automatically added to the routing and antispoofing configuration.

There is a default Network element called *Any network*, which is needed to define the default route. You must add a default route for packets whose destination IP address is not included anywhere else in the routing configuration.

On Single Engines, Single IPS engines, and Single Layer 2 Engines, if a dynamic IP address is added to an interface, there is an **Automatic Default Route** option that is selected by default. This option automatically adds a default route to the routing configuration.

There is also a default Host element for Bootp/DHCP clients in the antispoofing configuration of Engine, Master Engine, and Virtual Engine elements.

Adding routes for Master Engines and Virtual Engines

The need to configure routing can change depending on the role of the Security Engine and the types of interfaces that have been configured.

Basic routing information for networks directly connected to Master Engines and Virtual Engines is added automatically to both routing and antispoofing based on the IP addresses that you have defined for the interfaces. You must add a default route and any routes through next-hop gateways to networks that are not directly connected to the Master Engine or Virtual Engine.

On Master Engines, routing and antispoofing can only be configured for the Master Engine's system communications interfaces. No routes have to be defined if a Master Engine communicates only in its local IP network.

On Master Engines that host Virtual Engines, you can only add routes to interfaces that have IP addresses. Routing and antispoofing for Virtual Engines are configured in the same way as for Single Engines.

On Master Engines that host Virtual IPS engines or Virtual Layer 2 Engines, you can only add routes to Normal Interfaces that have IP addresses. It is not possible to add routes to Capture Interfaces or Inline Interfaces on Master Engines that host Virtual IPS engines or Virtual Layer 2 Engines.

Virtual IPS engines and Virtual Layer 2 Engines do not communicate directly with other SMC components. You cannot configure routing for Virtual IPS engines and Virtual Layer 2 Engines.

To transfer changes to the routing or antispoofing for a Master Engine, you must refresh the policy on the Master Engine. To transfer changes to the routing or antispoofing for a Virtual Engine, you must refresh the policy on the Virtual Engine.

Adding routes for IPS engines and Layer 2 Engines

Typically, only a default route through a Normal Interface is needed for IPS engines and Layer 2 Engines.

You might need to define a default route through a Normal Interface if SMC components are not on a directly connected network or if the Security Engine opens connections to a network that is not directly connected. You might need to add additional routes if one or more SMC components are not directly connected and cannot be reached through the default gateway. Capture Interfaces and Inline Interfaces on IPS engines and Layer 2 Engines do not have IP addresses. It is not possible to configure routing for these interfaces. You do not need to define routes if an IPS engine or Layer 2 Engine communicates only in its local IP network.

Adding routes for route-based VPNs

To configure routing for VPNs that use Route-based Tunnels elements, add remote networks that can be reached through the VPN tunnel to the tunnel interfaces. The routing defines which traffic is sent through the VPN tunnel.

Routing for route-based VPNs does not use Router or NetLink elements. Instead, you add the remote networks that are reachable through the VPN tunnel directly to the tunnel interfaces as if they were directly connected networks. Routes to local networks that are directly connected to the tunnel interface are automatically added if the tunnel interface has an IP address with a netmask other than the /32 host netmask.

For example, if a tunnel interface has 10.1.1.1/24 as the IP address and netmask, an automatic route to the 10.1.1.0/24 network is added for the tunnel interface.

Add routers

A Router element represents a next-hop gateway's IP address in routing configurations.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click a Security Engine, then select **Edit <element type>**.
3) Browse to **Routing**.
4) Right-click a Network element that is beneath an interface, then select **Add Router**.
5) Select ⋮ **More actions** > **New** > **Router**.
6) Enter a name and the IP address of the router, then click **OK**.
7) Select the Router you created, then select **Add**.
8) Click **OK**.

9) Click **Save**.

Add or view the default route

You can check which route is currently set as the default route for traffic leaving the Security Engine and set a new default route.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click a Security Engine, then select **Edit <element type>**.

3) Browse to **Routing**.



Tip

To view the full routing information for all interfaces, click **Expand All**.

4) To show the default route, click **Show Default Route** in the **Routing Tools** pane.
The current default route or routes are shown in bold font.

5) To set the default route, right-click a router or NetLink, then select **Set as Default Route**.
The Any network element is added beneath the router and is set as the default route.

6) Click **Save and Refresh**.

Add static routes

You can add a static route to a network in the routing configuration.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click a Security Engine, then select **Edit <element type>**.

3) Browse to **Routing**.



Tip

To view the full routing information for all interfaces, click **Expand All**.

4) Right-click the router or NetLink through which you want to route traffic, select **Add**, then select a Network element.

   For a route through a tunnel interface, add the Network element directly to the tunnel interface.

5) Click **Save and Refresh**.

Related tasks

Remove static routes on page 709

Using metrics or ECMP on multiple routes to the same destination

You can configure multiple routes through different next hop gateways to the same destination.

Route metrics

When you configure more than one route to a destination, the metric values configured for each route are used to determine the best path for traffic to use, with the lower value route winning. If the winning route becomes unavailable, the other configured route is automatically taken into use.

Equal-cost multi-path

If equal-cost multi-path (ECMP) is enabled on two or more routes to the same destination, traffic is sent over the routes using a hash-based balancing method. This approach can potentially increase the throughput of traffic through the use of multiple links. You can configure up to 16 routes.

Probing methods for route monitoring

You can configure probes that monitor the routes. Two methods are available:

- Ping — The Security Engine sends ICMP echo (ping) messages to a host in the destination network and expects responses.
- Next hop reach ability — The Security Engine checks that the next hop gateway can be contacted and that the link is up.

If the probing method fails, the route is considered to be unavailable.

The route monitoring continues after a route has been removed from the routing table, and if the route becomes available again, the route is added back to the routing table. The status of the routing table is shared by the dynamic routing protocols that the Security Engine is using. For example, if a route is unavailable, the Security Engine stops advertising the route.

When probing is enabled on routes, the **Allowed for Other Interfaces** option is automatically enabled in the Antispoofing configuration. When enabled through the right-click menu, the Security Engine does not block traffic coming from other interfaces.

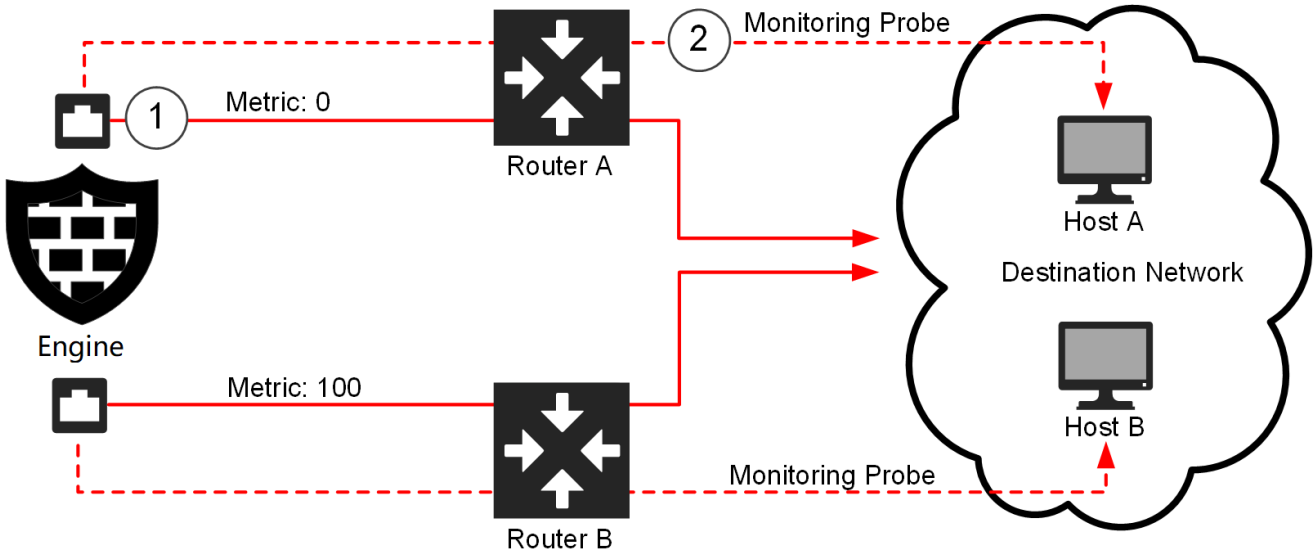
Multi-Link routes

With Multi-Link routes, you can use different outbound NetLinks to route traffic through different ISPs. You can also use route metrics or ECMP with NetLinks. For more information, see the section about defining Multi-Link routes.

Examples

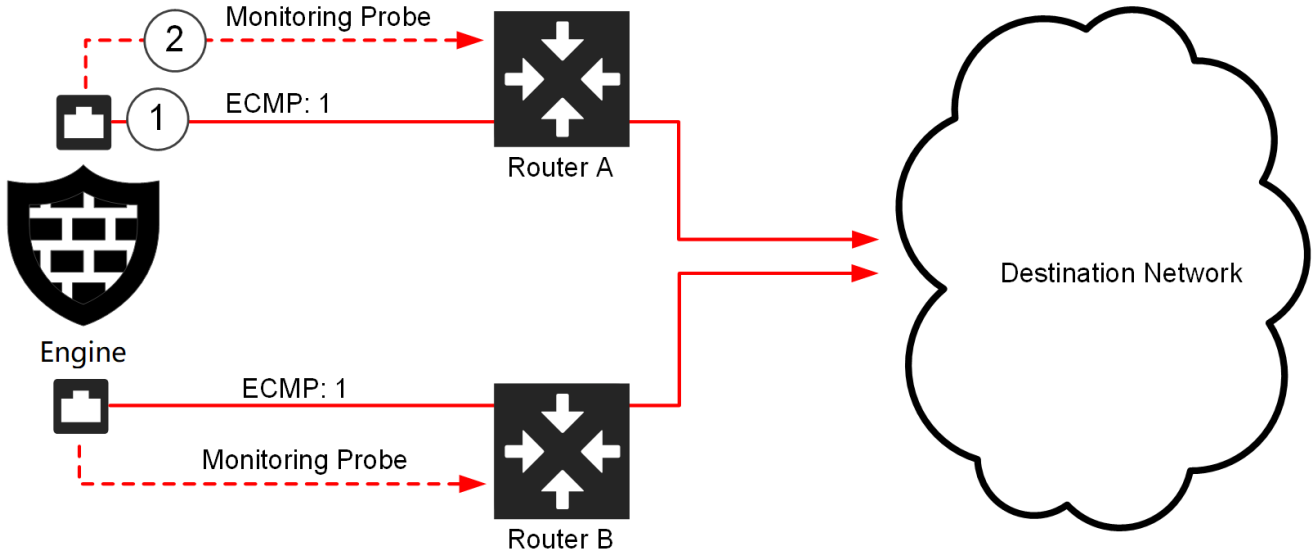
The following are examples of using route metrics or ECMP and the two route monitoring methods.

Route metrics and the ping method for route monitoring



1. The route through Router A has a metric value of 0, so it is the primary route. The route through Router B has a higher metric value, so it is the backup route.
2. The Security Engine pings a host in the destination network through Router A. If the host stops responding, connections failover to the route through Router B. Make sure the monitoring probes on each interface reach different hosts in the destination network that are reached through different routers.

ECMP and the next hop reach ability method for route monitoring



1. The routes through Router A and Router B have ECMP enabled, so traffic is sent over both routes using a hash-based balancing method.
2. The Security Engine checks that the links to Router A and Router B are up and that the routers are responding. If a link goes down and the router cannot be reached, the route is removed.

Related concepts
Defining Multi-Link routes on page 735

Add metrics or enable ECMP on a route

You can modify the metrics or enable ECMP on routes in the routing configuration.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click a Security Engine, then select **Edit <element type>**.
3) Browse to **Routing**.
4) Right-click a network that is beneath a router, NetLink, or tunnel interface, then select **Edit Route Metrics**.
5) Configure the settings, then click **OK**.
6) Click **Save and Refresh** to transfer the changed configuration.

Configure Forced Next Hop routing

The Forced Next Hop option can be used to force forward traffic to a specific tunnel, or next hop IP address overriding normal routing.

The advantage of using this option is that you can forward the traffic matching to a specific application, to a specific tunnel interface or next hop IP address.

To configure the **Forced Next Hop** option:

Steps

1) Select **Engine Configuration**.

2) Browse to **Policies** > **<Policy type>**.

3) Right-click a policy, then select **Edit <Policy name>**.

4) Double-click the **Action** cell or Right-click the **Action** cell and select **Edit Options** of an access rule. The **Select Rule Action Options** dialog-box is displayed.

5) Enter an IP Address in the **Forced Next Hop** field or click **Select** to select a host to force forward the matched traffic to the preferred destination.

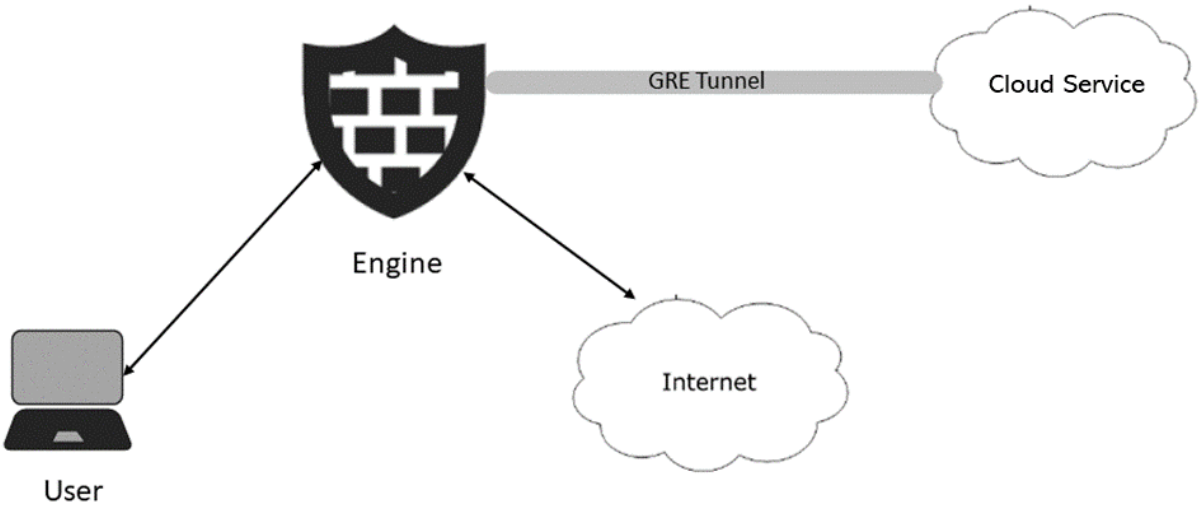
Note

For more details on **Action** options, refer to the *Define Action options in Access rules* section in the *Forcepoint Network Security Platform online help*.

Example scenario - Forced Next Hop

This section provides a simple example on how to configure Office 365 to route directly to the internet whereas other web traffics are routed through the cloud service using the GRE Tunnel.

Example of a Forced Next Hop routing



Access rule 1

Source	Destination	Services	Action
Intranet	Any	Office 365	**Allow**

Here, the Office 365 traffic is routed directly to the internet using the normal routing table.

Access rule 2

Source	Destination	Services	Action
Intranet	Any	http or https	**Allow** **Forced Next Hop** (IP Address behind the GRE Tunnel.

Here, the Forced Next Hop feature is configured for other web traffics, and hence other web traffics are routed to the cloud service using the GRE tunnel.

Note

When the traffic that match the Access rule which has Forced Next Hop configured:

1) Address specified in the Forced Next Hop is used in the route lookup instead of the destination address of the packet.

2) Reply packets of the connection are allowed in antispoofing, like they come from the IP Address that is specified in the Forced Next Hop configuration.

Configure policy routing

Policy routing allows you to route traffic based on both source and destination IP address. Policy routing is useful when you want to configure a different route for a destination based on the source address.

To configure alternative routes to the same destination, use route metrics, ECMP, or Multi-Link routing.

In most cases, there is no need to add policy routing entries. Policy routing entries are applied before the regular routes defined in the Routing tree (overriding those configurations if matches are found). They are processed exactly in the order specified in the **Policy Routing** view; the first matching policy routing entry is applied to a connection and any further entries are ignored.

Policy routing entries are not automatically added to the antispoofing configuration, so you might need to update antispoofing manually.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click a Security Engine, then select **Edit <element type>**.
3) Browse to **Routing** > **Policy Routing**.
4) Click **Add**.
5) Double-click the cells in the table to enter the routing information.

   Note

   Routing is done after NAT is applied.

6) Click **Save and Refresh** to transfer the changed configuration.

Related concepts

Modifying antispoofing on page 710

Configuring multicast routing

The Security Engine supports static multicast, IGMP-based multicast forwarding (IGMP proxying), and multicast routing using protocol-independent multicast (PIM). You can also configure dynamic multicast routing using the Free Range Routing (FRR) Suite on the command line.

IP multicasting is the transmission of an IP datagram to all hosts in a multicast host group, which is identified by a single destination IP address.

Static multicast is suitable for enduring configurations, while IGMP-based forwarding and PIM can save bandwidth and provide faster service.

Note

In addition to configuring the engine, routers and other network devices must be configured to allow IP multicasting along the path to the client machines.

Static multicast

Static multicast allows you to configure static routes for multicast traffic between a source IP address and Security Engine interface pair, and a destination (multicast) IP address and Security Engine interface pair. Static multicast is often used for enduring configurations, such as mutually agreed multicast traffic between organizations (for example, multicast news feeds and video conferences).

IGMP-based multicast forwarding

In IGMP-based multicast forwarding, the Security Engine maintains a list of subscriptions to the multicast host group and forwards multicast traffic to the subscribed hosts. The Security Engine periodically queries the downstream networks for hosts that want to join the multicast host group. The Security Engine also processes unsolicited IGMP join/leave requests received from downstream networks. As multicast traffic is only sent to the currently subscribed hosts, IGMP-based multicast forwarding can save bandwidth and provide faster service. IGMP-based multicast forwarding is only supported in tree topology networks. See RFC 4605 for more information.

Protocol-independent multicast (PIM)

Three variants of PIM are supported: PIM sparse mode (PIM-SM), PIM dense mode (PIM-DM), and PIM source-specific multicast (PIM-SSM), and the mode can be set separately for different multicast groups.

Limitations

- Only IPv4 addresses are supported in multicast routing.
- Multicast routing is not supported on Wireless Interfaces.
- IGMP Proxy mode for multicast routing is not supported on Tunnel Interfaces.

Related concepts

Multicasting vs. unicasting or broadcasting on page 1495

Related tasks

Configure PIM on page 726

Define static multicast

You can configure static routes for multicast traffic between a source IP address and engine interface pair and a destination (multicast) IP address and engine interface pair.

Note

Make sure your IPv4 Access rules allow the static multicast traffic to pass through the engine.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click a Security Engine, then select **Edit <element type>**.
3) Browse to **Routing** > **Multicast Routing**.
4) Select **Static** as the **Multicast Routing Mode**.
5) Click **Add**.
6) Configure the values for the entry.
7) Click **Save and Refresh** to transfer the changed configuration.

Related concepts

Getting started with Access rules on page 831

Define IGMP-based multicast forwarding

You can configure IGMP-based multicast forwarding for a specified Engine element.

IGMP-based multicast forwarding (IGMP proxying) is implemented on the Engine based on RFC 4605. IGMP-based multicast forwarding is only supported in tree topology networks. RFC 4605 includes support for source-specific multicast (SSM) with IGMP version 3. SSM is not supported with IGMP-based multicast forwarding. However, you can configure Access rules that filter multicast traffic based on the source.

The engine maintains a membership database of the subscriptions from the downstream networks and sends unsolicited reports or leaves on the upstream interface when the subscription database changes. It also sends IGMP membership reports when queried on the upstream interface.

Note

Make sure your IPv4 Access rules allow this traffic to pass through the engine.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click a Security Engine, then select **Edit <element type>**.

3) Browse to **Routing** > **Multicast Routing**.

4) Select **IGMP Proxy** as the **Multicast Routing Mode**.

5) (Optional) Select the **Upstream Interface** and the **IGMP Querier Settings** for it.
   - If the multicast servers and the hosts are in the local networks, or if you want to limit the multicast to the local networks, it is not necessary to define the Upstream Interface. In that case, leave **Not Set** selected for **Upstream Interface**.
   - (Engine Clusters only) You can only select as the **Upstream Interface** an interface that has a Cluster Virtual IP Address (CVI). You cannot select a Heartbeat Interface as the **Upstream Interface**.
   - You might need to select a specific IGMP Querier Settings element, for example, to troubleshoot multicast accessibility on hosts, or if some hosts use an earlier IGMP version.

6) Click **Add** to define **Downstream Interfaces**.
   The engine periodically queries the downstream networks for hosts that want to join or leave the multicast host group.
   A new entry appears in the table.

7) Click the **Interface** cell and select the Downstream Interface from the list.
   - You can use each interface only once in the IGMP proxy configuration.
   - (Engine Clusters only) The interface that you select as a Downstream Interface must have Node Dedicated IP Addresses (NDIs). It cannot be a Heartbeat Interface. It is recommended that the Node Dedicated IP Addresses increase in the same order on each node: for example, 192.168.1.10 and 192.168.2.10 for node A, and 192.168.1.11 and 192.168.2.11 for node B.

   Note

   The downstream interfaces must have the lowest IP addresses among all IGMP queries in the local networks.

8) Click the **IGMP Querier Settings** cell, then select the IGMP Querier Settings element that uses the IGMP version for the downstream interface.
   You might need to select a specific IGMP Querier Settings element, for example, to troubleshoot multicast accessibility on hosts, or if some hosts use an earlier IGMP version.

9) Click **Save and Refresh** to transfer the changed configuration.

Related concepts

Getting started with Access rules on page 831

Configure DHCP message routing

The engine can relay DHCP messages. If DHCP messages are routed through the engine (from a network segment to some other, isolated segment), you must enable DHCP relay on the engine interface properties for the interface where the DHCP requests are originating from (client's network).

Both IPv4 and IPv6 addresses are supported in DHCP relay. You can enable both IPv4 and IPv6 DHCP relay on the same engine interface.

This DHCP relay configuration does not relay DHCP messages from VPN Clients. You can configure DHCP message settings for VPN Clients in the **VPN** > **VPN Clients** branch of the Engine Editor.

Define a DHCP Server

A DHCP Server dynamically assigns IP addresses.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Network Elements**.
2) Browse to **Servers**.
3) Right-click **Servers**, then select **New** > **DHCP Server**.
4) In the **Name** field, enter a unique name.
5) Add one or more IP addresses.
   You can add both an IPv4 and an IPv6 address to the same DHCP Server element.
6) If there is a NAT device between a engine and the server so that the engine cannot connect directly to the IP address defined for the interface, select a location and add a contact address.

   Note

   Contact addresses are only supported for IPv4 addresses.
7) Click **OK**.

Related concepts

Define contact IP addresses on page 127

Related tasks

Create Location elements on page 127

Enable DHCP relay

You must select which interfaces perform DHCP relay. Activate the relay on the interface toward the DHCP clients.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click a Security Engine, then select **Edit <element type>**.

3) Browse to **Interfaces**.

4) Right-click a physical, VLAN, or Port Group interface, then select **Edit <interface type>**.
   - When configuring VLAN interfaces, set the DHCP relay separately for each VLAN.
   - When configuring Port Group interfaces, set the DHCP relay separately for each Port Group.

5) Enable IPv4 DHCP relay, IPv6 DHCP relay, or both.
   - On the **DHCPv4** tab, select **DHCPv4 Relay** from the **DHCP Mode** drop-down list.
   - On the **DHCPv6** tab, select **DHCPv6 Relay** from the **DHCP Mode** drop-down list.

6) On each tab where you enabled DHCP relay, select the DHCP server from the list of servers on the left, then click **Add**.

7) (Optional) From the **DHCP Relay** drop-down list, select the CVI or IP address that you want to use for DHCP Relay.

8) Click **OK**.

9) (Optional) Allow connections from interfaces on which DHCP relay is active to remote DHCP servers using automatic rules.
   a) Browse to **Policies** > **Automatic Rules**.
   b) Next to **Allow Connections from Local DHCP Relay to Remote DHCP Server**, select **Yes**.

   Note

   To relay DHCP messages through a policy-based VPN, you must add specific Access rules to allow the traffic. The Access rules must refer to the correct policy-based VPN.

10) Click **Save** to save and validate changes.

Check routes using the Route Query tool

The Route Query tool allows you to check where packets with a certain IP address are routed according to the current definitions. You can check that the routing is correct and quickly find a particular branch in the Routing tree.

The query does not take policy routing into consideration. The route query uses the configuration that is stored on the Management Server (shown in the SMC Client). You must refresh the policy of the affected Security Engines after completing the configuration to transfer the changed routing information.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click a Security Engine, then select **Edit <element type>**.
3) Browse to **Routing**.
4) In the **Routing Tools** pane, click the **Query Route** tab.
5) (Optional, NetLinks only) In the **Source** field, enter an IP address.
6) In the **Destination** field, enter an IP address.
7) Click **Query**.

Result

The route that matches the search criteria is shown.

Remove static routes

New networks are automatically added to routing when you change the properties of an interface. However, the Networks are not automatically removed, so you must check the routing configuration for obsolete entries and remove them manually.

Automatically added elements corresponding to a previous configuration are not automatically removed when the IP address of an interface is changed. Instead, the elements that belong to the old configuration are shown as invalid, and you must remove the obsolete elements manually.

You cannot remove elements that are added automatically based on the IP addresses of the interfaces while they are still in use. Use Access rules to control access to or from these addresses.

All additions and deletions in the routing configuration are automatically reflected in the antispoofing configuration. Manual definitions in the antispoofing configuration are preserved regardless of routing changes.

When you remove an element from the routing view, the element is not deleted from the SMC.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click a Security Engine, then select **Edit <element type>**.
3) Browse to **Routing**.
4) Right-click the element, then select **Remove**.
5) Click **Save and Refresh** to transfer the changed configuration.

Related concepts
Getting started with Access rules on page 831

Modifying antispoofing

IP address spoofing is an attack where the source IP address in a packet is changed to gain unauthorized access or to cause a denial-of-service. Such attacks can be prevented with antispoofing rules.

Antispoofing is intended to prevent malicious attempts to use a legitimate internal IP address to gain access from lower-security networks to higher-security networks by determining which addresses are valid source addresses for the networks connected to each interface. If an interface receives a packet with a source address that is not a valid source address for the networks that are connected to that interface, the packet is considered to come from a spoofed IP address.

Antispoofing is used on Engines, IPS engines, Layer 2 Engines, Master Engines, and Virtual Engines. Anti-spoofing rules are created automatically based on the static routing configuration for interfaces that have IP addresses. As long as no dynamic routing is used, there is usually no need to change the anti-spoofing configuration in any way.

Note

- An antispoofing configuration is not automatically generated for routes learned through dynamic routing protocols.
- Antispoofing related to dynamic routing is done by using the SMC GUI in the dynamic routing configuration in the Engine Editor. You must manually add hosts or networks to the **Additional Networks to Automatically Add to Antispoofing** table in the Engine Editor.

If you do modify the antispoofing configuration, manually changed entries are marked with a plus sign (+) for active entries or a minus sign (–) for disabled entries.

Limitations

Antispoofing cannot be configured for the following types of interfaces because they do not have IP addresses:

- Capture Interfaces and Inline Interfaces on IPS engines or Layer 2 Engines

- Master Engines that host Virtual IPS engines or Virtual Layer 2 Engines.
- Layer 2 physical interfaces on Engines.
- All interfaces on Virtual IPS engines and Virtual Layer 2 Engines.

Deactivate antispoofing for an IP address interface pair

In rare cases, you might need to change the default antispoofing definitions to make exceptions to antispoofing, for example, if you have defined policy routing manually.

Note

Errors in the routing configuration (in the SMC Client or in the surrounding network) can cause legitimate packets to be incorrectly identified as coming from a spoofed IP address. Always make sure that the routing is configured correctly before changing antispoofing. For example, routing loops generate log messages about spoofed packets. You cannot remove routing loops by changing antispoofing.

By default, the Security Engine interprets the antispoofing tree by selecting the most specific entry defined in the view. For example, a definition of a single IP address is selected over a definition of a whole network. If an IP address must be allowed access through two or more interfaces, the definition for each interface must be at the same level of detail for the IP address.

For example, if Interface A contains a Host element for 192.168.10.101 and Interface B contains a Network element for 192.168.10.0/24, connections from 192.168.10.101 are considered to be spoofed if they enter through Interface B, even though the address is included in the Network element. The antispoofing configuration must be changed to allow the address from Interface B.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click a Security Engine, then select **Edit <element type>**.
3) Browse to **Routing** > **Antispoofing**.
4) Right-click the interface, select **Add**, then select a Host or Network element.
5) (Optional) If you want to allow all connections from a network through a specific interface, right-click the network that is beneath the interface, then select **Absolute**.

   CAUTION

   Never mark the Any Network element as Absolute. Disabling antispoofing in this way is a security risk. Resolve large-scale antispoofing conflicts with specific antispoofing definitions or by changing routing.

   All IP addresses that belong to that network are now allowed for the interface. More specific antispoofing definitions for some addresses in the network can be defined for other interfaces.

6) Click **Save and Refresh** to transfer the configuration.

Activate antispoofing for routable IP addresses

In rare cases, you might need to disable or remove an allowed IP address for an Security Engine.

The antispoofing configuration shows the allowed addresses for each interface. There is rarely any need to change the automatically added entries. The preferred way of preventing routing for IP addresses is to make changes in the routing configuration.

Note

Disabling or removing elements in the antispoofing configuration prevents access. The IP addresses that a disabled or removed element represents are considered to be spoofed addresses.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click a Security Engine, then select **Edit <element type>**.

3) Browse to **Routing** > **Antispoofing**.

4) Right-click the element, then select **Remove** or **Disable**.
   If you disabled an element, you can re-enable the element in the same way.

5) Click **Save and Refresh** to transfer the configuration.

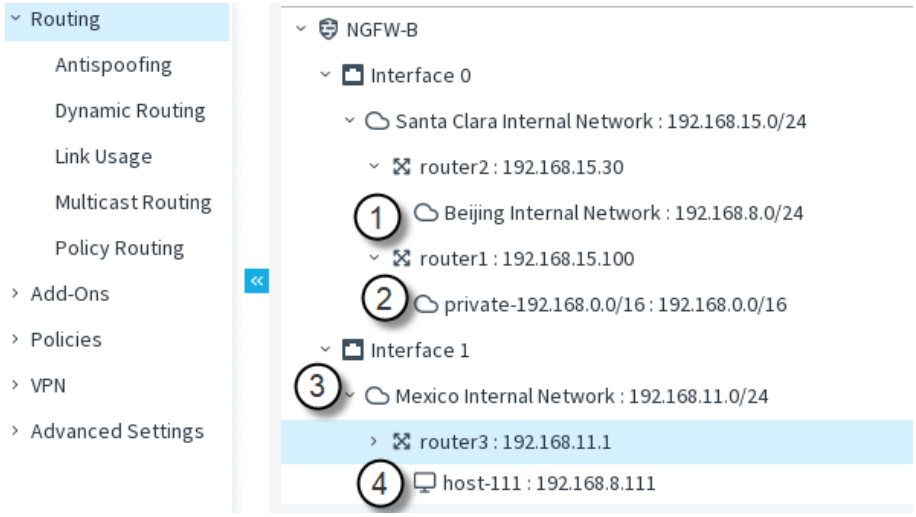
Examples of routing configuration

The SMC provides features that allow you to configure many routing scenarios.

Examples of route selection and antispoofing definitions

There are several considerations to take into account when configuring routing and antispoofing.

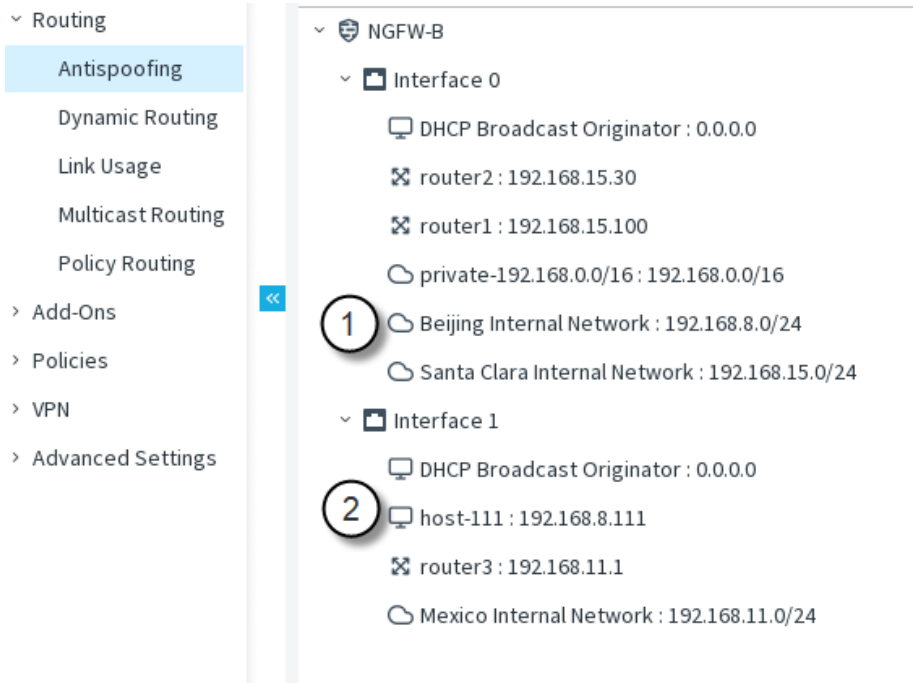
The more specific destination is considered first in routing



1. Traffic with a destination address from 192.168.8.0/24 is routed through router2 because it is the most specific route to those destinations.
2. All other traffic with a destination address from 192.168.0.0/16 is routed through router1 because it remains the most specific route to those destinations.
3. Interface 1 is directly connected to the 192.168.11.0/24 network. Traffic with a destination address from 192.168.11.0/24 is routed there because it is the most specific route to those destinations.
4. Traffic with a destination address of 192.168.8.111 is routed through router 3 because host-111 (192.168.8.111) has the most specific address.

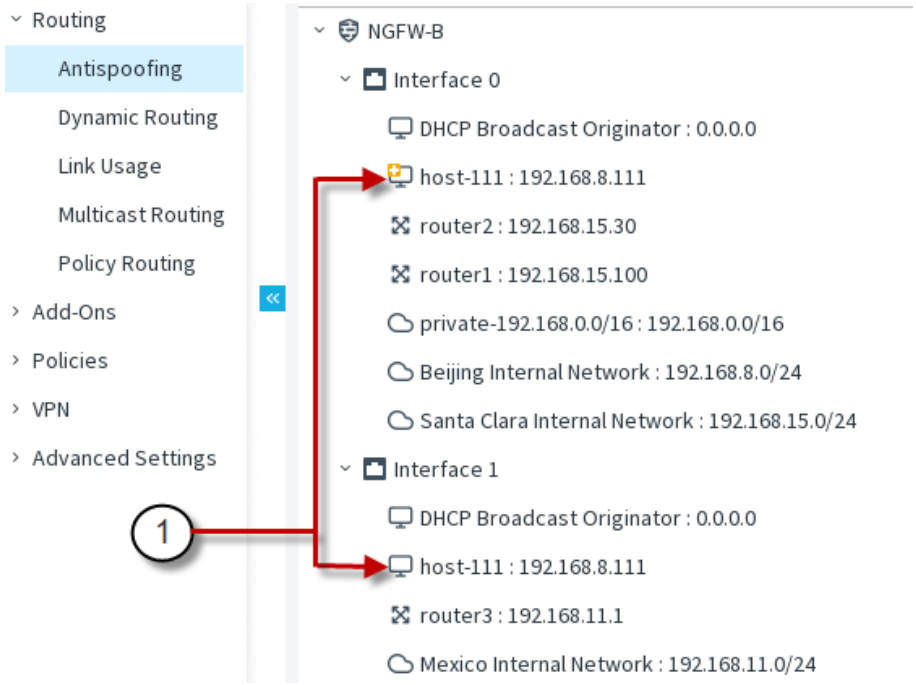
Only the most specific destination is considered valid in antispoofing

If an interface receives a packet with a source address that is not a valid address for the networks connected to that interface, the packet is discarded. This is the case, for example, when an external interface receives a packet with an internal source. The Security Engine selects the most specific antispoofing definition it finds for each packet. The following antispoofing configuration is based on the previous routing example.



1. Traffic from host-111 (192.168.8.111) is discarded if it originates from Interface 0 because it has the less specific definition for that address (network 192.168.8.0/24).
2. Traffic from host-111 (192.168.8.111) is only considered valid if it originates from Interface 1 because it has the most specific route to the address of the host.

Both interfaces are valid because they are equally specific



1 Both Interface 0 and Interface 1 are considered valid sources for host-111 (192.168.8.111) because the Host element is beneath both interfaces. The plus sign on the host on Interface 0 indicates that the host was manually added to the configuration. Traffic can originate from both Interface 0 and Interface 1.

Example: Routing traffic with two interfaces

To route traffic to two interfaces, add a router for each interface.

Company A needs to route traffic to the Internet as well as to the internal Network B, which is not directly connected to the company's Branch Office Engine. The company's administrators decide to create a separate route to the internal Network B and a default route for traffic to the Internet. The administrators:

1) Open the Engine Editor for the Branch Office Engine, then browse to **Routing**.

2) Add a Router and Network B beneath Interface 0.

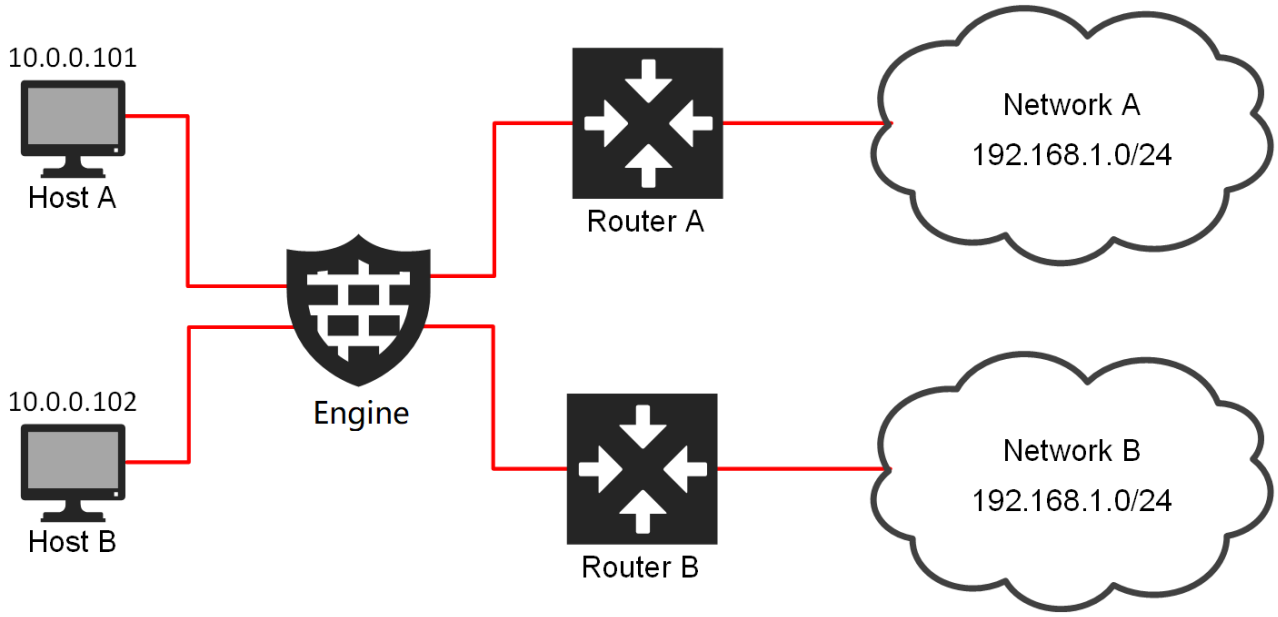
3) Add a Router beneath Interface 1, then set it as the default route.

4) Save and refresh the policy.

Example: Using policy routing to route traffic to networks that use the same address space

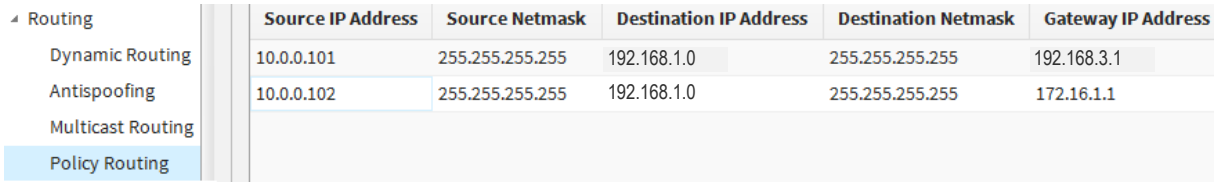
To manage traffic to networks that share the same address space, create a policy routing entry for each host.

The network is connected to two partners: Network A and Network B. The Network A and Network B partners use the same address space in their internal networks.



There are two hosts in the network. Host A works only with the Network A partner and Host B only with the Network B partner. The administrators decide to use policy routing to route the traffic between the network and the two partner sites. The administrators:

1) Create policy routing entries for Host A and Host B on the Engine HQ Cluster.



2) Change the antispoofing rules so that they take into account the routing defined with the policy routing entries.

3) Save and refresh the policy.

Chapter 43

Configuring dynamic routing

Contents

- Getting started with dynamic routing on page 717
- Dynamic routing configuration overview on page 718
- Creating elements for dynamic routing on page 718
- Configure BGP on page 722
- Configure OSPFv2 on page 724
- Configure PIM on page 726
- Preview the dynamic routing configuration in FRR syntax on page 727
- Using the command line to configure dynamic routing on page 728
- Restart dynamic routing processing on page 730

With dynamic routing, Security Engines automatically change their routing when the network topology changes. The Security Engines can also exchange information about appropriate routing paths.

Getting started with dynamic routing

You can configure dynamic routing in the SMC Client or on the command line of the Security Engine.

Note

We recommend that you configure dynamic routing only if you have experience and knowledge of the general principles of dynamic routing and Free Range Routing (FRR) syntax. Poorly configured dynamic routing can lead to adverse effects on networks, such as packets being directed to the wrong interfaces or data being exposed to undesired networks.

The two methods for configuring dynamic routing are separate, and using both methods on the same engine is not supported.

- If you configure dynamic routing for an Security Engine on the command line, that dynamic routing configuration is not shown in the SMC Client.
- If you configure dynamic routing for an Security Engine in the SMC Client, the configuration overwrites the dynamic routing configuration that you configured on the command line.

In the SMC Client, you can configure dynamic routing using border gateway protocol (BGP), open shortest path first (OSPFv2), and protocol-independent multicast (PIM).

When using FRR commands and syntax on the command line of the Security Engine, the supported dynamic routing protocols are BGP, OSPFv2, OSPFv3, PIM, and routing information protocol (RIP).

Dynamic routing is supported on Single Engines, Engine Clusters, and Virtual Engines. The rules that allow dynamic routing traffic are created automatically.

Note

An antispoofing configuration is not automatically generated for routes learned through dynamic routing protocols. You must manually add hosts or networks to the **Additional Networks to Automatically Add to Antispoofing** table in the Engine Editor.

Before installing a policy that includes a dynamic routing configuration, you can preview and save the configuration in FRR syntax.

Dynamic routing configuration overview

The steps for configuring dynamic routing depend on the protocol you use and whether you configure it in the SMC Client or on the engine command line.

Follow these general steps to configure dynamic routing in the SMC Client:

1) Create the elements for dynamic routing that you want to use in the configuration.
2) Configure the settings for dynamic routing for the Engine, Engine Cluster, or Virtual Engine in the Engine Editor.
3) (Recommended) Preview the dynamic routing configuration in Free Range Routing (FRR) format to check that the routing configuration is what you intended.

Follow these general steps to configure a dynamic routing protocol on the command line:

1) Create an empty protocol-specific configuration file.
2) Start the protocol-specific routing daemon.
3) Start the FRR command shell, then enter the dynamic routing commands.

Creating elements for dynamic routing

Before enabling dynamic routing in the SMC Client, you can optionally create customized elements for dynamic routing.

If the default elements meet your needs, it is not necessary to create custom elements.

BGP elements for dynamic routing

Element	Description
BGP Profile	This element contains administrative distance, redistribution, and aggregation settings. You can use this element in multiple Engines, Virtual Engines, and Engine Clusters. There is a default BGP Profile element that is used automatically.

Element	Description
Autonomous System	An autonomous system (AS) is a group of IP routing prefixes controlled by an administrative entity. Each AS has a unique identifying number. This element is used to define the AS number. The number determines whether internal BGP (BGP peers have the same AS number) or external BGP (BGP peers have different AS numbers) is used.
BGP Peering	An element that you place between a Engine and another Engine or an External BGP Peer element in the **Routing** view in the Engine Editor. It contains the parameters to define a BGP peering relationship, and implements inbound and outbound policies through access lists and the Route Map.
BGP Connection Profile	This element can be used to set a password for TCP MD5 authentication between BGP peers, and to set timers, such as the keepalive value. You can use this element in multiple BGP Peering elements. There is a default BGP Connection Profile element that is used automatically.
External BGP Peer (Optional)	Use this element to define the IP address and AS number of a BGP peer that is under the administrative control of a third party, such as another organization or SMC.

OSPFv2 elements for dynamic routing

Element	Description
OSPFv2 Domain Settings	Use this element to set the area border router (ABR) type, throttle timer settings, and the max metric router link-state advertisement (LSA) settings. There is a default OSPFv2 Domain Settings element that is used automatically.
OSPFv2 Profile	This element contains administrative distance and redistribution settings. There is a default OSPFv2 Profile element that is used automatically.
OSPFv2 Interface Settings	This element contains the interface settings. You can also select the type of authentication to use. There is a default OSPFv2 Interface Settings element that is used automatically.
OSPFv2 Area	Use this element to set the area ID and type. You can also configure the ABR settings and virtual links.
OSPFv2 Key Chains	This element contains the keys that you can use if the **Message Digest** authentication type is used in an OSPFv2 Interface Setting element.

PIM elements for dynamic routing

Element	Description
PIM Profile	This element contains the multicast groups and determines the PIM mode that is used. There is a default PIM Profile element that is used automatically.
PIM Interface Settings	This element contains various settings, such as the designated router (DR) priority and the zone boundary router (ZBR) group. You must also set the IGMP Querier Settings element to use. There is a default PIM Interface Settings element that you can use.
IGMP Querier Settings	This element determines the IGMP version to use. You can also set the query interval and robustness settings. You can use this element when configuring both PIM and IGMP Proxy multicast. There are three default IGMP Querier Settings elements; one element for each version of IGMP.

Common elements for dynamic routing

Element		Description
Route Map (Optional)		An element that contains rules to control or manipulate received or advertised routes based on matching conditions. The rules can be edited in the same way as in policy elements.
Access Lists (Optional)	IP Access List	A filtering element that you can use in the Matching Condition cell of a Route Map rule. You can filter by subnet. This element can also be used to implement inbound and outbound filtering policies in a BGP Peering element.
	IP Prefix Access List	Similar to an IP Access List, but includes the prefix length or netmask as a filter. This element can be used in the Matching Condition cell of a Route Map rule and in a BGP Peering element.
	Community Access List	A filtering element used for the BGP community attribute. If the Type is Expanded, you can use regular expressions in the Community Regular Expression cell. This element can be used in the Matching Condition cell of a Route Map rule and in a BGP Peering element.
	Extended Community Access List	A filtering element that represents the extended version of the Community Access List. If the Type is Expanded, you can use regular expressions in the Community Regular Expression cell. This element can be used in the Matching Condition cell of a Route Map rule.
	AS Path Access List	A filtering element for the BGP AS path attribute. You can enter regular expressions for matching AS paths (a sequence of AS numbers). This element can be used in the Matching Condition cell of a Route Map rule and in a BGP Peering element.

Create core elements for dynamic routing

Create the elements that contain dynamic routing configuration information. The elements can be used in multiple Engines, Virtual Engines, and Engine Clusters.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Browse to **Other Elements** > **Dynamic Routing Elements**.
3) Browse to **BGP Elements**, **OSPFv2 Elements**, or **PIM Elements**.
4) Right-click the type of element that you want to create, then select **New <element type>**.
5) Adjust the properties as needed, then click **OK**.

Create Access List elements

You can use Access List elements as a Matching Condition in a Route Map element. You can also use certain Access Lists in BGP Peering elements.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Browse to **Other Elements** > **Dynamic Routing Elements**.
3) Right-click the type of Access List that you want to create, then select **New <Access List>**.
4) Configure the settings, then click **OK**.

Create Route Map elements

Use Route Map elements in more complex networks to control or manipulate routes. You can use Access List elements as a Matching Condition in a Route Map rule.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Browse to **Other Elements** > **Dynamic Routing Elements**.
3) Right-click **Route Maps**, then select **New Route Map**.
4) Enter a name for the Route Map, then click **OK**.
5) To add a rule, right-click the Route Map, then select **Add Rule**.
6) (Optional) To add a Rule Section, right-click the Route Map, then select **Add Rule Section Before** or **Add Rule Section After**.

   You can use Rule Sections to group similar rules. Double-click the Rule Section to give it a name. Because you can jump to a Rule Section in the Goto action, use a unique name.
7) To add a Matching Condition for a rule, double-click the **Matching Condition** cell. Configure the rules, then click **OK**.
8) To select the action for the rule, right-click the **Action** cell, then select **Permit** or **Deny**.
9) To change the Route Entry Settings, double-click the **Route Entry Settings** cell, configure the settings, then click **OK**.

10) (Optional) To jump to a Route Map, double-click the **Call** cell, then select a Route Map.

11) To select how the rules processing continues, double-click the **Goto** cell.
- **Finish** — Further rules are not processed. This is the default action.
- **Goto Next Rule** — The next rule is processed.
- **Goto Rule Section** — If you have created rule sections, you can jump to a rule section.

Configure BGP

You can configure dynamic routing using the BGP protocol on Engines, Virtual Engines, and Engine Clusters.

Before you begin

We recommend that you first create elements for configuring dynamic routing.

Enable BGP on the Engine, Engine Cluster, or Virtual Engine

You must enable BGP for the Engine, Engine Cluster, or Virtual Engine in the Engine Editor.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click an engine, then select **Edit <element type>**.

3) In the navigation pane on the left, browse to **Routing** > **Dynamic Routing**.

4) In the BGP section, select **Enabled**.

5) (Optional) Enter the Router ID in the **Router ID** field.

6) If you do not want to use the default BGP Profile, select another BGP Profile element from the **BGP Profile** drop-down list.

7) Select an Autonomous System element from the **Autonomous System** drop-down list.

8) (Optional) To add Announced Networks, click **Add**.
You can add hosts, networks, or groups that contain both hosts and networks.

9) (Optional) You can configure the following BGP Monitoring Protocol (BMP) options for the Engine, Engine Cluster, or Virtual Engine in the Engine Editor. It is used to monitor BGP sessions and send the monitored data from BGP routers to the network management entities.

10) (Optional) To add a network to the Antispoofing pane, click **Add** next to the **Additional Networks to Automatically Add to Antispoofing** table.
You can add hosts, networks, or groups that contain both hosts and networks.

11) Click **Save**.

Next steps

You are now ready to add a BGP Peering element to the engine on the **Routing** branch.

Add the BGP Peering element

You must add a BGP Peering element to the routing configuration.

Before you begin

You must enable BGP on the engine. To establish a BGP peering relationship with another Engine or Virtual Engine, you must enable BGP on both Engines.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click an engine, then select **Edit <element type>**.

3) In the navigation pane on the left, browse to **Routing**.

4) Right-click a directly connected network under an interface or right-click a Tunnel Interface, then select **Add BGP Peering**.

5) Browse to the BGP Peering element, click **Add**, then click **OK**.
The element is added beneath the network or Tunnel Interface.

6) (Optional) To change the CVI address:
   a) Right-click the BGP Peering element.
   b) Select **IP addresses**, then select an IP address.

7) Right-click the BGP Peering element, then select **Add Engine** or **Add External BGP Peer**.

8) Browse to the Engine, Virtual Engine, Engine Cluster, or External BGP Peer element, click **Add**, then click **OK**.
   The Engine, Virtual Engine, Engine Cluster, or External BGP Peer element is added beneath the BGP Peering element.

9) Click **Save**.

Configure OSPFv2

You can configure dynamic routing using the OSPFv2 protocol on Engines, Virtual Engines, and Engine Clusters.

Before you begin

We recommend that you first create elements for configuring dynamic routing.

Enable OSPFv2 on the Engine, Engine Cluster, or Virtual Engine

You must enable OSPFv2 for the Engine, Engine Cluster, or Virtual Engine in the Engine Editor.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click an engine, then select **Edit <element type>**.

3) In the navigation pane on the left, browse to **Routing** > **Dynamic Routing**.

4) In the OSPFv2 section, select **Enabled**.

5) (Optional) Enter the Router ID in the **Router ID** field.

6) If you do not want to use the default OSPFv2 Profile, select another OSPFv2 Profile element from the **OSPFv2 Profile** drop-down list.

7) (Optional) To add a network to the antispoofing configuration, click **Add** next to the **Additional Networks to Automatically Add to Antispoofing** table.
   You can add hosts, networks, or groups that contain both hosts and networks.

8) Click **Save**.

Next steps

You are now ready to add an OSPFv2 Area element to the engine on the **Routing** branch.

Add an OSPFv2 Area element

You must add one or more OSPFv2 Area elements to the routing configuration.

Before you begin

You must have enabled OSPFv2 on the engine.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click an engine, then select **Edit <element type>**.
3) In the navigation pane on the left, browse to **Routing**.
4) Right-click a directly connected network under an interface or right-click a Tunnel Interface, then select **Add OSPFv2 Area**.
5) Browse to the OSPFv2 Area element, click **Add**, then click **OK**.
   The element is added beneath the network or Tunnel Interface.
6) (Optional) Override the OSPFv2 Interface Settings element that is used in the OSPFv2 Area element.
   a) Right-click the OSPFv2 Area element, then select **Override Interface Settings**.
   b) Browse to the OSPFv2 Interface Settings element, click **Add**, then click **OK**.
7) (Optional) To change the communication mode, right-click the OSPFv2 Area element, then select **Force Communication Mode** and one of the following options:
   - **Not Forced** — This is the default option. Use this option if you do not want to use a different communication mode.
   - **Point-to-Point** — Select this option if the network is between two routers only.
   - **Passive** — Select this option if you do not want the interface to send routing updates, but you do want the interface to be announced.
   - **Unicast** — Select a Host element, click **Add**, then click **OK**.
8) Click **Save**.

Configure PIM

You can configure dynamic multicast routing using a variant of the PIM protocol on Engines, Virtual Engines, and Engine Clusters.

Before you begin

We recommend that you first create elements for configuring dynamic routing.

Three variants of PIM are supported: PIM sparse mode (PIM-SM), PIM dense mode (PIM-DM), and PIM source-specific multicast (PIM-SSM), and the mode can be set separately for different multicast groups. The variant to use depends on many factors, such as the network topology, the number of recipients, and what the existing infrastructure supports.

- Use PIM-SM when multicast traffic needs to travel greater logical distances; over WAN links, for example. There is usually intermittent multicast traffic.
- Use PIM-DM when there are a lot of potential recipients and the logical distances are short; within a WAN, for example. There is usually a high, constant load of multicast traffic. Messages are flooded over the network, so this is not always the most efficient use of bandwidth.
- Use PIM-SSM when you want receivers to be able to specify the source IP address of the requested multicast stream. This is the most efficient use of bandwidth, but all devices and client applications must have support for IGMPv3 and this variant of PIM. However, SSM mapping allows IGMPv2 requests to be converted into IGMPv3 requests.

Enable PIM on the Engine, Engine Cluster, or Virtual Engine

You must enable PIM for the Engine, Engine Cluster, or Virtual Engine in the Engine Editor.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click an engine, then select **Edit <element type>**.
3) In the navigation pane on the left, browse to **Routing** > **Multicast Routing**.
4) From the **Multicast Routing Mode** drop-down list, select **PIM**.
5) If you do not want to use the default PIM Profile, select another PIM Profile element from the **PIM Profile** drop-down list.
6) (Optional, PIM-SM only) Expand the **Bootstrap Settings** section, then modify the settings.
7) Click **Save**.

Next steps

You are now ready to add a PIM Interface Settings element to the engine on the **Routing** branch.

Add a PIM Interface Settings element

You must add one or more PIM Interface Settings elements to the routing configuration.

Before you begin

You must have enabled PIM on the engine.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click an engine, then select **Edit <element type>**.
3) In the navigation pane on the left, browse to **Routing**.
4) Right-click a directly connected network under an interface or right-click a Tunnel Interface, then select **Add PIM Interface Settings**.
   When configuring a Engine Cluster, the interface must have a CVI address defined.
5) Browse to the PIM Interface Settings element, click **Add**, then click **OK**.
   The element is added beneath the network or Tunnel Interface.
6) (Optional) To change the IP address used, right-click the PIM Interface Settings element, then select **IP Addresses** and an IP address.
7) Click **Save**.

Preview the dynamic routing configuration in FRR syntax

To make sure the dynamic routing configuration done in the SMC Client is what you intended, preview the configuration in Free Range Routing (FRR) syntax.

You can also save the configuration as a .conf file, for testing with the FRR suite.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click the Security Engine, then select **Configuration** > **Dynamic Routing** > **Dynamic Routing Configuration Preview**.
   The **Dynamic Routing Configuration Preview** dialog box opens, and you can see the configuration in FRR syntax.

3) (Optional) To save the configuration, click **Browse**, select a folder on the local workstation, then click **Save**.

Using the command line to configure dynamic routing

You can enter dynamic routing commands through the vtysh shell of the Free Range Routing (FRR) suite on the command line of the Security Engine.

Important

If you make changes to dynamic routing using the SMC Client, all changes that you made on the command line are overwritten.

Dynamic routing is supported on Single Engines, Engine Clusters, and Virtual Engines.

For Virtual Engines, you must enter the following command on the Master Engine to access the command line of the Virtual Engine:

```
se-virtual-engine
```

Engine Clusters or Master Engines automatically synchronize the FRR configuration files between nodes, and the routing daemon is automatically stopped and started. Edit the FRR configuration on the node that is currently active for dynamic routing. To check that you are working on the correct node, enter the following command:

```
sg-dynamic-routing status
```

The output must include the following string:

```
FRR suite status (node is Active)
```

Detailed instructions for configuring the FRR suite can be found on the FRR Software Routing Suite website at https://docs.frrouting.org/en/latest/.

Note

Configuring dynamic routing requires SSH access to the command line. We recommend that you disable SSH access whenever it is not needed. Make sure that your Access rules allow SSH access to the Security Engines from the administrators' IP addresses only.

Related tasks
Send commands to Virtual Security Engines on page 368

Edit the FRR configuration file

To configure dynamic routing on the command line of the Security Engine, edit the Free Range Routing (FRR) configuration file.

The protocol-specific configuration file (for example, ospfd.conf) must exist in the /data/config/frr directory, and the routing daemon must be running before you configure dynamic routing using vtysh. Otherwise, vtysh cannot save protocol-specific configurations.

Steps

1) Open an SSH connection to the command line of the Security Engine.

2) To create empty protocol-specific configuration files, enter, for example:

```
sg-dynamic-routing sample-config ripd ospfd bgpd
```

3) To start the routing daemon, enter:

```
sg-dynamic-routing restart|force-reload
```

4) To start the FRR command shell, enter:

```
vtysh
```

Related concepts
Adjust Master Engine clustering options on page 681

Related tasks
Access the Security Engine command line on page 364
Adjust general Engine clustering options on page 677

Related reference
Security Engine commands on page 1445

Back up or restore command-line dynamic routing configurations

You can back up and restore the current dynamic routing configuration that has been configured on the command line of the Security Engine.

Note

This operation does not back up or restore configurations done in the SMC Client. Those configurations are included in Management Server backups.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click the Security Engine, select **Configuration** > **Dynamic Routing**, then select one of the following:
   - **Backup** — To save the current dynamic routing configuration. This action overwrites the previously saved backup.
   - **Restore** — To restore a saved dynamic routing configuration.

   A progress summary opens in a new tab.

3) Click **Close** when the backup or restoration is complete.

Related tasks

Edit the FRR configuration file on page 729

Restart dynamic routing processing

If routes are not being updated as expected, you can restart the dynamic routing processing.

Note

Instead of issuing the BGP stop and start commands, the Security Engine issues the BGP graceful restart command.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click the Security Engine, then select **Configuration** > **Dynamic Routing** > **Restart**.

3) In the **Confirmation** dialog box, add an optional comment for auditing, then click **Yes**.

Chapter 44

Outbound traffic management

Contents

- Getting started with outbound traffic management on page 731
- Defining Multi-Link routes on page 735
- Enable outbound traffic management using element-based NAT on page 739
- Manually configuring outbound traffic management on page 739
- Monitoring and testing outbound traffic management on page 744
- Examples of manually configuring Multi-Link on page 745

You can use Multi-Link to distribute outbound traffic between multiple network connections and to provide high availability and load balancing for outbound traffic.

Getting started with outbound traffic management

Multi-Link for VPN provides high availability for outbound connectivity so that business-critical traffic gets through even when one or more Internet connections fail.

Multi-Link distributes and balances the load of outbound traffic between multiple network connections.

Single Engines, Engine Clusters, and Virtual Engines can balance outbound traffic between two or more network links (NetLinks) using the Multi-Link feature. NetLinks are combined into Outbound Multi-Link elements. The NetLinks can represent different types of ISP connections and they can have different speeds. The NetLinks must be added under the appropriate interfaces in the Routing tree to support Multi-Link. You can also use Multi-Link with aggregated link interfaces.

Multi-Link allows you to:

- Balance outbound traffic between two or more alternative network links to increase the available bandwidth.
- Ensure that outbound network connectivity remains available even if network links fail. When a network link fails, the engine detects this and stops forwarding traffic through the failed link.

You can create multiple Outbound Multi-Link elements, and each NetLink can belong to more than one Outbound Multi-Link element at the same time. Multiple Outbound Multi-Link elements can be useful, for example, when you want a certain type of traffic to be balanced only between some of the NetLinks, and another type of traffic to be balanced between all NetLinks.

To improve redundancy, we recommend connecting each link through different physical network equipment (such as routers).

Outbound traffic management has the following limitations:

- Multi-Link is supported on Single Engines, Engine Clusters, and Virtual Engines.
- Multi-Link for outbound load balancing is only supported for IPv4 traffic.

- VPN traffic is balanced independently from the settings covered in this configuration when the engine is the VPN endpoint.

Related concepts

Getting started with inbound traffic management on page 749
Network address translation and how it works on page 119
VPNs and Multi-Link for VPN on page 1165

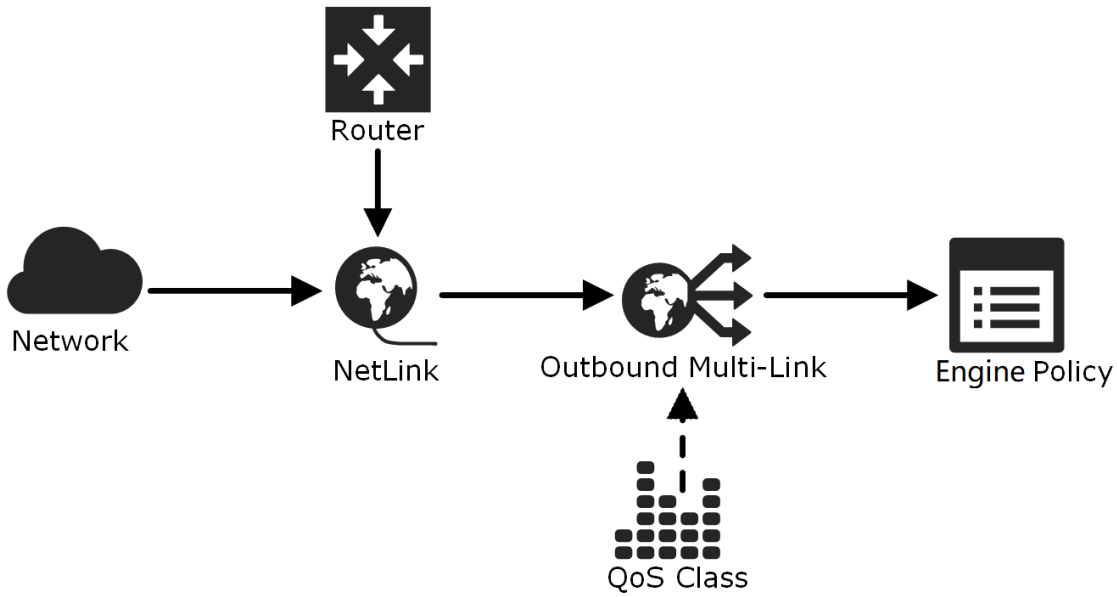
Outbound traffic management configuration overview

Multi-Link uses multiple NetLinks to balance the load of outbound traffic and ensure high availability of Internet connectivity. With each new outbound connection, the Security Engine selects the fastest route for the connection from the available NetLinks.

There are two ways to configure outbound traffic management. You can:

- Use element-based NAT to automatically create an Outbound Multi-Link element based on the routing configuration.
- Manually create an Outbound Multi-Link element and define specific NAT rules for outbound load balancing.

Elements in a Multi-Link configuration



This illustration shows the elements that are used to configure Multi-Link. Each NetLink element contains a Router element and a Network element. The Router element represents the router for that network connection. The Network element represents the set of public IP addresses allocated by the provider of the network connection. NetLinks are added to the Routing tree under the Interface IDs and the Modem numbers that represent the physical interfaces or the mobile broadband modems toward the routers used for the Internet connections.

Multiple NetLinks are combined into an Outbound Multi-Link element. Outbound Multi-Link elements are the central elements used to configure load balancing for outbound traffic. If you use element-based NAT, the default

NAT address works like an Outbound Multi-Link and the NAT rules are automatically generated. You can also use an Outbound Multi-Link element as a NAT address in a NAT definition. You can alternatively use the Outbound Multi-Link elements in the Engine Policy's NAT rules to implement outbound load balancing manually.

The configuration consists of these general steps:

1) Configure routing with at least two NetLinks.

2) If you want the Security Engine to select the NetLink based on the type of traffic, create QoS Classes and assign them to traffic.

3) Enable outbound traffic management in one of the following ways:
   - To use element-based NAT for outbound traffic management, enable automatic default NAT in the Engine Editor.
   - To manually configure outbound traffic management, create an Outbound Multi-Link element to group your NetLinks and define traffic management settings.

4) To manually configure outbound traffic management, create NAT rules for outbound load balancing in the Engine Policy.

Related concepts
Element-based NAT and how it works on page 633
Quality of Service (QoS) and how it works on page 973

Using standby NetLinks for high availability in outbound traffic management

Standby NetLinks allow you to define a NetLink as a backup that is only activated when all primary NetLinks are unavailable. Using standby NetLinks minimizes the use of NetLinks that are more expensive or otherwise less preferable, while still ensuring the high availability of Internet connectivity.

To test which NetLinks are available, the status of the NetLinks is monitored. As soon as one or more primary NetLinks become active again, the standby NetLinks are deactivated. The deactivated NetLink still handles the previously established connections, but new connections are no longer sent to the standby NetLink. You can define multiple active and standby NetLinks.



Note

You must configure probing settings for each NetLink when you use active and standby NetLinks.

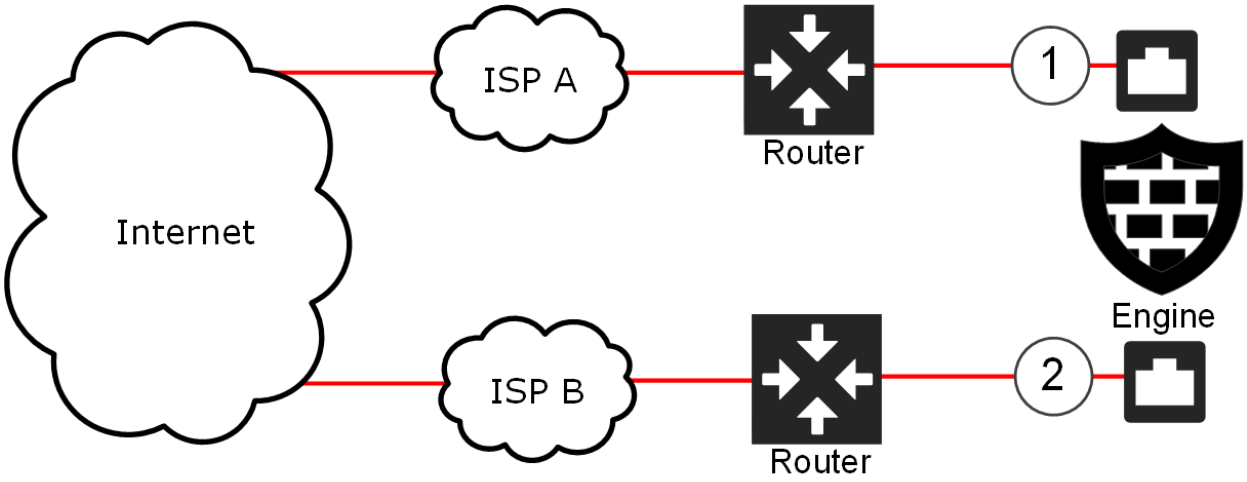
When load balancing is used with standby NetLinks, traffic is only distributed between the NetLinks that are currently active. Standby NetLinks are only activated when failure is detected, not to balance the load.

Using standby NetLinks on the same interface as other NetLinks affects the interface speed you enter in the configuration of QoS for engine interfaces.

Using Multi-Link with a Single Engine for outbound traffic management

You can configure Multi-Link with Single Engines where each interface is connected to a router from a different ISP.

Single Engine interfaces with Multi-Link



CAUTION

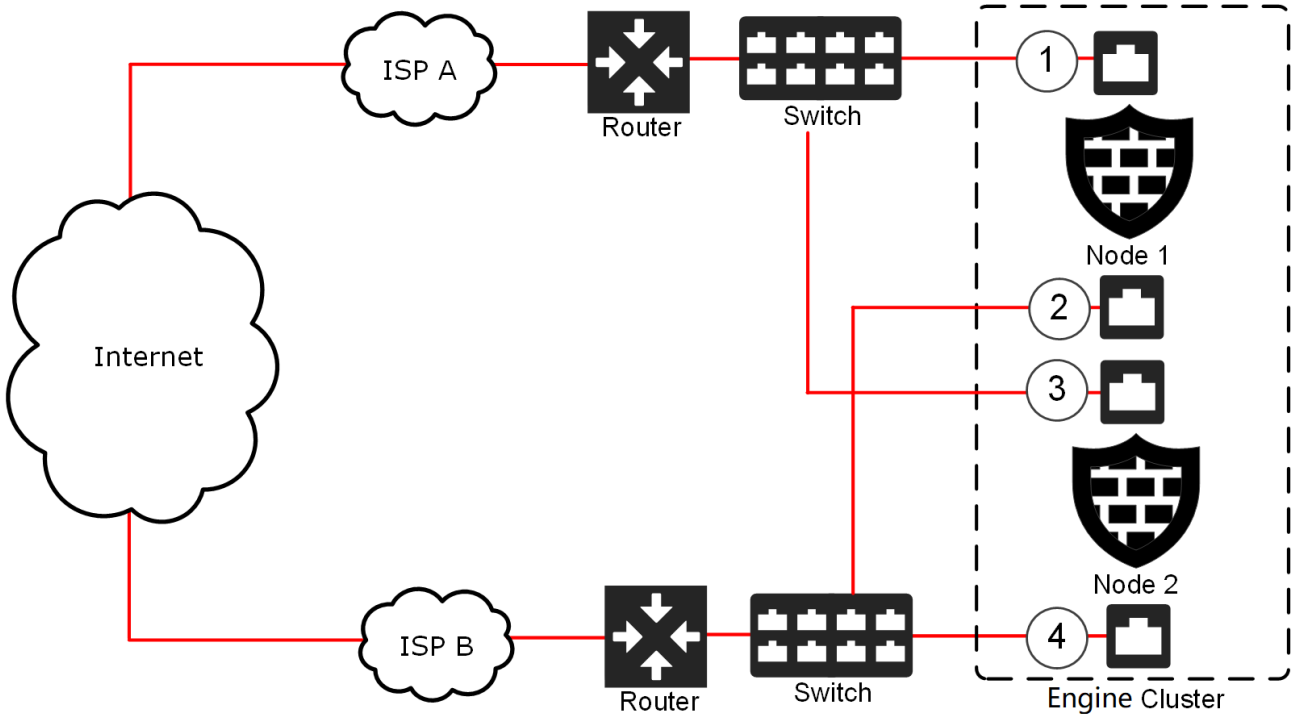
We do not recommend defining two or more IP addresses for a single physical interface. If two IP addresses are defined, the router behind the interface forwards the traffic to the different ISPs. This creates an additional single point of failure at the intermediate router, and the associated cabling and network cards.

1. interface 1 is used as the network interface for Internet traffic that is routed through ISP A.
2. Interface 2 is used as the network interface for Internet traffic that is routed through ISP B.

Using Multi-Link with a Engine Cluster for outbound traffic management

You can configure Multi-Link with a Engine Cluster consisting of two nodes by using two routers and two switches and physically connecting each node to each switch.

Cluster interfaces for Multi-Link



1. Interface 1 on Node 1 is used as the CVI for Internet traffic that is routed through ISP A.
2. Interface 2 on Node 1 is used as the CVI for Internet traffic that is routed through ISP B.
3. Interface 1 on Node 2 is used as the CVI for Internet traffic that is routed through ISP A.
4. Interface 2 on Node 2 is used as the CVI for Internet traffic that is routed through ISP B.

Both nodes have one physical interface for each CVI, so that both nodes are physically connected to both switches that are connected to the routers that lead to the Internet.

CAUTION

We do not recommend configuring Multi-Link by connecting two CVIs to a single router, which in turn connects to both ISPs. It creates a single point of failure.

Defining Multi-Link routes

When you use Multi-Link routing, traffic can use multiple network connections to reach its destination.

You can define Multi-Link routes for Security Engines and Virtual Engines and for both IPv4 and IPv6 traffic.

Note

Multi-Link for outbound load balancing is only supported for IPv4 traffic.

NetLink elements represent the network connections for Multi-Link. Usually, a NetLink element represents an ISP connection. However, NetLinks can also represent a leased line, xDSL, or any other type of network connection mediated by your engine.

There are two types of NetLinks: *static* and *dynamic* NetLinks.

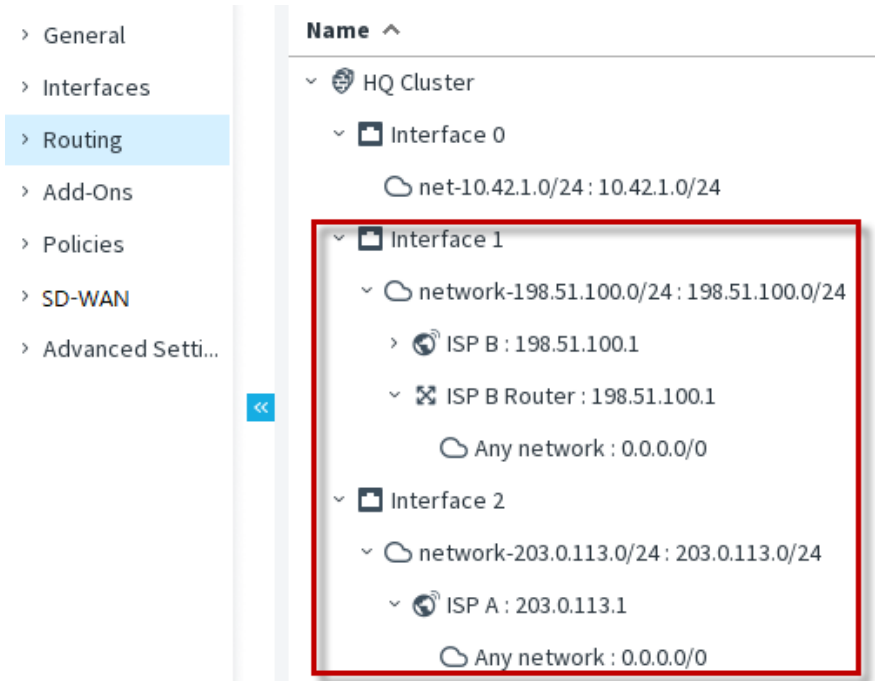
- Static NetLinks are supported in the routing configuration for both IPv4 and IPv6 traffic.
- Dynamic NetLinks are supported only with Single Engines. Dynamic IP addresses are not supported for Engine Clusters.

Note

If you configure wireless Multi-Link on a Modem Interface of a Single Engine, only Dynamic NetLinks are supported. Modem Interfaces always have dynamic IP addresses.

A Router or a NetLink element represents a next-hop gateway that forwards packets to networks that are not directly connected to the Security Engine. Tunnel interfaces for route-based VPNs do not use Router or NetLink elements. Instead, networks that are reachable through the VPN tunnel are added directly to the tunnel interface as if they were directly connected networks.

NetLinks in the Routing tree view



This illustration shows a Multi-Link default route to the Internet using the Any network element through the ISP A and ISP B NetLinks. We recommend using separate network interfaces for each NetLink.

For each NetLink, a range of IP addresses is defined for applying NAT to the source IP address of an outbound connection that goes through the NetLink. Element-based NAT or a NAT rule in the Engine Policy defines the Outbound Multi-Link element that is used for outbound Multi-Link connections.

Defining Multi-Link routes consists of these general steps:

1) Create a NetLink for each alternative route.

2) Add Networks under the NetLinks in the **Routing** tree to define a route.

3) (Optional) Configure route metrics or ECMP.

Monitoring the status of NetLinks

To monitor the status of the links, define the probe IP addresses in the NetLink properties.

The Security Engine sends ICMP messages to make sure that a link is still available. Only NetLinks that are used in an Outbound Multi-Link element are probed. Status monitoring is not available for NetLinks that are only used in Routing.

Probing is always recommended, and it is mandatory with the following features:

- The ratio-based load-balancing method
- Fail over to standby NetLinks
- Inbound traffic balancing with dynamic DNS updates

Make sure the Probe IP Addresses that you select produce a reliable measurement of the link performance. For example, probing the IP address of an ADSL router usually succeeds even if the ISP network is unreachable. Probing the default gateway provided to you by the ISP might succeed even if the ISP is not able to forward traffic anywhere outside the ISP's own network. You can add several alternative probe IP addresses to avoid probing failures caused by a probed host going down.

Create NetLink elements for Multi-Link configuration

NetLink elements usually represent Internet connections, but you can also use NetLinks to represent other network connections. You can use NetLinks to define alternative routes that lead to the same destination IP addresses.

Although some NetLink element settings are optional, we recommend that you configure all settings for a fully functional configuration.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.

2) Browse to **Traffic Handlers**.

3) Right-click **Traffic Handlers**, then select **New** > **Static NetLink** or **New** > **Dynamic NetLink**.

4) Configure the settings, then click **OK**.

Related concepts
Getting started with outbound traffic management on page 731

Add Multi-Link routes

You can add a Multi-Link route for a Engine or Virtual Engine.

Note

A Network element represents the IP addresses of a network or a sub network to which the router or the NetLink forwards the traffic, or the IP addresses of a network that is reachable through a Route-based Tunnels.

Note

The network interfaces for the NetLinks must have a node dedicated IP address (NDI) defined for all nodes in clusters. Otherwise, the IP address of the interface marked as the default IP address for outgoing connections is used, which can lead to incorrect load balancing.

CAUTION

If you use Multi-Link with IGMP proxy multicast routing, make sure that you do not create routing loops. If you add a NetLink to the upstream interface, do not add a NetLink to any downstream interface.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click a Security Engine, then select **Edit <element type>**.
3) Browse to **Routing**.
4) In the routing tree, browse to the NetLink through which you want to route traffic.
5) Add the element that contains the IP addresses that are routed through this interface:
   - To add a default route, right-click the NetLink element, then select **Set As Default Route**.
     This inserts the default element **Any Network** under the NetLink in the routing tree.
   - To add a route using a Host or Network element, right-click the NetLink element, select **Add**, then select the element.
6) Click **Save and Install**.
   If you are configuring routing for Master Engines, install or refresh the policy on the Master Engine and the Virtual Engines to transfer the changed configuration.

Enable outbound traffic management using element-based NAT

When you use element-based NAT for outbound traffic management, an Outbound Multi-Link element is automatically created based on the routing configuration.

Before you begin

Configure routing with at least two NetLinks.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click an engine element and select **Edit <element type>**.
3) In the navigation pane on the left, browse to **Policies** > **Element-based NAT**.
4) From the **Use Default NAT Address for Traffic from Internal Networks** drop-down list, select **Automatic**.
5) Click **Save and Refresh**.

Related concepts

Element-based NAT and how it works on page 633

Manually configuring outbound traffic management

Outbound Multi-Link elements combine NetLinks and set options for the high availability and load balancing features.

NetLink selection for load balancing can be based on either of two methods:

- **Round Trip Time** — The engine periodically probes the NetLinks to test them for speed and selects the fastest available active NetLink for each new outbound connection. NetLink performance is measured for each new TCP connection by sending the initial SYN request to the destination through all available NetLinks. When the destination host sends the SYN-ACK reply, the NetLink that receives the reply first is used to establish the TCP connection. The engine cancels the slower connection attempts by sending a TCP Reset (RST) to the destination through the other NetLinks.
  The fastest route is automatically selected for each connection. Information about the performance of each NetLink is cached, so no new measurement is made if a new connection is opened to the same destination within a short time period.

To use the round trip time method, you must configure the probing settings in the Static NetLink or Dynamic NetLink properties.

- **Ratio** — Traffic is distributed between all available active NetLinks according to the relative bandwidth of each NetLink. The NetLink with the highest bandwidth is assigned the largest portion of the traffic. The bandwidths of the other NetLinks are automatically compared to the bandwidth of the NetLink with the most bandwidth to produce a ratio for distributing the traffic.
  When the volume of traffic is low, the ratio of actual traffic distribution is approximate. When the volume of traffic is high, the ratio of traffic handled by each NetLink is closer to the ratio calculated from the link capacity.

You can optionally assign QoS Classes to NetLinks in the Outbound Multi-Link element to specify which traffic is routed through which NetLink. NAT rules can alternatively be used to select a particular link, but if you use QoS Classes, traffic can still fail over to other links if the selected link fails.

The same QoS class can be assigned to more than one NetLink in the same Outbound Multi-Link element to balance traffic through those selected NetLinks when those links are available. If you want to use QoS class to specify which traffic uses which NetLink, you must assign the QoS class to the traffic in an Access rule or with the QoS policies based on the DSCP codes in the traffic.

To manually configure outbound traffic management, follow these general steps:

1) Create an Outbound Multi-Link element to group your NetLinks and define traffic management settings.

2) In the Engine Policy, create NAT rules for outbound load balancing.

Create Outbound Multi-Link elements

Outbound Multi-Link elements group together NetLinks as a single entity and set options for load balancing.

Before you begin

Configure routing with at least two NetLinks.

To use the round trip time NetLink selection method for load balancing, you must configure the probing settings in the Static NetLink or Dynamic NetLink properties.

You can create multiple Outbound Multi-Link elements, and each NetLink can belong to more than one Outbound Multi-Link element at the same time.

To create the Outbound Multi-Link element, you must define the following:

- Which NetLinks are included
- The load-balancing method for determining which link is selected for each new outbound connection
- Whether each NetLink in the Outbound Multi-Link element is an Active or Standby NetLink

The Outbound Multi-Link elements you create do not work on their own; you must use them in the Engine Policy's NAT rules to select traffic for outbound load balancing.



Note

If you use element-based NAT, the default NAT address works like an Outbound Multi-Link element. You do not need to create an Outbound Multi-Link element when you use element-based NAT.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.
2) Browse to **Traffic Handlers**.
3) Right-click **Traffic Handlers**, then select **New** > **Outbound Multi-Link**.
4) In the **Name** field, enter a unique name.
5) From the **Method** drop-down list, select the method for link selection.

Next steps

Select NetLinks for the Outbound Multi-Link element.

Select NetLinks for Outbound Multi-Link elements

NetLink elements represent ISP connections in Outbound Multi-Link elements.

Before you begin

You must have NetLink elements that represent your ISP connections.

To use the round trip time NetLink selection method for load balancing, you must configure the probing settings in the Static NetLink or Dynamic NetLink properties.

Each NetLink element contains the IP addresses that are used for translating source IP addresses (NAT) so that outgoing connections receive the correct IP address depending on the ISP. This configuration allows the correct routing of the return packets. Each NetLink must have a unique IP address space.

A NetLink can be either static or dynamic. Dynamic NetLinks are supported only for Single Engines. In addition, only Dynamic NetLinks are supported if you configure Multi-Link using mobile broadband modems with a Single Engine. You can use the same NetLink with several engines. If you want to use NetLinks with a engine that has several interfaces with dynamic IP addresses, you must create a separate Dynamic NetLink element for each interface with a dynamic IP address.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.
2) Browse to **Traffic Handlers**.
3) Right-click the Outbound Multi-Link element, then select **Properties**.

4) On the **General** tab of the **Outbound Multi-Link Properties** dialog box, click **Add**.

5) From the **NetLink** drop-down list, select a NetLink element.

6) From the **Network** drop-down list, select the Network element that represents the IP address space in the directly connected external network of this network link.

7) In the **Selected Range** fields, specify the IP address range for dynamic source address translation (NAT) for the internal source IP addresses on this NetLink.
   To define a single IP address, enter the same address in both fields.

8) From the **Type** drop-down list, select an option to specify how traffic is routed through the NetLink.

9) (Optional) Select the QoS Classes for traffic handled by this NetLink, then click **Add**.
   You can use the QoS classes to assign the NetLink with or without activating the actual QoS features.
   - You can select the same QoS class for several NetLinks to balance the traffic between the NetLinks. If none of the NetLinks with the appropriate QoS class are available, or if the traffic has not been assigned a QoS class, the traffic is distributed between the NetLinks according to the Method you specify in the Outbound Multi-Link element properties.
   - QoS classes are assigned based on ToS codes in network traffic or in the Access rules. Traffic that has been assigned the selected QoS class uses this NetLink if the NetLink is available.

10) Click **OK**.
    The NetLink is listed in the **Multi-Link Members** list.

11) Click **OK** in the **Outbound Multi-Link Properties** dialog box.

Next steps

Continue in one of the following ways:

- (Optional) Define how information about the performance of each NetLink is cached.
- If you manually configured Outbound Multi-Link elements, define NAT rules using the source NAT addressing defined in the Outbound Multi-Link element.

Define destination cache settings for NetLinks in Outbound Multi-Link elements

Information about the performance of each NetLink is cached so that no new measurement is made if a new connection is opened to the same destination within a short time period. You can optionally define the duration of the cached information.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.

2) Browse to **Traffic Handlers**.

3) Right-click the Outbound Multi-Link element, then select **Properties**.

4) In the **Outbound Multi-Link Properties** dialog box, select the **Advanced** tab.

5) Deselect **Default**.

6) Enter the **Timeout** (in seconds) after which a new measurement of NetLink performance is made.
The default is 3600 seconds.

7) Enter the number of **Maximum Retries** for checking each NetLink.
The default is 2.

8) Click **OK**.

Create outbound load-balancing NAT rules

If you manually configured Outbound Multi-Link elements, define NAT rules using the source NAT addressing defined in the Outbound Multi-Link element.

Note

When you use element-based NAT, NAT rules are automatically generated.

You can create several Outbound Multi-Link elements to balance different types of traffic differently. Only the part of traffic that matches a NAT rule with the Outbound Multi-Link element is balanced between different links.

Some protocols cannot use dynamic NAT based on IP/port translation. To provide high availability and load balancing for these protocols, you can use static NAT with an Outbound Multi-Link element in an outbound load-balancing NAT rule. When static NAT is used, the size of the source network must be the same as the size of the Multi-Link network.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Browse to **Policies** > **Engine Policies**.

3) Right-click a Engine Policy and select **Edit**.

4) On the **IPv4 NAT** tab, add a rule.

5) Specify the source, destination, and service according to the traffic that you want to match.

6) Double-click the **NAT** cell.

7) Configure the settings, then click **OK**.

8) Click **Save and Install**.

Monitoring and testing outbound traffic management

You can monitor the status of the NetLinks in the SD-WAN Dashboard. NetLinks are shown as ISP connections.

Status is displayed based on the results of status probes to the Probing addresses configured in the NetLink elements.

You can test that the outbound traffic management configuration works as intended by generating traffic and then disabling network connections, for example, by unplugging cables. In a cluster, you can disable the same link on all active nodes to simulate ISP failure.

Note

By default, if you unplug a cable from a node in a cluster, the automatic engine tester recognizes this as a malfunction and turns the node offline unless it is the last node that remains online. To override this automatic test, command all nodes to “Lock Online” before the test. After the tests, command the notes to "Online".

Most connections that are currently open on the link that goes down are cut and must be reopened by the communicating applications because the external IP address used in the communications changes to an address of the new link. Any type of traffic can switch to using a different link transparently if it is transported through a Multi-Link VPN between two Engines.

If the Ratio-based load-balancing method is used, a link is used until a failure is detected based on the link status probes to the probing addresses configured in the NetLink elements. With the Round Trip Time method, new connections are directed to working links because the link that responds the fastest is chosen even if probing is not configured or working.

Examples of manually configuring Multi-Link

Multi-Link technology has many uses, some of which might apply to your situation.

Example: Using Multi-Link to prepare for ISP downtime

This scenario shows an example of using Multi-Link to configure connections to two ISPs. This configuration guarantees connectivity if one ISP fails.

Company A wants to make sure that their Internet connection remains available even when one ISP connection fails. The company has subscribed to one Internet connection each from ISP A and ISP B. The administrators decide to use Multi-Link to guarantee high availability of Internet connectivity.

The administrators do the following:

1) Create NetLink elements to represent connections to ISP A and ISP B.

2) Place the ISP A and ISP B NetLinks under the correct interfaces in the **Routing** tree.

3) Create an Outbound Multi-Link element and add the ISP A and ISP B NetLinks to it.

4) Define the following NAT rule in the Engine Policy so that traffic from the internal network to destinations that are not internal is handled by the Outbound Multi-Link element:

Source	Destination	Service	NAT
Internal Network element	Not Internal expression	ANY	Dynamic load balancing: My Multi-Link

Example: Excluding NetLinks from handling a QoS Class of traffic

Using Multi-Link technology and QoS classes, you can exclude VoIP traffic from a satellite connection.

Company B has three Internet connections: ISP A, ISP B, and ISP C, which is a satellite link. Because of the long latency in satellite connections, the administrators do not want any VoIP traffic to be routed through ISP C. They decide to use QoS classes so that VoIP traffic is only routed through ISP A and ISP B.

To do this, the administrators:

1) Create NetLink elements to represent connections to ISP A, ISP B, and ISP C.

2) Place the ISP A, ISP B, and ISP C NetLinks under the correct interfaces in the **Routing** view.

3) Define a QoS class and assign it to VoIP traffic.

4) Create an Outbound Multi-Link element and add the ISP A, ISP B, and ISP C NetLinks to it.

5) Select the QoS class for the ISP A NetLink and the ISP B NetLink in the Outbound Multi-Link element properties. No QoS class is assigned to ISP C.

6) Define the following NAT rule for outbound load balancing in the Engine Policy:

Source	Destination	Service	NAT
ANY	ANY	ANY	Dynamic load balancing: Multi-Link Element

Example: Balancing traffic according to link capacity

To manage ISP connections having very different bandwidths, you configure an Outbound Multi-Link element and select ratio load balancing.

Company B has three ISP connections that have different bandwidths:

- ISP A 20 Mbit/s
- ISP B 10 Mbit/s
- ISP C 4 Mbit/s

The administrators want the traffic to be divided between the NetLinks according to the ratio of their relative bandwidths. This means that ISP A handles twice as much traffic as ISP B and five times as much traffic as ISP C. The administrators have already created and configured NetLink elements to represent each ISP connection, so now they:

1) Combine the NetLinks for each ISP connection into an Outbound Multi-Link element and select the Ratio load-balancing method.

2) Define the following NAT rule for outbound load balancing in the Engine Policy:

Source	Destination	Service	NAT
ANY	ANY	ANY	Dynamic load balancing: Multi-Link Element

Example: Balancing traffic between Internet connections

This scenario shows an example of using Multi-Link to balance the load of traffic between two smaller Internet connections instead of one larger connection.

The administrator at Company B determines that a 4 megabyte Internet connection is required to handle the volume of traffic their network receives. However, Company B is a small company on a tight budget, and the cost of a single 4 megabyte connection is too high. The administrator decides to subscribe to one 2 megabyte

connection each from ISP A and ISP B, and use Multi-Link to balance the load of traffic between the two connections to reduce costs.

The administrator:

1) Creates NetLink elements to represent connections to ISP A and ISP B.

2) Places the ISP A and ISP B NetLinks under the correct interfaces in the **Routing** tree.

3) Creates an Outbound Multi-Link element and adds the ISP A and ISP B NetLinks to it.

4) Defines the following NAT rule in the Engine Policy so that traffic from the internal network to destinations that are not internal is balanced by the Outbound Multi-Link element (My Multi-Link):

Source	Destination	Service	NAT
Internal Network element	Not Internal expression	ANY	Dynamic load balancing: My Multi-Link

Chapter 45

Inbound traffic management

Contents

- Getting started with inbound traffic management on page 749
- Create Server Pool elements on page 753
- Configure server availability monitoring on page 755
- Create Access rules to allow the type of traffic that is handled by the Server Pool on page 756
- Enable Server Pool load balancing using NAT rules on page 757
- Enable Server Pool load balancing using Access rules on page 758
- Configuring dynamic DNS updates for Server Pools on page 759
- Using Server Pool Monitoring Agents on page 762
- Examples of Server Pools on page 780

Inbound traffic management ensures that services remain available even when one or more servers or NetLinks fail, and balances the load of incoming traffic more efficiently between a group of servers. Inbound traffic management is not supported on Layer 2 Engines or on layer 2 physical interfaces on Engines.

Getting started with inbound traffic management

Server Pool elements provide inbound traffic management for traffic to servers in the protected network.

A Server Pool is a built-in load balancer in the Security Engine that distributes incoming traffic between a group of servers. Server Pools primarily provide load balancing and high availability for two or more servers that offer the same service. You can also us Server Pools with Multi-Link to control which NetLink incoming traffic uses so that clients can access the Server Pool through multiple Internet connections.

Inbound traffic management can:

- load-balance incoming traffic between several servers to even out their workload
- monitor the status of each server so that the traffic is not directed to servers that are unavailable or overloaded
- send dynamic DNS (DDNS) updates to a DNS server to prevent incoming traffic from attempting to use a non-functioning NetLink in a Multi-Link configuration.

Clients make their incoming connections to the external addresses of the Server Pool. The Security Engine decides which server handles the connection and uses NAT to translate the public IP addresses to the private IP address of that server.

The Engine uses the internal IP address of each member of the Server Pool to select which server handles which traffic that arrives at the Server Pool's external address. When you define a Host element, you enter the IP address that the engine uses to contact the server.

The server load is distributed to the Server Pool members based on each server's availability. The server availability can be monitored by periodically sending ICMP echo requests (ping) or by periodically sending TCP

strings to check that the expected response is returned. You can also use Server Pool Monitoring Agents installed on each server for advanced monitoring of server availability and status.

If the server availability monitoring reports a server failure, the server is removed from the Server Pool and the connections are distributed to the remaining servers. When a server is removed from the Server Pool, traffic from existing connections might still be sent to the server, but new connections are not sent to the failed server. When a previously unavailable server comes back online, existing connections are not redistributed, but some of the new connections that are opened are again directed to the server that rejoins the pool.

Inbound traffic management has the following limitations:

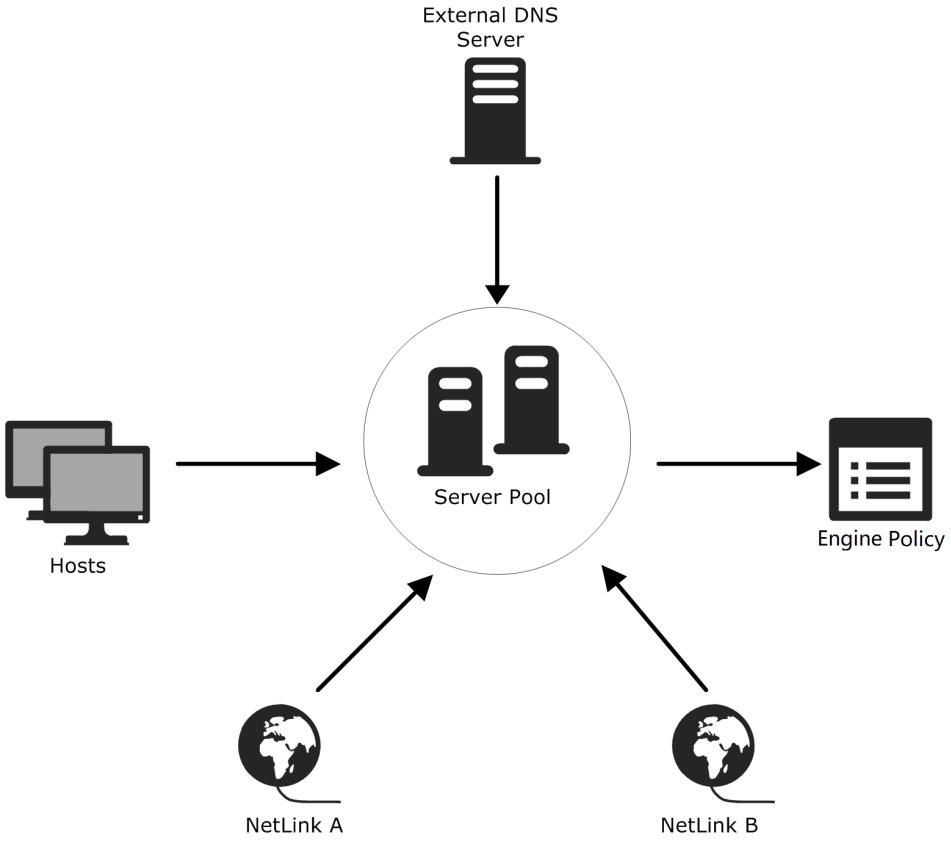
- DDNS updates have no access control, so the communications must be secured in other ways.
- Only IPv4 addresses are supported for DNS servers for DDNS updates.
- Standby servers cannot be defined for a Server Pool. Only load balancing between the servers in the Server Pool is supported.
- Only TCP and UDP protocols are supported for Server Pools.
- Server Pool Monitoring Agents only support IPv4 addresses.

Using Multi-Link with Server Pools in inbound traffic management

If you have configured Multi-Link, it can be used to improve Server Pool availability.

You can also use Multi-Link with just one server in the Server Pool to take advantage of dynamic DNS updates.

Multi-Link configuration for a Server Pool



As an addition to the basic configuration, the NetLinks and (optionally) the External DNS Server are also specified for the Server Pool.

When dynamic DNS updates are not used, Multi-Link is based on assigning an IP address for the Server Pool in each NetLink. The Server Pool's DNS entry on the external DNS server must be configured with an IP address for each NetLink so that clients can access the servers through the different NetLinks. When the connecting client requests the IP address for the Server Pool's DNS name, the DNS server sends the Server Pool's DNS entry with the IP addresses on the different NetLinks. The client connects to one of these addresses and the Security Engine allocates the connection to one of the Server Pool members. If the first Server Pool IP address is unreachable, the client can connect to the Server Pool's next IP address on a different NetLink (depending on the client application).

When dynamic DNS updates are used, the Security Engine updates the DNS entries automatically based on the availability of the NetLinks. When a NetLink becomes unavailable, the Server Pool's IP address for that link is automatically removed from the DNS entry on the external DNS server. When the NetLink becomes available, the IP address is again automatically added to the DNS entry.

Related concepts

Configuring dynamic DNS updates for Server Pools on page 759

Entering Server Pool IP addresses for your DNS Server

When using static DNS entries (recommended), you must make sure that the IP addresses for your Server Pool are properly entered into your DNS server's records.

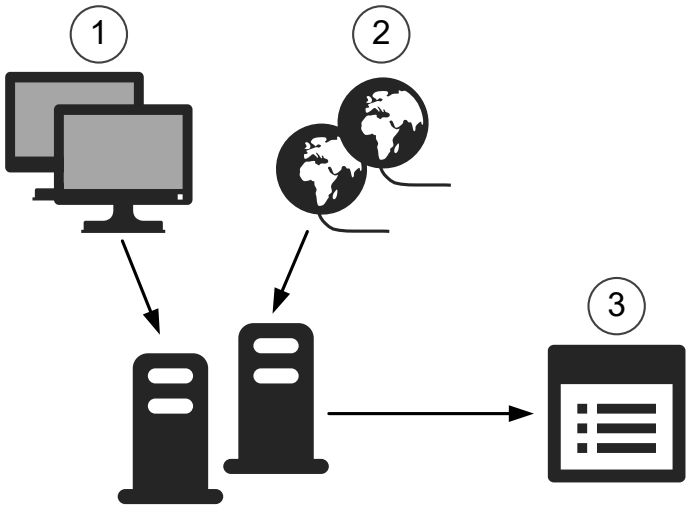
The Server Pool has one IP address for each NetLink (Internet connection) in a Multi-Link configuration and a single IP address in a single-link configuration. If you decide to use dynamic DNS updates, you must configure the DNS server to accept the DDNS updates from the engine. If the servers have multiple IP addresses, make sure the IP addresses are all defined on the DNS server to ensure operation as intended.

See the documentation of your DNS server software for instructions on how to enter the IP addresses.

Inbound traffic management configuration overview

The process of configuring inbound traffic management consists of the following overall steps.

Elements in the configuration



Server Pool

1. Host elements represent your servers in the SMC. Add one or more Host elements as Server Pool members to the Server Pool element.
2. NetLinks represent the network connections through which the external IP addresses of the Server Pool are reached. Select one or more NetLink elements as a property of the Server Pool element.

   You can also use the default **Not Specified** NetLink element to load-balance servers without using Multi-Link for inbound traffic management.
3. To route incoming traffic to the Server Pool members, use the Server Pool element in an a NAT rule or Access rule in the Engine Policy.

The configuration consists of the following general steps:

1) Create Host elements for all servers you want to include in the Server Pool.

2) Define a Server Pool element and define settings for inbound traffic management.
It is recommended that each Server Pool offers one type of service. If the servers provide several services, such as HTTP and HTTPS, create a separate Server Pool for each service.

3) Make sure that the servers' public, external IP addresses in the Server Pool properties correspond to their IP addresses in the DNS server.

4) Add Access rules to allow the type of traffic that is handled by the Server Pool

5) Add NAT rules or Access rules to specify which traffic uses the Server Pool.
NAT rules are the preferred way to enable Server Pool load balancing. For backward compatibility, it is still possible to enable Server Pool load balancing using Access rules.

6) (Optional) If you want to send dynamic DNS (DDNS) updates to a DNS server, configure the updates.

7) (Optional) If you want to use Server Pool Monitoring Agents, install, configure, and enable them.

Create Server Pool elements

The Server Pool element collects servers that provide a particular service into a single element and defines the settings for handling the inbound traffic.

Before you begin

You must have a Host element that represents the internal IP address of each server that you want to add to the Server Pool.



Note

Make sure that other NAT configurations do not overlap with the internal and external IP addresses of the Server Pool.

We recommend creating a separate Server Pool element for each type of service. Add servers to Server Pools based on the services that the servers provide.

The Server Pool can have up to 255 members. If you have only one server and you want to balance the inbound traffic between your NetLinks, you can define a Server Pool element with just one host. This approach allows dynamic DNS update information to be used to prevent contacting clients from attempting to use a NetLink that is out of service.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.

2) Browse to **Traffic Handlers**.

3) Right-click **Traffic Handlers**, then select **New** > **Server Pool**.

4) In the **Name** field, enter a unique name for the Server Pool.

5) Define the external IP addresses of the Server Pool.
   a) In the **External Addresses** section, click **Add**.

   b) Select a NetLink element.

   Tip

   To load-balance servers without using Multi-Link for inbound traffic management, select **Not Specified**.

   c) From the **Network** drop-down list, select the network to which the Server Pool's external IP address belongs.

   d) In the **IP Address** field, enter the external IP address for the Server Pool.

   Note

   The IP address you enter here must be reserved for NAT and it must not be used by any equipment in your network. Remember to update your DNS server with any changes in IP addressing.

   e) From the **Status** drop-down list, select **Enabled**.

6) Add servers to the list of Server Pool members.
   a) In the **Server Pool Members** section, click **Add**.

   b) Add the Host elements that represent the internal IP address of the servers to the list of Server Pool members.

   Tip

   For servers that have some special role in the SMC configuration, you can add the existing Server element.

   c) From the **Allocate Traffic to Servers by** drop-down list, select the granularity for the server selection.

   Consider the type of traffic when selecting the allocation method. Connections from the same source might be directed to different servers in the following cases:

   - Using the **Host** setting if the host's IP address apparent to the SMC can change.
   - Using the **Connection** setting in all cases.

   Depending on the services offered, directing connections from the same source to different servers might reduce the quality of service.

7) Click **OK**.

Configure server availability monitoring

There are different methods for monitoring whether a server or a service running on a server is available.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.
2) Browse to **Traffic Handlers**.
3) Right-click a Server Pool element, then select **Properties**.
4) In the **Frequency Check** field on the **Monitoring** tab, define how often you want the availability to be checked.
5) From the **Method** drop-down list, select the method for monitoring the availability of the servers in the Server Pool.
6) (TCP or HTTP methods) Specify the additional options.

   The text in the **Request** field is sent by the engine in the first packet after the TCP handshake. The text in the **Response** field is checked against the first data packet received from the server after the TCP handshake.

   Tip

   The response is evaluated using a `begins with` pattern. It is not necessary to enter the complete text of the response.

   When using TCP, if the **Request** and **Response** fields are left empty, the Engine only checks that a connection can be established.

   For example, to check for the HTTPS service, enter 443 as the port number, but leave the **Request** and **Response** fields empty. In this case, the engine only checks that a connection to port 443 can be established.
7) Click **OK**.

Related concepts

Install Server Pool Monitoring Agents on page 765

Related tasks

Enable Server Pool Monitoring Agents on page 778

Create Access rules to allow the type of traffic that is handled by the Server Pool

Before you can enable Server Pool load balancing using NAT rules, you must create Access rules to allow the type of traffic that is handled by the Server Pool.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Open the Engine Policy for editing and add an IPv4 or IPv6 Access rule.

Note

If the Server Pool uses both IPv4 and IPv6 addresses, you must create separate IPv4 and IPv6 Access rules.

2) Configure the **Source**, **Destination**, and **Service** of the rule to allow the type of traffic that is handled by the Server Pool.

If you want to restrict the destinations to which this type of traffic is allowed, we recommend creating Host elements that represent the external IP addresses of the Server Pool and using them in the Destination cell of the Access rule.

Note

Do not add the Server Pool element to the Destination cell of this rule. If the Destination cell of an Access rule contains a Server Pool element, the Access rule applies Server Pool load balancing, and the NAT rules are ignored.

3) Save and Install the Engine Policy to transfer the changes.

Next steps

Create NAT rules for inbound load balancing.

Related tasks

Add rules to allow dynamic DNS updates on page 762

Enable Server Pool load balancing using NAT rules

NAT rules specify which traffic is directed to the Server Pool. You can use NAT rules to apply both source and destination address translation for Server Pools.

Before you begin

Add Access rules that allow the type of traffic that is handled by the Server Pool.

If there are any existing Access rules that enable Server Pool load balancing, we recommend that you remove the Server Pool elements from the rules or delete the rules before you configure NAT rules for Server Pool load balancing.

Note

If the Destination cell of an Access rule contains a Server Pool element, the Access rule applies Server Pool load balancing, and the NAT rules are ignored.

When you use destination address translation with Server Pools, the NAT operation translates the external IP addresses of Server Pool elements to the internal IP addresses of the Host elements that are members of the Server Pool.

When you use source address translation with Server Pools, the return packets from the Server Pool servers are routed through the Security Engine to the client. These packets are recognized as part of the existing connection between the client and the server. This feature also allows you to use dynamic source NAT with Server Pool load balancing.

Note the following:

- Make sure that there are no overlapping NAT rules in the policy.
- If you want to balance traffic that arrives through a VPN using a Server Pool, NAT must be enabled in the properties of the VPN element (NAT is disabled by default for traffic that uses a VPN).
- You must create a separate NAT rule for each Server Pool.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Open the Engine Policy for editing and add an IPv4 or IPv6 NAT rule.

   Note

   If the Server Pool uses both IPv4 and IPv6 addresses, you must create separate IPv4 and IPv6 NAT rules.

2) Add the elements that represent the IP addresses of the clients that connect to the Server Pool to the **Source** cell.

3) Add the Server Pool element to the **Destination** cell.

   The Server Pool element must be the only element in the Destination cell. Destination address translation is automatically configured when you add a Server Pool element to the Destination cell. You cannot change the destination translation options.

4) Add the Service element that represents the service that the Server Pool offers to the **Service** cell.

   Note

   Each rule must contain only one service.

5) (Optional) Double-click the **NAT** cell, then define options for source address translation.

   Defining source address translation routes return packets from servers in the Server Pool to the clients through the Security Engine.

6) Save and Install the Engine Policy to transfer the changes.

Next steps

If you want the Security Engine to automatically update dynamic DNS (DDNS) entries for the Server Pool according to the available NetLinks, configure DDNS updates.

Related concepts

Configuring dynamic DNS updates for Server Pools on page 759

Enable Server Pool load balancing using Access rules

NAT rules are the preferred way to enable Server Pool load balancing. For backward compatibility, it is still possible to enable Server Pool load balancing using Access rules.

When you use a Server Pool element in the Destination cell of an Access rule, the rule enables Server Pool load balancing and specifies which traffic is directed to the Server Pool. When the rule matches traffic, the Server Pool uses NAT to change the destination IP address to the IP address of the server that the engine selects for the connection. Reverse NAT (for the replies the server sends back to the client) is handled automatically. No separate NAT rule is required.

If you use Access rules to enable Server Pool load balancing, note the following:

- The Server Pool does automatic NAT from the external addresses you configured in the Server Pool element to the addresses of the included servers. Make sure that there are no overlapping NAT rules in the policy. You can add a NAT rule that disables further NAT for matching connections (empty NAT cell), if necessary.
- If you want to balance traffic that arrives through a VPN using a Server Pool, NAT must be enabled in the properties of the VPN element (NAT is disabled by default for traffic that uses a VPN).
- You must create a separate rule for each Server Pool.
- If the same Server Pool provides more than one service, you must create a separate rule for each Service.
- You must enable Connection Tracking for the rule that directs traffic to the Server Pool. The Server Pool uses NAT, which does not work without Connection Tracking.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Open the Engine Policy for editing and add an IPv4 or IPv6 Access rule.

> **Note**
>
> If the Server Pool uses both IPv4 and IPv6 addresses, you must create separate IPv4 and IPv6 Access rules.

2) Configure the rule to match the **Source**, **Destination**, and **Service** of the traffic that you want to direct to the Server Pool.

> **Note**
>
> Each rule must contain only one **Service**.

3) Set the **Action** to **Allow**.

4) In the **Action Options**, enable **Connection Tracking**.
The following example rules direct traffic from external networks to the HTTP Server Pool and to the HTTPS Server Pool.

Source	Destination	Service	Action
Not Internal network Expression	HTTP Server Pool	HTTP	Allow Connection tracking: normal
Not Internal network Expression	HTTPS Server Pool	HTTPS	Allow Connection tracking: normal

5) If you are using static DNS entries, save and Install the Engine Policy to transfer the changes.

Next steps

If you want the Security Engine to automatically update dynamic DNS (DDNS) entries for the Server Pool according to the available NetLinks, configure DDNS updates.

Configuring dynamic DNS updates for Server Pools

The Security Engine can automatically update dynamic DNS (DDNS) entries for the Server Pool according to the available NetLinks.

The Security Engine removes the Server Pool IP addresses for NetLinks that are not available from the DNS entry, and adds the IP addresses back when the NetLink becomes available again. When the connecting client requests the Server Pool's IP address from the DNS server, the client receives a list of IP addresses that only contains IP addresses that work.

Security Engines support the Dynamic DNS protocol and can send DDNS updates to a specified DNS server. If a network connection specified by a NetLink element fails, the dynamic DNS updates notify the DNS, which then removes the corresponding IP address from its records.

To configure DDNS updates, you must have already defined the necessary NetLinks and the Server Pool element. To use DDNS updates, you must set up a DDNS-capable DNS server in your network. The DNS server must be configured as the primary DNS server for the domain. Only IPv4 addresses are supported for DNS servers for DDNS updates.

CAUTION

Although Security Engines support dynamic DNS updates, the protocol itself poses a security risk because there is no access control. If you must use dynamic DNS updates, do so only after careful research, planning, and testing.

There are actions you can take to improve the security of dynamic DNS updates:

- Always place the DNS servers behind the Security Engine for protection from IP address spoofing.
- Use BIND or an equivalent DNS server that allows you to define which hosts are allowed to send dynamic updates.
- Consider using static DNS entries instead, as DDNS is not necessarily needed with inbound load balancing. In that case, the DNS entries are not removed automatically from the DNS server if an ISP fails, but you can sometimes solve these problems by other means. For example, some web browsers can automatically try other IP addresses if one address does not respond.

The configuration consists of the following general steps:

1) Familiarize yourself with the security risks of implementing DDNS updates before proceeding with the configuration.
2) To enable monitoring of the status of your NetLinks, add probe IP addresses in the NetLinks' properties.
3) Define an External DNS Server element.
4) Edit the Server Pool element to include information about how the DNS server is updated.
5) Edit the Engine policy to allow DDNS updates.

Define External DNS Server elements

There are some cases in which you must define an External DNS Server element.

- (Engines only) For dynamic DNS (DDNS) updates with a Multi-Link configuration.
- (Engines only) If you want to use a DNS server for resolving malware signature mirrors.
- If you want to use a DNS server for resolving domain names and URL filtering categorization services on Engines, IPS engines, and Layer 2 Engines.

You can also optionally use External DNS Server elements to specify the DNS servers to which the engine forwards DNS requests when you configure DNS relay.

If the device has additional IP addresses, you can enter them as secondary IP addresses instead of creating additional External DNS Server elements. However, secondary IP addresses are only used in the **Source** and **Destination** cells in rules. They are ignored otherwise.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Network Elements**.
2) Right-click **Servers**, then select **New** > **External DNS Server**.
3) In the **Name** field, enter a unique name.
4) In the **IP address** field, enter the IP address of the server.



Note

Only IPv4 addresses are supported for DNS servers for DDNS updates.

5) Configure the other settings.
6) Click **OK**.

Define the dynamic DNS update information

You must define the settings for the Dynamic DNS updates sent from your Engine to the DNS server.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.
2) Browse to **Traffic Handlers**.
3) Right-click the Server Pool, then select **Properties**.
4) Select **Enable Dynamic DNS Updates**.
5) From the **DNS Server** drop-down list, select the External DNS Server element to which the DDNS updates are sent.
6) In the **FQDN** field, enter the fully qualified domain name for the Server Pool service.
Example: www.example.com
7) Click **OK**.

Add rules to allow dynamic DNS updates

You might need to add a rule that allows the dynamic DNS updates from your engine to the specified DNS server.

The Firewall Template allows this traffic by default. If your policy is based on an unmodified Firewall Template, there is no need to add a rule that allows Dynamic DNS updates.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click the Engine Policy, then select **Edit**.

2) Add the following type of Access rule:
   - **Source**: the address (toward the DNS server) of the Engine or NDIs of all nodes in the Engine Cluster
   - **Destination**: the DNS server
   - **Service**: the DNS Service Group
   - **Action**: **Allow**

3) Save and install the policy.

Related tasks

Create Access rules to allow the type of traffic that is handled by the Server Pool on page 756

Using Server Pool Monitoring Agents

Server Pool Monitoring Agents provide advanced features for monitoring the server load and status.

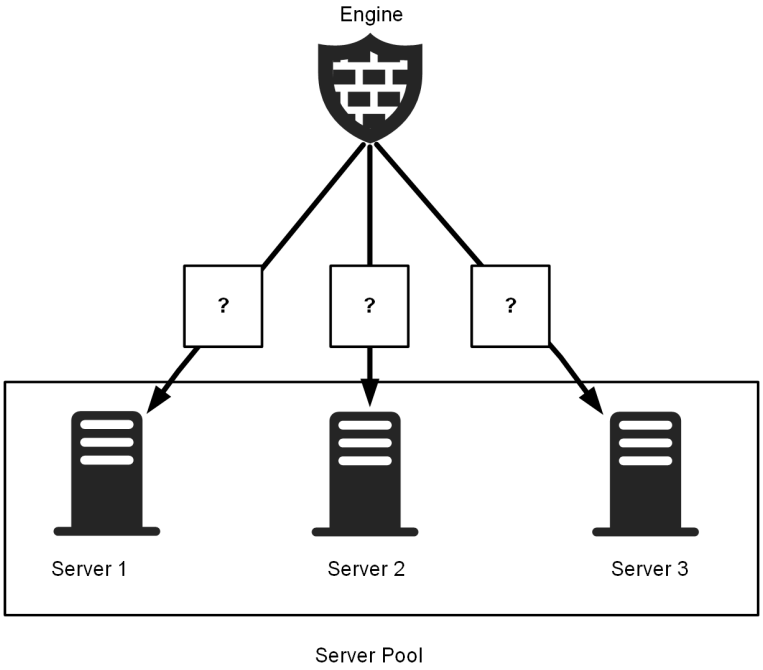


Note

Server Pool Monitoring Agents only support IPv4 addresses.

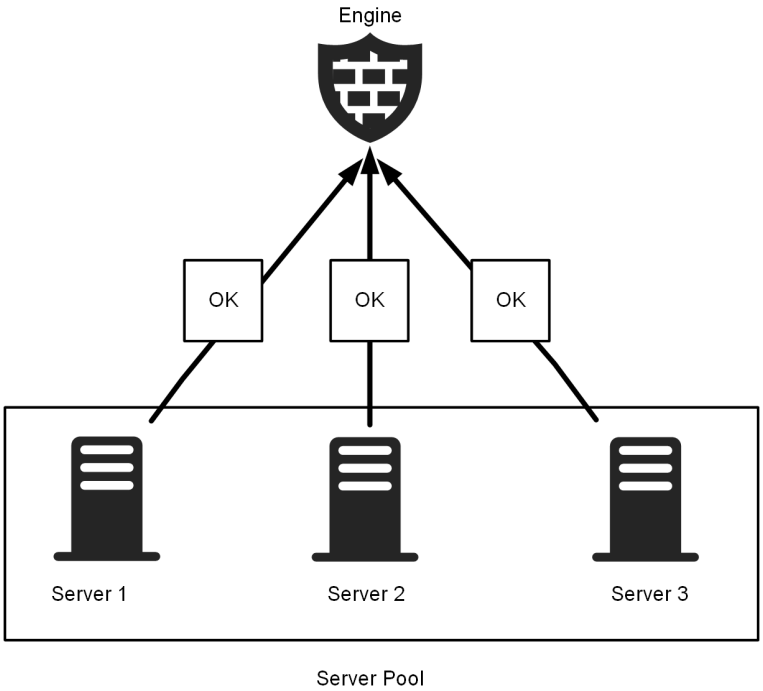
Monitoring Agents run on each server to check the server status and load. The Monitoring Agents can also run system checks on the servers and send log messages to the SMC.

Security Engine queries Server Pool Monitoring Agents on each server



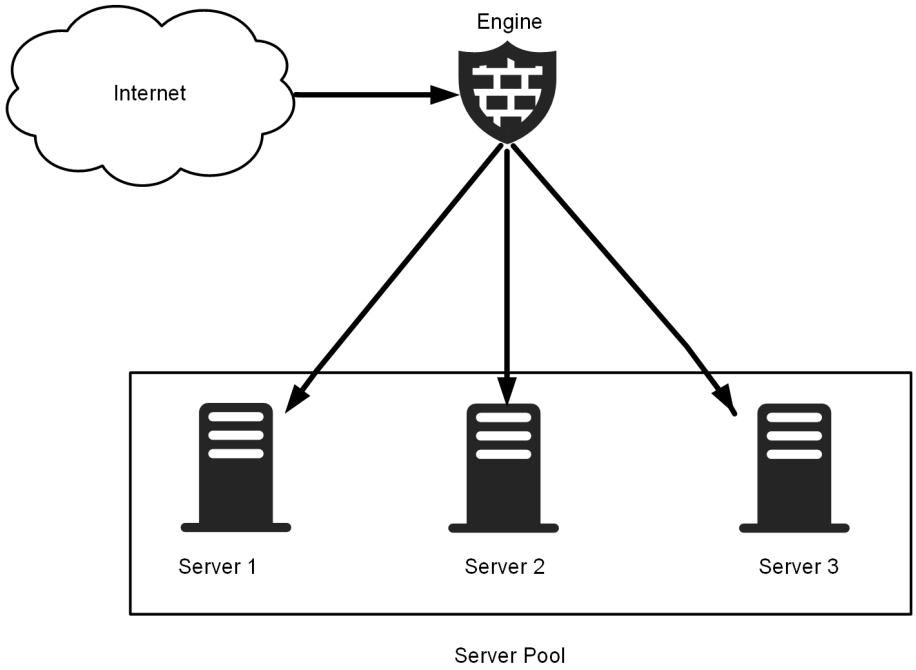
A Monitoring Agent runs as a service (port 7777/UDP by default) on the Server Pool member. The Security Engine queries the Monitoring Agent on each Server Pool member to check the status and load of the server.

Server Pool Monitoring Agents provide status information



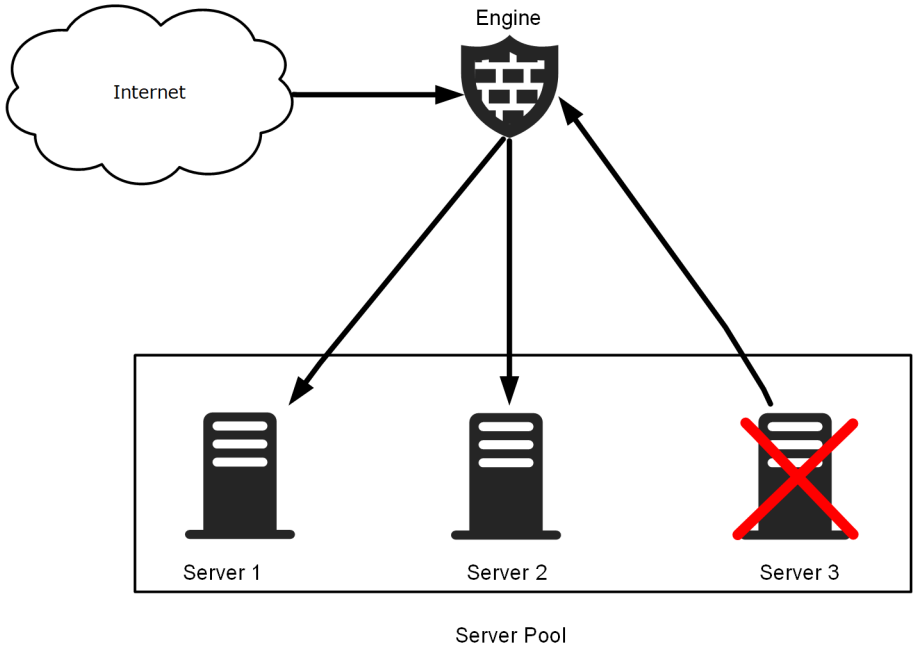
The Monitoring Agent on each Server Pool member provides information about the server load and status to the Security Engine.

Security Engine balances connections between Server Pool members



The Security Engine balances the incoming connections between the Server Pool members according to the status and load information it receives from the Monitoring Agents.

Server Pool Monitoring Agents: Test failure



The Server Pool Monitoring Agent also has a tester that is configured to run predefined tests or user-defined programs. Automatic action can be configured based on the results of the test. When a test fails, an alert is sent to the Security Engine. Optionally, the agent can also take the server out of the Server Pool by changing its status

from “OK” to “Excluded”. When a server is excluded from the Server Pool, all new connections are directed to other available servers in the pool. The excluded server does not process any connections.

Server Pool Monitoring Agents are available for Windows and Linux platforms.

Install Server Pool Monitoring Agents

Server Pool Monitoring Agents can be installed on the Server Pool servers to monitor the availability and load of the Server Pool members.

Server Pool Monitoring Agents are installed on the servers that form the Server Pool to report the server’s health and load. The Engine uses the information to send traffic to the server that can best handle the additional load and avoid sending new connections to a server that is inoperative. The Monitoring Agent has various in-built tests and it can also run external scripts or programs on the server to check that the server is functioning properly. You must install the Server Pool Monitoring Agents onto all members of the Server Pool.

Install Server Pool Monitoring Agents in Windows

Install Server Pool Monitoring Agents on all members of a Windows-based server pool. Perform this task on all Server Pool members.

Steps

1) Open Forcepoint Customer Hub.
2) Click the **Downloads** tab.
3) Navigate to **Next Generation Firewall (NGFW)** > **Add-on Features**.
4) Download the SPMA Windows installer to a suitable location on the server.
5) Run the EXE self-extracting installation program.

Install Server Pool Monitoring Agents in Linux

Install Server Pool Monitoring Agents on all members of a Linux-based Server Pool. Perform this task on all Server Pool members.

Steps

1) Open Forcepoint Customer Hub.
2) Click the **Downloads** tab.
3) Navigate to **Next Generation Firewall (NGFW)** > **Add-on Features**.
4) Download the appropriate SPMA Debian or SPMA RPM package to a suitable location on the server.

5) Run the `rpm` command to install the package.

```
For example: rpm -i sgagent-5.0.0.build503-1.x86_64.rpm
```

Uninstall Server Pool Monitoring Agents

You can uninstall the Server Pool Monitoring Agents you no longer need.

Steps

1) On the server, uninstall the Monitoring Agent according to the operating system of the server.
   - In Windows, in the Windows Control Panel, click **Add/Remove Programs** and select **Monitoring Agent**. Follow the instructions to remove the Monitoring Agent.
   - In Linux, run the following command: `rpm -e sgagent`.

Configuring Server Pool Monitoring Agents

Server Pool Monitoring Agents are configured in two files. The sgagent.local.conf file is specific for each server. The sgagent.conf file applies to all servers in the pool.

Both files are found in the following directory on each server:

- Windows: Program Files\Stonesoft\StoneGate Monitoring Agent\
- Linux: /etc/stonegate/

The sgagent.local.conf file contains server-specific information, and is different for each server. If you do not want to change the default values, there is no need to edit this file.

The sgagent.conf file contains information that refers to the entire Server Pool.

Editing sgagent.local.conf

You can edit the sgagent.local.conf file for each Server Pool member or use the default values.

The server-specific file sgagent.local.conf includes information to identify the Server Pool member if default values are not used. If the file is not configured, the Monitoring Agent uses the system host name as the default value to identify the member. If edited, this file must be different on each Server Pool member.

Statement	Description
`host host_name`	Specifies the `host_name` of the server. If no value is defined, the system host name is used.
`set alias value`	Defines the `value` of an `alias` that can be used as a variable in the global sgagent.conf configuration.

For example, to configure the member of the Server Pool where sgagent.local.conf resides to use server1 as its host name, and to set the alias hostip to 192.168.1.1:

```
host server1
set hostip 192.168.1.1
```

Editing sgagent.conf

You define attributes for the entire Server Pool in the sgagent.conf file.

The file common to all Server Pool members is called sgagent.conf. In this file, you define attributes for the entire Server Pool. In most cases, the Server Pool members are similar in configuration and function, so the tests performed on each member are also similar. Usually, the same sgagent.conf file is used on all servers. The sgagent.conf file contains two sections: the statement section and the test section.

An example sgagent.conf file

```
# --------------------------START OF STATEMENT SECTION
begin host server1
config boot-delay 5:00 # set the config delay on the slow member
end
begin host not server1
config boot-delay 20 # 20s config delay on other members
end
config alert-interval 1:00:00
config startup-script /etc/scripts/startup.sh
# Monitor one minute load average and set alert threshold at 5
load-index 500 load-average-1
# --------------------------START OF TEST SECTION
# Tests
test "webster listening"
interval 2:00
action exclude
recovery 10:00
command portlistening 80 %hostip%
# Check that this member's
# port 80 is listening.
test "server1 running baz"
interval 10:00
action exclude
host server1
script /etc/init.d/baz start
command servicerunning baz
test "external1"
interval 10:00
action alert
command external 3 700 /usr/lib/webster/test.sh 203.0.113.21
```

Statements in the sgagent.conf statement section

You can define general commands for the tests the Monitoring Agents perform in the statement section.

A statement can apply to all or only some members of the Server Pool.

Statements in the sgagent.conf statement section

Statement	Description

`begin host host_name`	If you want a statement to apply to certain Server Pool members only, the statement must start with the begin host command followed by the name of the Server Pool member.
	The `begin host` command can be followed by actual server names or an exclusion. In an exclusion, `begin host` is followed by `not` and the server name, instead of just a server name. This exclusion indicates that the setting is applied to all other servers except the one specified.
	For example, the following statement would apply to all members in the Server Pool except server1:
	`begin host not server1`
	`config settings`
	`end`

`config`
`[port port_number]`
`[boot-delay time]`
`[alert-interval time]`
`[alert-script path]`
`[startup-script path]`
`[overload-script path]`
`[listen IP_address [:port]`
`[load-index threshold`
`index_name`
`[[index-parameter]]`
`[load-index-action exclude time]`
	Defines the option in question.
	If a statement applies to all members in the Server Pool, you can enter the setting on a single line without using begin or end.

`end`	Ends the statement when `begin` is used.

Options in the sgagent.conf statement section

When you want to configure a particular setting, the syntax starts with the command `config`.

The available options are listed in the following table:

Options in the sgagent.conf statement section

Option	Description
`port port_number`	Specifies the port that the Server Pool uses for communication with the engines. All Server Pool members must use the same port number; the port setting cannot start with the `begin host` command. The port number can be from 0–65535. By default, the port is set to 7777. If you use another port instead of the default port, you must edit the IPv4 and IPv6 Access rules to allow connections to the new port.

Option	Description
`boot-delay time`	Defines how long the tester waits before starting the tests after the system has been booted up or restarted to ensure that all necessary processes have completed their startup. Time is entered in the format [(hh:)mm:]ss. The default is 30 seconds.
`alert-interval time`	Defines how long the system waits before sending a new alert in response to a test failure to prevent excessive alerts from being sent. Time is entered in the format [(hh:)mm:]ss. The default is 1 hour.
`alert-script path`	Specifies the path to a custom alert script. You must use quotation marks if the path contains spaces. This is only needed when you want to use a custom alert script. There is no default value.
`startup-script path`	Specifies the path to the custom script to be run when the Monitoring Agent starts. You must use quotation marks if the path contains spaces.
`overload-script path`	Specifies the path to the custom script to be run if the load index exceeds the threshold value defined with the `load-index` statement. You must use quotation marks if the path contains spaces.
`listen IP_address [:port]`	Specifies the IP address the Monitoring Agent listens on. If no IP address is specified, the first IP address of the interface is used by default. If no port is defined, the default port (7777) is used.
`load-index threshold index_name [index-parameter]`	Defines the method that is used to measure the load. We recommend using the same method of measurement for all members of the Server Pool. If the load measurement exceeds the defined threshold, a log message and an alert are generated. The `index_name` can be a complex expression combining multiple methods with basic arithmetic operations (+ - * /), such as "index-name + index-name".
`load-index-action exclude [time]`	Excludes the server from handling traffic for the specified time (in seconds) when the threshold value of the load-index is reached. When the server has switched to the Excluded state, the server does not return to the normal state until the specified length of time has elapsed, even if the load of the server has dropped to a normal level.

Index_name options

Option	Description
`processor-load time`	Measures the average processor load level over the specified time period. The time is entered in the format [(hh:)mm:]ss. The returned value is a percentage value from 0–100.
`processor-kernel-load time`	Measures the average processor time spent in the kernel mode over the specified time period. The time is entered in the format [(hh:)mm:]ss. The returned value is a percentage value from 0–100.
`load-average- minutes`	Measures the average number of processes waiting in the execution queue due to reserved resources over the specified number of minutes. The options for the value of minutes are 1, 5, or 15. The returned value is the load average for the specified time period multiplied by 100.

Monitoring Agent statement configuration examples

When you add options for Monitoring Agents to the statement section of the sgagent.conf file, you can use the following examples as a guide.

Port option example

To use port 5555 for the communication between the engines and the Server Pool members instead of the default port:

```
config port 5555
```

Boot-delay option example

To set a boot delay of 20 seconds on server1:

```
begin host server1
config boot-delay 20
end
```

Alert-interval option example

To set an alert interval of 1 hour on servers 1 and 2:

```
begin host server1 server2
config alert-interval 1:00:00
end
```

Alert-script option example

To run `alertscript.sh` every time an alert is generated:

```
config alert-script /etc/test/alertscript.sh
```

Startup-script option example

To run `startup.sh` when the Monitoring Agent starts on server1:

```
begin host server1
config startup-script /etc/test/startup.sh
end
```

Overload-script option example

To run `overload.sh` script when the threshold value of the load index is reached on server2:

```
begin host server2
config overload-script /etc/test/overload.sh
end
```

Listen option example

To make the Monitoring Agent of server2 listen to the IP address 192.168.10.31 using the default port 7777:

```
begin host server2
config listen 192.168.10.31
end
```

Load-index option example

This example demonstrates how the measurement of `load-average-5` is used on all servers except server1 to compare against the `threshold` value 4. When the load-average is higher than 4, a log message is generated:

```
begin host not server1
config load-index 400 load-average-5
end
```

Because server1 is ignored in the previous statement, it would typically have its own statement with different options, such as the following statement:

```
begin host server1
config load-index 400 load-average-1
end
```

Load-index-action option example

To set the server status to Excluded when the threshold value of the load index is reached on server2:

```
begin host server2
config load-index-action exclude 10
end
```

Parameters in the sgagent.conf test section

You can use the parameters described in the following table to configure internal and external tests for Monitoring Agents.

The second part of the sgagent.conf file specifies which tests the Monitoring Agents carry out on the servers they monitor. A test begins with the test definition, followed by optional parameter lines, and ends with the command definition. The `test` and `command` parameters are always required. The `interval`, `action`, `host`, `script`, and `recovery` parameters are optional and have their default values if not specified for the test.

Parameters in the sgagent.conf test section

Parameter	Description
`test name`	Begins the test definition. Specifies the name of the test. The name appears in the logs and alerts in the event of a failure. It should be descriptive to give the administrators enough information about what has happened. The test name can be any string. You must use quotation marks if the name contains spaces. This parameter is required and has no default value.
`interval time`	Defines how often the test is executed. Time is entered in the format [(hh:)mm:]ss. The minimum interval is 3 seconds and the maximum is 1 day. If no interval is defined, the default is 60 seconds.
`action alert|exclude`	Defines what to do if the test fails. ■ `alert` creates an alert entry and runs the alert script if one is defined for the test. ■ `exclude` sets the server status to Excluded and creates an alert. No new connections are directed to the excluded server and the current connections are moved to the other members of the Server Pool. ■ If no action is set for the test, `alert` is the default action.
`host server`	Defines the server to be tested. If the name of the server is specified in the host field of sgagent.local.conf, that name must be used here. Otherwise, the default host name of the server must be used. You can also list multiple server names, separated by spaces. If no specific server is defined, the test is performed on all Server Pool members by default.
`script [always|repeat] path`	Specifies how many times the script is executed after the test fails, and the path to the script file. You must use quotation marks if the path contains spaces. There are two options for how many times to execute the script after the test fails: ■ repeat defines the number of failures after which the script is no longer run. The counter is reset each time the server status changes, or the test succeeds. ■ `always` runs the script without restrictions on the number of failures. If no argument is given, the script is run only after the first failure by default.
`recovery always|time`	Determines how long the tester reruns a failed test to see if it succeeds on the next attempt. There are two options for specifying the time period: ■ `always` : There are no time limits on how long the tester waits for the test to succeed again. The server state is changed from Excluded to OK when the test succeeds. ■ time : If the test succeeds in the defined time period after a failure, the server state is changed from Excluded to OK. Time is entered in the format [(hh:)mm:]ss. If no recovery options are specified, the test is never rerun after a failure, and the server state remains Excluded.

Parameter	Description
`command external| test_expression [ retry ] [ timeout ] [ path ] [ parameters ]`	Specifies which test script to use, and the options for running the test. This parameter is required and has no default value. There are two options for specifying the test script: ■ `external` : Indicates that an external test is used. External tests can be created by combining predefined tests with AND, OR, and NOT operators, or you can write your own script. External tests must return an exit code of 0 (zero) to indicate success. Any non-zero return value is interpreted as a failure. ■ test_expression : Specifies the internal test to run. For more information about the internal tests, see *Internal tests for Monitoring Agents*. The other options are used as follows: ■ retry : Specifies how many times the script is run before the test is considered failed. ■ timeout : Specifies how long (in milliseconds) the tester waits for the test result. ■ path : Specifies the path to the external test script. You must use quotation marks if the path contains spaces. The path is required for external tests, and is not used for internal tests. ■ parameters : Specifies any possible parameters that a specific test might need.

Related concepts
Editing sgagent.local.conf on page 766

Related reference
Internal tests for Monitoring Agents on page 774

Monitoring Agent test configuration examples

When adding internal or external tests for Monitoring Agents to the sgagent.conf file, you can use the following examples as a guide.

Internal Test Example

In this example, the test checks that server1 is listening to traffic on port 80. In this example, the port information and the IP address of the server are required in the `command` definition. The test is performed every two minutes, and if there is a failure, the server is excluded from the Server Pool. If the test fails, the `/etc/init.d/port start` script is run. If the subsequent tests succeed within 10 minutes of the failure, recovery occurs and the server returns to the OK state:

```
test “port listening test”
interval 2:00
action exclude
host server1
script /etc/init.d/port start
recovery 10:00
command portlistening 80 192.168.1.1
```

External Test Example

In this example, the external test script test.sh is run. There are three attempts to run the test before it is considered failed, and the tester waits for one second before the test is timed out. The test is run every 10 minutes, and if there is a failure an alert is generated:

```
test “multiple_services”
interval 10:00
action alert
command external 3 1000 /usr/lib/server/test.sh 192.168.1.1
```

Internal tests for Monitoring Agents

You can add the internal tests for Monitoring Agents described in the following table to the test section of the sgagent.conf file.

All internal tests available for the Server Pool monitoring are listed in the following table.



Note

Some of the internal tests might not work on virtualization platforms.

Internal tests

Test	Description
`swapfree limit`	Checks the available swap space. If the current amount of free swap space drops below the limit (given in kilobytes), the test fails.
`processcount count`	Checks how many processes are currently running. If the number of processes exceeds the count, the test fails.
`multi-ping retry timeout ip_address [ip_address [...]]`	Sends ICMP Echo Request messages to the defined IP addresses and waits for the replies. The test fails if none of the IP addresses answers the ping query. ■ retry specifies how many ICMP Echo Request packets are sent to each IP address. ■ timeout defines how long (in milliseconds) the tester waits for the reply.
`filesystem partition free_space`	Checks whether there is enough space left on a particular partition. The test fails if the amount of free_space (in kilobytes) is lower than the free space on the partition. You can also identify the partition by its path. File system tests are quite taxing on the operating system. We recommend that you set an interval of at least 15 minutes when using this test.
`servicerunning service_name`	Checks whether the specified service is running on the server. The test fails if the specified service or process is not running. On Windows platforms, the test checks whether the service called service_name is running. In UNIX environments, the test checks whether the process named service_name is running.
`ip-listening ip_address`	Checks whether the specified ip_address is listening on the server. The test fails if the host does not have the specified IP address.
`portlistening port [ip_address [protocol]]`	Checks whether the specified port is listening on the server. The test fails if the port is not in the listen state. You can optionally specify the ip_address at which to check for the port. The protocol can be TCP or UDP.

Test	Description
`portanswer retry timeout query response port [ip_address]`	Checks whether the TCP port answers a TCP query with the expected response. The test fails if the port does not answer, or the answer differs from the expected response. ■ retry specifies how many attempts are made before the test is considered failed. ■ timeout specifies how long the tester waits after sending a query to the port. The timeout is entered in milliseconds. If no response is received within the time period, the test fails. ■ query specifies the TCP query to send. You must use quotation marks if the query contains spaces. If no query parameter is specified, nothing is sent to the port. ■ response specifies the expected response to the TCP query. You must use quotation marks if the response contains spaces. ■ port specifies the port to which the query is sent. ■ ip_address is the IP address to which the query is sent.
`httpanswer retry timeout URL file port [ip_address]`	Checks whether the server answers an HTTP request for the specified URL with the expected file. The test fails if the server does not answer the HTTP request or if the reply does not contain the contents of the specified file. ■ port specifies the port to which the request is sent. ■ retry specifies how many attempts are made before the test is considered failed. ■ timeout specifies how long the tester waits after sending a query to the port. The timeout is entered in milliseconds. If no response is received within the time period, the test fails. ■ ip_address is the IP address to which the request is sent. If no address is specified, the tester sends the query to the localhost address by default.
`networkinterface-up interface`	Checks that the specified interface is up. The test fails if the specified network interface does not exist or is down. In Windows, this test behaves similarly to the networkinterface-linkstatus test described next.
`networkinterface-linkstatus interface`	Checks the link-layer status of the specified interface. The test fails if the specified network interface is not linked.
`file-exists path`	Checks whether a file exists on the server. The test fails if the file is not found at the specified path.

Monitoring Agent internal test examples

When you add internal tests for Monitoring Agents to the sgagent.conf file, you can use the following examples as a guide.

Each test has an example configuration, which might differ from what you configure for your own environment.

swapfree

This example checks that there is at least 1500 kilobytes of free space on server 2. The test runs at one-hour intervals and sends an alert if the test fails.

```
test “swapfree test”
interval 1:00:00
host server2
action alert
command swapfree 1500
```

process_count

This example checks that there are a maximum of 2500 processes running on server2. The test runs at 60-second intervals and sends an alert if the test fails:

```
test process_count
interval 60
host server2
action alert
command prosesscount 2500
```

multi-ping

This example tests the connectivity of server2. The tester sends three ICMP Echo requests to the IP addresses 192.168.1.2 and 192.168.1.254, and waits 1000 milliseconds for a reply. The test runs at 30-second intervals. If the test fails, server2 is excluded from the Server Pool.

```
test “multi-ping test”
interval 30
host server2
action exclude
command multi-ping 3 1000 192.168.1.2 192.168.1.254
```

filesystem

This example checks that there is at least 1000 kilobytes of free space in the /var/tmp partition on server2. The test runs at 15-second intervals and sends an alert if the test fails.

```
test “free space in filesystem”
interval 15
host server2
action alert
command filesystem /var/tmp 1000
```

servicerunning

This example checks that the ntpd service is running. The test runs at 15-second intervals. If the test fails, the ntpd start script is run. If the test succeeds again within one minute of failing, recovery occurs and the server is returned to the OK state. Otherwise, the server is excluded from the Server Pool.

```
test “is ntp running -test”
interval 15
script /etc/init.d/ntpd start
action exclude
recovery 1:00
command servicerunning ntpd
```

ip-listening

This example checks that server2 is listening at the IP address 192.168.1.1. The test runs at 30-second intervals, and the server is excluded from the Server Pool if the test fails.

```
test ip-listening
interval 30
host server2
action exclude
command iplistening 192.168.1.1
```

port-listening

This example checks that server2 and server3 are listening on port 80. The test runs at 30-second intervals, and the server is excluded from the Server Pool if the test fails.

```
test port-listening
interval 30
host server2 server3
action exclude
command portlistening 80
```

httpanswer

This example checks that the server sends the contents of the file /web/file.html as a part of its response to a URL request for http://www.example.com/index.html on port 80. The request is attempted four times, and the tester waits 3500 milliseconds for a response before retrying. The test runs at 30-second intervals, and the server is excluded from the Server Pool if the test fails.

```
test http_answering
interval 30
action exclude
command httpanswer 4 3500 /www.example.com/index.html /web/file.html 80
```

networkinterface-up

This example checks that the eth0 interface is up. The test runs at 45-second intervals, and an alert is sent if the test fails.

```
test etc0_up
interval 45
action alert
command networkinterface-up eth0
```

networkinterface-linkstatus

This example checks that the eth0 interface is linked. The test runs at 45-second intervals, and an alert is sent if the test fails.

```
test “link status on eth0”
interval 45
action alert
command networkinterface-linkstatus eth0
```

file-exists

This example checks that the /important/index.html file exists. The test runs at 2-minute intervals, and an alert is sent if the test fails.

```
test “index -file exists”
interval 2:00
action alert
command file-exists /important/index.html
```

Related reference

Parameters in the sgagent.conf test section on page 771

Enable Server Pool Monitoring Agents

Enable Server Pool Monitoring Agents for one Server Pool element.



Note

Do not enable the Monitoring Agents before you configure them.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.

2) Browse to **Traffic Handlers**.

3) Right-click the Server Pool element for which you want to enable the Monitoring Agents and select **Properties**.

4) Click the **Monitoring** tab, then select **Agent** as the **Method**.

5) In the **Frequency Check** field, define how often the status is checked.

6) (Optional) Enter the **Port** if you want to use a port other than the default port (7777).



Note

Remember to edit the corresponding IPv4 and IPv6 Access rules to allow connections to the new port.

7) Click **OK**.

Test Monitoring Agents on the command line

You can test the Server Pool Monitoring Agents on the command line.

The monitoring tool allows you to query the agents locally or remotely to check their status.

Steps

1) On the command line, browse to the folder where the Monitoring Agent is installed.

2) Run one of the following commands:
   - **Linux**: `sgagentd -d test [file_name]`
   - **Windows**: `agent test [file_name]`

   Where file_name is a configuration you want to test. Giving the file name is optional, but if you omit the file name, the test is run with the active configuration.

Related reference

Server Pool Monitoring Agent commands on page 1454

Monitor Monitoring Agents on the command line

You can monitor the Server Pool Monitoring Agents on the command line.

Steps

1) On the command line, browse to the folder where the Monitoring Agent monitoring tool is installed:
   - In Windows, the program is installed in the folder where the Monitoring Agent itself is installed.
   - On Linux and Solaris, the program is installed in usr/bin.

2) (Optional) Copy the sgmon program to a remote host that you want to use for testing the Monitoring Agents.

3) Run the command `sgmon -h` to see usage instructions, then run your monitoring commands according to the syntax displayed.
   The host argument is mandatory. To get the status locally on the server where the Monitoring Agent is installed, use localhost as the host argument.

Examples of Server Pools

Using Server Pools, you can manage incoming traffic to your web servers.

Example: Load balancing for web servers

To configure load balancing for multiple web servers, you can set up a Server Pool.

Company A has three web servers to handle the large volume of traffic its website receives. The administrators have previously created Host elements to represent their web servers and created NAT rules to assign an external IP address to each web server.

Now the administrators want to distribute the load of the traffic between the servers. The administrators:

1) Create a Server Pool element and add the Host elements to it. Because they are not balancing incoming connections to the Server Pool between multiple Internet connections, the administrators select the Not Specified NetLink.

2) Add the following Access rule to the Engine policy to allow HTTP connections from addresses that are not internal (Not Internal expression) to the external IP addresses of the Server Pool.

Source	Destination	Service	Action
Not Internal network expression	Host elements that represent the external IP addresses of the Server Pool	HTTP	Allow



Note

The administrators do not use the Server Pool element in the Access rule because they want to use NAT rules to enable Server Pool load balancing. If the Destination cell of an Access rule contains a Server Pool element, the Access rule applies Server Pool load balancing, and the NAT rules are ignored.

3) Delete the existing NAT rules that translate the IP address of each server so that NAT rules for Server Pool load balancing do not conflict with them.

4) Add the following NAT rule to the Engine policy to enable Server Pool load balancing and specify which traffic is directed to the Server Pool:

Source	Destination	Service	NAT
Not Internal network expression	Server Pool element	HTTP	Destination: Server Pool External Addresses to Server Pool Members



Note

Destination translation is automatically configured when the administrators add a Server Pool element to the Destination cell.

5) Define options for source address translation so that return packets from servers in the Server Pool to the clients are routed through the Security Engine.

6) Save and Install the Engine Policy to transfer the changes.

Example: Setting up Multi-Link and dynamic DNS updates

When you set up dynamic DNS updates, the Server Pool NetLink addresses that correspond to the available Internet connections are updated automatically on the DNS server.

The administrators at Company A have already configured a Server Pool (see the previous example). Now they want to make sure that web services remain available even when an Internet connection fails. They also want the Server Pool's NetLink addresses to be automatically updated on the DNS server based on the availability of the Internet connections.

The administrators have already configured Multi-Link routing with the necessary NetLink elements to represent each of their Internet connections. A DNS server has also already been set up in the network. The administrators decide to add the NetLinks to the Server Pool and set up Dynamic DNS (DDNS) updates.

The administrators do the following:

1) Configure the DNS server to accept DDNS updates from the engine.

2) Edit the NetLink elements to add probe IP addresses for testing the NetLinks' status.

3) Edit the Server Pool element's properties and replace the default Not Specified NetLink with the NetLink elements that represent their Internet connections.

4) Define an External DNS Server element to represent the DDNS-capable server in the SMC.

5) Enable dynamic DNS updates and configure the dynamic DNS settings in the Server Pool element's properties.

Chapter 46

Dynamic link selection

Contents

- Getting started with dynamic link selection on page 783
- Using dynamic link selection with QoS Class elements on page 785
- Create Connection Type elements on page 788
- Create Link Usage Profile elements on page 789
- Select a Link Usage Profile element for an Security Engine on page 790
- Define exceptions to the Link Usage Profile for an Security Engine on page 790
- Select a Link Usage Profile for a policy-based VPN on page 791
- Select a Link Usage Profile for a Route-based Tunnels group on page 791
- Select a Link Usage Profile for a VPN Broker domain on page 792

When you use Multi-Link for outbound traffic management or Multi-Link VPNs, Forcepoint Network Security Platform in the Engine/VPN role can dynamically select the NetLink or VPN link that best matches the quality requirements of traffic.

Getting started with dynamic link selection

Some traffic is affected more easily by changes in the quality of the connection. The Security Engine can dynamically select the ISP link that best matches the quality requirements of traffic.

The best ISP link for one type of traffic might be different from the best link for another type of traffic. Criteria that affect the quality of a connection include:

- *Bandwidth* is the maximum rate of data transfer for the connection. The bandwidth of the connection is more important when the application transfers a large amount of data at once. For example, the transfer of a single large file using FTP requires higher bandwidth that the transfer of several smaller files.
- *Jitter* is a variation in the delay of received packets. Many applications are affected by jitter, but voice over IP (VoIP) is especially sensitive to jitter.
- *Latency* is a delay in packet transmission. Applications for which communication includes many sequential transactions that each require a round trip are the most affected by latency. For example, latency has more of an effect on VoIP applications.
- *Packet loss* means that one or more packets of data fail to reach their destination. Some applications, such as VoIP applications, are better able to tolerate minor packet loss.
- *Stability* means that the connection is reliably available and the other quality metrics do not vary too much. Applications that have real-time traffic or interactive use require higher stability.

Link selection options in the properties of Network Application, Protocol, and QoS Class elements specify how important different quality metrics are for traffic that is associated with the elements. Traffic uses the link that best

matches the link selection options. When links have similar quality, traffic is distributed between the links so that bandwidth is used in proportion to the quality of the connections.

Link Usage Profile elements define the connection types that are used unless a connection with significantly higher quality is available, are used only if necessary, or must not be used for specific types of traffic. You can use Link Usage Profile elements with the following configurations:

- Outbound Multi-Link — When you select a Link Usage Profile element in the properties of an Security Engine, the settings defined in the Link Usage Profile element are applied to outbound traffic that uses outbound Multi-Link.
- Multi-Link VPNs — When you select a Link Usage Profile element in the properties of a policy-based VPN, Route-based Tunnels group, or a VPN broker domain, the settings defined in the Link Usage Profile element are applied to all tunnels in the VPN according to their link types.

Dynamic link selection is supported on Security Engines, Master Engines, and Virtual Engines in the Engine/VPN role.

Dynamic link selection is only supported for layer 3 physical interfaces.

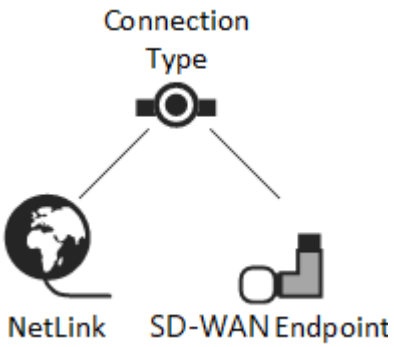
Dynamic link selection configuration overview

The configuration of dynamic link selection consists of several general steps.

1) Define link selection options in the properties of QoS Class elements to specify how important different quality metrics are for traffic that is associated with the QoS Class.

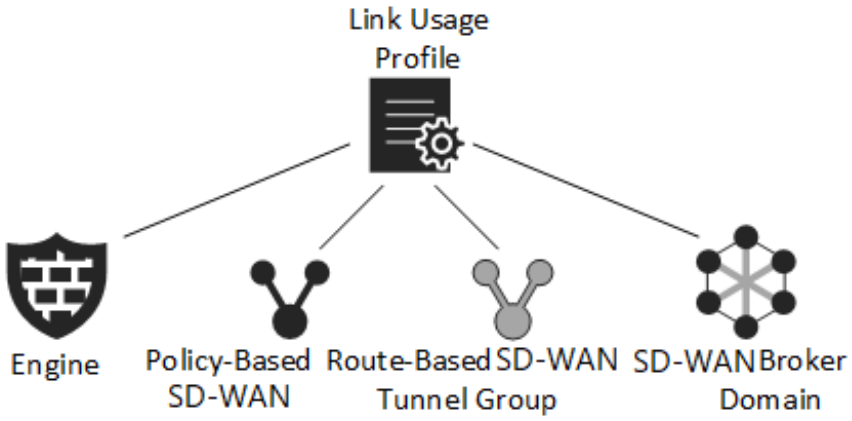
2) (Optional) Create Connection Type elements to define how NetLinks and VPN endpoints are used in a Multi-Link configuration.
   If the default Connection Type elements meet your needs, it is not necessary to create custom Connection Type elements.

   When you create NetLinks or VPN endpoints, select a Connection Type element.



3) Create a Link Usage Profile element to define which link types are preferred, avoided, or not used for specific types of traffic.

4) Activate dynamic link selection.
   - To use dynamic link selection for outbound traffic that uses Multi-Link routes, select a Link Usage Profile for an Security Engine.
   - To use dynamic link selection for Multi-Link VPNs, select a Link Usage Profile for a policy-based VPN, Route-based Tunnels group, or a VPN broker domain.



5) (Optional) To specify which traffic uses specific NetLinks, define exceptions to the Link Usage Profile for an Security Engine.

Using dynamic link selection with QoS Class elements

You can use QoS Class elements to specify link selection criteria for traffic that is associated with the QoS Class or to override the default link selection options in Network Application elements and Protocol elements.

Related concepts

Quality of Service (QoS) and how it works on page 973
Defining Protocol parameters on page 931
Getting started with Network Application elements on page 961

Define link selection options in QoS Class elements

Link selection options in the properties of QoS Class elements specify how important different quality metrics are for traffic that is associated with the QoS Class. Traffic uses the ISP link that best matches the link selection options.

Before you begin

Create one or more custom QoS Class elements.

Note

You cannot define link selection criteria for the default QoS Class elements.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.
2) Browse to **Other Elements** > **QoS Classes**.
3) Right-click a QoS Class element, then select **Properties**.
4) Select **Override Link Selection Preferences set in Network Applications and Protocol Agents**.
5) Configure the settings, then click **OK**.

Related tasks

Create QoS Class elements on page 981

Use QoS Class elements to apply custom link selection options to traffic

You can use QoS Class elements in Access rules to override the default link selection options in Network Application elements and Protocol elements.

Before you begin

You must create QoS Class elements and define link selection options in them.

Note

Only layer 3 physical interfaces on Security Engines, Master Engines, and Virtual Engines in the Engine/VPN role support QoS.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Browse to **Policies** > **Engine Policies**.
3) Right-click the Engine Policy, then select **Edit Engine Policy**.
4) On the **IPv4 Access tab** or to the **IPv6 Access** tab, add a rule.
5) Define the source and destination according to your needs, or set the value to ANY.
6) In the **Service** cell, add elements that represent the type of traffic to which you want to apply custom link selection options.
   - To override the default link selection options in a Network Application element, add the Network Application element.
   - To override the default link selection options in a Protocol element, add a Service element that uses the Protocol.
7) Set the action to **Allow** or **Continue**.
8) In the **QoS Class** cell, add the QoS Class element that defines the link selection options.
9) Save and refresh the policy to transfer the changes.

Related tasks

Create QoS Class elements on page 981
Apply QoS to traffic on page 985

Create Connection Type elements

Connection type elements allow you to define how NetLinks and VPN endpoints are used in a Multi-Link configuration, and which VPN endpoints can communicate with each other.

You can use the following default Connection Type elements for links in a Multi-Link configuration:

- **Active** — The link is always used. If there are multiple active links, the traffic is load-balanced between the links based on the load of the links. Traffic is directed to the link that has the lowest load.
- **Aggregate** — The link is always used, and each connection is load-balanced in round-robin fashion between all of the aggregate links. For example, if there are two aggregate links, a new connection is directed to both links.
- **Standby** — The link is used only when all active or aggregate links are unusable.

If the default Connection Type elements meet your needs, it is not necessary to create custom Connection Type elements.

The link type option in Connection Type elements is an identifier that allows you to group together similar types of ISP connections. You can use any link type to represent any type of ISP connection as long as you consistently use the same link type for the same type of ISP connection. The link type determines how the connection type is used in Link Usage Profile elements.

The connectivity group option in Connection Type elements defines which endpoints can communicate with each other. Only endpoints that belong to the same connectivity group can communicate with each other. The default Connection Type elements belong to connectivity group 1.

If you want to group VPN endpoints into multiple connectivity groups, you must create custom Connection Type elements. Grouping VPN endpoints into connectivity groups improves the efficiency of tunnel negotiation in VPNs, and reduces false positives related to failed tunnels in log entries and monitoring statistics. When you use multiple connectivity groups, tunnels are created only between VPN endpoints that belong to the same connectivity group. The SMC automatically disables tunnels between VPN endpoints that cannot communicate with each other. It is not necessary to manually disable unused gateway-to-gateway tunnels.

For example, you can group all endpoints that are connected to the Internet into one connectivity group, and group all endpoints that are connected to a private wide-area network that uses an MPLS connection into another connectivity group. Tunnels are not created between the endpoints that are connected to the Internet and the endpoints that are connected to a private wide-area network.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.
2) Browse to **Other Elements** > **Connection Types**.
3) Right-click **Connection Types**, then select **New Connection Type**.
4) Configure the settings, then click **OK**.

Create Link Usage Profile elements

Link Usage Profile elements define which link types are preferred, avoided, or not used for specific types of traffic.

Also, the Link Usage Profile elements can be used to enable the packet duplication functionality for Multi-Link tunnels or Forward Erasure Correction (FEC) functionality for any IPsec tunnels.

The link usage exceptions in the Link Usage Profile element override the link selection options in QoS Class elements. If you do not specify an exception for a particular QoS class, the connection is selected according to the link selection options in the QoS Class element.

For each link type, you can add QoS Class elements to one of the following categories:

- **Prefer** — The specified traffic uses the link type unless a connection with significantly higher quality is available.
  For VPN connections, links that are preferred on both ends of the connection are preferred over those that are preferred on just one end of the connection.
- **Avoid** — The specified traffic uses link type only if necessary.
- **Do Not Use** — The specified traffic must not use the link.

When you select a Link Usage Profile element in the properties of an Security Engine, the settings defined in the Link Usage Profile element are applied to outbound traffic that uses outbound Multi-Link.

When you select a Link Usage Profile element in the properties of a policy-based VPN, Route-based Tunnels group, or a VPN broker domain, the settings defined in the Link Usage Profile element are applied to all tunnels in the VPN according to their link types.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.
2) Browse to **Other Elements** > **Link Usage Profiles**.
3) Right-click **Link Usage Profiles**, then select **New Link Usage Profile**.
4) Configure the settings, then click **OK**.

Select a Link Usage Profile element for an Security Engine

When you select a Link Usage Profile element in the properties of an Security Engine, the settings defined in the Link Usage Profile element are applied to outbound traffic that uses outbound Multi-Link.

Before you begin

Create a Link Usage Profile element.

Note

The **Use Default NAT Address for Traffic from Internal Networks** option for element-based NAT must be **On** or **Automatic** to use a Link Usage Profile for an Security Engine.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click a Security Engine, then select **Edit <element type>**.
3) Browse to **Routing** > **Link Usage**.
4) From the **Link Usage Profile** drop-down list, select a Link Usage Profile element.
5) Click **Save and Refresh**.

Define exceptions to the Link Usage Profile for an Security Engine

Security Engine-specific exceptions to the Link Usage Profile allow you specify which traffic uses specific NetLinks.

Before you begin

Select a Link Usage Profile element for the Security Engine.

For example, if you do not want some traffic to be balanced between the NetLinks, add an exception to direct the traffic through a specific NetLink.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click a Security Engine, then select **Edit <element type>**.
3) Browse to **Routing** > **Link Usage**.
4) To add a row to the **Exceptions to Link Usage Profile** table, click **Add**, then define the matching criteria.
5) Click **Save and Refresh**.

Select a Link Usage Profile for a policy-based VPN

To override the default link selection options in QoS Class elements, select a Link Usage Profile element for the policy-based VPN.

The settings defined in the Link Usage Profile element are applied to all tunnels in the VPN according to their link types.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.
2) Browse to **Policy-Based VPNs**.
3) Right-click a **Policy-Based VPN element**, then select **Properties**.
4) From the **Link Usage Profile** drop-down list, select a Link Usage Profile element.
5) Click **OK**.

Select a Link Usage Profile for a Route-based Tunnels group

To override the default link selection options in QoS Class elements, select a Link Usage Profile element for the Route-based Tunnels group.

The settings defined in the Link Usage Profile element are applied to all tunnels in the VPN according to their link types.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.
2) Browse to **Other Elements** > **Tunnel Groups**.
3) Right-click a Tunnel Group element, then select **Properties**.
4) From the **Link Usage Profile** drop-down list, select a Link Usage Profile element.
5) Click **OK**.

Select a Link Usage Profile for a VPN Broker domain

To override the default link selection options in QoS Class elements, select a Link Usage Profile element for the VPN Broker domain.

Before you begin

To create VPN Broker domains, configure the VPN Broker in the Security Management Center.
For more information about VPN Broker, see the *Forcepoint Security Engine Manager and VPN Broker Product Guide*.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.
2) Browse to **VPN Broker Domains**.
3) Right-click a VPN Broker Domain element, then select **Properties**.
4) From the **Link Usage Profile** drop-down list, select a Link Usage Profile element.
5) Click **OK**.

Part IX

Traffic inspection policies

Contents

- Creating and managing policy elements on page 799
- Ethernet rules on page 825
- Access rules on page 831
- NAT rules on page 851
- Inspection Policy elements on page 867
- Snort inspection on Security Engines on page 877
- Editing policies on page 885
- Defining IP addresses on page 919
- Working with Service elements on page 927
- Defining Situation elements on page 941
- Using Network Application elements on page 961
- Defining User Response elements on page 969
- Quality of Service on page 973
- Anti-malware scanning on page 989
- File filtering on page 993
- Integrating Forcepoint One Endpoint with Forcepoint Network Security Platform on page 1011
- Filtering URLs on page 1019
- Protocol Agents on Security Engines on page 1031
- Sidewinder Proxies on page 1043
- Setting up TLS inspection on page 1063
- Setting up QUIC inspection on page 1083
- Forward traffic to a proxy service for external inspection on page 1087
- Configuring Explicit HTTP Proxy (Experimental) on page 794

- Block listing IP addresses on page 1093

Policies are key elements that contain rules for allowing or blocking network traffic and inspecting the content of traffic.

Configuring Explicit HTTP Proxy (Experimental)

You can configure the Explicit HTTP Proxy to allow clients to send traffic to the Engine before it is sent to the destination.

When the Engine receives traffic, it resolves the destination address for the follow-up connection where the traffic is sent to. Additionally, the Engine uses the resolved destination address to match the follow-up access rules for the traffic before sending the traffic to the destination.

Note

The Engine does not resolve the Source IP address to make it available for the follow-up access rule matching. Therefore, a NAT rule is required to ensure the follow-up connections are routed correctly.

Benefits:

- The client internet access is controlled by proxy settings to route through the explicit proxy IP or port of the Engine.
- Connection authentication occurs even if the connection cannot be decrypted without External Certificate Authority (ECA).

To configure the Explicit HTTP proxy, follow these general steps:

1) Create a service element for the HTTP Explicit proxy. For more details, refer to the *Create a service element for HTTP Explicit proxy* topic.

2) (Optional) Configure the Integrated Windows Authentication (IWA) to authenticate proxy users. For more details, refer to the *Configuring the Integrated Windows Authentication* topic.

3) Add access rules for the Explicit HTTP Proxy in the Engine policy. For more details, refer to the *Add access rules for Explicit HTTP Proxy* topic.

Related tasks

Create a service element for HTTP Explicit proxy on page 795
Configuring Integrated Windows Authentication on page 796
Add access rules for Explicit HTTP Proxy on page 797

Create a service element for HTTP Explicit proxy

You can create a Service element with the HTTP Explicit Proxy settings enabled to use with Explicit HTTP proxy rules.

Steps

1) Select **Engine Configuration**.

2) Browse to **Other Elements** > **Services**.

3) Create the Service element in one of the following ways:
   - To create an element with no settings predefined, right-click the branch for the type of Service you want to create, then select **New** > **TCP Service**.
   - To create a Service based on an existing TCP Service element, right-click the existing Service, then select **New** > **Duplicate**.

4) On the **General** tab:
   a) Provide the new Service a unique **Name** and write an optional **Comment**.

   b) Configure the settings as required.

   c) Select the HTTP protocol agent:
      i) Click the **Select** button against the **Protocol** field. The **Protocol Agent** dialog box is displayed.

      ii) Select the **TCP Protocol Agent** option.

      iii) Select the **HTTP** option.

      iv) Click the **Select** button.

5) On the **Protocol Parameters** tab:
   a) Configure the settings as required.
   b) Select the **Enable HTTP proxy** checkbox.
   c) (Optional) Select the **HTTP proxy auth by IWA** checkbox.



Note

You must configure the IWA settings before you select the **HTTP proxy auth by IWA** checkbox. For more details, refer to the *Configuring the Integrated Windows Authentication* topic.

   d) (Optional) Select the **Cache authenticated proxy users** checkbox.

6) Click the **OK** button.

Configuring Integrated Windows Authentication

You can configure Integrated Windows Authentication (IWA) to authenticate Explicit HTTP Proxy users via the browser.

Steps

1) Configure an active directory server element. For more details, refer to the *Create Active Directory Server elements* topic.
2) Bind the configured active directory server element to the LDAP domain. For more details, refer to the *Define LDAP domain elements* topic.
3) Join the Engine to the domain. For more details, refer to the *Join an Engine to a domain* topic.

Related tasks

Create Active Directory Server elements on page 1108
Define LDAP domain elements on page 1112
Join an Security Engine to a Domain on page 355

Add access rules for Explicit HTTP Proxy

To enable the Explicit proxy feature, you must add access rules for the Explicit HTTP Proxy in the Engine Policy.

Before you begin

You must have the following configured:

- The Explicit HTTP Proxy service element. For more details, refer to the *Create a service element for HTTP Explicit proxy* topic.
- The Proxy Follow-up Sub-policy. For more details on how to create a sub-policy, refer to the *How sub-policies work* topic.
- **Note**

  Before the traffic is sent to the destination, the Proxy Follow-up sub-policy rules are used to control the traffic that matches the Explicit HTTP Proxy access rule.

Note

The access rules for the Explicit HTTP Proxy must always be positioned at the top in the policy.

Steps

1) Select **Engine Configuration**.
2) Browse to **Policies** > **<Engine Policies>**.
3) Right-click a policy, then select **Edit <Policy name>**.
4) To select traffic for Explicit HTTP Proxy, you can add the Access rules similar to the following:

Source	Destination	Service	Action
ANY	Engine	Explicit HTTP Proxy service element	Allow **Proxy Jump:** Proxy Follow-up Sub Policy

5) 5. Right-click the **Source** cell in the Access rule, then select a source of the traffic.
6) Right-click the **Destination** cell in the Access rule, then select a destination for the traffic.
7) Right-click the **Service** cell in the Access rule, then select the Explicit HTTP Proxy service element.

8) For the **Action** cell:
   a) Right-click the **Action** cell in the Access rule, then select **Allow**.
   b) Right-click the **Action** cell, then select **Edit Options**. The **Action Options** dialog box is displayed.
   c) On the **General** tab, click the **Select** button against the **HTTP Proxy Policy** field to select a sub-policy.

9) Click **Save and Install**

Chapter 47

Creating and managing policy elements

Contents

- Getting started with policies on page 799
- Create template policies or policies on page 808
- How sub-policies work on page 810
- Install policies on page 813
- Using policy elements and rules on page 816
- Deleting policies templates or sub-policies on page 819
- Engine Policy elements examples on page 819
- IPS Policy example on page 821
- Local alternative policies on page 822

Policy elements are containers for the rules that determine how Security Engines, Master Engines, and Virtual Engines examine traffic. The policy elements for the engines include Template Policies, Policies, and Sub-Policies.

Getting started with policies

Policies organize traffic processing rules hierarchically, to make administration easier and to optimize traffic inspection performance.

What policy elements do

- Engine, IPS, and Layer 2 Engine Policies contain the rules according to which the Security Engines allow or block traffic.
- Layer 2 Interface Policies contain rules according to which Security Engines in the Engine/VPN role allow or block traffic detected by Capture Interfaces, Inline IPS Interfaces, and Inline Layer 2 Engine Interfaces on Security Engines in the Engine/VPN role.
- The same policy can be shared by several Security Engines that have the same role, several Master Engines, and several Virtual Engines that have the same role.
- Inspection Policies contain the rules according to which the Security Engines inspect traffic. The same Inspection Policy can be shared by several Engine Policies, IPS Policies, and Layer 2 Engine Policies.

  Note

  Inspection Policies are not supported in Layer 2 Interface Policies.

- Each policy must always be based on a Template Policy. Template Policies contain rules that are inherited into any template or policy below it in the policy hierarchy.

- You can also insert Sub-Policies in your policies. A Sub-Policy is a set of IPv4 or IPv6 Access rules that can be matched conditionally to a restricted part of the traffic. Using Sub-Policies can improve processing performance. Sub-Policies can also enforce administrative boundaries.
- Policies can share Policy Templates and Sub-Policies. In shared rules, Alias elements can represent IP addresses that depend on the environment, so that the actual values are defined separately for each component.

What do I need to know before I begin?

- Master Engines always use Engine Policies, regardless of the role of the Virtual Engines they host.
- Virtual Engines use Engine Policies.
- Virtual IPS engines use IPS Policies.
- Virtual Layer 2 Engines use Layer 2 Engine Policies.

Default policy elements

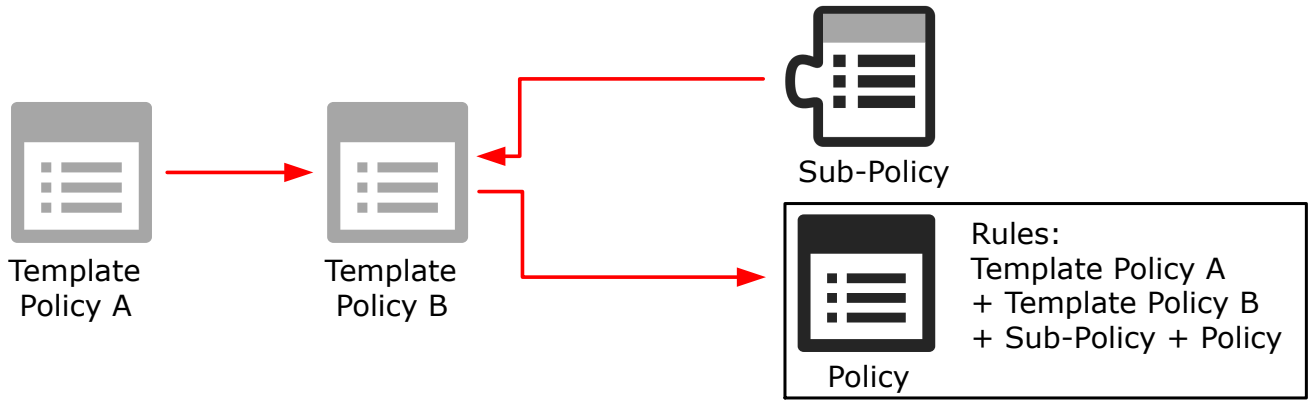
Predefined Policies, Template Policies, and Inspection Policies provide an easy starting point for determining what kinds of rules your system needs.

Four kinds of policy elements are used in the policy configuration for Security Engines, Master Engines, and Virtual Engines:

- A Template Policy is a policy that can be used as the basis for Policy or other Template Policy elements. The rules in the Template Policy are copied as *inherited rules* into the Policies and Template Policies that are based on the Template Policy. You can edit the inherited rules only by editing the original Template Policy from which the rules were inherited.
- An Inspection Policy element is a set of Inspection rules that are referenced from the **Inspection** tab of Policy and Template Policy elements. You can use the same Inspection Policy in multiple Policy and Template Policy elements.
- A Sub-Policy element is a section of IPv4 Access rules that you can insert into Policies and Template Policies. The rules in the sub-policy are conditional rules that allow you to define matching criteria that determines whether the sub-policy applies to a connection. You can edit the rules by editing the sub-policy where the rules belong.
- A Policy element gathers together all rules from the different policy elements. Policies include the rules added directly to the Engine, IPS, or Layer 2 Engine Policy, the rules from the higher-level Template Policy, and possibly rules from one or more sub-policies. For example, an IPS Policy is always based on an IPS Template Policy element, and a Layer 2 Engine Policy is always based on a Layer 2 Firewall Template Policy element.

The hierarchy of how rules are inherited between the main policy elements is shown in the following illustration.

Rule inheritance (without Inspection rules inherited from Inspection Policies)



In this illustration, Template Policy A is the basis for Template Policy B, so Template Policy B contains all rules defined in Template Policy A.

Template Policy B also contains all rules in a Sub-Policy, as well as rules defined directly in Template Policy B.

The example Policy inherits the following rules:

- All rules in Template Policy A.
- All rules in Template Policy B.
- All rules in the Sub-Policy.

In addition to the inherited rules, the example policy also contains any rules that the administrators add to it directly. In the policy, the administrators can only edit the rules that were added directly to the policy. To change rules inherited from Template Policy A, Template Policy B, or the Sub-Policy, they must edit the policy in which the rules were originally defined.

A hierarchy such as the one outlined here is useful to:

- Reduce the need for creating the same or similar rule in several policies. For example, any rule added to Template Policy A is also added to any policy created based on that template. The next time the policies based on Template Policy A are installed on the engines, the new rule is used on all those engines. There is no need to edit each individual policy separately.
- Restrict the editing rights of administrators. For example, administrators who are granted rights to only policies cannot edit the rules defined in the Template Policies on which the policies are based. Their actions have no effect on rules that are placed above the row where the Template Policy allows them to insert new rules. In the hierarchy shown in the illustration, the insert points for the Policy are defined in Template Policy B. Template B can be edited only in the place where there is an insert point in Template Policy A.
- Reduce the likelihood of mistakes affecting important communications. Template Policies can be reserved for defining only the rules for essential communications, so that most daily editing is done in the lower-level policies. If the Template Policy is properly designed, rules in the lower-level policy cannot override the rules in the Template Policy. Good organization also makes it easier to read policies, and reduces the risk of errors.
- Improve processing performance. Using sub-policies, whole blocks of rules can be skipped during processing when a connection does not match the rule that directs the traffic processing to the sub-policy. Using sub-policies reduces the processor load, which can lead to better throughput if the processor load is constantly high.

The default policy elements are introduced into the SMC when you import and activate a recent dynamic update package (for example, during the installation). The elements might change when you install newer update packages. None of the default policy elements can be changed. However, you can make copies of the default policies if you have to create a changed version.

Default policy elements for Security Engines

Element type	Default name	Description
Firewall Template Policy	Firewall Template	A Template Policy that contains the predefined Access rules necessary for the engine to communicate with the SMC and some external components. The Firewall Template Policy uses Inspection rules from the No Inspection Policy. The rules in the No Inspection Policy do not enforce inspection.
	Firewall Inspection Template	Firewall Template Policy that uses Inspection rules from the High-Security Inspection Template.
Engine Sub-policy	DHCP Relay	A Sub-Policy containing rules that allow the engine to relay DHCP requests from a host in one internal network to a DHCP server in a different network, as well as DHCP requests from VPN clients to an internal DHCP server.
IPS Template Policy	High-Security IPS Template	IPS Template Policy that uses Inspection rules from the High-Security Inspection Template. A Template Policy containing the predefined Access rules necessary for the IPS engine to communicate with the SMC and some external components. The High-Security IPS Template Policy provides an easy starting point for determining what kinds of rules your system needs.
	Medium-Security IPS Template	IPS Template Policy that uses Inspection rules from the Medium-Security Inspection Policy.
IPS Policy	Customized High-Security Inspection IPS Policy	Example of a customized IPS Policy that uses Inspection rules from the Customized High-Security Inspection Template. Used in testing IPS at ICSA Labs and NSS Labs.
	Default IPS Policy	Basic IPS Policy that uses Inspection rules from the High-Security Inspection Template. Can be used as a starting point for creating a customized IPS Policy. The Default IPS Policy does not add any rules to the rules defined in the IPS Template. It allows you to install the predefined rules in the IPS Template on the IPS engine right after installation. (Template Policies cannot be installed on the engines.)
Layer 2 Firewall Template Policy	Layer 2 Firewall Template	A Template Policy that contains the predefined Access rules necessary for the Layer 2 Engine to communicate with the SMC and some external components. The Layer 2 Firewall Template uses Inspection rules from the No Inspection Policy. The rules in the No Inspection Policy do not enforce inspection.
	Layer 2 Firewall Inspection Template	A Template Policy that is based on the Layer 2 Firewall Template. The Layer 2 Firewall Inspection Template uses Inspection rules from the High-Security Inspection Template. The Layer 2 Firewall Inspection Template enables deep inspection for all traffic.
Layer 2 Interface Policy	Layer 2 Interface Template	A Template Policy that contains rules for traffic detected by Capture Interfaces, Inline IPS Interfaces, and Inline Layer 2 Engine Interfaces on engines.
Inspection Policy	No Inspection Policy	Suitable for Engine deployments, in which only packet filtering is needed. Disables deep packet inspection.

Element type	Default name	Description
	Medium-Security Inspection Template	For Engines, Layer 2 Engines, inline IPS deployments in asymmetrically routed networks, and IPS deployments in IDS mode. Terminates reliably identified attacks and logs Situations that have some degree of inaccuracy. Low risk of false positives.
	High-Security Inspection Template	For Engine, Layer 2 Engine, and inline IPS use. Extended inspection coverage and evasion protection. Not for asymmetrically routed networks. Terminates reliably identified attacks, and Situations that have some inaccuracy. Moderate false positive risk.
	Customized High-Security Inspection Policy	This policy is an example of a highly customized Inspection Policy for network environments in which unconditional inspection coverage and evasion protection are required. The risk of false positives is high in production use.

You can add new Template Policies without basing them on any existing Template Policy.

However, in most cases we recommend using the IPS Template as the starting point for your customized IPS Template Policies and IPS Policies. We recommend using the Layer 2 Firewall Template as the starting point for your customized Layer 2 Firewall Template Policies and Layer 2 Engine Policies.

Situations are the central elements in the Inspection rules of your Inspection Policies. The Situation elements detect exploit attempts against known vulnerabilities and other commonly known security threats. Dynamic updates include new and updated Situations. New patterns in traffic might begin to match when a new dynamic update is activated and you refresh the Inspection Policy.

In most environments, we recommend using the High-Security Inspection Policy as the starting point for Inspection Policies. The High-Security Inspection Policy provides extended inspection coverage. It also protects the network against evasions, which are attempts to disguise attacks to avoid detection and blocking by network security systems. The only difference between the rules in the High-Security Inspection Policy and the Medium-Security Inspection Policy is in the way the Inspection rules handle Situations that are categorized as Suspected Attacks. The High-Security Inspection Policy terminates Suspected Attacks with an alert, whereas the Medium-Security Inspection Policy only logs Suspected Attacks.

Situations that belong to the Suspected Attacks category contain basic fingerprints. Suspected Attacks also contain traffic patterns that might indicate malicious activities but are not any verified attack patterns. Suspected Attacks can detect zero-day attacks (attacks that are not yet publicly known), but can sometimes block some legitimate traffic if the traffic pattern resembles malicious activities.

Policy hierarchies

The policy structure is a hierarchy based on templates.

The structure allows you to:

- Reuse rules without duplicating them.
- Assign and enforce editing rights of different parts of a single policy to different administrators.
- Reduce the resource consumption of the engines.
- Make policies easier to read.

The template and policy hierarchy is flattened when the Policy is transferred to the engines. The policy looks the same to the engines regardless of how it is organized on the Management Server (as long as the rules are in the same order). You can also create sections of conditional IPv4 Access rules that you can insert into the other policy elements. The engine can skip the processing of a conditional block of rules based on whether certain common matching criteria is found in the packet being examined.

If your environment is simple and you do not need the benefits outlined here, you can create a simple policy hierarchy. You can, for example, start with one Engine Policy built on the provided Firewall Template. The same Engine Policy can be used on more than one engine. Likewise, you can use the same IPS Policy on any number of IPS engines and Virtual IPS engines, and the same Layer 2 Engine Policy on any number of Layer 2 Engines and Virtual Layer 2 Engines.

How Security Engines examine traffic

Security Engines check the validity of packets, then process only the valid packets.

Engines, Master Engines, Virtual Engines, and Layer 2 Engines pass through only traffic that is explicitly allowed in the policy. All other traffic is discarded. IPS engines allow packets to pass if the engine's policy does not apply a more specific action. All connections are handled in the same way, including connections that the Security Engine itself opens, and the management connections that the Security Engine is intended to receive.

On Engine Clusters and clustered Master Engines, the load-balancing filter first determines which node in the cluster actually processes the received packet. The processing then begins on the selected node.

Security Engines check new connections against the policy, rule by rule. The header on each packet arriving on an interface is examined for the source and destination IP address, and protocol-related information, such as the port. The authentication status of the user attempting a connection and the current date and time can also be included as parameters in the examination process.

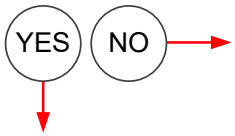
An IPS engine or Layer 2 Engine matches traffic to different protocols and then checks the definitions for known vulnerabilities and other threats for that protocol. The protocol is assigned first, before the deep inspection. An Security Engine in inline mode can also filter traffic based on protocols, IP addresses, and the interface that received the traffic without analyzing the traffic for threat patterns. IPS engines and Layer 2 Engines can be installed either in inline mode or in capture mode.

The following table describes the packet processing flow for Security Engines. The following abbreviations are used for the roles and interface types:

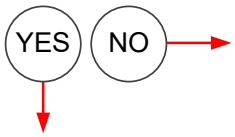
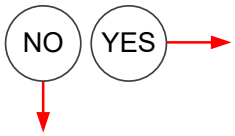
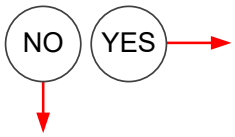
- FW — Engine/VPN role
- IPS — IPS role and layer 2 physical interfaces of the inline IPS interface type on Security Engines in the Engine/VPN role
- L2FW — Layer 2 Engine role and layer 2 physical interfaces of the inline Layer 2 Engine interface type on Security Engines in the Engine/VPN role

Some stages apply to all roles.

Packet processing

	Stage	Role	Event / Check	Action	Notes
1	PACKET IN	ALL	Incoming packet 		
2	Ethernet rules	IPS, L2FW	Allowed? 	Packet dropped	The Security Engine checks Ethernet frames against the Ethernet rules in the policy. The packet is processed until it matches an Ethernet rule that tells whether to allow or to discard the packet.

	Stage	Role	Event / Check	Action	Notes
3	Packet validity checks	ALL	Valid packet? YES NO →	Packet dropped	The Security Engine checks the validity of packets to protect the Security Engine from processing invalid packets. Depending on the action in the Inspection Policy, Security Engines in the IPS and Layer 2 Engine roles can either terminate invalid packets or bypass the traffic without processing the invalid packets. For self-protection reasons, Security Engines in the Engine/VPN role role always terminate invalid packets without establishing a connection. Packets dropped in this phase produce log entries with specific situations, such as TCP_Segment-SYN-No-Options. Not all invalid packet drops are logged unless Packet Filter diagnostics is enabled. **Note** Invalid Packet Situations are always log rate limited.
4	Antispoofing	FW	Spoofed? NO YES →	Packet dropped	Checks that the traffic is coming in through the correct interface or VPN tunnel as defined in the routing and antispoofing configuration or VPN configuration.
5	Connection tracking	ALL	Existing connection or new related connection? YES NO →	Access rule matching	Checks the current connection tracking information to see if the packet is part of an established connection (for example, a reply packet to a request that has been allowed). If TCP SYN rate limits or other DoS protection features are enabled, they are enforced at this stage. If the packet is not part of an existing connection, the packet is compared with the Access rules in the installed policy. The processing continues until the packet matches a rule that tells the Security Engine to allow or stop the packet. If there is no rule match anywhere else in the policy, the packet is allowed or discarded according to the final action of the policy.
6	Access rule matching	ALL	Allowed by Access rules? YES NO →	Packet dropped	If the packet is not part of an existing connection, the packet is matched to IPv4 or IPv6 Access rules according to the IP protocol used. ■ If the traffic is tunneled using IP-in-IP or Generic Routing Encapsulation (GRE), the traffic can be checked against the IPv4 or IPv6 Access rules several times. The traffic is re-matched according to the number and type of layers in the tunnel and the settings of the Security Engine. ■ The processing of the packet continues until the packet matches a rule that tells the Security Engine to allow or discard the packet. If the packet does not match any Access rule, the final action depends on the Security Engine role. An IPS engine allows the packet to pass through, while a Layer 2 Engine drops the packet.

	Stage	Role	Event / Check	Action	Notes
7	Protocol Validation	ALL	Valid for connection state? 	Packet dropped	If the packet is allowed as an existing connection or in an Access rule, checks that the packet is valid for the state of the connection. If not, the packet is dropped. For example, a TCP connection must always begin with a SYN packet (as defined in the protocol standards). The Security Engine checks that the first packet of a new connection is a valid SYN packet.
8	Inspection and File Filtering Rules	ALL	Harmful pattern or file? 	Action defined for rule	Applies Inspection rules to connections that are selected for deep packet inspection, Snort inspection, or file filtering in the Access rules. ■ Inspection applies to all packets in a connection, so the Inspection rules are applied even if the packet is a part of an existing connection. ■ If both Security Engine deep packet inspection and Snort inspection are enabled, the order in which inspection is applied depends on the direction of the traffic. For traffic from a client to a server, Snort inspection is applied before Security Engine deep packet inspection. For traffic from a server to a client, Security Engine deep packet inspection is applied before Snort inspection. ■ The Inspection rules look for patterns of interest in allowed connections. The patterns can indicate potential attacks, exploits, or other possible threats. They can also be any other patterns of interest, such as multiple logon attempts, use of peer-to-peer or instant messaging software, or protocol violations in the traffic. ■ If a pattern in traffic matches a pattern defined in a rule, the actions defined in the rule are taken.
9	NAT Modifications	FW	Address Translation (if needed for the connection) 		Applies Network Address Translation (NAT) rules to IPv4 and IPv6 connections. The source and destination addresses are translated according to the first matching NAT rule (or not done at all, if a NAT rule so defines). If none of the NAT rules match, the packet continues with the original addresses. By default, NAT is not applied to traffic to or from policy-based VPNs.
10	VPN	FW	Policy-based VPN processing 		Policy-based VPN processing includes finding the correct tunnel from the VPN configuration. If a VPN configuration does not match Access rules correctly, packets might be discarded with the message ”VPN tunnel selection failed”.
11	Routing / Route-based Tunnels	FW	Route Selection  GRE/IP-IP encapsulated with policy-based VPN? 	Go back to policy-based VPN processing	If a Zone is used as the Destination in an Access rule, the packet is matched against the Access rules after the final route has been selected. If the packet does not match a rule due to applied NAT, it is dropped. If the packet is routed to a tunnel interface, the packet is encapsulated according to the Route-based Tunnels configuration.

	Stage	Role	Event / Check	Action	Notes
12	QoS	ALL	QoS Processing 		The packet is let through the Security Engine according to its priority and any bandwidth limits or guarantees that might have been defined.
13	PACKET OUT	ALL	Outgoing packet		

Access rule matching based on the payload of connections

When you use some types of Service elements in Access rules, the Security Engine can only determine whether the connection matches a rule when the payload of the packets is checked against the Access rules.

When you use elements such as Network Applications, URL Categories, or URL List Applications in the Service field of an Access rule, matching is based on the payload of the packets. When the first SYN packet of a new connection is processed, the Security Engine cannot determine whether the connection matches the Access rule. The Security Engine can only determine whether the connection matches the Access rule when the Security Engine processes, for example an HTTP request in an HTTP connection.

The Security Engine checks traffic against the Access rules from the top down. Matching criteria that do not depend on the payload of the connection, such as the source and destination IP address and port, are always evaluated first. If a connection might still match another rule that allows traffic, the connection is considered potentially allowed. When enough of the payload has been processed, the number of rules that could potentially allow the connection gets smaller.

When traffic matches a rule that tells the Security Engine to allow or discard the packet, the Security Engine stops checking traffic against the Access rules. Because the first matching rule defines how the first packet is forwarded, connections might not match the intended rule.

Application routing

You must use network applications that have the Application Routing tag because the routing decision is made based on the application that is detected in the traffic. For other network applications, if the network application cannot immediately be identified, the routing decision is made according to the first rule that could potentially allow the connection.

Routing decisions are delayed until enough of the payload has been processed to identify the network application. If you use features that are not compatible with delaying the decision, use more specific source and destination criteria in the rules, or change the rule order.

If a rule that could potentially allow the connection activates a feature that is not compatible with delaying the routing decision, the decision is made according to the first rule that could potentially allow the connection.



Important

After the routing decision has been made, the Security Engine might later identify a different application in the connection. If the application that is detected would cause a different routing decision to be made, the connection might be discarded.

Snort inspection

We do not recommend using services that match based on the payload of connections, such as Network Applications, URL Categories, or URL List Applications, in Access rules that select traffic for Snort inspection. At the beginning of a connection, the Security Engine cannot determine whether the traffic should be selected for Snort inspection. The Security Engine selects all potentially matching traffic for Snort inspection. As a result, Snort inspection might be applied to traffic that was not intended to be selected for Snort inspection. Applying Snort inspection to this traffic can create false positive Snort rule matches.

Policy element configuration overview

After you decide on a policy hierarchy, you can populate the policy elements with rules for handling the traffic.

Policy elements are only containers for the actual traffic handling rules. When you have decided on a policy hierarchy, you can populate the policy elements with the rules for handling the traffic.

1) (Optional) Create a custom template and add the rules and insert points.

2) Create a custom policy and add the rules.

3) (Optional) Create sub-policies and add the IPv4 Access rules.

Related concepts

How sub-policies work on page 810

Related tasks

Create template policies or policies on page 808

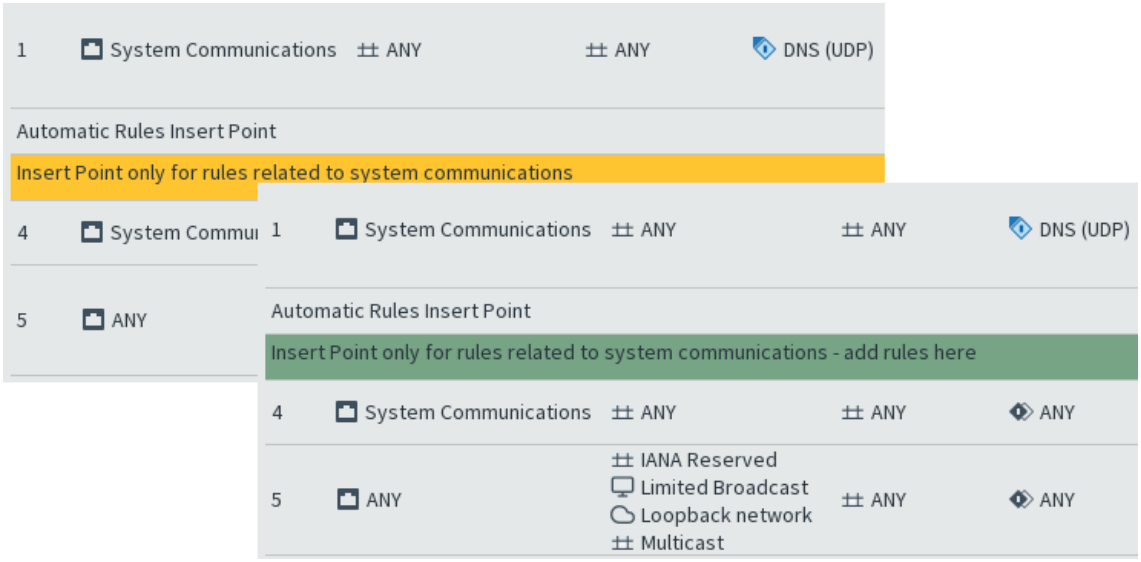
Create template policies or policies

Template Policy elements are used as a basis for Policies and other Template Policies.

Every Policy and Template Policy that you create is based on a Template Policy. You can base several policies on the same Template Policy. The Template Policy or a customized copy of the Template Policy is always at the highest level of the policy hierarchy. It is not mandatory to create any custom Template Policies if you feel that it is not necessary in your environment.

When editing policies, the main difference between Policies and Template Policies are the special rows called insert points. Insert points are shown in both Template Policies and in Policies, but you can add them only to Template Policies. The insert points added to Template Policies mark where new rules can be added to policies that are based on the templates. If you create a Template Policy and do not base the Template Policy on any predefined Template Policy, you must add insert points separately for Access rules, NAT rules, and Ethernet rules.

Insert point in a Template Policy and the inheriting (Template) Policy



This illustration shows what the same insert point looks like in a Template Policy and in the inheriting policy elements. The color of the insert point indicates whether the insert point has been added in the current Template Policy for inheritance to lower levels (orange) or whether it has been inherited from the higher-level Template Policy (green). Only the orange insert points are inherited to lower-level policy elements. You must add at least one new insert point at each Template Policy level to make the lower-level policies editable. When you add the first new rule to the green insert point, the rule replaces the insert point. Any number of rules can then be added directly above and below that first rule. The engine reads rules in order from the top down. The rules above the insert point in the higher-level Template Policy cannot be canceled by anything a lower-level policy adds into the insert point.

Rules defined in the Template Policy itself cannot be edited in lower-level policies that use the Template Policy. Such inherited rules are shown only on your request and they are displayed with a gray background. Only the actual rules are inherited from a higher-level Template Policy into the lower-level policies and Template Policies. The rights to edit policies and Template Policies are defined separately.

A Engine Policy, IPS Policy, Layer 2 Engine Policy, or Layer 2 Interface Policy is the element that gathers all rules from the different policy elements:

- Rules inherited from the Template Policy that is used as the basis of the policy
- Rules from one or more Sub-Policies added to the policy
- Rules added directly to the policy
- Rules from the Inspection Policy that is referenced from the Inspection tab in the policy

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click the **Policies** branch and select one of the following options:
   - **New > Engine Policy**.
   - **New > IPS Policy**.
   - **New > Layer 2 Engine Policy**.
   - **New > Layer 2 Interface Policy**

3) In the properties dialog that opens, enter a **Name** for the element.

4) Select the **Template** you want to base this template or policy on.

5) (Optional) Switch to the **Permissions** tab and grant rights for the template or policy.
   a) To add a permission, click **Add Permission**.
      A new row appears on the administrator list.
   b) Click the **Administrator** cell and select the Administrator.
   c) Double-click the **Administrator Role** cell to select the correct role.

6) Click **OK**.
   The new Template Policy or Policy opens in the **Policy Editing** view.
   If you changed administrator permissions for the policy, the changes are applied immediately. The permissions are also automatically updated in the properties of the administrator's account.

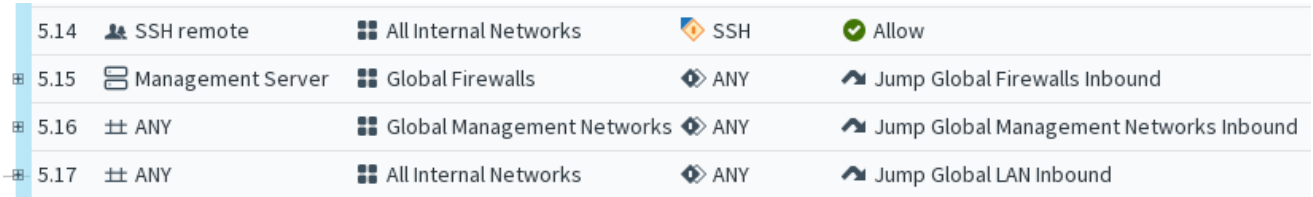
How sub-policies work

Sub-Policies are sections of Access rules that you can insert in the IPv4 or IPv6 Access rules of Policies, Template Policies, or other Sub-Policies.

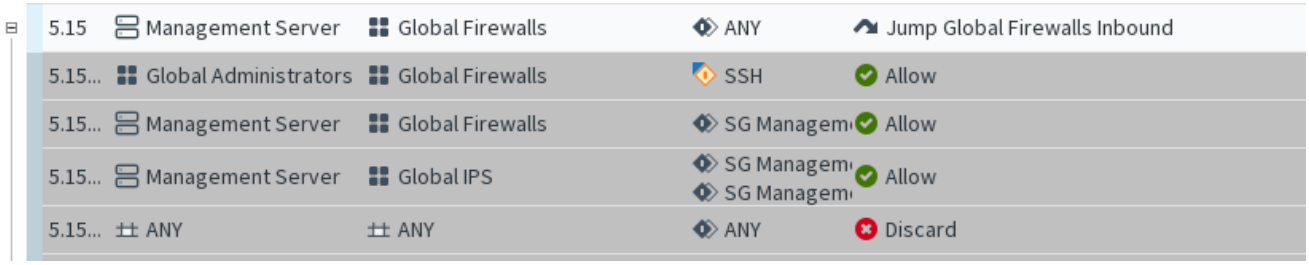
Sub-Policies can also be used to organize the policies and to delegate administrator rights. You can restrict specific administrators to edit, add, or delete rules within a limited section of IPv4 or IPv6 Access rules, which is more restrictive than giving access to a template or policy.

A Sub-Policy is inserted into some other policy element by adding a Jump rule to the policy element. The Jump rule directs connections that match the Jump rule for matching against the rules in the Sub-Policy.

A collapsed Sub-Policy



An expanded Sub-Policy



The illustrations show the same Jump rule in a policy in the collapsed and the expanded state. The rules of the Sub-Policy are shown on a gray background, because they can be edited only within the Sub-Policy itself, not in the Engine Policy that uses the rules.

You could use a Sub-Policy, for example, for examining traffic destined to a group of servers located in one particular network. The Jump rule could then use the destination network as a criteria for directing connections for matching against the Sub-Policy. Any connection that was destined to some other network would not get matched against any of the rules in the Engine Sub-Policy. This makes the matching process faster, because the engine can skip a whole Sub-Policy by comparing a connection to just one simple rule for any non-matching connection. If the Sub-Policy rules were inserted directly into the main Engine Policy, the engine would have to compare all connections to all those rules individually (because that is the only way to find out whether the connection matches the rules). The performance benefit gained depends on the number and complexity of the rules that can be placed in a Sub-Policy. It also depends on how heavy the engine load is to begin with.

The main goal of Sub-Policies is to match only part of the traffic against the rules in the Sub-Policy and allow other traffic to bypass the Sub-Policy. Rules in a Sub-Policy should have at least some identical matching criteria (source, destination, service, source VPN, or user authentication details). These criteria can be used to select only a portion of the traffic for the sub-policy checks.

There are two ways to create Sub-Policies: you can either create a Sub-Policy and add Access Rules to it manually or select Access rules in a policy and convert them into a Sub-Policy.

Create empty sub-policies

Before you can edit a sub-policy, you must create an empty sub-policy.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click the **Policies** branch and select one of the following options:
   - **New > Engine Sub-Policy**.
   - **New > IPS Sub-Policy**.
   - **New > Layer 2 Engine Sub-Policy**.
   - **New > Layer 2 Interface Sub-Policy**

3) Enter a unique **Name** for the element.

4) (Optional) Click the **Permissions** tab and adjust the Access Control Lists at the top part of the dialog box to include the Sub-Policy on one or more custom Access Control Lists.

5) Click **OK**.
   The new Sub-Policy opens in the **Policy Editing** view.

Create sub-policies from existing Access rules

You can convert IPv4 and IPv6 Access rules in an existing policy into a Sub-Policy.

The IPv4 and IPv6 Access rules do not have to be consecutive. However, if you add several references to a Sub-Policy in the same policy, all Sub-Policy rules are checked at each reference point, even if those rules were already checked at a previous reference point. This can be avoided, for example, by adding a rule at the end of the Sub-Policy that stops all connections that did not match the other rules.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click the policy or template and select **Edit <policy type>**.

3) On the **IPv4 Access** or **IPv6 Access** tab, select the rules that you want to add to the Sub-Policy.

4) Right-click one of the selected rules and select **Create Sub-Policy**.

5) Enter a **Name** for the Sub-Policy and click **OK**.
   The Sub-Policy element is created, a new Jump rule that references the Sub-Policy is automatically added to the policy, and the selected rules are moved to the Sub-Policy.

6) Edit the Jump rule cells to be as specific as possible, so that traffic is not unnecessarily matched to the sub-policy.
   If necessary, you can add more references to the Sub-Policy, for example, by copy-pasting the Jump rule.

7) (Optional) Add the Sub-Policy to a custom Access Control List:
   a) Right-click the **Action** cell in the Jump rule and select **Properties**.
      The **Properties** dialog box for the Sub-Policy opens.
   b) Switch to the **Permissions** tab and adjust the Access Control Lists at the top part of the dialog box.
   c) Click **OK**.

Install policies

After creating or editing a policy, you must install or refresh the policy on the engine.

Policy installation transfers any new engine configuration information in addition to the policy. Whenever you update the engine's configuration, you must reload the policy on the engine so that the changes take effect. These changes include, for example, changes in the routing configuration, the VPN configuration, and the properties of the Security Engine element itself. You must reload the policy even if the changes are not directly related to the rules in the policy.

Note

When you install a changed or a new Engine Policy, any existing connections that are not allowed by the new Engine Policy are dropped. The existing connections allowed by the new Engine Policy continue uninterrupted. These connections include related connections and authenticated connections on the engines.

If the policy installation fails, the system automatically rolls back to the previously installed configuration. By default, a rollback also occurs if the system detects that the new policy or related configuration (such as routing configuration) does not allow the Management Server to connect to the engines. This safety feature prevents you from inadvertently installing a configuration that would cause the critical management connections to fail.

You can only install Policy elements. Template Policy and Sub-Policy rules are installed as part of the main Policy. A Policy Snapshot is automatically created each time you install or refresh a policy. You can install a policy through the Policy element or through the engine element. The following procedure explains the first method.

Note

You cannot install Layer 2 Interface Policies on engines. Instead, you select the Layer 2 Interface Policy for the Security Engine in the Engine Editor.

For more details about the product and how to configure features, click **Help** or press **F1**.

Related concepts

Viewing policy validation issues on page 915

Related tasks

Validate rules automatically on page 912
Configure Layer 2 Settings for Security Engines in the Engine/VPN role on page 523

Check the currently installed policy

You can check the details of the currently installed policy.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click the engine, then select **Current Policy** > **Info**.
   A message is displayed on the screen with the following information:
   - Name of the installed policy
   - Name of the administrator who installed the policy
   - Date (year-month-day) and time of the policy installation

3) Click **OK** to close the message.

Preview the currently installed policy

You can preview the policy that is currently installed on the Management Server.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click the engine, then select **Current Policy** > **Preview**.
   A preview of the policy in its current format on the Management Server opens.

Note

The policy on the Management Server might be different from the actual currently installed policy if the policy has been edited after it was installed.

Check and compare Policy Snapshots

You can compare two policy snapshots, or check which policy was in place at a particular time.

Each time a policy is successfully installed, a record of that configuration is stored in the upload history. These policy snapshots enable you to the check which policies were uploaded and when, and allows you to run an automatic check for policy refresh needs.

Note

The policy snapshot is only stored when the policy is successfully installed. No snapshot is stored if the policy installation fails.

View Policy Snapshots

A policy snapshot captures a policy's configuration at the time it was installed.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Browse to **Other Elements** > **Policy Snapshots**.
A list of policy snapshots with the upload times and dates appears.

3) Double-click the policy snapshot you want to view.
The details of the selected policy snapshot open.



Tip

You can also view the snapshot in HTML format: right-click a snapshot and select **More actions** > **Save as HTML File**.

4) Select elements in the other pane to see their details. The elements are organized in the following groups:
- Policy: The policy whose installation created the snapshot.
- Target: Properties of the engine on which the policy has been saved.
- Elements: Elements included in the policy.



Tip

Select **Show XML** to view the policy snapshot in XML format.

Compare Policy Snapshots

You can compare any two policy snapshots from the list to check for changes between policy installations.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Open the comparison in one of the following ways:
- To compare any two snapshots to each other, select two policy snapshots, right-click, and select **Compare Snapshots** .
- To compare a policy snapshot to the engine's current policy, right-click a policy snapshot and select **Compare** > **Compare Snapshot to Engine's Current Policy**.
- To compare a policy snapshot to the most recently saved policy on the Management Server, right-click the policy snapshot and select **Compare** > **Compare Snapshot to Most Recently Saved Policy**

The changes are highlighted in the policy and shown as a summary.

2) Select objects in the summary to focus on them.
For example, if you select a network element, its properties are shown. The objects are organized in the following groups:
- **Targets** — The engine on which the policy has been saved.
- **New Elements** — Elements added to the policy.
- **Removed Elements** — Elements that have been removed from the policy.
- **Modified Elements** — Elements that exist in both policies but that have been edited.

Change templates of policies

You can change the template of a Engine, IPS, or Layer 2 Engine policy.

For example, if you created policy A based on template X, you can later change the parent template for policy A so that policy A inherits rules from template Y.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Expand the **Policies** tree and select the appropriate policy type.

3) Right-click the correct policy and select **Properties**.
   The **Properties** dialog box for the policy opens, showing the **Policy** tree.

4) Select the new parent template.
   The parent template must have insert points for the types of rules that the Policy contains.

5) Click **OK**.
   The policy is moved under your chosen template in the tree view, and starts using rules from the new template.

Using policy elements and rules

Names, comments, and the Continue action make policy rules easier to use. You can also enable policy validation and configure user responses and connection tracking.

Validating policies

You can automatically validate the policy while editing and during policy installation.

The number of rules in a policy can grow large over time. It can become difficult, for example, to notice configuration errors in a policy. To make policy management easier and to make sure that the policy does not contain misconfigured rules, you can automatically validate the policy while editing and during policy installation. You can select different criteria for validating the policy. You can, for example, check the policy for duplicate and empty rules or check if there are rules that cannot match traffic.

Additionally, the engines automatically count how many times each Access rule has matched. Engines also count the number of matches to NAT rules. You can run an analysis over a selected time frame in the policy editing view to display rule counter hits for each rule (in the **Hits** cell). Rule analysis allows you to find otherwise valid rules that are unnecessary because they match traffic that does not appear in your networks.

Configuring connection tracking

Connection tracking means that the engine keeps a record of all currently open connections (stateful inspection).

With connection tracking, the engine can verify that the connection proceeds according to the protocol standards. Connection tracking is also required for changing addresses using NAT and enforcing some other connection-related settings. By default, connection tracking is on.

However, it is not necessary to track the state of certain kinds of connections (for example, SNMP traps). Some applications can also communicate in a non-standard way that prevents them from being allowed through the engine when connection tracking is used. For those connections, you can disable connection tracking in the Access rules, which allows the engine to function as a simple packet filter for those connections. However, disabling connection tracking also prevents the use of some features that require connection tracking.

When connection tracking is off, each packet that you want to allow must match an Access rule that allows the packet. This means that even if two-way communications are always opened one way, both directions of communication must be explicitly allowed in Access rules. Reply packets are not recognized, so they are not automatically allowed through. If done carelessly, turning off connection tracking reduces the benefits you gain from your engine and might even weaken security. You might have other options: sometimes using the correct Protocol Agent helps.



Note

Before disabling connection tracking, always check if there is a Protocol Agent for the protocol in question. The Protocol Agents can pass connections that require special handling when connection tracking is on, which is always a more secure option.

When connection tracking is enabled in an Access rule, the **Service** cell of the rule must contain one of the protocols supported for connection tracking (ICMP, TCP, UDP, or another protocol that belongs to the IP protocol suite). ICMP and UDP are stateless protocols that do not contain any connection data. However, ICMP and UDP packets contain data that the engine can use to build virtual connections. The engine can also build virtual connections based on the IP address and IP protocol data in other types of IP packets.

You can choose between several connection tracking modes, depending on the traffic type and how strictly you want the connections to be tracked. The effect of the connection tracking setting in the Access rules depends on the traffic type.

If Connection Tracking is on, you can also set the **Idle Timeout** for connections. The timeout is meant for clearing the engine's records of old connections that the communicating hosts leave hanging. The timeout concerns only idle connections, so connections are not cut because of timeouts while the hosts are still communicating. The timeout defined for an Access rule overrides the default idle timeout value that is set for the protocol in the engine's properties.



CAUTION

Setting excessively long timeouts for many connections can consume engine resources and degrade engine performance and stability. Be especially careful when defining timeouts for ICMP and UDP. The ICMP and UDP virtual connections do not have closing packets. The engine keeps the records for the ICMP and UDP connections until the end of the timeout.

Connection Tracking options in Access rules also allow you to override the limit for connections from a single source or destination IP address defined in the **Traffic Handling** settings for Security Engine and in the properties of some interface types. When the set number of connections is reached, the engine blocks the next connection attempts until a previously open connection is closed.

Changes in the Connection Tracking mode affect how existing connections are handled when you install or refresh the policy. When you install or refresh the policy on an engine, the engine checks if the existing connections are still allowed in the policy. If the connection tracking mode changes from Loose to Strict, existing virtual ICMP connections are only allowed if they began with a valid packet (for example, not with a response

packet). In addition, if the mode changes from Normal to Strict, existing TCP connections are only allowed if all packets in the connection have been seen. In all other cases, changes in connection tracking mode do not affect existing ICMP, TCP, and UDP connections at policy installation.

Policy Snapshots

A Policy Snapshot is a stored record of a policy configuration.

A Policy Snapshot is stored in an engine's upload history whenever a policy is installed or refreshed on the engine. The Policy Snapshots allow you to check which policies and other configuration information were uploaded, and when they were uploaded. You can also compare any two Policy Snapshots and see the differences between them.

Continue rules

The Continue action for a rule sets default options for traffic matching.

The Continue action for a rule sets default options (such as logging options and idle timeout) for the traffic matching process. Options set in Continue rules are used for subsequent rules that match the same criteria as the Continue rule, unless the rules are set to override the options. Continue rules are also useful in the hierarchical structure of the policies. Template Policies are convenient for setting options with a Continue rule, because all Policies and Template Policies that use the template inherit the option settings you have specified.

Related concepts

Configuring default settings for several Access rules on page 839

Adding comments to rules

You can add comments to rules to provide useful information to administrators and to make policies easier to read.

Each policy can include a large number of rules. Adding comments provides administrators with useful information and also makes it easier to read policies. You can add comments to all types of rules. In rule tables, you can add comments in the rule's **Comment** cell. You can also add a Rule Section, which begins with a comment row and can include one or more rules.

The Rule Section automatically includes all rules below the Rule Section until the next Rule Section in the policy. You can expand and collapse the Rule Sections as necessary. The comment row in a Rule Section is displayed against a colored background (with configurable colors). Rule Sections are useful in visually separating the sections of rules within a single policy.

Naming rules

You can specify an optional name or short description for Access rules, NAT rules, Ethernet rules, Exceptions in Inspection Policies, and rules in QoS Policies and Alert Policies.

Names help administrators identify individual rules in large rule tables. You can also search for a rule by its name. If a rule has been named, the name is displayed in the **Logs** view as well.

Deleting policies templates or sub-policies

If no policies are based on a policy template, you can delete the template. Similarly, you can delete a Sub-Policy as long as no Jump rules use it.

To delete a template, there must be no policies that are based on that template. The deletion is not allowed until you either delete the policies or switch them to use a different template.

To delete a Sub-Policy, there must be no policies that have Jump rules that use the Sub-Policy. The deletion is not allowed until you edit or remove the Jump rules in the policies.

Related concepts

How the Trash works on page 198

Related tasks

Change templates of policies on page 816

Engine Policy elements examples

Using Engine Policy elements, you can protect essential communications, make Engine rules more readable and Engine Policies more efficient, and manage network administrator rights.

Protecting essential communications example

You can make sure that essential communications are protected and cannot be cut off.

Company A has a engine system administered by multiple administrators of various degrees of familiarity with networking, engines, and Forcepoint Network Security Platforms. The administrators must often make quick changes to respond to the needs of the company and attend to any problems detected.

In this situation, it is possible that someone might accidentally change the Engine Policy in such a way that important services are cut off. The administrators decide to separate the rules allowing the most important business communications from rules that deal with non-essential traffic to minimize this risk. The administrators:

1) Create a Firewall Template Policy and select the predefined Firewall Template as the basis of the policy.

2) Cut and paste the rules allowing essential communications from their current Engine Policy into the new Firewall Template Policy.

   In this case, all administrators are allowed to edit the new Firewall Template Policy as well.

3) Add an insert point below the copied rules in the Firewall Template Policy.

   Having the insert point below the essential rules prevents the rules added to the inheriting Engine Policy from affecting the essential communications.

4) Re-parent their current Engine Policy to use the new template, moving it down a step in the policy hierarchy.

5) After validating the policy and making sure that the rules are correct, refresh the current Engine Policy.

   Most daily editing is done in the Engine Policy. There is less risk of someone accidentally changing the essential rules in the Firewall Template Policy, because the rules are not editable in the Engine Policy.

Improving readability and performance example

You can make Engine rules more readable and improve the performance of Engine Policies.

Company B has two separate DMZs, one for the extranet and one for other web services. The number of services offered is large. The company also has many partners and customers that have varying access rights to the different services. The administrators realize that many of the rules in their policies are related to the DMZ connections. The rest of the rules govern access to and from the company's internal networks. Many of the rules have been entered over time by inserting them at the beginning of the rule table, so rules governing access to the different networks are mixed. Finding all rules that govern access to a particular network takes time.

The administrators decide that they want to make their Engine Policy more readable and at the same time optimize the way the engine handles traffic, so they:

1) Create two new Engine Sub-Policies: one for each DMZ.

2) Cut and paste the rules from the current Engine Policy into the correct Engine Sub-Policy.

3) Add Jump rules to the Engine Policy, to direct the examination of traffic to/from the different networks to the correct Engine Sub-Policy.

4) Refresh the Engine Policy.

Restricting administrator editing rights in Engine Policies example

You can restrict the editing rights of network administrators in your organization, as needed.

Company C is implementing a distributed network with multiple sites: one central office where most of the administrators work, and several branch offices in different countries. The branch offices mostly have IT staff with only limited networking experience, but who are still responsible for the day-to-day maintenance of the network infrastructure at their site. They must be able to, for example, add and remove Access rules for testing purposes without always contacting the main administrators.

The administrators decide to limit the permissions of the branch office IT staff so that they are not able to edit the policies of the engines at any of the other sites. The administrators:

1) Create a Firewall Template Policy and select the predefined Firewall Template as the basis of the policy.

2) Add rules to the Firewall Template Policy using Alias elements to cover the essential services that each of these sites has, such as the VPN connections to the central site.

Using a common Firewall Template Policy for all branch offices also eliminates the need to make the same changes in several policies, easing the workload.

3) Create a Engine Policy based on the new Firewall Template Policy for each of the branch office sites.

   Although the same Engine Policy might work for all sites, in this case the administrators decide against it. Separate policies are needed for the separation of editing rights. The policies are based on the same Firewall Template Policy, so rules can still be shared without duplicating them manually.

4) Grant each Engine Policy to the correct Engine element.

   After this, only the correct policy can be installed on each engine. No other policy is accepted.

5) Create administrator accounts with restricted rights for the branch office administrators and grant the correct Engine element and Engine Policy to each administrator.
   - The branch office administrators are now restricted to editing one Engine Policy and can install it on the correct engine.
   - The branch office administrators are not allowed to edit the Firewall Template Policy the policy is based on. They also cannot install any other policies on any other engines.

IPS Policy example

You might want to restrict the policy editing rights of administrators to the IPS engines at their sites.

Restricting administrator editing rights in IPS Policies example

You can restrict the policy editing rights of administrators to the IPS engines at their sites.

Company A is implementing a distributed network with multiple sites: one central office where most of the administrators work, and several branch offices in different countries. The branch offices mostly have IT staff with only limited networking experience, but who are still responsible for the day-to-day maintenance of the network infrastructure and the IPS engines at their site. They must be able to, for example, add and remove Access rules for testing purposes without always contacting the main administrators.

The administrators decide to limit the permissions of the branch office IT staff so that they are not able to edit the policies of the IPS engines at any of the other sites. The administrators:

1) Create an IPS Template Policy based on the predefined IPS Template.

2) Add rules to the IPS Template Policy using Alias elements to cover the essential services that each of these sites have.

   Using a common IPS Template Policy for all branch offices eliminates the need to make the same changes in several policies, easing the workload.

3) Create an IPS Policy based on the new template for each of the branch office sites.

Although a single IPS Policy for all sites could work, in this case the administrators decide against it. Separate policies are needed for the separation of editing rights. The policies are based on the same template, so rules can still be shared without duplicating them manually.

4) Grant each IPS Policy to the correct IPS engine elements.

   After this, only the correct IPS Policy can be installed on each IPS engine. No other policy is accepted.

5) Create accounts with restricted rights for the branch office administrators and grant the correct IPS engine element and IPS Policy to each administrator.
   - The branch office administrators are now restricted to editing one IPS Policy and can install it on the correct IPS engine.
   - The branch office administrators are not allowed to edit the template the IPS Policy is based on. They also cannot install any other policies on any other IPS engines.

Local alternative policies

SMC administrators can now define up to three local alternative policies which can be activated in cases where the connectivity between the Security Engine and the Management Server is lost. The administrators can select whether the normal policy (one pushed from the Management Server) or one of the local alternative policies is active on the Security Engine.

This capability provides limited support for Security Engine backup use case where the centralized server may not be available and the local administrators may have a need to change the policy due to some tactical reason. However, there are few limitations with the current implementation:

- Only single engines are supported (Engine, IPS or Layer2 Engine).
- Master and virtual contexts are not supported.
- ECA Client configuration may not work properly when switching active policy.

Uploading policies as local alternative policies

Up to three local alternative policies can be uploaded from SMC to the Security Engine and not applied.

Use the steps below to upload policies to the Security Engine.

Steps

1) Select **Engine Configuration**.

2) Right-click the Security Engine, then select **Configuration** > **Install Policy**.

   The **Upload Policy Task Properties** dialog box displays:
   - In the **Policy** field, select the policy to be uploaded.
   - In the **Upload as** drop down, the administrator can select up to three **Local Alternative Policy**.

Upload Policy Task Properties dialog box

Use this dialog box to upload the selected policy to the selected engines.

Option	Definition
Elements	Shows the engines and Domains that you can refresh the policy on.
Search	Opens a search field for the selected element list.
Up	Navigates up one level in the navigation hierarchy. Not available at the top level of the navigation hierarchy.
More actions	■ **New** — Creates an element of the specified type. ■ **Show Deleted Elements** — Shows elements that have been moved to the Trash.
Add	Adds the selected elements to the **Target** list.
Remove	Removes the selected elements from the **Target** list.
Target	Shows the elements that you have selected.
Policy	Shows the policy to upload.
Select	Opens the **Select Element** dialog box.
Upload as	Specify whether to upload as an alternative policy or a normal policy.
Keep Previous Configuration Definitions	When selected, the previous configuration is kept.
Validate Policy before Upload	When selected, validation checks are performed.
Select Settings	Opens the **Validate Policy** dialog box.
Upload Comment	An optional comment for your own reference.

Activating an alternate policy

Activate one or all of the available local alternative policies on the Security Engine.

On the Security Engine run the following command to activate anyone of the local alternative policies:

Steps

1) `activate-alternative-policy`:

   `Enter the local alternative policy number:` Select the policy that need to be activated. Once the selected policy is activated, a confirmation message is displayed.

Monitoring for local policies

View the currently applied policies on the selected Security Engine on the **Info** view.

Steps

1) Select **Engine Configuration**.

2) On the Security Engine list, select a specific row.
   The **Info** view of each engine shows the installed local alternative policies.

Refresh local alternative policies

On policy refresh, you can also refresh local alternative policies.

Steps

1) Select **Engine Configuration**.

2) Right-click the Security Engine, then select **Current Policy** > **Refresh**.
   The **Refresh Policy Task Properties** dialog box displays:
   - Select the **Refresh policies on alternative slots** checkbox.
     When this check box is selected, upon policy refresh, the local alternative policies will also get refreshed.

Reverting to normal policy

Revert to the latest policy uploaded from the Management Server from the Security Engine command line.

On the Security Engine run the following command to revert to the latest policy uploaded from the Management Server:

Steps

1) `restore-normal-policy`
   Upon revert to the last policy uploaded from the Management Server, a confirmation message is displayed after reverting to the policy that was updated from the Management Server.

Removing local policies

Policies can be removed from either SMC or the Security Engine command line.

From the SMC, do the steps below to remove a local alternative policy. To remove the local policy from the Security Engine command line, run the `remove-alternative-policy`.

Steps

1) Select **Engine Configuration**.

2) Right-click the Security Engine, then select **Commands** > **Remove all Alternative Policies**.

Chapter 48

Ethernet rules

Contents

- Getting started with Ethernet rules on page 825
- Configuration of Ethernet rules on page 826
- Examples of Ethernet rules on page 829

Getting started with Ethernet rules

Ethernet rules are lists of matching criteria and actions that define whether Ethernet protocol traffic is allowed or discarded.

Ethernet rules are used by IPS engines, Layer 2 Engines, and layer 2 physical interfaces on Engines.

The traffic matching in Ethernet rules is based on the Source and Destination MAC Address in the packets. Any Ethernet network traffic, such as ARP, RARP, IPv6, Cisco Discovery Protocol (CDP), and Spanning Tree Protocol (STP), can be checked against the Ethernet rules. Ethernet traffic can be allowed or discarded. Regardless of the action taken, a matching rule can also create a log or alert entry.

The following types of interfaces can stop traffic when the Discard action is used:

- Inline IPS Interfaces on Engines
- Inline Layer 2 Engine Interfaces on Engines
- Inline Interfaces on IPS engines
- Inline Interfaces on Layer 2 Engines

For the following types of interfaces, only the Allow action is available:

- Capture Interfaces on Engines
- Capture Interfaces on IPS engines
- Capture Interfaces on Layer 2 Engines

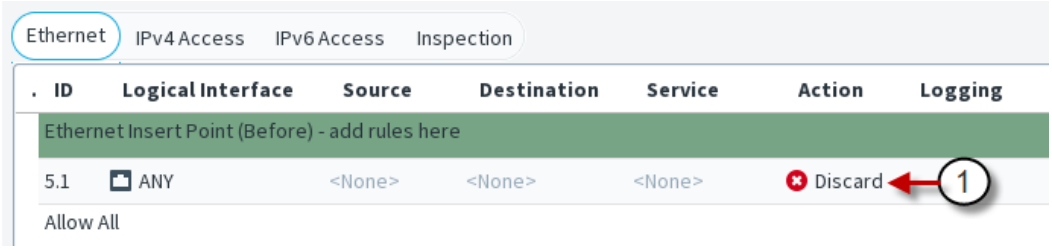
If your policy is based on the IPS Template or the Layer 2 Firewall Template, the Ethernet rules direct IPv4 and IPv6 traffic to the Inspection Policy for inspection, and let ARP, RARP, and STP traffic through. You can use the first Insert Point in the template to make exceptions to this behavior for certain MAC addresses or Logical Interfaces. We recommend that you insert any other changes at the second insert point.

Make sure that your Ethernet rules direct IP traffic for inspection against Access rules by applying the default IPv4 and IPv6 Services to traffic. When traffic does not match any Ethernet rule, the traffic is let through without further inspection.

Configuration of Ethernet rules

You can use Ethernet rules in IPS, Layer 2 Engine, and Layer 2 Interface Policy and Template Policy elements. Sub-Policies cannot contain Ethernet rules.

Newly inserted Ethernet rule - Main cells

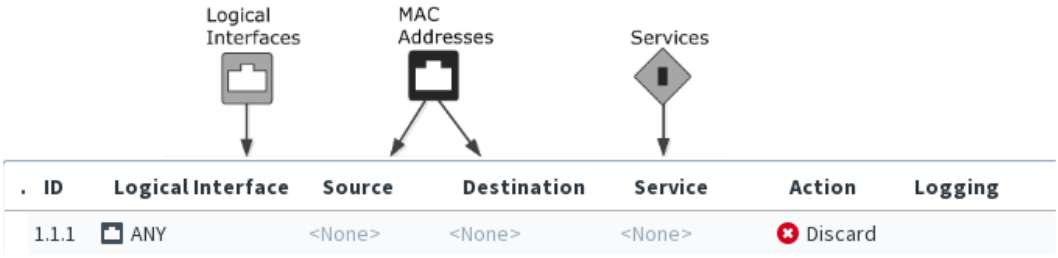


1 Security Engine applies action when it finds a match

This illustration shows an Ethernet rule that has been inserted into the policy. The Source, Destination, and Service cells are set to NONE, so this rule does not until they are changed to ANY or a more specific value. The Logical Interface cell is also matched against traffic, but it is not mandatory to change it if you want the rule to apply regardless of the interface. The other editable cells specify further conditions and options for handling matching connections. It is not necessary to edit all cells in each rule. Ethernet rules are read from the top down. More specific rules must be placed above more general rules that match the same traffic.

The illustration below shows the types of elements that you can use in Ethernet rules.

Elements in Ethernet rules



Cells in Ethernet rules

Cell	Description
Logical Interface	Accepts Logical Interface elements. Traffic matches the rule based on which interface the traffic is picked up from.
Source	Accept MAC address elements. You can set these cells to ANY to make the rule match all possible source or destination MAC addresses. You can also add more than one element in each cell to make the rule match multiple MAC addresses. The source and destination MAC addresses specified in an Ethernet rule are compared to the MAC addresses in each packet's header. Based on these and other criteria, the rule is applied to matching packets.
Destination	

Cell	Description
Service	Accepts Ethernet Services elements. Defines which protocols the Ethernet rule applies to. You can set the Service to ANY to make the rule match all protocols.
Action	Defines what happens when a packet matches the Ethernet rule. The following actions are available: ■ **Allow** — The traffic is let through the engine. ■ **Discard** — The traffic is silently dropped if going through an Inline interface. These actions stop the processing from continuing down the rule table for any packet that matches the rule. Rules with these actions must be placed so that the more limited the rule is in scope, the higher up in the rule table it is. If the traffic does not match any of the Ethernet rules by the end of the policy, the final action is applied. **Note** When the Allow action is used for IPv4 or IPv6 traffic in the Ethernet rules, the traffic is then checked against the IPv4 or the IPv6 Access rules. The final action for the IPv4 and IPv6 traffic is determined according to traffic type by the IPv4 or the IPv6 Access rules.

Default elements for Ethernet rules

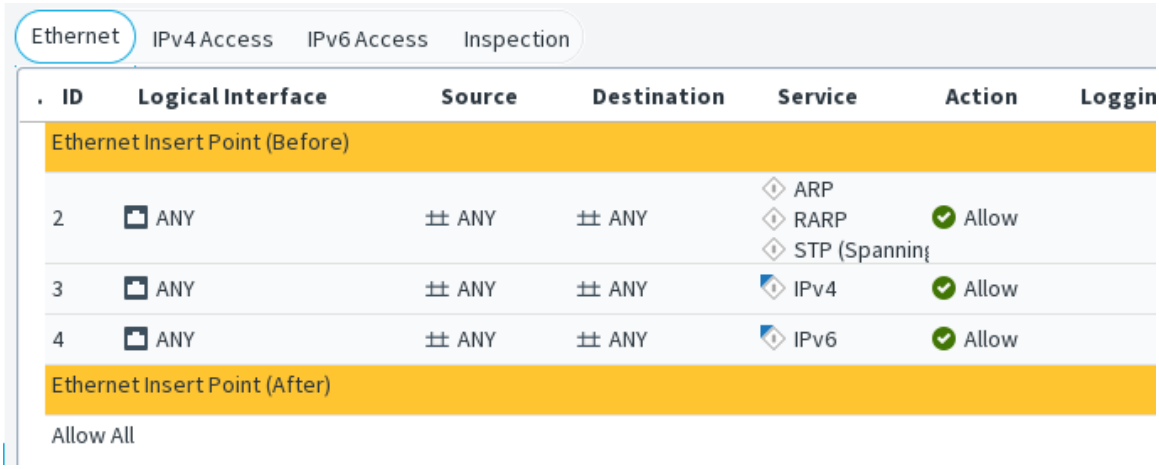
The predefined Ethernet rules in default template policies allow the most common types of Ethernet traffic.

The following default template policies contain predefined Ethernet rules:

- IPS Template
- Layer 2 Firewall Template
- Layer 2 Interface Template

Because the template policies are added and updated through dynamic update packages, your templates might look different from the example here.

IPS template - Ethernet rules



This illustration shows a green insert point at the top of the rule table, three default rules below it, and then another insert point.

- The first rule contains the IPv4 protocol and allows the matching traffic to pass through.
- The second rule contains the IPv4 protocol and allows IPv4 traffic with further inspection against the IPv4 Access rules.
- The third rule contains the IPv6 protocol and allows IPv6 traffic with further inspection against the IPv6Access rules.

The two insert points indicate where you can add Ethernet rules to a policy that uses the template policy. The first insert point above the default rules allows you to make exceptions to how traffic that matches the three default rules is checked. For example, you could add a rule defining that no IPv4 or IPv6 traffic is allowed between certain MAC addresses.

The second insert point below the default rules allows you to define how traffic that matches other protocols is checked. The final action depends on the type of template policy.

- IPS Template — Allow all
- Layer 2 Firewall Template — Discard All
- Layer 2 Interface Policy — Discard All for Inline Layer 2 Engine Interfaces. Allow all for Capture Interfaces and Inline IPS Interfaces.

Examples of Ethernet rules

These examples illustrate some common modifications to the default Ethernet rules and general steps on how each example is configured.

Example: Logging protocol use in Ethernet rules

An example of configuring Ethernet rules to log the use of Ethernet protocols.

The administrators at Company A have installed an IPS engine in Transparent Access Control mode and they want to create some custom Ethernet rules. The administrators know that most traffic uses the IPv4 protocol, but they are not sure which other Ethernet protocols are being used in the company's network. They decide to temporarily log the use of Ethernet protocols, excluding IPv4.

To log all Ethernet protocol traffic excluding IPv4, the administrators:

1) Create an IPS Policy based on the IPS Template.

2) Add a rule in the Ethernet rules to exclude IPv4 traffic from logging:

Ethernet rule for excluding IPv4 traffic from logging

Source	Destination	Service	Action	Options
ANY	ANY	IPv4	Allow	Logging: None

3) Add a rule to log the use of other Ethernet protocols:

Ethernet rule for logging Ethernet protocol use

Source	Destination	Service	Action	Options
ANY	ANY	ANY	Allow	Logging: Stored

4) Save and install the policy on the IPS engine.

5) View the logs generated by the matches to the Ethernet rules in the Logs view.

6) Disable the logging Ethernet rule to prevent excess log data from being generated.

Example: restricting the use of Ethernet protocols

An example of configuring Ethernet rules to restrict which Ethernet protocols are allowed.

Now that the administrators at Company A from the previous example have a clear picture of which Ethernet protocols are being used, they want to restrict allowed protocols. The administrators determine that ARP and

Spanning Tree Protocol (STP) must be allowed. Because most traffic will use these protocols, the administrators do not want to log matches to the rules that allow specific protocols.

They decide to block the Cisco Discovery Protocol (CDP) on the logical interface named Inline Interface because of security problems, and log detected CDP use.

To block the use of the CDP protocol, the administrators:

1) Add a new rule in the Ethernet rules to allow ARP, Spanning Tree Protocol (STP), and IPv4 without producing any logs:

Ethernet rule for allowing ARP and STP use

Logical Interface	Source	Destination	Service	Action	Options
ANY	ANY	ANY	ARP, STP, IPv4	Allow	Logging: None

2) Add another rule to block the use of Cisco Discovery Protocol (CDP) on the Inline Interface, and produce logs that will be stored:

Ethernet rule for blocking CDP use

Logical Interface	Source	Destination	Service	Action	Options
Inline Interface	ANY	ANY	CDP	Discard	Logging: Stored

3) Add a rule on the last line of the Ethernet rules to block the use of other Ethernet protocols without producing logs:

Ethernet rule for blocking other Ethernet protocols

Logical Interface	Source	Destination	Service	Action	Options
ANY	ANY	ANY	ANY	Discard	Logging: None

4) Save and install the policy on the IPS engine.

Chapter 49

Access rules

Contents

- Getting started with Access rules on page 831
- Overview of Access rules on page 831
- Configuring Access rules on page 833
- Using Access rules on page 838
- Examples of engine Access rules on page 845
- Examples of IPS Access rules on page 848

Access rules are lists of matching criteria and actions that define how the engine treats different types of network traffic. They are your main configuration tool for defining which traffic is stopped and which traffic is allowed.

Getting started with Access rules

Access rules filter traffic by defining matching criteria and an action that is applied to packets that match all criteria defined in the rule.

Access rules are used by Engines, IPS engines, Layer 2 Engines, Master Engines, Virtual Engines, Virtual IPS engines, and Virtual Layer 2 Engines:

- In Engine and Layer 2 Engine policies, the Access rules are the most important type of rules. The criteria you define in the Access rules determines which connections are allowed. By default, Engine and Layer 2 Engine Access rules stop traffic that you do not specifically allow.

  Note

  Master Engines always use Engine Policies regardless of the role of the Virtual Security Engines they host. Virtual Engines use Engine Policies. Virtual IPS engines use IPS policies. Virtual Layer 2 Engines use Layer 2 Engine Policies.

- In IPS policies, Access rules can be used to optionally filter out some traffic and to exclude some traffic from further inspection. Only traffic on Inline Interfaces can be filtered with Access rules. IPS engines allow all traffic that you do not specifically deny. For IPS policies based on the IPS Template, all traffic allowed by rules placed after the inherited rules is inspected against the Inspection Policy by default.

Overview of Access rules

Access rules are traffic handling rules in which you define the details of traffic examination and which action to take when matching details are found.

The IPv4 and IPv6 Access rules are stored in policy elements.

Access rules apply to all network interfaces, unless you use Zone elements to match traffic based on which interfaces traffic passes through.

The traffic matching is based on the information contained in the packets:

- Source and destination IP addresses.
- Protocol-specific information, such as the port information for protocols that use ports.
- Payload data in the packets, such as HTTP requests in an HTTP connection.

Additional matching criteria that is not based on information in the packets includes:

- The interface the traffic is coming from or going to. This allows you to restrict which traffic is allowed through which interfaces in more detail than basic antispoofing.
- (Engines only) The VPN the traffic is coming from (on an engine where that VPN terminates). This criteria allows creating rules that apply to VPN traffic only, or rules that apply to all traffic except VPN traffic.
- (Engines only) User authentication. This criteria allows you to create rules that define the end users who are allowed to make connections and the authentication methods for the end users.
- The User or User Group of a user who has logged on to an integrated Microsoft Active Directory domain (allowing you to create user-specific rules without configuring authentication).
- The day of the week and the time of day (allowing you to enforce rules only during certain times, such as working hours).

The Access rules provide several different ways to react when some traffic is found to match a rule. You can:

- Specifically allow the traffic.
- (Engines only) Allow the traffic on the condition that the user has passed authentication.
- (Engines only) Allow the traffic on the condition that a VPN is established.
- (Engines only) Allow the traffic on the condition that the same source or destination IP address does not have an excessive number of connections already open (concurrent connection limit).
- Allow the traffic with inspection against the Inspection Policy.
- Allow the traffic without further inspection.
- (Engines and inline interfaces only) Specifically stop the traffic.

Regardless of which of the above actions is taken, a matching rule can also create a log or alert entry.

In addition to traffic allowed by the Access rules, Engines allow the following types of traffic:

- Traffic that is allowed automatically based on the Security Engine configuration or by static rules generated by the Management Server.
- Traffic that is allowed by automatic rules for traffic to and from the engine.
- Traffic that is allowed by the default template on which the Engine Policy is based.

Engines automatically allow the following types of traffic with specific configurations:

- DHCP requests and replies for an interface for which a DHCP server is enabled.
- DHCP requests and replies for an interface that has a dynamic IP address.
- State synchronization between cluster nodes.
- IPv6 Neighbor Discovery traffic.

Related concepts

Getting started with policies on page 799

Configuring Access rules

Access rules filter traffic by defining matching criteria and an action that is applied to packets that match all criteria defined in the rule.

IPv4 Access rules are configured on the **IPv4 Access** tab, and IPv6 Access rules are configured on the **IPv6 Access** tab inside the following elements:

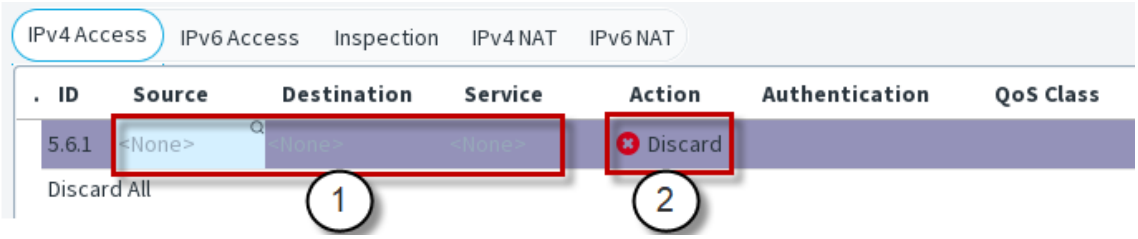
- Engine Policy
- IPS Policy
- Layer 2 Engine Policy
- Layer 2 Interface Policy
- Template Policy
- Sub-Policy

You can create new Access rules in the Policy Editing View. You can also create IPS and Layer 2 Engine Access rules in the Logs view. Use one or more selected log entries to create these rules (only available for IPS and Layer 2 Engine IPv6 Access rules, not for Engine IPv6 Access rules).

Before starting to build policies, make sure that you understand the network element types available and how you can use them efficiently to define the resources that you want to protect and control.

Configuring Engine Access rules

Newly inserted Engine IPv4 Access Rule - Main cells



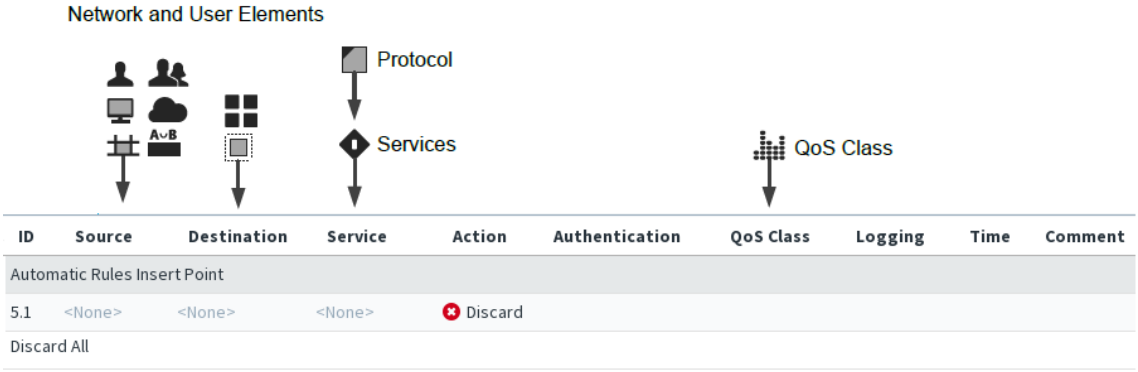
1 Mandatory cells for matching traffic

2 Engine applies this Action when it finds a match

This illustration shows an Access rule that has been inserted into the policy. The matching cells are set to **<None>** and the action is set to **Discard**. These settings prevent the rule from having any affecting in case a new rule is added to the policy accidentally. It is not necessary to edit all cells in each rule. However, the mandatory cells for traffic matching (**Source**, **Destination**, and **Protocol**) must be set to some value other than **<None>** for the rule to be valid. The **Source VPN** cell is also matched against traffic in the inspection process, but it can be left empty to match all traffic. The other editable cells specify further conditions and options, such as logging.

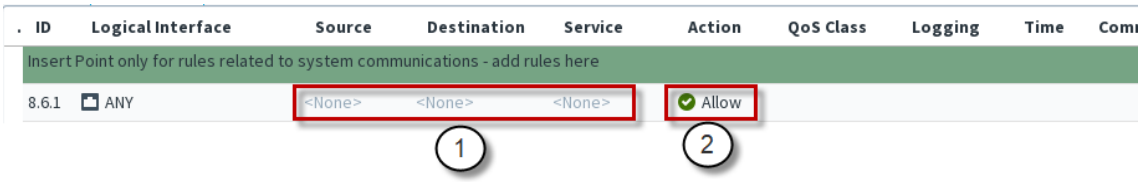
The following illustration shows the types of elements that you can use in IPv4 and IPv6 Access rules.

Elements in Engine Access rules



Configuring IPS and Layer 2 Engine Access rules

Newly inserted IPS or Layer 2 Engine Access Rule - Main cells



1 Mandatory cells for matching traffic

2 Engine applies this Action when it finds a match

This illustration shows an Access rule that has been inserted into the policy. The matching cells are set to **<None>** to prevent the rule from affecting traffic in case a new rule is added to the policy accidentally. It is not necessary to edit all cells in each rule, but the mandatory cells for traffic matching (**Source**, **Destination**, and **Service**) must be set to some value other than **<None>** for the rule to be valid. The **Logical Interface** cell is also matched against traffic, but it is not mandatory to change its value if you want the rule to apply regardless of the interface. The other editable cells specify further conditions and options for handling connections that match the cells that are used for traffic matching.

The following illustration shows the types of elements that you can use in IPv4 and IPv6 Access rules.

Elements in IPS and Layer 2 Engine Access rules



Related concepts

Getting started with directory servers on page 1103
Getting started with user authentication on page 1127
Overview of external user authentication on page 1130
Bandwidth management on page 974
Traffic prioritization on page 975

How Access rules match traffic

Access rules match traffic based on the Source, Destination, and Service cells. You can also specify other optional matching criteria.

The Source and Destination cells specify the IP addresses that are compared to the IP addresses in each packet's header. Based on these and other criteria, the rule is applied to matching packets.

The Service cell defines which protocols the Access rule applies to. The Service also determines the applications protocol used in the Inspection Policy for matching traffic (the protocol that is detected and selected for traffic by an Access rule is a matching criteria in the Inspection Policy). By default, the Service is set to <None>, and you must change the value to make the rule valid.

In addition to more specific matching criteria, the matching cells can be set to two more settings:

- ANY (available by right-clicking in a cell and selecting Set to ANY) matches all valid values for the cell, for example, all IPv4 addresses.
- NONE is the default value for mandatory traffic matching cells that have no matching criteria in them. As long as any cell in a rule contains NONE, the whole rule is invalid and is ignored.

Using Zones in the Destination cell of Access rules

Due to the processing order of Access and NAT rules, the interface through which the packet will be sent out is not yet determined when Access and NAT rules are processed. During the matching against Access and NAT rules, the destination Zone is matched based on the current routing decision for the packet. NAT and VPN operations can change the route that is used when the packet is sent out. Because of this possibility, the packet is checked against the Access rules again before being forwarded. If the changed destination Zone still matches,

traffic is processed according to the original rule. If the changed destination Zone does not match the Access rule, the traffic is discarded. Carefully consider how the rules will be applied when using Zones in the Destination Cell of Access rules when NAT and VPN operations can change the routing decision.

To define how an Access rule matches traffic, fill in the cells with elements.

Related concepts

Types of traffic inspection interfaces for IPS engines on page 579
Getting started with directory servers on page 1103
Getting started with Service elements on page 927
Getting started with user authentication on page 1127

Related tasks

Example VPN configuration 1: create Access rules on page 1226
Define Source, Destination, and Service criteria in rules on page 891
Define Authentication options for Engine Access rules on page 902
Specify rule validity times on page 912

Considerations for designing Access rules

One of the crucial issues in designing policies is the order of the rules.

Rules are read from the top down. The actions **Allow**, **Refuse**, and **Discard** stop the processing from continuing down the rule table for any connection that matches the rule. You must place rules with any of these actions so that the more limited the rule is in scope, the higher up in the rule table it is.

Example: An Access rule that allows connections to the IP address 192.168.10.200 must be put above an Access rule that stops all connections to the network 192.168.10.0/24.

In Engine and Layer 2 Engine policies, traffic that does not match any of the Access rules by the end of the policy is discarded by default. In IPS policies, traffic that does not match any of the Access rules by the end of the policy is allowed by default. In Layer 2 Interface Policies, the final action depends on the type of interface. Inline Layer 2 Engine Interfaces discard all traffic. Capture Interfaces and Inline IPS Interfaces allow all traffic.

Related concepts

Example: exempting traffic from inspection in IPS Access rules on page 848

Default elements for Access rules

There are several predefined elements for configuring Access rules.

Default elements for Engine Access rules

There are two predefined Firewall Template Policies called Firewall Template and Firewall Inspection Template. They contain the basic Access rules that allow communications between the engine on which the policy is installed and other SMC components. The Firewall Inspection Template is based on the Firewall Template.

You must use one of the predefined Firewall Template Policies as the basis for defining your own templates and policies. It is not possible to create a new template at the highest level of the policy hierarchy. No changes can be made directly to the predefined Firewall Template Policies. However, you can create your own copies of the predefined Firewall Template Policies if you have a specific need to edit the rules in them.

Note

If you use a copy of a predefined Firewall Template Policy, you may have to adjust your rules manually when the system is upgraded to account for changes in system communications. Upgrades can change only the predefined Firewall Template Policies, not the copies.

There is a yellow row near the end of the list of rules on the IPv4 Access and IPv6 Access tabs in the predefined Firewall Template Policies. The yellow row marks the insert point, where rules can be added in the inheriting Engine Policy and Firewall Template Policy elements.

The rules above the insert point detail the various kinds of system communications. Most of the IP addresses are defined using Aliases to make the rules applicable on any system where they are installed. These Aliases are default elements. The Service cell is the best starting point for understanding in greater detail what these rules do.

There are two rules below the insert point. The rule directly below the insert point has the action **Refuse** for the Ident protocol traffic. This action stops the traffic and sends an ICMP error message to the Ident request sender. This rule exists to prevent Ident requests from being silently dropped (as the next rule specifies for all other traffic). Silently dropping Ident requests might cause delays in legacy environments where the Ident protocol is used. The last rule before the end of the policy is a rule that discards all traffic and creates a log entry that is stored. This rule's purpose is to make sure that this connection dropping is logged. This rule is important because the engine silently drops the connection without creating a log entry if the matching process reaches the end of the policy.

Default elements for IPS and Layer 2 Engine Access rules

The IPS Template and the Layer 2 Firewall Template have predefined Ethernet rules, IPv4, and IPv6 Access rules. Because the default policy elements are introduced when you import and activate a recent dynamic update package, the templates you currently have in your SMC might look slightly different from the one that is presented in this section. Newer versions of the templates work in the same way as described below. Any changes to the templates are documented in the Release Notes document for each dynamic update package.

There are several IPv4 and IPv6 Access rules with various Services defined with Continue as the action and a yellow insert point indicating the place where a Policy that uses the template can be edited.

The Access rules that you add at the insert points in custom policies based on the IPS Template or the Layer 2 Firewall Template are usually specific exceptions to the rules inherited from the template. For example, you could insert a rule there that allows a connection between two particular hosts to continue without any further inspection. Or, you could instert a rule for inline IPS engines to always stop traffic between particular IP addresses and ports.

Related reference

Forcepoint Security Management Center ports on page 1457
Security Engine ports on page 1460

Restricting Engine Access rule matches based on the source VPN

You can match Engine Access rules based on whether the traffic is coming from a particular policy-based VPN.

You can define that the rule matches only non-VPN traffic, or only traffic from a particular policy-based VPN.

Using Access rules

There are some additional concepts that are useful when working with Access rules.

Related concepts

Using policy elements and rules on page 816

Allowing system communications in Access rules

You must add Access rules for some types of communication between SMC components.

The necessary communications between the engine and other SMC components are allowed in the predefined Firewall Template Policy, IPS Template, and Layer 2 Firewall Template. However, the predefined templates do not allow other SMC components to communicate through the engine to some third SMC component.

For example, when you have a engine and a Log Server at a remote site that are managed by a Management Server behind a engine at a central site, you must create rules in the Engine Policy at the central site to allow:

- Management and monitoring connections to/from the remote engine.
- Monitoring and log browsing connections from the central site to the remote Log Server.
- Any remote-site SMC Client connections to the Management Server at the central site.

If NAT is applied to the connections, Access rules alone are not enough. You must also create Location elements and add Contact Addresses for the elements to define which translated addresses are necessary for making contact.

If you have inline IPS engines or Layer 2 Engines, be careful that you do not define rules that would prevent other SMC components from communicating with each other.

There are predefined Service elements for all system communications. You can use these elements to create Access rules.

Related concepts

NAT and system communications on page 856

Related reference

Forcepoint Security Management Center ports on page 1457
Security Engine ports on page 1460

Configuring default settings for several Access rules

The **Continue** action allows you to set default values for some settings in rules to avoid defining the same settings for several rules individually.

When a connection matches a rule with Continue as the action, some of the rule's settings are written in memory but the matching continues until another rule that matches is found. This matching rule uses the defaults set in the Continue rule *unless* the rule specifically overrides the defaults with different settings. This way, you do not have to define the settings for each rule separately.

You can use Continue rules to set default settings for a general type of traffic and define settings for individual rules only when required. There are also default values that are used for rules that are set to use the values of a Continue rule, but there is no previous matching Continue rule.

The options that can be set using Continue rules in Access rules include:

- The Connection Tracking option:
  - For Engines, the default is on.
  - Idle Time-out also overrides the global defaults set in the engine's properties.
  - The concurrent connection limits define the maximum number of connections allowed from a single source or destination IP address.
- The logging options (for Engines, the default is **Stored**).
- The Protocol option inside the Service used (for Engines, this option is used to apply a Protocol to rules with ANY as their Service).
- The QoS Class (default is that no specific QoS Class is assigned).

Continue rules are defined the same way as other rules. However, when any of the options listed above is set in the Continue rule, many or even all rules below can be affected. The Continue rule options are used by the rules below if the source, destination, service port, and the optional source VPN match the same connection as the Continue rule. Continue rules are inherited from Template Policies into lower-level templates and policies like any other rules.

Continue rules behave in the same way as any other rules. A Continue rule can be overridden by:

- A later Continue rule that has an identical scope (such as Source and Destination)
- Partially overridden by a Continue rule that partially overlaps with the previous Continue rule
- A rule with the Allow action that has an identical scope and specifies different settings.

When you define Continue rules with different matching criteria, you can have several Continue rules one after another without them interfering with each other in any way at all.

Sub-Policies might require special attention with Continue rules: the Sub-Policies can have different options when you insert them into different policies if the Sub-Policy rules do not override the options set by preceding Continue rules. Also, when a Sub-Policy contains a Continue rule, the options are then used for further matching in the higher-level policy (if the processing returns to the higher-level policy).

Using Continue rules to set Access rule logging options

Instead of setting the log level for all rules individually, you can set a Continue rule in a template or in a policy to set the default log level.

One common use for the Continue action is to set the default log level for all subsequent rules. The log level for any subsequent matching rules can be left undefined. The rules trigger logging as defined in the Continue rule.

Setting the default log level



The default log level is set to Transient for any source, destination, or service. All subsequent rules in this policy and any sub-policies log Transient by default. Individual rules can still override this option with specific log levels, such as Essential or Stored.

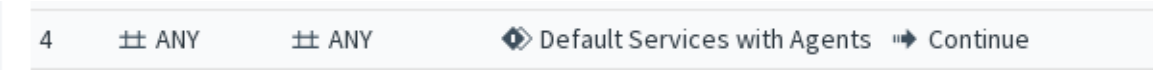
If logging is not defined for a rule and there is no prior Continue rule that sets logging options, the default log level is **Stored**.

Using Continue rules to set the Protocol in Access rules

The default Protocol can be set using the Continue action.

This way, the correct Protocol is also used for traffic that is allowed by rules that are set to match any Service. These rules have no particular Service element that would set the correct protocol). This setting can even be mandatory, for example, if you want to allow certain protocols that allocate ports dynamically. On Engines, the Protocol is needed to associate the traffic with the correct protocol for further inspection and to handle some types of traffic, such as FTP, correctly. The IPS Template and the Layer 2 Firewall Template include several Continue rules that associate all traffic with Protocols according to standard ports. The Firewall Template includes one Continue rule that defines that Protocols of the type Protocol Agent are used for the Service Group **Default Services with Agents**.

Protocol Agent rule in Firewall Template



You can set a Protocol for more types of protocols or override the Default rule shown in the illustration above for some or all connections. Do this by placing one or more Continue rules at the top or some other suitable place in your own template or policy. The Source and Destination can be some specific addresses if you want to limit the scope of your Continue rules.

Note

A rule that matches the same source, destination, and service port as the Continue rule and specifies an Protocol of the types Protocol Agent, Protocol Tag, or SSM Proxy for the Service overrides the default set in the Continue rule. To avoid this limitation, do not add rules that specify these types of Protocols for the same matching criteria as the Continue rules.

If you have TCP and UDP services set up in your network under non-standard ports, the traffic might not be associated to the correct protocol. The traffic could therefore inspected at a more general (TCP or UDP) level. In this case, you can create your own custom Service and add it in your policy to have the traffic inspected with the correct protocol information. Only some protocols and some of their parameters are supported in the services that are used in IPS policies.

You can also add your own rules for the opposite purpose: to have some traffic not inspected as a particular protocol, but more generally as TCP or UDP traffic. In this case, you add a rule in your policy that includes the general TCP or UDP Service element from the **IP-proto** branch of the Services tree.

Using Continue rules with Sidewinder Proxies

When you use Continue rules to specify a default Protocol, including Sidewinder Proxies, rules later in the policy can override the defaults set in Continue rules.

These limitations are due to the way that rules are processed, and are not specific to Sidewinder Proxies.

If you use a Continue rule to specify a Sidewinder Proxy as a default Protocol, a rule later in the policy overrides the Continue rule that specifies the Sidewinder Proxy when:

- The rule matches the same source, destination, and service port as the Continue rule.
- The rule specifies a Protocol of the type Protocol Agent or Protocol Tag for the Service.

To avoid this limitation, do not add rules that specify a Protocol of the type Protocol Agent or Protocol Tag for the same matching criteria as the Continue rules for Sidewinder Proxies.

Example of a rule that overrides a Continue rule that specifies a Sidewinder Proxy

ID	Source	Destination	Service	Action
14.1	Internal network	External	SSM HTTP on port 80	Continue
14.7	Internal network	External	HTTP with Protocol Agent on port 80	Allow

In this example, the second rule overrides the defaults set in the Continue rule because it specifies a Protocol Agent for the same matching criteria as the Continue rule. HTTP traffic on port 80 from the internal network to external destinations matches the second rule. The traffic does not use the Sidewinder Proxy.

A rule later in the policy does not override the default Protocol set in the Continue rule when:

- The rule matches the same source, destination, and port as the Continue rule.
- The rule does not specify a Protocol of the type Protocol Agent or Protocol Tag for the Service. For example, the rule specifies a Service element without a Protocol Agent or Protocol Tag.

Example of a rule that does not override a Continue rule that specifies a Sidewinder Proxy

ID	Source	Destination	Service	Action
14.1	Internal network	External	SSM HTTP on port 80	Continue
14.7	Internal network	External	HTTP on port 80	Allow

In this example, the second rule does not override the first because it specifies a Service element without a Protocol Agent or Protocol Tag. Because there is no more specific rule for the same matching criteria, the traffic uses the Sidewinder Proxy specified in the Continue rule.

Rematching tunneled packets in Access rules

You can rematch encapsulated traffic against the Access rules.

If an engine inspects traffic that is tunneled using IP-in-IP tunneling or Generic Routing Encapsulation (GRE), the traffic can be checked against IPv4 or IPv6 Access rules several times. The number of checks depends on the number and type of layers in the tunnel.

For example, when an IPv4 datagram contains an IPv6 datagram, the IPv4 datagram is first matched according to Access rules. If the tunneling Service in the Access rule specifies that the encapsulated IPv6 datagram should be matched again, the contents are then matched against the IPv6 Access rules.

To limit the number of encapsulating layers, the engine properties define the maximum rematch count. By default, the maximum rematch count is 1. If this count is exceeded, the packet is allowed or discarded according to the setting specified in the engine properties and a log or an alert is generated.

Using Access rules for application routing

Access rules for application routing match based on the network application that is detected in the traffic.

When you use Access rules for application routing, you can select which VPN traffic uses depending on the network applications detected in the traffic. For example, you can:

- Route traffic from specific network applications through the local Internet connection, and route other business traffic through a VPN to a data center using another connection, such as MPLS.
- Direct all traffic related to a specific network application to one ISP connection, and reserve the other ISP connection for more important traffic.
  For example, you can direct YouTube traffic to a low-cost ISP connection, and direct business-critical traffic to a faster, but more expensive ISP connection.



Important

After the routing decision has been made, the Security Engine might later identify a different application in the connection. If the application that is detected would cause a different routing decision to be made, the connection might be discarded.

Using Access rules for application routing has the following limitations:

- You can only use Network Application elements that have the Application Routing tag.
- You cannot use rules that match based on the network application to apply Sidewinder Proxies to traffic.

Related concepts

Access rule matching based on the payload of connections on page 807

Enforcing safe search features in Access rules

The safe search feature helps schools and other organizations to limit web searches and filter out potentially offensive content from search results. If end users change the safe search settings in the browser, safe search is not disabled.



Note

To use the safe search feature, an Inspection Policy must be selected on the **Inspection** tab of the Engine Policy. If **No Inspection Policy** is selected, the safe search feature does not work.

You can enforce safe search in two ways:

DNS method

When you use the DNS (TCP with SafeSearch) or DNS (UDP with SafeSearch) services, the Security Engine redirects searches by changing the DNS response. You can see the supported search engines on the **Protocol Parameters** tab in the properties of the Service element. The supported search engines might be updated when a new dynamic update package is activated.

Note

The DNS method is not compatible with the DNS relay feature.

HTTP request modification method

When you use the HTTP (SafeSearch) or HTTPS (SafeSearch with decryption) services, the Security Engine modifies the search request by adding a restriction parameter to the search URL. For safe search enforcement to apply to HTTPS, TLS inspection must be configured, and you must enable decryption in the Access rule for HTTPS traffic. The Microsoft Bing and Yahoo search engines are supported.

Using Alias elements in Access rules

You can use **Alias** elements to create a single rule that changes in meaning depending on where it is installed.

Alias elements are one of the most useful tools for reducing the complexity of a policy. Alias elements are like variables in a mathematical equation—their value changes depending on the component on which they are installed. Because Alias elements are able to change their meaning to adapt to local contexts, they can be used to create a single rule. That rule then changes in meaning depending on where it is installed. With Alias elements, you can avoid creating multiple, near-duplicate rule sets when you have several engines. The Alias element is used like any other network element. However, the IP addresses that the Alias element represents depends on the engine where the rules are installed. The IP address to engine mapping is defined in the Alias element.

For example, a company has its headquarters in Helsinki and branch offices in Atlanta, Munich, Tokyo, and Montreal. Each office has its own web server. The web server rules could be put in a single Sub-Policy, but each location's web server has a different IP address. Normal rules would require allowing access to all IP addresses on all engines, which is not only unnecessary, but can also be a security risk. Using Alias elements, the company can create a single set of rules that are still valid when applied to multiple engines. These rules would not, however, allow access to IP addresses that are not in use on a particular engine.

The administrator of the example company can create a web server alias, `$WebServers`. In the Alias element's properties, the administrator defines what `$WebServers` means for each component. For the IPS engine in Helsinki, the web server would be defined as 192.168.1.101, for the IPS engine in Tokyo as 192.168.2.101, and so on.

When the administrator installs a policy containing the web server rules with the Alias element, the addresses are translated to the correct address on that component. Therefore, when the policy is installed on the Helsinki IPS engine, the Alias element translates to an IP address of 192.168.1.101. The other addresses are not included in the policy that is transferred to that particular engine.

In this way, Alias elements simplify policies without reducing security.

Creating user-specific Access rules

You can use User and User Group elements as the source or destination of a rule to create user-specific rules.

You can optionally use the Forcepoint User ID Service or the Integrated User ID Service with Forcepoint Network Security Platform to associate IP addresses with users in an Active Directory database. This makes it possible to use User and User Group elements as the source or destination of a rule to create user-specific rules without requiring user authentication. The Integrated User ID Service is primarily meant for demonstration purposes and proof-of-concept testing of user identification services.

Note

For Forcepoint Network Security Platform version 6.4 or higher, we recommend that you use the Forcepoint User ID Service.

User-specific rules do not replace user authentication; they are a tool to simplify the configuration of access control, and improve the end-user experience by allowing transparent access to services. They are intended to be used for trusted users in a trusted environment where strong authentication is not required. User-specific rules can be used together with user authentication rules to allow some user groups to access a service, while otherwise requiring authentication for the same service.

Using Domain Name elements in Engine Access rules

You can use Domain Name elements in Access rules to represent a fully qualified domain name (FQDN) that might be associated with multiple IP addresses.

If you have specified one or more DNS servers in the engine's properties, the engine periodically queries the DNS server to automatically resolve domain names to IP addresses. This makes it possible to create rules that are valid even if new addresses are added to the domain or the domain's IP addresses change. If the DNS server returns multiple IP addresses for the same domain name, the engine associates all the IP addresses with the domain name. However, if there are a large number IP addresses associated with the same domain name, the DNS server might only reply with a few of the IP addresses at a time. In this case, the engine might need to make more queries to the DNS server to resolve all the IP addresses for the domain name. By default, the engine queries the DNS server every six minutes. Resolved IP addresses are kept in the engine's DNS cache for a maximum of one hour by default.

Domain Name elements also enable use of custom dynamic elements in Access rules. FQDN domain suffix `.namedb.local` is recognized by Engine and handled with specific custom resolver scripts. For more information, see Knowledge Base article 33503

Note

The DNS cache is not synchronized between nodes of a cluster. Each node separately queries the DNS server using the node's NDI address. It is possible that the DNS cache can be different on different nodes of a cluster.

Using Zone elements for interface matching in Access rules

Access rules apply to all network interfaces, unless you use Zone elements to match traffic based on which interfaces traffic passes through.

Zone elements are interface references that can combine several network interfaces of an engine into one logical entity. Using Zones in the Source or Destination cells allows you to restrict traffic according to which interfaces the traffic passes through. Zones can be useful, for example, when a type of traffic is only valid it when it passes through a specific interface, but basic Anti-Spoofing allows the traffic on any interface.

Examples of engine Access rules

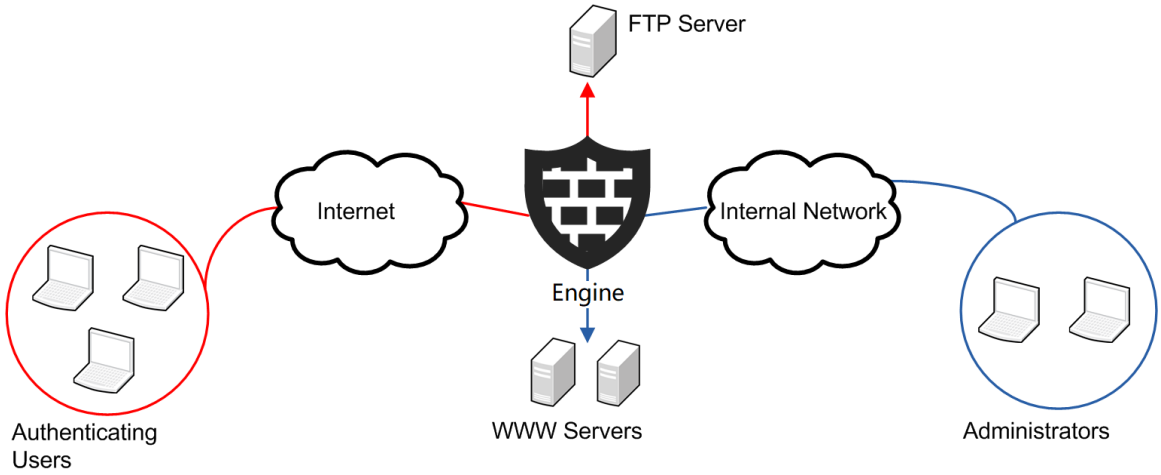
These examples illustrate some common uses for Engine Access rules and general steps on how each example is configured.

Example: engine Access rule order

An example of how Access rule order affects traffic matching.

Company A has an office network, a DMZ for WWW servers, and a second DMZ for an FTP server. The administrators only need to add rules for the DMZ traffic.

Company A's communication of special interest



The WWW servers must be accessible to anyone from both internal and external networks. HTTP traffic is inspected against the Inspection rules, excluding the administrators' own PCs (on the right in the illustration) because they often test the servers for vulnerabilities. The FTP server is accessible to all users in the general office network, but only to certain external users (on the left) that authenticate using an external authentication server.

The administrators:

1) Create Host elements for the WWW servers, the FTP server, and the administrators' PCs.

2) Create a Group element that contains the WWW server Host elements.

3) Create a Group element that contains the administrator PCs' Host elements.

4) Configure an external authentication server for use with the Engine.

5) Create User and User Group elements for the allowed external FTP users.

6) Add IPv4 Access rules with the Allow action for access to the DMZs:

Access rules for the DMZ

Source	Destination	Service	Authentication	Action
"Administrator PCs" Group	"WWW Servers" Group	"HTTP" Service		Allow (Deep Inspection Off)
ANY	"WWW Servers" Group	"HTTP" Service		Allow (Deep Inspection On)
Network element for Office Network	"FTP Server" Host	"FTP" Service		Allow (Deep Inspection Off)
ANY	"FTP Server" Host	"FTP" Service	**Users** tab: "External Users" User Group **Authentication Methods** tab: A suitable authentication method	Allow (Deep Inspection Off)

- As seen in the rule table, there are two rules for traffic to both the WWW servers and the FTP server.
- The rules are arranged so that the more specific rules are above the more general rules.
  For example, the rule allowing administrators to connect to the WWW servers without checking against the Inspection rules is above the more general rule. The general rule allows any connection to the servers as subject to the Inspection rules.
- If the first two rules were in the opposite order, the rule specific to administrators would never match, as the rule with the source as ANY would be applied first. The connection would be allowed according to that general rule, and the engine would stop checking the rule table.

Example: Continue rules in engine Access rules

An example of using Continue rules to set default options for several Engine Access rules.

Company B has decided to implement QoS Policies. The administrators want to set the QoS Class for traffic using a classification of high, medium, and low for all traffic depending on the sender. High priority is assigned to a few hosts in different networks, medium priority to one internal network, and low priority to all other hosts. The administrators want to follow how much traffic is allowed using the highest priority. Because of this, they also want to make sure that this traffic is logged with the accounting option turned on. They decide that the lower priorities of traffic don't need to be permanently logged at this point, so the administrators:

1) Configure the QoS features.

2) Create elements for all high-priority hosts.

3) Add the following Access rules to the top of their policy:

Continue rules for Logging and QoS Class

Source	Destination	Service	Action	Logging	QoS Class
Important Hosts	ANY	ANY	Continue	Stored with accounting	High priority
Network element for Important Network	ANY	ANY	Continue	Transient	Medium priority
All other Hosts	ANY	ANY	Continue	Transient	Low priority

After adding these rules, individual rules can override the settings as needed. However, most of the existing rules governing access from internal networks to the Internet now use the QoS Class and Logging options as set in these rules.

4) Transfer the policy to the engine.

Example: User-specific engine Access rules

An example of using access control by user to define Engine Access rules that only apply to specific users.

Company C has an existing Microsoft Active Directory server that it uses for user accounts in its Windows domain. Users are divided into groups according to the department they work in. The administrators have already integrated the Active Directory user database with the SMC to be able to view and manage Users in the SMC Client.

There is already an Access rule that blocks access to a video sharing site. However, the marketing team needs to be able to publish videos for its new marketing campaign on the site. The administrators want to allow users in the marketing group to access the site, but do not want to require user authentication.

Because the video sharing site has multiple servers with different IP addresses, the administrators decide to use a Domain Name element. This element dynamically resolves the IP addresses of servers in the video sharing site's Internet Domain.

The administrators:

1) Integrate a Forcepoint User ID Service server with Forcepoint Network Security Platform.

2) Add the following Access rule before the rule that blocks access to the video sharing site:

User-Specific Access Rule

Source	Destination	Service	Action
Marketing user group	Domain Name element that represents the video sharing site	■ HTTP ■ HTTPS	Allow

3) Install the policy on the engine.

Examples of IPS Access rules

These examples illustrate some common uses for IPS Access rules and general steps on how each example is configured.

Example: exempting traffic from inspection in IPS Access rules

An example of using IPS Access rules to allow specific traffic without deep inspection.

At Company A, there is an IPS engine deployed between the general office network and a subnetwork.

Company A's networks



In the subnetwork, there are several servers that provide services to the general office network as well as the Management Server and Log Server. There is also a Engine deployed between the internal and external networks. There is heavy traffic to the subnetwork where the internal servers are. The administrators decide to exempt the log transmissions between the Engine and the Log Server from being inspected against the Inspection Policy to reduce the IPS engine's workload. The administrators:

1) Create an IPS policy based on the IPS Template to replace the Default IPS Policy that they have currently installed.

2) Add a rule in the Access rules for their IPS engine:

Access Rule for exempting traffic from inspection against the Inspection policy

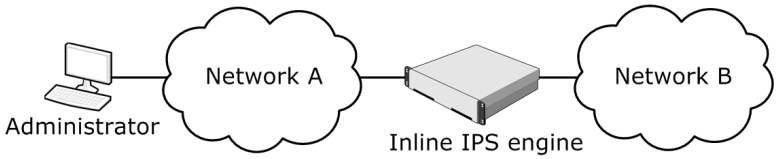
Source	Destination	Service	Action	Options
Engine	Log Server	SG Engine to Log	Allow	Deep inspection: Off

Example: using Access rules to filter traffic on an inline IPS engine

An example of using IPS Access rules to filter traffic between internal networks.

Administrators at company B decide that they want more control over which hosts and ports can be used between two networks.

Company B's network



Hosts in the two networks must be able to communicate between each other using certain specific ports. Also, one of the administrators has a workstation connected to Network A. The administrator's workstation must have unrestricted access to Network B. The administrators decide that the inline IPS engine provides an acceptable level of security between two internal networks.

The administrators:

1) Create elements for network A, network B, and administration host.

2) Add new Access rules for their inline IPS engine:

Access rules for filtering traffic

Source	Destination	Service	Action	Options
Administrator Network B	Administrator Network B	ANY	Allow	Logging: Undefined Deep inspection: On
Network A Network B	Network A Network B	Service elements for allowed services	Allow	Logging: Undefined Deep inspection: On
ANY	ANY	ANY	Allow	Logging: Stored Deep inspection: (irrelevant, because dropped traffic is never inspected further)

- Each of the first two rules allows traffic between the Source and the Destination in both directions. The order of the elements within the Source, Destination, and Service cells makes no difference to the outcome of the matching process.
- The order of the rules is important. The rules above proceed correctly from most specific to the least specific. The two first rules must be in this order, because the administrators want all connections from the Administrator host (which is in Network A) to always match the first rule and never the second one because the rules have different logging options.
- The last of the added rules stops all traffic that is not allowed in the rules above to prevent unauthorized traffic from passing.

Note

If the inline interfaces are on a fail-open network card, traffic passes freely whenever the IPS engine is offline regardless of what the Access rules state.

Chapter 50

NAT rules

Contents

- Getting started with NAT rules on page 851
- Configuring NAT rules on page 853
- NAT and system communications on page 856
- Outbound load-balancing NAT on page 859
- Proxy ARP and NAT on page 859
- Protocols and NAT on page 859
- Examples of NAT configuration on page 860
- Examples of NAT rules on page 862

Getting started with NAT rules

Network address translation (NAT) replaces the source or destination IP addresses in packets with other IP addresses. NAT rules define how NAT is applied to traffic.

NAT rules are matched to allowed connections after Access rule matching. NAT is applied before a routing decision is made, so the address translation can affect how the traffic is routed. NAT can be applied to IPv4 and IPv6 traffic. Engines, Master Engines, and Virtual Engines can use NAT rules.



Note

Master Security Engines always use Engine Policies regardless of the role of the Virtual Engines they host. Virtual Engines use Engine Policies.

In addition to manually configured NAT rules, you can also use element-based NAT to define how engines translate network IP addresses. The NAT rules generated from NAT definitions are applied only after the NAT rules that you have added manually to the policy. This means that the NAT rules that are generated from NAT definitions do not override the rules that you have manually added to the policy.

You can define the following types of NAT:

- Static source NAT — Typically translates the internal (“real”) IP address of an internal host to a different IP address in the external network.
- Static destination NAT — Typically translates the public IP address of an internal host to the private IP address, so that the host (server) can receive new connections from external hosts. Allows IP address or port translation (PAT), or both.
- A combination of both static source NAT and static destination NAT — Typically translates both the Source and Destination IP address in the same connection. Used, for example, to allow internal hosts to access your organization’s public servers using the public IP addresses of both the client and the server.
- Dynamic source NAT — Typically translates the internal IP addresses of several internal hosts to one or a few external IP addresses. Used to hide the internal network structure from outsiders and to avoid acquiring a separate public IP address for each of the hosts.

Note

Dynamic destination NAT can be configured for IPv4 traffic as part of the Server Pool feature, but not in NAT rules.

The following general guidelines apply when you add NAT rules:

- NAT rules only apply to connections that are handled statefully (Connection Tracking option is enabled in the Access rule that allows the connection).
- NAT rules are applied to whole connections. Reverse NAT for reply packets is automatic, so you do not need to define rules for replies within a connection.
- Connections are matched against NAT rules with the same type of matching criteria as other types of rules. The first matching NAT rule is applied and any other NAT rules are ignored. To prevent a NAT rule from matching some connections, create a NAT rule that specifies no translation for those connections and place it above the rule that matches.
- By default, NAT rules are ignored for traffic that enters or leaves a VPN tunnel. To match such traffic against NAT rules, enable NAT in the VPN Gateway element's properties.
- Routing decisions are made after NAT, so remember that translating the destination address can affect how the traffic is routed. If the translated IP addresses are not included in existing definitions, you might need to add the translated IP addresses to the Routing tree.
- If you install the Engine Policy with the **Keep Previous Configuration Definitions** option selected, previous NAT rules are kept until all currently open connections that use those rules are closed. In some cases, the old and the new rules can conflict and prevent policy installation until the option is deactivated.

Application routing

NAT rules for application routing match based on the network application that is detected in the traffic. When you use NAT rules for application routing, you can apply different NAT rules to traffic, and redirect traffic to different proxy servers depending on the network applications detected in the traffic. For example, you can:

- Exclude specific network applications from being redirected to proxy servers.
- Direct some network applications to one proxy server, and direct the rest of the traffic to another proxy server.

Note

The supported protocols depend on the proxy server to which traffic is forwarded.

When using NAT rules for application routing, you can only use Network Application elements that have the Application Routing tag.

Important

After the routing decision has been made, the Security Engine might later identify a different application in the connection. If the application that is detected would cause a different routing decision to be made, the connection might be discarded.

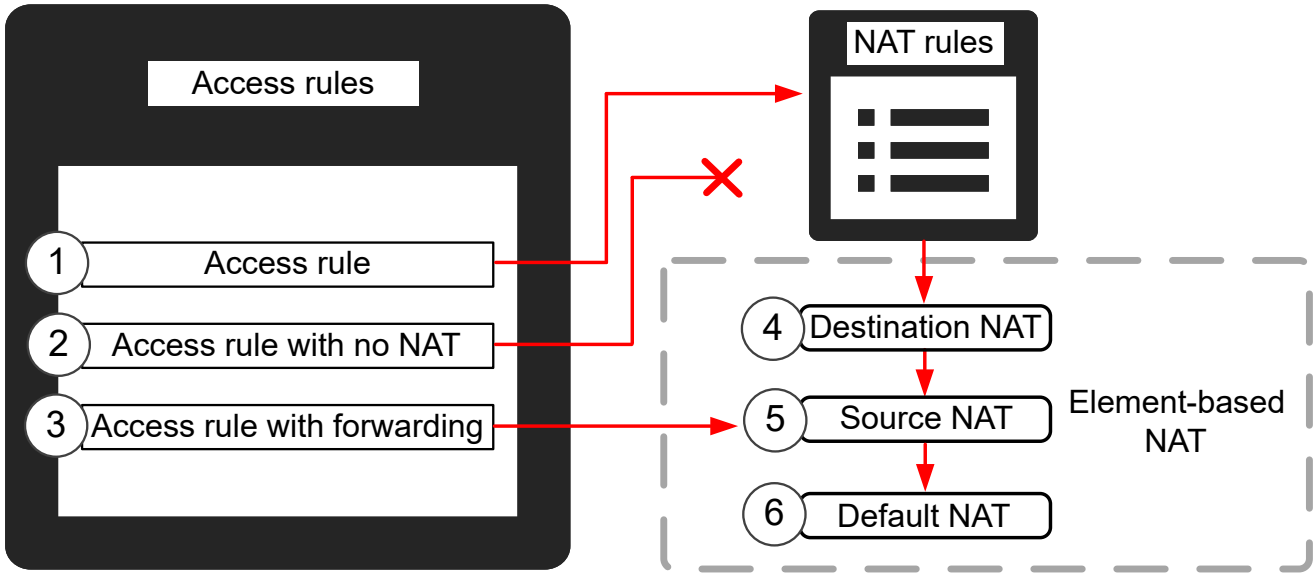
Interaction between Access rules and NAT

Access rules can have an impact on how NAT is applied to traffic.

The Access rules are always processed first. Depending on the type of Access rule, there can be the following processing behavior:

- Both NAT rules and element-based NAT definitions are processed.
- Both NAT rules and element-based NAT definitions are not processed.
- NAT rules are ignored, but some element-based NAT definitions are processed.

Access rules and NAT processes



1. After an Access rule is processed, it is processed by NAT rules. If there are no matches in the NAT rules, the processing continues with element-based NAT definitions, if they have been defined in the properties of the Security Engine.
2. If an Access rule has no NAT defined, such as an Access rule for Server Pool load balancing or for a policy-based VPN that does not have NAT applied to it, NAT rules and element-based NAT definitions are not processed.
3. If the Access rule forwards traffic to a proxy or a host, NAT rules and destination NAT definitions are not processed. Source NAT definitions and default NAT are processed.
4. If there are no matches in the NAT rules, the processing continues with element-based NAT definitions. Destination NAT definitions are processed first, then source NAT definitions are processed.
5. If there are no matches for source NAT definitions, the processing continues with default NAT.
6. If there are no matches for default NAT, NAT is not applied to the traffic.

Configuring NAT rules

Address translation is configured as part of the Engine Policy using NAT rules.

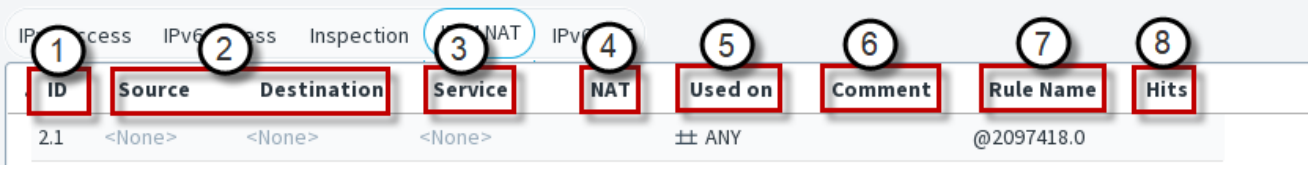
NAT rules are configured on the **IPv4 NAT** and **IPv6 NAT** tabs in Engine Policy and Firewall Template Policy elements. Engine Sub-Policies cannot contain NAT rules.

Note

NAT rules are applied only after a packet matches an Access rule and is allowed by the engine. The Access rule must have connection tracking enabled (default).

The following illustration shows a NAT rule that has just been inserted into a policy. The **Source**, **Destination**, and **Service** cells are set to **<None>** and they must be changed to something else for the rule to match any traffic. The **Used on** cell is also used for traffic matching: you can add specific Engine elements to this cell to make the rule valid only on those engines, or you can leave it to the default **ANY** to make the rule valid on all engines where the policy is installed. The columns are in the default order, but you can drag and drop them to your preferred order.

Newly inserted NAT rule



1. A non-editable ID that indicates the order of the rules in the policy. The rules are matched against traffic in order. For example, rule 4.3 is the third rule added in this Engine Policy element to the insert point that is the fourth NAT rule in the upper-level Template Policy element.
2. Matching criteria that defines the IP addresses and interfaces that the rule matches. You can use elements in the Network Elements category, as well as User and User Group elements, and elements related to endpoint information. You can also use Proxy Server elements in the Destination field. Both the Source and the Destination must match for the packet to match the rule. The addresses you insert must be valid. For example, static source address translation requires that the Source cell contains a single continuous IP address space.
3. Limits the scope of NAT rules to a specific protocol. You can use Service and Network Application elements.
4. The NAT that is applied to connections that match the rule. You can also set outbound load-balancing parameters. If left empty, address translation is not applied to matching connections. Use this to specify that NAT is not applied to matching connections, as an exception to other NAT rules.
5. Limits the scope of NAT rules based on the engine where the rule is installed. This takes into account different addressing in different networks when a shared policy is used on several different engines.
6. You can add a comment for your own reference. You can also add separate comment rows between rules.
7. The optional name or description for the rule is displayed before the rule tag. The non-editable tag is a unique identifier for the rule that links log entries to the rule that generated the log entries. The rule tag has two parts separated by a period. The first part is permanent and belongs to only that rule. The second part changes when the rule is changed.
8. Shows the number of connections that have matched the rule. Only shown if you run a Rule Counter Analysis. Shows “N/A” if there is no information available.

Related concepts

Outbound load-balancing NAT on page 859

How NAT rules match traffic

NAT rules are matched based on IP addresses and services.

Each address translation operation places specific restrictions on what you can put in the cells. Consider the design of your NAT rules and Access rules separately, because attempting to match the different types of rules one-to-one is usually not effective or even possible.

To overwrite both the source and destination IP address in the same packet (for example, to achieve hairpin NAT), configure both address translations in the same NAT rule.



Tip

With element-based NAT, the same connection can separately match the source and destination NAT. Hairpin NAT is automatic.

Considerations for designing NAT rules

NAT rules are processed from the top down, and more specific rules must be placed above more general rules that match the same traffic.

The traffic is matched based on the **Source**, **Destination**, **Service**, and **Used on** cells. The Source and Destination addresses in the cells must be valid for the address translation operation (the Source cell for source address translation and the Destination cell for destination address translation). When the first matching rule is found, the NAT defined for the rule is applied and the rest of the NAT rules are ignored. All NAT operations for the same connection must be defined in the same NAT rule (if you want to apply both source and destination translation to a connection).



Note

NAT is applied after an Access rule has allowed the connection but before a routing decision is made. Make sure that the Access rules allow the connection with the original (before NAT) addresses. Make sure that the translated (after NAT) addresses are included under the correct interface in the Routing pane of the Engine Editor, if necessary.

If you use element-based NAT, the NAT rules generated from NAT definitions are applied only after the NAT rules that have been added manually to the policy. This means that the NAT rules that are generated from NAT definitions do not override the rules that you have manually added to the Engine policy. Remember, however, that a more specific NAT rule can prevent traffic from matching the automatically generated NAT rules.

Default elements for NAT rules

The Firewall Template contains NAT rules that exclude some system communications from address translation.

The NAT rules exclude from address translation the communications between the engine and the Management Server that controls it. Also the communications from the engine to the Log Server where the engine sends its log data are excluded. Do not use NAT rules to translate the addresses in these system communications. Define Locations and Contact Addresses instead.

Defining address translation

The NAT cell of the NAT rules allows you to define that the source address, the destination address, or both addresses are translated..

If a connection matches the rule, the address translation defined in the NAT cell is applied. You can leave the NAT cell empty if you do not want to apply NAT to any connections that match the rule.

Static translation creates a one-to-one relationship between the original IP addresses and the translated IP addresses. When you use static source or destination translation, the translated address space must be as large as the original address space.

Dynamic source translation creates a many-to-one relationship between the original IP addresses and the translated IP addresses, so that several hosts can use the same IP address. In dynamic translation, a port is reserved for each host that is communicating. The number of ports in the port range determines how many hosts can communicate simultaneously using a single IP address. If the number of hosts exceeds the number of ports in the port range, translation cannot be applied and some of the communications fail. If failures happen, you might need to divide the dynamic translation rule and use an extra IP address for the translation. Dynamic translation can only be applied to communications that use ports (TCP and UDP-based protocols).

How NAT affects other engine configurations

Translated IP addresses are used in routing, in VPN site definitions, and system communications.

After adding or editing NAT rules, you must consider how these areas of communications are affected and what changes are needed. If you are using Multi-Link, Outbound Multi-Links have their own NAT configurations that must not overlap with the NAT rules you define.

In particular, check that:

- Access rules and Inspection rules use the addresses that are seen in the packets as they arrive on the engine (as they are before any NAT operation is done).
- Routing decisions are made after NAT, so the routing decision is made using the translated address. Make sure that the translated address is included in the Routing pane of the Engine Editor under the correct interface, unless the packets are forwarded to the default gateway.
- If you translate addresses of communications going through VPN tunnels, the translated addresses must be included in the VPN site definitions.



Note

By default, NAT is disabled with connections traversing a VPN (NAT rules are completely ignored for VPN traffic). If you want the NAT rules to apply to connections traversing a VPN, enable NAT in the properties of the VPN element.

NAT and system communications

If NAT is needed between SMC components, you must define Contact Addresses for the communications so that the components use the correct address for contact when needed.

Contact Addresses are used in a NAT environment for communications of SMC components with each other and with some external components that function as a part of the system. An example of these is a RADIUS server used for authenticating Administrators. Contact Addresses might also be needed for site-to-site VPNs and mobile VPNs.

The Firewall Template includes NAT rules which define that NAT is not done for communications between the engine where the policy is installed and the Management Server and Log Server that the engine uses. If these connections require NAT, the configuration must be done as explained here. Other system communications traversing the engine can be translated as any other connections. However, Location and Contact Address definitions are usually still needed for those components so that they know the correct addresses to use with each other. See *Situation where contact addresses are needed scenario*.

Contact Addresses are defined for *Locations*, which is an element that represents all devices that are routable behind a particular interface of a NAT device. The components that need Contact Addresses are placed in the Locations according to the Contact Address they use.

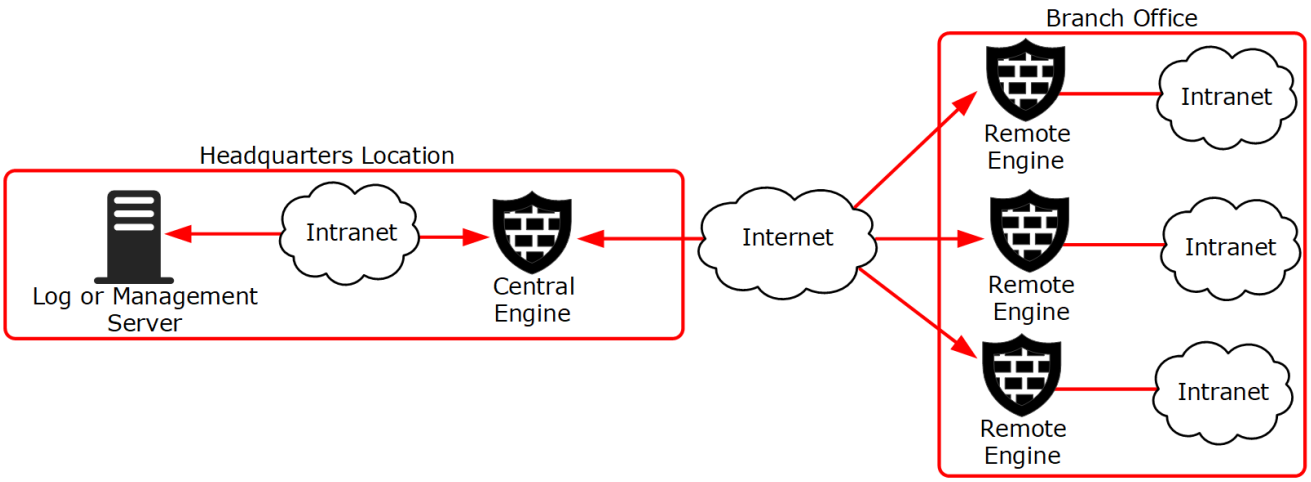
Related concepts

Situations where contact addresses are needed on page 857

Situations where contact addresses are needed

An example of a situation in which Contact Addresses are needed.

Contact address example



In this illustration, there are several remote engines that are managed through Management and Log Servers at a central site. NAT is typically applied at the following points:

- The central site engine or an external router can provide the SMC servers external IP addresses on the Internet. The external addresses must be defined as Contact Addresses so that the remote engines can contact the servers across the Internet.
- The central engine's IP address can be translated by an external router. The external IP address must be defined as a Contact Address to allow VPN connections from the remote engines to the central site using that address.
- NAT can also be applied at the remote sites (by external equipment) to translate the remote engines' IP address. In this case, you must define Contact Addresses for the remote engines so that the Management Server can contact them. The communications between the remote engines and the Management Server can also be reversed, so that the remote engines open the connections to the Management Server and maintain the connections open while waiting for commands.

When Contact Addresses are needed, a single Location to group all remote sites might be enough. The SMC servers' and the central engine's definitions must include a Contact Address for the "Remote Engines" Location. However, if VPN communications between engines from different remote sites are allowed, it is necessary to create a Location for each remote engine and to add further Contact Addresses for the engines.

Contact addresses, Location elements, and NAT

Contact Addresses represent the translated address of a component. Location elements group components together, so that there is no NAT between them.

You can specify an IP address or Fully Qualified Domain Name (FQDN) as the contact address to enable Security Engine to contact SMC management server or log server. If FQDN is specified, then you must also specify a DNS server in the engine configuration. The DNS server is used to match server host names to their corresponding IP addresses.

Note

FQDN resolves to IPv4 or IPv6 address. If FQDN resolves to multiple addresses, then all the addresses are attempted and the first IP address that works is used.

Contact Addresses are defined for each Location element. The Location element is a way to group components together, in effect telling them that there is no NAT device between them.

The SMC components on each side of a NAT device are grouped into two separate Location elements (if necessary, more Location elements can be used). The Contact Address is defined in each element's properties for the other Location. When contacting some other component in their own Location, the components always use the untranslated address. When contacting some component outside their own Location, the contacting component checks if the other component has a Contact Address defined for the contacting element's Location. If it finds one, it uses the Contact Address. If there is no Location-specific Contact Address defined, the contacting component checks if the element has a Default Contact Address that components belonging to any other Location use for contacting the element. If the element does not have a Default Contact Address, the connection is attempted using the element's untranslated address.

For example, when a Management Server contacts a engine node through NAT, the Management Server uses the translated Contact Address instead of the engine node's real Control IP address. The NAT device in between translates the NAT address to the engine's real IP address as usual.

We recommend dividing elements into different Locations based on NAT and the communications the components have, and not just based on actual physical sites. For example, you might have one central site and several remote sites, and the system communications take place only from each remote site to the central site (not between the remote sites). In this case only two Locations are needed no matter how many of the engines use a translated address.

Note

If NAT is performed between a Log Server and a SMC Client, you might need to select the correct Location for the SMC Client as well.

Outbound load-balancing NAT

You can use NAT for outbound load-balancing.

In addition to source and destination translation, another main use for NAT is outbound load balancing. It is used in a Multi-Link configuration where the Engine balances outbound traffic between two or more network connections. To be able to direct traffic to the faster connection, the engine translates outgoing connections to an address from a pool assigned to each available NetLink. In this case, the IP addresses used for the NAT are defined in the properties of the Outbound Multi-Link element.

Related concepts

Defining Multi-Link routes on page 735
Getting started with outbound traffic management on page 731

Proxy ARP and NAT

Automatic proxy ARP requires an explicit route to the host or network to be configured in the Routing pane of the Engine Editor.

Proxy ARP (Address Resolution Protocol) is a specification that allows a device to respond to ARP requests on behalf of some other device on the network. When network address translation is used on a engine, the engine is typically configured to use proxy ARP so that it can accept packets for the translated addresses. The engine uses proxy ARP instead of requiring the administrator to assign all the translation addresses to the engine's network interface.

In the Engine, proxy ARP is handled automatically. Proxy ARP is enabled by default in the **NAT** cell in NAT rules for each translation type, although you can deselect it, if necessary.

Protocols and NAT

Protocols of the Protocol Agent type help with problems related to certain complex protocols and NAT.

In some protocols, such as FTP, IP address information is included in the data payload of the packets, which do not normally include information for routing. Protocols of the Protocol Agent can examine the data payload of packets arriving to the engine and also edit it. For example, when the source address is included in a packet's data, the engine can translate the original source address and also the address embedded in the data.

Related concepts

Protocol Agents overview on page 1031

Examples of NAT configuration

These examples illustrate some common uses for NAT rules and the general steps on how each example is configured.

Example: configuring NAT for dynamic source address translation

An example of configuring NAT for dynamic source address translation.

Company A uses private IP addresses that are not routable on the Internet in their internal network. The administrators need to translate the internal IP addresses to IP addresses that are routable on the Internet to make it possible to use external services. The administrators:

1) Create an Address Range element “External Addresses” for two consecutive IP addresses from the pool of addresses that they have been assigned by their Internet service provider.

2) Add a NAT rule to their Engine Policy:

Dynamic translation rule for opening connections to the Internet

Source	Destination	Service	NAT
“$ Local Protected Sites” Alias	“NOT $ Local Protected Sites” Expression	ANY	Source: Dynamic to External Addresses 1024–65535

- The administrators use the whole range of high ports (1024–65535) for translating the addresses in this case.
- Return address translation is done automatically. Therefore, a single rule suffices to cover all (client) hosts that only open connections themselves, and do not need to accept new connections coming from external networks.

3) Refresh the Engine Policy. All internal addresses are now hidden behind two IP addresses and a range of ports.

Example: configuring NAT for static address translation

An example of configuring NAT for static address translation.

Company A has set up the engine to translate the IP addresses of all communications between the internal and the external network dynamically. However, the company also has a mail server, which must be able to accept connections from external networks. For this, it must have a fixed translated IP address. The administrators:

1) Create the Host element “Mail Server” to represent the mail server’s private IP address.

2) Create the Host element “Mail Server NAT” to represent the mail server’s public IP address.

3) Add two new NAT rules above the general dynamic translation rule.
   - In this case, new connections can be opened both from the mail server and from external hosts, so two rules are necessary.

4) Change the newly added NAT rules as follows:

Static translation rules for opening connections both ways

Source	Destination	Service	NAT
“Mail Server” Host element	“NOT $ Local Protected Sites” Expression	“SMTP” Service element	Source: Static from Mail Server to Mail Server NAT
“NOT $ Local Protected Sites” Expression	“Mail Server NAT” Host	“SMTP” Service element	Destination: Static from Mail Server NAT to Mail Server

   - The first rule is for connections that the mail server opens to external hosts.
   - The second rule is for connections that external hosts open to the mail server.
   - Return address translation is done automatically, so if the connection would always be opened from one end, a single rule would suffice.

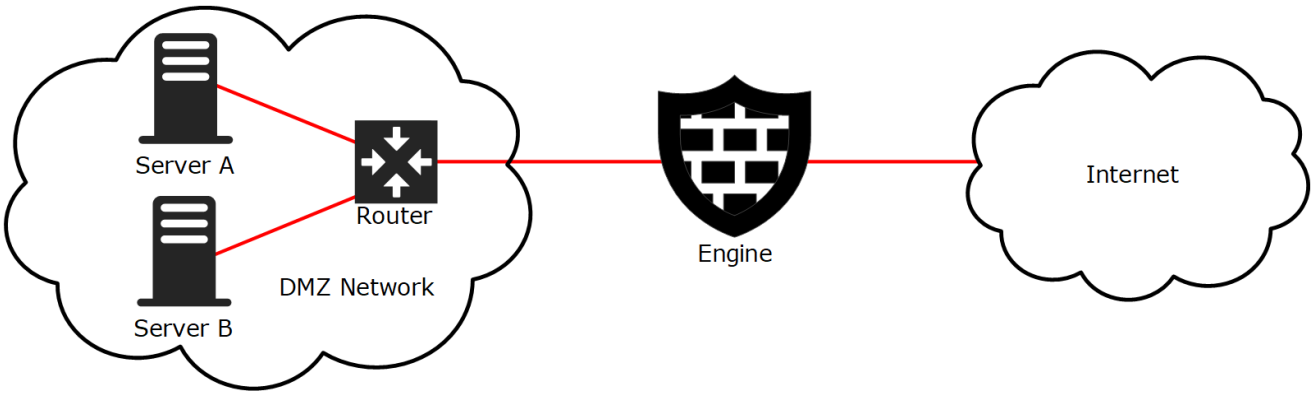
5) Refresh the Engine Policy.

Example: configuring NAT with hosts in the same network

An example of configuring NAT to route traffic between hosts in the same network through the engine.

Company B has two servers running in the same DMZ network. The servers keep contact with each other to exchange some information. The administrators want to route the traffic through the engine so that it is logged for reporting purposes instead of letting the servers communicate with each other directly.

Company B’s network setup



The administrators first intend to just configure the servers to use the external (NAT) address of the other server as a destination and configure the related static destination NAT rule. However, they soon realize that the receiver

would see the real source address in the communications and the replies would be sent directly, bypassing the engine for the reply communications. This action would obviously prevent the connections. A static source NAT is required in addition to the static destination NAT.

The administrators:

1) Create Host elements to represent the private addresses of the two servers.

2) Create Host elements to represent the public addresses of the two servers.

3) Add two new NAT rules before any other NAT rule that would match these connections:

Static translation rules for opening connections both ways

Source	Destination	Service	NAT
“Server A Private” Host	“Server B Public” Host	ANY	Source: Static from Server A Private to Server A Public Destination: Static from Server B Public to Server B private.
“Server B Private” Host	“Server A Public” Host	ANY	Source: Static from Server B Private to Server B Public Destination: Static from Server A Public to Server A private.

- When the servers are configured to contact each other using the public IP addresses, the communications are routed through the engine.
- The Engine translates the destination to the other server’s private IP address and the private IP address of the source to the public IP address to “hide” the private source address from the receiving host. This way, the replies are sent to the public IP address and routed correctly through the engine.

Examples of NAT rules

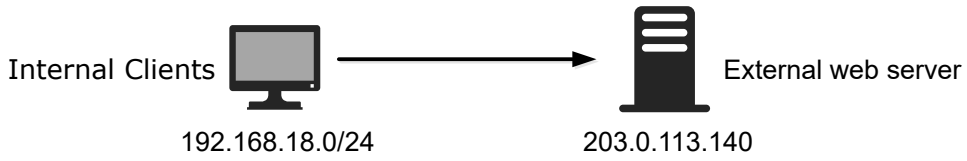
Read the following examples of NAT rules.

Example of a static source translation NAT rule

This example shows a static address translation that translates the addresses in one network to IP addresses in another network.

In this example, the NAT is done to access a particular server on the Internet.

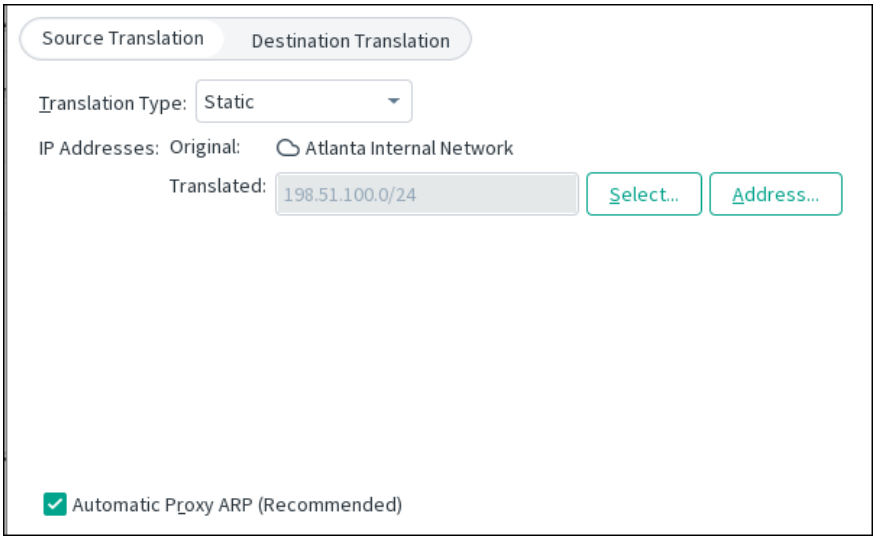
Example scenario



Example NAT rule matching cells

Source	Destination	Service
192.168.18.0/24	203.0.113.140	HTTP

Example NAT settings

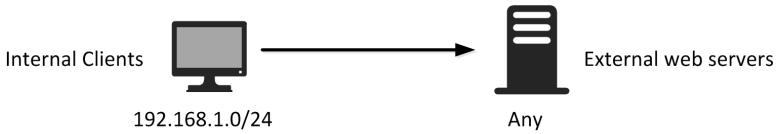


In static address translation using whole networks, each original source IP address has a static translated pair. For example here, the host 192.168.1.6 is always translated to 198.51.100.6 and host 192.168.1.11 is always translated to 198.51.100.11.

Example of a dynamic source translation NAT rule

This example shows a dynamic address translation that translates the addresses in one internal network to a single external address for general web browsing.

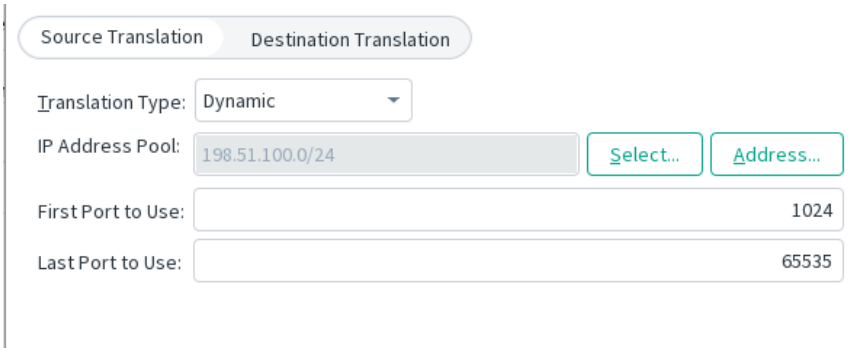
Example scenario



Example NAT rule matching cells

Source	Destination	Service
192.168.1.0/24	ANY	HTTP

Example NAT settings



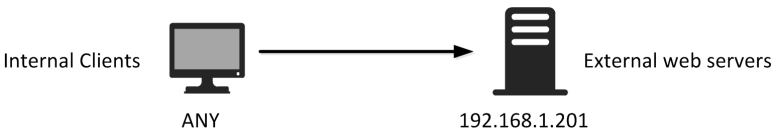
In dynamic address translation, several source IP addresses are translated using a smaller pool of translated addresses with the help of port translation. Each client connection uses a different port on an IP address that is shared between several different connections. Because each client reserves a port, the maximum number of simultaneous connections can be calculated by multiplying the number of IP addresses by the number of ports in the range. Every port and IP address pair must be free from any other use (duplicate connections cannot successfully cross the engine).

Example of a destination translation NAT rule

This example shows a static address translation that translates the external IP address of a web server to the server's internal address.

The external IP address (203.0.113.140) of the web server is translated to the server's internal address (192.168.1.201).

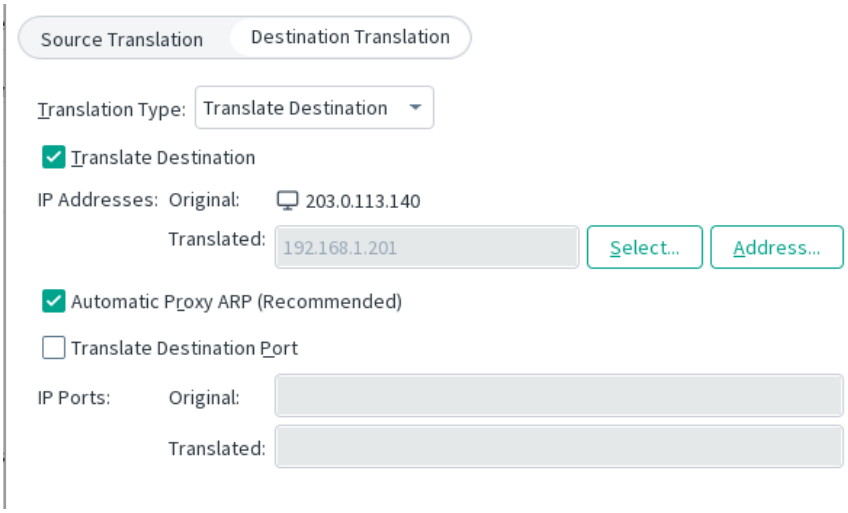
Example scenario



Example NAT rule matching cells

Source	Destination	Service
ANY	203.0.113.140	HTTP

Example NAT settings



Example of a combined source and destination translation NAT rule

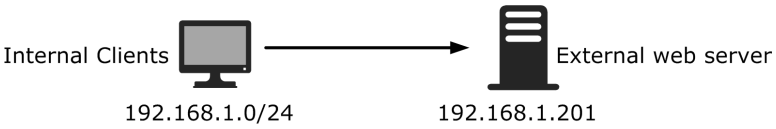
In this example, *hairpin NAT* is configured.

Tip

With element-based NAT, the same connection can separately match the source and destination NAT. Hairpin NAT is automatic.

Clients in the internal network (192.168.1.0/24) contact the organization's own public web server using the public IP address (203.0.113.140). The server's external address is translated to an internal address (192.168.1.201) that belongs to the same internal network address space as the contacting clients. Source address translation is used to prevent the server replies to the client's original IP address. Such replies would be routed directly within the local network instead of through the engine, and the connections do not work without the reverse NAT that the engine provides.

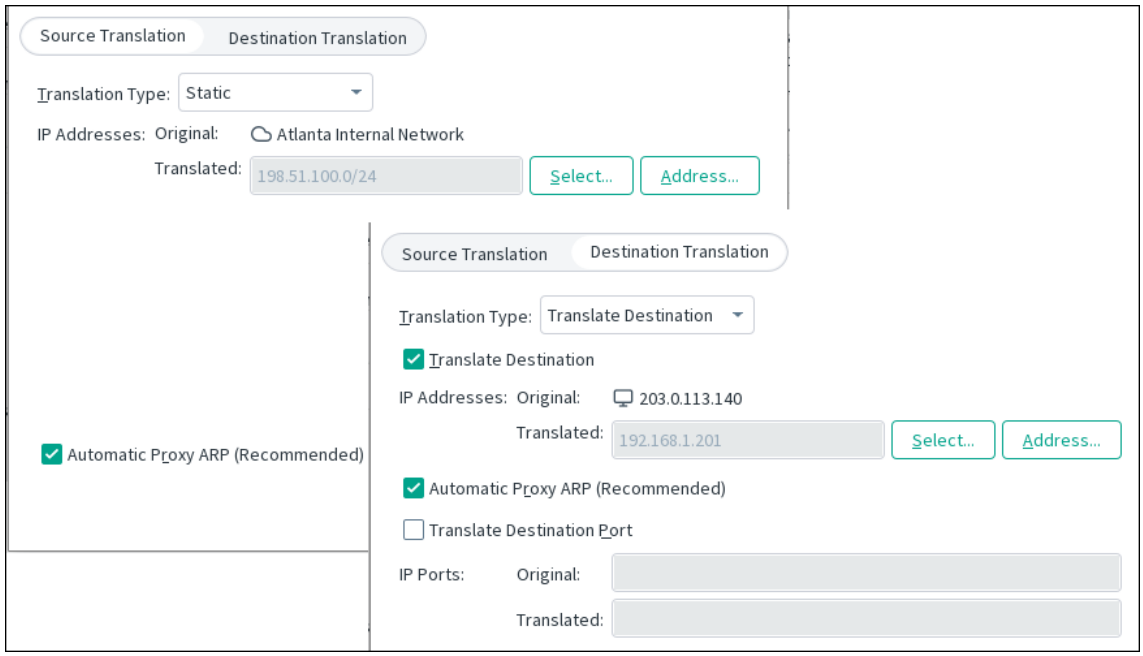
Example scenario



Example NAT rule matching cells

Source	Destination	Service
192.168.1.0/24	203.0.113.140	HTTP

Example NAT settings



The NAT settings on each tab are not any different than when you apply only source translation or only destination translation to matching connections. Both definitions must be defined in the same NAT rule, because none of the other NAT rules are considered after the first match is found.

Chapter 51

Inspection Policy elements

Contents

- Inspection Policy elements and how they work on page 867
- How Inspection Policy elements are designed on page 868
- Set default options for Exception rules in Inspection Policy elements on page 874
- Example: Tuning an Inspection Policy element to eliminate false positives for a engine on page 875

Inspection Policy elements define how the engines look for patterns in traffic allowed through the Access rules and what happens when a certain type of pattern is found.

Inspection Policy elements and how they work

Inspection Policy elements define how the main traffic analysis is done for traffic that has been allowed and selected for deep inspection in the Access rules. They define what action the engine takes when a match is found.

The Inspection Policy elements are selected in Engine, IPS, and Layer 2 Engine Policy elements. The IPS Template and Layer 2 Firewall Template enable deep inspection for all IP traffic. Deep inspection is not automatically enabled in the Firewall Template.

Deep inspection examines the packet payload throughout whole connections, and acts when something threatening is discovered.

Security Engines examine IPv4 and IPv6 traffic against traffic patterns defined in Situation elements. Security Engines and Log Servers process the detected events using *Correlation Situation* criteria. Dynamic update packages are the main source of Situation elements. If you want to detect a specific traffic pattern (for example, a particular internal file server in your network being accessed) or if you want to create an edited version of some existing Situation element, you can also define new patterns as custom Situation elements. You can add your custom Situation elements to the Rules tree by selecting a Situation Type for them.

There are three general types of cases for using Inspection Policy elements:

- You can detect attempts to exploit known vulnerabilities in your systems and prevent such attempts from succeeding if the system is not patched against it.
- You can monitor traffic that does not cause alarm on the surface, but when examined for certain patterns, can turn out to conceal actual threats. For example, you can detect if a series of occasional service requests are someone secretly scanning the network structure or if a spike in traffic is a denial-of-service attack.
- You can also detect other sequences in traffic, such as the use of certain applications or even access to a particular file.

Based on the detection results, the Inspection Policy element provides several different ways to react when some traffic is found to match a pattern of interest:

- Stop the traffic if it is going through a Engine.
- Stop the traffic if it is going through an IPS engine or Layer 2 Engine with inline interfaces.

- Reset the connection.
- Block list the connection on one or more Security Engines.
- Allow the traffic.

Regardless of which action is taken, a match can also create:

- A log entry with or without recording some of the detected traffic.
- An alert with or without recording some of the detected traffic.

Engines can inspect all protocols. Virtual Security Engines do not individually inspect traffic. One shared inspection process running on the Master Engine handles the inspection and correlation for all Virtual Security Engines associated with the Master Engine. To prevent excessive resource consumption on the Master Engine, take care when configuring Inspection policies for use on Virtual Security Engines.

Related concepts
Getting started with policies on page 799
Getting started with Situation elements on page 941

How Inspection Policy elements are designed

Inspection Policy elements activate checks for specific traffic patterns and define what action the engine takes when a match is found.

Engines, IPS engines, and Layer 2 Engines inspect traffic based on Situation elements, which contain the information about traffic patterns. Patterns can trigger immediate responses or be recorded. Detected events can be matched against Correlation Situations, which combine and further analyze the traffic-based findings to detect more threats and produce an easy-to-read event stream.

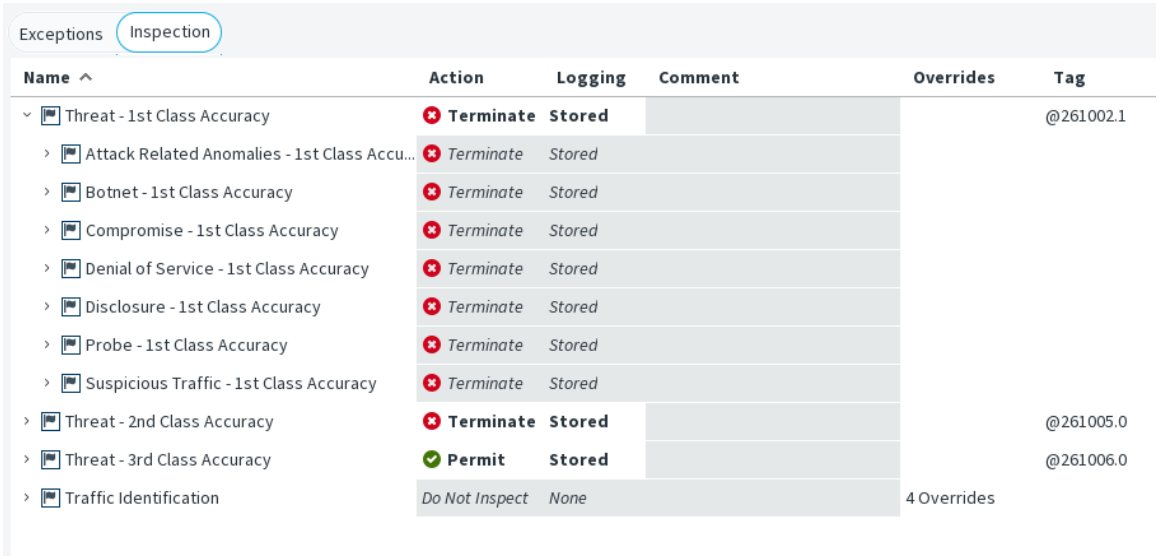
Inspection Policy elements are selected on the **Inspection** tab inside the Engine, IPS, and Layer 2 Engine Policy and Template Policy elements. Sub-Policies cannot contain Inspection Policy elements. You can add new rules to the Inspection Policy elements in the Policy Editing View and also in the Logs view based on log entries.

The Inspection Policy elements has two parts:

- The **Inspection** tab contains the main rules for finding traffic patterns. The Rules tree is applied to all traffic that is not handled as Exceptions.
- The **Exceptions** tab contains rules that match specific portions of the traffic based on Logical Interface, IP addresses, and Ports. Exceptions have some additional options, and can also set some of those options for the main Rules by using *Continue* rules.

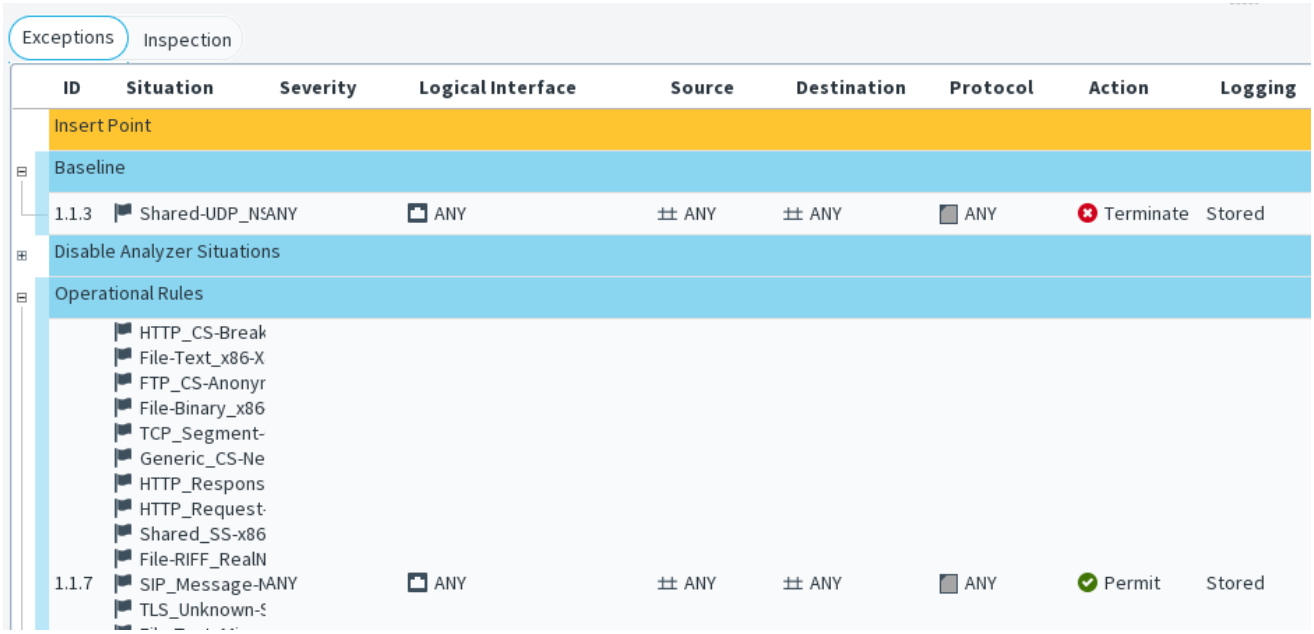
The main Rules tree on the Inspection tab contains a tree of Situations, which are organized under Situation Types. This tree allows you to control which inspection checks trigger a reaction and which checks are ignored. The Rules tree defines general checks that are applied to all patterns that are not handled by a more specific definition. It is not possible to limit the scope of the checks to certain IP addresses or Logical Interfaces in the Rules tree.

Inspection tab - Rules tree



The Exceptions are matched before the main rules. The most frequent use of Exceptions is to eliminate false positives. This typically requires permitting a pattern for some part of the traffic while the same pattern still triggers a reaction when it is encountered in any other traffic.

Exceptions tab



The main matching cell is the Situation cell, which contains the actual patterns. The other matching cells are Logical Interface, Source, Destination, Protocol, and Time. The role of the other matching cells is to limit the scope of the rule to some specific traffic. For example, the engine can take different action based on which host is the sender or receiver of traffic identified as malicious.

Verifying and tuning Inspection Policy elements

Tuning increases the relevancy and accuracy of the findings that the system generates.

The most common way to introduce inspection is to start with a default Inspection Policy element. Tuning the policy is important, because a general policy that is meant to work in all environments is not necessarily suited to your particular network environment. A tuning period is needed to activate and deactivate inspection checks based on the findings and your needs.

To help policy tuning, you can use the *passive termination* feature. When passive termination is used, the engine creates a special log entry that notes that a certain connection would have been terminated. However, the engine does not actually terminate the connection. This allows you to check the logs and adjust your policy without the risk of cutting important business communications. There are two levels of activating this feature:

- Passive termination can be activated globally in the engine's properties for the initial policy tuning.
- Later on, you can test newly added Situations by setting individual Exception rules to passive termination mode.

For cautious introduction of new Situations introduced in dynamic update packages, you can use the Tags for the five most recent updates (**Situations** > **By Tag** > **By Situation Tag** > **Recent Updates**).

Rule order for Inspection Policy elements

The rules in Inspection Policy elements are read from the top down. More specific rules must be placed above more general rules that match the same traffic.

The detailed rules specific to some IP addresses and Protocols are defined on the **Exceptions** tab. The general rules that are applied to remaining traffic are defined in the **Rules** tree on the **Inspection** tab.

The traffic matching in Inspection rules and exceptions is different from other types of rules because it is done based on the traffic pattern definitions in Situation elements. The Security Engines inspect the traffic for all patterns included in the policy. When a pattern is found, the Inspection rules and exceptions match based on the Situation element that contains the detected pattern. Inspection rules and exceptions match certain patterns only. Non-matching traffic is allowed through without taking any actions.



Note

Each Situation element is a unique pattern. Avoid defining the same pattern in different Situation elements. Duplicate situations in the policy can create unintended results and makes the policies difficult to manage.

Inspection rules and exceptions can look different even if they refer to the same Situation because Situations can be grouped using Situation Tag and Situation Type elements. However, the rules match patterns in the same way whether you add the Situation as a single element or together with other Situations through a Situation Tag or Situation Type.

Because traffic matching is based on the traffic pattern definitions in Situation elements, the behavior of the Inspection rules and exceptions can change without anyone editing the policy directly. For example, creating a Situation element can include the Situation in the policy if the Situation is associated with a Situation Tag or Situation Type element that is used in the policy.

The Permit and Terminate actions in Inspection rules and exceptions have different effects on policy processing when a rule matches.

- **Permit** — Allows traffic that matches the traffic pattern. A Permit action does not unconditionally allow the traffic because processing continues to look for other patterns. However, a Permit match does prevent the same Situation from matching again if it appears at any point further down in the policy.
- **Terminate** — Stops traffic that matches the pattern. The Terminate action prevents the same Situation from matching again if it appears at any point further down in the policy, but does not prevent other Situations from matching.

For example, there is a rule that contains Situation A with Permit as the action and the logging level set to “None”. There is a second rule that contains Situation A below the first rule with Terminate as the action and the logging level set to “Stored”. Because traffic already matched the rule that permits traffic that matches Situation A, no log entries are generated for Situation A and the traffic that matches the pattern continues uninterrupted.

Related concepts

Getting started with Situation elements on page 941

Inspection on Master Engines and Virtual Engines

One shared inspection process running on the Master Engine handles the inspection and correlation for all Virtual Engines associated with the Master Engine.

Virtual Engines do not individually inspect traffic. To prevent excessive resource consumption on the Master Engine, take care when configuring Inspection policies for use on Virtual Engines.

Default elements for Inspection Policy elements

Default Inspection Policy elements are introduced when you import and activate a dynamic update package. The rules in the Inspection Policy Templates can change when you activate new update packages.

To customize inspection, you must have a custom Inspection Policy element. The predefined templates are a good starting point for your own customization.

Note

Keeping your system up to date with latest dynamic updates is an essential part of maintaining your Inspection Policy elements.

Default Inspection Policy elements

Template	Description
No Inspection Policy	Suitable for Engine deployments, in which only packet filtering is needed. Disables deep packet inspection.
Medium-Security Inspection Template	For Engines, Layer 2 Engines, inline IPS deployments in asymmetrically routed networks, and IPS deployments in IDS mode. Terminates reliably identified attacks and logs Situations that have some degree of inaccuracy. Low risk of false positives.

Template	Description
High-Security Inspection Template	For Engine, Layer 2 Engine, and inline IPS use. Extended inspection coverage and evasion protection. Not for asymmetrically routed networks. Terminates reliably identified attacks, and Situations that have some inaccuracy. Moderate false positive risk.
Highest-Security Inspection Template	For Engine, Layer 2 Engine, and inline IPS use. Highest level of inspection coverage and evasion protection. Occasional false positives are accepted.

Activating deep inspection in Access rules

Action options in the Access rules define which traffic is inspected against the Inspection Policy.

Typically, you introduce deep inspection after creating and testing initial Access rules. You must specifically activate deep inspection for the portion of traffic that you want to deep inspect. This activation is done in the Access rules. You also select which Inspection Policy element is used for deep inspection on the **Inspection** tab of the Engine, IPS, or Layer 2 Engine policy.

Related concepts
Getting started with Access rules on page 831

Activating the relevant inspection checks

The Rules tree is the main tool for controlling deep inspection.

Traffic patterns of interest are defined in Situation elements. The inspection checks are based on selecting the reaction to the Situations when the pattern is found. It is not mandatory to create any additional Situation elements to activate inspection checks, because there are many default Situation elements and they are continuously updated through dynamic update packages.

The Rules tree on the **Inspection** tab is the main tool that allows you to select which traffic patterns are permitted and stopped. You can also select whether a log entry or an alert is triggered, and whether matching traffic is recorded. All Rules in the Rules tree can be edited, including overrides that have been set in a higher-level template. The Rules tree can contain a maximum of one instance of each Situation to prevent the definitions within the Rules tree from overlapping.

Inspection categories

Inspection categories are grouped into Threat and Traffic Identification categories.

Threat categories

Situations categorized under Threat categories are generated from traffic that exhibits malicious patterns. Threat categories are classified as 1st class accuracy, 2nd class accuracy, and 3rd class accuracy.

Template	Description
1st Class Accuracy	The most reliable indication of malicious traffic. The situations categorized under 1st Class Accuracy are terminated in all inspection policy templates.
2nd Class Accuracy	A reliable indication of malicious traffic. In rare conditions these situations may be triggered from normal traffic. The situations categorized under 2nd Class Accuracy are terminated in Highest-Security Inspection Template and High-Security Inspection Template, and are logged but not terminated in Medium-Security Inspection Template.
3rd Class Accuracy	Moderately reliable indication of malicious traffic. There is a higher risk of false positives from normal traffic. The situations categorized under 3rd Class Accuracy are only terminated in the Highest-Security Inspection Template, and logged but not terminated in High-Security Inspection Template and Medium-Security Inspection Template.

Traffic identification categories

Situations categorized under Traffic Identification categories are generic situations used for identifying various details about network traffic, such as protocol versions, message types, browser versions (including identification for obsolete versions) and file types. These situations are useful for debugging or blocking specific traffic types. Inspection Engine uses these situations for file inspection and application identification purposes. These situations are not terminated in the inspection policy templates.

An exception to this are the “Packet Validity Situations” categories under Traffic Identification. These categories contain all packet validity situations that the engine generates, including threat related packet validity situations. For more information about the packet validity situations, see Knowledge Base article 36144 .



Note

The Logging setting for **Content Identification** cannot be changed to **None** (default value), otherwise the functionality of the inspection engine is compromised.

Defining Exception rules in Inspection Policy elements

Exception rules in Inspection Policy elements allow you to make changes to the Inspection Policy that are not applied to all connections.

The **Exceptions** tab allows you to create detailed rules, which are processed before the Rules tree definitions on the Inspection tab. The Exceptions have additional features compared to the Rules tree:

- You can make exceptions to the general Rules tree definitions based on Source, Destination, and Protocol information.
- You can set options for connection termination (including User Responses) in addition to the options that are available in the Rules tree. The Response options define an automatic client notification for any HTTP connection that is terminated.
- You can create Continue rules to set Action Options and general rule Options for other Exceptions and the Rules tree. The Rules tree contains specific definitions for logging, so the logging options set with Continue rules do not affect traffic that matches the Rules tree.
- You can create rules in Inspection Policy Template elements that cannot be changed in the inheriting policies.
- You can create rules that are applied only on certain days or times of day.

In addition to individual Situation elements, the Situation cell can contain Tag and Situation Type elements. These elements are shown as branches in the Situations tree and allow adding the whole branch of Situations at once to a rule. Most of the Situations you add to the Exceptions are those that you regard as false positives in your environment. An example might be Situations for exploit attempts against an operating system that is not used in your organization).

In the Exceptions, it is highly unusual to set the Situation cell to ANY. This is not useful in most cases because the patterns that Situations define range widely. There are Situations that detect something as benign as the use of particular applications and Situations that detect something as malicious as suspicious traffic on a server. The ANY setting also creates unnecessary load on the engines, as a high number of Situations is checked in each matching connection.

Tuning Inspection Policy elements to eliminate false positives

False positives are matches to rules and exceptions in Inspection Policy elements that are incorrect or irrelevant in your environment.

As the Inspection rules and exceptions are matched to traffic, there are always some false positives. By tuning the Inspection Policy element to the actual traffic and applications in your network environment, you can increase the relevance of inspection results greatly. To eliminate a false positive, you adjust either the Inspection Rules tree or the Exception rules depending on whether the change should be applied globally or to traffic between specific hosts. An easy way to create new Exceptions is to use an existing log entry as the basis: you can create Exceptions through the right-click menu of log entries.

Related concepts

Example: Tuning an Inspection Policy element to eliminate false positives for a engine on page 875

Set default options for Exception rules in Inspection Policy elements

The **Continue** action allows you to set default values for some settings in Exception rules to avoid defining the same settings for several rules individually.

You might want to set default settings for some Exception rules to avoid defining the same settings for several rules individually. The Continue action in Exception rules is used to set such default options in the same general way as in the Access rules. In Exception rules, all settings in the Action Options and the Logging cell can be set using Continue rules. However, the Rules tree on the Inspection tab ignores any logging options set with Continue rules. In the Rules tree, the rules either inherit the logging settings from a higher level in the tree or define a specific logging option as an override.

Related concepts

Configuring default settings for several Access rules on page 839

Example: Tuning an Inspection Policy element to eliminate false positives for a engine

An example of using Exception rules in the Inspection Policy element to eliminate a false positive.

The administrators in this example have started using inspection. They have installed a policy that includes only the rules defined in the Loose Inspection policy. When they install the Engine Policy, they soon start receiving alerts.

After some investigation, the administrators realize that a custom-built application causes the alert. This application communicates in a way that happens to match the pattern of how an attacker would carry out a certain exploit. The custom-built application is only used by a specific server and a few clients in the internal network. The administrators quickly edit the Inspection policies to exclude those particular hosts for the Situation in question. The administrators:

1) Create Host elements to represent the server and the clients.

2) Create a Group element that includes the client's Host elements.
   - The administrators name the Group so that it is immediately clear from the name that the Group contains those hosts that must contact the server running their custom-built application. This makes the new rule easier to read than if they included the hosts directly in the rule.

3) Add the following rule on the Exceptions tab in their Inspection Policy element:

   Rule for eliminating a false positive

Situation	Source	Destination	Action	Logging
The Situation element that is mentioned in the alerts in the Logs view.	The Group defining the clients.	The Host for the internal server.	Permit	None

   - If the Situation matches traffic between any other hosts than those included in the Group, the IP address does not match the hosts defined in the new rule. The processing will continue to the next rule, which terminates the traffic and triggers an alert.
   - The logging would not have to be set to None, because it is the default option. However, the administrators want to do so anyway to make sure any rules that they add in the future cannot accidentally set logging on for this rule.

4) Refresh the policy on the Engines.

Chapter 52

Snort inspection on Security Engines

Contents

- Getting started with Snort inspection on Security Engines on page 877
- Prepare Snort configuration files on page 878
- Import Snort configuration files globally for all Security Engines on page 879
- Enable Snort inspection for Security Engines on page 880
- Override settings in the global Snort configuration for individual Security Engines on page 881
- Add Access rules for Snort inspection on page 881
- Logging for Snort inspection on page 882

The Snort open source intrusion prevention system is integrated into Forcepoint Security Engine. You can import externally created Snort configurations into Forcepoint Security Engine to use Snort rule sets for inspection.

Getting started with Snort inspection on Security Engines

Snort inspection on Security Engines allows you to use externally created Snort rule sets for inspection.

You can use both Security Engine deep inspection and Snort inspection for the same traffic, or you can use only Security Engine deep inspection or only Snort inspection.

Snort inspection is supported for Security Engines in the Engine/VPN role, IPS, and Layer 2 Engine roles.

Limitations of Snort inspection on Security Engines

These limitations apply to Snort inspection on Security Engines.

- Snort inspection is currently not supported for Master Engines and Virtual Engines.
  If you install a policy that includes Access rules for Snort inspection on Master Engines and Virtual Security Engines, the rules are ignored.
- Snort inspection is not supported for Capture interfaces.
- Snort inspection is supported for VLAN interfaces, but the same Snort rules apply to the traffic regardless of the VLAN tag. Snort inspection is only applied to the IP datagrams without Ethernet headers. It is not possible to apply different Snort rules to traffic from different VLANs.

- If you use Logical Interfaces that have overlapping IP address spaces as matching criteria in Access rules that select traffic for Snort inspection, traffic might not match Snort rules as intended.
- We do not recommend using services that match based on the payload of connections, such as Network Applications, URL Categories, or URL List Applications, in Access rules that select traffic for Snort inspection. At the beginning of a connection, the Security Engine cannot determine whether the traffic should be selected for Snort inspection. The Security Engine selects all potentially matching traffic for Snort inspection. As a result, Snort inspection might be applied to traffic that was not intended to be selected for Snort inspection. Applying Snort inspection to this traffic can create false positive Snort rule matches.
- Snort inspection cannot be applied to traffic that has been decrypted for TLS inspection.
- If Snort inspection fails, the traffic is allowed by default.
- Security Engines do not receive automatic updates for Snort rule sets. When new Snort rule sets are available, you must import new Snort configuration files and refresh the policy on the Security Engine to start using the new Snort rule sets.

Snort inspection configuration overview

The configuration of Snort inspection consists of these general steps.

1) Prepare Snort configuration files.

2) Import Snort configuration files globally to configure default settings for Snort inspection for all Security Engines.

3) (Optional) Import Snort configuration files for individual Security Engines to override settings in the global Snort configuration for specific Security Engines.
Settings in the Snort configuration .zip file for an individual Security Engine are combined with the settings in the global Snort configuration .zip file. If any configuration files in a Snort configuration .zip file for an individual Security Engine have the same files name and paths as configuration files in the global Snort configuration .zip file, the overlapping files in the global Snort configuration .zip file are ignored.

4) Enable Snort inspection for Security Engines.

5) Create Access rules to select traffic for Snort inspection.

Prepare Snort configuration files

Snort configuration files are externally created compressed archive files that contain the Snort rule sets and Snort settings.

Before you begin

Create or obtain Snort rule sets. For example, if you subscribe to a Snort rule updates, download the latest Snort rule update file.

The Security Engine uses data from the Snort configuration file to configure Snort inspection.

The `snort.conf` file must be in the root directory of the Snort configuration. The configuration in the `snort.conf` file can also reference other configuration files. The recommended practice is to use only relative references. For example, `./<directory name>/<filename>.conf` rather than `/etc/snort/<directory name>/<filename>.conf`

Steps

1) On a computer where the Snort rule sets are saved, copy the `snort.conf` file to the root directory of the Snort configuration.

2) Copy the Snort rule files to the `rules` directory.

3) Compress the Snort configuration as a .zip file.

4) Copy the Snort configuration file to a location that you can access from the computer where you use the SMC Client.

Next steps

Import the Snort configuration file globally for all Security Engines or for individual Security Engines.

Import Snort configuration files globally for all Security Engines

Import Snort configuration files globally to configure default settings for Snort inspection for all Security Engines.

Note

Security Engines do not receive automatic updates for Snort rule sets. When new Snort rule sets are available, you must import new Snort configuration files and refresh the policy on the Security Engine to start using the new Snort rule sets.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Settings** > **Global System Properties**.

2) Click the **Global Options** tab.

3) Click **Browse** next to the **Snort Configuration** field, then select the Snort configuration file.

4) Click **OK**.

Next steps

Enable Snort inspection for each Security Engine where you want to use Snort inspection.

Enable Snort inspection for Security Engines

Enable Snort inspection for each Security Engine where you want to use Snort inspection.

Before you begin

Import a Snort configuration file in the Global System Properties dialog box.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click a Security Engine, then select **Edit <element type>**.
3) Browse to **Add-Ons** > **Snort**.
4) Select **Enable**.

   Note

   To apply Snort inspection to traffic, you must also create Access rules to select traffic for Snort inspection.

5) Click **Save and Refresh**.

Next steps

Continue the configuration in one of the following ways:

- If you want to override settings in the global Snort configuration for specific Security Engines, import Snort configuration files for individual Security Engines.
- Create Access rules to select traffic for Snort inspection.

Override settings in the global Snort configuration for individual Security Engines

You can optionally import a Snort configuration .zip file for an individual Security Engine to override settings in the global Snort configuration for specific Security Engines.

All Security Engines for which Snort inspection is enabled use the global Snort configuration by default. If you do not want to override settings in the global Snort configuration, it is not necessary to import a Snort configuration file for an individual Security Engine.

Settings in the Snort configuration .zip file for an individual Security Engine are combined with the settings in the global Snort configuration .zip file. If any configuration files in a Snort configuration .zip file for an individual Security Engine have the same files name and paths as configuration files in the global Snort configuration .zip file, the overlapping files in the global Snort configuration .zip file are ignored.

Note

Security Engines do not receive automatic updates for Snort rule sets. When new Snort rule sets are available, you must import new Snort configuration files and refresh the policy on the Security Engine to start using the new Snort rule sets.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click a Security Engine, then select **Edit <element type>**.
3) Browse to **Add-Ons** > **Snort**.
4) Click **Browse** next to the **Snort Configuration** field, then select the Snort configuration file.
5) Click **Save and Refresh**.

Next steps

Create Access rules to select traffic for Snort inspection.

Add Access rules for Snort inspection

Create Access rules to select traffic for Snort inspection.

You can use the following matching criteria in Access rules that select traffic for Snort inspection:

- Source and destination IP addresses, networks, and interfaces

Note

If you use Logical Interfaces that have overlapping IP address spaces as matching criteria in Access rules that select traffic for Snort inspection, traffic might not match Snort rules as intended.

- Service elements that match based on the port that the traffic uses

Note

We do not recommend using services that match based on the payload of connections, such as Network Applications, URL Categories, or URL List Applications, in Access rules that select traffic for Snort inspection.

Note

If Snort inspection fails, the traffic is allowed by default.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) To select traffic for Snort inspection, add the following type of Access rules:

Access rules for Snort inspection

Source	Destination	Service	Action
IP address, Network, or Zone elements, or **ANY**	IP address, Network, or Zone elements, or **ANY**	One or more Service elements that match based on the port	**Allow** or **Continue**

2) Right-click the **Action** cell, then select **Edit Options**.

3) From the **Snort** drop-down list, select **On**.

4) Click **Save and Install**.

Logging for Snort inspection

Log entries are generated when traffic matches a Snort rule that sends a message or an alert.

By default, log entries are produced when traffic matches the following Situation elements for Snort inspection:

- Snort_Alert
- Snort_Drop
- Snort_Message
- Snort_Reject
- Snort_Timeout

You can optionally use the following Situation elements in Inspection Exception rules to create log entries when traffic matches a Snort rule that does not a message or an alert:

- Snort_Drop-Silent
- Snort_Reject-Silent

In the Logs view of the SMC Client, the **Snort** facility shows log entries related to Snort inspection.

The following log fields show information about Snort inspection:

- **Snort Message** — Shows the message or alert that Snort sends when traffic matches a Snort rule.
- **Snort Rule ID** — Shows the rule identifier of the Snort rule that the traffic matched.

Chapter 53

Editing policies

Contents

- Getting started with editing policies on page 885
- The different parts of the policy editing view on page 887
- Editing rules in a policy on page 889
- Add Insert Points in Policy Templates on page 896
- Automatic rules and how they work on page 896
- Configure settings for Automatic rules on page 897
- Add Ethernet rules on page 897
- Add Access rules on page 899
- Add NAT rules on page 903
- Add Inspection rules on page 904
- Add Exception rules on page 908
- Specify rule validity times on page 912
- Validate rules automatically on page 912
- How default rules can be changed on page 917

The rules in Engine, IPS, Layer 2 Engine, and Layer 2 Interface Policies allow you to control how the engines inspect and filter network traffic, and how NAT (network address translation) is applied on Engines, Master Security Engines, and Virtual Engines.

Getting started with editing policies

Rules in policies are instructions to the engines for handling traffic.

What rules do

There are five main types of rules.

- Ethernet rules (IPS, Layer 2 Engine, and Layer 2 Interface Policies only) filter traffic based on MAC addresses and low-level network protocols. These rules can be used to segment a network.
- Access rules filter traffic based on IP addresses and IP-based protocols. These rules control access to resources. There are separate Access rules for IPv4 and IPv6 traffic.
- NAT rules (Engine Policy only) change source or destination IP addresses in traffic that is allowed to pass through the engine. NAT (network address translation) can hide the network structure and allows several computers to use the same IP address on the Internet. There are separate NAT rules for IPv4 and IPv6 traffic.
- Inspection rules in Inspection Policies filter traffic based on patterns in any of the information that is transferred. These rules log complex traffic use patterns and find network attacks, network worms, or other worrying or unwanted traffic like the use of peer-to-peer file transfer applications.

- Exceptions in Inspection Policies create detailed exceptions to the Inspection rules to eliminate false positives and to activate block listing or User Responses for specific traffic patterns.

The engines process the rules one type at a time in the order previously listed. IPv4 and IPv6 traffic can be matched to both IPv4 and IPv6 Access rules in any order if traffic is tunneled, possibly several times.

Basic rule design considerations

Keep the following in mind when editing rules:

- Rule tables are read from the top down, so the order of the rules is important. Make sure that the rules advance logically from specific rules at the top toward more general rules at the bottom whenever the matching criteria in rules overlap.
  Example: A rule that denies access to your server from a particular network must be placed above a general rule that allows access from any source address.
- Any two rules that have identical matching criteria are redundant and should be merged. Automatic rule validation can be used to find such mistakes.
- When rules are matched to traffic, the traffic is compared to each rule one by one until a match is found. What happens when the end of the rule table is reached without any matches varies by the component and the type of rules.
- If you use element-based NAT, the NAT rules generated from NAT definitions are applied only after the NAT rules that you have added manually to the policy. This means that the NAT rules that are generated from NAT definitions do not override the rules that you have manually added to the policy. Remember, however, that a more specific manually created NAT rule can prevent traffic from matching the automatically generated NAT rules.

What do I need to know before I begin?

There are different policy types for different Security Engine elements.

Policy types

Policy type	Security Engine elements
Engine Policy	Single Engine, Engine Cluster, Master Engine, Virtual Engine **Note** Master Engines always use Engine Policies regardless of the role of the Virtual Engines they host.
IPS Policy	Single IPS, IPS Cluster, Virtual IPS
Layer 2 Engine Policy	Single Layer 2 Engine, Layer 2 Engine Cluster, Virtual Layer 2 Engine
Layer 2 Interface Policy	Single Engine, Engine Cluster **Note** You select the Layer 2 Interface Policy for the Security Engine in the Engine Editor.

Related tasks

Create template policies or policies on page 808

The different parts of the policy editing view

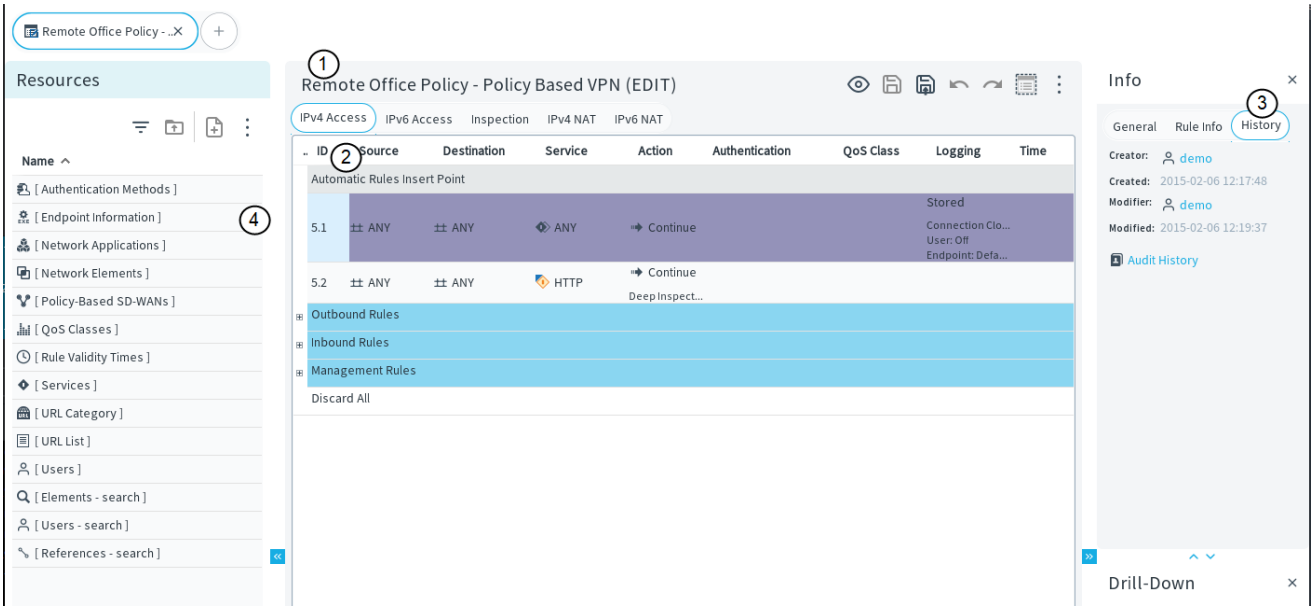
The tabs and options shown in the policy editing view depend on the type of policy you are editing.

Note

Only one administrator at a time can edit a policy. Save your changes and close the policy editing view when you are finished.

All policy editing views have by default the **Resources** pane on the left, the rule table on the right, and the policy toolbar at the top of the page. The available elements in the **Resources** pane can be used as matching criteria in different rule cells. There are tabs for the different rule types that are available for each type of policy.

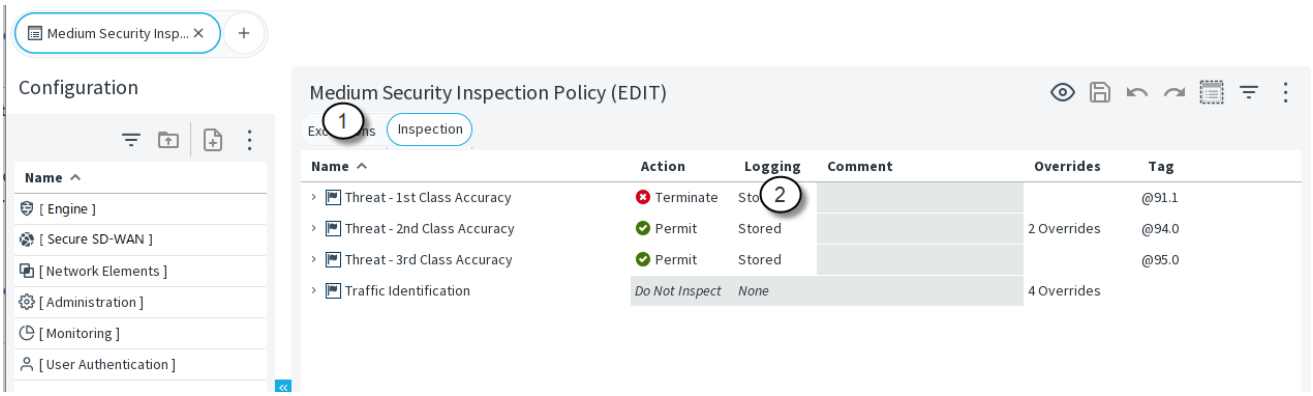
Policy editing view (IPv4 Access tab)



1. Policy toolbar
2. Rule table
3. History for selected rule
4. Search tool

In the Inspection Policy Editing view, there are only two tabs: **Exceptions** and **Inspection**. Global Inspection rules are configured on the **Inspection** tab, and exceptions to global Inspection rules on the **Exceptions** tab.

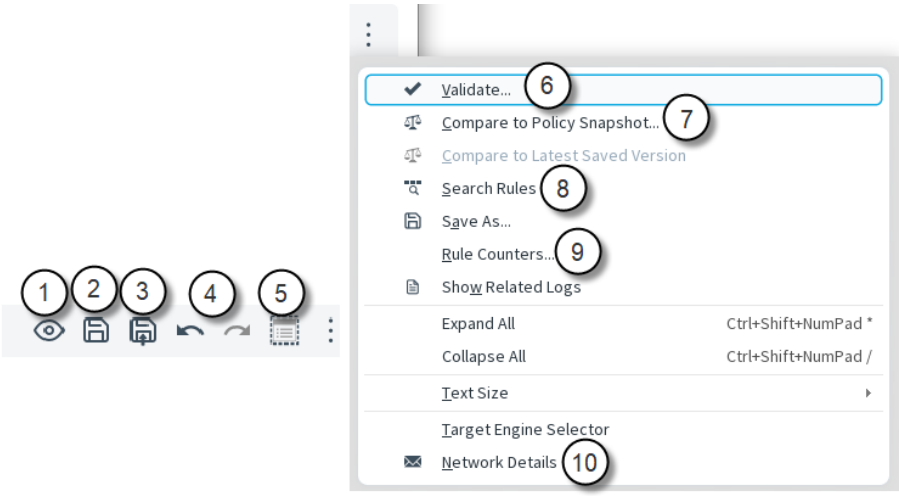
Inspection Policy Editing view



1	Detailed Exceptions to the main Inspection rules
2	The main rules tree

The policy toolbar contains tools for managing the policy.

Policy toolbar



1	Preview the policy in read-only mode
2	Save
3	Save changes and install policy on Security Engines
4	Undo or Redo
5	Show Inherited rules passed down from higher-level templates
6	Automatic validation finds rules that are clearly incorrect
7	A Snapshot is made at each policy installation to allow change tracking
8	Search tool for finding rules
9	Display the number of hits for each rule
10	Toggle between showing element names and IP addresses

Related tasks

Create template policies or policies on page 808

Editing rules in a policy

Editing rules consists of defining matching criteria for different rule cells and making sure that the rules are organized in a logical way in the rule table.

Editing rule tables

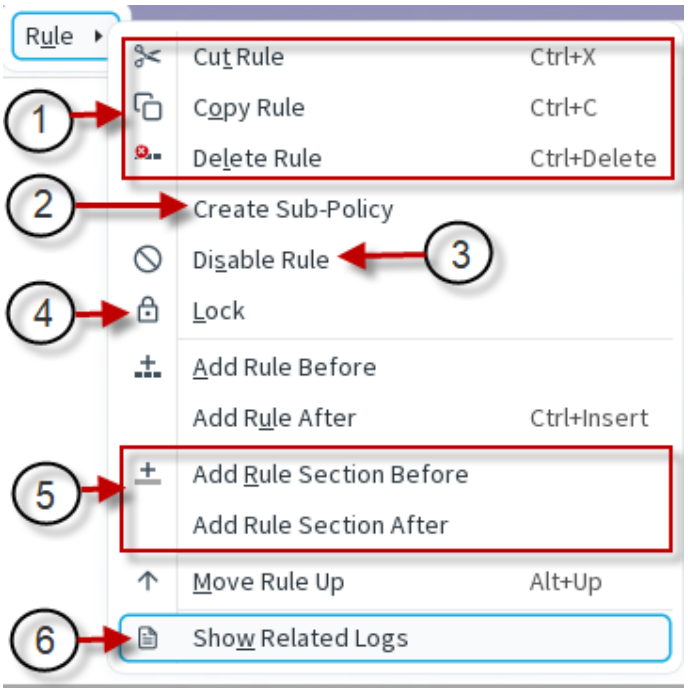
Use the right-click menu options to add, remove, and organize rules.

Use the actions in the right-click menu, such as add, cut, copy, and paste. If you right-click a cell that has cell-specific actions, the rule-specific actions are under the **Rule** submenu.

Tip

Drag and drop (move) whole rules by dragging the rule ID cell.

Example of a right-click menu for a rule



1. Standard editing actions.
2. Convert the selected rules into a Sub-Policy.
3. Temporarily disable rules without deleting them.
4. Prevents editing the rule until the rule is explicitly unlocked.
5. Create collapsible rule sections.
6. View logs related to the rule.

Editing rule cells

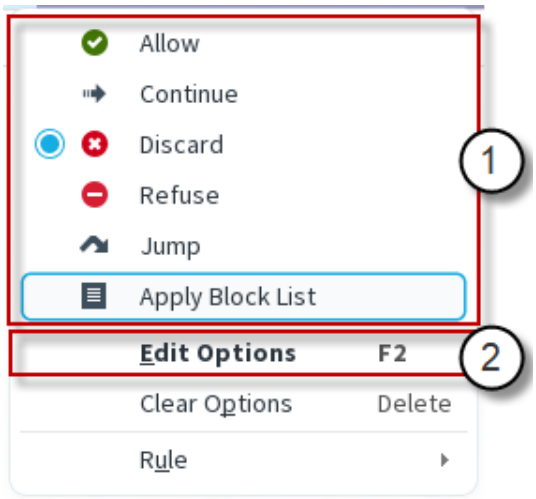
Rule cells define the matching criteria for the rules.

Most rule cells require you to insert elements of specific types:

- When you select the cell, the Resources pane on the left shows the types of elements that you can insert.
- You can drag and drop elements from the Resources pane, from another cell, or even between tabs.
- Define detailed sets of matching criteria in the Definitions dialog box for the Source, Destination, and Service cells.
- You can create new elements in the Resources pane.

To edit rule cells that do not accept elements, right-click the cell and select an item from the right-click menu.

Right-click menu for the Action cell in an Access rule



1 Main options for the **Action** cell.

2 Opens a dialog box that contains more settings for the selected Action.

Define Source, Destination, and Service criteria in rules

You can create detailed sets of matching criteria for the rule in the Source, Destination, and Service cells.

You can create Source and Destination Definitions for the following types of rules:

- All types of rules in Engine Policies.
- IPv4 and IPv6 Access rules in IPS, Layer 2 Engine, and Layer 2 Interface Policies.

The following types of items can be used as matching criteria:

Matching criteria for Source and Destination Definitions

User	IP Address	Domain Name	Zone
■ User names and groups of user names of users that have authenticated to the Engine. ■ User and User Group elements for users stored on an integrated Active Directory server in an environment with a Forcepoint User ID Service server installed and configured.	Any element from the Network Elements branch that directly represents an IP address.	Domain Name elements. If DNS Server IP addresses have been defined in the engine properties, the engine automatically resolves the Internet domain names to IP addresses.	Zone elements for interface matching

Note

VPN and NAT operations can change the routing of packets, potentially causing packets that no longer match the Destination Zone of an Access rule to be discarded.

You can create Service Definitions for the following types of rules:

- IPv4 and IPv6 Access rules, and IPv4 and IPv6 NAT rules in Engine policies.
- IPv4 and IPv6 Access rules in IPS, Layer 2 Engine Policies, and Layer 2 Interface Policies.

The following types of items can be used as matching criteria:

Matching criteria for Service Definitions

Network Application	Service (Port)	TLS Match
Network Application elements for application detection and application routing.	TCP and UDP Service elements In NAT rules that forward traffic to a proxy server, the supported protocols depend on the proxy server to which traffic is forwarded. If the row contains both a Network Application element and a Service element, the ports specified in the Service element override the ports specified in the Network Application elements. When the row contains a Network Application element, you can also specify which ports traffic matches without adding a Service element.	(IPv4 and IPv6 Access rules only) TLS Match elements for application detection. TLS Match elements must be used with Network Application elements that contain a TLS Match.

Note

You cannot use Network Application elements and Service elements on different rows of the same Service Definition.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click the **Source**, **Destination**, or **Service** cell, then select **Edit Source**, **Edit Destination**, or **Edit Service**.

2) For each row of matching criteria that you want to add:

   a) Click **Add Row**.

   b) Drag and drop elements from the list on the left to the correct cell in the row.

   Note

   All items on the same row must match the traffic for the row to match. You do not have to insert elements into all cells on the same row.

3) Click **OK**.

Related concepts

Enabling access control by user on page 1113
Defining Domain Name elements on page 920
How Access rules match traffic on page 835
Getting started with Network Application elements on page 961

Related tasks

Create TLS Match elements for network application detection on page 963

Adding comments in policies

You can add comments to the Comment cell in each rule or you can insert Rule Sections to add comment rows to the policy.

There are two ways to add comments in policies:

- The Comment cell in each rule allows you to write rule-specific comments, for example, to record why the rule was added. The History in the Info pane shows when the rule was added and last changed and through which administrator account. Double-click the cell to edit the comment text.
- In rule tables, you can insert Rule Sections to visually structure the policy under collapsible sections of rules that are preceded by a comment row. Double-click the row to edit the comment text. You can set each comment row's color through the **Colors** submenu in the comment row's right-click menu.

The maximum length of both types of comments is 4096 characters.

Rule identifiers and how they work

The **ID** and **Tag** for each rule are automatically created and updated. You can optionally specify a name for each rule.

Each rule has two non-editable identifiers:

- The **ID** cell shows the order of the rules. For example, the ID 14.1.2 shows that the rule is the second rule in the policy. It is in an insert point that is the first rule in the parent template. That insert point is the 14th rule in the top-level parent template. The number changes as you add, remove, and move rules.
- The **Tag** is the unique identifier of the rule in this policy. It contains a static part that does not change when rules are added, removed, or moved, and a changing part that indicates the version of the rule. For example, in Tag "@274.12", "274" is the unchanging part and "12" indicates that the rule is in its 12th revision. The tag is used, for example, to provide links from logs to rules that created the log entries.

In addition to non-editable identifiers, you can specify an optional name for each rule.

Naming rules

You can optionally add a name or short description to a rule to help identify it.



Note

You cannot add a name or short description to the rules on the Inspection tab of Inspection Policies.

In the policy editing views, the name is displayed in the Rule Name cell with the rule tag, and on the General tab of the Info pane. In the Logs view, in statistics, and in reports, the rule name is displayed instead of the rule tag in the Rule Tag column. If no name is specified, the rule tag is shown.

The maximum length of a name or description is 254 characters. The name does not have to be unique. Double-click the **Rule Name** cell in the editing view of a policy to edit the name text.

You can search for a specific rule by its name in the Search Rules view.

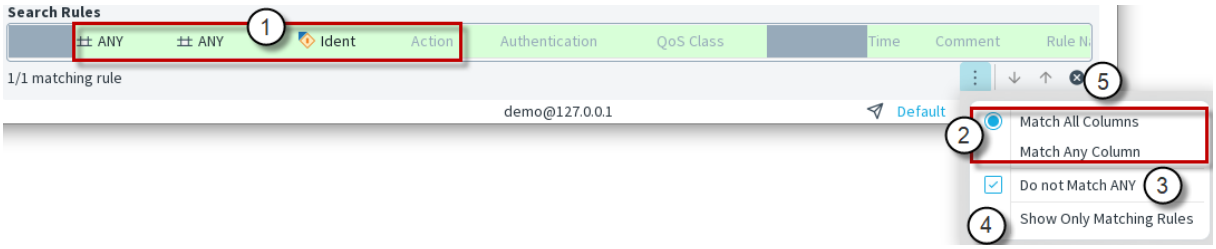
Searching in rules

You can search rules in rule tables and on the Inspection tab of Inspection Policies.

In rule tables, you can search rules based on most of the cells. Select ⋮ **More actions** > **Search Rules** to display the search at the bottom of the rule table.

In the Rules tree on the Inspection tab of Inspection Policies, you can search the Situations through type-ahead searching. When you type a part of the name of the Situation, the tree shows only matching Situations (and their parent Situation Types, which make up the tree branches). The currently active type-ahead search is shown at the bottom of the tree pane.

Rule search for rule tables



1. Drag and drop or enter the search criteria in selected cells.
2. Search rules that match all or any defined criteria.
3. Do not find rules that have ANY in a cell that is used as search criterion.
4. Hide rules that do not match the defined criteria.
5. Remove all search criteria.

You can add values in different ways:

- Drag and drop elements from the rule table, from different windows and tabs, or from the resource pane (shown in edit mode).
- Right-click a cell, then choose **Select** to browse for elements.
- In the **Source** and **Destination** search cells, you can manually type in IPv4 or IPv6 addresses, networks, or address ranges. Use standard notations (for example, 192.168.1.0/16, or 192.168.10.0 - 192.168.10.101 for IPv4 networks, or 2001:0db8:1234::/48, or 2001:0db8:1234:: - 2001:0db8:1234::100 for IPv6 networks.
- In the **Comment** and **Rule Name** search cells, you can manually type in any part of the comment or name.

Fill in the relevant cells. The cells you leave empty are ignored in the search. The first rule that matches your search is shown on a dark green background and all other matching rules are highlighted on a light green background. Click ↓ **Next** or ↑ **Previous** to move up or down from the currently selected rule to a matching rule above or below.

Run a rule counter analysis

Each rule contains a Hits cell that shows how many times each rule in your policy has matched network traffic. Viewing the rule hits allows you to find valid rules that match traffic that the engine does not encounter in the network.

This feature complements the rule validation checks, which can find rule design errors. Engines count rule hits automatically for all rules of supported types. The hits are stored as statistical counter data on the Log Servers.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Open the policy for preview or editing, then click the tab for the type of rules that you want to examine.

3) Select an engine from the **Target Engine** drop-down list.
   If the drop-down list is not visible, select ⋮ **More actions** > **Target Engine Selector**.

4) Select ⋮ **More actions** > **Rule Counters**.

5) From the **Period** drop-down list, select the period for which you want to check the rule matches.
   - Select one of the existing options.
   - To define a custom period, select **Custom**.

6) (Optional) Click **Add** to add other engines to the **Target** list.

   Tip

   You can run a rule counter analysis on several engines at the same time.

7) (Optional) To select Management or Log Servers for this operation, or to include archived data, click the **Storage** tab, then change the selection.
   Make sure that you include the Log Servers and folders that contain data for the target engine and the period you selected.

8) Click **OK** to display the rule hits. The Hit information is displayed until you close the view.
   The Hit cell in each rule is filled in with the number of connections that matched the rule during the chosen period.
   If there is no statistical information about the rule with your selected criteria, the Hit cell shows N/A (for example, for rules added after the period analyzed).

Related tasks

Validate rules automatically on page 912

Add Insert Points in Policy Templates

Insert Points mark the positions where rules can be added.

When you edit a Template Policy, add at least one new yellow Insert Point on all tabs if you want their inheriting Policy or Policy Template to be editable. Green Insert Points are inherited insert points from the previous level, and they are not inherited further down in the hierarchy. They only show you where the higher-level template allows you to add rules and disappear when you add a rule or a new (yellow) insert point.

You can add as many insert points in the Template as your rule structure calls for.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Open the Policy template for editing.

3) Right-click the green insert point and select **Add Insert Point** or a rule that is editable in this Template and select **Rule** > **Add Insert Point Before** or **Add Insert Point After**.

4) Give the insert point a descriptive **Name** and click **OK**. An inheritable (yellow) insert point is added to the Template.

Related tasks

Create template policies or policies on page 808

Automatic rules and how they work

When you enable a feature that requires traffic between certain components to be allowed, rules allowing the traffic are automatically created.

Automatic rules are created for traffic to and from the engine, never for traffic that passes through the engine. Some features require more specific control over what traffic is allowed between specific components, and in those cases you still have to configure Access rules manually.

Automatic rules are not visible in rule tables, but you can view a summary of currently used Automatic rules in the **Automatic Rules** section of the Engine Editor. You can also change some settings for Automatic rules in the Engine Editor.

Automatic rules are only created if the policy that is installed on the engine contains the Automatic Rules Insert Point. The default Template Policies in the SMC Client already contain this insert point. No further action is needed for Automatic rules to be created if you base your Template Policies and Security Policies on the recommended default Template Policies.

Configure settings for Automatic rules

View a summary of Automatic rules and manage related settings in the Engine Editor.

Before you begin

The Template Policy used on the engine must contain the **Automatic Rules Insert Point**.

To add the Automatic Rules Insert Point manually, open the Template Policy for editing, right-click the **ID** cell of any rule, and select **Add Automatic Rules Insert Point**. You can add the Automatic Rules Insert Point anywhere in the Template Policy.

In the Automatic Rules section of the Engine Editor, you can set the log level and possible Alert element for Automatic rules. For Engines, Virtual Engines, and Master Security Engines, you can also define whether traffic from the engine to authentication ports is allowed.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click an engine, then select **Edit <element type>**.
3) Browse to **Policies** > **Automatic Rules** in the navigation pane on the left.
4) Configure the settings.
5) Click **Save**.

Note

If you enable the ZTNA connector in the Engine Editor, you will view "Traffic from the Engine to the Forcepoint ONE ZTNA connection". For more information, see *Engine Editor* > *Add-Ons* > *ZTNA Connector*.

Add Ethernet rules

Ethernet rules define which Ethernet protocol packets the engines stop, and which packets are allowed through.

Ethernet rules are used in IPS, Layer 2 Engine, and Layer 2 Interface Policies.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Expand the **Policies** branch and select one of the following types of policies:

- **IPS Policies**
- **Layer 2 Engine Policies**
- **Layer 2 Interface Policies**

3) Open a Template Policy or Policy for editing, and switch to the **Ethernet** tab.

4) Add the rule in one of the following ways:
   - Right-click the **ID** cell of an existing rule and select **Add Rule Before** or **Add Rule After**.
   - Copy and paste an existing rule.

5) Specify the matching criteria.

6) Right-click the Action cell, then select the Action.

7) (Optional) Define options for triggering logs and alerts.

Define logging options for Ethernet rules

Ethernet rules can create a log or alert entry each time they match.

By default, logging options set in a previous rule with Continue as its action are used. If no such rule exists, the default logging options defined in the template policy are used.

- Layer 2 physical interfaces on Engines log connections by default.
- Layer 2 Engines and Virtual Layer 2 Engines log connections by default.
- IPS engines and Virtual IPS engines do not log connections by default.

Each individual rule can be set to override the default values.

When the Log Server is unavailable, log entries are temporarily stored on the engine. When the engine is running out of space to store the log entries, it begins discarding log data in the order of importance. Monitoring data is discarded first, followed by log entries marked as Transient and Stored, and finally log entries marked as Essential. The Alert entries are the last log entries to be discarded.

The settings for storing the logs temporarily on the engine are defined in the engine's log spooling policy.



Note

A log entry is generated for each packet that matches an Ethernet rule. Use careful consideration when setting the logging options to avoid producing an excessive amount of log data.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Double-click the **Logging** cell in the rule.

2) Define the options.

Related tasks
Configure log handling settings on page 685

Define MAC addresses for Ethernet rules

MAC Address elements are used to match a certain source or destination MAC address in the Ethernet rules.

The MAC Address element defines the MAC (Media Access Control) address of a network card.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Expand the **Other Elements** branch.
3) Right-click **MAC Addresses** and select **New MAC Address**.
4) **Name** the element.
5) Enter the **MAC Address**. You can enter any valid MAC address including, for example, the broadcast address (ff:ff:ff:ff:ff:ff).
6) Click **OK**.

Add Access rules

Access rules are used in Engine, IPS, Layer 2 Engine, and Layer 2 Interface Policies.

Before you begin

You must have a custom Policy element and permissions to edit it.

Engines, IPS engines, Layer 2 Engines, Virtual Engines, Virtual IPS engines, and Virtual Layer 2 Engines use both IPv4 Access rules and IPv6 Access rules. Master Engines only use IPv4 Access rules.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Expand the **Policies** tree and select a type of policy (for example, **Engine Policies**).

3) Open a Template Policy, Policy, or Sub-Policy for editing.

4) On the **IPv4 Access tab** or to the **IPv6 Access** tab, add a rule in one of the following ways:
   - Right-click the **ID** cell of an existing rule and select **Add Rule Before** or **Add Rule After**.
   - Copy and paste an existing rule.

5) Specify the matching criteria.

6) Right-click the Action cell, then select the Action.

7) (Optional) Define options for triggering logs and alerts.

Define Action options in Access rules

Action options define additional specific options for various features.

If no options are specified, the settings defined in Continue rules higher up in the policy are used.

- **Allow** — You can define what traffic to allow through the Engine. Also, this option lets you configure the following:
  - Forward traffic to a proxy, a host, or into a VPN.
  - Force forward the matched traffic to a preferred destination.
  - Control stateful inspection by setting options for connection tracking, including idle timeouts and TCP segment size enforcement.
  - Enable or disable rate-based DoS protection and scan detection.
  - You can enable or disable deep inspection for the matched traffic against an Inspection Policy. This includes the following:
    - You can configure deep inspection and anti-malware options to check IPv4 traffics for malware.
    - By default, the deep inspection is enabled for all supported protocols with Continue rules if you use the IPS Template or Layer 2 Firewall Template as the base for a policy.
    - You can disable deep inspection for a specific rule.

    Note

    If deep inspection is not disabled, make sure that the custom template policy directs all necessary protocols for inspection.
- **Continue** — You can set the default options for multiple rules. Options specified in the Continue rule are applied to any other Access rule that the same packet matches. However, if the Access rules have rule-specific definitions, those will be used instead.
- **Discard** — You can define a User Response to be shown to the user when an HTTP connection is discarded.
- **Refuse** — You can define a User Response to be shown to the user when an HTTP connection is refused.
- **Jump** — The rule processing jumps to a Sub-Policy to continue processing rules.
- **Apply Block list** — You can configure options that affect the reception of block list entries.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Browse to **Policies** > **<Policy type>**.
3) Right-click a policy, then select **Edit <Policy name>**.
4) Right-click the **Action** cell in an Access rule, then select **<Action>**.
5) Right-click the **Action** cell, then select **Edit Options**.
6) Configure the settings, then click **OK**.

Related concepts

Configuring default settings for several Access rules on page 839
Getting started with example VPN configurations on page 1223
Getting started with anti-malware scanning on page 989
Configuring connection tracking on page 817

Define logging options for Access rules

Access rules can create a log or alert entry each time they match.

By default, logging options set in a previous Access rule with Continue as its action are used. If no such rule exists, Engines, Virtual Engines, Layer 2 Engines, and Virtual Layer 2 Engines log the connections by default. IPS engines and Virtual IPS engines do not log the connections by default. Each individual rule can be set to override the default values.

Note

Log pruning might override the logging options by deleting any number of generated log entries when they are received at the Log Server.

Logging for the closing of the connection can be turned on or off, or on with accounting information. You must collect accounting information if you want to create reports that are based on traffic volumes.

When the Log Server is unavailable, log entries are temporarily stored on the engine. When the engine is running out of space to store the log entries, it begins discarding log data in the order of importance. Monitoring data is discarded first, followed by log entries marked as Transient and Stored, and finally log entries marked as Essential. The Alert entries are the last log entries to be discarded. The settings for storing the logs temporarily on the engine are defined in the engine's log spooling policy.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Double-click the **Logging** cell.

2) Set the options.

3) Click **OK**.

Related concepts
Enabling access control by user on page 1113

Related tasks
Configure log handling settings on page 685

Define Authentication options for Engine Access rules

The Authentication options define which users can authenticate and the type of authentication required.

A mobile VPN always requires some form of authentication, but you can also add an authentication requirement to non-VPN rules.

The authentication requirements are configured in the **Authentication** cell. The cell accepts User and User Group elements to define the end users who are allowed to make connections allowed by the rule, and Authentication Method elements to define the type of authentication required for connections that match the rule.

If the authentication fails, the connection is discarded. If the authentication succeeds, the connection is allowed through.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Double-click the **Authentication** cell.

2) Configure the settings, then click **OK**.

Related tasks
Create User Group elements on page 1120
Create User elements on page 1120

Add NAT rules

NAT rules define how NAT is applied to matching connections.

Before you begin

You must have a custom Policy element and permissions to edit it.

NAT rules are only available in Engine Policy elements. NAT is not supported for layer 2 physical interfaces on Security Engines in the Engine/VPN role role.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Add the NAT rule in one of the following ways:
   - Right-click the **ID** cell of an existing NAT rule and select **Add Rule Before or Add Rule After**.
   - Copy and paste an existing NAT rule.
   - Copy and paste an Access rule to match the rule to the same Source, Destination, and Service.
2) Specify the portion of the traffic to which you want to apply NAT.
3) Define the translation you want to apply.

Translate source addresses in packets

There are two types of source address translation: static source translation and dynamic source translation.



Tip

With element-based NAT, you do not need to use separate static source and static destination NAT rules. Static NAT is bidirectional.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Double-click the **NAT** cell in the NAT rule.
2) Select the translation type.
3) If you selected an address translation operation, configure the additional options according to the type of operation.

Translate destination addresses in packets

Destination translation is typically used to translate new incoming connections from a server's public IP address to the server's private IP address.

You can also use destination translation to forward traffic to a proxy server.

Note

Destination translation can change the routing of packets and potentially cause packets that no longer match the Destination Zone of an Access rule to be discarded.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Double-click the **NAT** cell in the NAT rule.
2) On the **Destination Translation** tab, select the translation type.
3) Configure the options according to the selected translation type.
4) Click **OK**.

Related concepts

How Access rules match traffic on page 835

Add Inspection rules

Inspection rules filter traffic based on traffic patterns. Inspection rules are stored in Inspection Policy elements.

Before you begin

You must have a custom Policy element and permissions to edit it.

The rules tree on the **Inspection** tab is the main tool that allows you to select which traffic patterns are permitted and stopped. You can also select whether a log entry or an alert is triggered, and whether matching traffic is recorded.

The rules table on the **Exceptions** table allows you to define detailed exceptions to the Inspection rule. The main uses for Exceptions are to eliminate false positives and to activate block listing or User Responses for specific traffic patterns.

Note

For layer 2 physical interfaces on Security Engines in the Engine/VPN role, you select the Inspection Policy in the Engine Policy.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Browse to **Policies** > **Inspection Policies**.
3) Right-click a Template Inspection Policy or Inspection Policy element, then select **Edit Inspection Policy**.
4) On the **Inspection** tab, adjust the rules.
5) (Optional) On the **Exceptions** tab, define exceptions.

Inspection rules tree and how it works

The rules tree is the main tool for controlling deep packet inspection in Inspection Policy elements.

The rules tree on the **Inspection** tab in Inspection Policies allows you to define what action the engine takes when Situation matches are found and how they are logged. To edit these rules, click the Action cell or the Logging cell of a rule and select the suitable option. The definitions on the **Exceptions** tab are matched before the Rules tree on the **Inspection** tab.

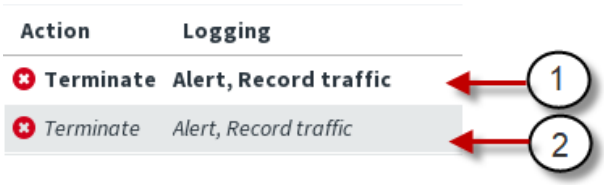
In the rules tree, items that have subitems are Situation Type elements. The items that have no subitems are individual Situation and Correlation Situation elements. The rules tree contains all Situation Types and the Situations associated with them.

All levels of the rules tree are editable. By default, subitems inherit the Action and Logging options from their parent item. If a subitem has any setting that differs from the parent item's settings, this is regarded an *override*. If you change a value in an item that has subitems, all subitems that are set to use the default value inherit this change. Any subitems that are set to an override continue to use that override.

Example

The parent item and 10 of the subitems are set to use the “Permit (Default)” action. Two of the subitems are set to use the “Permit” action. You change the parent to use the “Terminate” action. Ten subitems change to “Terminate (Default)”. Two subitems continue to use “Permit”.

How overrides are highlighted in the rules tree.



1 Overrides
2 Default values shown in italics

In the list of options available in the right-click menu, “default” is included in the label. For example, “Permit (default)” means that this action is the default action for the selected Situation Type or Situation.

Regardless of the settings in the rules tree in a higher-level Template Policy, it is still possible to change any rules tree values in the inheriting policy. To add to a Template Policy rules that cannot be edited in the inheriting policies, add the rules as Exceptions.

Related tasks

Configure log handling settings on page 685

Add Situations to the Inspection rules tree

The Inspection rules tree can contain a maximum of one instance of each Situation to prevent the definitions within the Rules tree from overlapping.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select the most appropriate Situation Type (these are all included at some level of the tree) as the **Situation Type** in the properties of a custom Situation element.

Related tasks

Create custom Situation elements on page 943

Remove overrides from the Inspection rules tree

You can reset branches in the rules tree to their default values.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click an item in the tree and select **Reset Branch**.
   The item you right-clicked and all its subitems are reset to the Default value.

Define what action Inspection rules take

The Action defines the command for the engine to carry out when a connection matches the Inspection rule.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click the **Action** cell and select the **Action Options** from the context menu to display the Action Options dialog box for the following actions:
   - Permit
   - Terminate

Define Action options for the Permit action in Inspection rules

The options for the Permit action in Inspection rules allow you to set options for traffic that has been allowed.

On the Engine, the options for the Permit action in Inspection rules allow you to control the inspection options in further detail and set a User Response for malware scanning or Situation matches.

Note

Malware scanning is not supported on Virtual Engines.

The options for the Permit action in Inspection rules allow you to generate block list event that block list executors can act on.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click the **Action** cell and select **Action Options** from the context menu.
   Alternatively, you can click the **Action** cell to open a combo-box. Scroll-down the combo-box to select **Action Options** option.

Define Action options for the Terminate action in Inspection rules

The Terminate action options control connection termination, notifications, and the creation of block list entries.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click the **Action** cell and select **Action Options** from the context menu.
   Alternatively, you can click the **Action** cell to open a combo-box. Scroll-down the combo-box to select **Action Options** option.

Define logging options for Inspection rules

Inspection rules can create a log or alert entry each time they match.

By default, an Inspection Policy uses the logging options set in a previous Exception rule with Continue as its action. If no such rule exists, Engines, Virtual Engines, Layer 2 Engines, and Virtual Layer 2 Engines log connections by default. IPS engines and Virtual IPS engines do not log connections by default.

Each individual Inspection rule can be set to override the default values of the engine role.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Switch to the **Inspection** tab.

2) Click the **Logging** setting of a rule and select **Logging Options**.

3) Set the options.

Note

Storing or viewing the packets' payload can be illegal in some jurisdictions due to laws related to the privacy of communications.

Related tasks

Configure log handling settings on page 685

Add Exception rules

Inspection Exceptions allow you to make changes to the Inspection Policy that are not applied to all connections.

Exception rules also allow you to set some options (using the Continue action) for Exceptions and rules that are processed later. Exception rules also contain some additional options that are not available in the rules tree.

- You can match specific connections based on the IP addresses of the communicating hosts, the Service used, and the Logical Interfaces of IPS engines and Layer 2 Engines. For example, an Exception can be used to eliminate a false positive in traffic between two internal hosts without disabling inspection.
- You can set more responses to matches that are found. You can block list connections on an engine, and you can add User Responses as notifications to some types of events.

Note

Inspection Policies are not supported for layer 2 physical interfaces on Security Engines in the Engine/VPN role.

Related tasks

Create template policies or policies on page 808

Add Exception rules in Inspection Policy elements

The **Exceptions** tab allows you to create detailed rules that are processed before the Rules tree definitions on the Inspection tab.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Add the Exception rule in one of the following ways:
   - Right-click a generated log entry and select one of the options in the **Create Rule** submenu. The rule is created as an Exception with matching details from the log entry.
   - Right-click the **ID** cell in an existing Exception and select **Add Rule Before** or **Add Rule After**.

- Copy and paste an existing Exception rule.
- Copy and paste a rule from the Inspection tab to the Exceptions tab to match the same Situations and options.
- Copy and paste an Access rule to match the Exception to the same Source, Destination, and Service.

2) Match the Exception rule to traffic.

3) Define the exception you want to apply.

4) (Optional) Define options for triggering logs and alerts.

Define what traffic Exception rules match

Exception rules in Inspection Policy elements are matched based on the patterns defined in Situation elements.

The traffic is checked against all patterns in the Inspection Policy. When a match is found, the Situation element is used in looking up the rule that determines what happens to the traffic. If none of the Exceptions match, the matching continues in the Rules tree.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Fill in the cells.

   Setting a value in the **Time** cell is optional. All other cells must always contain a value.

Related concepts

Defining IP addresses as elements on page 919
Getting started with directory servers on page 1103

Related tasks

Add logical interfaces on page 592
Define Source, Destination, and Service criteria in rules on page 891
Specify rule validity times on page 912

Define what action Exception rules take

The Action defines the command for the engine to carry out when a connection matches the Exception rule.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click the **Action** cell and select the action.

Define Action options for the Continue action in Exception rules

The Continue action can set options for the Permit and Terminate actions in subsequent rules.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click the **Action** cell and select **Continue**.
2) Double-click the **Action** cell.
3) Set the options, then click **OK**.

Define Action options for the Permit action in Exception rules

The options for the Permit action in Exception rules allow you to set additional options for traffic that has been allowed.

On the Engine, the options for the Permit action in Exception rules allow you to control the inspection options in further detail and set a User Response for malware scanning or Situation matches.

Note

Malware scanning is not supported on Virtual Engines.

On the IPS and the Layer 2 Engine, the options for the Permit action in Exception rules allow you to block list traffic.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click the **Action** cell, then select **Permit**.
2) Double-click the **Action** cell.
3) Set the options, then click **OK**.

Define Action options for the Terminate action in Exception rules

The Terminate action options control connection termination, notifications, and the creation of block list entries.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click the **Action** cell and select **Terminate**.
2) Double-click the **Action** cell.
3) Select the action options, then click **OK**.

Related concepts

User Response elements and how they work on page 969

Define logging options for Exception rules

Inspection Exception rules can create a log or alert entry each time they match.

Engines, Virtual Engines, Layer 2 Engines, and Virtual Layer 2 Engines log connections by default. You can override the default logging options in an Exception rule with Continue as its action. IPS engines and Virtual IPS engines do not log connections by default.

Each individual Exception rule can be set to override the default values of the engine role.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Switch to the **Exceptions** tab.
2) Double-click the **Logging** cell of an Exception rule.
3) Set the options, then click **OK**.

Note

Storing or viewing the packets' payload can be illegal in some jurisdictions due to laws related to the privacy of communications.

Related concepts

Add Exception rules on page 908

Related tasks

Configure log handling settings on page 685
Add Situations to the Inspection rules tree on page 906
Remove overrides from the Inspection rules tree on page 906

Specify rule validity times

You can specify when Access rules, inspection Exception rules, and rules in Alert policies are enforced.

You specify rule validity times using Rule Validity Time elements. Rule Validity Time elements allow you to:

- Specify when each rule starts being enforced, and when each rule automatically expires.
  When a rule automatically expires, traffic can no longer match the rule. The rule is disabled, but it is still shown in the policy. If you do not specify when a rule expires, the rule never expires.
- Specify when each rule is active.
  For example, you might create a rule that allows access only during business hours on weekdays. If you do not specify when a rule is active, the rule is always active.

You can use the same Rule Validity Time element in multiple rules and policies.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Browse to **Other Elements** > **Rule Validity Time**.
3) Right-click **Rule Validity Time**, then select **New Rule Validity Time**.
4) Configure the settings, then click **OK**.
5) Drag and drop the Rule Validity Time element to the **Time** cell in a rule.

Validate rules automatically

You can automatically validate the rules in a policy at any time.

You can also validate the policy when you install or refresh the policy on an engine. In both cases, you can also select which issues are checked in the policy.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Start the policy validation in one of the following ways:
   - If the policy is open in the Policy Editing view, select ⋮ **More actions** > **Validate**.

- If you are installing or refreshing a policy and the Task Properties dialog box is open, make sure that the **Validate Policy before Upload** option is selected, then click **Select Settings**.

2) (Optional, available in the Policy Editing view) Select the **Target** engine on which you want to install the policy to get more accurate results.
   - The Target engine selection is used to resolve Alias elements when a policy is validated.
   - If no Target engine is selected, all issues related to the engine configuration cannot be checked (for example, parts of the VPN configuration).

3) (Optional) Edit the **Validation Settings** (the types of issues that are checked).

4) (Optional) Click **Save as Default** if you want to save the selected settings as the default set for future policy validations.

5) Click **OK**.



Tip

Any issues that are found are displayed in the Issues pane.

Start the policy validation

You can validate the rules in the Policy Editing view or when you install or refresh the policy on an engine.

In both cases, you can also select which issues are checked in the policy.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Start the policy validation in one of the following ways:
   - If the policy is open in the Policy Editing view, select ⋮ **More actions** > **Validate**.
   - If you are installing or refreshing a policy and the Task Properties dialog box is open, make sure that the **Validate Policy before Upload** option is selected, then click **Select Settings**.

2) (Optional, available in the Policy Editing view) Select the **Target** engine on which you want to install the policy to get more accurate results.
   - The Target engine selection is used to resolve Alias elements when a policy is validated.
   - If no Target engine is selected, all issues related to the engine configuration cannot be checked (for example, parts of the VPNs configuration).

3) (Optional) Edit the **Validation Settings** (the types of issues that are checked).

4) (Optional) Click **Save as Default** if you want to save the selected settings as the default set for future policy validations.

5) Click **OK**.

Result

Any issues that are found are displayed in the **Issues** pane.

Override default validation settings for rules

You can define rule-specific settings for policy validation.

The rule properties allow you to view some rule-specific information and select the settings that are applied to the selected rule when the policy is validated. The rule-specific settings override the default validation options for your administrator account.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Double-click the rule's **ID** cell.

2) Click the **Validate** tab.

3) Edit the **Validation Settings** (the types of issues that are checked).

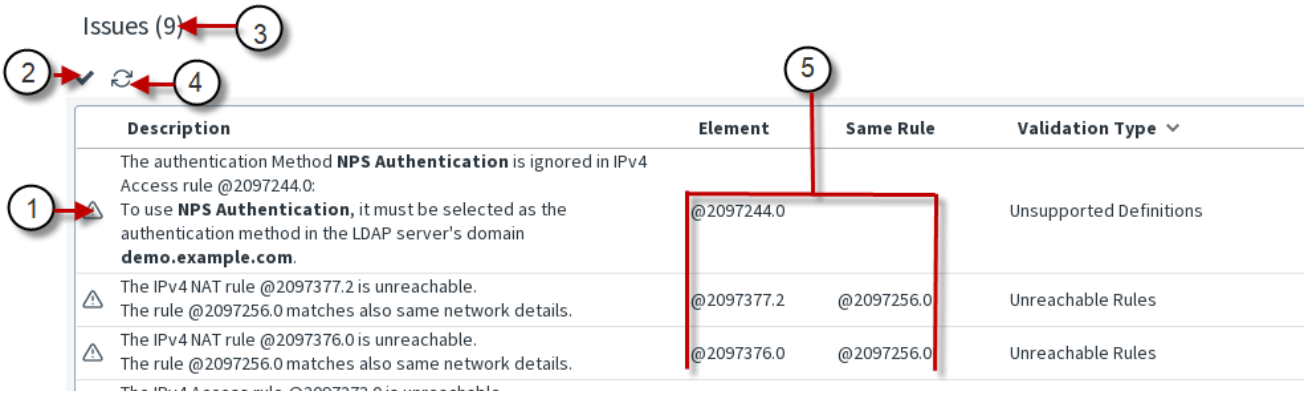
4) Click **OK**.

   The selected Validation Settings are now applied to this rule when you next validate the policy. A green checkmark is added to the rule's **ID** cell in the rule table. This indicates that the Validation Settings of the rule are different from those of the whole policy.

Viewing policy validation issues

If policy validation finds issues, the issues are displayed in the **Issues** pane in the **Policy Editing** view or on the tab that shows the progress of installation.

Issues pane



1. Descriptions of found issues
2. Validate policy and select validation properties
3. Number of found issues
4. Re-validate policy with the same validation settings
5. Rules with issues are listed by rule name

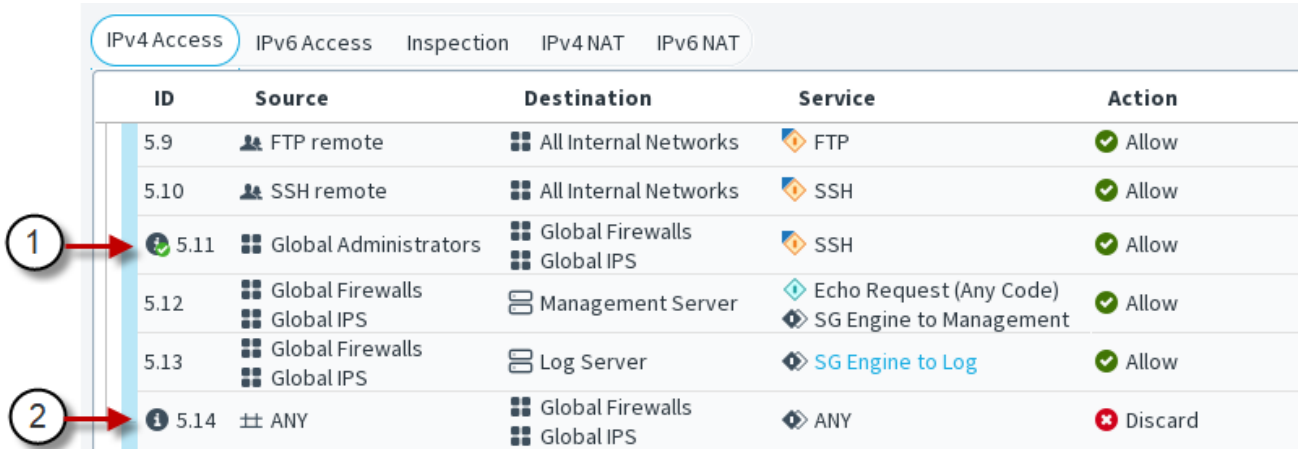
View the configuration in which a validation issue was found

You can view the part of the configuration that caused a policy validation issue.

In rules, the **ID** cell shows the status of validation issues in the rule.

- If issues are found for a rule, the rule's **ID** cell contains ⚠.
- If a rule's Validation Settings override the Validation Settings of the whole policy, the rule's **ID** cell contains ✔.

Rules with validation issues or rule-specific validation settings



1 Rule with modified Validation Settings

2 Rule in which a validation issue was found

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Double-click the issue. The relevant part of the configuration is shown.

2) Fix the issues that are indicated.

Disable validation warnings for rules

You can optionally set an issue listed in the Issues pane to be ignored for a specific rule.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click the rule in the Issues pane and select **Ignore Issue Type for Rule**.
   This issue type is no longer checked for this rule. A green checkmark is added to the rule's ID cell in the rule table. It shows that the rule's Validation Settings are not the same as the policy's Validation Settings.
   You can change the overall selection of validation issues for a rule in the rules's properties.

Exclude rules from policy validation

If you do not want to validate a certain rule in the policy, you can exclude the rule from policy validation.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click the rule in the Issues pane and select **Disable All Issues for the Rule**.

   The rule is no longer checked when you validate the policy. A green check mark is added to the rule's **ID** cell in the rule table. It shows that the rule's Validation Settings are not the same as the policy's Validation Settings.

How default rules can be changed

Rules inherited from the default Template policies cannot be edited directly. Instead, you can create a copy of the default Template policies.

In most cases, templates and policies are inherited from the default Template policies. However, it is not possible to edit these system elements. If you must edit the default templates, you can create a copy.

Version upgrades and dynamic updates might require changes to the default Template policies. These changes are not applied to any copies of the templates. You must manually edit the copies to make sure that the system communications continue to be allowed and all necessary inspection continues to be performed.

Create a custom version of a default Template policy

It is not possible to directly edit a default Template policy. Instead, you must create a copy of the default Template policy.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click the template and select **New** > **Duplicate**.

2) Save the copy of the template under a different name.

3) Right-click existing policies that you want to use the edited template and select **Properties**.

4) Select your copy of the template in the **Template** pane and click **OK**.

CAUTION

Incorrect modifications to the default Template policies can seriously disturb the operation of the engines.

Chapter 54

Defining IP addresses

Contents

- Defining IP addresses as elements on page 919
- Access and modify network elements on page 924
- Edit Expression elements on page 924
- Using SMC elements to represent IP addresses in policies on page 925

When you define IP addresses as elements, you can use the same definitions in multiple configurations for multiple components.

Defining IP addresses as elements

There are several types of elements in the SMC that represent IP addresses.

The elements that you can use for defining IP addresses are called network elements (not to be confused with the Network element, which defines an IP network). Each element can be inserted in several places in the Access, Inspection, or NAT rules (as source and destination of traffic). Network elements are also used in many other places where you have to define IP addresses (for example, in routing and log filtering).

The primary tools for defining IP addresses are elements in the SMC whose only role is to define an IP address. But elements created for configuring a feature in the SMC can also be used in policies (with some limitations) if they represent an IP address.

Different types of elements allow you to flexibly define any set of IP addresses:

- *Address Range* elements allow you to define any continuous range of IP addresses.
- *Alias* elements represent a variable value in policies. The IP address value is filled in based on the engine on which the policy is installed. Aliases make using the same policy on several engines practical.
- *Country* elements contain lists of IP addresses that are registered in a particular country.
- *Domain Name* elements represent all IP addresses that belong to a particular domain.
- *Expression* elements allow you to define any set of IP addresses in a single element. They are especially suited for excluding some IP addresses from otherwise continuous ranges.
- *Group* elements allow you to combine different types of elements into a single element.
- *Host* elements represent a single device in the network. Each Host can represent one or more individual IP addresses in policies.
- *IP Address List* elements contain IP addresses.
- *Network* elements represent complete network segments.
- *Router* elements represent a gateway device in the network and are primarily meant for configuring routing. Each Router can represent one or more IP addresses in policies.
- *Zone* elements allow you to combine engines' network interfaces into a single element.

Defining Address Range elements

An Address Range element can specify any continuous range of IP addresses.

On the Address Range element's **NAT** tab, you can view or edit the NAT definitions for the NAT configuration in which the element is included. However, you primarily configure NAT definitions in the properties of a Engine.

Defining Alias elements

Alias elements represent different IP addresses depending on the engine on which they are used.

Alias elements are like variables: they can be given different IP address values in the policy depending on the engine on which the policy is installed. This makes it possible to create rules that are valid on several engines without including all IP addresses in the policies of all elements. Alias elements are especially useful in Template Policies.



Tip

You can view the IP addresses that the Alias represents in the policy of each component: in a policy, select ⋮ **More actions** > **Network Details**. Then select a component from the list that is added to the toolbar.

Some of the default system Aliases always receive their values directly from other parts of each engine's configuration and cannot be edited. These Aliases start with two $$ symbols. There are also some default Aliases that either do not receive any value without your action, or allow you to add to and change the default values. These Aliases start with one $ symbol, as do all Alias elements you create yourself.

Defining Country elements

Country elements are IP address lists based on country-level geolocation information. They are grouped within continents.

Country elements can be used to filter traffic in Access rules, based on the source or destination country or an entire continent. They can also be used in NAT rules, Inspection rules, and File Filtering rules.



Note

You cannot edit or create Country elements or continents. Country elements are system elements that are imported and updated when you activate new dynamic update packages.

Defining Domain Name elements

A Domain Name element represents all IP addresses that belong to a particular domain.

If you have entered the IP addresses of one or more DNS servers in the engine properties, the Engine, IPS, and Layer 2 Engines periodically query the DNS server to automatically resolve domain names to IP addresses. The use of DNS servers makes it possible to create rules that are valid even if new addresses are added to the domain or the domain's IP addresses change.

If the DNS server returns multiple IP addresses for the same domain name, the engine associates all IP addresses with the domain name. However, if there are a many IP addresses associated with the same domain

name, the DNS server might only reply with a few of the IP addresses at a time. In this case, the engine might need to make more queries to the DNS server to resolve all IP addresses for the domain name.

By default, the engine queries the DNS server every six minutes. Resolved IP addresses are kept in the engine's DNS cache for a maximum of one hour by default.

Note

The DNS cache is not synchronized between nodes of a cluster. Each node separately queries the DNS server using the node's NDI address. It is possible that the DNS cache might be different on different nodes of a cluster.

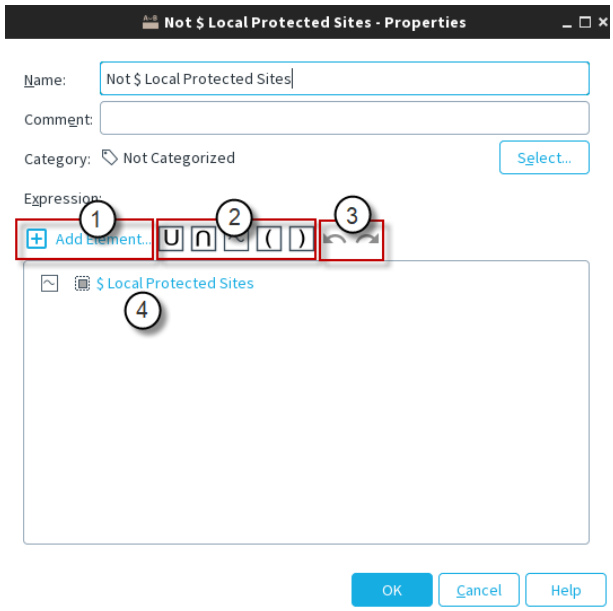
Related concepts

Creating Security Engine elements on page 504

Defining Expression elements

Expression elements allow you to combine other network elements with logical operators to represent complex sets of IP addresses.

Expression element properties



1. Click **Add Element** to select an element to use in the Expression.
2. Logical operators
3. Undo and redo actions
4. The Expression

Logical Operators in Expressions are resolved in the following order (however, items inside parentheses are evaluated before any items outside parentheses are considered):

1) Negations

2) Intersections

3) Unions

Defining Group elements

The Group element can be used to combine any number of previously defined elements into a single element.

The elements can be of different types. You can use Group elements in policies to make the policies clearer to read and to simplify the editing of configurations where the same elements always appear together. You can also use Groups to add monitored elements that are displayed in the **Dashboard** view.

Example: Host elements for file servers could be gathered together in a Group element, which is used in several rules in different policies. When a new file server is introduced, it is added to the Group. The change is propagated to all rules in all policies in which the Group is used.

On the Group element's **NAT** tab, you can view or edit the NAT definitions for the NAT configuration in which the element is included. However, you primarily configure NAT definitions in the properties of a Engine.

Defining Host elements

A Host element represents the IP addresses of any single device.

Host elements are used to represent individual devices that have no additional special role in the SMC configuration (such as being a next-hop router or an external authentication server).

You can optionally configure on the **Monitoring** tab that a Log Server monitors the device.

If you selected any of the options on the **Monitoring** tab, a new Monitoring rule for the Host is added on the selected Log Server.

On the Host element's **NAT** tab, you can view or edit the NAT definitions for the NAT configuration in which the element is included. However, you primarily configure NAT definitions in the properties of a Engine.

Defining IP Address List elements

IP Address List elements contain large lists of IP addresses, IP address ranges, or networks that can be used to filter traffic in Access rules, NAT rules, Inspection rules, and File Filtering rules.

IP Address List elements are imported and updated when you activate new dynamic update packages. You can also create your own IP Address List elements.

You can use IP Address List elements in Access rules to block the IP addresses used by specific services, such as Tor, or anonymous proxies. You can also block IP address ranges used by known botnets. If there are IP address ranges specific to your company that you want to control access to, create custom IP Address List elements.

Defining Network elements

A Network element represents the IP address space of a complete network or subnetwork.

On the Network element's **NAT** tab, you can view or edit the NAT definitions for the NAT configuration in which the element is included. However, you primarily configure NAT definitions in the properties of a Engine.

Defining Router elements

A Router element represents a next-hop gateway's IP address in routing configurations when the Router has a fixed IP address.

The element can also be used to represent IP addresses in rules and other configurations as needed.

Note

If the interface toward the next-hop gateway has a dynamic IP address, a special Gateway (DHCP Assigned) element must be added directly through the right-click menu of the automatically added Network (DHCP assigned) element in the Routing tree. The Gateway (DHCP Assigned) element is not valid in policies. Use a corresponding Alias element instead.

You can optionally configure on the **Monitoring** tab that a Log Server monitors the device.

If you selected any of the options on the **Monitoring** tab, a new Monitoring rule for the Router is added on the selected Log Server.

On the Router element's **NAT** tab, you can view or edit the NAT definitions for the NAT configuration in which the element is included. However, you primarily configure NAT definitions in the properties of a Engine.

Defining Zone elements

Zone elements allow you to group network interfaces of Engine, IPS, and Layer 2 Engines.

You can use Zones to specify the receiving or sending interfaces in policies. The Zone element represents all interfaces that belong to the Zone. All rules that include a Zone element also apply to any new interfaces that you associate with the same Zone.

There are several predefined System Zones available:

- **DMZ**: interfaces connected to DMZ networks.
- **External**: interfaces connected to the Internet or other external networks.
- **Guest**: interfaces connected to guest networks.
- **Internal**: interfaces connected to internal networks.
- **Node-internal**: Engine, IPS, and Layer 2 Engine nodes themselves. This Zone is automatically assigned to interfaces through which traffic to or from the engine node travels. It cannot be assigned to other interfaces, but it can be used in policies.

Related concepts

How Access rules match traffic on page 835

Related tasks
Define Source, Destination, and Service criteria in rules on page 891

Access and modify network elements

Network elements are different types of elements that represent IP addresses.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Network Elements**.

2) To create or edit elements:
   - In the right pane, click **New** and select an element from the drop-down list.
   - In a tree view, right-click the **Network Elements** category or a subcategory and select **New** <element>.
   - To create a copy of an existing element, right-click the element, then select **New** > **Duplicate**.
   - To edit an existing element, right-click the element, then select **Properties**. Default system elements cannot be edited.

Edit Expression elements

The criteria in an expression combine network elements (IP addresses) with logical operators.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Network Elements**.

2) Browse to **Expressions**.

3) Right-click an Expression element, then select **Properties**.

4) Click inside the expression to display the cursor.
   You can move the cursor using the arrow keys on your keyboard.

5) Click **Add Element** to add a new element, or click a logical operator.
   Expressions can start with a negation operation or a section in parentheses.

6) To delete an element or logical operator, use the Delete or Backspace key on your keyboard.

7) Click **OK** to save your changes.

Examples

- Expression that selects any network other than one internal network (172.16.1.0/24):



- Expression for selecting Hosts in the Network 172.16.1.0/24 from a Group element that contains Hosts from many different networks:



Using SMC elements to represent IP addresses in policies

There are special considerations for using SMC elements in policies.

Many elements are created as part of configuring a particular feature. When such elements define an IP address for a device, the element can also be used to represent the IP address in policies. However, there are some special issues that might have to be considered depending on the element type.

Tip

To view the actual IP addresses that the element adds to a policy, insert the element in a rule, then select ⋮ **More actions** > **Network Details**.

SMC components

Using elements that represent SMC components as a source or destination IP address in policies can produce unexpected results. Be careful especially when you use engine elements as source or destination IP addresses in policies:

- Engine, Single IPS, IPS Cluster, Single Layer 2 Engine, and Layer 2 Engine Cluster: These elements represent all static IP addresses defined for all interfaces. Create separate Host elements to represent individual IP addresses.
- Engine Cluster: Represent all CVI IP addresses of all interfaces, but not the NDI addresses. Create separate Host elements to represent individual CVI addresses and NDI addresses.
- Engines with dynamic IP addresses: The Engine element does not represent any of the dynamic IP addresses. There are default Aliases that can be used to represent the engine's own dynamic IP addresses in the engine's own policy. Fixed IP address definitions are needed for the dynamically assigned IP addresses when they have to be defined in the policies of any other components.
- SMC servers: Represent the single primary IP address defined for the element.

- Contact addresses are not taken into account when the element is used in a policy. Consider which IP address has to be added to the rule and create separate Host elements for the contact addresses as necessary.

External Servers

Several types of external servers can be integrated with the SMC when configuring different features. In general, each server element simply represents the single primary IP address defined in the element when used in a policy. Some elements have additional considerations when used in policies:

- Secondary IP addresses: Many server elements can contain one or more secondary IP addresses in addition to the primary address displayed for the element. The secondary addresses are equally valid in policies.
- Contact addresses: Some server elements can have a contact address. Contact addresses are not taken into account when the element is used in a policy. Consider which IP address has to be added to the rule and create separate Host elements for the contact addresses as necessary.
- Server Pools: The Server Pool element represents the external addresses that the clients contact. Use the Server Pool in rules that allow clients' traffic to the servers whenever you want to use the Server Pool features. Elements that represent the individual members of the pool can be used to allow connections to individual pool members (for example, to allow remote administration of each server).

Traffic handlers

Traffic handlers are used in Engine Policies when configuring Multi-Link for Engines. They can be used in rules in the following ways:

- In **Source** and **Destination** cells: A NetLink element represents the whole network address space that is associated with the NetLink element. An Outbound Multi-Link element represents the network address spaces of all NetLinks included in the Outbound Multi-Link element.
- In the **NAT** cell in NAT rules: When the source address is translated using the Outbound Multi-Link element as the address space, the traffic is balanced between the included NetLinks according to the options selected for the Outbound Multi-Link element.

Chapter 55

Working with Service elements

Contents

- Getting started with Service elements on page 927
- Creating Service elements on page 928
- Protocol elements and how they work on page 931
- Defining Protocol parameters on page 931

Service elements match traffic based on protocol or port and set options for advanced inspection of traffic. Service elements are used in Engine Policies, IPS Policies, Layer 2 Engine Policies, and Layer 2 Interface Policies.

Getting started with Service elements

Service elements specify a network protocol, as well as source or destination ports for TCP and UDP traffic.

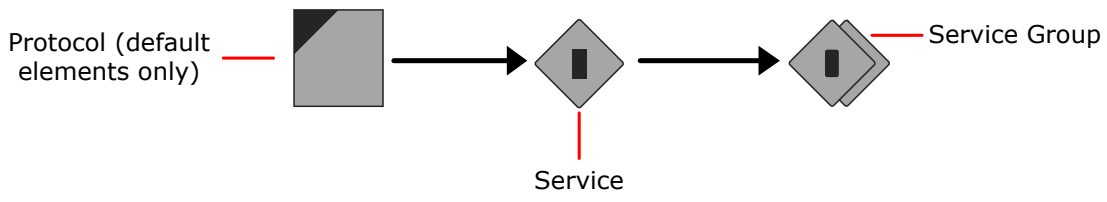
You can use Service elements to match rules to traffic in Ethernet rules (Ethernet Services), Access rules, and NAT rules.

Services can refer to **Protocol** elements, which activate further inspection checks and advanced traffic handling. Some Protocol elements have additional options that you can set in the Service element's properties.

Most of the time, you can use the default Service elements to represent standard protocols and ports. For example, you can enforce safe search features on the Security Engine by using predefined Services elements in Access rules. However, you might need to create a custom Service in the following cases:

- If none of the default Service elements match the type of traffic you want to allow, for example, if some TCP or UDP service in your network uses a non-standard port.
- If you want to set options for advanced traffic handling, for example:
  - Access rules that disallow the use of either the active or passive FTP connection mode
  - Engine Access rules for redirection to proxy services
  - Engine Access rules and Engine NAT rules for protocols that assign ports dynamically inside the packet payload

Elements in the Services configuration



The configuration of Service elements consists of the following general steps:

1) Create a Service element that matches the correct protocol number and port number (if applicable).

2) *(Optional)* Select one of the default Protocol elements if you want the traffic to be inspected further or if you want to use a Protocol Agent.

3) *(Optional)* Add the Service to a Service Group to make it easier to insert several related Services into configurations.

Related concepts
Enforcing safe search features in Access rules on page 842

Creating Service elements

There are predefined Service elements that correspond to reserved and commonly used protocols and ports. You might also need to add your own custom Service elements for any non-standard ports in use or if you want to define options for advanced traffic handling (Protocol Agents).

Create custom Service elements

Create a custom Service element if you need to match a protocol or port number that is not represented by the default Service elements. You can also use a custom Service element to change the properties of a Service element.

IP-based services are used in Access rules and NAT rules. Make sure that know which underlying protocol the traffic you want to allow uses, and be aware of whether you must define a protocol number or a port number. Usually, the Services you define yourself are TCP-based or UDP-based and are identified by the port number they use. However, there are many common protocols that are not TCP-based or UDP-based (for example, ICMP and RPC) and are identified by other information.

Example: The GRE protocol is transported directly over IP as protocol number 47 - on the same layer as TCP (#6) and UDP (#17). Therefore, any custom Services created for TCP and UDP ports 47 do not allow GRE to pass the Engine.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Expand the **Other Elements** tree and select **Services**.

3) Create the Service element in one of the following ways:
   - To create an element with no settings predefined, right-click the branch for the type of Service you want to create, then select **New** > **[Service type] Service**.
   - To create a Service based on some other Service element, right-click the existing Service, then select **New** > **Duplicate**.

4) Give the new Service a unique **Name** and write an optional **Comment**.

5) Configure the following options depending on the protocol:

Required options for each protocol

Protocol	Option
TCP and UDP	**Dst. Ports** (Optional)
	Src. Ports (Optional)
ICMP	**Type**
	Code (Optional)
SUN RPC	**Program Number**
	Version (Optional)
	Allow TCP (Optional)
	Allow UDP (Optional)
IP	**Code**

IANA assigns the protocol codes. See https://www.iana.org for a list.

6) (Optional) To associate the Service with a Protocol element, click **Select** next to the **Protocol** field and select a Protocol element.

Selecting the Protocol is mandatory if the Service is used in an Access rule that directs packets to deep inspection against the inspection rules. Some types of traffic might require a Protocol element of the type Protocol Agent.

7) (Optional, not available for all Protocol elements) Set more options on the **Protocol Parameters** tab.

8) Click **OK**.

Related concepts

How Access rules match traffic on page 835

Create Ethernet Service elements

Create a custom Ethernet Service element if you need to match Ethernet-level traffic that is not represented by the default Ethernet Service elements. You can also use a custom Ethernet Service element to change the properties of an Ethernet Service element.

There are predefined Ethernet Service elements that correspond to commonly used Ethernet services. You can use the predefined Ethernet Services if they meet your needs.

CAUTION

Match any IP traffic you allow in Ethernet rules to the default IPv4 and IPv6 Services. These Services match the traffic using the correct Protocol element. Only IP traffic matched to the correct Protocol element is inspected further against the Access rules. Non-IP traffic is never inspected any further.

Ethernet Services are used in IPS Policies, Layer 2 Engine Policies, and Layer 2 Interface Policies.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Expand the **Other Elements** tree.
3) Right-click **Ethernet Services** and select **New** > **Ethernet Service**.
4) In the **Name** field, enter a unique name for the new Ethernet Service.
5) Select the Ethernet-level protocol and enter the details depending on the protocol:

   Required options for each protocol

Protocol	Option
Ethernet 2 (DIX)	**MAC type**
LLC	**SSAP**
	DSAP
SNAP	**Vendor**
	Type

   Note

   IEEE assigns the protocol codes. The Raw IPX and Protocol protocols do not have any configurable options.

6) Click **OK**.

Create Service Group elements

Grouping several Service elements into Service Group elements simplifies your policies.

You can use the Service Group element in your policies instead of several individual Service elements. You can group both default Service elements and custom Service elements.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Expand the **Other Elements** tree.
3) Right-click the **Services** or (IPS only) **Ethernet Services** branch according to the type of service you want to create and select one of the following:

- **New** > **[Service type] Service Group** to create a Service Group under the branch for the selected type of services (for example, TCP).
- **New** > **Service Group** to create the Service Group under the Group branch (item not available if you right-clicked Ethernet Services).

4) In the **Name** field, enter a unique name for the Service Group.

5) From the list on the left, select the Service elements that you want to include in the group and click **Add**.

6) Click **OK**.

Protocol elements and how they work

Protocol Elements identify traffic as being of a certain network protocol.

Protocol elements can be inserted directly in Inspection rule exceptions. In Access rules, the Protocol elements are always contained in a Service element, which can then be inserted into the Service cell in rules. Some Protocols add options that you can adjust to custom Service elements that you create. You cannot add or modify the Protocol elements directly.

A Protocol element in Access rules identifies the protocol for inspection against Inspection rules. In Inspection rules, the Protocol can be used to limit the scope of exception rules according to the Protocol (as identified in the Access rules) in rules that otherwise match many Protocols. Also, the Protocols might activate some additional software modules on the engines.

This action depends on the type of the Protocol element:

- **Protocol Tag**: a Protocol element that does not activate additional modules.
- **Protocol Agent**: a Protocol element that activates an additional module on the engines to provide advanced application layer features.

Defining Protocol parameters

Many Protocols provide options for you to set in the Service that uses them.

Some of the parameters are only used by a specific type of component.

To set options for a Protocol, you must attach it to a custom Service element: either open the properties of a custom Service you have created previously or create a Service.

Link selection options in the Protocol Parameters specify how important different quality metrics are for traffic that is associated with the Protocol. Traffic uses the VPN link that best matches the link selection options.

The default link selection values are predefined and cannot be changed even if you duplicate the default elements. You can optionally use QoS Class elements in Access rules to override the default link selection values.

Define DNS Protocol parameters

The DNS protocol parameters control DNS protocol enforcement.

When you activate this feature, the engines can determine if traffic on the DNS port is actual DNS traffic or some other application that misuses this commonly open port. For example, peer-to-peer file transfer applications might use this port).

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the properties of a custom Service you have created, click **Select** next to the **Protocol** field and select **DNS**.

2) On the **Protocol Parameters** tab, set the parameters for the Protocol.

3) Click **OK**.

Define FTP Protocol parameters

The FTP Protocol Agent keeps track of the ports used in File Transfer Protocol (FTP) sessions. You can also use this Protocol Agent to redirect FTP traffic to a proxy service.

An FTP session starts with a control connection (by default, TCP port 21), and the communications continue using a dynamically allocated port. The FTP Protocol Agent can open the actual ports used in FTP sessions as needed so that the whole range of possible dynamic ports does not need to be allowed in the policy.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the properties of a custom Service you have created, click **Select** next to the **Protocol** field and select **FTP**.

2) On the **Protocol Parameters** tab, set the parameters for the Protocol.



Note

Do not change options marked as **Engine Only** from their default values when you use the Service element on IPS engines or Layer 2 Engines.

3) Click **OK**.

Define GRE Protocol parameters

The Generic Routing Encapsulation (GRE) protocol is a tunneling protocol that allows the encapsulation of network layer packets inside IP tunneling packets.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the properties of a custom Service you have created, click **Select** next to the **Protocol** field and select **GRE**.

2) On the **Protocol Parameters** tab, set the parameters for the Protocol.

3) Click **OK**.

Define H323 Protocol parameters

The H323 Protocol Agent tracks H.323 connections.

H.323 consists of a series of different types of standards relating to, for example, video and audio services, real-time transport, control channels and security. The H323 Protocol Agent can allow H.323 traffic through a engine when NAT is used.

Note

T.120 connections, used for instance for file transfer and whiteboard drawing, are not allowed by the H.323 Protocol Agent. To allow T.120, use the H.323 Service Group or the T.120 Service element.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the properties of a custom Service you have created, click **Select** next to the **Protocol** field and select **H323**.

2) On the **Protocol Parameters** tab, set the parameters for the Protocol.

3) Click **OK**.

Define HTTP or HTTPS Protocol parameters

You can use the HTTP and HTTPS Protocol Agents to redirect traffic to a proxy service and to log the URLs from HTTP requests.

You can also use the HTTPS agent to identify encrypted HTTPS traffic for decryption and inspection in the Access rules, and to identify encrypted HTTPS traffic for inspection in the Inspection Policy.

You can configure parameters for the HTTP and HTTPS Protocol elements.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the properties of a custom Service you have created, click **Select** next to the **Protocol** field, then select **HTTP** or **HTTPS**.
2) On the **Protocol Parameters** tab, set the parameters for the Protocol.
3) Click **OK**.

Related concepts

Getting started with forwarding traffic on page 1087
TLS inspection and how it works on page 1063

Define IPv4 Encapsulation Protocol parameters

The IPv4 Encapsulation Agent provides protocol inspection for tunneled IPv4 traffic.

The parameters define if IPv4 packets encapsulated in IPv6 packets are rematched to the policy.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the properties of a custom Service you have created, click **Select** next to the **Protocol** field and select **IPv4 Encapsulation**.
2) On the **Protocol Parameters** tab, set the parameters for the Protocol.
3) Click **OK**.

Define IPv6 Encapsulation Protocol parameters

The IPv6 Encapsulation Agent provides protocol inspection for tunneled IPv6 traffic.

The parameters define if IPv6 packets encapsulated in IPv4 packets are rematched to the policy.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the properties of a custom Service you have created, click **Select** next to the **Protocol** field and select **IPv6 Encapsulation**.
2) On the **Protocol Parameters** tab, set the parameters for the Protocol.
3) Click **OK**.

Define MSRPC Protocol parameters

The MSRPC (Microsoft RPC) Protocol Agent allows related connections for the endpoint mapper (EPM) protocol. It also handles NAT modifications for communications between Microsoft Outlook clients and Microsoft Exchange servers.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the properties of a custom Service you have created, click **Select** next to the **Protocol** field and select **MSRPC**.

2) (Engine only) On the **Protocol Parameters** tab, set the parameters for the Protocol.

3) Click **OK**.

Define NetBIOS Protocol options

The NetBIOS Protocol Agent can be used to make NAT modifications in IP addresses transported in the payload of Windows NetBIOS Datagram Service connections through the Engine. It can also be used for deep inspection on IPS engines and Layer 2 Engines.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the properties of a custom UDP Service you have created, click **Select** next to the **Protocol** field and select **NetBIOS (UDP)**.

2) On the **Protocol Parameters** tab, set the parameters for the Protocol Agent.

3) Click **OK**.

Define Oracle Protocol parameters

The Oracle Protocol Agent handles Oracle Transparent Network Substrate (TNS) protocol-based SQL*Net, Net7, and Net8 connections.

CAUTION

The Oracle Protocol Agent is for cases where TCP port 1521 is used only for negotiating the port number for Oracle database connections, and the port number for the actual connection is assigned dynamically. It must not be used in any other cases.

The Oracle Protocol Agent is meant only for non-SSL connections where the port number is assigned dynamically. If TCP port 1521 is used for the actual database connection, do not use a Service that contains this Protocol element. This can consume excessive resources on the engine and lead to performance problems. Instead, use a Service that matches TCP port 1521 without any Protocol element.

If you plan to use NAT for Oracle connections, you must configure the Oracle listener so that the listener tells the client its original non-NATed IP address. This configuration, and the Protocol Agent itself, is necessary only if the database is located on a different computer than the Oracle listener. The Oracle Protocol Agent does not modify payload data because the database service connections could go through a different route than the listener connection.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the properties of a custom Service you have created, click **Select** next to the **Protocol** field and select **Oracle**.

2) On the **Protocol Parameters** tab, set the parameters for the Protocol.

3) Click **OK**.

Define RTSP Protocol parameters

The RTSP Protocol Agent allows RTP (Real-time Transport Protocol) and RTCP (Real-time Control Protocol) media streaming connections initiated with RTSP through the engine.

The RTSP (Real Time Streaming Protocol) network control protocol is used for establishing and controlling media sessions between clients and media servers. On engines, the Protocol Agent also handles NAT modifications to the protocol payload.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the properties of a custom Service you have created, click **Select** next to the **Protocol** field and select **RTSP**.

2) On the **Protocol Parameters** tab, set the parameters for the Protocol.

3) Click **OK**.

Define Shell (RSH) Protocol parameters

The Shell (RSH) Protocol Agent manages Remote Shell connections and allows NAT modifications to the standard output (stdout) stream.

Remote Shell (RSH) is a widely used remote management protocol. The Shell (RSH) Protocol Agent also manages RExec connections.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the properties of a custom Service you have created, click **Select** next to the **Protocol** field and select **Shell**.

2) On the **Protocol Parameters** tab, set the parameters for the Protocol.

3) Click **OK**.

Define SIP Protocol parameters

The Session Initiation Protocol (SIP) agent can be used to handle multimedia connections that use SIP as their transfer protocol.

Using the agent allows SIP to be used across a engine that uses NAT. SIP uses TCP or UDP port 5060 to initiate the connection, after which the traffic is allocated a dynamically assigned port. The Protocol Agent monitors the actual ports used, so that the range of dynamic ports does not need to be allowed in the engine policy.

The SIP agent can be configured to force the client or server address used within the SIP transport layer to be used also for the media stream carried over SIP. This configuration is by default set for both client and server.



Note

Connections used for file transfer and whiteboard drawing are not allowed by the SIP Protocol Agent. Allow them in a different rule as necessary.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the properties of a custom Service you have created, click **Select** next to the **Protocol** field and select **SIP**.

2) On the **Protocol Parameters** tab, set the parameters for the Protocol.

3) Click **OK**.

Define SMTP Protocol parameters

On Engines, the Simple Mail Transfer Protocol (SMTP) Protocol Agent can be used to redirect connections to a proxy service. On Layer 2 Engines and IPS engines, the Protocol Agent can be used for protocol validation and deep inspection.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the properties of a custom Service you have created, click **Select** next to the **Protocol** field, then select **SMTP**.

2) (Engine only) On the **Protocol Parameters** tab, set the parameters for the Protocol Agent.

3) Click **OK**.

Related concepts
Getting started with forwarding traffic on page 1087

Define SSH Protocol parameters

The SSH Protocol Agent validates the communications to make sure the protocol used really is SSH.

Secure Shell (SSH) is an encrypted remote use protocol. You can create custom SSH agents with different settings, if necessary. The SSH Agent validates SSHv1 only. This Protocol Agent is available on Engines, IPS engines, and Layer 2 Engines.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the properties of a custom Service you have created, click **Select** next to the **Protocol** field and select **SSH**.

2) (Engine only) On the **Protocol Parameters** tab, set the parameters for the Protocol.

3) Click **OK**.

Define SunRPC Proxy parameters

The Sun Remote Procedure Call (RPC) Protocol Agent assists the Engine, Layer 2 Engine, or IPS engine in Portmapper connections.

There are both UDP and TCP-based Protocol Agents for Sun Remote Procedure Call (RPC) protocol. On the engine, these agents only assist the engine in Portmapper connections. They make the handling of RPC program numbers used in the Access rules more rapid. On IPS engines and Layer 2 Engines, these protocol agents provide deep inspection.

Note

The Protocol Agent is meant only for Portmapper connections. Allow other RPC services using Service elements without the Protocol Agent.

The Portmapper Protocol Agents collect information about RPC services by interpreting the GET PORT and DUMP PORTS requests and their respective answers. All information it collects is stored in the Portmapper cache.

When the packet filter needs to evaluate RPC matches, it consults the Portmapper cache to check if the destination of the packet has the appropriate service defined in the rule. If the cache does not have the requested information available, the packet under evaluation is not let through and a query is sent to the destination host for RPC information. The information received is stored in cache.

We recommend following these precautions with the RPC protocol:

- Attach the Portmapper Protocol Agent only to Portmapper connections passing through the engine.
- Allow the engine to send RPC queries.
- Optimize the structure of your security policy. See Knowledge Base article 10086 for more information.

RPC queries are sent from the engine to TCP port 111 of the external host. You can use the SunRPC (TCP) Service element or the SunRPC (UDP) Service element, or you can use the Portmapper Service element with both TCP and UDP. We recommend adding the following rule above any other Portmapper rules to allow connections without the Protocol Agent:

Rule for RPC Queries

Source	Destination	Service	Action
Engine IP address (NDIs on clusters)	Any	SunRPC (TCP) SunRPC (UDP)	Allow

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the properties of a custom Service you have created, click **Select** next to the **Protocol** field and select **SunRPC** ([TCP|UDP]).

2) (Engine only) On the **Protocol Parameters** tab, set the parameters for the Protocol Agent.

3) Click **OK**.

Define TFTP Protocol parameters

The Trivial File Transfer Protocol (TFTP) Protocol Agent transfers data using dynamically assigned ports.

A TFTP Agent is attached to a UDP connection established between the client and the server. The client opens the control connection from a dynamically selected source port to the fixed destination port 69/UDP on the server. A separate UDP data connection is established between the client and the server after the client has sent a “read” of “write” command to the server. The server opens a data connection from a dynamic source port to the client’s destination port. This port is same as the one used as the source port of the control connection.

The TFTP protocol (RFC 1350) does not limit the port range that can be used. This Protocol Agent is available on Engines, IPS engines, and Layer 2 Engines. The TFTP Protocol Agent supports NAT operations (Engine only).

The parameters in this protocol are for engines only.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the properties of a custom Service you have created, click **Select** next to the **Protocol** field and select **TFTP**.

2) (Engine only) On the **Protocol Parameters** tab, set the parameters for the Protocol.

3) Click **OK**.

Chapter 56

Defining Situation elements

Contents

- Getting started with Situation elements on page 941
- Situations configuration overview on page 942
- Create custom Situation elements on page 943
- Context options for Situation elements on page 944
- Defining Context Options for Correlation Situation elements on page 946
- Default elements for Situation elements on page 952
- Using Tags with Situation elements on page 953
- Vulnerability elements and how they work on page 955
- Using Situation elements on page 957
- Examples of custom Situation elements on page 957

Situation elements contain the context information that defines the pattern that the Security Engine looks for in the inspected traffic. Situation elements also define the patterns that match events in the traffic.

Getting started with Situation elements

Situation elements define a pattern in traffic that the engine looks for.

The patterns and events are defined by selecting a *Context* for the Situation element. The Context contains the information on the traffic to be matched, and the options you can set for the matching process.

The Inspection Policy defines how the Situation elements are matched to traffic and what action the engine takes when a match is found.

Correlation Situation elements are Situation elements that group event data to find patterns in that data.

Situation elements also provide a description that is shown in the logs, and a link to relevant external information (CVE/BID/MS/TA) in the form of a *Vulnerability* element attached to the Situation.

You can group Situations together using *Tags*. The Tag elements are shown as branches in the Situations tree and they can be used in policies to represent all Situations that are associated with that Tag. For example, using the Tag **Windows** in a rule means that the rule matches all Situations that are classified as concerning Windows systems.

Associating a Situation with a *Situation Type* includes the Situation in the Rules tree in the Inspection Policy, which is grouped according to the Situation Type.

Depending on the Usage Context properties of a Correlation Situation, correlation can be done only on the Security Engine, only on the Log Server, or on both the Security Engine and the Log Server. When correlation is done only on the Security Engine, the Correlation Situation only matches when all correlated events are detected by the same Security Engine. The following table lists the Usage Contexts for predefined Correlation Situations:

Usage Contexts for predefined Correlation Situations

Correlation Context	Usage Context
Compress	Engine Only
Count	Engine Only
Group	Engine Only
Match	Engine Only
Sequence	Log Server Only

By default, correlation is done on both the Security Engine and the Log Server for custom Correlation Situations.

Situations configuration overview

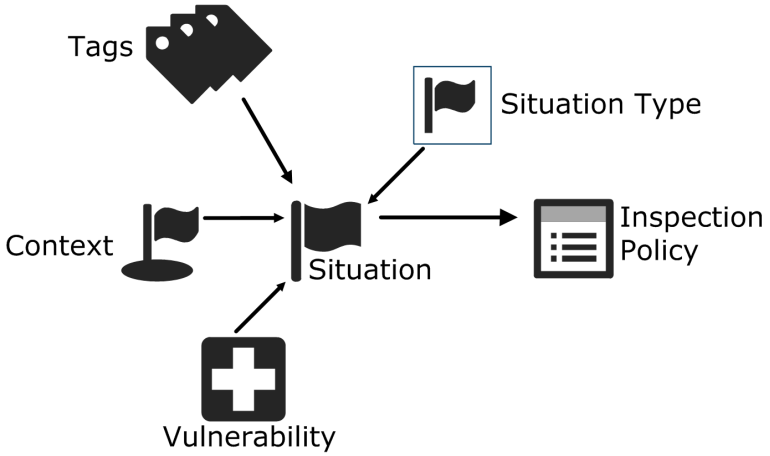
Configuring Situation elements involves several main steps.

The Situation element uses different elements to form a representation of the traffic that you want to detect in your Inspection Policy. The purpose of these elements is as follows:

- The **Tag** elements help you to create simpler policies with less effort. Tag elements represent all Situations that are associated with that Tag. For example, using the Tag “Windows” in a rule means that the rule matches all Situations that concern Windows systems.
- The **Situation Type** elements define the general category of the Situation and the branch of the Rules tree under which the Situation appears (such as Threats or Suspicious Traffic.). One Situation Type can be associated with each Situation.
- The **Context** element defines the traffic patterns the Situation detects. The Context binds the Situation to a certain type of traffic and gives you a set of options or a field for entering a regular expression.
- The **Vulnerability** element associates your custom Situation with a commonly known vulnerability. It allows you to attach a description of the Vulnerability and references to public vulnerability databases (which are shown in the Logs view if a match is found).

The Context is the only mandatory element in a Situation. However, it is recommended to consistently associate all relevant Tags with each custom Situation you create. The vulnerability description is not mandatory, but it is helpful to have it for Situations that detect some publicly known issue.

Elements in the configuration



1) Create a Situation element.

2) Give the Situation a Context, and fill in the context information according to the patterns in traffic that you want to match.

3) (Optional) Associate the Situation with the relevant Tags.

4) (Optional) Associate the custom situation description with a relevant Vulnerability.

5) Use the Situation in the Inspection Policy.

Related concepts
Inspection Policy elements and how they work on page 867

Create custom Situation elements

You can create custom Situation elements in addition to using the predefined ones.

Before you begin

Creating new Situation elements requires detailed knowledge of the protocols that you want to inspect and the traffic patterns related to their use.

You can create a Situation element to detect individual events or a Correlation Situation element to detect a group of related events.

A Situation element collects together the related elements and settings and sets the severity value for the Situation. The severity value can be set between Info (the least severe) to Critical (the most severe). You can use the severity value to restrict which Situations added to the Situations cell are considered in Inspection Exceptions and Alert Policies. For example, if a rule matches a large range of Situations you can create separate rules for less severe and more severe Situations.



Note

Avoid defining the same pattern in different Situation elements. Duplicate situations in the policy can create unintended results and makes the policies difficult to manage.

The predefined Situation elements are updated through dynamic update packages. You can also create new Situation elements to fine-tune the patterns that the engines look for in the traffic.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Browse to **Other Elements** > **Situations**.

3) Right-click **Situations**, then select **New** > **Situation** or **New** > **Correlation Situation**.

4) In the **Name** field, enter a unique name.

5) (Optional) Click **Select** and select the Situation Type with which to associate this Situation.
You can only select one Situation Type for each Situation. The Situation Type specifies the branch of the Rules tree under which the Situation is included.

6) In the **Description** field, enter a description of the traffic pattern that the Situation represents.
This description is shown in log entries and other places where statistics related to the Situation appear.

7) In the **Severity** drop-down list, select the severity value for the Situation.
The Severity is shown in the logs and can be used in Alert Policies as a criterion for alert escalation.

8) (Optional) In the **Attacker** and **Target** drop-down lists, select how the attacker and target are determined when the Situation matches.
This information is used for block listing and in log entries.

Related tasks
Add URL List Application elements to manually block or allow URLs on page 1025

Context options for Situation elements

Context elements define what the Situation element matches. Adding a context to a Situation allows you to define what kinds of patterns you want to look for in the traffic.

For example, you can specify that you want to look for a certain character sequence in an HTTP stream from the client to the server.

When you select a context, a set of options or a field for entering a regular expression as parameters for the context is added to the Situation element. The parameters define the pattern you want to look for in the traffic.

The following types of contexts are available:

Context	Description
Anti-Malware	Anti-Malware contexts are used to detect malware.
DoS Detection	DoS detection contexts provide parameters for detecting DoS (Denial of Service) events in network traffic.
File	File contexts are used to detect malicious or suspicious content in transferred files regardless of the transport protocol used. When a file is detected, the file is inspected to identify the file type. When the file type is identified, more specific inspection can be applied to the file.

Context	Description
Protocol-specific contexts	Protocol-specific contexts are used to detect a particular characteristic in the network traffic. For example, you can detect a certain option number used in IP packets, or set the maximum length for particular arguments in FTP commands. For contexts that have particular values to be filled in (instead of a regular expression), the parameters you define in the contexts often actually determine what is regarded as normal. Anything above/below/outside/not matching these values is regarded as a match for the Situation. In some cases, you might define what the Situation *does not* match. Using protocol-specific contexts requires basic knowledge of the underlying network protocols and how the traffic in your network uses those protocols. For more information about what a particular context is used for, see the Properties dialog box of the context.
Scan Detection	Scan detection contexts provide parameters for detecting attempts to scan which IP addresses are in use or which ports are open in your systems.
System	System contexts are used for errors and other system events. System Contexts are internal to the SMC, and they cannot be edited in any way.

Related tasks

Add URL List Application elements to manually block or allow URLs on page 1025

Define Context options for Situation elements

The Context gives the Situation the information on which patterns you want it to match in the traffic.

For example, you might want to look for a certain character sequence in an HTTP stream from the client to the server.

The Content gives you a set of options or a field for entering a regular expression that you can use to define the pattern you want to look for in the traffic.

Note

Avoid defining the same pattern in different Situation elements. Duplicate situations in the policy can create unintended results and makes the policies difficult to manage.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) On the **Context** tab of the **Situation Properties** dialog box, click **Select**.
   The available Context categories are shown.

2) Select the Context you want to associate with this Situation.

Note

- The details related to the Contexts can be different from what is described here because the Contexts might have been updated through dynamic update packages. Read the Release Notes of each update package you import to see which elements are affected.
- For many Contexts, type in a regular expression.
- In other cases, open the Properties dialog box for the Context element for more information.

The options for the selected Context are added to the Situation Properties.

Defining Context Options for Correlation Situation elements

Correlation Contexts define the patterns for matching groups of related events in traffic.

Correlation Situations are used by Security Engines and Log Servers to conduct further analysis of detected events. Correlation Situations do not handle traffic directly. Instead they analyze the events generated by matches to Situations found in traffic. Correlation Situations use Event Binding elements to define the log events that bind together different types of events in traffic.

Correlation Context types

Correlation Context Type	Description
Compress	Combines repeated similar events into the same log entry, reducing clutter in the Logs view. Example: There is a custom Situation for detecting suspicious access to a file server. An attacker is likely to browse through many files, triggering an alert entry for each file. An Event Compress Situation can be used to combine Situations together when the suspect's IP address is the same.
Count	Finds recurring patterns in traffic by counting how many times certain Situations occur within the defined period, so that action can be taken if the threshold values you set are exceeded. Example: A Situation that detects access to a system could normally trigger just a log entry. The Event Count Situation could be used to block list connections when access by any single host is too frequent.
Group	Finds event patterns in traffic by keeping track of whether all events in the defined set of Situations match at least once in any order within the defined time period. Example: Individual attempts to exploit different vulnerabilities in a software product in use on your server might not be too alarming if you know that your system is patched against those vulnerabilities. However, when several such events are found in a short period, it becomes more likely that someone is trying to systematically attack the server. They might also already knows that the server is running that particular piece of software. A Situation that belongs to the Group Context can detect this kind of attack.
Match	Allows you to use Filters to filter event data produced by specific Situations.

Correlation Context Type	Description
Sequence	Finds event patterns in traffic by keeping track of whether all events in the defined set of Situations match in a specific order within the defined time period. Example: Clients might use a certain type of request (for example, “give file X”) to fetch a file from a file server. When administrators log on to the same server, a successful administrator logon can be seen in the traffic as a certain type of response (for example, “full access granted”). However, a vulnerability in the server software can allow an attacker to send a specially crafted file fetch request. This kind of request might look like a valid “give file x” command, but actually causes the server to give the attacker administrator rights. This action is seen as a normal-looking “full access granted” response from the server. The Event Sequence Situation can detect when a “give file X” Situation match is followed by a “full access granted” Situation match, which cannot be any legitimate traffic.

Depending on the Usage Context properties of the Correlation Situation, correlation can be done only on the Security Engine, only on the Log Server, or on both the Security Engine and the Log Server. By default, correlation is done on both the Security Engine and the Log Server for custom Correlation Situations. For more information about selecting the Usage Context, see the following Context-specific sections.

CAUTION

In custom Correlation Situations, logging might be automatically enabled for the correlated Situations even if the correlated Situations do not normally have logging enabled. If the Situations produce a large amount of log data and correlation is done on the Log Server, the increased amount of log data might overload the network or the Log Server even if no correlation matches occur.

Define Context options for Correlation Situation elements

Correlation Contexts define the patterns for matching groups of related events in traffic.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) On the **Context** tab of the **Correlation Situation Properties** dialog box, click **Select**.

2) Select the Context you want to associate with this Correlation Situation.

Related tasks

Add Tag elements to Situation elements on page 953

Define Compress Context parameters for Correlation Situation elements

The Compress Context combines repeated similar events into the same log entry, reducing clutter in the Logs view.

CAUTION

Be careful when defining the Compress Context options. You must make sure that all event data you compress is part of the same event. Otherwise you risk losing valuable event information.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Browse to the Situations you want to compress in the left pane of the dialog box and drag and drop them into the **Correlated Situations** field.

   CAUTION

   In custom Correlation Situations, logging might be automatically enabled for the correlated Situations even if the correlated Situations do not normally have logging enabled. If the Situations produce a large amount of log data and correlation is done on the Log Server, the increased amount of log data might overload the network or the Log Server even if no correlation matches occur.

2) Enter the **Time Window Size** in seconds. The matches to the Situations selected are combined to a common log entry when they are triggered within the defined time period.

3) Enter the **Events per Window**. This defines the maximum number of events that are forwarded within the Time Window defined.

4) Select a **Log Fields Enabled** option.

5) Double-click the **Event Binding** field and select the Event Binding that is used by the matching option you selected in the previous step.

6) Select a **Location** to determine the execution order of this Compress operation in relation to other Compress operations. Operations that share the Location are executed in parallel; each compress operation receives the same events as the other compress operations in the same Location. The next Location receives only the events that are left after the compression.

   CAUTION

   Be careful when using the Early or Very Early Locations. The compression can affect the other types of correlation tasks.

7) Click **Edit** and define a **Compress filter** in the Local Filter Properties dialog box. The filter is used for filtering data to be included in the compression.

8) Make sure that **Engine Only** is selected as the **Usage Context**.



Note

The purpose of the Compress context is to reduce the amount of data that is sent to the Log Server. Including the Log Server in the Usage Context of a Compress context actually increases the amount of data that is sent to the Log Server.

Related concepts

Creating and editing local filters on page 337

Define Count Context parameters

The Count Context finds recurring patterns in traffic by counting how many times certain Situations occur within the defined period. Action can then be taken if the threshold values you set are exceeded.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Browse to the Situations you want to count in the left pane and drag and drop them to the **Correlated Situations** field.



Note

In custom Correlation Situations, logging might be automatically enabled for the correlated Situations even if the correlated Situations do not normally have logging enabled. If the Situations produce a large amount of log data and correlation is done on the Log Server, the increased amount of log data might overload the network or the Log Server even if no correlation matches occur.

2) Enter the **Time Window Size** in seconds. All events must occur during this length of time for the Correlation Situation to match.

3) Enter the **Alarm Threshold** number. This is the number of times that the event must occur for the Correlation Situation to match.

4) Select a **Log Fields Enabled** option.

5) Double-click the **Event Binding** field and select the Event Binding that is used by the matching option you selected in the previous step.

6) (Optional) Select the Usage Context to define where correlation is done.



Note

If you select a Usage Context that does not include the Log Server, events only match if they are all detected by the same Security Engine or Security Engine Cluster.

Define Group Context parameters for Correlation Situation elements

The Group context finds event patterns in traffic by keeping track of whether all events in the defined set of Situations match at least once in any order within the defined time period.

The Group context has a table that allows you to define local filters and log fields for selecting which details are considered when events are grouped. In this context, the order in which the events occur is not relevant. If you would like the order of the events to matter, use the Sequence context instead.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Double-click the **Event Match** cells for each **Member** cell and define a local filter. The local filter selects the events for examination.

   You can add and remove members using the buttons to the right (to remove a member, first select a cell within that member's column).

2) Double-click the **Needed Number** cell of each member and enter the number of occurrences of the Event Match that are required for the events to be grouped.

3) Double-click the **Event Binding** field and select the Event Binding that defines the set of log events to match.

4) Drag and drop the relevant Situations to the **Correlated Situations** field.

5) Select whether you want to **Keep and Forward Events**.

6) Enter the **Time Window Size** in seconds. All events must occur during this length of time for the Correlation Situation to match.

7) Select whether you want to trigger **Continuous Responses**.

8) (Optional) Select the Usage Context to define where correlation is done.



Note

If you select a Usage Context that does not include the Log Server, events only match if they are all detected by the same Security Engine or Security Engine Cluster.

Related tasks

Create or edit Filter elements on page 340

Define Match Context parameters for Correlation Situation elements

The Match context allows you to use Filters to filter event data produced by specific Situations.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Browse to the Situations you want to count in the left pane of the dialog box and drag and drop them into the **Correlated Situations** field.

CAUTION

In custom Correlation Situations, logging might be automatically enabled for the correlated Situations even if the correlated Situations do not normally have logging enabled. If the Situations produce a large amount of log data and correlation is done on the Log Server, the increased amount of log data might overload the network or the Log Server even if no correlation matches occur.

2) Click **Edit** and define a local filter.

3) (Optional) Select the Usage Context to define where correlation is done.

Note

If you select a Usage Context that does not include the Log Server, events only match if they are all detected by the same Security Engine or Security Engine Cluster.

Related concepts

Creating and editing local filters on page 337

Define Sequence Context parameters for Correlation Situation elements

The Sequence context finds event patterns in traffic by keeping track of whether all events in the defined set of Situations match in a specific order within the defined time period.

The Sequence context has a table you can use to define, in order from left to right, the events that comprise a sequence. This context allows detecting events such as when a request from a client triggers one pattern and the response from a server triggers a second pattern.

The table has gray and white cells; white cells must be filled, gray cells are left empty. If you would like to match events regardless of the order in which they occur, use the Group context instead.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Double-click the **Event Match** cell and define a local filter.

2) Double-click the **Event Binding** field and select the Event Binding that defines the set of log events to match.

3) Click Add Event Before or Add Event After to add more Event rows.
   - Define the Event Match for each row.
   - Define the Event Binding for each row.

4) When you are finished defining the sequence, drag and drop the relevant Situations in the **Correlated Situations** field below the table.

5) Select whether you want to **Keep and Forward Events**.

6) Enter the **Time Window Size** in seconds. All events must occur during this length of time for the Correlation Situation to match.

7) Select one of the following options as the **Usage Context**:
   - Engine and Log Server.
   - Log Server Only.



Note

If you select a Usage Context that does not include the Log Server, events only match if they are all detected by the same Security Engine or Security Engine Cluster.

Related concepts

Creating and editing local filters on page 337

Default elements for Situation elements

There are many predefined Contexts, Situations, Tags, and Vulnerabilities.

Predefined Contexts, Situations, Tags, and Vulnerabilities are imported and updated from dynamic update packages. This also means that the set of elements available changes whenever you update your system with new definitions. Both Situation elements and Context elements have a comment and a longer description that you can view in the SMC Client. This information is shown in the Info pane or in the Properties dialog for the element and explains what each element is meant for.

The Release Notes of each dynamic update package list the new elements that the update introduces. If your Management Server can connect to the Forcepoint website, you can view the Release Notes directly through the SMC Client.

Using Tags with Situation elements

Tag elements collect together several Situations that have something in common.

For example, you can use Tag elements to group Situations that detect threats against a particular Operating System. Tags are shown as branches in the Situations tree. Tag elements help you organize the tree, and you can use the Tags in the Inspection Policy to easily match the rule to all Situations that reference the Tag.

Create new Tag elements for Situation and Vulnerability elements

Tag elements collect together several Situations and Vulnerabilities that have something in common.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Browse to **Other Elements** > **Situations** or **Other Elements** > **Vulnerabilities**.
3) Right-click the **By Tag** branch, then select **New** > **<Tag Type>**.
4) Enter a **Name** and optionally a **Comment** for the new Tag, then click **OK**.
   The new Tag is added to the tree under the main category you selected in the previous step.

Add Tag elements to Situation elements

You can use Tag elements to group Situation elements and Situation Type elements to classify Situations.

You can use predefined Tags or create new ones according to any criteria (for example, create a Tag for grouping together related services). Situation Types are predefined, and you cannot create new Situation Types. You can associate multiple Tags with one Situation, but only one Situation Type can be associated with each Situation.

You can use the Tags and/or Situation Types to represent a group of Situations in the Rules and Exceptions of the Inspection Policy. This allows you to match a rule to all Situations that contain the Tag or Situation Type. Situations that are associated with a Situation Type are automatically included in the Rules tree.

When you are editing the Situation properties, you can add tags.

Note

If a Tag or Situation Type you add to a Situation is in use in some Inspection Policy, the new Situation is automatically included in the policy when you save the Situation. The engines start matching traffic to the Situation when you refresh the policy.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the Situation properties, switch to the **Tags** tab.
2) Click **Add Tags** and select a Tag type from the list that opens.
3) Select the Tags you want to use with this Situation and click **Select**.
4) Click **OK** to confirm the Situation properties change.

Add Tag elements to several Situation elements from the element tree

You can add the same Tag element to several Situation elements.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Expand the **Other Elements** > **Situations** branch of the element tree and browse to the Situations you want to tag.
3) Ctrl-select or Shift-select the Situations that you want to tag.
4) Right-click one of the selected Situations and select **Add Tag**.
   The tag types are shown in a submenu.
5) Select the type of Tag you want to add.
6) Select the Tags you want to attach to the Situations.

Remove Tag elements from Situation elements

You can remove custom Tag elements that administrators have added.

The default Tags in System Situations (provided in Update Packages) cannot be removed.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Expand the **Other Elements** > **Situations** branch of the element tree and browse to the Situations you want to edit.

3) Right-click the Situation and select **Properties**.

4) Switch to the **Tags** tab.

5) Right-click the Tag you want to remove and select **Remove**.

6) Click **OK**.

Vulnerability elements and how they work

Vulnerability elements provide a short description of the event that has matched and a reference to external vulnerability information (CVE/BID/MS/TA).

When a Situation element refers to a Vulnerability element, the vulnerability information is included in the log entries generated when the Situation matches traffic. You can also use the vulnerability IDs in the element search to find Situation and Vulnerability elements that refer to a particular vulnerability.

Vulnerability information is included in dynamic update packages, so Situations provided by Forcepoint are already linked to a Vulnerability when appropriate. You can associate Situations with an existing Vulnerability or add a custom Vulnerability element.

Create Vulnerability elements

You can create custom Vulnerability elements to associate Situations with vulnerabilities that are not included in the default Vulnerability elements.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Expand the **Other Elements** > **Situations** branch of the element tree.

3) Right-click the **By Vulnerability** branch in the tree view and select **New** > **Vulnerability**.

4) Give the Vulnerability a descriptive name and optionally a comment.

   The Comment is not shown in the Logs view. Use the **Description** field to enter information to be shown in the logs.

5) To create a reference to external vulnerability information, select one or more reference systems in the **Reference System** section and enter the **ID** this vulnerability has in that system:

   - **Mitre**: vulnerability ID format is CVE-YYYY-XXXX

- **SecurityFocus**: vulnerability ID format is BID-XXXXX
- **Microsoft**: vulnerability ID format is MSYY-XXX
- **Us-Cert**: vulnerability ID format is TAYY-XXXX

6) After you have entered the vulnerability ID, click **Show** next to the ID field to view the information about the vulnerability in the reference system.

7) Type or copy-paste a short description of what the vulnerability is about into the **Description** field.

8) Under **Situations**, browse to the correct Situation elements, select them (one or several at a time) and click **Add** to associate them with this Vulnerability.
   The selected Situations are added to the **Content** field on the right.

9) When you are finished adding Situations, click **OK**.
   The selected Situations are now associated with this vulnerability, and a link to this Vulnerability is added on the Situations' properties dialog box.

Associate Vulnerability elements with Situation elements

You can associate Situation elements with an existing Vulnerability element or add a custom Vulnerability element.

You can add up to four references to public vulnerability databases to your custom Vulnerabilities (CVE/BID/MS/TA). System vulnerabilities can have an unlimited number of references to any reference system, and can have multiple references to the same reference system. The reference information is also shown in the Logs view.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Expand the **Other Elements** > **Situations** > **By Vulnerability**.

3) Right-click the correct Vulnerability in the tree and select **Properties**.

4) Select the Situation elements (one or several at a time) and click **Add** or **Remove** to change the selection in the **Content** field on the right.

5) When you are finished adding Situations, click **OK**.
   The selected Situations are now associated with this vulnerability, and a link to this Vulnerability is added on the Situations' properties dialog box.

Related tasks

Create custom Situation elements on page 943

Using Situation elements

You use Situation elements to define what you want to detect with the Inspection Policy.

Situations are generally used for:

- Detecting malicious patterns in traffic. The Situations supplied in dynamic update packages concentrate on such known vulnerabilities and exploits.
- Reducing the number of alert and log entries you receive (using Correlation Situations).
- Detecting some other traffic patterns that you want to record. For example, you might be interested in the use of certain applications.

Although the general workflow requires making sure that a Situation you want to use is included in the Inspection Policy, you might often not actually insert the Situation into the rule. Instead, you might use a Tag or Situation Type element to represent a whole group of Situations.

Examples of custom Situation elements

These examples illustrate some common uses for Situations and the general steps on how each example is configured.

Example: detecting the use of forbidden software

Engine-specific examples of using custom Situations to detect traffic patterns associated with specific software.

Engine

Company A has a Engine that inspects all outgoing web traffic against the Inspection Policy. The use of instant messaging clients across the Internet is forbidden in the company. The Inspection Policy is set to detect and log Situations with the **Instant Messaging** Tag.

The company's administrators have found out that some internal users have started chatting using a new little-known instant messaging client that does not have a default Situation yet. The communications seem to be standard HTTP directly from client to client. The administrators find one distinctive characteristic in the software: when started, the software in question always connects to a particular address to check for updates using HTTP.

The administrators:

1) Create a custom Situation element with the name "Software X".
2) Add the **HTTP Request URI** Context to the Situation and type in a regular expression that contains the address they want the Situation to find using the SMC regular expression syntax.
3) Add the default system Tag **Instant Messaging** to the Situation.
4) Refresh the Engine's policy.
5) Open the Logs view and filter the view using the "Software X" Situation as the filtering criteria.

6) See which computers use the forbidden software and remove the software from the computers shown in the logs.

IPS engine

Company A has an IPS engine deployed in between their internal network and the Internet. The IPS engine uses a policy that is based on the IPS Template policy.

The administrators find out that some of the internal users have installed a piece of software on their computers that the company's security policy forbids. They consider this software a security risk.

The administrators decide that they would like to detect the use of the software so that they can find out which users have installed it. The administrators find one simple but distinctive characteristic in the software: when started, the software in question always connects to a particular address to check for updates using HTTP.

The administrators:

1) Create a custom Situation element with the name "Software X".

2) Add the **HTTP Client Stream** Context to the Situation and type in a regular expression that contains the address they want the Situation to find using the SMC regular expression syntax.

3) Add one of the default Situation Types under **Traffic Identification** to the Situation.

4) Select the correct options for logging the traffic in the Rules tree in the Inspection Policy and install the policy on the IPS engine.

5) Open the Logs view and filter the view using the "Software X" Situation as the filtering criteria.

6) See which computers use the forbidden software and take action based on which IP addresses are shown in the logs.

Example: counting events to reduce number of repeated queries to a server

An example of using the Count context in a Correlation Situation.

Company B has a Engine and an IPS engine that monitor traffic to a DMZ network. The DMZ contains a server that provides information to Company B's partners. A while ago, users started complaining that the service had slowed down.

Upon investigation, Company B's administrators found out that the traffic had grown dramatically even though the number of users and the data available had stayed the same. They found out that one of the partners had made a misconfiguration script that frequently copied several large catalogs from Company B's server to their own server. Furthermore, they had given the script to a few other partners as well. As a first step, the administrators decide to immediately stop excessive queries to the server.

The administrators:

1) Create a custom Situation for detecting access to the catalog files.

2) Create a custom Correlation Situation and attach the **Count** Context to it. Then define the settings for the Count Context to detect when there are more than 5 requests per minute to any of the files from the same source address.

Context settings for the example Correlation Situation

Field	Option
Correlated Situations	Custom Situation
Time Window	60
Alarm Threshold	5
Log Fields Enabled	Select
Event Binding	Src Addr

3) Insert the Correlation Situation in the Inspection Policy with block listing as the Action. The traffic from the offending hosts is stopped at the Engine.

4) Refresh the Inspection Policy on the IPS engine.

Example: preventing access to forbidden websites

An example of using Situations to block access to specific websites.

The Administrators at Company C have noticed that employees frequently visit websites that are not related to their work. They want to block access to these websites to prevent employees from accessing them at work. To prevent access, they:

1) Create a Situation element.

2) Add the **Website Access Control** Context to the Situation.

3) Specify the addresses they want to prevent access to. Access to the specified addresses is blocked.

4) Refresh the Inspection Policy on the IPS engine.

Chapter 57

Using Network Application elements

Contents

- Getting started with Network Application elements on page 961
- Create TLS Match elements for network application detection on page 963
- Access rules for network application detection on page 964
- Example: blocking network application use on page 967

Network Application elements collect combinations of identified characteristics and detected events in traffic to dynamically identify traffic related to the use of a particular network application.

Getting started with Network Application elements

Network Application elements provide a way to dynamically identify traffic patterns related to the use of a particular network application.

Network Application elements allow you to more flexibly identify traffic beyond specifying a network protocol and ports for TCP and UDP traffic with a Service element. Matching is done based on the payload in the packets, making it possible to identify the protocol even when non-standard ports are used. First, the protocol is identified, then a protocol-specific pattern matching context is applied to identify the network applications.

Link selection options in the properties of Network Application elements specify how important different quality metrics are for traffic that is associated with the network application. Traffic uses the VPN link that best matches the link selection options.

The default link selection values are predefined and cannot be changed even if you duplicate the default elements. You can optionally use QoS Class elements in Access rules to override the default link selection values.

Keep the following in mind when working with Network Application elements:

- There are several predefined Network Application elements available that define the criteria for matching commonly used network applications. No configuration is required to be able to use Network Application elements in Access rules.
- Predefined TLS Match elements are used in the properties of some predefined Network Application elements to allow the Network Application to match the use of the TLS protocol in traffic.
- You cannot edit the predefined Network Application elements. However, Access rules can override the properties of a predefined Network Application element.
- Creating Network Application elements requires detailed knowledge of the network applications you want to detect and the traffic patterns related to their use. Creating Network Application elements is not recommended.

- If a certificate for TLS inspection has been uploaded to the engine, adding a Network Application that allows or requires the use of TLS to an Access rule enables the decryption of all TLS traffic. So does enabling the logging of Application information in the Access rules.

Network Applications configuration overview

No configuration is required to be able to use Network Application elements in Access rules.

There are several predefined Network Application elements available that define the criteria for matching commonly used network applications. Creating Network Application elements or duplicating existing elements is not recommended. To override the settings of a predefined Network Application, edit the Service Definition of the rule in which you use the Network Application.

1) (Optional) Create TLS Match elements to override the properties of predefined Network Application elements.

2) Use the Network Application element in the Access rules.

Default elements for network applications

There are several predefined elements for working with network applications.

Application Type elements define general categories of network applications. One Application Type element can be associated with each Network Application element. Application Type elements are predefined, and you cannot create Application Type elements.

Tag elements help you to create simpler policies with less effort. Tag elements represent all Network Application elements that are associated with that Tag. For example, the Media Tag includes several web-based image, music, and video applications. Several Tags can be associated with each Network Application element.

Dependencies for network applications define other network applications that must also be allowed when the network application is allowed. When you use a network application that has dependencies in a rule with the Allow or Jump action, or in a NAT rule, the rule also applies to the related network applications. When you use the network application that has dependencies in a rule with the Continue, Discard, or Refuse action, the rule does not apply to the related network applications.

TLS Match elements define matching criteria for the use of the TLS protocol in traffic. When a connection that uses the TLS protocol is detected, the server certificate for the connection is compared to the TLS Match in the Network Application definition. TLS connections are allowed only to sites that have trusted certificates that meet the following criteria:

- The certificate domain name must match the domain name in the TLS Match element.
- The certificate must be signed by a valid certificate authority.
- The certificate must be valid (not expired or revoked).

TLS Match elements can also specify whether to decrypt TLS traffic to particular Internet domains for inspection. The default TLS Match elements deny decryption of only the following types of traffic:

- Traffic for Network Applications that do not work correctly if the traffic is decrypted.
- Traffic that is functionally critical, such as connections to the Forcepoint Advanced Malware Detection service, or to services for automatic dynamic updates and engine upgrades.

For more information, see Knowledge Base article 18074.

The predefined elements are imported and updated from dynamic update packages. The set of elements available changes whenever you update your system with new definitions. The Release Notes of each dynamic update package list the new elements that the update introduces.

Create TLS Match elements for network application detection

TLS Match elements define matching criteria for the use of the TLS (transport layer security) protocol in traffic.

In addition to the predefined TLS Match elements used in predefined Network Application elements, you can optionally define your own TLS Match elements.

TLS Match elements can match traffic based on the following criteria:

- Whether certificate validation succeeded, failed, or was not performed.
- The server domain name in a valid certificate.
- Specific reasons a certificate is regarded as invalid if certificate validation failed.
- The domain name in the Server Name Indication (SNI) field of the TLS Client Hello packet.

TLS Match elements also specify whether to decrypt TLS traffic to particular Internet domains for inspection. TLS Match elements that deny decryption are applied globally. Even if the TLS Match element is not used in the properties of any Network Application elements or in the Access rules, matching connections are never decrypted. Denying decryption in a TLS Match prevents network applications from being detected in encrypted connections to the specified domains. If the server certificate provides sufficient information to identify the network application without decrypting the client communications, you can alternatively specify that decryption is not necessary for network application identification in the Network Application properties.

A Network Application element matches a TLS connection only if a TLS Match element in the Network Application also matches. However, TLS Match elements used in Service Definitions override the TLS Match of a Network Application. In this case, the rule matches when the TLS Match elements specified in the rule match.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Expand the **Other Elements** branch.

3) Right-click **TLS Matches**, then select **New TLS Match**.

4) In the **Name** field, enter a unique name.

5) (Optional) Select **Deny Decrypting** to prevent connections from being decrypted for inspection.

Note

Selecting this option prevents the network application where the TLS Match is used from being identified if the traffic is encrypted. If you want to specify that decryption is not necessary for identifying the network application, use the **Application Identifiable by TLS Match Alone** option in the Network Application properties instead.

6) To define whether the server certificate validity is checked and what to match, select an option from the **Match Certificate Validation** drop-down list.

7) Configure the additional settings depending on the option that you selected from the **Match Certificate Validation** drop-down list:
   - **Validation Succeeded** — Click **Add**, then enter the fully qualified domain name to match in the server certificate.
   - **Validation Failed** — (Optional) Select the specific types of invalid certificates to match.
   - **No Validation** — There are no additional settings to configure.

8) Click **OK**.

Access rules for network application detection

To detect network application use, you must create Access rules that define the matching criteria.

The Service cell defines the protocols that are compared to the protocol-related information in each packet's header. You can use Network Application elements directly in the Service cell, or as part of the matching criteria in the **Service Definition** dialog box. Any other criteria in the service definition override the properties of the Network Application element.

Depending on the options that you select in the **Service (Port)** cell of the service definition, you can specify which ports traffic matches:

- **Automatic Port Selection** — The ports that traffic matches are selected automatically depending on the action specified in the rule.
  For rules that allow traffic and for rules with the Continue action, traffic matches on the standard ports defined in the Network Application element. For rules that stop traffic, traffic matches any port where the application can be detected.
- **Any Port** — Traffic matches any port where the application can be detected.
- **Standard Ports** — Traffic matches only the standard ports defined in the Network Application element.

When you add new Access rules, **Automatic Port Selection** is selected by default.

Alternatively, you can use Application Type elements and Tag elements directly in the Service cell. Application Type elements represent general categories of network applications. Tag elements represent all Network Application elements that are associated with that Tag.

Some network applications can open several related connections. If an Access rule that detects network application use identifies a related connection, the related connection is matched against the Access rules again. If the rule that detected the network application use has deep inspection enabled and the related connection matches a rule that has deep inspection enabled, the related connection is matched against the Inspection Policy. No NAT payload modifications are done for the connection that matches the rule that detected the network application use. NAT payload modifications can be done for the related connections according to the policy.

Create Access rules for network application detection

To detect network application use, create Access rules that define the matching criteria.

Note

If a TLS Credentials or Client Protection Certificate Authority element has been uploaded to the engine, adding a Network Application element that allows or requires the use of TLS to an Access rule might enable the decryption of the following TLS traffic:

- TLS traffic from network applications that cannot be identified based on cached network application information.
- TLS traffic that matches an Access rule that enables deep inspection if the **Service** cell contains a Network Application or Service element that does not include a Protocol Agent.
- TLS traffic for which there is no TLS Match with the **Deny Decrypting** option that excludes the traffic from TLS Inspection.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Browse to **Policies**, then browse to the policies of the type that you want to edit.

3) Right-click a policy, then select **Edit <policy type>**.

4) On the **IPv4 Access** or **IPv6 Access** tab, add a rule in one of the following ways:
   - Right-click the last row in an empty rules table, then select **Add Rule**.
   - Right-click the **ID** cell of an existing rule, then select **Add Rule Before** or **Add Rule After**.

5) Drag and drop elements from the **Resources** pane to the **Source** and **Destination** cells, or define source and destination criteria.

6) In the **Action** cell, select the action according to your needs.

7) Define the value of the **Service** cell in one of the following ways:
   - Drag and drop a Network Application, Application Type, or Tag element to the **Service** cell.
   - Right-click the **Service** cell, select **Edit Service**, then add a Network Application, Application Type, or Tag element to the **Network Application** cell.

     Note

     If you add a Service element to the same row, the ports specified in the Service elements override the ports specified in the Network Application elements. You cannot use Network Application elements and Service elements on different rows of the same Service Definition.

8) (Optional) If you created a service definition, right-click the **Service (Port)** cell, then specify which ports traffic matches.
   By default, the ports that traffic matches are selected automatically depending on the action specified in the rule.

9) In the **Logging** cell, select options according to your needs.
   If you want to include information about network application use in the logs, select **Default** or **Enforced** for **Log Application Information**.

10) Click **Save and Install**.

Log network application use

You can optionally log the use of network applications without using network application detection for access control.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Browse to **Policies**, then browse to the policies of the type that you want to edit.

3) Right-click a policy, then select **Edit <policy type>**.

4) On the **IPv4 Access** or **IPv6 Access** tab, add a rule in one of the following ways:
   - Right-click the last row in an empty rules table, then select **Add Rule**.
   - Right-click the **ID** cell of an existing rule, then select **Add Rule Before** or **Add Rule After**.

5) Drag and drop elements from the **Resources** pane to the **Source** and **Destination** cells, or define source and destination criteria.

6) (Optional) Drag and drop a Network Application, Application Type, or Tag element to the **Service** cell.
   It is not necessary to add a Network Application element if you only want to log the use of network applications.

7) In the **Action** cell, select **Continue**.

8) Double-click the **Logging** cell.

9) Select **Override Settings Inherited from Continue Rule(s)**.

10) In the **Log Level** drop-down list, select the log level for traffic that matches the rule.

11) Select **Override Recording Settings Inherited from Continue Rule(s)**.

12) In the **Log Application Information** drop-down list, select **Enforced**.



Note

If a TLS Credentials element or a Client Protection Certificate Authority element has been uploaded to the engine, selecting **Enforced** might enable the decryption of the following TLS traffic:

- TLS traffic from network applications that cannot be identified based on cached network application information.
- TLS traffic that matches an Access rule that enables deep inspection if the **Service** cell contains a Network Application or Service element that does not include a Protocol Agent.
- TLS traffic for which there is no TLS Match with the **Deny Decrypting** option that excludes the traffic from TLS Inspection.

13) Click **OK**.

14) Click **Save and Install**.

Example: blocking network application use

An example of using Network Application elements to block the use of specific network applications.

The administrators at Company A want to allow the use of HTTP in general, but block the use of social media applications from its corporate network. When social media use is detected, the administrators want to redirect users to the corporate security policy page on the company intranet.

The administrators:

1) Create a User Response element to redirect dropped connections to the corporate security policy intranet page.

2) Add the following Access rules:

Source	Destination	Service	Action
Internal networks	Not internal networks expression	Social Media Application Tag	Discard Response: User Response to redirect connections to the intranet page
Internal networks	Not internal networks expression	HTTP	Allow

3) Refresh the engine's policy.

Chapter 58

Defining User Response elements

Contents

- User Response elements and how they work on page 969
- Create User Response elements on page 971

With the User Response element, you can send customized replies to users, instead of just closing an HTTP or HTTPS connection.

User Response elements and how they work

User Response elements allow you to define custom responses that are sent to the user when an HTTP or HTTPS connection is closed.

User Responses make it possible to explain to the user why the connection was closed instead of simply closing the connection with no notification. They help administrators differentiate cases where the Security Engine closes a connection from cases where a technical problem prevents the connection from going through.

When you combine User Responses with browser-based user authentication, you can also redirect users to their original destination after they have authenticated to a Engine. The redirection can be automatic or require the users to click a link to the original HTTP destination address on the user authentication page after they have authenticated.

You can use User Responses in Access rules and in Inspection Policies. Redirection to the user's original HTTP destination after authentication must be configured in the Inspection Policy.



Note

You can also redirect some part of HTTP or HTTPS traffic to an additional security processing service like Forcepoint Remote Browser Isolation. For more information about the configuration to redirect user web traffic to Forcepoint RBI, see Configuring the redirect for Forcepoint Security Engine in Remote Browser Isolation Documentation.

Limitations of User Responses

User Responses have the following limitations:

- To use User Responses with HTTPS traffic, you must enable decryption of HTTPS traffic.
- Some web browsers, such as Mozilla Firefox and Google Chrome, use HTTP Strict Transport Security (HSTS) to enforce the use of HTTPS by default. The end user's web browser might not accept the certificate for TLS inspection when HSTS is used.

User Response configuration overview

Configuring a User Response involves creating a User Response element, defining the response content, and adding the User Response element to a policy.

Follow these general steps to configure User Responses:

1) Create a User Response element.

2) Define the responses that are sent to the users when a connection is closed.

3) Use the User Response element in the Access rules and Inspection rules as required.

Related tasks

Add Access rules on page 899
Add Inspection rules on page 904
Create User Response elements on page 971

Variables for User Response elements and how they work

Variables allow you to dynamically add details about the connection to the message that is shown to users. Users can provide the information to their administrator when they report problems with the connection.

You can add variables to the **Message Response** and **Custom HTML** response types. If the information that the variables represent is available for the connection, the details are automatically added as text to the message that is shown to users. For example, if the client's IP address is 192.168.22.33, the {{SrcIP}} variable is replaced with 192.168.22.33 in the message. If the information is not available, N/A is shown in the message.

The following variables are supported:

Variables for User Response elements

Variable	Description
{{SrcIP}}	The source IP address of the connection.
{{DstIP}}	The destination IP address of the connection.
{{SrcPort}}	The source port of the connection.
{{DstPort}}	The destination port of the connection.
{{User}}	The user name and LDAP domain of the user that made the connection in the format user@domain. If the LDAP domain is not available, only the user name is shown.
{{Group}}	A comma-separated list of User Groups to which the user belongs.
{{Url}}	The URL requested in the connection.
{{UrlCategory}}	The URL Category for category-based URL filtering that the requested URL matches.

Variable	Description
{{Application}}	The Network Application identified in the connection.
{{RuleTag}}	The rule tag of the rule that blocked the connection.
{{TimeStamp}}	The date and time of the connection.
{{FileName}}	The name of the file that was detected in the connection.
{{ArchiveName}}	The file name of the compressed file that was detected in the connection.
{{MalwareName}}	The name of the malware that was detected in the connection.

Create User Response elements

You can define a different User Response entry for each case in which an HTTP or HTTPS connection matches a rule that closes the connection.

The HTTP or HTTPS connection is not allowed to continue in the following cases:

- **Connection Block listed (HTTP only)** — The connection was closed according to a rule with the Apply Block list action.
- **Connection Discarded by Access Rule (HTTP only)** — The connection was discarded according to an Access rule with the Discard action.
- **Connection Terminated by Inspection Rule** — The connection was terminated according to the Inspection Policy.
- **URL Not Allowed** — The connection was closed by rules for URL filtering.
- **Malware Found** — The anti-malware feature detected malware in the connection.



Note

In some cases, such as when inspecting a large file transfer, it is not possible to apply a User Response to HTTPS traffic. In this case, the engine applies the default action for matching traffic. If an HTTPS connection is discarded before the payload is decrypted, no User Response is sent.

You can also redirect users to their original HTTP destination after they have authenticated to a Engine. You automatically redirect the users after they have authenticated, or you can require the users to click a link to the original HTTP destination address on the user authentication page after authentication.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Expand the **Other Elements** branch, then select **Engine Properties**.

3) Right-click **User Responses**, then select **New User Response**.

4) In the **Name** field, enter a unique name.

5) For the entry that you want to change, expand the section.

6) From the **Type of Response** drop-down list, select one of the following responses:
   - **Message Response** — (Optional) Customize the message to display to the user.
   - **URL Redirection** — In the **HTTP or HTTPS URL Destination** field, enter the URL to which the connection is redirected. The URL must begin with `http://` or `https://`.
   - **Custom HTML** — (Optional) Customize the HTML source code for the message to display to the user. The HTML source code must be a complete HTML page, including the <html> and <body> tags.

7) (Optional) To dynamically add details about the connection to the message for the **Custom HTML** and **Message Response** response types, add variables.
   a) Click the location in the HTML source code or message text where you want to add the variable.
   b) Click **Add Variable**, then select the variable that you want to add.
   c) (Optional) To preview the message in your default web browser, click **Preview in Browser**.

8) (Optional) To redirect users to their original destination after they have authenticated to the Engine for the **URL Redirection** response type, define the settings for the redirection.
   a) To enable redirection to the original destination, select **Enable Manual Redirection to Original URL After Authentication**.
   b) (Optional) To automatically redirect the user to the original destination, select **Enable Automatic Redirection to Original URL After Authentication**.

9) Click **OK**.

Chapter 59

Quality of Service

Contents

- Quality of Service (QoS) and how it works on page 973
- Create QoS Class elements on page 981
- Define QoS Policy elements on page 982
- Apply QoS to traffic on page 985
- Examples of bandwidth management and traffic prioritization scenarios on page 986

The Quality of Service (QoS) features allow you to manage bandwidth and prioritize connections on the Security Engines. QoS features are available on Engines, IPS Security Engines, Layer 2 Engines, Master Engines, Virtual Engines, Virtual IPS Security Engines, and Virtual Layer 2 Engines.

Quality of Service (QoS) and how it works

QoS (Quality of Service) allows you to manage the available network bandwidth and make sure that important network services are given priority over less important traffic.

QoS consists of bandwidth management and traffic prioritization. You can use both bandwidth management and traffic prioritization together or bandwidth management or traffic prioritization individually for any given type of traffic.

The QoS features help you in the following ways:

- You can set up a *Guarantee* for a type of traffic that must always be given a certain minimum share of the available bandwidth.
- You can set up a *Limit* for maximum bandwidth for a type of traffic that must never use more than a certain share of the available bandwidth.
- You can set a *Priority* value for the traffic. Higher priority traffic is sent forward to its destination before lower priority traffic if the Security Engine queues packets due to congestion.
- *Active Queue Management* (AQM) reduces the volume of dropped or retransmitted packets when there is network congestion. AQM monitors the average queue size and uses a scheduling algorithm to determine the statistical probability for dropping incoming packets.
- On Security Engines, Master Engines, and Virtual Engines in the Engine/VPN role, you can use QoS Class elements to apply criteria for dynamic link selection in Multi-Link VPNs to traffic.

  Note

  Only layer 3 physical interfaces are supported.

- The Security Engine can read or write DiffServ Code Point (DSCP) type of service (ToS) field markers. The markers allow the Security Engine to be aware of the priorities set by other network equipment. Other equipment is also aware of the priorities set in the QoS Policy. The markers allow you to integrate the Engine with other network equipment that implements QoS management in your own or your ISP's network.

- The Security Engine can collect statistics about traffic that matches Access rules that apply a QoS Class to the traffic. QoS Class-based statistics items are used in Overviews and Reports.

The QoS features have the following limitations:

- QoS is only available on some interface types:

Security Engine role	Interface types
Engine/VPN role	■ Layer 3 physical interfaces ■ Layer 2 physical interfaces of the Inline IPS Interface and Inline Layer 2 Engine type ■ VLAN interfaces ■ Tunnel interfaces ■ SSID interfaces ■ Port group interfaces of an integrated switch **Note** QoS is also available in the properties of policy-based VPNs
IPS, Layer 2 Engine	Physical interfaces of the Inline Interface type

  Bandwidth management and traffic prioritization are not supported on Modem interfaces of Single Engines.
- It is not possible to apply a bandwidth guarantee to incoming Internet traffic on your Internet link. By the time the Security Engine processes the traffic, the bandwidth has already been used. If you want guaranteed bandwidth for a specific portion of your incoming Internet traffic, contact your ISP and ask if they can enforce this guarantee for you.
- If you want to create QoS rules for both incoming and outgoing traffic, you must assign a QoS Policy to at least two interfaces. Incoming traffic is processed according to the Engine, IPS, or Layer 2 Engine policy, and then the QoS Policy is applied to the allowed traffic on the outgoing interface.
- When you use the DSCP Match/Mark rules of a QoS Policy to assign a QoS Class based on the DSCP code in incoming traffic, custom link selection options in the QoS Class elements are not applied to the traffic. Instead, the traffic uses the settings in QoS Class elements in Access rules that override the default link selection options defined in the Network Application or Protocol elements. If there are no matching Access rules that override the default link selection options defined in the Network Application or Protocol elements, the traffic uses the default settings.

Related concepts
Defining Policy-Based VPN elements on page 1193

Bandwidth management

Bandwidth management means giving a guaranteed minimum portion of the available bandwidth to some types of traffic. It also sets limits for how much of the total available bandwidth each type of traffic is allowed to consume at any given time.

You can set a limit, a guarantee, or both for any given type of traffic. These features can be used to guarantee the quality of time-critical communications (even under normal network load), prepare for malfunctions, and limit the total bandwidth needed.

Note

Bandwidth management applies to outbound traffic only. The engine can only indirectly limit the bandwidth use of incoming traffic.

Related concepts

Managing the bandwidth of incoming traffic on page 978

Traffic prioritization

You can ensure that essential or time-critical traffic flows without delays by giving it a higher priority value.

Even under normal traffic conditions, temporary traffic peaks sometimes occur. With many communications, slight delays caused by queuing traffic are not noticeable to the user of the service. However, some connections, such as streaming audio or video, are time-critical, and even relatively minor delays cause noticeable reduction in service quality.

Normally, when packets are queued, they are sent onwards in the same order in which the packets were received. To change this behavior, you can assign priority values to the traffic. For example, you can assign time-critical connections a high priority. High-priority packets are placed before any lower-priority packets in the queue, allowing the fastest possible delivery.

Active Queue Management (AQM) reduces the volume of dropped or retransmitted packets when there is network congestion. AQM monitors the average queue size and uses a scheduling algorithm to determine the statistical probability for dropping incoming packets. If the queue is almost empty, all packets are accepted. As the queue size increases, the probability for dropping incoming packets also increases. When the queue is full, all packets are dropped.

Effects of bandwidth management and prioritization

Bandwidth management and traffic prioritization improve the quality of service for important traffic. However, the quality of service for traffic that you define as less important can decrease.

Both features are configured using the same tools. You can use both bandwidth management and traffic prioritization together, or bandwidth management or traffic prioritization individual for any given type of traffic. Bandwidth management and traffic prioritization are not supported on the Modem interfaces of Single Engines.

Usually the traffic management process allows all connections to proceed, although some traffic can occasionally slow down when the bandwidth limits are reached. If there is prolonged congestion in the network, lower priority traffic eventually starts to time out. If you set priorities without setting any maximum limits or minimum guarantees for bandwidth, high-priority traffic can even use all available bandwidth, blocking all lower-priority traffic.

In most situations, the guaranteed minimum bandwidths given to important connections allow traffic to proceed. However, even traffic with bandwidth guarantees might not get through, if network links are unable to maintain throughput or if the volume of traffic continuously exceeds the throughput. Make sure that your bandwidth limits and guarantees are granular enough to account for losing bandwidth, for example, due to ISP failure in a Multi-Link environment. Track the total bandwidth use, so that you can increase the throughput before problems appear.

The Engine can also read and write DiffServ Code Point (DSCP) markers in type of server (ToS) fields. The markers allow you to integrate the Engine with other network equipment that implements QoS management in your own or your ISP's network.

CAUTION

Inappropriate bandwidth limits and guarantees only disturb traffic. Make sure the guarantees and limits you set are appropriate for the volume of each type of traffic.

Default QoS elements

There are three default QoS Classes: High Priority, Normal Priority, and Low Priority. These QoS Classes are used in the default QoS Policy, *Prioritize*.

The Prioritize QoS Policy is a sample policy that contains simple rules for prioritizing traffic according to the three default QoS Classes. High Priority traffic is assigned the highest possible priority value of 1. The Normal Priority value is 8, and Low Priority is assigned the lowest possible value of 16. The default Prioritize policy does not provide any bandwidth guarantees or limits.

CAUTION

If you set priorities without setting any bandwidth limits or guarantees, high-priority traffic can use all available bandwidth, blocking all lower-priority traffic.

If the default Prioritize policy is sufficient for you, you can use the default QoS Classes and the Prioritize policy as they are. Add the QoS Classes to Access rules, then configure the interfaces to use the Prioritize QoS Policy.

Related tasks

Apply QoS to traffic on page 985
Define QoS Policy elements on page 982

Using bandwidth management and traffic prioritization

You can use bandwidth management and traffic prioritization to ensure critical communications, prepare for network link failures, and restrict non-essential traffic.

Bandwidth management and traffic prioritization are used for the following purposes:

- To ensure the quality of service for time-critical communications during occasional traffic peaks. Even if there is ample bandwidth available, short periods of congestion can degrade the quality of some types of communications.
- To prepare for severe congestion, caused by the loss of network links when there are technical problems. In a Multi-Link environment, you can have several NetLinks. Ideally, the throughput of these links is large enough that each link can alone handle the traffic. However, if it is not a viable option, it might become necessary to choose which connections are given priority if network connections are lost.
- To reduce the total bandwidth needed, if it is not possible to increase throughput of the network links or add new links. For example, important services (such as VPN connections and clients' connections to extranets) can be given priority over Internet browsing (all HTTP connections or based on IP addresses).

Designing QoS policy elements

The purpose of the QoS Policy is to determine which limit, guarantee, or priority is given to traffic marked with a certain QoS Class.

Each QoS Class can appear in only one (active) rule on each tab of a QoS Policy. The same QoS Class can be used on both the **QoS** tab and the **DSCP Match/Mark** tab of the same QoS Policy. The order of the QoS rules does not matter. The classification of the traffic is made using Access rules.

Except for the QoS Class, all other cells for rules on the **QoS** tab are optional. However, at least one of the other cells must be filled for the rule to affect the traffic. None of the cells exclude any of the other cells, so you are free to select which cells you want to use for any given QoS Class. It is not necessary to define the use of all available bandwidth in your QoS Policy. The bandwidth outside the guarantees, as well as bandwidth within the guarantees that is not used for traffic, is used to handle the traffic that has no specific QoS rule. The traffic is handled on the normal first-come-first-served-basis, using the medium priority of 8.

CAUTION

If your guarantees are equal to the total throughput of an interface, any traffic without a guarantee is blocked if all guarantees are fully used.

When you save the QoS Policy, the system checks for contradictions within each rule, such as a rule that sets a limit that is lower than the guarantee for the same rule. When you refresh the engine's configuration, the QoS Policies defined for the engine's interfaces are checked again, comparing the QoS rules to the throughput values set for the interfaces. The values can be automatically scaled down if the sum of all guarantees in the QoS Policy exceeds the interface's throughput. However, the values are never scaled up.

The values for the bandwidth limits and guarantees can be entered either in kilobits or as percentages of the total throughput of the interface. Technically, nothing prevents you from using both ways of entering values even in the same rule. However, it is recommended to use one way of entering the values consistently throughout each QoS Policy. Using mixed methods of entering the values makes it more difficult for the administrators to read the QoS policy. Mixed methods can also prevent the system from checking for issues when you save the QoS Policy, as the throughputs of the interfaces are not known. If the QoS Policy cannot be checked when you save it, it is checked when the Engine Policy is installed. Mistakes in the QoS Policy can prevent you from installing or refreshing the Engine Policy.

Communicating DSCP markers to other network equipment to prioritize traffic

DSCP (DiffServ type of service field) markers in the traffic are a standard way to indicate priorities in network traffic. You and your ISP might have routers that decide how to handle packets based on the priority of the traffic.

It is possible to read or write DSCP markers for a particular type of traffic without configuring Access rules to apply a QoS Class to the traffic. The matching is done based on the QoS Policy. When a packet that matches a particular protocol comes in, the Security Engine reads the DSCP markers and assigns a QoS Class according to the DSCP Match/Mark rules of the QoS Policy. When the packet is sent out, the Security Engine writes a DSCP mark in the packets. The marking is based on the QoS Class according to the DSCP Match/Mark rules of the QoS Policy on the interface through which the traffic leaves the Security Engine.

Important

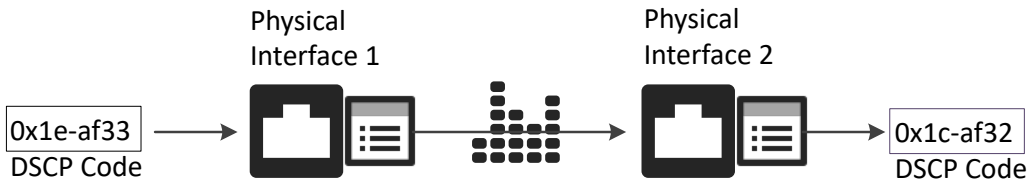
- If a value is specified in the **DSCP Match** cell, the engine considers all incoming traffic matching the **DSCP Match** cell to belong to the QoS Class that is mentioned in the **QoS Class** cell.
- If a value is specified in the **DSCP Mark** cell, the engine marks all traffic matching the defined QoS Class with the value specified in the **DSCP Mark** cell.

The markers allow you to:

- Communicate the priority of this traffic to other devices that support QoS.
- Convert the packet to use a different classification scheme, if the QoS Class was originally assigned to matching traffic by a DSCP match in the source interface's QoS Policy.
- Remove the DSCP classification set by other devices by entering 0 as the value (shown in the policy as 0x00).

Two QoS Policies on two Physical Interfaces can be used together to translate between two different DSCP schemes as shown in the illustration.

Translating between two DSCP schemes



In the illustration, the packets arrive at Physical Interface 1. The engine reads the existing DSCP value and compares it to the QoS Policy assigned to Physical Interface 1. The policy has a DSCP Match rule for the DSCP marker with an associated QoS Class, which is then assigned to this traffic.

Note

The same traffic must not match any engine Access rule with a QoS Class definition. The QoS Class in the Access rule overrides the QoS Class that is assigned based on the DSCP marker.

When the packets are sent out through Physical Interface 2, the Engine checks the QoS Policy assigned to this Physical Interface. In this QoS Policy, a DSCP Match/Mark rule defines that traffic with the assigned QoS Class is marked with a DSCP marker specified in the rule. The engine overwrites the original DSCP marker before sending the packets onwards.

- By default, the DSCP mark for the encrypted ESP packet in VPN traffic is inherited from the plaintext packet. Selecting a QoS Policy in the properties of the policy-based VPN makes it possible to mark the ESP packet after encryption.
- Priorities, limits, and guarantees are applied. DSCP codes are written to outgoing packets on the interface that the traffic uses to exit the Security Engine according to the QoS Policy and interface speed defined for that interface.
- For packets entering the Security Engine, the QoS Policy on that interface is only used for reading DSCP codes and matching them to QoS Classes for further use. It is the only QoS operation that is done on the interface that the traffic uses to enter the Security Engine.
  Example: A new packet enters a Engine through interface A and leaves the Engine through interface B. The priorities, guarantees, and limits configured on interface A are ignored for packets in this direction. Any priorities, guarantees, and limits are configured and applied on interface B. If the packet contains a DSCP code when entering the Engine, the DSCP code is read and matched to a QoS Class on interface A. If a new DSCP code is (over)written in the packet, the new code is written on interface B.

Managing the bandwidth of incoming traffic

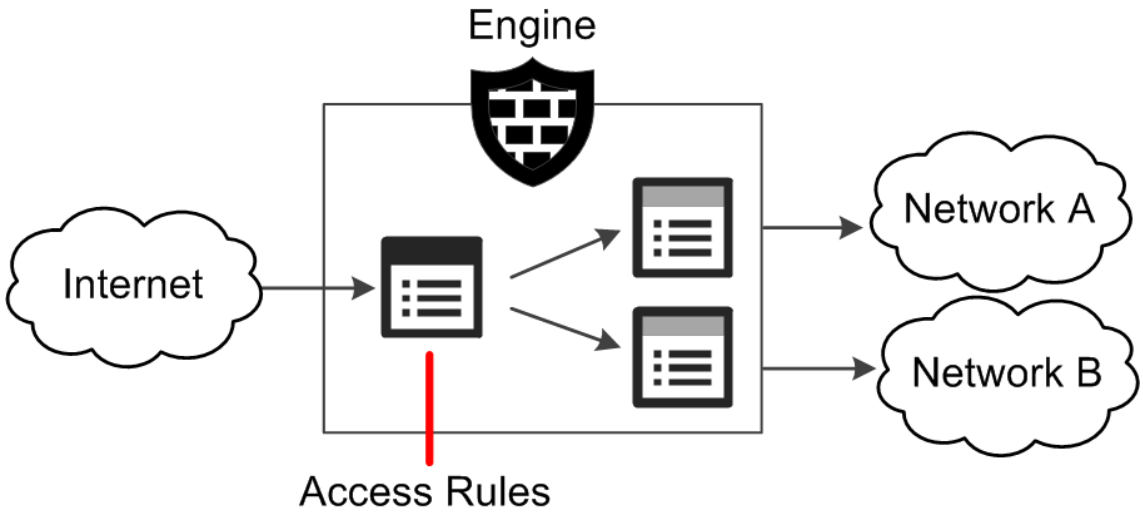
Bandwidth management and traffic prioritization are most useful for managing outgoing traffic.

Bandwidth management and prioritization usually help manage the quality of service for traffic going out through Internet links. These links are often the choke point in a corporate network due to the costs associated with increasing the bandwidth.

Controlling incoming traffic is more difficult because when the Security Engine processes the traffic, the packets have already traveled through congested links and used bandwidth. Still, you might be able to limit some types of incoming traffic in a limited way. In this case, only limits apply. To set guarantees and priorities for traffic, consider other solutions, such as arranging for your ISP to implement traffic management before the traffic is passed to your Internet links.

To limit the bandwidth incoming traffic consumes, you can apply the QoS Policy on the Security Engine's interfaces that are connected to the internal network. This configuration is shown in the following illustration.

Applying QoS to incoming traffic



In the illustration, traffic is checked against the Access Rules, and allowed traffic is assigned a QoS Class. At the interfaces connected to the internal network, the QoS Policies limiting the bandwidth use are enforced as the traffic is sent onwards.

Limiting the bandwidth in this way requires that the application that is the source of the traffic scales down the transmissions to match the available bandwidth. If an application does not scale down its bandwidth use, any limits you set have no effect. The only option is to control the traffic at your ISP before it reaches the Security Engine.

Collecting statistics based on QoS Classes

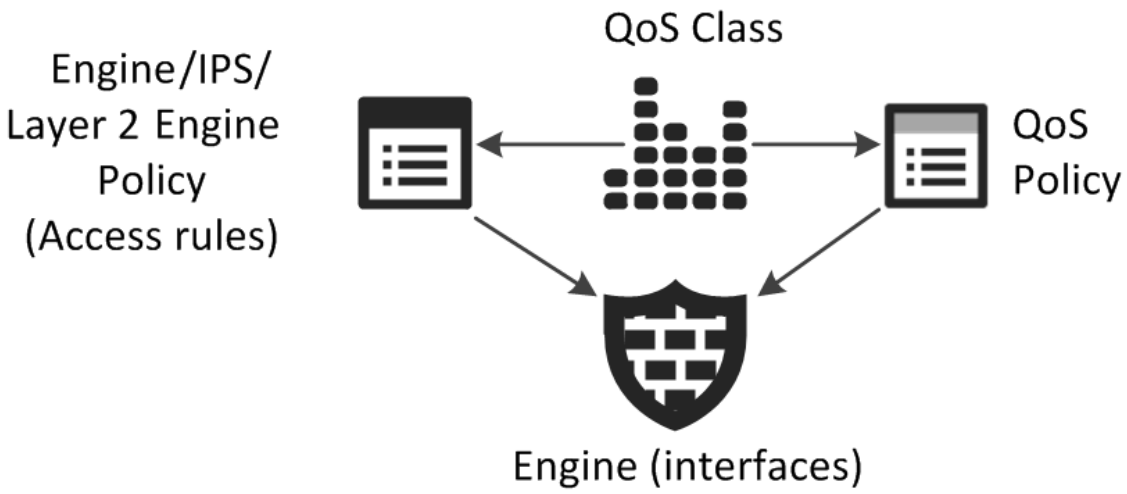
QoS Class-based statistics items are available even when QoS is not used for bandwidth management and traffic prioritization.

Selecting the "QoS Statistics Only" QoS Mode for an interface allows the collection of QoS Class-based counters without activating any other QoS feature. No QoS Policy is needed in this case, but you must define Access rules to apply QoS Classes to traffic. QoS Class-based statistics items can be used in Overviews and Reports.

QoS configuration overview

To apply QoS to traffic, you must create QoS classes and assign them to different types of traffic. You must also create QoS Policies that determine how traffic is handled, and define a QoS Mode for each interface.

Elements in the QoS configuration



Follow these general steps to configure QoS:

1) (Optional) Create a QoS Class element for each type of traffic that you want to handle differently on any single network interface.

2) (Optional) Create one or more QoS Policies to define how each type of traffic is handled on the interfaces.

3) Assign QoS Classes to different types of traffic in your Access rules.

4) Define the QoS Mode of each interface.

Note

You can select a QoS Mode and define a bandwidth for traffic in the properties of a Physical, VLAN, ADSL, Tunnel, SSID Interface, or Port Group Interface of an integrated Switch. Each Physical, VLAN, Tunnel, ADSL, SSID, or Port Group Interface has separate QoS settings.

Bandwidth management and traffic prioritization are configured in QoS Policies. The policies contain rules for the bandwidth guarantees, limits, and priorities you want to set for each type of traffic. The QoS Policies do not contain traffic profiles: to define which QoS rule affects which type of traffic, the same QoS Class element is used in QoS Policies and Access rules, to link them.

The QoS Mode for each interface defines how QoS is applied to the interface. By default, No QoS is selected. You can select a QoS Mode and define a bandwidth for traffic in the properties of a Physical, VLAN, ADSL, Tunnel, SSID Interface, or Port Group Interface. You can select different QoS Modes for each interface. It is not mandatory to use QoS on all interfaces of the same Security Engine. QoS is not supported on Capture Interfaces.

There are two ways the QoS Class can be applied to a packet:

- If traffic contains a DSCP code when entering the Security Engine, and DSCP handling and throttling or full QoS are enabled, the Security Engine checks if the interface has a QoS Policy. If the DSCP Match/Mark tab of the QoS Policy defines a QoS Class match for that code, the selected QoS Class is applied to the traffic.
- When traffic is inspected against the policy, the traffic might match an Access rule that includes a QoS Class. The QoS Class specified in the QoS Class cell is always applied to the traffic, overwriting any other QoS Class the traffic might have been assigned. Access rules are not needed if you only want to use DSCP handling and throttling.

Using the QoS Class as a matching criterion, the Security Engine checks if the interface that the packet uses to exit the Security Engine has a QoS Policy. If the QoS Policy contains a rule with the same QoS Class defined, the next steps depend on the QoS rules and the current traffic situation. The QoS rule is applied to the connection and packets are dropped, sent directly, or sent into the queue.

Related concepts
Network interfaces for Security Engine on page 543

Related tasks
Create QoS Class elements on page 981
Define QoS Policy elements on page 982
Apply QoS to traffic on page 985
Define DSCP Match/Mark rules in QoS Policy elements on page 984

Create QoS Class elements

QoS Classes are used to collect QoS statistics about traffic or to create a link between the Access rules and the QoS Policy.

When traffic matches an Access rule, the QoS Class defined in the rule is applied to the traffic. You can also use QoS Classes in Outbound Multi-Link elements to adjust the load balancing of different types of connections.

There is a default QoS Policy and three default QoS Classes. The classes can be used to set priorities for high, normal, and low priority traffic, without any bandwidth guarantees or limits.

You must create one QoS Class for each rule you plan to add in any single QoS Policy, as the QoS Policy cannot contain any overlapping rules. You can create as many QoS Classes as necessary. The QoS Policy must not have overlapping rules. You must create one QoS Class for each rule you plan to add to the QoS Policy (each type of traffic must have its own QoS Class). The same QoS Class can be used in multiple Access rules, so several Access rules can point matching traffic to the same QoS rule.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.
2) Browse to **Other Elements** > **QoS Classes**.
3) Right-click **QoS Classes**, then select **New QoS Class**.
4) Configure the settings, then click **OK**.

Result

The new QoS Class is added to the resources and can now be used in policies.

Define QoS Policy elements

QoS policies determine the rules that the Security Engine follows when it decides which traffic is given priority and how the available bandwidth is allocated.

- One QoS Policy can be assigned for each Physical Interface, VLAN Interface, Tunnel Interface, ADSL Interface, SSID Interface, Port Group Interface of an integrated Switch, and policy-based VPN. You can assign the same QoS Policy to several interfaces.
- QoS Policies are tables of QoS rules and DSCP Match/Mark rules. If you only want to collect QoS statistics about traffic, you do not need to define a QoS Policy.
- Because the QoS rules are separate from the Access rules, you can flexibly design the rules. For example, you can create different QoS Policies for different interfaces of the same Security Engine.
- All cells in the QoS rules are applied to outgoing packets. When using Full QoS, packets that do not match a QoS rule are handled with priority 8 (middle of the scale) without bandwidth guarantees or limits.
- All cells in the DSCP Match/Mark rules except the DSCP Match cell are applied to outgoing packets. If packets do not match a DSCP Match/Mark rule, DSCP markers in the traffic are preserved, but do not affect how the Security Engine handles traffic.

The *QoS Mode* for each interface defines how QoS is applied to the interface. By default, No QoS is selected. You can select a QoS Mode and define a bandwidth for traffic in the properties of a Physical, VLAN, ADSL, Tunnel, SSID, or Port Group Interface. You can select different QoS Modes for each interface. It is not mandatory to use QoS on all interfaces of the same Security Engine.

When using Full QoS, define the available throughput in the properties for each Physical, VLAN, ADSL, Tunnel, SSID, or Port Group Interface or whose throughput you want to manage. There is no way to automatically find out how much bandwidth each interface has. The throughput must correspond to the actual throughput that interface offers to clients, that is, the outbound bandwidth of an Internet link that is connected to the interface. If there are VLANs on a Physical Interface, the settings are only available in the properties of each VLAN.



Note

When you define the throughput of an interface, the Security Engine always scales the traffic to this throughput. Take special care that you set the throughput value correctly.

If you use Multi-Link where more than one NetLink is connected to the same Physical Interface, the throughput setting depends on the Multi-Link configuration:

- If you are using load-balancing Multi-Link, set the throughput to the combined outbound bandwidth of all Internet links behind the Physical Interface.
- If you are using standby NetLinks, set the throughput to the outbound bandwidth of the primary (active) NetLink. When the bandwidth of the backup NetLink is lower, set the throughput to the speed of the primary NetLink, as it is the most used link.

Policy-based VPNs can optionally use a QoS Policy to define how DSCP matching or marking is done for VPN traffic. In policy-based VPN traffic, the DSCP mark for the encrypted ESP packet is normally inherited from the plaintext packet. Selecting a QoS Policy for the policy-based VPN makes it possible to mark the ESP packet after encryption. Because the total throughput is undefined, Guarantees and Priorities cannot be used for policy-based VPN traffic.

Related concepts

Collecting statistics based on QoS Classes on page 979
Types of VPNs in Forcepoint Network Security Platform on page 1157

Create QoS Policy elements

Create separate QoS Policy elements for QoS rules.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Browse to **Policies** > **QoS Policies**.
3) Right-click **QoS Policies**, then select **New QoS Policy**.
4) In the **Name** field, enter a unique name.
5) Click **OK**.
   The new QoS Policy opens for editing.

Define QoS rules in QoS Policy elements

After creating a QoS Policy, add QoS rules to it.

Follow these general guidelines when editing QoS rules:

- The order of the rules in a QoS Policy does not affect how the Security Engine handles traffic, as the match is made based on the QoS Class.
- If you want to use the same QoS Policy on interfaces that have different types of throughput and use the Full QoS Mode, enter Guarantees and Limits as percentages.
- If you want to use the QoS Policy on interfaces that use DSCP Handling and Throttling QoS Mode, enter Guarantees and Limits in kilobits per second.
- Operations are made according to the matching rule in the QoS Policy that is assigned to the interface that the traffic uses to exit the Security Engine.
- The rules do not need to cover all traffic. When Full QoS is used, traffic that is not covered is given a priority of 8 without limits or guarantees.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Browse to **Policies** > **QoS Policies**.
3) Right-click the policy you want to edit, then select **Edit QoS policy**.
4) Right-click the **ID** cell of a rule in the QoS Policy, select **Add Rule Before** or **Add Rule After** or the **ID** cell of the Not Classified rule, then select **Add Rule**.
   A new blank rule is added.

5) Click the **QoS Class** cell in the new rule, then select the QoS Class.

6) Define how the Security Engine handles traffic in this QoS Class using any combination of the available options.
   Each one is optional.

7) Save the QoS Policy.

8) Refresh the policy to transfer the changes.

Related concepts
Network interfaces for Security Engine on page 543

Define DSCP Match/Mark rules in QoS Policy elements

Read or write DSCP markers to communicate traffic priorities with other network equipment in your environment.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Browse to **Policies** > **QoS Policies**.

3) Right-click the policy you want to edit, then select **Edit QoS policy**.
   The QoS Policy opens in the **Policy Editing** view.

4) Click the **DSCP Match/Mark** tab in the QoS Policy.

5) Right-click the **ID** cell of an existing rule, then select **Add Rule Before** or **Add Rule After** or the **ID** cell of the Not Classified rule, then select **Add Rule**.
   A new blank rule is added.

6) Click the **QoS Class** cell in the new rule, then select the QoS Class.

7) Define how the Security Engine reads or writes DSCP codes for traffic in this QoS Class using the available options.

8) Save the QoS Policy.

9) Refresh the Security Engine's policy to transfer the changes.

Related concepts
Network interfaces for Security Engine on page 543

Apply QoS to traffic

You can apply QoS to traffic by selecting a QoS Class for an Access rule. The same QoS Class can appear in several Access rules.

You can insert a QoS Class in an Access rule that allows traffic or in an Access rule that uses the Continue action to set the same QoS Class for several rules. This way, you can assign a specific QoS Class to any traffic that you can match with a single Access rule. If you only want to collect QoS statistics about traffic, define Access rules to assign a QoS Class to the traffic.

The rules on the **QoS** tab of the QoS Policy are linked to different types of traffic using the QoS Classes. QoS Classes are matched to traffic in the Access rules with the following actions:

- Access rules with the **Allow** Action set a QoS Class for traffic that matches the rules.
- Access rules with the **Continue** Action set a QoS Class for all subsequent matching rules that have no specific QoS Class defined.
- If a **VPN Action** setting is selected in the Action options, a QoS Class is set for VPN traffic. If there is no **VPN Action** setting selected, incoming VPN traffic can also match after decryption. Otherwise, for outgoing traffic, encryption is done after the QoS Policy is checked. For incoming traffic, decryption is done before the QoS Policy is checked.

If you want to read and use DSCP markers set by other devices, the QoS Class is assigned according to rules on the **DSCP Match/Mark** tab of the QoS Policy.

Note

If traffic is assigned a QoS Class using a DSCP Match rule, the same traffic must not match Access rules that assign a different QoS Class to the same traffic. Such Access rules override the QoS Class that has been set with a DSCP Match rule.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Open the Engine, IPS, Layer 2 Engine, or Layer 2 Interface Policy for editing.

2) Click the **QoS Class** cell of a rule that allows traffic or a Continue rule, then drag and drop a QoS Class element into the cell.
   - The QoS Class links connections to a rule on the **QoS** tab of the QoS Policy. There can be different rules in different QoS Policies for the same QoS Class.
   - Packets in both directions of a connection are assigned the same QoS Class (when connection tracking is active for the rule).
   - The applied QoS Class is shown in the logs. You can also generate reports based on this information.

3) Refresh the policy to transfer the changes.

Related concepts
Getting started with Access rules on page 831

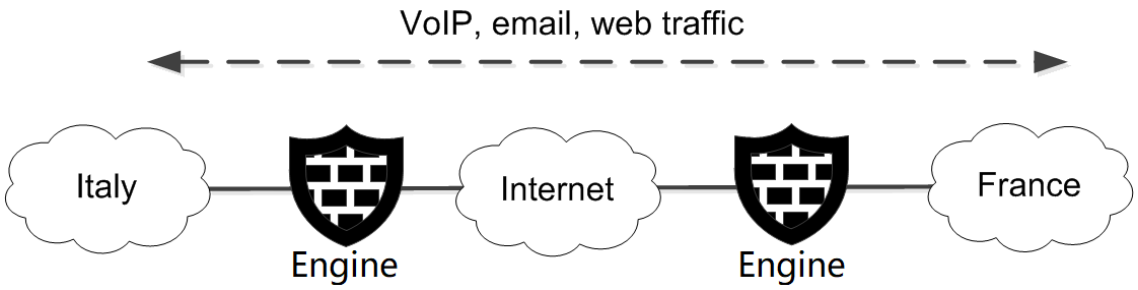
Examples of bandwidth management and traffic prioritization scenarios

There are common uses for the bandwidth management features and general steps on how each scenario is configured.

Example: ensuring quality of important communications

In this example, Company A has two offices, one in Italy and one in France. The company has decided to replace phone lines with VoIP telephony.

Company A networks



The illustration shows the two offices and the traffic between them. Telephone and email connections are an important tool for the employees, who use these services to communicate with team members at the other office. Also, employees at the Italian office must be able to use web-based tools at the French site. The administrators determine the priorities as follows:

- The VoIP streaming audio is not only important, but it is also a time-critical service. VoIP streaming audio must have the highest priority.
- Even though business email is important, email does not need to be delivered immediately after it has been sent. This traffic can be assigned a lower priority.
- The web-based services are not time-critical, but delays and time-outs can annoy employees. The company decides to give these services a lower priority than VoIP, but a higher priority than email. It is not necessary to define a specific QoS Class for the medium-priority traffic because all traffic that is not classified is assigned a medium priority.

The internal networks are fast, so there is no need to implement QoS Policies for those interfaces. Only interfaces connected to the Internet need a QoS Policy. The administrators decide that the same QoS Policy can be used at both sites, and that the default elements and the default Prioritize policy are suitable.

So now they:

1) Add the QoS Class “High Priority” to Access rules that permit VoIP traffic.

2) Add the QoS Class “Low Priority” to Access rules that permit email traffic.

3) Define the QoS Policy Prioritize to be used on the interfaces connected to the Internet at both sites. They also define the types of interface throughput.

4) Refresh the policies of both engines.

Example: using QoS to prepare for ISP breakdown

Company B decides to use Multi-Link to ensure high availability of network connections for important business communications.

The company, an engineering subcontractor, is concerned about two types of connections:

- A VPN connection they have for accessing the internal tools and resources of an important client when doing work for them.
- HTTPS connections to the extranet server that the company’s clients use to check the status of projects.

The company has a tight budget, and the cost of having enough bandwidth in both links even during peak hours is deemed too high. They decide that only the two most important types of traffic must get through if one ISP link goes down during peak hours. The company determines that 500 kbps is enough to handle those connections, so they subscribe to 512 kbps links from two different ISPs. None of the communications are especially time-critical, so the company decides not to prioritize the traffic.

Then the administrators:

1) Create a QoS Policy and two QoS Classes, called VPN and Extranet.

2) Create the QoS rules for the important connections by filling in the following cells:

QoS rules in QoS Policy for Company B

QoS Class	Guarantee
VPN	400
Extranet	100

3) Add the QoS Class “VPN” to the VPN rule for outbound traffic in the Engine’s Access rules.

4) Add the QoS Class “Extranet” to the Access rule that allows outbound connections from the company extranet.

5) Define the types of throughput and select the custom QoS Policy to be used for the Physical Interfaces that correspond to the ISP links on the engine.

6) Refresh the policy of the engine.

Example: limiting bandwidth

Company C has experienced a radical increase in the amount of network traffic and wants to limit non-essential traffic to prioritize business communications.

It seems that many employees use bandwidth-intensive services, download large files, and listen to high-quality Internet radio streams. The situation is starting to slow down business communications. Management would rather prohibit connections that are not directly work-related than fund the required increase in bandwidth.

However, the administrators suggest a different approach: limiting the portion of the bandwidth that non-essential traffic can use, so some employees can still listen to Internet radio, while business communications are guaranteed the bandwidth they need. To ensure the quick delivery of time-critical business communications, they also decide to prioritize the traffic using the three default QoS Classes.

The administrators:

1) Create a custom QoS Policy with the following rules:

QoS rules in QoS Policy for Company C

QoS Class	Priority	Guarantee	Limit
High Priority	1	35%	90%
Normal Priority	8	55%	90%
Low Priority	16	5%	50%

   - Normal Priority traffic gets the largest guaranteed portion of the bandwidth because it has the largest volume.
   - High Priority and Normal Priority traffic can each use up to 90% of the bandwidth. Low Priority traffic cannot consume more than 50% of the available bandwidth even if there is more bandwidth available. In this configuration, there must be traffic in at least two of the classes for the bandwidth to be used up to 100%.
   - Even Low Priority traffic is given 5% of the bandwidth to avoid total loss of service, which can cause more complaints from users than slowed-down service.

2) Place a Continue rule at the top of the engine Access rules that includes the Normal Priority QoS Class. This way, all traffic that is not classified as High Priority or Low Priority is classified as Normal Priority.

3) Edit the Access rules to assign QoS Classes to traffic:
   - Place the High Priority QoS Class into Access rules that permit important traffic.
   - Place the Low Priority QoS Class into Access rules that permit low-importance traffic.

4) Define the types of throughput and select the new custom QoS Policy to be used for the Physical Interfaces connected to the Internet on the engine.

5) Refresh the policy of the engine.

Chapter 60

Anti-malware scanning

Contents

- Getting started with anti-malware scanning on page 989
- Selecting traffic for anti-malware scanning on page 990
- Enable anti-malware on the Security Engine on page 990
- Manually update the anti-malware database on page 991
- View the status of the anti-malware database on page 992

An anti-malware scanner compares network traffic against an anti-malware database to search for viruses and other malware. If malware is found, the traffic is stopped or content is stripped out.

Getting started with anti-malware scanning

Anti-malware scanning checks files against an anti-malware database to find harmful content, such as viruses. If the scanner detects infected files, it strips them out.

Note

This feature requires a separate license.

The anti-malware scanner can inspect IPv4 traffic. The supported protocols in anti-malware inspection are FTP, HTTP, HTTPS, IMAP, POP3, and SMTP.

Anti-malware is supported on all Single and Clustered Security Engines and Virtual Engines (configured on the Master Engine).

Because the data being inspected is not synchronized between the nodes, connections that are undergoing anti-malware scanning at the time of a failover are dropped. The applications must reopen the connections.

Related concepts

Getting started with forwarding traffic on page 1087

Selecting traffic for anti-malware scanning

The File Filtering Policy used on the engine policy determines which traffic is inspected for malware.

Activating the scanning always also activates deep packet inspection for the same traffic (the traffic is also checked against the Inspection rules). You can deactivate the anti-malware scanning if the download source is trusted and the download process takes too long. You must define the services individually in the Service cell to enable deep inspection and anti-malware scanning for them.

To not scan for certain destinations, create a more specific IPv4 Access rule before a more general one that does not have the file filtering option defined.

For some content delivered through HTTP or HTTPS, anti-malware scanning might not be feasible. For example, you might want to prevent videoconferencing sessions from being scanned for malware, to avoid any increase in latency.

Enable anti-malware on the Security Engine

To use anti-malware, you must enable the anti-malware feature in the Engine Editor.

The anti-malware settings in the Engine Editor allow you to set a schedule for downloading updates to the anti-malware database and change the settings for logging the malware found in network traffic.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click an Security Engine, then select **Edit <element type>**.

3) In the navigation pane on the left, browse to **Add-Ons** > **Anti-Malware**.

4) Select **Enable**.

5) Select the log level from the **Malware Log Level** drop-down list.
   - The log levels are the same as used in Access rules.
   - If you selected **Alert** in the **Malware Log Level** drop-down list, select the Alert element from the list.

6) In the **Malware Signature Update Settings** section, select how often the engine checks for updates to the anti-malware database.



Note

The engine queries DNS servers to resolve the anti-malware database URLs. Define at least one DNS IP address on the **General** branch of the Engine Editor.

7) Enter the URL of the anti-malware database mirror in the **Mirror(s)** field.
   The engines contact the mirror to update the anti-malware database. Separate multiple addresses with commas.

8) Continue the configuration in one of the following ways:
   - If you have not yet defined when to use anti-malware inspection, edit the rules in the File Filtering Policy.
   - Otherwise, click **Save and Refresh** to transfer the changes.

Manually update the anti-malware database

If you want the engine to have an up-to-date anti-malware database before you upgrade the engine, you can manually update the anti-malware database on an Security Engine or individual nodes of an Security Engine Cluster.

The update command in the SMC Client uses the default anti-malware database mirror URL.

Note

You must update the anti-malware database by running the `avdbfetch` script on the engine command line in the following cases:

- To use an alternative mirror
- To use an HTTP proxy
- To update the anti-malware database by copying the database to the engine from a USB memory stick

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Dashboard** > **Engines Dashboard**.

2) Expand the nodes of the engine for which you want to update the anti-malware database.

3) Right-click an engine node, then select **Commands** > **Update Malware Database**.

Result

The anti-malware database update begins. The update can take a while.

View the status of the anti-malware database

You can view the progress of the anti-malware database update or check when the anti-malware database was last updated.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Dashboard** > **Engines Dashboard**.
2) Expand the nodes of the engine which you want to view the status of.
3) In the **Info** pane, click the **Status** tab. Expand the **Anti-Malware** branch to view the details.

Chapter 61

File filtering

Contents

- How file filtering works on page 993
- Integrate on-premises DLP servers with Forcepoint Network Security Platform on page 995
- Integrate file reputation services and sandboxes on page 997
- Integrate McAfee GTI file reputation with Forcepoint Network Security Platform on page 997
- Forcepoint Advanced Malware Detection and how it works on page 998
- Integrate a Forcepoint Advanced Malware Detection appliance with Forcepoint Network Security Platform on page 1003
- Define Sandbox Service elements on page 1003
- Connect Forcepoint Network Security Platform to a sandbox service on page 1004
- View sandbox analysis reports on page 1006
- Restrict file types with file filtering on page 1007
- Support for McAfee Advanced Threat Defense on page 1009

Monitoring and restricting what data is sent out is an important part of data loss prevention (DLP). File filtering allows you to restrict the file types that are allowed in and out through the engine, and to apply malware detection to files.

How file filtering works

Use file filtering to scan files for malware and to restrict which file types are allowed through the Security Engine.

Note

File filtering is only available for the following protocols: FTP, HTTP, HTTPS, IMAP, IMAPS, SMTP, POP3, and POP3S.

You can configure one or more malware detection methods that are applied to the traffic that matches the rules in the File Filtering Policy. Scanning is done in the following order:

1) The file is checked against file reputation information in the engine's cache.
   If no match is found, the configured malware detection methods are applied to the traffic in the order listed here.

2) If DLP scanning is enabled, the Security Engine forwards the file to an external DLP server.

Note

DLP scanning is typically used for outbound file transfers to prevent sensitive data from being sent out. DLP scanning and other file filtering methods are not typically applied to the same traffic.

3) The file is scanned using one of the following file reputation services:
   - McAfee® Threat Intelligence Exchange (TIE)
   - McAfee® Global Threat Intelligence™ (McAfee GTI)

4) The file is scanned using anti-malware scan on the Security Engine.

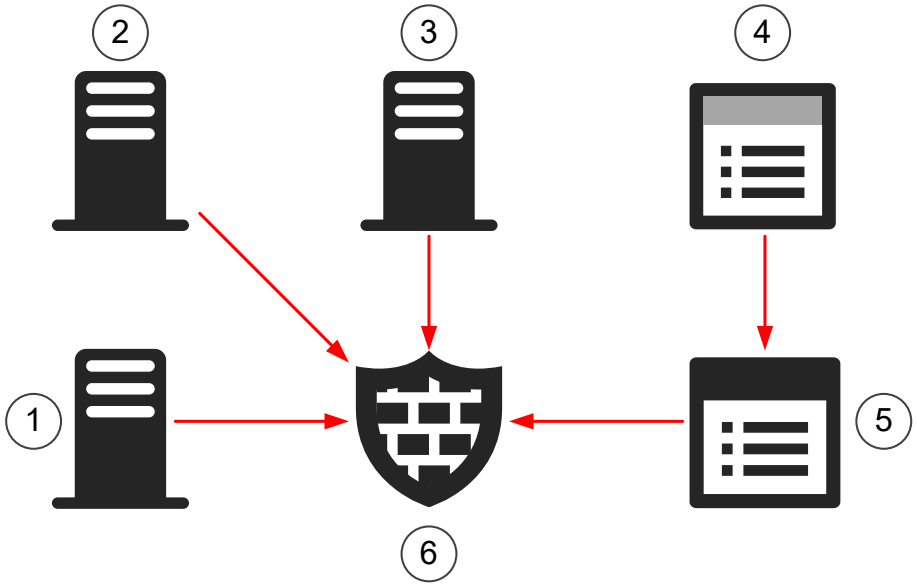
5) The file is scanned using one of the following sandboxes:
   - Cloud Sandbox — Forcepoint Advanced Malware Detection
   - Local Sandbox — Forcepoint Advanced Malware Detection

The Security Engine allows or blocks the file according to the action defined in the File Filtering Policy.

File filtering configuration overview

To use file filtering, integrate or configure one or more malware detection or data protection methods and add rules for file filtering.

Components and elements in the configuration



1. On-premises ICAP DLP server
2. McAfee Global Threat Intelligence file reputation service or McAfee Threat Intelligence Exchange (TIE) file reputation service
3. Forcepoint Advanced Malware Detection sandbox service
4. File Filtering Policy
5. Engine, IPS, Layer 2 Engine, or Layer 2 Interface Policy
6. Security Engine

The configuration of file filtering consists of the following general steps:

1) (Optional) Integrate one or more on-premises DLP servers for data protection.

2) Integrate or configure one or more malware detection methods.
   - (Optional) Integrate Forcepoint Network Security Platform with one of the following file reputation services:
     - McAfee Global Threat Intelligence
     - McAfee Threat Intelligence Exchange (TIE)
   - (Optional) Integrate Forcepoint Network Security Platform with the Forcepoint Advanced Malware Detection sandbox service.
   - (Optional) Enable **Anti-Malware** for the engine.

   Note

   Each malware detection method is optional, but you must integrate or configure at least one malware detection method to use file filtering.

3) Create a File Filtering Policy element and add rules to the File Filtering Policy.

4) Enable file filtering in a Engine, IPS, Layer 2 Engine, or Layer 2 Interface Policy element.

   Note

   To enable file filtering in a Layer 2 Interface Policy, you must enable file filtering and select the File Filtering Policy in the Engine Policy.

Integrate on-premises DLP servers with Forcepoint Network Security Platform

You can integrate on-premises DLP servers, such as Forcepoint DLP, with Forcepoint Network Security Platform and use them as a scanning method in the file filtering policy.

Before you begin

To use TLS to secure the connection to the DLP server, you must:

- Create a TLS Profile element that specifies the settings for cryptography, trusted certificate authorities, and the TLS version used in TLS-protected communication with the DLP server.
- Configure TLS on the ICAP server or in the environment in which the TLS server is deployed. See the documentation for your DLP server for more information.

DLP scanning is typically used for outbound file transfers to prevent sensitive data from being sent out. DLP scanning is supported for the following protocols: FTP, HTTP, HTTPS, IMAP, IMAPS, POP3, POP3S, and SMTP.

Security Engines communicate with the integrated DLP servers using the ICAP protocol. ICAP Server elements represent the DLP servers. You can integrate one or more ICAP servers with the Security Engine. When you integrate multiple ICAP servers, traffic is balanced between the ICAP servers.

The Security Engine sends files to the ICAP server, then allows or blocks the file transfers depending on the response it receives from the ICAP server. The Security Engine can optionally add headers to the request to communicate the user and IP address from which the original request came to the ICAP server. You can specify

the header names to use for each of these headers. By default, the standard names are used. If you leave the name of the header blank, the specified header is not sent to the ICAP server.

Integrating on-premises DLP servers with Forcepoint Network Security Platform has the following limitations:

- Only the ICAP protocol is supported. The DLP server must support ICAP.
- Only the REQMOD method is supported for sending files to the DLP server.
- Only on-premises DLP servers are supported. Cloud-based DLP services are not supported.

Each Security Engine node is counted as a separate client of each ICAP server. The same Security Engine node can make several connections to the same ICAP server, up to the Max-Connections value returned in the server's OPTIONS response. Make sure that the Max-Connections value for the ICAP server is large enough to allow all connections from all Security Engine nodes with which it is integrated. For more information about adjusting the Max-Connections value, see the documentation of your DLP server.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Create an ICAP Server element to represent the DLP server.
   a) Select **Network Elements**.
   b) Right-click **Servers**, then select **New** > **ICAP Server**.
   c) (Optional) To enable TLS for ICAP connections, select **Secure ICAP**, then select a TLS Profile element.
   d) Configure the settings, then click **OK**.
2) Enable ICAP for data protection on the Security Engine.
   a) Select **Engine Configuration**.
   b) Right-click an engine, then select **Edit <element type>**.
   c) Browse to **Add-Ons** > **Data Protection**.
   d) Select **Enable ICAP for data protection**.
   e) Click **Add** next to the **ICAP Servers** field, then add one or more ICAP Server elements.
   f) Click **Save and Refresh** to transfer the changed configuration.

Result

You can now use the DLP scan for data protection in the File Filtering Policy.

Integrate file reputation services and sandboxes

Integrating Forcepoint Network Security Platform with file reputation services and sandboxes improves the malware detection coverage of Forcepoint Network Security Platform when you use file filtering.

Integrate McAfee GTI file reputation with Forcepoint Network Security Platform

Integrating Forcepoint Network Security Platform with McAfee Global Threat Intelligence file reputation services allows access control based on the scan results.

Before you begin

Note

This feature requires a separate license.

Integrating McAfee GTI requires enabling McAfee GTI File Reputation and authorizing the use of the McAfee GTI service.

The McAfee GTI database contains classifications of files. When a file transfer matches a rule in the File Filtering Policy that applies McAfee GTI file reputation, the Security Engine sends a hash of the file to the McAfee GTI cloud. McAfee GTI file reputation compares the hash of the file against the McAfee GTI database.

Only a hash of the file is sent to the McAfee GTI cloud. No other data or telemetry information is sent to the McAfee GTI cloud.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Authorize the use of McAfee GTI in the SMC Client.
   a) Select **Settings** > **Global System Properties**.
   b) On the **Global Options** tab, select **Enable McAfee Global Threat Intelligence (GTI) and Threat Intelligence (TIE) usage**.
2) Enable McAfee GTI file reputation checks.
   a) Select **Engine Configuration**.
   b) Right-click an engine, then select **Edit <element type>**.

c) Browse to **Add-Ons** > **File Reputation**.

d) In the **File Reputation Service** drop-down list, select **McAfee Global Threat Intelligence (GTI)**.

e) Click **Save and Refresh**.

Result

McAfee GTI file reputation scan can now be used for malware detection in the File Filtering Policy.

Forcepoint Advanced Malware Detection and how it works

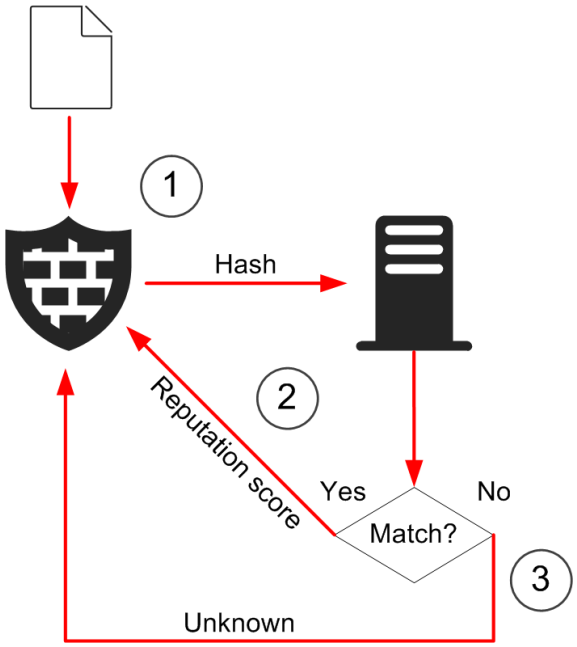
Forcepoint Advanced Malware Detection detects advanced threats by analyzing the behavior of files in the a restricted operating system environment.

Two types of sandbox servers are available for Forcepoint Advanced Malware Detection:

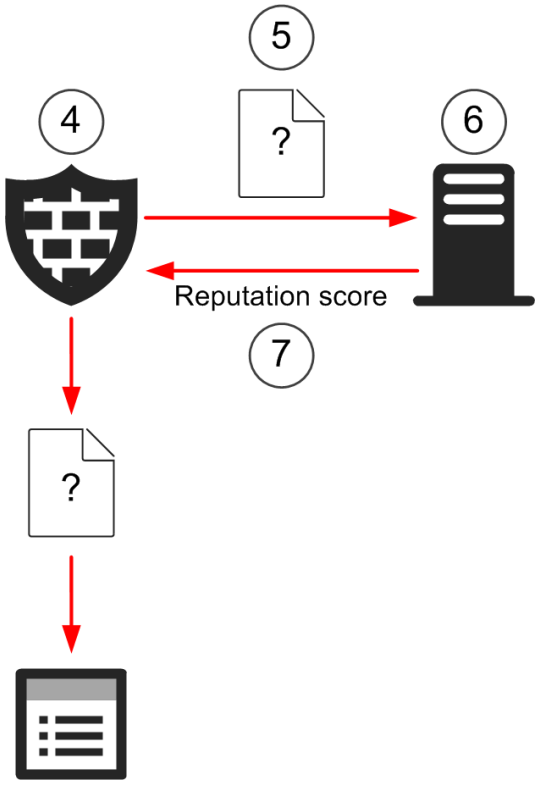
Sandbox servers for Forcepoint Advanced Malware Detection

Type of server	Description
Cloud Sandbox — Forcepoint Advanced Malware Detection	Files are analyzed externally on a cloud sandbox server.
Local Sandbox — Forcepoint Advanced Malware Detection	Files are analyzed locally on a Forcepoint Advanced Malware Detection appliance.

File filtering using Forcepoint Advanced Malware Detection follows this process:



1. When a file transfer matches a rule in the File Filtering Policy that applies the advanced malware sandbox scan, the Security Engine sends a hash of the file to the sandbox server. If the file is a .zip archive, the Security Engine sends a hash of each file in the archive to the sandbox server.
2. If the hash matches a file that has previously been analyzed, the sandbox server returns a file reputation to the Security Engine. The Security Engine allows or blocks the file according to the File Filtering Policy.
3. If the hash of the file does not match a previously analyzed file, the sandbox server returns the **Unknown** file reputation for the file.



4 The Security Engine allows, blocks, or delays the file transfer according to the **Allow After** options for the rule in the File Filtering Policy.

5 If the file has not previously been analyzed, the Security Engine uploads a copy of the unknown file to the sandbox server. If any of the files in a .zip archive have not previously been analyzed, the Security Engine uploads a copy of the whole .zip archive to the sandbox server.

Note

When you use the cloud sandbox for Forcepoint Advanced Malware Detection, unknown executable, document, and archive files, including HTML and JavaScript, are uploaded to the cloud sandbox servers. Do not use the cloud sandbox in countries where transferring files or other data outside of the country is prohibited. Binary files that are uploaded to the cloud sandbox might be stored in the cloud sandbox.

6 The sandbox server analyzes the behavior of the file in a restricted operating system environment. If the file is a .zip archive, the sandbox server analyzes the behavior of each file in the archive.

7 When the analysis is complete, the sandbox server sends an updated file reputation to the Security Engine. The updated file reputation is cached on the Security Engine that requested the scan and stored on the sandbox server.

Note

The updated file reputation does not affect files that have already been allowed or discarded.

If the same file is transferred again, the sandbox server returns the stored file reputation for the file.

Licenses for Forcepoint Advanced Malware Detection

To use Forcepoint Advanced Malware Detection, you must download and install a license that includes the Forcepoint Advanced Malware Detection feature.

For more information about how to obtain licenses for Forcepoint Advanced Malware Detection, see Knowledge Base article 12514.

CAUTION

The license keys and license tokens allow access to confidential analysis reports. Handle the license key and license token securely.

You can enter the license key and license token in two ways:

- Globally in the properties of the Sandbox Service element
  The settings in the Sandbox Service element apply to all Security Engines for which Forcepoint Advanced Malware Detection is enabled.
- Locally in the Engine Editor for individual Security Engines
  The settings for an individual Security Engine override the settings defined in the Sandbox Service element.

Note

If you do not enter the license key in the properties of the Sandbox Service element, you must enter the license key in the Engine Editor for each Security Engine for which Forcepoint Advanced Malware Detection is enabled.

The license for the cloud sandbox includes two pairs of license keys and license tokens:

- A license key and license token for EU data centers
- A license key and license token for US data centers

For the cloud sandbox, the license keys and license tokens for the cloud sandbox define the home data center where the cloud sandbox analyzes and stores files. Make sure that you use the license key and license token for the region that you want to use as your home data center. If the data center that the engine contacts is in a different region than the home data center that is specified in the license, files are forwarded to the home data center.

For the local sandbox, the license includes a license key and license token for use with your Forcepoint Advanced Malware Detection appliance.

Limitations of Forcepoint Advanced Malware Detection

There are some limitations when you use Forcepoint Advanced Malware Detection.

- Each engine communicates separately with the sandbox service. If different engines detect the same file before an analysis result is stored on the sandbox server, the engines might upload the same file more than once. If the hash of the file matches a stored result, the engine does not upload the file again.
- To generate permanent links to sandbox analysis reports in log entries, the SMC makes an API query to the sandbox service. Make sure that traffic from the SMC to the API for the sandbox service is allowed. If necessary, add Access rules that allow traffic from the SMC to the sandbox data centers on TCP port 443. If

the API query to the sandbox service does not succeed, the SMC generates a unique dynamic link for each sandbox analysis report.

If you use Forcepoint Advanced Malware Detection AirGap, the following additional limitations apply:

- Forcepoint Advanced Malware Detection AirGap does not receive automatic periodic updates. You must manually update the reputation database.
- You must manually update licenses for Forcepoint Advanced Malware Detection AirGap.

Rules for Forcepoint Advanced Malware Detection

Rules to allow communication with the sandbox service are automatically generated based on the Security Engine configuration.

The Default File Filtering Policy applies the advanced malware sandbox scan settings defined for the engine to some traffic by default. If you want to exclude specific traffic from these scans, add a more specific rule to a custom File Filtering Policy, or add an Access rule that excludes the traffic from file filtering.

Logging for Forcepoint Advanced Malware Detection

Log entries related to allowed or blocked files contain information about the advanced malware sandbox scan and about the status of the connection to the sandbox service.

The **File Filtering** Log Data Context shows all log entries related to file filtering events.

The following information is available in log entries:

- The **Sandbox Reputation** and **Scan Result** fields show the file reputation provided by the sandbox service.

  Note

  Logging for allowed files is not enabled by default in the File Filtering Policy. You must enable logging of allowed files to see the sandbox reputation of allowed files in the logs.

- The **Scan Report** field shows a link to the sandbox analysis report in log entries related to cloud sandbox or local sandbox scans when file analysis is complete.
- When the Security Engine receives an updated file reputation from the sandbox service, the **File reputation updated** Situation matches and a log entry is created.
- If the file reputation is not trustworthy, information messages can include a description of suspicious or malicious behavior observed in the analysis.
- The information messages in log entries related to allowed or blocked files show errors related to the connection to the sandbox service.

Related tasks

View sandbox analysis reports on page 1006

Using Forcepoint Advanced Malware Detection with Master Security Engines and Virtual Engines

In environments with Master Engines and Virtual Security Engines, you configure the settings for the connection to the sandbox service on the Master Security Engine.

All Virtual Engines that are hosted by the same Master Security Engine use the same configuration. The File Filtering Policy for each Virtual Security Engine defines whether the engine applies Forcepoint Advanced Malware Detection to traffic.

Integrate a Forcepoint Advanced Malware Detection appliance with Forcepoint Network Security Platform

To use a local sandbox for Forcepoint Advanced Malware Detection, integrate a Forcepoint Advanced Malware Detection appliance with Forcepoint Network Security Platform.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In your network environment, install and configure the Forcepoint Advanced Malware Detection appliance.
   For detailed instructions, see the documentation that is included in your appliance delivery.

2) Download and install licenses for Forcepoint Advanced Malware Detection.
   For detailed instructions, see Knowledge Base article 12514.

Next steps

Create a Sandbox Service element to represent the Forcepoint Advanced Malware Detection appliance.

Define Sandbox Service elements

To use Forcepoint Advanced Malware Detection, you must create a Sandbox Service element that defines the settings for the connection to the cloud sandbox or the local sandbox.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Browse to **Other Elements** > **Engine Properties** > **Sandbox Services**.

3) Right-click **Sandbox Services**, then select **New Sandbox Service**.

4) In the **Name** field, enter a unique name.

5) From the **Data Centers** drop-down list, select the data center that the engine contacts to request file reputation scans.
   To use a local sandbox, select the applicable local sandbox option from the **Data Centers** drop-down list.

6) (Local Sandbox only) Enter the host name of the sandbox server in the **Host Name** field.
   The host name is used to automatically generate the default values in the **Server URL** and **Portal URL** fields. You can optionally change the URLs.

7) (Local Sandbox only) Click **Select** next to the **TLS Profile** field, then select a TLS Profile element

8) Click **OK**.

Next steps

Connect Forcepoint Network Security Platform to a sandbox service.

Connect Forcepoint Network Security Platform to a sandbox service

Configure the settings that Forcepoint Network Security Platform uses to connect to the cloud sandbox or the local sandbox.

Before you begin

You must have a Sandbox Service element that defines the settings for the connection to the cloud sandbox or the local sandbox.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click an engine, then select **Edit <element type>**.

3) Browse to **Add-Ons** > **Sandbox**.

4) From the **Sandbox Type** drop-down list, select one of the following options:
   - **Cloud Sandbox - Advanced Malware Detection & Protection**
   - **Local Sandbox - Advanced Malware Detection & Protection**

- **Cloud Sandbox - Advanced Malware Detection**
- **Local Sandbox - Advanced Malware Detection**

Note

- To use a local sandbox, you must have a Forcepoint Advanced Malware Detection appliance.
- The field options change as per the sandbox type that is selected from the **Sandbox Type** drop-down list.
- The License Key and License Token are only used with **Cloud Sandbox - Advanced Malware Detection** and **Local Sandbox - Advanced Malware Detection** services.
- The **Advanced Malware Detection & Protection** cloud sandbox service does not use a License Key or License Token. Instead, the cloud service automatically identifies the caller engine license and the subscription status for the Advanced Malware Detection & Protection sandbox service.
- The **Advanced Malware Detection & Protection** local sandbox service requires an API key from the local AMDP server to authenticate the connections.
- The **Advanced Malware Detection & Protection** cloud sandbox service is only supported on engine version 7.0.2 and higher.
- The **Advanced Malware Detection & Protection** local sandbox service is only supported on engine version 7.1.1 and higher.

5) Click **Select** next to the **Sandbox Service** field, then select a Sandbox Service element.
   - For the cloud sandbox, select the Sandbox Service element that represents the data center that the engine contacts to request file reputation scans.
   - For the local sandbox, select the Sandbox Service element that represents your Forcepoint Advanced Malware Detection appliance.

6) In the **License Key** field, enter or paste the license key for the connection to the sandbox service.

Note

This field is only displayed when the **Cloud Sandbox - Advanced Malware Detection** or **Local Sandbox - Advanced Malware Detection** option is selected from the **Sandbox Type** drop-down list.

7) In the **License Token** field, enter or paste the license token for the connection to the sandbox service.

Note

This field is only displayed when the **Cloud Sandbox - Advanced Malware Detection** or **Local Sandbox - Advanced Malware Detection** option is selected from the **Sandbox Type** drop-down list.

8) (Optional) Click **Add** next to the **HTTP Proxies** field, then select a http proxy element to add to the list.

9) Click **Save and Refresh** to transfer the changed configuration.

Result

You can now use the Forcepoint Advanced Malware Detection scan for malware detection in the File Filtering Policy.

View sandbox analysis reports

In an external portal or in the local portal provided by your Forcepoint Advanced Malware Detection appliance, you can view detailed reports for files that have been analyzed by sandbox services.

When a file has been analyzed, log entries related to cloud sandbox or local sandbox scans include a link to the analysis report in the **Scan Report** field.

Note

You cannot view reports for log entries where the **Scanner Details** cell shows the message "Sandbox Analysis Pending".

By default, the SMC generates permanent links to sandbox analysis reports. Viewing sandbox analysis reports using permanent links does not require separate authentication in the external portal.

Note

To generate permanent links to sandbox analysis reports in log entries, the SMC makes an API query to the sandbox service. Make sure that traffic from the SMC to the API for the sandbox service is allowed. If necessary, add Access rules that allow traffic from the SMC to the sandbox data centers on TCP port 443.

If the API query to the sandbox service does not succeed, the SMC generates a unique dynamic link for each sandbox analysis report. To view sandbox analysis reports using a dynamic link, you must separately authenticate in the external portal.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Logs**.

2) From the Log Data Context drop-down list on the **Query** pane, select **File Filtering**, then click **Apply**.
   Only log entries related to file filtering events are shown.

3) Right-click a log entry that shows **Report Available** in the **Scan Report** cell, then select **Open Sandbox Report in Default Browser**.
   The portal opens in a web browser.

Restrict file types with file filtering

Configure file filtering if you want to restrict the file types that are allowed through the engine, and to apply malware detection to files.

Before you begin

Integrate or configure one or more malware detection methods.

Rules for file filtering are defined in the File Filtering Policy. When a file transfer is detected, the traffic is checked against the File Filtering Policy. The first rule that matches the traffic is applied. If no matching rule is found, the file transfer is allowed.

If you do not want to create a custom File Filtering Policy, you can use one of the following default File Filtering Policy elements:

- **Anti-Malware All** — Applies the anti-malware and file reputation scanning methods that are defined in the Engine Editor to all traffic. Rematches archive content for the following file types: Memory Dumps, Media File, Data File, Text, Empty.
- **Anti-Malware Legacy** — Applies the anti-malware and file reputation scanning methods that are defined in the Engine Editor to all traffic. Rematches archive content for all file types.
- **Default File Filtering** — Applies specific scanning methods and options depending on the file source, file destination, and file type.

Related concepts

Getting started with anti-malware scanning on page 989

Related tasks

Integrate McAfee GTI file reputation with Forcepoint Network Security Platform on page 997

Create File Filtering Policy elements

File Filtering Policy elements contain the rules for file filtering.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click **Policies** and select **New** > **File Filtering Policy**.
3) Configure the settings, then click **OK**.

Result

The File Filtering Policy opens for editing.

Add rules to File Filtering Policy elements

The rules in the File Filtering Policy allow you to define rule-specific options for malware detection.

Before you begin

You must create a File Filtering Policy element.

Rules are read from the top down. Place more specific rules above more general rules that match the same traffic. For example, if there is a rule that allows a file type without scanning above a rule that applies scanning, the matching files are allowed without scanning.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Add a rule in one of the following ways:
   - Right-click the last row in an empty policy, then select **Add Rule**.
   - Right-click the **ID** cell of an existing rule, then select **Add Rule Before** or **Add Rule After**.

2) Drag and drop elements from the **Resources** pane to the **Source** and **Destination** cells, or define source and destination criteria.

   Note

   The **Source** and **Destination** fields are the source and destination of the file transfer, not the source and destination of the connection.

   A client in the internal network downloads a file from a web server on the Internet. The source is the web server that served the file. The destination is the client computer.

3) Drag and drop **File Type Situations** from the **Resources** pane to the **File Type** cell.

4) Right-click the **Action** cell, then select the action.

5) If you selected **Allow After**, select options for malware detection scans.
   The scanning methods are applied in the order in which they are listed. If a file transfer is not blocked by an earlier scan, the action specified for the last scanning method determines whether the file transfer is allowed or blocked. If none of the enabled malware detection scanning methods are available, the action specified for the **Action When No Scanners Are Available** option determines whether the file transfer is allowed or blocked.

6) (Optional) To configure the logging options, double-click the **Logging** cell in the rule.

7) Click **Save**.

Enable file filtering in policy elements

When you enable file filtering in Engine, IPS, Layer 2 Engine, or Layer 2 Interface Policy elements, traffic is checked against the File Filtering Policy.

Note

To enable file filtering in a Layer 2 Interface Policy, you must enable file filtering and select the File Filtering Policy in the Engine Policy.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click a policy and select **Edit <policy type>**.

3) On the **IPv4 Access** or **IPv6 Access** tab, enable file filtering in existing rules or add new rules.
   - Enable **File Filtering** in the action options of individual Access rules.
   - Add rules with the **Continue** action to set defaults for file filtering.

   To use file filtering for the IMAPS, HTTPS, or POP3S protocols, you must use a Service element that enables decryption and inspection, and enable TLS decryption on the Security Engine.

4) (Not available in Layer 2 Interface Policies) On the **Inspection** tab, select the File Filtering Policy.
   If there is no custom File Filtering Policy, the default File Filtering Policy is used.

5) To save and install the policy, click **Save and Install**.

Related concepts

Getting started with policies on page 799

Support for McAfee Advanced Threat Defense

McAfee Advanced Threat Defense is no longer supported in Security Engine version 6.4.0 and higher. We recommend that you use Forcepoint Advanced Malware Detection instead.

Chapter 62

Integrating Forcepoint One Endpoint with Forcepoint Network Security Platform

Contents

- Forcepoint One Endpoint and how it works on page 1011
- Create ECA Configuration elements on page 1014
- Enable Forcepoint Endpoint Context Agent (ECA) on the Security Engine on page 1015
- Define Endpoint Application elements on page 1015
- Create Endpoint Settings elements on page 1016
- Use endpoint information in Access rules on page 1017
- Enable logging of endpoint information on page 1017

If you have installed Forcepoint One Endpoint clients on the endpoints in your network, you can collect information about endpoint clients, and use the information for access control in the SMC.

Forcepoint One Endpoint and how it works

Integrating Forcepoint One Endpoint enables you to collect per-connection user and application information about Windows endpoint clients that connect through an Security Engine managed by the SMC.

To use Forcepoint One Endpoint, the Forcepoint One Endpoint client must be installed on the endpoints. For more information about Forcepoint One Endpoint clients, see the *Installation and Deployment Guide for Forcepoint One Endpoint* .

The endpoints send metadata to the Security Engine, and you can use the information as criteria for access control in policies. This information about the endpoints can also be viewed in log data and used in Report elements.

On the home page for an Security Engine, you can see the number of endpoint clients that are connected and sending information. You can also use the drill-down menu to see which users are connected.

Forcepoint One Endpoint is supported on Engines, Layer 2 Engines, IPS engines, and on Virtual Engines. The Security Engine license includes support for Forcepoint One Endpoint integration.

You cannot use Forcepoint One Endpoint if there is a NAT device between the Security Engines and the endpoints.

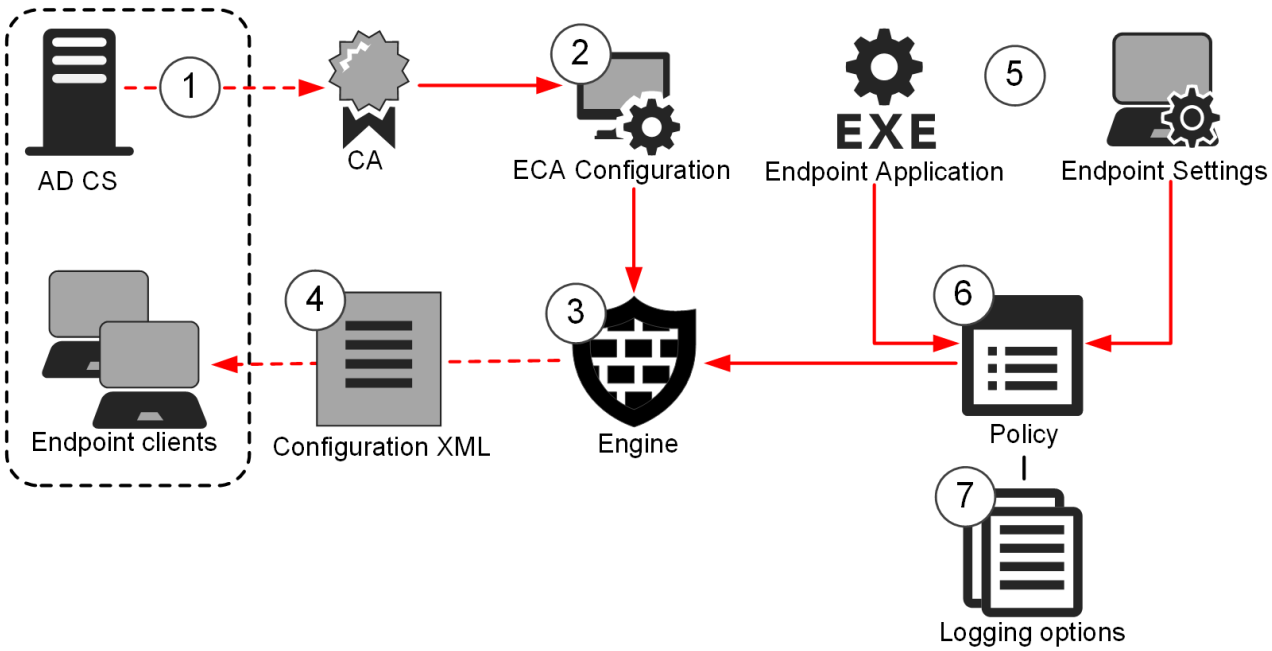
Use cases

An example use case is a point of sale (PoS) terminal. For example, you can:

- Allow a certain browser version to access the corporate intranet, only if the local engine on the endpoint is enabled and the operating system was updated within the past 30 days.
- Allow the PoS application to access corporate servers
- Allow the Windows Update service
- Block all other applications

Forcepoint One Endpoint integration configuration overview

The integration of Forcepoint One Endpoint consists of several general steps.



1. Create or use a certificate authority from the domain where the endpoint clients are located, then import the CA to the SMC as a Trusted Certificate Authority element. For more information, see Knowledge Base article 14099.
2. In the SMC Client, create an ECA Configuration element that uses the created CA.
3. Enable Forcepoint Endpoint Context Agent (ECA) on the engine, and use the ECA Configuration element that you created.
4. Export the configuration XML file, and use the file when installing the Forcepoint One Endpoint client on the endpoints. The file contains the details of all the Security Engines that use the same ECA Configuration element. If additional Security Engines are added to the configuration, the updated XML configuration file is automatically sent to the endpoint clients when they connect to an Security Engine.
5. (Optional) To use endpoint client information for access control, define Endpoint Application and Endpoint Settings elements.
6. (Optional) In the policy of the Security Engine, configure Access rules using Endpoint Application and Endpoint Settings elements as matching criteria.

   Access rules to allow communication between Forcepoint One Endpoint components are automatically generated.
7. (Optional) To view endpoint information in log data and reports, enable endpoint information logging in the Access rules.

Evaluate Forcepoint One Endpoint

If you are interested in testing how Forcepoint One Endpoint works in your environment, see Knowledge Base article 16193 for information about how to easily deploy Forcepoint One Endpoint to a limited set of users for evaluation purposes.

Create ECA Configuration elements

ECA Configuration elements contain the Trusted Certificate Authority element used to secure communication between the Security Engine and the Forcepoint One Endpoint clients.

Before you begin

Create or use a certificate authority from the domain where the endpoint clients are located, then import the CA to the SMC as a Trusted Certificate Authority element. For more information, see Knowledge Base article 14099.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Browse to **Other Elements** > **Engine Properties** > **ECA Configurations**.
3) Right-click **ECA Configurations**, then select **New ECA Configuration**.
4) Configure the settings, then click **OK**.

Note

If **Advertise Engine's Contact Address to ECA Clients** is selected, the Security Engine can send ICMP or ICMPv6 discovery messages to endpoint clients that are not aware that the contact address for the Security Engine has changed or that the Security Engine can receive Forcepoint One Endpoint metadata. The ICMP message is Destination Unreachable, and the type is Communication Administratively Prohibited.

Next steps

Enable Forcepoint Endpoint Context Agent (ECA) on the Security Engines, and select the ECA Configuration element that you created.

Related tasks

Create Trusted Certificate Authority elements on page 155

Enable Forcepoint Endpoint Context Agent (ECA) on the Security Engine

Enable Forcepoint Endpoint Context Agent (ECA) on the Security Engines on which you want to receive Forcepoint One Endpoint client information.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click a Security Engine, then select **Edit <element type>**.
3) Browse to **Add-Ons** > **Endpoint Integration**.
4) From the **Endpoint Service** drop-down menu, select **Forcepoint Endpoint Context Agent (ECA)**, then configure the settings.
5) Click **Save**.
6) Click **Export Configuration for Endpoint Clients**, then save the XML file that contains the configuration.
   The details of all the Security Engines that use the same ECA Configuration element are included in the exported XML file. You must have finished configuring all the Security Engines before you export the file.

Next steps

Use the exported XML configuration file when installing the Forcepoint One Endpoint clients on the endpoints. For more information, see the *Installation and Deployment Guide for Forcepoint One Endpoint* .

Define Endpoint Application elements

Use Endpoint Application elements for access control in the SMC.

Endpoint Application elements contain information about the applications used by Forcepoint One Endpoint clients. The elements are delivered in dynamic update packages, but you can also create elements manually. If you upgraded the SMC, you might have legacy Executable elements. We recommend using Endpoint Application elements instead.

Endpoint Application elements are categorized by usage type with Tag elements. You can use a Tag in a policy to allow or block the use of a category of applications, such as instant messaging applications.

Create Endpoint Application elements

Create Endpoint Application elements by defining the details of the applications used on the endpoints.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Browse to **Other Elements** > **Endpoint Information** > **All Endpoint Applications**.
3) Right-click **All Endpoint Applications**, then select **New Endpoint Application**.
4) Configure the settings, then click **OK**.

Create Endpoint Settings elements

Endpoint Settings elements define what information to collect about the endpoint clients.

For example, you can check:

- The client operating system and whether it has received security updates
- Whether anti-virus software is enabled on the client
- The status of the local engine on the client

You can use Endpoint Settings elements for access control based on endpoint information. Depending on the type of information that you are checking for, you might want to allow or discard traffic from clients where the criteria matches.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Browse to **Other Elements** > **Endpoint Information** > **Endpoint Settings**.
3) Right-click **Endpoint Settings**, then select **New Endpoint Settings**.
4) Configure the settings, then click **OK**.

Use endpoint information in Access rules

In Access rules, you can configure the engine to allow or discard connections based on endpoint client information.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Browse to **Policies** > **<policy type>**.
3) Right-click the policy you want to edit, then select **Edit <policy type>**.
4) Add a rule.
5) In the navigation pane, select **Endpoint Information**, browse to the Endpoint Application or Endpoint Settings elements, then drag them to the **Source** cell.

   Tip

   To select all Endpoint Application elements that have been categorized using the same tag, drag the appropriate Tag element.

6) Define the other options for the rule as needed.
7) Save and install the policy to transfer the changes to the Security Engines.

Enable logging of endpoint information

Logging of endpoint information allows you to view endpoint client information and users in log data and Report elements.

By default, endpoint information is logged when it is used for matching in the Access rules. You can optionally set endpoint information to be logged whenever it is received.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Open the policy for editing.
2) In an Access rule where Endpoint Application or Endpoint Settings elements are used, double-click the **Logging** cell.

3) Select **Override Recording Settings Inherited from Continue Rule(s)**.

4) In the **Log User Information** and **Log Endpoint Information** drop-down lists, select the logging options.



Note

To log user information, both **Log User Information** and **Log Endpoint Information** must be set to **Enforced**.

5) Click **OK**.

6) Save and install the policy to transfer the changes to the engines.

Chapter 63

Filtering URLs

Contents

- URL filtering and how it works on page 1019
- Enable ThreatSeeker on page 1022
- Use an HTTP proxy to connect to the ThreatSeeker Intelligence Cloud server on page 1023
- Add Access rules for category-based URL filtering on page 1024
- Add URL List Application elements to manually block or allow URLs on page 1025
- Add Access rules for custom URL List Applications on page 1027
- Examples of URL filtering on page 1028

URL filtering allows you to filter URLs based on categories of content or lists of individual URLs.

URL filtering and how it works

URL filtering compares the URLs that end-users attempt to access to URL categories or lists of URLs.

You can use URL filtering to prevent users from accessing websites that provide content that is objectionable, potentially harmful, or not work-related. This kind of content filtering can increase network security and enforce an organization's policy on acceptable use of resources.

In URL filtering, the engines compare the URLs in HTTP and HTTPS requests against URL categories or lists of URLs. There are two ways to define the URLs:

- You can use URL Category and URL Category Group elements to filter URLs based on URL categorization.
- You can use URL List Application elements to filter specific URLs.

You can use both methods together. You can also define allowed URLs manually if a URL that you want to allow is included in a category of URLs that you otherwise want to block.

The URL categorizations are provided by the external Forcepoint™ ThreatSeeker® Intelligence Cloud service. ThreatSeeker Intelligence Cloud (ThreatSeeker) provides categories for malicious websites and several categories for different types of non-malicious content you might want to filter or log. The Security Engine sends categorization requests using HTTPS to ThreatSeeker.

URL Category Group elements contain several related URL Categories. When you use URL Category Group elements in the Access rules, the rule matches if any of the URL Categories included in the URL Category Group match.

The Security Engine can use the server name indication (SNI) in HTTPS traffic for URL categorization without decrypting the HTTPS connection. When a web browser contacts a server to request a page using HTTPS, the browser sends the server name in an unencrypted SNI field. However, the requested URL is not known when HTTPS connections are not decrypted.

Note

When traffic is allowed based on URL categorization, the domain name in the SNI is validated by checking the server certificate. If the certificate is not valid for the domain name in the SNI, the certificate is not within the certificate validity period, or the certificate is not issued by a CA that is trusted by the Security Engine, the initial match based on SNI is invalidated.

You can use the Private Data application usage tag in Access rules to disallow decryption for traffic in predefined categories. You can also create Access rules to disallow decryption for specific categories. For example, you can create rules that prevent traffic to online banking services from being decrypted.

Note

If the end user's browser does not use SNI and the traffic does not match rules for category-based URL filtering, the traffic might be decrypted.

You can configure the engine to respond in various ways when a match is found. For example, you can log the matches or block the traffic. If you decide to block traffic, the engine can notify end users with a message in their browsers. You can define customized User Response elements for URL filtering matches, such as a custom HTML page that is displayed to the end user when a connection is blocked.

If the engine detects that it cannot connect to ThreatSeeker, all URLs match the **Data Provider Error** URL Category. You can optionally add Access rules to discard all traffic that cannot be categorized.

Limitations

- Category-based URL filtering using the ThreatSeeker Intelligence Cloud service is a separately licensed feature.
- Category-based URL filtering is based on the categorizations of the external service, so it is not possible to manually add or directly edit the URL filtering categories. Add URL List Applications to Access rules if you want to create exceptions to the category-based URL filtering.

Note

The engine can only use one kind of URL categorization at a time. If you used legacy URL filtering with Forcepoint Network Security Platform version 6.0 or lower, you must migrate the engine to use ThreatSeeker URL filtering when you upgrade the engine to version 6.1 or higher. See Knowledge Base article 16868.

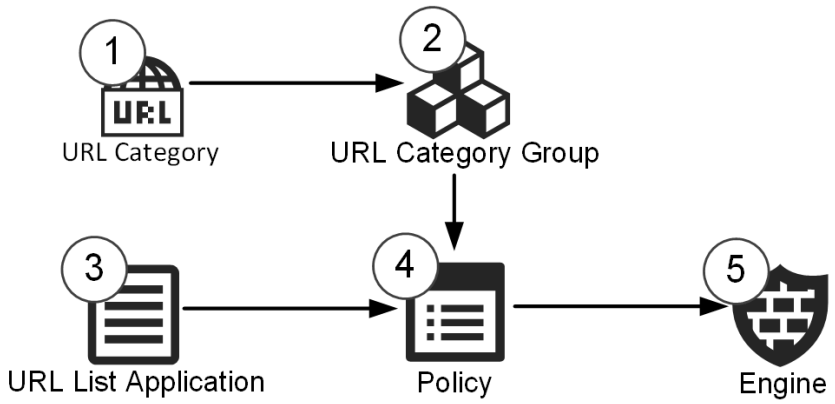
Related concepts

Getting started with forwarding traffic on page 1087

URL filtering configuration overview

The URL filtering feature is configured using category-based URL filtering and custom URL List Applications.

Elements in the configuration



1. The URL Category is referenced in the URL Category Group
2. The URL Category Group is referenced in the Access rules in a policy
3. The URL List Application is referenced in the Access rules in a policy
4. The policy contains the Access rules that reference the URL elements.
5. The policy is installed on the Security Engine

The Access rules define how URL Categories and URL List Applications are matched to traffic and what reaction a match triggers. URL List Applications can override category-based URL filtering to allow some URLs manually.

Because an external service defines the URLs that are included in category-based URL filtering, it is not possible to add or edit categories. The category names for category-based URL filtering are updated through dynamic update packages.

1) (Category-based URL filtering) Enable ThreatSeeker for the engine.

2) (Category-based URL filtering) Make sure that the engine and the network are set up so that the engine can contact the ThreatSeeker Intelligence Cloud servers to request URL categorizations.
   - Make sure that DNS servers are defined for the engine in the **General** branch of the Engine Editor.
   - Make sure that the engine can access the DNS server (UDP port 53, Service element "DNS (UDP)"). The connections are automatically allowed from the engine on which the policy is installed, but not from other components.
   - (Optional) Configure an HTTP proxy for the connection to the ThreatSeeker Intelligence Cloud server in the **Add-Ons** > **ThreatSeeker** branch of the Engine Editor.

3) (Optional) Create User Responses to notify users about matches that are found.

4) (Optional when using category-based URL filtering) Create URL List Applications to block or allow URLs manually.

5) Add Access rules to define how URL filtering is applied.
   - Create IPv4 or IPv6 Access rules to define how URL Categories are matched to traffic and what kind of reaction a match triggers.

- Create IPv4 or IPv6 Access rules to define how URL List Applications are matched to traffic and what kind of reaction a match triggers.
  For example, you can add rules for URL List Applications to create exceptions to rules for category-based URL filtering.

Related concepts

User Response elements and how they work on page 969

Related tasks

Use an HTTP proxy to connect to the ThreatSeeker Intelligence Cloud server on page 1023
Add Access rules for category-based URL filtering on page 1024
Add URL List Application elements to manually block or allow URLs on page 1025
Add Access rules for custom URL List Applications on page 1027

Default elements for URL filtering

There are default elements for the categories you can use in category-based URL filtering and for creating custom lists of URLs.

- URL Category elements represent the categories for category-based URL filtering.
- URL Category Group elements contain several related URL Categories.
  You can find these elements under **Network Applications** > **By Type** > **URL Category Group** in the element tree.

  Note

  The URL Categories and URL Category Groups are predefined. It is not possible to add or edit URL Categories or URL Category Groups.
- URL List Application elements allow you to manually define lists of URLs. You can find these elements under **Network Applications** > **By Type** > **URL List** in the element tree.

Enable ThreatSeeker

To start using ThreatSeeker categories for URL filtering, enable ThreatSeeker for the engine.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click an engine, then select **Edit <element type>**.
3) In the navigation pane on the left, browse to **Add-Ons** > **ThreatSeeker**.

4) Select **Enable**.

5) Select from the following:
   - To validate the changes, select ⋮ **More actions** > **Validate**.
   - To validate and save the changes, click **Save**.
   - To validate and save the changes and refresh the security policy on the engine, click **Save and Refresh**.



Note

Validation issues are displayed in the **Issues** pane. Double-click an issue to return to the section in which the issue can be fixed.

Use an HTTP proxy to connect to the ThreatSeeker Intelligence Cloud server

You can optionally configure an HTTP proxy for the engine's connection to the external ThreatSeeker Intelligence Cloud server.

When you use an HTTP proxy, categorization requests are sent through the proxy instead of the engine accessing the external network directly.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Add one or more HTTP Proxy elements.
   a) Select **Engine Configuration**.
   b) Browse to **Other Elements** > **Engine Properties**.
   c) Right-click **HTTP Proxies**, then select **New HTTP Proxy**.
   d) Configure the settings.
   e) Click **OK**.
2) Browse to **Engine** > **Engines**.
3) Right-click an engine, then select **Edit <element type>**.
4) In the navigation pane on the left, browse to **Add-Ons** > **ThreatSeeker**.
5) Next to the **HTTP Proxies** list, click **Add**.
6) Select one or more HTTP Proxy elements, then click **Select**.

7) Click **Save**.

Add Access rules for category-based URL filtering

Use category-based URL filtering in IPv4 or IPv6 Access rules in a Engine Policy, IPS Policy, Layer 2 Engine Policy, or Layer 2 Interface Policy to define which traffic is logged or blocked when a URL match is found.

Before you begin

Category-based URL filtering requires that the engine is licensed to use the ThreatSeeker categorization service. You must also define DNS server addresses in the Security Engine elements so that the engines can send categorization requests to ThreatSeeker.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Browse to **Policies**.

3) Right-click a policy and select **Edit <policy type>**.

4) On the **IPv4 Access** or **IPv6 Access** tab, add a rule.

Tip

As a general guideline, we recommend placing rules that allow traffic above rules that block traffic.

5) Drag and drop elements from the **Resources** pane on the left to the **Source** and **Destination** cells.

6) Add URL Category or URL Category Group elements for category-based URL filtering in one of the following ways:

- Drag and drop one or more elements from the **Resources** pane on the left to the **Service** cell.
- Add elements to the Service definition.

When you use URL Category Group elements in a rule, the rule matches if any of the URL Categories in the group match.

7) In the **Action** cell, select the action depending on the purpose of the rule.

- To allow matching traffic, select **Allow**.
- To block matching traffic, select **Discard**.

8) (Optional) In the **Logging** cell, configure the logging options for the rule.

9) Click **Save and Install**.

Next steps

If you want to make exceptions to the category-based URL filtering, add rules to manually block or allow URL List Applications.

Related concepts
Getting started with Access rules on page 831
User Response elements and how they work on page 969

Related tasks
Define Source, Destination, and Service criteria in rules on page 891

Add URL List Application elements to manually block or allow URLs

URL List Application elements allow you to define custom lists of URLs to block or allow web traffic.

Important

- When you block HTTPS traffic with URL lists:
  - If the URL list only has simple domains, then it works without decryption. For example, example.com.
  - If the URL list has URLs with paths, then it requires TLS decryption to work. For example, example.com/path.
- When you allow HTTPS traffic with URL lists:
  - If the URL list only has simple domains, then it works. For example, example.com.
  - If the URL list has URLs with paths, then it does not work. For example, example.com/path.

Note

You cannot currently all ow HTTPS URLs with paths by using the URL List Application.

For information about the workaround to allow HTTPS traffic with paths, refer to the *Using a URL List on NGFW to Allow HTTPS Connections to the Specific URL Path Does Not Work* Knowledge Base Article.

The action that you select in the Access rules determines whether the URLs in the URL List Application are blocked or allowed.

When you use URL List Applications in combination with category-based URL filtering, you can allow individual URLs that are included in a blocked category. Using URL List Applications to allow URLs only affects other URL-based filtering. It does not exclude the traffic from other inspection checks. Traffic to allowed URLs might still be terminated if deep inspection is enabled and the traffic matches Situations in the Inspection Policy.

There is no limit on the number of URL List Applications that you can create or on the number of URLs that you can add to each URL List Application. You can enter URLs as whole URLs or partial URLs. Partial URLs must

end with a slash (/). The URLs in the list can match all URLs in a domain, all URLs in a specified path, or exact URLs.

When you add an exact URL, only the specified URL matches. Other URLs in the same domain or path do not match. For example, if you add the exact URL `www.example.com/index.html`, connections to `www.example.com/main.html` do not match.

When you add domains or paths, connections might match more than one URL. For example:

- If you add an exact URL and a path that is part of the exact URL, both URLs match if the engine detects a request to the exact URL.
  For example, if you add the exact URL `www.example.com/path/index.html` and the path `example.com/path/`, both URLs match if the engine detects a request to `http://www.example.com/path/index.html`.
- If you add a domain name followed by a slash, the URL List also matches connections to subdomains of the domain name.
  For example, if you enter `example.com/`, the URL List also matches connections to `login.example.com` and `www.example.com`.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Browse to **Other Elements** > **Network Applications** > **By Type** > **URL List**.

3) Right-click **URL List**, then select **New** > **URL List Application**.

4) Add one or more URLs.
   - To add the first URL, double-click the empty row.
   - To add a URL above an existing URL, right-click the URL and select **Add**.
   - To add a URL below an existing URL, right-click the empty space in the list and select **Add**.

   Tip

   You can also paste a URL or a list of URLs into the field. Each line of text is automatically added on a separate line.

5) Enter the URL without the http:// protocol.
   - To match all subdomains of a domain, enter the domain name followed by a slash (/).
     Example — `example.com/`
   - To match all URLs in the specified path, enter the path followed by a slash (/).
     Example — `example.com/path/`
   - To match an exact URL, enter the URL.
     Example — `www.example.com/index.html`

6) (Optional) On the **Tags** tab, select the Tags that you want to use with this URL List Application.

7) Click **OK**.

Next steps

To use the URL List Application element for URL filtering, add it to an Access rule.

Add Access rules for custom URL List Applications

Use URL List Application elements in the Access rules to block or allow individual URLs.

Before you begin

If you want to add a User Response element to the rule, you must have a custom User Response element.

You can use URL List Applications in Engine, IPS, Layer 2 Engine, and Layer 2 Interface Policies.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Browse to **Policies**.

3) Right-click a policy and select **Edit <policy type>**.

4) On the **IPv4 Access** or **IPv6 Access** tab, add a rule.

Tip

As a general guideline, we recommend placing rules that allow traffic above rules that block traffic.

5) Drag and drop elements from the **Resources** pane on the left to the **Source** and **Destination** cells.

6) Add URL List Application elements in one of the following ways:
   - Drag and drop one or more elements from the **Resources** pane on the left to the **Service** cell.
   - Add elements to the Service definition.

7) In the **Action** cell, select the action depending on the purpose of the rule.
   - To allow matching traffic, select **Allow**.
   - To block matching traffic, select **Discard**.

8) (Optional, **Discard** action only) To display warnings or notes in the users' browsers when URL filtering prevents access, add a User Response element to the rule.
   a) Right-click the **Action** cell and select **Edit Options**.
   b) On the **Response** tab, select **Override Settings Inherited From Continue Rule(s)**.
   c) Next to the **User Response** field, click **Select**.

d) Select your custom User Response element and click **Select**.

e) Click **OK**.

9) (Optional) In the **Logging** cell, configure the logging options for the rule.

10) Click **Save and Install**.

Related concepts
Getting started with Access rules on page 831
User Response elements and how they work on page 969

Related tasks
Define Source, Destination, and Service criteria in rules on page 891

Examples of URL filtering

These examples show some common uses for URL filtering, and general steps for how each example is configured.

Example: Allowing a blocked URL

In this example, a company is using category-based URL filtering and wants to make an exception to categorized URLs that are blocked by default.

A social media website that users in the marketing department need to access is blocked by category-based URL filtering. To make an exception for users in the marketing department, the administrators:

1) Create a URL List Application element and enter the URL of the social media website.

2) Add the following type of Access rule above the rules that apply category-based URL filtering.

Rule for allowing URLs

Source	Destination	Service	Action
Marketing department users' workstations	ANY	Custom URL List Application element	Allow

3) Save and install the policy.

Example: Discarding connections when URL categorization is not available

In this example, a company is using category-based URL filtering and wants to block traffic that cannot be categorized.

When URL categorization is not available, it might be possible for users to access URLs that are usually blocked. To prevent users from accessing potentially harmful URLs, the administrators want to discard connections to external servers that cannot be categorized.

Note

Depending on the matching criteria in the Access rules, discarding connections when URL categorization is not available might block all HTTP and HTTPS traffic.

The administrators:

1) Add the following type of Access rule after the rules that apply category-based URL filtering.

Rule to discard connections that cannot be categorized

Source	Destination	Service	Action
Internal Network element	Not Internal expression	**Data Provider Error** URL Category	Discard

Tip

If the engine detects that it cannot connect to ThreatSeeker, all URLs match the **Data Provider Error** URL Category.

2) Save and install the policy.

Example: Preventing decryption of private connections

In this example, a company uses category-based URL filtering to prevent connections that could contain users' personal information from being decrypted.

To identify connections that could contain users' personal information and prevent them from being decrypted, the administrators:

1) Add the following type of Access rule.

Rule to prevent decryption of private communications

Source	Destination	Service	Action
ANY	ANY	**Private Data** Application Usage Tag	Allow

2) Edit the action options and select **Disallowed** for the **Decryption** option.

3) Save and install the policy.

Chapter 64

Protocol Agents on Security Engines

Contents

- Protocol Agents overview on page 1031
- Configuring Protocol Agents on page 1033
- Using Protocol Agents on page 1034
- Examples of Protocol Agents on page 1040

Protocol elements of the **Protocol Agent** type are special modules for some protocols and services that require advanced processing. Protocol Agents can enforce policies on the application layer.

Protocol Agents overview

Protocol Agents are software modules for advanced processing of protocols that require special handling on the Engine, Layer 2 Engine, or the IPS engine.

Special handling might be required due to the complexity of the protocols, address information in the data payload, related connections, or other consideration. Protocol elements also associate the traffic with a certain protocol for inspection against the Inspection Policy.

Protocol Agents on Engines can:

- Validate application-level protocol use (for example, FTP command syntax).
- Open related connections when required (for example, FTP data connections).
- Modify application data when required (for example, NAT in H.323 data payload).
- Redirect FTP, HTTP, HTTPS, and SMTP connections to proxy services.

Protocol Agents on Layer 2 Engines and IPS engines can:

- Validate application-level protocol use (for example, FTP command syntax).
- Open related connections when required (for example, FTP data connections).

Some protocols require the use of the correct Protocol Agent to pass inspection by the Engine, Layer 2 Engine, or the IPS engine when traffic is handled using stateful inspection.

Using Protocol Agents for connection handling

When related new connections are opened based on information exchanged in an initial connection, Protocol Agents might be needed.

Protocol Agents are provided to handle the following protocols:

- FTP with related active and passive data connections

- H.323 conferencing protocol communications
- Microsoft RPC (MSRPC) for Microsoft Exchange and Outlook communications
- NetBIOS for the Windows NetBIOS datagram services
- Oracle TNS protocol communications
- Remote Shell protocol communications
- Session Initiation Protocol (SIP) communications
- Sun RPC Portmapper communications
- TFTP file transfers

File transfer protocol (FTP)

FTP uses two related connections: a control connection and a separately established data connection. If the control connection is allowed without the Protocol Agent, the engine does not recognize that the data connection is part of an existing connection. The connection is handles as a new connection, which usually leads to the data transfer failing.

Using Protocol Agents for protocol validation

Protocol Agents can be used to validate communications against standards of specific protocols. Exactly how this works depends on the protocol in question.

Here are a few examples:

- The FTP Protocol Agent can be set to strictly limit the allowed commands within the control connection to standard commands as listed in the FTP specifications. If other commands are sent in the control connection, the connection is dropped.
- The Oracle Protocol Agent can control the size of the Oracle TNS packets, or the location of the Listener service regarding the database services.
- The SSH Protocol Agent can ensure that the SSH handshake is performed at the beginning of an SSH connection.

Using Protocol Agents for NAT in application data

Protocol Agents on Engines can be used to assist with network address translation (NAT) in the application data.

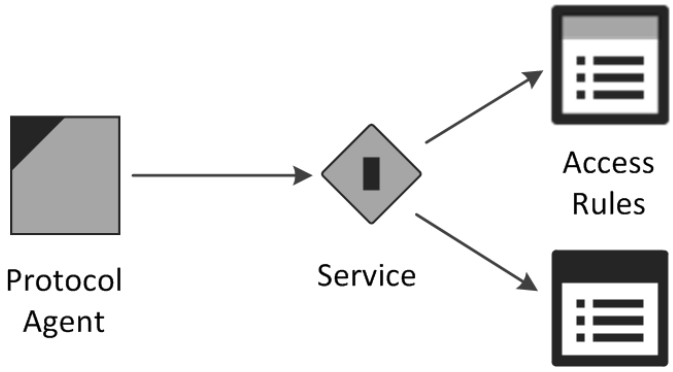
For example, the H.323 conferencing protocol includes the source and destination address information in the data payload of the packets. In ‘normal’ traffic, all IP address information relevant to the communications is in reserved spaces in the packet headers.

The H323 Protocol Agent can examine the data payload and change the addresses according to the network address translation as needed. Therefore, when the source address is included in the protocol data, the source address is also translated in the data payload. The receiving system then responds to the proper address.

Configuring Protocol Agents

Protocol Agents are represented in the SMC Client by **Protocol** elements that have *Protocol Agent* as their type. Other Protocol elements are of the type *Protocol Tag*.

Using Protocol Agents



You do not add Protocol Agent elements directly to policies. You select Protocol Agents in custom Service elements that you create. The custom Service elements are used in Access rules. Whenever traffic matches a rule that contains a Service element with an associated Protocol Agent, the Protocol Agent is automatically activated.

All Protocol Agents are default elements, and you cannot change them or add any new ones. There are also default Service elements for most supported protocols that you can use to activate the Protocol Agents. However, some Protocol Agents have parameters and options you can set by creating customized Services.

Creating custom Service elements with Protocol Agents

In addition to default Service elements that are associated with specific Protocol Agents, you can create custom Services.

There are default Service elements that link to Protocol Agents. These default Services can be used directly in the Access rules. However, the default Services do not allow you to change the default parameters of Protocol Agents that have configurable parameters.

If you want to change the way a Protocol Agent behaves, create a custom Service and attach the correct Protocol Agent to that Service. The Service element contains the identifying information, such as a port number, that determines which traffic the Service matches. In most cases, the identifying information makes sure that the Protocol Agent is not applied to the wrong type of traffic.

Setting parameters for Protocol Agents

When you create a custom Service that uses a Protocol Agent with configurable parameters, you specify the parameters in the properties of the Service.

Related concepts
Using Protocol Agents on page 1034

Inserting the Service with Protocol Agent in Engine Access rules

Whether you create a custom Service or use a predefined Services that have a Protocol Agent attached to them, define the traffic in the Access rules in your policies.

A Protocol Agent can be set either on a rule-by-rule basis or you can create a rule with Continue as the Action. When there is a Continue rule, rules further down in the rule table that match traffic (same source and destination) use the Protocol Agent defined in the Continue rule.

With Protocol Agents, the Continue rule affects only rules where the Service cell is set to ANY. More specific Service definitions override the Continue rule, as all Service elements specify that either some particular Protocol Agent or no Protocol Agent is used.

Some protocols might require a Protocol Agent if the Connection Tracking option is enabled for the rule. Those protocols might not be allowed by a rule that has ANY as its Service unless a Protocol Agent is configured using a previous matching Continue rule. The Firewall Template Policy contains a Continue rule that sets a Protocol Agent to be used with Services in the Service Group called Default Services with Agents.

Protocol Agents validate traffic against the specifics of a particular protocol, so make sure that a Service is not applied to traffic that does not use that protocol. Also, Protocol Agents are designed for particular types of uses, so they might not always be appropriate even if the protocol matches.

Using Protocol Agents

There are Protocol Agents for many different protocols.

Most of the available Protocol Agents have configurable parameters that they add to the Services that use them.

If there are no configurable parameters for a Protocol Agent, there are no options, but there can still be a control for turning the Protocol Agent on or off in the Service.

FTP Protocol Agent

One of the most common ways to transfer files across networks is using FTP. An FTP session starts with a control connection (by default, TCP port 21), and the data connection continues using a dynamically allocated port.

The Protocol Agent tracks the actual ports used so that ports can be opened only as needed for specific connections. This way, the whole range of possible dynamic ports does not need to be allowed in the policy.

The FTP Protocol Agent inspects protocol validity. There are two selectable levels of inspection: *loose* (default) and *strict*.

- In loose mode, the Protocol Agent tries to identify information for opening the data connection. Loose mode is needed with some non-standard FTP applications.
- With strict mode, protocol integrity can be enforced. All connections with commands that do not comply with the RFC 959 FTP standard are dropped.

The FTP Protocol Agent can change payload data, if necessary. The change can be required to handle NAT correctly.

The FTP Protocol Agent can also redirect traffic to a proxy service.

This Protocol Agent has parameters you can set in the Service properties.

Related concepts
Getting started with forwarding traffic on page 1087

GRE Protocol Agent

The Generic Routing Encapsulation (GRE) protocol is a tunneling protocol that allows the encapsulation of network layer packets inside IP tunneling packets.

The GRE Protocol Agent provides protocol inspection for tunneled GRE traffic. This Protocol Agent specifies rematching parameters for GRE-encapsulated packets, and defines which traffic is tunneled. This Protocol Agent has parameters that you can set in the Service properties.

GTP Protocol Agent

The GTP Protocol Agent provides protocol inspection for GTP traffic.

The GPRS Tunneling Protocol (GTP) is used to carry GPRS (general packet radio service) packets in GSM, UMTS, and LTE networks. There are no configurable parameters for this Protocol Agent.

H323 Protocol Agent

H.323 defines a set of protocols as well as the components and procedures for real-time multimedia communication.

H.323 consists of a series of different types of standards related to video and audio services, real-time transport, control channels, and security.

On Engines, this Protocol Agent has parameters you can set in the Service properties:

- H.323 can open several related connections, which places demands on access control and NAT. The H323 Protocol Agent's task is to track the related connections that are opened within the same session. Particularly, if you want the Engine to apply NAT to H.323 connections, you must make sure that the connections use this Protocol Agent.
- The H323 Protocol Agent examines Call Signaling Channel (Q.931/H.225.0) connections and allows the related Control Channel (H.245) connection to open. It also examines the H.245 connection and allows further related connections (RTP and RTCP) to open, based on the port negotiations on the parent connection.

- When NAT is applied to Q.931 connections, the Protocol Agent performs the same NAT to the related H.245 connection and changes the payload data of the parent connection. The same NAT operation is performed also on the opened RTP and RTCP connections.

There are no configurable parameters for this Protocol Agent on IPS engines or Layer 2 Engines.

HTTP Protocol Agent

The HTTP Protocol Agent can be used on Engines to redirect traffic to a proxy service, and in all engine roles to log the URLs from HTTP requests.

This Protocol Agent has parameters that you can set in the Service properties.

HTTPS Protocol Agent

The HTTPS Protocol Agent can be used for identifying encrypted HTTPS traffic for decryption and inspection in the Access rules, for identifying encrypted HTTPS traffic for inspection in the Inspection Policy, and for redirecting connections to a proxy service.

This Protocol Agent has parameters that you can set in the Service properties.

IPv4 Encapsulation Protocol Agent

The IPv4 Encapsulation Protocol Agent provides protocol inspection for tunneled IPv4 traffic on IPS engines and Layer 2 Engines.

This Protocol Agent specifies rematching parameters for IPv4 packets encapsulated in IPv6 packets. This Protocol Agent has parameters that you can set in the Service properties. This Protocol Agent is not available on Engines.

IPv6 Encapsulation Protocol Agent

The IPv6 Encapsulation Protocol Agent provides protocol inspection for tunneled IPv6 traffic on IPS engines and Layer 2 Engines.

This Protocol Agent specifies rematching parameters for IPv6 packets encapsulated in IPv4 packets. This Protocol Agent has parameters that you can set in the Service properties. This Protocol Agent is not available on Engines.

MGCP Protocol Agent

The MGCP (Media Gateway Control Protocol) Protocol Agent provides support for related RTP (Real-time Transport Protocol) connections in Voice over IP (VoIP) traffic.

There are no configurable parameters for this Protocol Agent.

MSRPC Protocol Agent

The MSRPC (Microsoft RPC) Protocol Agent allows related connections for the endpoint mapper (EPM) protocol. It also handles NAT modifications for communications between Microsoft Outlook clients and Microsoft Exchange servers.

The MSRPC Protocol Agent supports TCP as the EPM connection method. By default, the Microsoft RPC/ EPM service is available on port 135/tcp and the communications continue using a dynamically allocated port. The Protocol Agent monitors the ports used to dynamically allow the connections based on the port allocation. This removes the need to allow the full range of ports.

If the traffic is Outlook/Exchange traffic, the Protocol Agent can also be used to support NAT for related connections by changing the payload data of the control connection.

On Engines, this Protocol Agent has parameters that you can set in the Service properties. On IPS engines and Layer 2 Engines, there are no configurable parameters for this Protocol Agent.

NetBIOS Protocol Agent

This Protocol Agent provides deep inspection for Windows NetBIOS Datagram Service connections. This Protocol Agent is also used to allow Windows NetBIOS Datagram Service connections through the Engine or Layer 2 Engine.

There are no configurable parameters for this Protocol Agent.

Oracle Protocol Agent

This Protocol Agent handles Oracle Transparent Network Substrate (TNS) protocol-based SQL*Net, Net7, and Net8 connections.

It is for cases where TCP port 1521 is used only for negotiating the port number for Oracle database connections and the port number for the actual connection is assigned dynamically.

This Protocol Agent is needed only if the database is on a different computer than the Oracle listener. The Oracle Protocol Agent does not change payload data because the database service connections might go through a different route than the listener connection. You can create custom Oracle agents with different settings when required.

RTSP Protocol Agent

The RTSP (Real Time Streaming Protocol) network control protocol is used for establishing and controlling media sessions between clients and media servers.

The RTSP Protocol Agent allows RTP (Real-time Transport Protocol) and RTCP (Real-time Control Protocol) media streaming connections initiated with RTSP through the engine. On Engines, it also handles NAT modifications to the protocol payload. This Protocol Agent has parameters you can set in the Service properties.

SCCP Protocol Agent

The SCCP (Skinny Call Control Protocol) provides support for related RTP (Real-time Transport Protocol) connections in VoIP traffic.

There are no configurable parameters for this Protocol Agent.

Services in Engine Protocol Agent

This Protocol Agent is intended for system services running on Engines. On Layer 2 Engines and IPS engines, it is also used with services running on Engines managed by the same Management Server as the Layer 2 Engine or IPS engine.

This Protocol Agent is only intended for the system's internal use. There are no configurable parameters for this Protocol Agent.

Shell Protocol Agent

This Protocol Agent manages Remote Shell connections and RExec connections.

Remote Shell is a widely used remote management protocol. RExec is a remote protocol with which it is possible to run commands on another computer.

This Protocol Agent has parameters that you can set in the Service properties.

SIP Protocol Agent

The Session Initiation Protocol (SIP) Protocol Agent can be used with Engines to handle multimedia connections that use SIP as their transfer protocol (such as VoIP).

Using the SIP Protocol Agent allows SIP to be used across a Engine that uses NAT. SIP uses TCP or UDP port 5060 to initiate the connection, after which the traffic is allocated a dynamically assigned port. The Protocol Agent tracks the actual ports used, so that the whole range of dynamic ports does not need to be allowed in the Engine Policy.

The SIP Protocol Agent can be configured to force the client or server address used within the SIP transport layer to also be used for the media stream carried over SIP (by default, enforced for both the client and the server).

This Protocol Agent has parameters you can set in the Service properties. This Protocol Agent is not available on IPS engines or Layer 2 Engines.

SMTP Protocol Agent

On Engines, the Simple Mail Transfer Protocol (SMTP) Protocol Agent can be used to redirect connections to a proxy service. On Layer 2 Engines and IPS engines, the Protocol Agent can be used for protocol validation and deep inspection.

On Engines, this Protocol Agent has parameters that you can set in the Service properties. On IPS engines and Layer 2 Engines, there are no configurable parameters for this Protocol Agent.

SSH Protocol Agent

Secure Shell (SSH) is an encrypted remote use protocol.

This Protocol Agent validates the communications to make sure the protocol used really is SSH. The SSH Agent validates SSHv1 only. On Engines, this Protocol Agent has parameters you can set in the Service properties. On IPS engines and Layer 2 Engines, there are no configurable parameters for this Protocol Agent.

SunRPC Protocol Agent

The Sun Remote Procedure Call (RPC) Protocol Agent assists the Engine, Layer 2 Engine, or IPS engine in Portmapper connections. It makes the handling of RPC program numbers used in the Access rules faster.

Only Portmapper connections going through the engine are assigned this Protocol Agent. This Protocol Agent is not intended for other communications.

The SunRPC Protocol Agent collects information about RPC services by interpreting the GET PORT and DUMP PORTS requests and their respective answers. All information it collects is stored in the Portmapper cache.

When the packet filter evaluates RPC matches, it checks the Portmapper cache to see if the destination of the packet has the appropriate service defined in the rule. If the cache does not have the requested information available, the packet under evaluation is not let through. A query is sent to the destination host for RPC information. The reply information is stored in cache.

There are no configurable parameters for this Protocol Agent.

TFTP Protocol Agent

The Trivial File Transfer Protocol (TFTP) Protocol Agent performs data transfer from a server to a client using dynamically selected ports.

There are no specific limits to the port range in the TFTP protocol (RFC 1350). Apart from Access rules, the TFTP Protocol Agent is also useful in NAT operations on Engines.

A TFTP Agent is attached to a UDP connection established between the client and the server. The client opens the control connection from a dynamically selected source port to the fixed destination port 69/UDP on the server. A separate UDP data connection is established between the client and the server after the client has sent a read or write command to the server. The server opens a connection from a dynamic source port to the client's destination port. The destination port is the same as the source port in the control connection.

On Engines, this agent has parameters you can set in the Service properties. On IPS engines and Layer 2 Engines, there are no configurable parameters for this Protocol Agent.

Examples of Protocol Agents

These examples illustrate some common uses for Protocol Agents and the general steps on how each example is configured.

Example: preventing active mode FTP with a Protocol Agent

Company A has an FTP server that allows access from the Internet. The Engine must only allow users to make passive mode FTP connections.

The administrators:

1) Create a Service element for passive FTP.
2) Attach the FTP Protocol Agent to the Service.
3) Change the active mode FTP setting to **No** in the Service properties.
4) Create an Access rule that allows users to connect to the FTP server using their custom-made Service element.
5) Refresh the policy on the IPS engine.

Example: logging URLs accessed by internal users with a Protocol Agent

Company B has decided to track which webpages the employees visit. In addition to logging the connections, the administrators also want to log URLs.

An Access rule allows all outbound connections from the internal networks to the Internet, regardless of the service. The administrators decide to add the HTTP Protocol Agent in a Continue rule.

The administrators:

1) Add the Continue rule above the existing Access rule, as follows.

Source	Destination	Service	Action
Internal Networks	Expression “NOT Local Protected Sites”	“HTTP (with URL Logging)” default Service	Continue
Internal Networks	Expression “NOT Local Protected Sites”	ANY	Allow

Note

Using the “NOT Local Protected Sites” expression requires the Alias “Local Protected Sites” to be configured with a translation value for the engine.

2) Refresh the engine’s policy.

Chapter 65

Sidewinder Proxies

Contents

- Sidewinder Proxies and how they work on page 1043
- Using Sidewinder Proxies on page 1048
- Change logging options for Sidewinder Proxies on page 1050
- Enable Sidewinder Proxy on page 1050
- Configure Sidewinder SSH Proxy on page 1051
- Create custom Service elements for Sidewinder Proxies on page 1056
- Add rules for Sidewinder Proxy on page 1057
- Advanced settings for Sidewinder Proxies on page 1057
- Supported advanced Sidewinder Proxy settings on page 1058

Sidewinder Proxies are software modules that provide network level proxies, protocol validation, and configurable application level protocol filtering and translation on Forcepoint Network Security Platform.

Sidewinder Proxies and how they work

On Sidewinder engines, proxies provide high assurance protocol validation. On Forcepoint Network Security Platform, Sidewinder Proxies are software modules on engines that enable some of the proxy features that are available on Sidewinder.

Note

Not all features supported by proxies on Sidewinder are supported by Sidewinder Proxies on Forcepoint Network Security Platform.

Sidewinder Proxies terminate connections at the engine, and make separate connections with each of the communicating hosts. Hosts do not communicate with each other directly. Each host communicates only with the Sidewinder Proxy. The proxy does not forward the original packet from one host to the other. Only the data from the original packets is forwarded. For TCP connections, the TCP sequence numbers, ACKs, packet sizes, windows, and other parameters are different on the two sides of the connection.

On Forcepoint Network Security Platform, Sidewinder Proxies are supported for both Engines and Virtual Engines.

Benefits of Sidewinder Proxies

Sidewinder Proxies provide these benefits.

Benefits of Sidewinder Proxies

Benefit	Description
More extensive attack prevention	Sidewinder Proxies can prevent most attacks that involve modifying IP, TCP, or UDP headers to evade detection. These kinds of attacks cannot always be detected using signature-based detection alone.
Covert channel prevention	Sidewinder Proxies can prevent attempts to covertly send data in the header bits of IP, TCP, or UDP packets.
More detailed control of application protocols	Sidewinder Proxies provide more granular control of application protocols. For example, you can use the SSM HTTP Proxy for detailed HTTP header filtering and control.

Limitations of Sidewinder Proxies

These limitations apply to Sidewinder Proxies.

- Sidewinder Proxies are not supported with legacy role-specific licenses. An Security Engine license is required.
- Only the Engine/VPN role is supported.
- On Engine Clusters, connection failover is not supported.
- Master Engines only support advanced Sidewinder Proxy settings for the Virtual Engines that they host. Master Engines cannot use Sidewinder Proxies for their own traffic.
- Sidewinder Proxies can be used with both IPv4 and IPv6 traffic. However, they do not support translation between IPv4 and IPv6, or tunneling of IPv4 or IPv6 traffic in an IPv6 or IPv4 connection.
- Sidewinder Proxies do not support multicast traffic.
- In the Engine Policy, you cannot place Access rules that match based on Endpoint Context Agent information above Access rules for Sidewinder Proxies.
- You cannot use McAfee Threat Intelligence Exchange (TIE) file reputation with Sidewinder Proxies. If the File Filtering Policy uses these file reputation scans, they are ignored for traffic that uses Sidewinder Proxies. When you restrict file types with file filtering, the action specified for the **Action When No Scanners Are Available** option determines whether the file transfer is allowed or blocked if no other scanners are available.

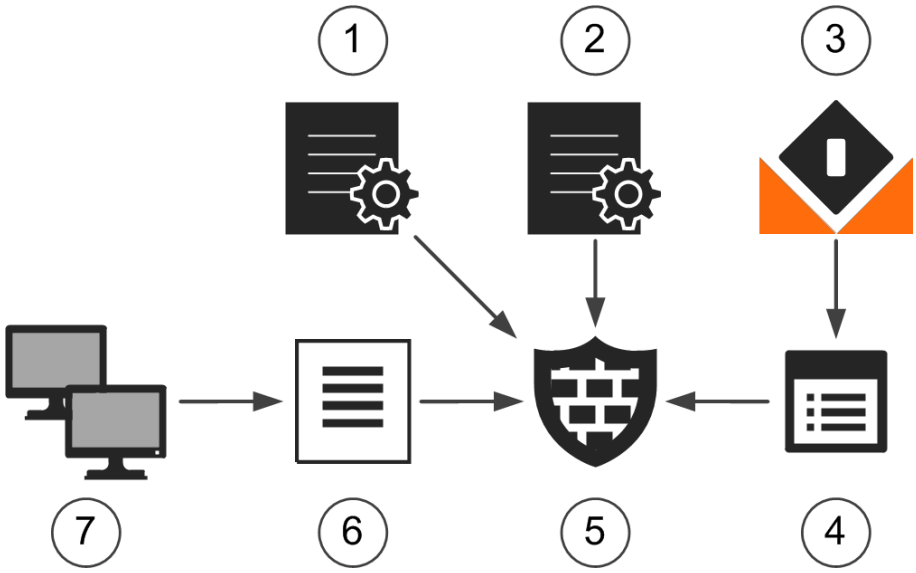
Related tasks

Restrict file types with file filtering on page 1007

Sidewinder Proxy configuration overview

To use Sidewinder Proxy, you must enable the feature and add Access rules. You can optionally add other elements to customize how Sidewinder Proxies work.

Elements in the configuration



1. SSH Profile
2. Sidewinder Logging Profile
3. Service with Proxy
4. Access Rules
5. Security Engine
6. SSH Known Host List
7. SSH Known Hosts

Follow these general steps to configure Sidewinder Proxies:

1) (Optional) To customize logging options for Sidewinder Proxies, create Sidewinder Logging Profiles.

2) Enable Sidewinder Proxy for the engine in the Engine Editor.

3) (Optional) Create elements and configure settings for specific Sidewinder Proxies.
   a) Create SSH Profiles for Sidewinder SSH Proxy.
   b) Create SSH Known Hosts for Sidewinder SSH Proxy and group SSH Known Hosts using SSH Known Hosts Lists.
   c) Configure settings for Sidewinder SSH Proxy in the Engine Editor.

4) (Optional) Create custom Service elements and configure Protocol Parameters for Sidewinder Proxies.

Note

There are no configurable Protocol Parameters for the Sidewinder TCP Proxy or the Sidewinder UDP Proxy.

5) Add Service elements to Access rules to specify which traffic uses Sidewinder Proxies.

Default elements for Sidewinder Proxy

The SMC has default elements that you can use to configure Sidewinder Proxy.

You cannot edit the default elements, but you can duplicate the default elements and edit the duplicated elements.

Default elements for Sidewinder Proxy

Type of Element	Element	Description
Protocol	**Engine Configuration** > **Other Elements** > **Protocols** > **By Protocol Type** > **Proxy**	
	SSM HTTP Proxy	Allows you to use Sidewinder HTTP Proxy in Service elements, and provides Protocol Parameters for detailed control of the HTTP and HTTPS protocols. This Protocol element is automatically selected in the properties of the default SSM HTTP and SSM HTTPS Proxy Service elements.
	SSM SSH Proxy	Allows you to use Sidewinder SSH Proxy in Service elements, and provides Protocol Parameters for detailed control of the SSH protocol. This Protocol element is automatically selected in the properties of the default SSM SSH Service element.
	SSM TCP Proxy	Allows you to use Sidewinder TCP Proxy in Service elements. **Note** This Protocol element has no configurable Protocol Parameters.
	SSM UDP Proxy	Allows you to use Sidewinder UDP Proxy in Service elements. **Note** This Protocol element has no configurable Protocol Parameters.
	[Protocol] with [Proxy Protocol]	These Protocol elements combine a standard Protocol element with a Sidewinder Proxy Protocol element. Combined Protocol elements make the Protocol Parameters from both Protocol elements available in the same Service element.
Service	**Engine Configuration** > **Other Elements** > **Services** > **With Proxy**	

Type of Element	Element	Description
	SSM DNS Proxy (TCP)	Allows you to use DNS traffic with the Sidewinder TCP Proxy. This Service element has the following default settings: ■ **Dst. Ports** — 53 ■ **Protocol** — SSM DNS Proxy (TCP)
	SSM DNS Proxy (UDP)	Allows you to use DNS traffic with the Sidewinder UDP Proxy. This Service element has the following default settings: ■ **Dst. Ports** — 53 ■ **Protocol** — SSM DNS Proxy (UDP)
	SSM FTP Proxy	Allows you to use FTP traffic with the Sidewinder TCP Proxy. This Service element has the following default settings: ■ **Dst. Ports** — 21 ■ **Protocol** — SSM FTP Proxy
	SSM HTTP Proxy	Allows you to use HTTP traffic with the Sidewinder HTTP Proxy. This Service element has the following default settings: ■ **Dst. Ports** — 80 ■ **Protocol** — SSM HTTP Proxy
	SSM HTTPS Proxy	Allows you to use HTTPS traffic with the Sidewinder HTTP Proxy. This Service element has the following default settings: ■ **Dst. Ports** — 443 ■ **Protocol** — SSM HTTP Proxy
	SSM SSH Proxy	Allows you to use SSH traffic with the Sidewinder SSH Proxy. This Service element has the following default settings: ■ **Dst. Ports** — 22 ■ **Protocol** — SSM SSH Proxy
	SSM TFTP Proxy	Allows you to use TFTP traffic with the Sidewinder UDP Proxy. This Service element has the following default settings: ■ **Dst. Ports** — 69 ■ **Protocol** — SSM TFTP Proxy
Sidewinder Logging Profile	Sidewinder Default	Contains default settings that define which events detected by Sidewinder Proxies are logged, and how often logs are created.
SSH Profile	High Compatibility Profile	Contains default settings for key exchange, encryption algorithm, and message authentication that are compatible with a variety of SSH server software.
SSH Host Keys	ECDSA 256, DSA 1024, and RSA 1024	When you enable Sidewinder Proxy in the Engine Editor, these SSH Host Keys are automatically created for the engine.

Using Sidewinder Proxies

You can use Sidewinder Proxies on Forcepoint Network Security Platform to enforce protocol validation and to restrict the allowed parameters for each protocol.

Sidewinder Proxies are primarily intended for users in high assurance environments, such as government or financial institutions. In environments that limit access to external networks or access between networks with different security requirements, you can use Sidewinder Proxies for data loss protection.

Sidewinder HTTP Proxy

You can use the Sidewinder HTTP Proxy with HTTP and HTTPS traffic to enforce strict protocol standards, log URLs in requests, validate requests, and block some types of content in requests.

You can use the Sidewinder HTTP Proxy with or without decryption.

- When decryption is enabled, the Security Engine decrypts HTTPS traffic, then applies the Sidewinder HTTP Proxy and optionally inspection to the encapsulated HTML. After inspection, the Security Engine re-encrypts the HTTPS traffic.
- When decryption is not enabled, the Sidewinder HTTP Proxy only validates HTTPS traffic to make sure that the traffic contains valid HTTPS protocol messages.

Using the Sidewinder HTTP Proxy with decryption for HTTPS traffic provides the following benefits compared to the standard TLS inspection feature:

- The Sidewinder HTTP Proxy can remove TCP options from HTTPS traffic.
- The Sidewinder HTTP Proxy can present a configurable warning page to inform users that their traffic is being decrypted.

Decrypting and re-encrypting HTTPS traffic requires the following configurations:

- You must have a Client Protection CA and other elements required for TLS inspection. To avoid certificate-related warnings in end users' web browsers, the client protection CA certificate must be imported as a trusted certificate in the browsers.

  Note

  The Sidewinder HTTP Proxy only provides client protection. The Sidewinder HTTP Proxy does not provide server protection for servers in the internal network. The Sidewinder HTTP Proxy is not compatible with servers that use client certificates for authentication.

- You must configure an external DNS resolver, and select one or more DNS IP addresses in the Engine Editor. The DNS resolver must be functioning and available to provide DNS results to the engine.
- To allow the certificate manager to communicate with an external certificate revocation list (CRL) server, you must add an Access rule that allows HTTP traffic on port 80 between the engine and the Internet for making online certificate status protocol (OCSP) queries and fetching CRLs.

The **Decryption** option in the Allow Action Options in Access rules defines whether traffic that matches the rule is decrypted. To exclude specific traffic from decryption by the SSM HTTP Proxy, add the following type of Access rule:

Source	Destination	Service	Action
Source IP address	Destination IP address	One or more of the following Service elements: ■ SSM HTTPS Proxy ■ A custom Service element that uses the SSM HTTP Proxy Protocol	Allow Decryption: Disallowed

Related tasks

Configure TLS inspection for client protection on page 1068

Sidewinder SSH Proxy

The Sidewinder SSH Proxy allows you to restrict the types of traffic and the commands that can be used with SSH connections.

For example, you can use the Sidewinder SSH Proxy to block port forwarding or to restrict the commands allowed in file transfers using the SSH protocol.

You can also use the Sidewinder SSH Proxy to enforce encryption strength for the connections. In the Protocol Parameters, you can separately specify the key type and key length for each side of the connection. The **Client Advanced Settings** define settings for connections between the Sidewinder SSH Proxy and the client. The **Server Advanced Settings** define settings for connections between the Sidewinder SSH Proxy and the server.

Sidewinder TCP Proxy and Sidewinder UDP Proxy

Sidewinder TCP Proxy and Sidewinder UDP Proxy provide protocol validation for TCP and UDP traffic.

There are no default Service elements for Sidewinder TCP and Sidewinder UDP Proxy. If you want to apply the Sidewinder TCP Proxy or the Sidewinder UDP Proxy to TCP or UDP services, you must create custom Service elements. There are no configurable Protocol Parameters for the SSM TCP Proxy or the SSM UDP Proxy Protocol elements.

Using Sidewinder Proxies in combined Protocol elements

Combined Protocol elements allow you to apply a standard Protocol element and a Sidewinder Proxy Protocol element to the same traffic.

Combined Protocol elements make the Protocol Parameters from both Protocol elements available in the same Service element.

Combined Protocol elements have the following benefits:

- Sidewinder Proxy Protocol elements provide improved security compared to using non-proxy Protocol elements alone.

- Protocol elements provide protocol-specific inspection and protocol validation that is not possible using Sidewinder Proxy Protocol elements alone.

Combined Protocol elements appear under the **Engine Configuration** > **Other Elements** > **Protocols** > **By Protocol Type** > **Proxy** branch in the Configuration view.

Change logging options for Sidewinder Proxies

Sidewinder Logging Profiles define which events detected by Sidewinder Proxies create log entries, and how often log entries are created.

If the settings in the default Sidewinder Logging Profile element meet your needs, there is no need to create a custom Sidewinder Logging Profile element.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Browse to **Other Elements** > **Sidewinder Elements** > **Sidewinder Logging Profiles**.
3) Right-click **Sidewinder Logging Profiles**, then select **New Sidewinder Logging Profile**.
4) In the **Name** field, enter a unique name.
5) Select logging options for Situation Categories or individual Situations.
6) Click **OK**.

Enable Sidewinder Proxy

To start using Sidewinder Proxy, you must enable the feature in the Engine Editor.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click a Single Engine or Engine Cluster and select **Edit <element type>**.
3) In the navigation pane on the left, browse to **Add-Ons** > **Sidewinder Proxy**.
4) Select **Enable**.

5) (Optional) Select a customized Sidewinder Logging Profile.
   a) Next to the **Sidewinder Logging Profile** field, click **Select**.
   b) Select a Sidewinder Logging Profile element, then click **Select**.
6) Click **Save**.

Configure Sidewinder SSH Proxy

Follow these general steps to configure Sidewinder SSH Proxy.

1) (Optional) To define custom settings for key exchange, encryption algorithm, and message authentication, create an SSH Profile element.
2) (Optional) Allow connections only to specific trusted servers.
   a) Create SSH Known Host elements.
   b) Group the SSH Known Hosts using SSH Known Hosts Lists.
   c) Select SSH Known Hosts Lists in the Engine Editor.
3) (Optional) Add host keys for Sidewinder SSH Proxy in the Engine Editor.
4) Add custom Service elements for Sidewinder SSH Proxy.

Create SSH Profiles for Sidewinder SSH Proxy

SSH Profiles define custom settings for key exchange, encryption algorithm, and message authentication for SSH connections that use the Sidewinder SSH Proxy.

If the default High Compatibility Profile SSH Profile element meets your needs, it is not necessary to create a custom SSH Profile.

The lists of selected key exchange methods, cipher methods, and message authentication code (MAC) methods are organized in order of preference. You can move selected items up or down in the list.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Browse to **Other Elements** > **Sidewinder Elements** > **SSH Profiles**.
3) Right-click **SSH Profiles**, then select **New SSH Profile**.

4) In the **Name** field, enter a unique name.

5) On the **Key Exchange** tab, select one or more key exchange methods, then click **Add**.

6) On the **Cipher** tab, select one or more cipher methods, then click **Add**.

7) On the **MAC** tab, select one or more MAC methods, then click **Add**.

8) Click **OK**.

Next steps

If you want to allow connections only to specific trusted servers, create SSH Known Hosts. Otherwise, select your custom SSH Profile in the Engine Editor.

Create SSH Known Hosts for Sidewinder SSH Proxy

SSH Known Host elements contain the contact information, public key, and fingerprint that you use to authenticate servers that you trust.

Tip

The Sidewinder SSH Proxy logs the SSH server key for each connection attempt. You can optionally add SSH Known Hosts from log entries using the right-click menu.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Browse to **Other Elements** > **Sidewinder Elements** > **SSH Known Hosts** > **SSH Known Hosts**.

3) Right-click **SSH Known Hosts**, then select **New SSH Known Host**.

4) In the **Name** field, specify the name of the element in one of these ways:
   - Enter the domain name of the server.
   - Enter a unique name for the server.

5) Enter one or more IP addresses.

   Note

   You can enter both an IPv4 address and an IPv6 address. You can only enter one address of each type.

   - In the **IPv4 Address** field, enter the IPv4 address of the server.

Tip

To automatically resolve the IP address from the domain name, enter the domain name, then click **Resolve**.

- In the **IPv6 Address** field, enter the IPv6 address of the server.

6) If the server communicates on a port other than the default port (TCP 22), enter the port number in the **Port** field.

7) To manually enter the SSH key for the server, enter the key as text.
   a) From the **Key Type** drop-down list, select the algorithm used for the key.
   b) In the **SSH Key** field, type or paste the key.

8) To use an existing key file as the SSH key for the server, import the key file.
   a) Click **Import**.
   b) Select the key file, then click **Open**.

9) To retrieve a public key from a host, follow these steps.
   a) From the **Key Type** drop-down list, select the algorithm used for the key.
   b) Click **Retrieve**.
   c) Select a engine, then click **Select**.

   The engine sends a public key request to the host, and retrieves the SSH key and fingerprint.

10) Click **OK**.

Next steps

Group SSH Known Host elements using SSH Known Hosts Lists.

Group SSH Known Hosts in SSH Known Hosts Lists

To allow connections only to specific trusted servers, group SSH Known Hosts in SSH Known Hosts Lists.

You select SSH Known Hosts Lists for each Security Engine in the Engine Editor.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Browse to **Other Elements** > **Sidewinder Elements** > **SSH Known Hosts** > **SSH Known Hosts Lists**.

3) Right-click **SSH Known Hosts Lists**, then select **New SSH Known Hosts List**.

4) In the **Name** field, enter a unique name.

5) Click **Add**, select one or more SSH Known Host elements, then click **Select**.

6) Click **OK**.

Next steps

Select SSH Known Hosts Lists for engines

Select SSH Known Hosts Lists for Security Engine

To allow connections only to specific trusted servers, select SSH Known Hosts Lists in the Engine Editor.

Note

When **Use Strict Known Hosts List** is selected for the **Server Host Key Validation** option in the properties of the custom Service element for SSM SSH Proxy, you must select SSH Known Hosts Lists for the engine. If you do not select an SSH Known Hosts List, connections are not allowed to any hosts.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click a Single Engine or Engine Cluster and select **Edit <element type>**.

3) In the navigation pane on the left, browse to **Add-Ons** > **Sidewinder Proxy**.

4) Next to the **SSH Known Hosts Lists** table, click **Add**.

5) Select one or more SSH Known Hosts List elements, then click **Select**.

6) Click **Save**.

Next steps

If the default SSH host keys do not meet your needs, or if you want to specify which host keys are used for specific SSH Proxy Services, add host keys for Sidewinder SSH Proxy.
Otherwise, add custom Service elements for Sidewinder SSH Proxy.

Add host keys for Sidewinder SSH Proxy

In the Engine Editor, you can add host keys for Sidewinder SSH Proxy and specify which host keys are used for specific SSH Proxy Services.

When you enable Sidewinder proxy in the Engine Editor, 3 default SSH Host Keys are automatically created for the engine. You can optionally replace the automatically created host keys and create more host keys. You can import existing host keys or generate host keys. You can also associate host keys with specific SSH Proxy Services. Each engine on which you use the SSH Proxy Service must have one host key of each type.

Note

You can only associate one key of each key type with the same Service. This limitation includes the default Proxy Service that is used if you do not select a specific Service.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click a Single Engine or Engine Cluster and select **Edit <element type>**.
3) In the navigation pane on the left, browse to **Add-Ons** > **Sidewinder Proxy**.
4) (Optional) Generate host keys.
   a) Next to the **Host Keys** table, click **Add**.
   b) From the **Host Key Type** drop-down list, select the algorithm to use for the key.
   c) From the **Host Key Length** drop-down list, select the length of the key.
   d) Click **Add**.
5) (Optional) Import existing host keys.
   a) Next to the **Host Keys** table, click **Import**.
   b) Select the key file, then click **Import**.
6) Specify which host keys are used for specific SSH Proxy Services.
   a) In the **Host Keys** table, double-click the **SSH Proxy Services** cell.
   b) From the **Resources** list, select one or more Service elements, then click **Add**.
   c) Click **OK**.
7) Click **Save**.

Next steps

Add custom Service elements for Sidewinder SSH Proxy.

Create custom Service elements for Sidewinder Proxies

If the default Service elements for Sidewinder Proxies do not meet your needs, add custom Service elements for Sidewinder Proxies.

Add a custom Service element in the following cases:

- You want to change the Protocol Parameters of the default Service elements for Sidewinder Proxies.
- You want to use combined Protocol elements.
- You want to apply the Sidewinder TCP Proxy or the Sidewinder UDP Proxy to TCP or UDP services.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Browse to **Other Elements** > **Services**.

3) Create the Service element in one of the following ways:
   - To create an element with no settings predefined, right-click the branch for the type of Service you want to create, then select **New** > **TCP Service** or **New** > **UDP Service**.
   - To create a Service based on an existing TCP or UDP Service element, right-click the existing Service, then select **New** > **Duplicate**.
   - To create a Service based on one of the default Service elements for Sidewinder Proxies, browse to **With Proxy**, right-click the existing Service, then select **New** > **Duplicate**.

4) In the **Name** field, enter a unique name.

5) If you did not duplicate one of the default Service elements for Sidewinder Proxies, click **Select** next to the **Protocol** field, browse to **TCP Proxy** or **UDP Proxy**, then select an SSM Proxy Protocol element or a combined Protocol element.

6) On the **Protocol Parameters** tab, select options according to your needs.



Note

There are no configurable Protocol Parameters for the Sidewinder TCP Proxy or the Sidewinder UDP Proxy.

7) Click **OK**.

Next steps

Use the custom Service element in the Access rules.

Add rules for Sidewinder Proxy

Use one of the default Service elements for Sidewinder Proxy or a custom Service element in the Access rules to specify which traffic uses Sidewinder Proxies.

In some cases, connections might not use the Sidewinder Proxy.

- If the policy contains rules that match traffic based on the payload, such as Applications or category-based web filtering, connections might not match rules that specify an Sidewinder Proxy.
  To avoid this limitation, do not add rules that match traffic based on the payload to the same policy where you use SSM Proxies.
- If you use a rule with the Continue action to specify a Sidewinder Proxy as a default Protocol, rules later in the policy can override the defaults set in Continue rules.
  To avoid this limitation, do not add rules that specify a Protocol of the type Protocol Agent or Protocol Tag for the same matching criteria as the Continue rules for SSM Proxies.

Because many websites use a combination of HTTP and HTTPS, users might not be able to connect to these websites if you add only one rule that applies the Sidewinder HTTP Proxy to HTTPS traffic. We recommend that you add a separate rule that allows unencrypted HTTP traffic or allow HTTP traffic in the same rule that applies the Sidewinder HTTP Proxy to HTTPS traffic.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) To specify which traffic uses Sidewinder Proxies, add the following type of Access rules to the Engine Policy:

Access rules for Sidewinder Proxy

Source	Destination	Service	Action
The elements that represent hosts in the internal network, or **ANY**.	The elements that represent external servers, or **ANY**.	One or more custom Service elements, or one or more default Service elements for Sidewinder Proxy.	**Allow** or **Continue**

2) Save and install the policy to start using the new configuration.

Advanced settings for Sidewinder Proxies

Advanced Sidewinder Proxy settings enable optional behavior or set parameters.

Each setting has a name and a value. Some settings are specific to individual proxies. Most settings can be used with any proxy, or for all proxies if you set them using shared proxy properties. If you do not specify a value, the default values are used.

Note

Changing the advanced settings requires detailed knowledge of the protocols involved and how the parameters affect your environment. We do not recommend changing the advanced settings unless you are instructed to do so by Forcepoint Technical Support.

In an environment with Master Engines and Virtual Engines, some advanced settings apply only to Master Engines, and some advanced settings apply only to Virtual Engines. Settings that do not apply to the type of engine on which they are configured are ignored. For example, if you configure a Master Security Engine advanced setting for a Virtual Engine, the setting has no effect on the Virtual Engine.

Advanced settings that you configure for a Virtual Engine affect only the individual Virtual Engine. Advanced settings that you configure for a Master Engine affect all Virtual Engines that are hosted by the Master Security Engine.

Supported advanced Sidewinder Proxy settings

This table lists the most commonly used advanced settings for Sidewinder Proxies.

Note

All advanced Sidewinder Proxy settings can be configured for Engines. Some settings can be configured for Master Engines or for Virtual Engines. Settings that do not apply to the type of engine on which they are configured are ignored.

Supported advanced Sidewinder Proxy settings

Property	Supported proxy types	Accepted values	Default value	Supported engine types	Description
allow_client_half_close	HTTP	0 or 1	1	Engine	If 1, allows clients to receive data after indicating that they will send no more.
debug_level	HTTP SSH TCP UDP	0–4	0	Engine, Master Engine	If any value other than 0, enables debugging messages. Higher values produce more output. See also send_debug_to_log.
display_user_warning_ttl	HTTP	Numerical values in seconds	43200	Engine	The default time an entry stays in the decryption warning page cache.
display_user_warning_dest	HTTP	0 or 1	0	Engine	If 1, the decryption warning page is displayed for each unique combination of source and destination address. If 0, the decryption warning page is displayed for each unique source address.

Property	Supported proxy types	Accepted values	Default value	Supported engine types	Description
enable_certificate_revocation_check	HTTP	0 or 1	1	Engine	If 1, the HTTP proxy validates the status of server certificates using certificate revocation lists (CRLs) or on-line certificate status protocol (OCSP).
encoded_url_max	HTTP	Numerical values in kilobytes	100000 (100 megabytes)	Engine, Master Engine	Maximum size of an encoded URL that can be decoded in normalization. Normalization can make up to 6 copies of a URL.
header_waiting	HTTP	0–100	25	Engine, Master Engine	Limit for the percentage of proxy sessions waiting for additional HTTP header information. If this limit is reached, half of the waiting sessions are discarded.
max_header_total_size	HTTP	Numerical values	65536	Engine	Maximum size of all HTTP header data (not just individual lines).
net.inet.ip.random_id	HTTP SSH TCP UDP	0 or 1	0	Engine	If 1, assigns random ip_id values to outgoing IPv4 packets. The default behavior is to assign a random initial value for each proxy instance, and increment for each outgoing packet.
net.inet.ip.ttl	HTTP SSH TCP UDP	Numerical values in the number of hops	64	Engine	The maximum time to live (TTL) in hops for IPv4 packets that are sent.
net.inet.tcp.always_keepalive	HTTP SSH TCP	0 or 1	1	Engine	If 1, enables use of TCP keepalive probes on all connections.
net.inet.tcp.drop_synfin	HTTP SSH TCP	0 or 1	1	Engine	If 1, drops TCP packets that have SYN+FIN set.
net.inet.tcp.keepidle	HTTP SSH TCP	Numerical values in milliseconds	7200000 (2 hours)	Engine	Time, in milliseconds, that the connection must be idle before keepalive probes are sent.
net.inet.tcp.keepinit	HTTP SSH TCP	Numerical values in milliseconds	75000 (75 seconds)	Engine	Time allowed to establish connection.

Property	Supported proxy types	Accepted values	Default value	Supported engine types	Description
net.inet.tcp.keepintvl	HTTP SSH TCP	Numerical values in milliseconds	75000 (75 seconds)	Engine	Time between keepalive probes.
net.inet.tcp.msl	HTTP SSH TCP	Numerical values in milliseconds	15000 (15 seconds, TCP TIME_WAIT time 30 seconds)	Engine	Maximum segment lifetime. The default TCP TIME_WAIT time is double this value.
net.inet.tcp.recvbuf_auto	HTTP SSH TCP	0 or 1	1	Engine	If 1, enables automatic receive buffer sizing.
net.inet.tcp.recvbuf_inc	HTTP SSH TCP	Numerical values in bytes	16K	Engine	Incrementor step size of automatic receive buffer. Use the following suffixes to specify larger values: ■ K — Kilobytes ■ M — Megabytes ■ G — Gigabytes
net.inet.tcp.recvbuf_max	HTTP SSH TCP	Numerical values in bytes	96K	Engine	Maximum size of automatic receive buffer. Use the following suffixes to specify larger values: ■ K — Kilobytes ■ M — Megabytes ■ G — Gigabytes
net.inet.tcp.recvspace	HTTP SSH TCP	Numerical values in bytes	64K	Engine	Size of the initial TCP receive window. Use the following suffixes to specify larger values: ■ K — Kilobytes ■ M — Megabytes ■ G — Gigabytes
net.inet.tcp.rfc1323	HTTP SSH TCP	0 or 1	1	Engine	If 1, enables the TCP timestamp option and window scaling option specified in RFC 1323, which allows per-packet timestamps, protection against wrapped sequences, and windows larger than 65535 bytes.
net.inet.tcp.sendbuf_auto	HTTP SSH TCP	0 or 1	1	Engine	If 1, enables automatic send buffer sizing.

Property	Supported proxy types	Accepted values	Default value	Supported engine types	Description
net.inet.tcp.sendbuf_inc	HTTP SSH TCP	Numerical values in bytes	8K	Engine	Incrementor step size of automatic send buffer. Use the following suffixes to specify larger values: ■ K — Kilobytes ■ M — Megabytes ■ G — Gigabytes
net.inet.tcp.sendspace	HTTP SSH TCP	Numerical values in bytes	32K	Engine	Size of the initial TCP send window. Use the following suffixes to specify larger values: ■ K — Kilobytes ■ M — Megabytes ■ G — Gigabytes
net.inet.udp.checksum	UDP	0 or 1	1	Engine	If 1, requires checksums on incoming UDP packets.
net.inet6.ip6.hlim	HTTP SSH TCP UDP	Numerical values	64	Engine, Virtual Engine	The hop limit for IPv6 packets that are sent.
reserved_allowed	SSH	0 or 1	1	Engine	If 1, allows messages in the reserved range.
send_debug_to_log	HTTP SSH TCP UDP	0 or 1	1	Engine, Master Engine	If 1, debugging messages are sent to the Log Server. If 0, messages are written to a file. **Note** Change this value only if instructed to do so by Forcepoint Customer Hub.
server_requests_allowed	SSH	0 or 1	1	Engine	If 1, allows global requests from the server.
server_channels_allowed	SSH	0 or 1	1	Engine	If 1, allows the server to open channels.
sftp_extensions_allowed	SSH	0 or 1	1	Engine	If 1, allows local SFTP extension commands.
ssh_extensions_allowed	SSH	0 or 1	1	Engine	If 1, allows local SSH extension messages.

Property	Supported proxy types	Accepted values	Default value	Supported engine types	Description
tls_cipher_override	HTTP	A single valid OpenSSL cipher string	ALL:-SEED: -RC4: -CAMELLIA: -PSK: -MD5: -SRP:-DES: -ADH: -AECDH: -kDH: -kECDH: -IDEA@ STRENGTH	Engine	The list of cipher algorithms that the HTTP Proxy negotiates with its peers. The default cipher list includes only cipher algorithms that are allowed in FIPS mode. Minus signs (-) exclude the specified ciphers from the ALL list. **Tip** You can use this setting to restrict the default cipher list or to add more cipher algorithms.
tls_curves_override	HTTP	A colon-separated list of OpenSSL elliptic curve names	P-521:P-384: P-256	Engine	The list of the elliptic curves supported by the HTTP Proxy. The default list includes only elliptic curves that are allowed in FIPS mode.
tls_key_curve_override	HTTP	A single OpenSSL elliptic curve name	P-521	Engine	The default curve that the HTTP Proxy uses to generate the elliptic curve private key for substitute certificates.
tls_protocol_override	HTTP	A colon-separated list of TLS version strings. Valid version strings are SSLv3, TLSv1.0, TLSv1.1, and TLSv1.2.	TLSv1.0: TLSv1.1: TLSv1.2	Engine	The TLS protocol versions supported by the HTTP Proxy. The default list includes only TLS protocol versions that are allowed in FIPS mode. **Tip** You can use this setting to restrict the default list or to add TLS versions, such as SSLv3, that are not included in the default list.
undefined_allowed	SSH	0 or 1	1	Engine	If 1, allows messages for which the proxy does not have a protocol handler.

Chapter 66

Setting up TLS inspection

Contents

- TLS inspection and how it works on page 1063
- Configure TLS inspection for server protection on page 1066
- Configure TLS inspection for client protection on page 1068
- Define trusted certificate authorities for TLS inspection on page 1075
- Exclude traffic from decryption for TLS inspection on page 1076
- Active destination server certificate probing on page 1079
- Examples of TLS inspection on page 1080

The TLS inspection feature decrypts TLS connections so that they can be inspected for malicious traffic and then re-encrypts the traffic before sending it to its destination.



Note

Security Engine versions 6.11 and higher support TLS 1.3 and decrypts all supported TLS 1.3 cryptographic algorithms.

TLS inspection and how it works

TLS inspection allows you to decrypt traffic uses the TLS protocol to secure connections, such as HTTPS traffic, so that it can be inspected.

The TLS inspection feature consists of server protection and client protection:

- Server protection decrypts incoming TLS connections from external clients to servers in the protected network.
- Client protection decrypts outgoing TLS connections initiated by clients in the protected network to external servers.

You can use client protection alone, server protection alone, or client and server protection together. After decrypting the traffic, you can apply normal HTTP inspection and optionally malware scanning to the traffic. If the traffic is allowed to continue, the Security Engine re-encrypts the traffic and forwards it to its original destination.

Certificates for TLS inspection

TLS Inspection requires two separate secure connections: one from the client to the Security Engine and one from the Security Engine to the server. Substituting its own certificate for the original server certificate allows the Security Engine to decrypt and re-encrypt the traffic.

HTTPS uses the TLS protocol to secure HTTP connections. When a client browser connects to a server that uses HTTPS, the server sends a certificate to the browser. The certificate contains the server's public key and a digital signature from a certificate authority that verifies the server's identity. The browser and the server negotiate an encryption algorithm, which is used to create the encrypted connection.

When a client in the protected network initiates a TLS connection to an external server, the Security Engine checks whether the server's certificate was signed by a certificate authority that is considered trusted. If the certificate was signed by a trusted certificate authority, the Security Engine makes a new certificate that matches the server's certificate. From the point of view of a user in the protected network, the process is invisible: the connection is established in the same way as a TLS connection made directly to the server.

When a server's certificate is self-signed or has not been signed by a trusted certificate authority, the Security Engine cannot trust the server certificate. In this case, the Security Engine makes a new self-signed certificate. This certificate is presented to the user in the protected network. The user's browser shows the same warning that it would show if it received a self-signed certificate directly from a server. In this case, the user must decide whether to accept the certificate.

When a server in the protected network is the destination of an incoming TLS connection, the Security Engine uses the server's credentials to decrypt and re-encrypt the traffic.

Limitations and considerations for using TLS inspection

Consider these limitations and other important information before configuring TLS inspection.

TLS inspection has the following limitations:

- TLS inspection for client protection cannot be done for traffic picked up through Capture interfaces.
- TLS inspection for server protection can be done for traffic picked up through both Capture interfaces and Inline interfaces.

  Note

  Due to security features of the TLS protocol, TLS decryption for traffic picked up through Capture interfaces can only be done when RSA key exchange negotiation is used between the client and the server.

- TLS inspection is not supported on Single IPS engines or on Single Layer 2 Engines if they are deployed alongside a Engine Cluster that uses dispatch clustering.
- Default Trusted Certificate Authority elements are automatically added to the SMC from dynamic update packages and cannot be edited or deleted.
- TLS inspection is not supported on Master Engines.

Consider this important information before configuring TLS inspection:

- Traffic that uses TLS might be protected by laws related to the privacy of communications. Decrypting and inspecting this traffic might be illegal in some jurisdictions.
- When a certificate for client or server protection has been uploaded to the Security Engine, it is possible to unintentionally enable TLS decryption for all traffic in one of the following ways:
  - Adding a Network Application that allows or requires the use of TLS to an Access rule
  - Enabling the logging of Application information in the Access rules
  - Enabling Deep Inspection in an Access rule with the Service cell of the rule set to ANY
- Strict TCP inspection mode is automatically applied to TCP connections when TLS Inspection is used.

Security considerations for TLS inspection

You must carefully consider security precautions when using TLS Inspection.

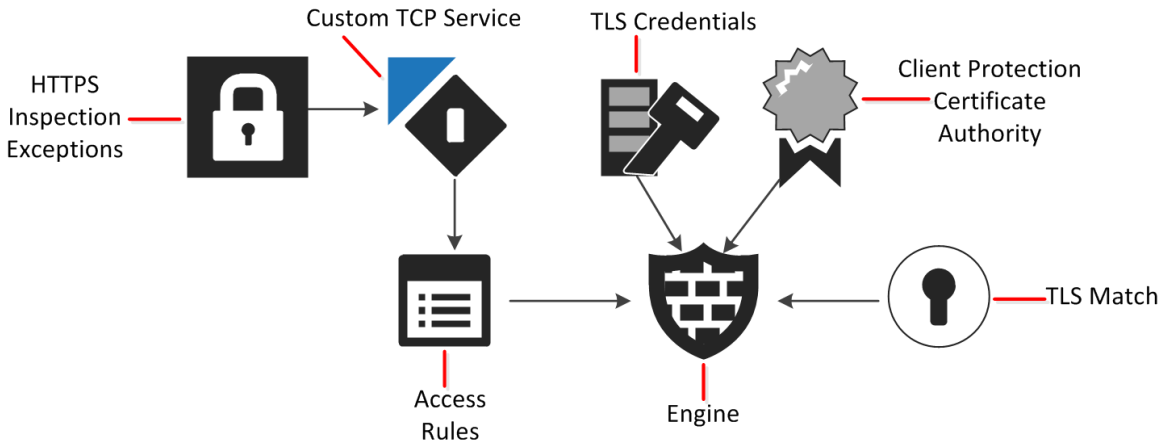
The TLS communications mediated by the Security Engine are decrypted for inspection, and the private keys of the servers are stored in the TLS Credentials elements on the Management Server. For these reasons, you must carefully consider security precautions when using TLS inspection. The following recommendations are general guidelines for ensuring the security of the Security Engine and the SMC:

- Run the Management Server on a hardened operating system.
- Disable SSH access to the engine's command line if it is not needed regularly.
- Make sure that the engine's Control IP address is in a protected network.
- Save Management Server backups as encrypted files.

TLS inspection configuration overview

To use TLS inspection, you must configure TLS Credentials elements and Client Protection Certificate Authority elements. You must also activate client protection or server protection in the engine properties and enable TLS inspection in the Access rules.

Elements in the configuration



The TLS Credentials and the Client Protection Certificate Authority elements are specified in the properties of the engine that provides TLS Inspection. The engine uses the private key and certificate stored in the TLS Credentials to decrypt traffic to and from TLS servers in the protected network for inspection.

The Client Protection Certificate Authority element contains a private key and a certificate. The engine uses the private key stored in the Client Protection Certificate Authority element to sign the certificates presented to the end user, and the certificate to negotiate encrypted connections with TLS servers.

TLS Match elements define matching criteria for the use of the TLS protocol in traffic, and allow you to prevent specified traffic from being decrypted. TLS Matches that deny decrypting are applied globally, even if the TLS Match elements are not used in the policy.

The HTTPS Inspection Exceptions element is a list of domains that are excluded from decryption and inspection. The HTTPS Inspection Exceptions can be specified in the Protocol Parameters of a custom HTTPS Service, which is used in the Access rules to select HTTPS traffic for inspection.

The Access rules define which traffic is decrypted and inspected. You can select specific traffic for decryption and inspection, or you can enable the decryption and inspection of all TLS traffic.

When a certificate for client or server protection has been uploaded to the engine, it is possible to unintentionally enable TLS decryption for all traffic in one of the following ways:

- Adding a Network Application that allows or requires the use of TLS to an Access rule
- Enabling the logging of Application information in the Access rules
- Enabling Deep Inspection in an Access rule with the Service cell of the rule set to ANY

TLS inspection configuration overview:

1) To configure server protection, create TLS Credentials elements.
2) To configure client protection, create Client Protection Certificate Authority elements.
3) (Optional) Define custom Trusted Certificate Authority elements in addition to the default system elements.
4) (Optional) To exclude certain domains from decryption and inspection, define a TLS Match element or an HTTPS Inspection Exceptions element.
5) Activate client protection or server protection in the properties of the engine and enable TLS inspection in the Access rules.

Default TLS inspection elements

The SMC has default elements that can be used for TLS inspection. The default elements cannot be edited, but can be duplicated and then edited.

- The Default HTTPS Inspection Exceptions element is an HTTPS Inspection Exceptions element that excludes domains used by the SMC and the engines from decryption and inspection. You cannot edit the Default HTTPS Inspection Exceptions element. If you have to make changes, you can duplicate the Default HTTPS Inspection Exceptions element and edit the copy.
- The default HTTPS (with decryption) Service element enables the decryption of HTTPS traffic that uses the default port 443, excluding the domains that are specified in the Default HTTPS Inspection Exceptions. You cannot edit the default HTTPS (with decryption) Service element. If you have to make changes, you can duplicate the HTTPS (with decryption) Service element and edit the copy.
- There are predefined Trusted Certificate Authority elements that represent the signing certificates of major certificate authorities. Default Trusted Certificate Authority elements are automatically added from dynamic update packages and cannot be edited or deleted. When client protection is used, the engine checks whether the certificate of an external server was signed by one of the Trusted Certificate Authorities. You can also create your own Trusted Certificate Authority elements to represent other certificate authorities that the engine should consider trusted.

Configure TLS inspection for server protection

If you want to inspect TLS traffic for which an internal server is the destination, you must create a TLS Credentials element to store the private key and certificate of the server.

The private key and certificate allow the engine to decrypt TLS traffic for which the internal server is the destination so that it can be inspected. The TLS Credentials element stores the private key and certificate of an

internal server. The certificate and the associated private key must be compatible with OpenSSL and be in PEM format. Make sure that the server's private key and certificate are accessible from the computer where you use the SMC Client.

When a TLS Credentials element is used in TLS inspection, the private key and certificate allow the Engine, IPS engine, Layer 2 Engine, or Virtual Engine to decrypt TLS traffic for which the internal server is the destination so that it can be inspected.

Create TLS credentials using an existing certificate

If you have an existing private key and certificate, create a TLS Credentials element and import the private key and certificate.

Before you begin

The private key and certificate must be saved in a location that you can access from the computer where you run the SMC Client.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.
2) Select **Certificates** > **TLS Credentials**.
3) Right-click **TLS Credentials**, then select **Import Private Key**.
4) In the **Name** field, enter a unique name.
5) Next to the **Private Key** field, click **Import**, then browse to the private key.
   If the private key is encrypted, you are prompted to enter the password.
6) Next to the **Certificate** field, click **Import**, then browse to the certificate.
7) (TLS Credentials for Web Access Server or SSL VPN Portal) Next to the **Intermediate Certificate** field, click **Import**, then browse to the certificate from an intermediate CA that was used to sign the server certificate.

   Note

   If you do not import a certificate from the intermediate CA, the Security Engine does not send the issuer CA to the client.

8) Click **OK**.

Create TLS credentials from a certificate request

If you do not have an existing private key and certificate, create a certificate request and self-sign it to create a TLS Credentials element.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.
2) Select **Certificates** > **TLS Credentials**.
3) Right-click **TLS Credentials** and select **New TLS Credentials**.
4) Complete the certificate request details.
   a) Enter a name for the certificate.
   b) In the **Common Name** field, enter the IP address or domain name of the server.
   c) Complete the remaining fields as needed.
   d) Click **Next**.
5) Select **Self Sign**.
6) Click **OK**.

Result

The TLS Credentials element is added to **Administration** > **Certificates** > **TLS Credentials**. The **State** column shows that the certificate has been signed.

Configure TLS inspection for client protection

Client Protection Certificate Authority elements are used to inspect TLS traffic between an internal client and an external server.

When an internal client makes a connection to an external server that uses TLS, the engine generates a substitute certificate that allows it to establish a secure connection with the internal client. The Client Protection Certificate Authority element contains the credentials the engine uses to sign the substitute certificate it generates. If the engine does not use a signing certificate that is already trusted by users' web browsers when it signs the substitute certificates it generates, users receive warnings about invalid certificates. To avoid these

warnings, you must either import a signing certificate that is already trusted, or configure users' web browsers to trust the engine's signing certificate.

Note

Traffic that uses TLS might be protected by laws related to the privacy of communications. Decrypting and inspecting this traffic might be illegal in some jurisdictions.

Create Client Protection Certificate Authority elements

Client Protection Certificate Authority elements contain the credentials the engine uses to sign the certificate it generates.

If you want to inspect TLS traffic between a client in the internal network and an external server, you must create a Client Protection Certificate Authority element.

You must configure users' browsers to trust certificates signed using the credentials in the Client Protection Certificate Authority element to avoid excessive warnings or error messages about invalid certificates.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.

2) Browse to **Certificates** > **Certificate Authorities** > **Client Protection Certificate Authorities**.

3) Right-click **Client Protection Certificate Authorities**, then select **New Client Protection Certificate Authority**.

4) Enter a unique name **Name**.

5) (Optional) Enter the **Validity time** (in minutes) for the substitute certificates the engine creates.

   Each substitute certificate expires at the end of the validity time, and the engine automatically generates a new certificate. This process might produce warnings or error messages in the users' web browsers. To avoid excessive warnings, define a sufficiently long validity time, for example, several hours.

   Note

   All fields except the **Name** and **Validity time** on the **General** tab are grayed out. The grayed out fields are always filled in automatically based on information contained in the certificate you generate or import, and you cannot change them.

6) On the **Certificate** tab, import an existing private key and certificate or generate a new private key and signing certificate.

7) Click **OK**.

Import private keys and signing certificates for client protection

If you already have a certificate authority that is trusted by users' web browsers, you can import its private key and signing certificate. The engine uses them when it signs the substitute certificates it creates.

Importing a private key and certificate removes the need to separately configure users' web browsers to trust the engine's signing certificate. You can also import a private key and signing certificate that you generated outside of the SMC even if you do not already have a certificate authority that is trusted by users' web browsers. The certificate and the associated private key must be compatible with OpenSSL and be in PEM format. Make sure that the private key and certificate are accessible from the computer where you use the SMC Client.

Note

The imported certificate must have the *pathLenConstraint* value of 1 or greater. The Engine uses the imported certificate to sign an intermediate certificate authority, which is then used to sign the actual server certificates used for the TLS inspection. If the *pathLenConstraint* value is 0, then the Engine cannot sign the intermediate certificate authority and the TLS inspection fails.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.
2) Select **Certificates** > **Certificate Authorities** > **Client Protection Certificate Authorities**.
3) Right-click **Client Protection Certificate Authorities**, then select **New Client Protection Certificate Authority**.
4) Click the **Certificate** tab.
5) Next to the **Private Key** field, click **Import**, then browse to the private key.
6) Next to the **Certificate** field, click **Import**, then browse to the certificate.

   If users' web browsers are not already configured to trust the certificate authority whose signing certificate you imported here, add it to the list of certificate authorities that are trusted by users' web browsers when you are finished configuring TLS inspection in the SMC.

Generate private keys and signing certificates for client protection

If you do not already have a private key and signing certificate for the engine, you can generate a private key and signing certificate.

When you generate a private key and signing certificate in the SMC, you must export the certificate. You must also add it to the list of certificate authorities that are trusted by users' web browsers.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.

2) Select **Certificates** > **Certificate Authorities** > **Client Protection Certificate Authorities**.

3) Open the **Client Protection Certificate Authority Properties** dialog box in one of the following ways:
   - Right-click **Client Protection Certificate Authorities**, then select **New Client Protection Certificate Authority**.
   - Right-click a Client Protection Certificate Authority element, then select **Properties**.

4) Click the **Certificate** tab.

5) Click **Generate**.

6) In the **Common Name** field, enter the certificate authority's common name.

7) From the **Public Key Length** drop-down list, select the length of the public key to be generated.

8) In the **Valid Until** field, enter the date and time until which the signing certificate is valid.

   Tip

   You can also select the date by clicking **Open Calendar**.

   By default, the signing certificate is valid for one year after the creations date and time.

9) Click **OK**.
   A private key and signing certificate are generated.

Export Client Protection CA certificates

To make the users' web browsers trust the engine's signing certificate, you must add the Client Protection CA certificate to trusted certificates.

If the users' web browsers are not configured to trust the engine's signing certificate, users receive warnings about invalid certificates. If you generated the signing certificate for client protection in the SMC, you must export the certificate. You must also add it to the list of certificate authorities that are trusted by the users' web browsers.

These instructions assume that you already have the **Client Protection Certificate Authority Properties** dialog box open.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.

2) Select **Certificates** > **Certificate Authorities** > **Client Protection Certificate Authorities**.

3) Open the **Client Protection Certificate Authority Properties** dialog box in one of the following ways:
   - Right-click **Client Protection Certificate Authorities**, then select **New Client Protection Certificate Authority**.
   - Right-click a Client Protection Certificate Authority element, then select **Properties**.

4) Click the **Certificate** tab.

5) Click the **Export** option for the **Certificate** field, then browse to the location where you want to save the file.
   When you are finished configuring TLS inspection in the SMC, add the exported certificate to the list of certificate authorities that are trusted by users' web browsers.

Activating TLS inspection

To activate TLS inspection, you must configure client or server protection on the engine and define the inspected traffic in Access rules. You might also need to create a custom HTTPS Service element.

In the Engine Editor, you specify the Client Protection Certificate Authority element for client protection, or the TLS Credentials element for server protection. Depending on the options you specify, you can configure only client protection, only server protection, or both client and server protection.



CAUTION

Uploading TLS Credentials or a Client Protection Certificate Authority elements to the engine might enable decryption of TLS traffic that is not excluded from TLS inspection. The following configurations might enable decryption of TLS traffic:

- Adding a Network Application that allows or requires the use of TLS to an Access rule
- Selecting the **Enforced** option for **Log Application Information** in the Access rules
- Enabling Deep Inspection in an Access rule if the **Service** cell contains a Network Application or a Service that does not include a Protocol Agent

If the default HTTPS (with decryption) Service element meets your needs, you can use the default HTTPS (with decryption) Service element in the Access rules without modification. You must create a custom HTTPS Service in the following cases:

- To enable decryption for HTTPS traffic that uses a different port
- To select a different HTTPS Inspection Exceptions element
- To log the URLs in matching traffic
- To change any of the other settings in the Service Properties

The Access rules define which traffic is decrypted and inspected. Access rules that enable Deep Inspection and use a custom HTTPS Service or the default HTTPS (with decryption) Service element select specific traffic for decryption and inspection. To enable the decryption and inspection of all TLS traffic, you enable Deep Inspection in an Access rule with the Service cell of the rule set to ANY. Traffic that matches the Access rule is decrypted and inspected in the same way as unencrypted HTTP traffic according to the Inspection rules.

Activating TLS inspection consists of the following steps:

1) Activate client protection or server protection and upload certificates to the engine.

2) (Optional) Define a custom HTTPS Service element and enable TLS Inspection in the Protocol Parameters.

3) Create Access rules to select specific traffic for decryption and inspection or enable decryption and inspection of all TLS traffic.

Activate TLS inspection on Security Engines

Depending on the elements you select in the engine properties, you can activate client protection alone, server protection alone, or client and server protection together.

CAUTION

Uploading TLS Credentials or a Client Protection Certificate Authority elements to the engine might enable decryption of TLS traffic that is not excluded from TLS inspection. The following configurations might enable decryption of TLS traffic:

- Adding a Network Application that allows or requires the use of TLS to an Access rule
- Selecting the **Enforced** option for **Log Application Information** in the Access rules
- Enabling Deep Inspection in an Access rule if the **Service** cell contains a Network Application or a Service that does not include a Protocol Agent

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click an engine element, then select **Edit <element type>**.

3) From the navigation pane on the left, select **Add-Ons** > **TLS Inspection**.

4) (For client protection) From the **Client Protection Certificate Authority** drop-down list, select a Client Protection Certificate Authority element.
   - To select an existing element, click **Select** and select the element.
   - To create an element, click **New**.

5) (For server protection) Click **Add**, then select one or more TLS Credentials elements and click **Select**.

6) Click **Save and Refresh** to transfer the configuration changes and upload the certificates.

Enable TLS inspection in a custom HTTPS Service

Create a custom HTTPS Service element and use it for TLS inspection.

The default HTTPS (with decryption) Service element enables the decryption of HTTPS traffic that uses the default port 443, excluding the domains that are specified in the Default HTTPS Inspection Exceptions. If the default HTTPS (with decryption) Service element meets your needs, you can use the default HTTPS (with decryption) Service element in the Access rules without modification.

You must create a custom HTTPS Service in the following cases:

- To enable decryption for HTTPS traffic that uses a different port

- To select a different HTTPS Inspection Exceptions element
- To log the URLs in matching traffic
- To change any of the other settings in the Service Properties

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Select **Other Elements** > **Services** > **TCP**.
   A list of TCP Services opens on the right.
3) Right-click the default **HTTPS (with decryption)** Service, then select **New** > **Duplicate**.
   The **TCP Service Properties** dialog box opens with the properties of the HTTPS Service.
4) Enter a unique **Name** for the custom Service.
5) (Optional) If traffic uses a different port than the default port 443, enter the port number in the first **Dst. Ports** field.
6) Click the **Protocol Parameters** tab.
7) (Optional) Click **Select** next to the **HTTPS Inspection Exceptions** field, then select an HTTPS Inspection Exceptions element.
8) (Optional) To log the URLs in matching traffic, select **Yes for Logging of Accessed URLs**.
9) Click **OK**.

Apply TLS inspection to traffic

Use the default or custom HTTPS Service element in Access rules to define which traffic is decrypted and inspected.

CAUTION

Uploading TLS Credentials or a Client Protection Certificate Authority elements to the engine might enable decryption of TLS traffic that is not excluded from TLS inspection. The following configurations might enable decryption of TLS traffic:

- Adding a Network Application that allows or requires the use of TLS to an Access rule
- Selecting the **Enforced** option for **Log Application Information** in the Access rules
- Enabling Deep Inspection in an Access rule if the **Service** cell contains a Network Application or a Service that does not include a Protocol Agent

To select specific traffic for decryption and inspection, you create Access rules that use a custom HTTPS Service or the default HTTPS (with decryption) Service element. To enable the decryption and inspection of all TLS traffic, you enable Deep Inspection in an Access rule with the **Service** cell of the rule set to **ANY**.

You must enable Deep Inspection in the Action options of the Engine Access rules to enable TLS inspection. Deep Inspection is enabled by default in the IPS, Layer 2 Engine, and Layer 2 Interface Access rules. Traffic that matches the Access rules for TLS inspection is decrypted and matched against HTTP Situations in the Inspection rules in the same way as unencrypted HTTP traffic. Any traffic that is allowed to continue by the Inspection Policy is re-encrypted and sent to its destination.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) (Client Protection) Add a rule with the following properties to select traffic from clients in the internal network for inspection.

Access rules for client protection

Source	Destination	Service	Action
The elements that represent clients in your internal network or **ANY**.	The elements that represent the HTTPS servers to which internal clients connect, or **ANY**.	Your custom HTTPS Service, the default **HTTPS (with decryption)** Service, or set to **ANY**.	Allow **Deep Inspection** selected in the Action options.

2) (Server Protection) Add a rule with the following properties to select traffic to internal servers for inspection.

Access rules for server protection

Source	Destination	Service	Action
The elements that represent the clients that connect to your HTTPS server, or **ANY**.	The elements that represent your internal HTTPS servers.	Your custom HTTPS Service, the default **HTTPS (with decryption)** Service, or set to **ANY**.	Allow **Deep Inspection** selected in the Action options.

3) Save and install the policy to start using the new configuration.

Define trusted certificate authorities for TLS inspection

If you are using client protection and users must connect to domains whose certificates are not signed by one of the default Trusted Certificate Authorities, define your own Trusted Certificate Authority element to represent it.

Trusted Certificate Authority elements represent the certificates that identify certificate authorities. When a client in the protected network connects to an HTTPS server, the engine checks whether the certificate authority that signed the server's certificate is one of the Trusted Certificate Authorities. If the certificate was signed by one of the Trusted Certificate Authorities, the engine makes a substitute certificate that matches the server's certificate. The engine then signs the substitute certificate with the Client Protection Certificate Authority signing certificate. If the server's certificate is not signed by a Trusted Certificate Authority, the engine makes a new self-signed certificate. In this case, users receive a warning that the issuer of the certificate is not trusted. In both cases, client protection continues to function normally.

When you define a CA as trusted, all certificates signed by that CA are considered valid until their expiration date.

Related tasks
Create Trusted Certificate Authority elements on page 155

Exclude traffic from decryption for TLS inspection

Some traffic is automatically excluded from decryption. You can also exclude traffic from decryption globally, add rules to exclude specific traffic from decryption, or create specific lists of domains to exclude from decryption.

Traffic to and from some servers that use TLS can contain users' personal information that is protected by laws related to the privacy of communications. Decrypting and inspecting this traffic might be illegal in some jurisdictions. Some connections or network applications also might not work correctly if the traffic is decrypted.

You can exclude traffic from decryption and inspection in several ways:

- Globally with a TLS Match element.
- For specific matching traffic with an HTTPS Inspection Exception element.
- For network applications that match the URL categories specified in the Private Data application usage tag. You can use the Private Data application usage tag in Access rules to prevent the decryption of all traffic that matches the specified URL categories.

  Note

  To use the Private Data application usage tag to exclude traffic from decryption, you must have a license for category-based URL filtering using the ThreatSeeker Intelligence Cloud service.

  For more information, see Knowledge Base article 18074.

In all cases, traffic to the specified domains is allowed to pass through the engine without being decrypted.

The Security Engine mainly matches the specified domains based on the server name information (SNI) in the TLS Client Hello packet before the server certificate is sent. You can also use category-based URL filtering to exclude all traffic in the selected URL Categories from decryption based on the SNI in the traffic.

Note

If the end user's browser does not use SNI, the traffic might be decrypted.

TLS Matches define matching criteria for the use of the TLS protocol in traffic, and allow you to prevent specified traffic from being decrypted. TLS Matches that deny decrypting are applied globally, even if the TLS Match elements are not used in the policy. However, TLS Match elements that are used in specific Access rules can override globally applied TLS matches.

In most cases, TLS Matches are the recommended way to prevent traffic from being decrypted and inspected. Globally excluding domains from decryption might also prevent some Network Applications from being detected in encrypted connections. In this case, you can use HTTP Inspection Exceptions exclude the domain from TLS inspection.

The **Decryption** option in the Allow Action Options in Access rules defines whether traffic that matches the rule is decrypted. To exclude specific traffic from decryption for TLS inspection, add the following type of Access rule:

Source	Destination	Service	Action
Source IP address	Destination IP address	One or more of the following Service elements: ■ HTTPS (with decryption) ■ HTTPS (SafeSearch with decryption) ■ A custom Service element that uses the HTTPS Protocol	Allow Decryption: Disallowed

HTTPS Inspection Exceptions are used in a custom HTTPS service to define a list of domains for which HTTPS traffic is not decrypted. The custom HTTPS service must be used in a rule, and only traffic that matches the rule is excluded from decryption and inspection. HTTPS Inspection Exceptions are primarily intended for backwards compatibility.

Starting from version 6.11 the Security Engine supports fetching the destination server certificates actively when inspecting TLS traffic. If the option **Active destination server certificate probing** is enabled in SMC and a client tries to open a TLS connection through the Security Engine to a destination server the Security Engine will first check whether it already has a cached copy of the server certificate available. If the certificate is not available in cache, it will open a separate connection to the destination server and fetch the certificate. If the Security Engine is successful in obtaining the server certificate it will cache the server certificate along with the server endpoint information for future TLS connections to the same server endpoint.

Globally exclude domains from decryption

TLS Match elements define matching criteria for the use of the TLS protocol in traffic, and allow you to prevent the specified traffic from being decrypted.

TLS Matches that deny decrypting are applied globally, even if the TLS Match elements are not used in the policy. However, TLS Match elements that are used in specific Access rules can override globally applied TLS matches.

In most cases, TLS Matches are the recommended way to prevent traffic from being decrypted and inspected. Globally excluding domains from decryption can also prevent some Network Applications from being detected in encrypted connections. In this case, you can exclude the domain from TLS inspection.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Browse to **Other Elements** > **TLS Matches**.
3) Right-click **TLS Matches**, then select **New TLS Match**.
4) In the **Name** field, enter a unique name.
5) Select **Deny Decrypting**.
6) From the **Match Certificate Validation** drop-down list, select **Validation succeeded**.

7) Click **Add**, then specify the **Matching Domains** to exclude from decryption.
We recommend adding the domain names that users access to guarantee that traffic is not decrypted. If no domains are specified, any connection for which validation succeeded is excluded from decryption.

8) Click **OK**.
Connections are excluded from decryption as specified in the TLS Matches.

Exclude domains from inspection of HTTPS traffic

The HTTPS Inspection Exceptions element is a list of domains that are excluded from decryption and inspection.

HTTPS Inspection Exceptions are used in a custom HTTPS service to define a list of domains for which HTTPS traffic is not decrypted. The custom HTTPS service must be used in a rule, and only traffic that matches the rule is excluded from decryption and inspection. HTTPS Inspection Exceptions are primarily intended for backwards compatibility.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Expand the **Other Elements** branch.

3) Right-click **HTTPS Inspection Exceptions**, then select **New HTTPS Inspection Exceptions**.

4) Enter a unique **Name**.

5) Click **Add**.

6) Enter the domain name or NetBIOS name, then click **OK**.

Note

The domain name or NetBIOS name must exactly match the domain names that users access. Otherwise, the domain is not excluded from decryption and inspection. Wildcard characters (*) are not allowed.

Related concepts

Example: Preventing decryption of private connections on page 1029

Active destination server certificate probing

Use the option **Active destination server certificate probing** to enable active fetching of server certificates for TLS connections.

The Security Engine now supports decrypting TLS 1.3 connections, prior to this version, the TLS 1.3 connections were downgraded to TLS 1.2 if decrypting was needed.

The Security Engine now supports active fetching of destination server certificates when inspecting TLS traffic, when the option **Active destination server certificate probing** is enabled on SMC.

If a client tries to open a TLS connection through the Security Engine to a destination server (“server endpoint”) the engine first checks whether it already has a cached copy of the server certificate available.

If the server endpoint is already known to the Security Engine it uses a cached copy of the server certificate. If the server endpoint is not yet known to the Security Engine, it opens an additional TLS connection to the destination server for fetching the server certificate.

Also, to be able to fetch the server certificate for a TLS connection through an inline interface pair the Security Engine must have an additional interface with a valid route to the server.

Depending on the currently active policy the client connection may be blocked until the Security Engine has finished processing the server certificate.

If the Security Engine was successful in obtaining the server certificate it will cache the server certificate along with the server endpoint information to be used for future TLS connections to the same server endpoint. The Server certificate cache timeout value determines how long the engine may rely on the cached certificate before it should be discarded.

The engine will also cache server certificates obtained from the server response during TLS Handshake processing for each unique server endpoint.

The major benefit of this is that when not decrypting, the Security Engine may not be able to observe the server certificate and needs to rely on an unreliable TLS SNI (“Server Name Indication”) and server address information for identifying the applications in the TLS connection. With support for active destination server certificate probing, reliability is improved by automatically fetching the server certificate for each new TLS connection.



Note

The server endpoints and their respective cached certificates are not shared between different virtual Security Engines, logical interfaces, or between the master Security Engine and virtual Security Engines.

Examples of TLS inspection

These examples illustrate some common uses for TLS Inspection and general steps on how each scenario is configured.

Example: server protection in TLS inspection

Company A's server offers HTTPS services to their customers. The administrators want to be able to detect and block attacks targeting the HTTPS server, even if the attacks are encrypted inside an SSL tunnel.

They decide to configure TLS Inspection to decrypt and inspect traffic to and from the HTTPS server.

The administrators do the following:

1) Create a TLS Credentials element and import the private key and certificate of the HTTPS server.

2) Select the TLS Credentials in the Engine Editor.

3) Create Access rules with the default HTTPS (with decryption) Service as the Service.

4) On Engines, use the Medium-Security Inspection Policy to look for attacks in HTTP traffic and check the HTTP traffic against the anti-malware signatures. On IPS engines, use the Inspection rules from the IPS Template to look for attacks in HTTP traffic.

5) Save and install the policy.

Example: client protection in TLS inspection

The administrators also want to detect and block -based attacks targeting the browsers of users in Company A's network to protect the workstations and internal networks.

In addition to searching for attacks, the administrators also want to enable malware scanning. However, the employees at Company A often use online banking services that are secured with HTTPS, and these connections should not be inspected. The administrators decide to configure TLS Inspection to detect and block attacks that are encrypted inside an SSL tunnel and use a TLS Match element to globally exclude the online banking domains from decryption and inspection.

The administrators do the following:

1) Create a Client Protection Certificate Authority element and generate a new certificate and private key. In their network environment, the administrators add the certificate of the Client Protection Certificate Authority element to the list of trusted certificate authorities in the users' browsers.

2) Enable TLS inspection and select the Client Protection Certificate Authority element in the Engine Editor.

3) Create a TLS Match element that prevents decryption when certificate validation succeeds for the domain names for the online banking sites that are excluded from decryption. Because the TLS Match is applied globally, the administrators do not have to use it in any specific rules.

4) Create Access rules with the default HTTPS (with decryption) Service as the Service.

5) On Engines, use the Inspection rules from the Medium-Security Inspection Policy to look for attacks in HTTP traffic and check the HTTP traffic against the anti-malware signatures. On IPS engines, use the Inspection rules from the IPS Template to look for attacks in HTTP traffic.

6) Save and install the policy.

Chapter 67

Setting up QUIC inspection

Contents

- QUIC inspection and how it works on page 1083
- Verify QUIC inspection settings on Security Engine on page 1083
- Examples of QUIC Inspection on page 1085

QUIC is a secure general-purpose transport protocol. QUIC combines encryption and transport layer data stream processing into one protocol, thereby, reduces latency and improves security.

Note

QUIC inspection is used with engine version 7.0 and later versions.

QUIC inspection and how it works

QUIC is an UDP based transport protocol. Compared to TCP / TLS, it gives faster connection setup and eliminates the head-of-line blocking that is present in TCP. The reason behind faster connection setup is the reduced number of handshake messages. QUIC uses TLS 1.3 for the handshake and generation of encryption keys. It takes only one round trip to establish a path for communication.

QUIC is an integral part of HTTP/3 protocol which is widely used by web applications.

Verify QUIC inspection settings on Security Engine

From Forcepoint Security Engine version 7.0 onwards QUIC protocol is always inspected and, by default, is matched for web traffic rules.

However, while upgrading from any previous versions if you observe in the logs that QUIC traffic is discarded by "inspection" facility, the following Security Engine configurations allow setting the QUIC protocol to match to the web traffic rules or not, and whether to discard QUIC traffic when TLS inspection rules require decryption for the traffic.

Decryption of QUIC traffic is not supported but discarding the QUIC traffic causes most of the standard web clients fall back to earlier versions of HTTP, for which decryption by TLS inspection is supported.

Dialog box	Verification steps
Engine Editor > Add-Ons > QUIC Inspection For more information, see the section *Engine Editor > Add-Ons > QUIC Inspection.*	Make sure that the following options are selected in the Engine properties, based on your requirements: ■ ***Include QUIC ports for Web Traffic*** This option determines if the QUIC port 443/UDP should be matched to the access rules for web traffic (for example, URL Categories/Lists, Network Applications). **Note:** If you unselect this option, then access rules allowing or blocking URL Categories/Lists, or Network Applications will not be matched for QUIC traffic. ■ ***Discard QUIC if TLS inspection is required by access policy*** If you select this option, then any TLS inspection rules matching the web traffic, causes QUIC traffic to be discarded. As a result, most web browsers fall back to earlier versions of HTTP, which can be decrypted.
UDP Service Group Properties dialog box For more information, see the section *Working with Service elements > Create Service Group elements > UDP Service Group Properties dialog box*	In UDP Service element, make sure that QUIC service parameter is selected in the **Protocols Parameters** tab and **Discard QUIC if TLS inspection is required by access policy** field is set to "No". **Note:** For networks that do not support QUIC inspection, the **Discard QUIC if TLS inspection is required by access policy** field is set to "Yes".
TCP Service Group Properties dialog box For more information, see the section *Working with Service elements > Create Service Group elements > TCP Service Group Properties dialog box*	As QUIC decryption is currently not supported, it is not recommended for decrypted TLS traffic to use QUIC. In such scenario, you can set the **Strip QUIC support from server replies** option to "Yes" in the **Protocols Parameters** tab for HTTPS Service.
Network Application Properties dialog box and **URL List Application Properties dialog box** For more information, see the following sections: ■ *Using Network Application elements > Getting started with Network Application elements > Default elements for network applications > Network Application Properties dialog box* ■ *Filtering URLs > Add URL List Applications to block or allow URLs > URL List Application Properties dialog box*	While creating a new custom Network Application or URL List Application, if QUIC is selected in the **Protocol** list, access rules containing URL lists, URL categories, and Network Applications inspect the QUIC traffic in a similar manner as HTTP/2 and HTTP/1.1 traffic.

Related reference
Engine Editor > Add-Ons > QUIC Inspection
UDP Service Group Properties dialog box
TCP Service Group Properties dialog box
Network Application Properties dialog box
URL List Application Properties dialog box

Examples of QUIC Inspection

These examples illustrate some common uses for QUIC Inspection and general steps on how each scenario is configured.

When TLS inspection is not configured
The administrator in Company A allows users to safely browse the internet; however, the content must be considered safe and approved. The administrator uses URL Categories, URL lists, and Network Applications for allowing traffic from Company A's network, and lets all other traffic discarded. The administrator did not configure TLS inspection for the traffic, so TLS traffic is let through without performing decryption in the Engine.

The administrator initially discarded the QUIC traffic, as a result the web browsers revert to using TLS when QUIC is not permitted. However, as QUIC provides desirable improvements over TCP based TLS and the Security Engine supports QUIC, the administrator want to enable web browsing using QUIC as well.

The administrator performs the following steps:

1) Navigate to **Add-Ons > QUIC Inspection** in the engine properties.
2) Select *Enable QUIC ports for Web Traffic*.
3) Unselect *Discard QUIC if TLS inspection is required by access policy*.
4) Save and refresh the policy.

When TLS inspection is configured
The administrator still uses URL Categories, URL lists, and Network Applications for allowing approved traffic, but has enabled TLS inspection in the Engine for a subset of the traffic. However, some TLS traffic is still let through without decryption. The administrator wants to make sure that for the traffic that needs to be decrypted, QUIC is discarded, as decryption is not yet supported for QUIC traffic in the Security Engine. As a result, web browsers revert to using TLS if QUIC is not permitted.

The administrator performs the following steps:

1) Navigate to **Add-Ons > QUIC Inspection** in the engine properties.
2) Select *Enable QUIC ports for Web Traffic*.
3) Select *Discard QUIC if TLS inspection is required by access policy*.
4) Save and refresh the policy.

Chapter 68

Forward traffic to a proxy service for external inspection

Contents

- Getting started with forwarding traffic on page 1087
- Create a Proxy Server element on page 1089
- Add Access rules to forward traffic on page 1089
- Add NAT rules to forward traffic on page 1090
- Example: Using Access rules to forward traffic on page 1091
- Example: Using NAT rules to forward traffic on page 1092

In addition to inspecting traffic on the Security Engine, you can transparently forward traffic to a proxy service in the cloud or on premises. For example, you can forward all HTTP and HTTPS traffic to the Forcepoint Web Security Cloud service.

Web Security Cloud is a service that offers real-time protection against advanced threats and data theft. You need a separate license to use the service. For more information, see https://www.forcepoint.com.

For more information and to learn about using EasyConnect services to forward traffic, see the document *How to forward web traffic from Forcepoint Network Security Platform to Forcepoint Web Security Cloud* in Knowledge Base article 10582. Also see the Forcepoint Web Security Cloud documentation at https://support.forcepoint.com/s/article/Documentation-Featured-Article.

Getting started with forwarding traffic

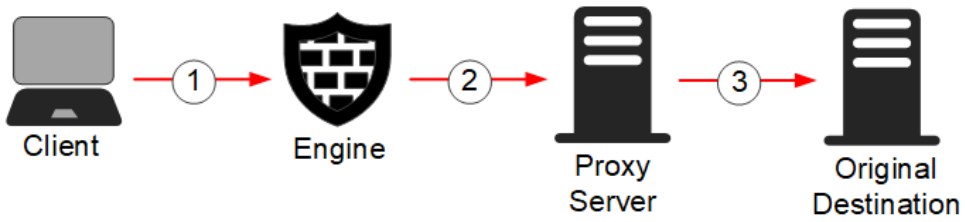
You can forward traffic to a proxy service where the traffic can be, for example, scanned for malware before the traffic continues to its final destination.

You can configure a Single Engine or Engine Cluster to forward FTP, SMTP, HTTP, and HTTPS traffic. The main benefit in using the engine to forward traffic to a proxy service is that the forwarding works transparently: the communicating hosts need no additional proxy configuration when the forwarding is done for them at the engine.

Proxies are typically used for malware scanning and content filtering, but there are more applications. Using an external service allows you to expand the capabilities of the engine with many other types of content screening. For example, the service can strip certain types of attachments out of emails without blocking the message itself. This type of feature is available directly on the engine as well, but an external service is a better option in medium to high throughput environments.

For the FTP or SMTP protocols, you cannot use inspection and forwarding. You must either forward the traffic or inspect the traffic locally. You cannot do both.

Forwarding traffic



1. Traffic from the client arrives at the Security Engine.
2. Access rules or NAT rules in the Engine policy determine which connections are forwarded to the proxy service for inspection.
3. The traffic is inspected and forwarded to the original destination. Reply packets are received with the IP address of the proxy service, so they are also forwarded to the proxy service.

Forwarding traffic configuration overview

Use Access rules or NAT rules to forward FTP, SMTP, HTTP, and HTTPS traffic to a proxy service, such as Forcepoint Web Security Cloud.

Elements in the configuration



1. The Proxy Server element is referenced in the Engine Policy.
2. The Engine Policy is installed on the Security Engine.

Follow these general steps to configure forwarding traffic:

1) Create the Proxy Server element that represents the proxy service.

2) Define the Access rules or NAT rules that select traffic for forwarding.

NAT considerations when using Access rules

We recommend that you use Access rules to forward traffic. However, if you have a more complex environment and existing NAT rules, forward traffic using the NAT rules method instead. When you use Access rules to forward traffic, all existing NAT rules in the policy are ignored, but element-based NAT is taken into account. All destination NAT definitions are ignored. If element-based source NAT definitions have been defined and if default NAT has been enabled in the properties of the Security Engine, those NAT definitions are processed.

Element-based NAT is sufficient in most cases, but if you need to use NAT rules to have greater flexibility, you must forward traffic using the NAT rules method. Use NAT rules if you want to, for example, forward traffic while using Outbound Multi-Link elements to select the network link for the traffic.

Create a Proxy Server element

Create a Proxy Server element that represents the proxy service.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Network Elements**.
2) Browse to **Servers**.
3) Select **New** > **Proxy Server**.
4) Configure the settings.
5) On the **Services** tab, configure the details of the service to which traffic is forwarded.
6) Click **OK**.

Related tasks

Create Location elements on page 127
Define contact addresses for Security Engines on page 138

Add Access rules to forward traffic

Define the connections that you want to forward in Access rules.

If the proxy service is in the cloud, a rule is needed to forward the matching traffic to the proxy service. If the proxy service is on premises, a rule is needed to forward the matching traffic to the proxy service and a second rule is needed to allow the connection from the proxy service to the original destination.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Browse to **Policies** > **Engine Policies**.
3) Right-click a policy, then select **Edit Engine Policy**.

4) Add a rule that forwards traffic to the proxy service.

Source	Destination	Service	Action
Original source address of the traffic. For example, clients in the internal network.	Original destination address of the traffic. For example, a web server.	The HTTP and TLS Network Application elements.	**Allow**. Action options: Proxy Server selected for the **Forward Traffic to** option.

5) (If the proxy service is on premises) Add a rule that allows traffic from the proxy service to the original destination.

Source	Destination	Service	Action
Proxy Server	Original destination address of the traffic. For example, a web server.	The HTTP and TLS Network Application elements.	**Allow**.

Make sure that you add this rule above the rule in the previous step to avoid potential loops if the proxy service is located in the same internal network as the clients.

Add NAT rules to forward traffic

For more complex environments where you already use NAT rules, use NAT rules to forward traffic.

If the proxy service is in the cloud, a rule is needed to forward the matching traffic to the proxy service. If the proxy service is on premises, a rule is needed to forward the matching traffic to the proxy service and a second rule is needed to allow the connection from the proxy service to the original destination.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Browse to **Policies** > **Engine Policies**.

3) Right-click a policy, then select **Edit Engine Policy**.

4) Add a rule that forwards traffic to the proxy service.

Source	Destination	Service	NAT
Original source address of the traffic. For example, clients in the internal network.	Original destination address of the traffic. For example, a web server.	The HTTP and TLS Network Application elements.	On the **Source translation** tab, select **Dynamic** as the **Translation Type**, then select the Outbound Multi-Link element that represents your public IP addresses. If you have only one IP address, click **Address**, then enter the address. On the **Destination translation** tab, select **Forward to Proxy** as the **Translation Type**, then select your Proxy Server element.

5) (If the proxy service is on premises) Add a rule that allows traffic from the proxy service to the original destination.

Source	Destination	Service	NAT
Proxy Server	Original destination address of the traffic. For example, a web server.	The HTTP and TLS Network Application elements.	On the **Source translation** tab, select **Dynamic** as the **Translation Type**, then select the Outbound Multi-Link element that represents your public IP addresses. If you have only one IP address, click **Address**, then enter the address.

Make sure that you add this rule above the rule in the previous step to avoid potential loops if the proxy service is located in the same internal network as the clients.

Example: Using Access rules to forward traffic

The example company has decided to screen HTTP and HTTPS connections using a proxy service hosted on premises.

The administrators have already installed the proxy on premises and configured it to process HTTP and HTTPS traffic according to the company's policy.

To configure the forwarding, the administrators:

1) Create a Proxy Server element to represent their proxy service.

2) Create an Access rule that forwards traffic to the proxy service.

3) Create another Access rule that forwards traffic from the proxy service to the Internet.

ID	Source	Destination	Service	Action
14.1	Proxy Server	ANY	The HTTP and TLS Network Application elements.	**Allow**.
14.2	Internal Network	ANY	The HTTP and TLS Network Application elements.	**Allow**. Action options: Proxy Server selected for the **Forward Traffic to** option.

Connections opened from the internal network are forwarded to the proxy in rule 14.2. The proxy then connects to the actual destination, which is allowed in rule 14.1. Rule 14.1 is higher in the rules to avoid potential loops if the Proxy Server is in the same internal network that is reference in rule 14.2

Example: Using NAT rules to forward traffic

The example company has decided to screen HTTP and HTTPS connections using a proxy service hosted in the cloud.

The company uses NAT rules to control application-specific link selection, to send traffic over high-quality network links if the traffic is critical to their business and to less expensive network links if the traffic is not critical.

The administrators have already set up the proxy and configured it to process HTTP and HTTPS traffic according to the company's policy. To configure the forwarding, the administrators:

1) Create a Proxy Server element to represent their proxy service.

2) Create a NAT rule that forwards traffic to the proxy service using a high-quality link if the traffic contains network applications that are critical to the business.

3) Create a second NAT rule that forwards traffic to the proxy service using another, low-cost link if the traffic contains non-critical network applications.

ID	Source	Destination	Service	Action
2.1	Original source address of the traffic. For example, clients in the internal network.	Original destination address of the traffic. For example, a web server.	Network Applications that are critical to the business. For example, Salesforce.	On the **Source Translation** tab, select **Dynamic** as the **Translation Type**, then select a high-quality Outbound Multi-Link element. On the **Destination translation** tab, select **Forward to Proxy** as the **Translation Type**, then select your Proxy Server element.
2.2	Original source address of the traffic. For example, clients in the internal network.	Original destination address of the traffic. For example, a web server.	Network Applications that are not critical to the business. For example, YouTube.	On the **Source Translation** tab, select **Static** as the **Translation Type**, then select a low-cost Outbound Multi-Link element. On the **Destination translation** tab, select **Forward to Proxy** as the **Translation Type**, then select your Proxy Server element.

Chapter 69

Block listing IP addresses

Contents

- Block listing traffic and how it works on page 1093
- Add Access rules for block listing on page 1096
- Configure automatic block listing of traffic on page 1097
- Block list traffic manually on page 1098
- Monitoring Block listing on page 1099

Block listing is a way to temporarily block unwanted network traffic either manually or automatically with block list requests from an Security Engine or Log Server. Engines, IPS engines, Layer 2 Engines, and Virtual Engines can use a block list for blocking traffic.

Block listing traffic and how it works

Block lists contain entries for blocking traffic temporarily based on traffic patterns that the engines detect or on administrator commands.

Block listing allows you to temporarily stop traffic:

- Without editing and installing policies (manual block listing only)
- Based on events detected by engines
- Based on correlation of detected events
- On a different engine than the one that detects an event
- On multiple engines with a single administrator command or a single detected event

Block listing makes it possible to block unwanted network traffic for a specified time. Engines can add entries to their own block lists based on events in the traffic they inspect. Security Engines and Log Servers can also send block list requests to other Security Engines. You can also block list IP addresses manually.

Example

A rule in the Inspection Policy detects a serious attack against a single host in your internal network. You can configure the rule to trigger automatic block listing of connections from that host to any other host in your internal networks.

Keep these limitations in mind when planning your block listing strategy:

- Layer 2 Engines can only block list IPv4 traffic.
- Engines and Layer 2 Engines do not enforce the block list by default. To enforce the block list, you must define the points at which the block list is checked in the Access rules.
- If a connection is allowed by a rule placed above the block list rule, the connection is allowed regardless of the block list entries.

Automatic block listing can have unintended consequences that could disrupt business-critical traffic. Use automatic block listing with careful consideration. The following two categories represent the typical risks associated with block listing:

Risks of block listing

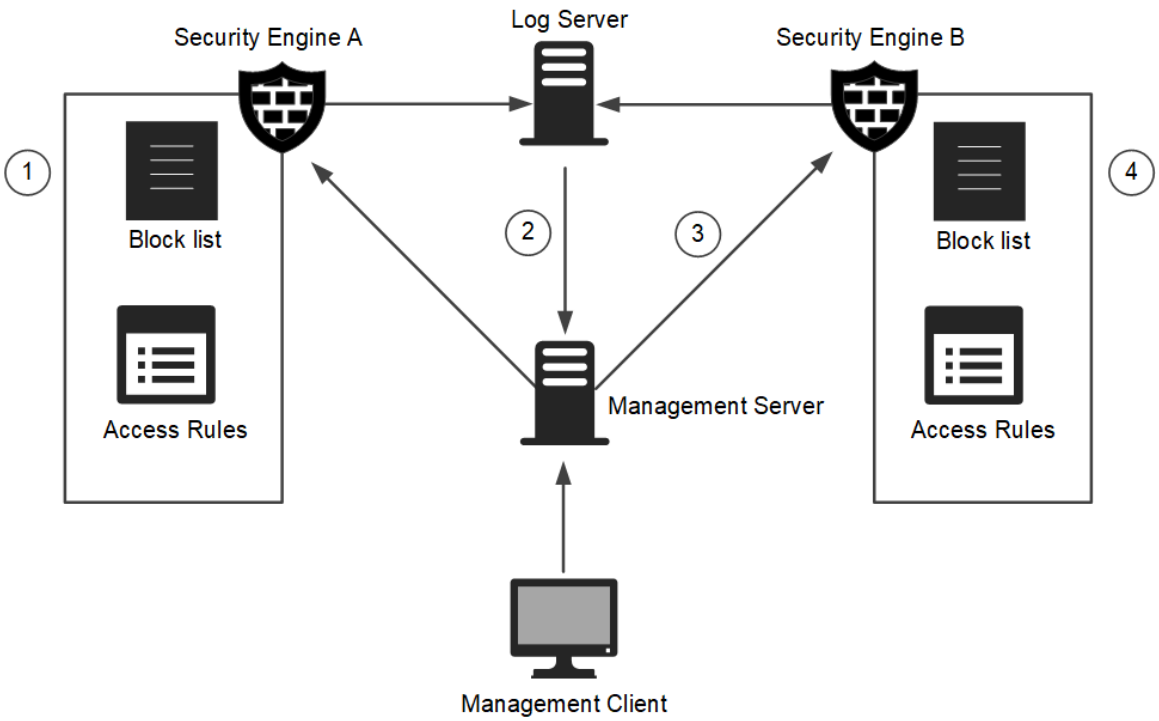
Risk	Explanation
Block listing legitimate connections (false positive)	If the defined pattern for detecting malicious traffic is inaccurate, legitimate traffic might sometimes be block listed. Block listing legitimate connections causes service downtime for hosts that are incorrectly identified as a source of malicious traffic.
Causing self-inflicted denial-of-service (DoS)	When an attacker uses spoofed IP addresses, a different (legitimate) IP address might be block listed instead of the attacker's IP address. Block listing spoofed IP addresses might cause a self-inflicted denial-of-service of legitimate traffic.

You can minimize these risks with good planning. Identify and evaluate the threats carefully before you configure block listing.

Block listing process

Block listing is executed as defined in the Access rules. Automatic block listing requests are sent as defined in the Inspection Policy.

Block listing process



1. Engines add entries to their own block lists for traffic they inspect.
   - There is one block list for each Engine, Layer 2 Engine, IPS engine, or Virtual Engine.
   - In engine clusters, there is one block list for each cluster. The nodes in the cluster exchange block list information in their synchronization communications.
2. Log Servers send block listing requests as a response to correlation of detected events. When one Security Engine sends a block listing request to another Security Engine, the Log Server relays the block listing request to the Management Server.
3. Management Servers relay manual block listing commands from administrators, and block listing requests sent by Log Servers to the Security Engines.
   There is no direct communication between different Virtual Engines or between Virtual Security Engines and the Management Server. For this reason, Virtual Engines cannot send block listing requests to other Virtual Engines.
4. Engines enforce the entries on their block lists according to their Access rules.
   - Each block list entry exists only for a defined duration, after which the entry is removed from the block list, and matching connections are again allowed. The duration of the blocking is defined when the block list entry is created.
   - Access rules check connections against the block list. If the IP addresses and ports in one of the block list entries match, the connection is discarded.
   - If the connection does not match a block listing Access rule or its related block list entries, the next Access rule in the policy is checked as usual.

Allow listing traffic

Allow listing means defining a list of IP addresses that must never be block listed.

Allow listing is implemented by following general Access rule design principles. Block listing applies only at the position of the block listing Access rules in the policy. Connections that have already been allowed or discarded before the block listing rules are not affected by block listing. If an Access rule allows a connection, an Access rule that refers to the block list further down in the policy cannot block list the connection.

Block listing configuration overview

Configuring block listing consists of several general steps.

Follow these general steps to configure block listing:

1) Define which traffic is matched against the block list in the Access rules.

2) (Automatic block listing only) Define the traffic that you want to block list automatically in the Exceptions in the Inspection Policy.

After block listing is configured, you can monitor the currently active block list.

Related concepts
Monitoring connections, block lists, VPN SAs, users, routing, SSL VPNs, and neighbors on page 236

Add Access rules for block listing

Access rules define which connections are checked against the block list.

By default, Engines and Layer 2 Engines do not enforce the block list. To enforce the block list, you must define the points at which the block list is checked.

The default High Security IPS Template and Medium Security IPS Template contain Access rules that apply the Security Engine's block list. If your IPS policy is based on these templates, it is not necessary to add Access rules for block listing. You can optionally add more Apply Block list rules with different matching criteria at different points in the policy.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Open the Engine, IPS, Layer 2 Engine, or Layer 2 Interface Policy for editing.

   Block list enforcement for Virtual Engines is configured in the Engine Policy, IPS Policy, or Layer 2 Engine Policy that is used on the Virtual Engine.

2) On the **IPv4 Access** or **IPv6 Access** tab, define which **Sources**, **Destinations**, and **Services** are compared with the block list.

3) Right-click the **Action** cell and select **Apply Block list**.

4) (Optional) Restrict which engines and servers are allowed to send block list requests.
   a) Right-click the **Action** cell and select **Edit Options**.
   b) On the Block listing tab, select **Restricted** for the **Allowed Block listers for This Rule** setting.
   c) From the **Available Block listers** list, select the elements that you want to add to the **Allowed Block listers** list and click **Add**.
      Add the Management Server to allow manual block listing through the SMC Clients. Add the Log Server to allow it to relay block listing requests from other Security Engines.
   d) Click **OK**.



Note

By default, engines are allowed to add entries directly to their own block lists for traffic they inspect.

5) Install the policy on the engine to activate the changes.

Next steps

No further configuration is needed if you want to block list connections manually.

Related concepts
Getting started with Access rules on page 831

Configure automatic block listing of traffic

Engines trigger automatic block listing based on the Block list Scope options in the Exceptions in the Inspection Policy.

Engines add entries directly to their own block lists for traffic they inspect. Engines can also send block listing requests to other Security Engines. In this case, the engine sends the block listing request to the Log Server. The Log Server relays the block listing request to the Management Server. The Management Server relays the block listing request to the other Security Engines that enforce the block listing.

Engines generate block list entries based on the patterns they detect in the traffic flow. The block list entry that is sent identifies traffic based on IP addresses and optionally the Protocol and port. The block list entries can include whole networks, even if the events that trigger them are related to a single source or destination IP address.

Automatic block list entries are created using the detected event's source and destination IP addresses, and optionally the TCP or UDP ports. If the event does not contain this information, a block list entry cannot be created. Netmasks can optionally be used to block list the detected event's network.

When the block list entry is created, the actions taken depend on the options you set. You can define Block listing scope options for any type of Exception, including rules that use Correlation Situations.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) In the navigation pane on the left, browse to **Policies** > **Inspection Policies**.
3) Right-click the Inspection Policy, then select **Edit Inspection Policy**.
4) On the **Exceptions** tab, add a rule, then specify the matching criteria for traffic that you want to block list.
5) Right-click the **Action** cell, then select **Terminate**.
6) Right-click the **Action** cell, then select **Edit Options**.
7) On the **Block list Scope** tab of the **Select Rule Action Options** dialog box, select **Override collected values set with "Continue" rules**. .
8) Select the type of Block list entry to create:
   - To create a Block list entry that terminates only the current connection using the default options, select **Terminate the Single Connection**, then click **OK**.
   - To block the traffic for defined duration and configure the settings, select **Block Traffic Between Endpoints**.
9) In the **Block list Executors** list, select the engines where the block list entry is sent, then click **Add**.
10) (Optional) To include the engine that detects the situation in the list of block list executors, select **Include the Original Observer in the List of Executors**.
11) Click **OK**.
12) Click **Save and Install**.

Block list traffic manually

You can block list traffic manually on Engines, IPS engines, and Layer 2 Engines.

For example, you can temporarily block a suspicious or disruptive source of communications while you conduct further investigations.

There are three ways to create new block list entries manually.

- Block list a connection found in the log data.
- Define a new block list entry for an Security Engine element.
- Create new block list entries in the **Block list** view, **Connections** view, **Monitoring** view, and **Logs** view.

The block list is not necessarily applied to all traffic. The Access rules determine how the block list is used.

Note

If a connection is allowed by a rule placed above the block list rule in the Access rules, the connection is allowed regardless of the block list entries. Check the logs to see which connections are discarded based on block listing.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Create a new block list entry in one of the following ways:
   - In the **Block list** view, **Connections** view, or **Logs** view — Right-click a row in the table and select **New Block list Entry** or **New Entry**.
   - To create a block list entry for a specific Security Engine — Right-click the Security Engine element in the **Connections** view, **Monitoring** view, or **Logs** view, and select **New Block list Entry** or **Block list** > **New Entry**.
2) Select the **Duration** for how long this entry will be kept.
   - If you leave the value as `0`, the entry only stops the current connections. Otherwise, the entry is enforced for the specified period of time.
3) Select the **Address** to block list for Endpoint 1 and Endpoint 2.
4) (Only if the protocol is TCP or UDP) Select the **Port** to block list for Endpoint 1 and Endpoint 2.
5) Select the **Block list Executors** that enforce the block list entry.
6) Click **OK**.
   The block list entry is sent to the executor and the traffic is blocked.

Monitoring Block listing

You can monitor active block listing entries in the **Block list** view.

The currently active block listing entries on the engine can be monitored in the **Block list** view. Block list monitoring does not show you which connections are dropped. Block list monitoring only shows you the IP addresses that are currently on the Block list. The **Logs** view can show which connections are dropped, depending on the logging options you have set. The Block list can be sorted and filtered in the same way as log entries.

Part X

Users and authentication

Contents

- Setting up directory servers on page 1103
- Setting up user authentication on page 1127

User accounts are stored in internal databases or external directory servers. You can use Forcepoint Network Security Platform in the Engine/VPN role or external authentication servers to authenticate users.

Chapter 70

Setting up directory servers

Contents

- Getting started with directory servers on page 1103
- Integrating external directory servers on page 1106
- Enabling access control by user on page 1113
- Defining user accounts on page 1119
- Add Users to User Group elements on page 1122
- Remove Users from User Group elements on page 1122
- Import user information on page 1123
- Export user information on page 1123
- Change user passwords on page 1124
- Remove the authentication settings of a user on page 1124
- Reset a engine's local user database on page 1125
- Set user database replication on or off for Engines and Master Engines on page 1125
- Examples of Directory Servers on page 1126

A directory server provides access to information about user accounts in a user database. Both internal and external directory servers can be used. Directory servers can be used for user authentication with Forcepoint Network Security Platform in the Engine/VPN role.

Getting started with directory servers

A directory server is a server that contains a user database that is queried during the user authentication process.

You can store the user accounts in the Management Server's internal user database, or on an external directory server. Different users can be stored in different directories. Authentication is based on the user information, but is a separate operation and is not necessarily done on the same server that stores the user information.

You can optionally use an integrated external Active Directory Server with the Forcepoint User ID Service or the Integrated User ID Service to provide transparent user identification for access control by user. Access control by user allows the use of Active Directory users as the source and destination of rules. The Integrated User ID Service is primarily meant for demonstration purposes and proof-of-concept testing of user identification services.



Note

For Forcepoint Network Security Platform version 6.4 or higher, we recommend that you use the Forcepoint User ID Service.

To implement user authentication, you must define user accounts either in the internal user database or on an external directory server.

If you use an external third-party authentication server and do not need to define different access rights for different users, it is not necessary to integrate an external directory server with the SMC. You can create a special

User element with the name *external* in the internal user database to represent any user that authenticates using the external authentication service.

Limitations

- The internal LDAP user database does not allow external authentication servers to query user information.
- The internal LDAP database limits the length of the User and User Group DN (distinguished name) to a maximum of 254 characters. Check the restrictions of external LDAP servers from the external server's documentation.
- If administrative Domains are configured, the internal user database is always in the Shared Domain. The user accounts stored in the internal database are also always in the Shared Domain. If you want to limit the visibility of end-user accounts, you must configure external LDAP databases separately for each Domain.
- User authentication is only supported on Engines. User authentication is not supported on layer 2 physical interfaces on Engines.

What do I need to know before I begin?

- The Management Server has an internal LDAP user database.
- Alternatively, you can use external LDAP user databases (including Active Directory).
- Different users can be stored in different databases.

Depending where user information is stored, different authentication options are available. The following table explains the possible combinations of internal and external directory servers and authentication servers:

Combinations of internal and external directory servers and authentication servers

	Internal authentication server	External authentication server
Internal directory server	User and User group information are maintained in the Management Server's internal user database. User and User Group information can be managed using the SMC Client and can be used for creating rules. Authentication can be done with password, IPsec certificate, or preshared key.	A second, external user database is required because the external authentication server has no access to the internal database. The same user information must be maintained separately in the Management Server's internal user database and in the external user database. User and User Group information can be used for creating rules. Any authentication method supported by the external authentication server can be used.

	Internal authentication server	External authentication server
External directory server	The Management Server is defined as an LDAP client for the external directory server. User and User Group information is shown in the SMC Client, and can be used for creating rules. Authentication can be done with password, IPsec certificate, or preshared key.	If you define Security Engines as LDAP clients for the external directory server, the Security Engine can send the user name and password to the external directory server for authentication. The external directory server checks the user name and password against the user's credentials in the external directory server, then replies to the Security Engine whether authentication succeeded or failed. You can optionally define the Management Server as an LDAP client for the external directory server. You can also duplicate and manually maintain the same user information separately in the Management Server's internal user database and in the external user database. Otherwise, you can create a single User element named *external* to represent all externally stored users. In this case, it is not possible to create different rules for different externally stored users. Each authentication rule includes all external users. There can be several rules, but any user that can authenticate in one rule can also authenticate when any of the other rules is triggered. Any authentication method supported by the external authentication server can be used.

Related concepts
Getting started with user authentication on page 1127

How internal user databases work

The Management Server includes an integrated LDAP directory for storing user information.

The Management Server's internal user database can be used for authenticating users with passwords. Using an internal LDAP directory is the simplest choice when there is no specific need to have an external LDAP server.

When the Management Server's internal LDAP directory is used, the user and user group information is stored on the Management Server. Each engine node stores a replica of the user database, and any changes to the main database are replicated immediately to the engines. This way, the engines can access their local directories instead of constantly communicating user information over the network.

Note

It is not possible to give external components (such as external authentication servers) access to the Management Server's internal LDAP directory.

If Domain elements have been configured, the Internal LDAP directory belongs to the Shared Domain. This means that the administrators who log on to some other Domain are allowed to view the contents of the Internal

LDAP directory. If all user information should not be available to administrators in all Domains, you must use an external LDAP directory in each Domain.

Directory server configuration overview

User information is stored in an internal or an external LDAP (Lightweight Directory Access Protocol) directory.

The standard LDAP user management model consists of three different levels: LDAP domains, user groups, and users. All three levels are represented as elements in the SMC Client.

Follow these general steps to configure the directory server:

1) (Optional) Integrate an external LDAP directory server:
   a) Define an LDAP Server or Active Directory Server element. The Active Directory Server element also contains the settings for NPS authentication; generic LDAP Server elements define a directory server only.
   b) (Optional) Define an LDAP Domain.
2) Define the User Group and User information:
   a) Import existing user information from some other Management Server.
   b) Create new accounts or edit accounts stored in an external LDAP database.

Integrating external directory servers

You can use an external directory server to store user group and user information instead of or in addition to the internal user database.

The external directory server can be an LDAP server, or a Microsoft Active Directory server that provides LDAP services.

You can use an external directory server without integrating it with the SMC components. You can view user information and use it for authentication against an external authentication service simply by allowing the SMC components to connect to the LDAP database.

The Management Server and the Security Engines each use their own integrated LDAP client to query the external LDAP directory directly. The external LDAP directory is not replicated into the internal directory on the Management Server or into the local directory of the Security Engines. Instead, the external LDAP directory is queried separately each time by the Security Engines each time a user attempts to authentication. The external LDAP directory is also queried separately Management Server when you view the User elements in the SMC Client.

You can configure access to the directory server for both the Management Server and the Security Engines, or for the Security Engines only. To take full advantage of user authentication features, we recommend configuring access to the directory server for both the Management Server and the Security Engines.

Configuring access to the external directory server for both the Security Engines and the Management Server allows the following:

- There is no need to manually duplicate user account information. User and User Group elements are automatically added to the SMC from the external directory.

- Externally stored user accounts are shown in the SMC Client and can be used to create different rules for different users.
- In most cases, users can be also added, removed, and edited through the SMC Client.
- Internal authentication methods can be used to authenticate externally stored users.

If only the Security Engines can access the external directory server, the following restrictions apply:

- You can authenticate externally stored users only against authentication methods provided by an external authentication server. Internal authentication methods are not available for externally stored users.
- A single element (User element named *external*) is used to represent all externally stored users in the Engine Policy. It is not possible to create different rules for different externally stored users.

Configuring schema files on external directory servers

A schema file defines the attributes (individual pieces of data) that an account can contain.

Updating the external server's schema with SMC-specific attributes is optional. Updating the schema also allows you to add SMC-specific information to Users and User Groups through the SMC Client.

You must update the schema file in the following cases:

- To be able to configure authentication requirements for specific Users or User Groups. Otherwise, you can configure authentication only at the LDAP domain level.
- To use the **User password** authentication method to authenticate users using user names and passwords. Alternatively, you can use the **LDAP Authentication** authentication method to authenticate users using user names and passwords without updating the schema.
- To be able to edit information in the LDAP directory through the SMC Client.

The method of configuring Schema files varies depending on which LDAP server you are using. The Schema update is done outside the SMC Client. In general, the schema update means that you add the SMC-specific attributes to the existing user information on the external LDAP server. These include attributes for the SMC-specific user name, password, and allowed authentication methods for the user.

Create LDAP Server or Active Directory Server elements

If you want to use an external directory server, you must create an element to define the parameters for contacting the server and the LDAP object classes.

You define the directory parameters in an LDAP Server element. If you are using an Active Directory server, you can define both the directory and the authentication parameters in the same Active Directory Server element. If you use administrative Domains, use a separate external LDAP Server or Active Directory Server in each Domain to create user accounts that are specific to each Domain.

Create Active Directory Server elements

The Active Directory Server element contains both the user database and authentication service options for using a Microsoft Active Directory server to store and authenticate users.

The settings in the Active Directory Server element can optionally include an account that the Management Server and Security Engines use to connect to query the directory. The user account must exist in the Active Directory Server's user database. Make sure the account you use has the permissions to manage other user accounts. The Security Engines and Management Server can also connect to query the directory without a user name and password if the Active Directory Server allows anonymous bind requests.

There are also values for some attributes that the Management Server and Security Engines look for in the directory. We recommend that you do not use special characters or accented letters in the Distinguished Names or user ID attributes. Currently, Active Directory has a limit of 160 characters for the Base DN, 24 characters for a UID (user ID) and 64 characters for the OU (organizational unit).

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **User Authentication**.

2) Right-click **Authentication Servers**, then select **New** > **Active Directory Server**.

3) On the **General** tab, configure the general settings.

4) (Optional) On the **Advanced** tab, add one or more secondary IP addresses for the server.

   These IP addresses are included when the Active Directory Server element is used in rules, but SMC components never use the addresses when initiating contact with the server.

5) (Optional) If you want to enable the Integrated User ID Service on an Security Engine and use the Active Directory Server in the Integrated User ID Service configuration to monitor logon information from Domain Controllers or Exchange Servers, configure the settings for the Domain Controllers and Exchange Servers on the **Monitored Servers** tab.

   Note

   The Integrated User ID Service is primarily meant for demonstration purposes and proof-of-concept testing of user identification services.

   Note

   The Security Engine uses the BIND user that is configured in the Active Directory Server element properties to monitor logon events.

6) (Optional) If you want the Active Directory Server to be monitored by the Log Server, configure the monitoring settings on the **Monitoring** tab.

7) Click **OK**.

Related concepts

Define contact IP addresses on page 127

Related tasks
Create Location elements on page 127
Activate monitoring of third-party devices on page 277

Create LDAP Server elements

The LDAP Server element can be used to configure access to any LDAP server as a user database for the Engines and the Management Server.

The settings in the LDAP Server element can optionally include an account that the Management Server and Security Engines use to connect to query the directory. The user account must exist in the LDAP Server's user database. Make sure the account you use has the permissions to manage other user accounts. The Security Engines and Management Server can also connect to query the directory without a user name and password if the LDAP server allows anonymous bind requests.

There are also values for some attributes that the Management Server and Engines look for in the directory. We recommend that you do not use special characters or accented letters in the Distinguished Names or user ID attributes. For character limits for these settings, see the documentation of your LDAP server.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **User Authentication**.
2) Right-click **Authentication Servers**, then select **New** > **LDAP Server**.
3) On the **General** tab, configure the general settings.
4) (Optional) On the **Advanced** tab, add one or more secondary IP addresses for the server.
   These IP addresses are included when the LDAP Server element is used in rules, but SMC components never use the addresses when initiating contact with the server.
5) (Optional) If you want the LDAP Server to be monitored by the Log Server, configure the monitoring settings on the **Monitoring** tab.
6) Click **OK**.

Related concepts
Define contact IP addresses on page 127

Related tasks
Create Location elements on page 127
Activate monitoring of third-party devices on page 277

Add LDAP object classes

If your Active Directory or LDAP server has LDAP object classes that are not defined in the SMC by default, add those object classes to the LDAP Object classes.

This way, the existing classes on the Active Directory or LDAP server can also be used for authentication.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **User Authentication**.
2) Right-click **Authentication Servers** and select **New** > **Active Directory Server** or **New** > **LDAP Server**, or right-click an existing Active Directory Server or LDAP Server element, then select **Properties**.
3) Click the **Object Classes** tab.
4) Enter the name of the **User Object Class** or **Group Object Class** and click **Add**.
   The object class appears in the list.

Configure LDAP attribute mapping

On the **Attributes** tab, you can define how attributes in the Active Directory or LDAP directory are mapped to user properties in the SMC.

You might need to change or fill in these values according to the server's configuration. Enter the same values for the attributes that are defined in the Active Directory Server's or LDAP Server's schema file.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **User Authentication**.
2) Right-click **Authentication Servers** and select **New** > **Active Directory Server** or **New** > **LDAP Server**, or right-click an existing Active Directory Server or LDAP Server element and select **Properties**.
3) Click the **Attributes** tab.
4) Configure the Attribute settings.

Related concepts

Configuring schema files on external directory servers on page 1107

Add authentication methods to LDAP Server or Active Directory Server elements

Authentication Methods specify the allowed authentication methods for the users stored on the Active Directory or LDAP server.

You can optionally use LDAP authentication for simple password authentication against LDAP database on the external directory server where user accounts are stored. When users authenticate to the Security Engine, the Security Engine sends the user name and password to the external directory server for authentication. The external directory server checks the user name and password against the user's credentials in the directory, then responds to the engine whether authentication succeeds or fails.

Note

Because the user name and password are sent through the LDAP connection, we recommend using LDAPS or Start TLS when you use LDAP Authentication.

You can optionally use the Internet Authentication Service (IAS) in previous Windows Server versions or the Network Policy Server (NPS) in Windows Server 2008 to authenticate end users. You must configure the IAS/NPS as a RADIUS server, and define each Engine that authenticates users as a separate RADIUS client for IAS/NPS. Use the NDI addresses when you define Engine Clusters as RADIUS clients for IAS/NPS. The IAS/NPS must have access to user information in the Active Directory. The user accounts must have remote access permissions. Set up the IAS/NPS as explained in the Microsoft Server documentation. The SMC does not support the Message-Authenticator attribute option available in the IAS/NPS, and is not NAP-capable. Only PAP authentication is supported.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) On the **Authentication** tab of the Active Directory Server or LDAP Server properties, configure the settings according to the type of server.

2) Click **OK**.

Related tasks

Define Authentication Method elements for external servers on page 1135

Add monitored servers to Active Directory server elements

To use an Active Directory Server in the configuration of the Integrated User ID Service on an Security Engine, add the Domain Controller servers and Exchange Servers from which the Integrated User ID Service receives information about users' IP addresses as monitored servers in the Active Directory server element.

Note

The Integrated User ID Service is primarily meant for demonstration purposes and proof-of-concept testing of user identification services.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) On the **Monitored Servers** tab of the Active Directory Server properties, click **Add**. The Monitored Server properties dialog box opens.

2) Select **Domain Controller** or **Exchange Server** as the **Server Type**, then configure the server settings.

3) Click **OK**.

Define LDAP domain elements

If you use an external LDAP directory for user management, you must create an LDAP Domain.

After the LDAP Domain is associated with the external server, the Management Server contacts the LDAP directory server or Active Directory Server. You can then view and edit users and user groups through the SMC Client.

Note

If you use the Management Server's internal user database, the users and user groups are always stored and managed in the default **InternalDomain** LDAP Domain.

You can select one LDAP Domain as the global **Default LDAP Domain**. You can also specify the default LDAP domain for each Security Engine in the Engine Editor. Selecting a default LDAP domain allows users belonging to that LDAP Domain to authenticate without specifying the LDAP Domain information. Users in other LDAP Domains must specify their LDAP Domain whenever they authenticate themselves.

If you use administrative Domains, create a separate LDAP Domain in each administrative Domain to create user accounts that are specific to each Domain. You can also use LDAP Domains in different administrative Domains to point to different parts of the directory hierarchy in the same LDAP directory. The internal LDAP directory is always in the Shared Domain, which makes its contents visible in all administrative Domains. You can select one Default LDAP Domain in each administrative Domain. You can also select an LDAP Domain in the Shared Domain as the Default LDAP Domain for all administrative Domains.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **User Authentication**.

2) Right-click **Users** and select **New External LDAP Domain**.

3) In the **Name** field, enter a name for the LDAP Domain.

4) Select **Default LDAP Domain** if this LDAP Domain is used for all authentication unless otherwise specified in the IPv4 or IPv6 Access rules.

Note

If the LDAP Domain you are creating is not the default LDAP Domain, users must type in the domain name when they authenticate.

Only one LDAP Domain can be the default LDAP Domain. The previous default LDAP Domain is automatically deselected.

5) Select a server, then click **Add** to bind the LDAP Server to the LDAP Domain.

6) (Optional) On the **Default Authentication** tab, click **Select** to define the allowed authentication methods for all accounts in this LDAP Domain.

Tip

You can override the default setting by selecting different authentication methods in the User Group or User properties.

We recommend that you set a default authentication method. If the authentication method is not defined yet, you can return to this dialog box to complete the configuration after you create the authentication method.

7) Click **OK**.

Related tasks

Activate monitoring of third-party devices on page 277

Enabling access control by user

Access control by user lets you use User and User Group elements as the source or destination of rules to create user-specific rules without user authentication.

You can use user-specific rules and user authentication rules to allow some user groups to access a service, while otherwise requiring authentication for the same service.

Note

User-specific rules do not replace user authentication. User-specific rules are a tool to simplify the configuration of access control, and improve the end-user experience by allowing transparent access to services. They are intended to be used for trusted users in a trusted environment where strong authentication is not required.

Access control by user requires the Forcepoint User ID Service or the Integrated User ID Service for transparent user identification. The Forcepoint User ID Service and the Integrated User ID Service monitor logon events from the Domain Controller servers and from Microsoft Exchange Servers to associate users with IP addresses. The Integrated User ID Service is primarily meant for demonstration purposes and proof-of-concept testing of user identification services.

Note

For Forcepoint Network Security Platform version 6.4 or higher, we recommend that you use the Forcepoint User ID Service.

Related concepts

Getting started with editing policies on page 885

Integrate Forcepoint User ID Service with Forcepoint Network Security Platform

Integrating Forcepoint Network Security Platform with Forcepoint User ID Service provides transparent user identification for access control by user.

Before you begin

You have installed and configured the components that send the user, group, and IP address information to the Security Engines. For information about integrating the Forcepoint User ID Service with other Forcepoint products, see the document *How to integrate Forcepoint User ID Service with other Forcepoint products* and Knowledge Base article 14100.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the SMC Client, create a Forcepoint User ID Service element.
   a) Select **Engine Configuration**.
   b) Browse to **Other Elements** > **Engine Properties** > **User Identification Services**.
   c) Right-click **User Identification Services**, then select **New** > **Forcepoint User ID Service**.
   d) In the **Name** field, enter a unique name for the Forcepoint User ID Service element.
   e) In the **IP Addresses** field, enter the IP address of the server on which the Forcepoint User ID Service is installed.
   f) Enter the port for communication between the Forcepoint Network Security Platform and the Forcepoint User ID Service server.

      Note

      The default port number is 5000. Use the same port that is used in the Forcepoint User ID Service configuration on the Forcepoint User ID Service server.

   g) In the **Monitored User Domains** section, click **Add** to define an Active Directory domain from which the Security Engine receives user information.

2) Enable TLS protection for the communication from the Security Engine to the Forcepoint User ID Service server.

   a) On the **Certificate** tab, click **Select**, then select a TLS Profile.

   Note

   The minimum supported TLS version for Forcepoint User ID Service 2.1 and higher is TLS 1.2. The TLS Profile element must use TLS 1.2.

   b) From the **TLS Server Identity** drop-down list, select a TLS server identity.

   c) In the **Identity Value** field, enter a value for the TLS server identity.

   Note

   If the TLS server identity is **Distinguished Name**, **SHA-1**, **SHA-256**, **SHA-512**, or **MD5**, click **Fetch Certificate** to fetch the value of the TLS server identity from a certificate.

3) Enable the Log Server to receive log data from the Forcepoint User ID Service.

   a) On the **Monitoring** tab, select the Log Server that receives the log data from the Forcepoint User ID Service.

   b) (Optional) To receive status information from the Forcepoint User ID Service, select **Status Monitoring**, then select a probing profile.

   c) To receive log data from the Forcepoint User ID Service, select **Log Reception**, then select the logging profile.

   The logging profile defines in which log fields the log data from the Forcepoint User ID Service is stored. If you create a new Logging Profile element that uses the default settings, all the log data is stored in the **Syslog message** field.

4) Click **OK**.

   The Forcepoint User ID Service element is created.

5) Select a Forcepoint User ID Service element for Security Engines.

   a) Select **Engine Configuration**.

   b) Right-click an engine, then select **Edit <element type>**.

   c) Browse to **Add-Ons** > **User Identification**.

   d) In the **User Identification Service** list, select a Forcepoint User ID Service element.

   If the Forcepoint User ID Service element that you want to use is not listed, select **Select**, then select a Forcepoint User ID Service element.

   e) Click **Save and Refresh**.

Next steps

If you want the Forcepoint User ID Service server to authenticate the Security Engine with the Management Server's internal certificate authority, export the certificate of the Management Server's active internal certificate authority.

Related tasks

Create Trusted Certificate Authority elements on page 155
Create TLS Profile elements on page 157
Export certificate of the active internal certificate authority on page 155
Create Logging Profile elements on page 263
Create Probing Profile elements for monitoring third-party devices on page 275

Configure the Integrated User ID Service

You can use the Integrated User ID Service on the Security Engine to provide transparent user identification for access control by user. The Integrated User ID Service is primarily meant for demonstration purposes and proof-of-concept testing of user identification services.

Before you begin

- You have created an Active Directory server element and added the Domain Control servers and Microsoft Exchange Servers from which the Active Directory server receives information in the Active Directory Server properties.
- You have created an External LDAP Domain element and bound the Active Directory Server element that you created to the External LDAP Domain element.

Note

The Integrated User ID Service requires that the external authentication method of the Active Directory Server and the authentication method of the External LDAP Domain is user password or LDAP authentication.

Note

You cannot use the Integrated User ID Service with Virtual Engines.

Steps

1) Select **Engine Configuration**.
2) Browse to **Other Elements** > **Engine Properties** > **User Identification Services**.
3) Right-click **User Identification Services**, then select **New** > **Integrated User ID Service**.
4) Enter a name for the service.

5) In the **Active Directory Domain** field, select the External LDAP Domain from which the Integrated User ID Service element receives information about users, groups, and IP addresses.

Select the External LDAP Domain element to which you bound the Active Directory Server that you want to use in the Integrated User ID Service configuration. If several Active Directory Servers are bound to the External LDAP Domain element, the Integrated User ID Service uses the first Active Directory Server element that is listed in the External LDAP Domain element.

6) Enter the time range for the first query of user, IP address, and group information from the Security Engine to the Active Directory Server.

The time range for the first query defines how far back in time the Security Engine queries for the user, IP address, and group information. The time range for the first query must be between one minute and seven days.



Note

The Security Engine uses the defined time range for the first query only when the Integrated User ID Service starts after it has first been configured or when the Integrated User ID Service starts after the Security Engine has been rebooted.

7) Define how often the Security Engines polls for the user, group, and IP address information from the AD server.

8) (Optional) Define user names and IP addresses that the Integrated User ID Service does not monitor.

- User names to be ignored are typically user names that are associated with service accounts that do not represent actual users.
- IP addresses to be ignored typically represent multi-user servers such as terminal servers.
- If you define an entry that contains both a user name and an IP address, the entry matches only if both the user name and the IP address are detected.

9) Click **OK**.

The User ID Service element is created.

10) Enable the Integrated User ID Service on the Security Engine.

a) Select **Engine Configuration**.

b) Right-click an engine, then select **Edit <element type>**.

c) Browse to **Add-Ons** > **User Identification**.

d) In the **User Identification Service** list, select an Integrated User ID Service element.

If the Integrated User ID Service element that you want to use is not listed, select **Select**, then select the Integrated User ID Service element.

e) If the LDAP domain for the External LDAP Domain is not the default LDAP domain, browse to **Advanced Settings** > **Authentication** , then select **Allow lookup from known User Domain matching to client email domain or UPN suffix** to allow the Active Directory Server to query the user information from the external LDAP domain.

f) Click **Save and Refresh**.

11) On the Domain Controller servers and Exchange Servers that provide information about users' IP addresses to the Integrated User ID Service, configure permissions for the user accounts that are used to query the IP address information.



Note

The Security Engine uses the BIND user that is configured in the Active Directory Server element properties to monitor logon events.

a) Open the command prompt, type `wmimgmt.msc`, then press **Enter**.

b) Right-click **WMI Control**, then click **Properties**.

c) Switch to the **Security** tab.

d) Browse to **Root** > **CIMV2**.

e) Make sure that **Execute Methods**, **Remote Enable**, **Read Security**, and **Enable Account** are selected.

f) Click **OK**.

Result

The Integrated User ID Service is enabled on the Security Engine.

Select User Identification Service for Security Engines

You can select the User Identification Service for each Security Engine in the Engine Editor.

Each User Identification Service can be associated with one or more Security Engines, but only one User Identification Service can be selected for each Security Engine.



Note

The Integrated User ID Service is primarily meant for demonstration purposes and proof-of-concept testing of user identification services.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click the Security Engine for which you want to select a Forcepoint User ID Service or an Integrated User ID Service element, then select **Edit <element type>**.

3) In the navigation pane on the left, browse to **Add-Ons** > **User Identification**.

4) Select the User Identification Service element that represents the server with which this Security Engine communicates.



Note

For Forcepoint Network Security Platform version 6.4 or higher, we recommend that you use the Forcepoint User ID Service.

5) (Optional) Configure the additional settings.

6) Click **Save**.

Defining user accounts

User Group and User elements define the user account information for end users.

You can use User Group and User elements in Engine IPv4 and IPv6 Access rules to add a requirement for authentication. If you have enabled the Forcepoint User ID Service or the Integrated User ID Service on the Security Engine, you can also use User Group and User elements as the source and destination of Access, Inspection, and NAT rules without user authentication. The Integrated User ID Service is primarily meant for demonstration purposes and proof-of-concept testing of user identification services.



Note

For Forcepoint Network Security Platform version 6.4 or higher, we recommend that you use the Forcepoint User ID Service.

Options for adding user accounts

If you are using the Management Server's internal user database:

- If you have existing user accounts stored in an internal user database on another Management Server, you can export or import the information between the databases.
- Otherwise, you must create the User Groups and Users individually.

If you are using an external directory server:

- If the LDAP database is integrated with the Management Server, you can view the user information in the SMC Client. However, for the accounts to be valid in Access rules, you must configure at least one Authentication Method for the users. You can configure Authentication Methods as default settings for the LDAP Domain and for the User Groups and Users.
- If the LDAP database is not integrated with the Management Server, the user accounts are not shown in the SMC Client and are not available for configuration.

Related concepts

Enabling access control by user on page 1113

Related tasks

Define LDAP domain elements on page 1112
Import user information on page 1123

Create User Group elements

If you have many users, you might want to organize your users into several different User Groups.

You can organize the groups, for example, according to different services. A single user can belong to several groups at the same time. You must have at least one User Group element.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **User Authentication**.

2) Expand the branch of the LDAP Domain that represents the correct user database. The default InternalDomain LDAP Domain represents the internal database.

3) Right-click the parent group under the LDAP Domain (this is called **stonegate** for the internal database), then select **New** > **Internal User Group**.

4) In the **Name** field, enter a name for the Internal User Group.
   - The name is used as the common name (**CN**) for the group.
   - The distinguished name (**DN**) is inherited from the LDAP Domain to which this Internal User Group belongs.

5) (Optional) In the **Expiration After** field, enter a number of days after which the Internal User Group expires.
   When the Internal User Group expires, it stays in the system but is invalid and does not allow authentication until it is re-enabled.

6) Click the **Authentication** tab.

7) Click **Add**, then select one or more **Authentication Methods**.
   If you select several Authentication Methods, you can restrict the Authentication Methods allowed for each user in the User element properties and in Access rules that require authentication.

8) Click **OK**.

Create User elements

The User element defines who your users are and how they can identify themselves to get access to networks and services as defined in your Engine Access rules.

You create Users as members of a User Group. You do not have to specify all user parameters separately for each individual User. A User that is a member of a User Group can inherit, for example, the Authentication

Method and account expiration time from the User Group. Each User Group must belong to an LDAP Domain. We recommend creating a separate user account used for each user. Each user can belong to several User Groups within the LDAP Domain. User-specific properties can override properties defined at the User Group level.

You can import and export Users and User Groups through an LDIF file to or from some other Management Server.

Note

Although you cannot edit User Group memberships in the User element properties, each user can belong to several User Groups. After creating the User element, drag and drop it to other User Groups to add more group memberships.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **User Authentication**.

2) Browse to **Users**.

3) Add a user to a User Group in one of the following ways:
   - Right-click a User Group and select **New** > **Internal User** (for the internal **stonegate** parent group).
   - Right-click a User Group and select **New Internal User** (for a User Group under the internal **stonegate** parent group).

4) In the **Name** field, enter a unique name to identify the User in the directory.
   The name is used as the common name (CN) for the User. The distinguished name (DN) is inherited from the LDAP Domain to which the User belongs.

5) (External directory only) Enter additional user information.

6) (Optional) Change the Activation settings for the user account.

7) Click the **Authentication** tab.

8) Click **Add** to select the **Authentication Methods** for the user.
   - You can add more than one authentication method for each user. This way, you can put the User in more than one User Group when the User Groups have different authentication methods.
   - If you have not configured any Authentication Methods yet, you can create them in this dialog box.

9) Define the properties for the selected Authentication Method.

CAUTION

Use strong passwords that are at least eight characters long and that contain numbers, letters, and special characters. Do not base passwords on personal information such as names, birthdays, ID numbers, phone numbers, street names, registration plate numbers, or relatives' names.

10) Click **OK**.

Result

The user account is created. If the user is stored in the internal LDAP database, the information is automatically synchronized to the local databases on the Engines unless user database replication has been disabled.

Related concepts
Getting started with user authentication on page 1127

Related tasks
Import user information on page 1123

Add Users to User Group elements

When you add a User to a group, they remain a member of the current User Group and also become a member of the new group.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Drag and drop the User element to a User Group in the element tree.

Remove Users from User Group elements

If Users do not belong in a User Group, remove them.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Browse to the User Group from which you want to remove the User.

2) Right-click the User element, then select **More actions** > **Remove**.
   A confirmation message is shown.

3) To remove the user from the User Group, click **Yes**.

Import user information

You can import the user information from one Management Server's internal LDAP user database to another Management Server.

The internal LDAP user database is represented by the default LDAP Domain **InternalDomain**.

To be able to import user information successfully, the information must meet two conditions:

- The distinguished name (DN) must always be of type **dn**: cn=cn-of-the-user-or-group,dc=stonegate.
- All user groups must be directly attached under the dc=stonegate top-level group.

Note

The import feature is only meant for importing users exported from another Management Server's internal LDAP database. The import is not meant to work with other .ldif files.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Menu** > **File** > **Import** > **Import Users**.

2) Select the correct file, then click **Import**.
   A new tab opens, showing the progress of the import.

3) When the import is finished, check the messages for any warnings before you close the tab.
   If warnings were displayed, select **Engine Configuration**, then browse to **User Authentication** > **Users** > **InternalDomain** to check the properties for those Users or User Groups.

Export user information

You can export the user information stored in the internal user database on one Management Server to an .ldif file, to transfer it to another Management Server.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **User Authentication**.

2) Browse to **Users** > **InternalDomain**.

3) Right-click **InternalDomain** and select **Export Users**.

4) Select the location, enter the file name, and click **Export**.
   A new tab opens showing the progress of the export.

5) When the export is finished, check the messages for any warnings before you close the tab.

Change user passwords

If you use LDAP password authentication, you can reset and change a user's password as necessary.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **User Authentication**.
2) Browse to **Users** and then to the correct LDAP Domain and User Group.
3) Right-click the User element and select **More actions** > **Change Password**.
4) In the **New Password** and **Confirm New Password** fields, enter and confirm the password.
5) Click **OK**.
   The new password becomes valid immediately.

Remove the authentication settings of a user

If you need to revoke a user's access rights quickly, you can remove all authentication settings of a user without deleting the user.

For example, you might want to remove the authentication settings of a user if the user's account has been compromised.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **User Authentication**.
2) Browse to **Users**, then to the correct LDAP Domain and User Group.
3) Right-click the User element, then select **More actions** > **Clear Authentication Attributes**.
4) In the Confirmation message, click **Yes**.

Reset a engine's local user database

You can replace the engine's local copy of the user database with a copy of the internal user database on the Management Server.

If you use the internal user database of the SMC, the user information is stored centrally on the Management Server. When changes are made, they are incrementally replicated to each engine node to guarantee fault-tolerant authentication. If necessary, you can perform a full synchronization manually.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Single Engine or individual nodes of a Engine Cluster and select **Commands** > **Reset User Database**.

2) Click **Yes** in the Confirmation message.

   The node's local copy of the user database is replaced with a copy of the internal user database on the Management Server.

Set user database replication on or off for Engines and Master Engines

The internal user database on the Management Server is replicated automatically to Engines and Master Engines.

The engine uses the local copy of the internal user database to authenticate users without a connection to the Management Server. If you have a reason to do so, you can turn off the replication.

CAUTION

If you want to prevent users from authenticating, remove the authentication settings of a user instead of turning off the replication. Turning off the replication prevents any new users you add after the operation from authenticating, but might not prevent existing users from doing so.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click a Single Engine, Engine Cluster, or Master Engine and select or deselect **Options** > **User DB Replication**.

Related tasks

Remove the authentication settings of a user on page 1124

Examples of Directory Servers

These examples illustrate some common uses for Directory Servers and general steps on how each scenario is configured.

Example: using the internal user database

This scenario shows an example of separating users into groups for accessing different resources.

Company A has a general office network and a separate HR network for servers that contain HR information, such as employee records and payroll information. The servers already restrict which users have access. For auditing reasons, the administrators want to separate the users into groups and require authentication to access the HR network. The administrators:

1) Create a User Group “HR Users” in the InternalDomain and assign one of the default internal authentication methods.

2) Create User elements for each person with access rights under the HR Users group.

3) Define Access rules for user authentication on the engine.

Example: integrating Microsoft Active Directory Servers

This scenario shows an example of integrating Microsoft Active Directory servers.

For more information about configuring the Network Policy Server (NPS), see Microsoft’s documentation at https://technet.microsoft.com.

Company B has an existing Microsoft Active Directory server that stores user information. They decide to integrate this existing server’s directory services.

The administrators:

1) Define an Active Directory Server element.

2) Add the SMC-specific classes and attributes into the Active Directory server’s configuration to be able to fully manage the user accounts through the SMC Client.

3) Define the Management Server as an LDAP client for the Active Directory server.

4) Define the Engine as an authentication client for the NPS.

5) Add an LDAP Domain element for the Active Directory server in the SMC Client.

Chapter 71

Setting up user authentication

Contents

- Getting started with user authentication on page 1127
- Integrating external authentication services on page 1130
- Define Access rules for authentication on page 1136
- Enable browser-based user authentication on page 1137
- Configure client certificate authentication for browser-based user authentication on page 1145
- Authenticate to the Security Engine on page 1147
- Customize the User Authentication Pages for browser-based user authentication on page 1148
- Monitoring and testing user authentication on page 1150
- Examples of user authentication on page 1150

You can implement user authentication to control which resources different end users can access. You can use authentication as an access requirement in IPv4 Access and IPv6 Access rules in Engine Policies. You can use both internal and external user authentication servers.

Getting started with user authentication

User authentication means requiring the users to prove their identity before giving access to a network resource.

Authentication requires a user database that stores the user information and an authentication method that inspects credentials and grants or denies access.

You can use the following kinds of authentication methods:

- The engine's internal authentication methods
- Authentication methods provided by external RADIUS or TACACS+ authentication servers, such as NPS or RSA Authentication Manager (SecurID)
- LDAP authentication for simple password authentication against the LDAP database on an external LDAP server or Active Directory server

Alternatively, if strong authentication is not required, you can allow specific users to access services in a trusted environment without requiring user authentication.

User authentication proceeds as follows:

1) The user opens an authentication connection to the engine.

2) The engine checks if the user exists and which authentication method the user can use.

3) The user-supplied credentials are verified.
   - When you use the engine's internal authentication methods, the engine checks user credentials against its own replica of the user database.

- When you use authentication methods provided by an external server, the external server verifies the user's credentials, then responds to the engine whether authentication succeeds or fails.

4) If authentication succeeds, the engine lists the user as an authenticated user, taking note of both user name and authentication method.

5) When the user opens new connections, IPv4 and IPv6 Access rules that contain an authentication requirement can now match. The user name and authentication method are both separately checked as matching criteria.

6) When the configured timeout is reached, the authentication expires and the user is removed from the list of authenticated users. Access rules that require authentication no longer match the user's connections.

With user authentication, you can:

- Maintain separation of internal networks that have different security levels when the confidentiality of the information that is accessed does not need to be strictly enforced. For example, user authentication can provide an extra access control measure for applications that already exchange information securely.
- Allow secure and confidential access from any location to any internal resources for Forcepoint VPN Client users.
- Authenticate Administrator and Web Portal User logons.

The following limitations apply to user authentication:

- User authentication is only supported on Security Engine in the Engine/VPN role.
- User authentication is not supported on layer 2 physical interfaces on Engines.

Related concepts

Integrating external directory servers on page 1106
Enabling access control by user on page 1113

Related tasks

Authenticate administrators using RADIUS or TACACS+ methods on page 392

Default elements for user authentication

There are predefined Authentication Method elements for user authentication on the engine and for external user authentication.

There are four predefined Authentication Method elements for user authentication on the engine.

- **Client Certificate** is for certificate-based authentication.
- **LDAP Authentication** is for simple password authentication against LDAP databases on external LDAP or Active Directory servers.
- **Pre-Shared Key Method** is for use with some third-party VPN clients.
- **User Password** is for simple password authentication against the internal LDAP database, including user authentication in Forcepoint VPN Client hybrid authentication.

To use the engine for user authentication, you must use one of the predefined Authentication Method elements.

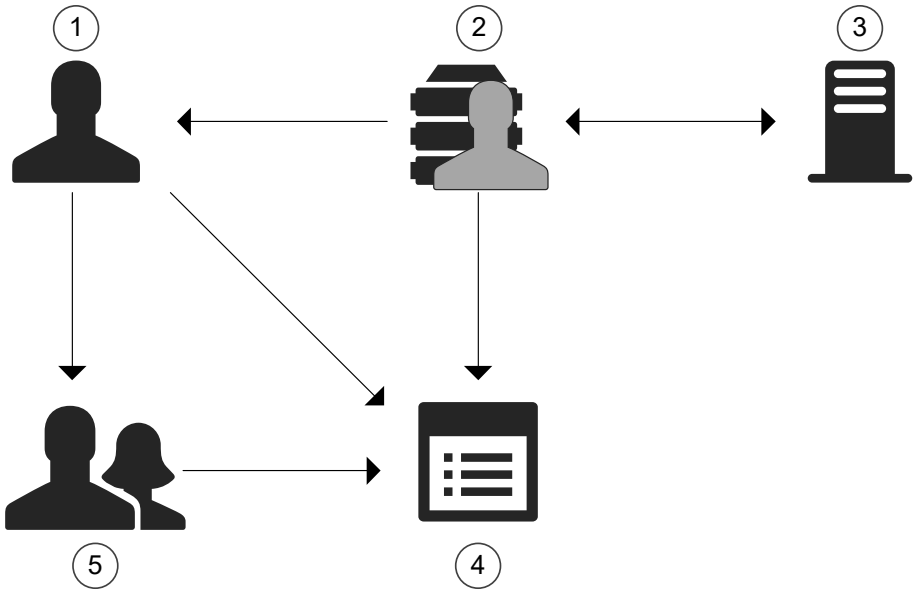
There are three predefined Authentication Method elements for use with RADIUS Authentication Server or TACACS+ Authentication Server elements.

- **Network Policy Server** is for use with an external Network Policy Server (NPS) server.
- **Pre-Shared Key Method** is for use with some third-party VPN clients.
- **User Password** is for simple password authentication against the internal LDAP database.

User authentication configuration overview

Authentication methods define the authentication method used by particular users and user groups.

Elements in the configuration



1. User
2. Authentication Method
3. Active Directory Server, LDAP Server, RADIUS Authentication Server, or TACACS+ Authentication Server
4. Engine Policy
5. User Group

External RADIUS or TACACS+ authentication servers are configured as RADIUS Authentication Server or TACACS+ Authentication Server elements. RADIUS or TACACS+ authentication servers can be located in any network that allows them to communicate with the engine that has an authentication rule in its policy. Authentication Method elements are associated with authentication servers to define the allowed authentication methods for the server, or the servers that use a particular authentication method.

Authentication Method elements define the allowed authentication methods for IPv4 and IPv6 Access rules, and for the Users and User Groups. Both User and User Group elements can be used in IPv4 and IPv6 Access rules to define rules that only match connections from specific, successfully authenticated users. A specific Authentication Method definition is needed in each rule especially when the Users and User Groups have several allowed Authentication Methods. Otherwise, the rules can allow any defined Authentication Method that is allowed for the included users.

Follow these general steps to configure user authentication:

1) (Optional) Create server elements and Authentication Method elements for external authentication services.

2) Add an authentication requirement to the relevant IPv4 or IPv6 Access rules.

3) (Forcepoint VPN Client authentication) Install the Forcepoint VPN Client software on the end users' computers. See the Forcepoint VPN Client documentation for more information.

4) (Browser-based Authentication) Configure the authentication prompt:
   - Enable end users to authenticate and re-authenticate using a browser-based authentication prompt.
   - Customize the authentication prompt.

Related concepts
Integrating external authentication services on page 1130

Related tasks
Define Access rules for authentication on page 1136
Enable browser-based user authentication on page 1137
Customize the User Authentication Pages for browser-based user authentication on page 1148

Integrating external authentication services

An external authentication server can be any server that supports either the RADIUS or the TACACS+ protocol, including Microsoft Active Directory servers.

Active Directory Server, RADIUS Authentication Server, and TACACS+ Authentication Server elements define the settings necessary for connecting to an external authentication server.

Authentication Method elements define the allowed authentication methods for the server, or the servers that use a particular authentication method.

In addition to the server element configuration in the SMC Client, you must configure the external authentication server to allow the engines to use the authentication services.

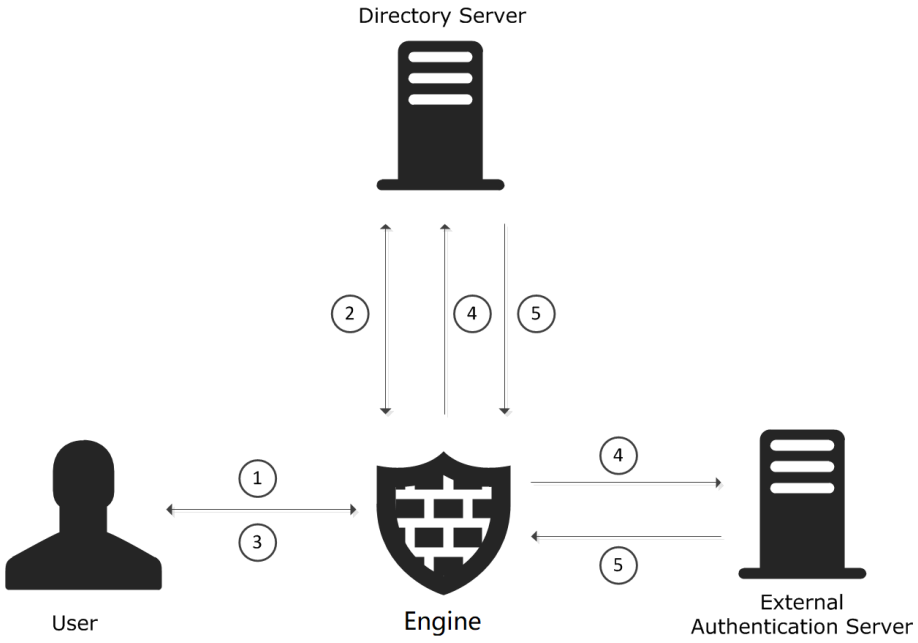
Overview of external user authentication

External user authentication means that authentication services are provided by an authentication server outside of the SMC.

You can use the following kinds of external authentication services:

- Authentication services that support the RADIUS or TACACS+ protocol, such as RSA Authentication Manager or the NPS (Network Policy Server) of a Windows (Active Directory) server.
- LDAP authentication for simple password authentication against the LDAP database on the external directory server where user accounts are stored.

External user authentication process



1. The user opens an authentication connection to the engine.
2. The engine queries the directory server to check if the user exists and which authentication method the user should use.
3. The engine prompts the user to authenticate, then the user enters the credentials required for the authentication method.
4. The engine relays the user credentials to one of the following components depending on the authentication method:
   - For RADIUS or TACACS+ authentication methods, the engine relays the user credentials to the external authentication sever.
   - For LDAP authentication, the engine relays the user credentials to the directory server.
5. Depending on the authentication method, one of the following components verifies the user credentials and responds to the engine whether authentication succeeds or fails:
   - For RADIUS or TACACS+ authentication methods, the external authentication server verifies the user credentials.
   - For LDAP authentication, the directory server verifies the user credentials.

Directory servers for external user authentication

Storing the user information and authenticating the users are two separate concepts with separate options

You can use the same server for storing and authenticating the users, such as when you use a Microsoft Active Directory server or an LDAP server for both tasks.

To define different IPv4 or IPv6 Access rules for different users and user groups with external authentication, you must integrate an external directory server with the SMC.

It is also possible to use an external authentication server without integration of an external directory server. In this configuration, the user information is not available to the engine. You cannot add different rules for different users and user groups. Instead, you add the user *external* with the external authentication methods into the internal user database, and use it to define rules for authentication.

Related concepts
Getting started with directory servers on page 1103

Using RADIUS in user authentication

Remote Authentication Dial-in User Service (RADIUS) is a protocol for carrying authentication, authorization, and configuration information.

RADIUS is a widely supported standard. For example, Microsoft NPS and RSA Authentication Manager support the protocol and can be used for user authentication in the SMC.

RADIUS uses UDP as its transport protocol. The exchanges between the client and the RADIUS server are authenticated by using a shared secret, which is never sent over the network. User passwords transferred between the client and the RADIUS server are encrypted using the MD5 message digest algorithm. The rest of the RADIUS communications are in cleartext.

Servers that provide RADIUS-based authentication methods can also be used for authenticating administrators' SMC Client logons and wireless client connections to wireless interfaces on engines.

RADIUS authentication servers support both IPv4 and IPv6 addresses.

Using TACACS+ in user authentication

Terminal Access Controller Access Control System Plus (TACACS+) is a protocol used for similar purposes as RADIUS.

In general, TACACS+ provides a more secure method of user authentication than RADIUS. TACACS+ uses TCP as the transport protocol instead of UDP, so transport is more reliable and less sensitive to disruption at the network layer.

TACACS+ also separates authentication, authorization, and accounting services, whereas RADIUS provides a user profile defining all user-specific parameters with the authentication. This separation of services allows TACACS+ to use other forms of authentication, such as Kerberos, together with its own authorization.

TACACS+ uses a pre-shared key to authenticate exchanges. TACACS+ encrypts all traffic between the authentication server and the device requesting authentication. User information, such as IDs and passwords, are secured with the MD5 message digest algorithm.

TACACS+ authentication servers support both IPv4 and IPv6 addresses.

Using LDAP authentication

When you use LDAP authentication, the external directory server where user accounts are stored verifies the user credentials.

When users authenticate to the Security Engine, the Security Engine sends the user name and password to the external directory server for authentication. The external directory server checks the user name and password against the user's credentials in the directory, then responds to the Security Engine whether authentication succeeds or fails.

Note

Because the user name and password are sent through the LDAP connection, we recommend using LDAPS or Start TLS when you use LDAP Authentication.

You can use LDAP authentication with the following features:

- IPsec and SSL tunnels in mobile VPNs
- The SSL VPN Portal
- Browser-based user authentication.

LDAP authentication has the following limitations:

- You cannot use LDAP authentication for users stored in the Management Server's internal LDAP user database.
- LDAP authentication is not supported for the WPA enterprise security mode on SSID Interfaces.

Note

WPA enterprise security mode always requires an external RADIUS server that has EAP support.

User authentication methods

Authentication Method elements define the authentication method used by particular authentication servers, and by particular users and user groups.

The SMC and engine supports many internal and external authentication methods.

The following authentication methods can be used to authenticate users:

- Client certificates.
- External authentication provided by servers that support the RADIUS (Remote Authentication Dial-in User Service) protocol.
- External authentication provided by servers that support the TACACS+ (Terminal Access Controller Access Control System Plus) protocol.
- LDAP authentication is used for simple password authentication against external LDAP databases.
- Pre-shared keys (for use with some third-party VPN clients).
- User passwords stored in internal or external LDAP databases.

Note

The user authentication methods are used for authenticating users who connects through the engine, or authenticating VPN client or SSL VPN Portal users.

The following authentication methods can be used to authenticate admin users:

- Client certificates.
- External authentication provided by servers that support the TACACS+ (Terminal Access Controller Access Control System Plus) protocol.
- External authentication provided by servers that support the RADIUS (Remote Authentication Dial-in User Service) protocol.
- LDAP authentication is used for simple password authentication against external LDAP databases.
- OpenID authentication by using an OpenID provider.
- SAML authentication by using a SAML based identity provider.

- User passwords stored in internal or external LDAP databases.

Note

1) The admin user authentication method is for authenticating admin users to grant access to SMC management tools, i.e. via SMC GUI client, or SMC Web Access.

2) The SAML and OpenID authentication methods can only be used with the SMC Web Access.

Create RADIUS or TACACS+ Authentication Server elements

You can authenticate end-user access through Engines and administrator's logons to the SMC against external authentication servers that support either the RADIUS or TACACS+ protocol.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **User Authentication**.

2) Right-click **Authentication Servers** and select **New** > **RADIUS Authentication Server** or **New** > **TACACS+ Authentication Server**.

3) Configure the settings on the **General** tab.

4) (Optional) Click the **Secondary IP Addresses** tab and add more IP addresses.
   These IP addresses are only used in Access rules and routing. Engines always use the main IP address of the RADIUS or TACACS+ authentication server when they contact the authentication server.

5) (Optional) If you want the RADIUS or TACACS+ authentication server to be monitored by the Log Server, click the **Monitoring** tab and configure the monitoring settings.

6) Click **OK**.

7) Configure the RADIUS or TACACS+ authentication server to accept connections from your Engines:
   - Make sure that the shared secret is entered identically in the SMC Client and on the RADIUS or TACACS + authentication server.
   - The identity that the Engine provides to the server is the IP address of the interface that has been selected as the value for the **IPv4 Identity for Authentication Requests or IPv6 Identity for Authentication Requests** or the **IPv6 Identity for Authentication Requests or IPv6 Identity for Authentication Requests** in the Engine's Interface Options.

     Note

     The IP address used as the identity is a name only. The interface and IP address used for the authentication-related connections is selected based on the Engine's routing information just like for any other connection.

Result

The connections to RADIUS or TACACS+ authentication servers are allowed in the predefined Firewall Template. Make sure your Access and NAT rules for user authentication are configured correctly for these connections.

Related concepts
Define contact IP addresses on page 127

Related tasks
Create Location elements on page 127
Activate monitoring of third-party devices on page 277

Define Authentication Method elements for external servers

Authentication Method elements represent the types of authentication provided third-party external authentication servers in user properties, and IPv4 or IPv6 Access rules.

You can use the following predefined Authentication Method elements without configuring a custom Authentication Method element:

- **Network Policy Server** — This Authentication Method for NPS is automatically used for authentication with Microsoft NPS and Active Directory.
- **LDAP Authentication** — You can use this Authentication Method in the properties of LDAP Server or Active Directory Server elements that provide simple password authentication against the external LDAP database.

To use other external authentication servers, or to use an Active Directory server for RADIUS-based authentication, you must define custom Authentication Method elements.

Each Authentication Method element can be associated with one or more servers, but each RADIUS Authentication Server or TACACS+ Authentication Server can only be associated with one Authentication Method element. When multiple servers are associated with the same Authentication Method element, the servers are used as alternative servers if the first contacted server does not respond. All servers associated with the same Authentication Method element must contain identical information for each authenticating user. It is not possible for the user to determine which of the alternative servers is contacted.

Authentication Method elements are used in the following configurations:

- In directory server element properties and in User and User group properties to specify the allowed authentication methods for the users.
- In the properties of a RADIUS or TACACS+ Authentication Server, or an Active Directory Server to specify the authentication method offered by the server.
- In the IPv4 or IPv6 Access rules to specify which authentication method users are required to use.

The RADIUS and TACACS+ protocols are generic communications protocols for user authentication. You can use many different types of authentication methods that use the RADIUS and TACACS+ authentication protocols, such as static passwords, one-time passwords, or any other user name/passcode-type authentication schemes.

Engines can also use the Internet Authentication Service (IAS) in previous Windows Server versions or the Network Policy Server (NPS) in Windows Server 2008 to authenticate end users. If you use a Windows Server's IAS/NPS service for authentication, define Active Directory Server elements instead.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **User Authentication**.
2) Browse to **Authentication Methods**.
3) Select **New** > **Authentication Method**.
4) Configure the settings, then click **OK**.

Define Access rules for authentication

The IPv4 and IPv6 Access rules in a engine policy can be configured to match only when the user is authenticated.

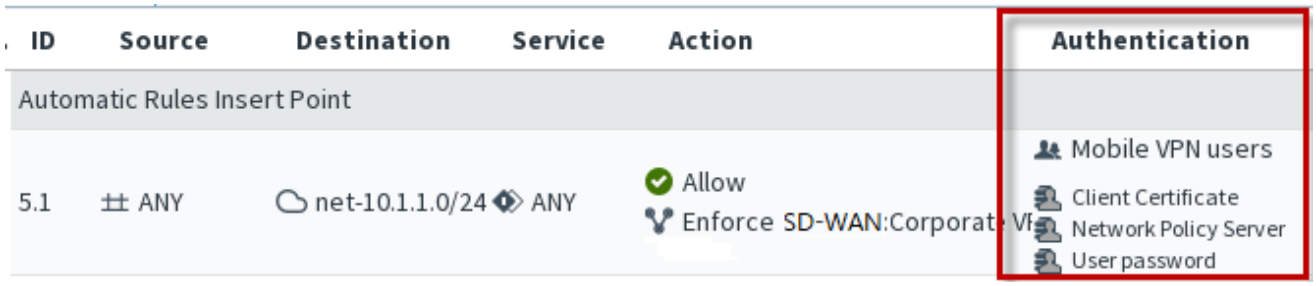
In the Access rules of a Engine policy, the Authentication cell specifies matching criteria for accessing a particular service and for setting options for the authentication. Authentication rules can be used to require authentication to access services and for authenticating VPN client users. With mobile VPNs, authentication is always mandatory. You can also require authentication for non-VPN access. Mobile VPN user authentication does not require specific rules for clients to authenticate. Browser-based authentication requires Access rules that allow access to the engine interface.

CAUTION

Only a VPN guarantees confidential information exchange. A rule that only requires authentication does not significantly increase the security of access from external networks to internal networks.

The authentication settings in a rule are configured in the same way regardless of whether a VPN is used. You define the authentication parameters in the **Authentication** cell.

Authentication field in the IPv4 Access rules



The User, User Group, and Authentication Method elements are only used as matching criteria. Any of the other rules above or below the rule for authentication can also match the authenticated user's connections. If necessary, you can define rules that discard connections from some combinations of Users and Authentication methods.

An authentication method is activated when at least one rule that contains the corresponding Authentication Method element is installed on the engine. The authentication is granted for a specific duration based on source IP address.

After the user successfully authenticates, the engine adds the user to a list of authenticated users. The next connection that the user opens can match an Access rule that requires authentication if the user and authentication method match the parameters of the rule.

Connections from users who have not successfully authenticated, or whose authentication has expired do not match rules that require authentication. The connection matching continues to rules further down in the policy.

It is especially important to consider whether other rules might match VPN client connections. If necessary, you can define rules that discard connections from some combinations of Users and Authentication methods. You can use the Source VPN cell in Access rules to match VPN traffic or non-VPN traffic. The VPN Client can be configured to receive an IP address from the organization's internal IP address space.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Browse to **Policies** > **<Policy type>**.
3) Right-click a policy, then select **Edit <Policy name>**.
4) Add an IPv4 or IPv6 Access rule, then define the **Source**, **Destination**, and **Service**.
5) Right-click the **Action** cell, then select the action.
6) Double-click the **Authentication** cell.
7) Configure the settings, then click **OK**.
8) Click **Save and Refresh**.

Related tasks

Enable browser-based user authentication on page 1137

Enable browser-based user authentication

As an alternative to authenticating using a VPN client, end users can authenticate themselves on an authentication page in a web browser.

The users can authenticate using encrypted HTTPS connections as well as plain HTTP connections. If the authentication will expire before the user has completed their tasks, the user can re-authenticate without disruption to any connections.

Browser-based user authentication is configured in the properties of the Security Engine. The IPv4 or IPv6 Access rules for allowing authentication connections are not included in the Firewall Template Policy. You must add rules that allow this traffic to the Security Engine's policy. You must also add Access and Inspection rules to enable redirection of unauthenticated HTTP connections to the logon page.

Related concepts
Getting started with example VPN configurations on page 1223

Related tasks
Customize the User Authentication Pages for browser-based user authentication on page 1148

Enable browser-based user authentication on the Security Engine

Browser-based user authentication is configured in the properties of the Security Engine.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click a Security Engine, then select **Edit <element type>**.

3) Browse to **Add-Ons** > **User Authentication**.

4) Select **HTTPS** to allow authentication using encrypted HTTPS connections and **HTTP** to allow authentication using plain HTTP connections.

CAUTION

Plain HTTP connections are unsecured and transfer the user name and password in cleartext. Use encrypted HTTPS connections to avoid loss of sensitive information.

5) (Optional) Change the port settings if you want to use a different port for the authentication interface.
You must use the same port settings when you define the IPv4 or IPv6 Access rules that allow the authentication connections.

6) (Optional) Select **Always Use HTTPS** if the Security Engine also listens on other ports and you want to redirect the connections to the HTTPS port and enforce the use of HTTPS.
Example: The Security Engine listens on port 80, but you want to redirect connections to port 443.

7) (Recommended) To prevent unauthorized access to resources, select the interfaces through which users can authenticate to the Security Engine in the **Listen on Interfaces** section.

8) From the **User Authentication Page** drop-down list, select a User Authentication Page element.
The User Authentication Page element defines the look of the logon page, challenge page, and status page shown to end users when they authenticate.

9) (Optional) Select **Enable Session Handling** to enable cookie-based strict session handling.
If the option is selected, the end user must keep the status page open. If the status page is closed or cannot be refreshed, the connection is terminated.

10) (Optional) Select **Refresh Status Page Every**, then define how often the status page is automatically refreshed.
- The option is automatically selected when you select **Enable Session Handling**.
- If the option is selected, the end user must keep the status page open. If the status page is closed or cannot be refreshed, the connection is terminated.

11) (Optional) Browse to **Advanced Settings** > **Authentication**, then configure advanced settings for browser-based user authentication.

Create and sign HTTPS certificates for browser-based user authentication

If HTTPS is enabled for Browser-Based User Authentication, you must have a signed HTTPS certificate.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click a Security Engine, then select **Edit <element type>**.

3) Browse to **Add-Ons** > **User Authentication**.

4) If **HTTPS** is not selected, select **HTTPS**.

5) Click **HTTPS Settings**.

6) Enter the certificate information.

7) Select how you want to sign the certificate:
- Select **With External Certificate Authority** if you want to create a certificate request for an external certificate authority to sign.
- Select **Internally with** to sign the certificate using the Internal CA for Gateways of the SMC.
If more than one valid internal certificate authority is available, select which internal CA signs the certificate request.

8) Click **Generate Request**.

9) (External certificate authorities only) When the certificate request is displayed, click **Export** and sign the certificate with an external certificate authority.

10) Click **Import Certificate** to import the signed certificate.

11) Click **OK** to close the **Certificate Request** dialog box.

12) Click **OK** to close the **Browser-Based User Authentication** dialog box.

Add Access rules for browser-based user authentication

Browser-based user authentication is not allowed by default in the Firewall Template policy. You must add Access rules that allows this traffic to the Engine Policy.

To reduce the risk of resource consumption or DoS (denial of service) attacks, we recommend limiting the number of connections from each source IP address. Under normal conditions, there should only be one connection at a time from each source IP address. However, incomplete connections or other network errors might temporarily result in more than one simultaneous connection attempt from the same IP address. Set the limit for your simultaneous connections according to your network environment so that the limit does not interfere with legitimate connection attempts.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Browse to **Policies** > **<Policy type>**.

3) Right-click a policy, then select **Edit <Policy name>**.

4) Add the following IPv4 or IPv6 Access rule:

Access rule for browser-based user authentication

Source	Destination	Service	Action
ANY	$$Local Cluster (CVI addresses only) or $$Interface ID X. (If specific listening interfaces are selected on the General tab in the Browser-Based User Authentication Properties.)	HTTP, HTTPS, or both (Port settings must be the same as defined on the General tab in the Browser-Based User Authentication Properties.)	Allow Connection tracking: Default Connection limit by Source: the number of simultaneous connection attempts you want to allow

5) Click **Save and Install**.

Enable redirection of unauthenticated HTTP or HTTPS connections

To use browser-based user authentication, you must define some IPv4 or IPv6 Access rules.

Note

To redirect HTTPS traffic, you must enable TLS decryption for the traffic.

You must define the following IPv4 or IPv6 Access rules:

- An Access rule that allows all clients to access the logon page.
- An Access rule that allows authenticated users to establish HTTP or HTTPS connections.
- An Access rule that redirects unauthenticated HTTP or HTTPS traffic to the logon page.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Browse to **Policies** > **<Policy type>**.

3) Right-click a policy, then select **Edit <Policy name>**.

4) Add the following Access rules:

Example Access rules for unauthenticated HTTP connections

Source	Destination	Service	Action	Authentication
ANY	IP addresses of interfaces through which users can authenticate.	HTTP HTTPS (Port settings must be the same as defined in the User Authentication settings for the Security Engine.)	Allow	
ANY	IP addresses of network services that require authentication.	HTTP HTTPS	Allow	Users or User Groups who are allowed to access services, and appropriate Authentication Methods.
ANY	IP addresses of network services that require authentication.	HTTP HTTPS	Refuse Connection tracking: Default Response: redirect to the logon page.	

5) Click **Save and Install**.

Related concepts
User Response elements and how they work on page 969
TLS inspection and how it works on page 1063

Enable redirection to the original destination

Add the necessary Access rules to configure the redirection of unauthenticated HTTP or HTTPS connections from the status page to the destination that the user originally wanted to access.

Note

To redirect HTTPS traffic, you must enable TLS decryption for the traffic.

You must define the following IPv4 or IPv6 Access rules:

- An Access rule that allows all clients to access destinations that do not require authentication.
- An Access rule that allows authenticated users to establish HTTP or HTTPS connections.
- An Access rule that redirects unauthenticated HTTP or HTTPS traffic to an Inspection rule.
- An Access rule that refuses all HTTP or HTTPS traffic.

Using the HTTP_Request-with-redirect-capability Situation, you must also define the following Inspection Exceptions in the Inspection Policy:

- An Exception that permits all matching connections to access destinations that do not require authentication.
- An Exception that permits authenticated users to establish HTTP connections.
- An Exception that redirects unauthenticated HTTP traffic to the logon page using the original destination URL as a parameter in the redirection.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Browse to **Policies** > **<Policy type>**.

3) Open the Engine Policy for editing, then add the following Access rules:

Example Access rules for redirecting unauthenticated HTTP connections to the original HTTP destination

Source	Destination	Service	Action	Authentication
ANY	IP addresses of services that do not require authentication.	HTTP HTTPS	Allow	
ANY	ANY	HTTP HTTPS	Allow	Users/User Groups who are allowed to access services, and appropriate Authentication Methods.
ANY	IP addresses of network services that require authentication.	HTTP HTTPS	Allow Deep Inspection: on	
ANY	ANY	HTTP HTTPS	Refuse	

Note

Deep Inspection must be enabled in the Access rules for redirecting unauthenticated HTTP or HTTPS connections to the original destination. The redirection must be configured in the Inspection Policy using the HTTP_Request-with-redirect-capability Situation.

4) Click **Save**.

5) Open the Inspection Policy for editing.

6) Add the following Inspection Exceptions, then specify a User Response that redirects traffic terminated by the Inspection rules to the URL of the logon page and onwards to the original destination.

Example Inspection Exceptions for redirecting unauthenticated HTTP connections to the original HTTP destination

Situation	Severity	Source	Destination	Protocol	Action
HTTP_Request-with-redirect-capability	ANY	ANY	IP addresses of services that do not require authentication	ANY	Permit
HTTP_Request-with-redirect-capability	ANY	Users/User Groups who are allowed to access services, and appropriate Authentication Methods.	ANY	ANY	Permit
HTTP_Request-with-redirect-capability	ANY	ANY	ANY	HTTP	Terminate Response: redirect to the logon page, including the original URL as a parameter in the redirection

7) Click **Save and Install**.

Related concepts

User Response elements and how they work on page 969
TLS inspection and how it works on page 1063

Configure client certificate authentication for browser-based user authentication

In environments that require multi-factor user authentication, you can configure certificate-based authentication using X.509 certificates for browser-based user authentication.

Before you begin

Before configuring client certificate authentication, configure the following:

- Integrate an external Active Directory server or LDAP server with the SMC.
- Enable browser-based user authentication.

Users can authenticate to the engine using an X.509 certificate stored on their computers or on a smart card, such as a Common Access Card (CAC). The Security Engine verifies that the certificate is valid and that the value configured to be checked in certificate matches the value for the user in the LDAP server.

Note

Enabling client certificate authentication prevents the use of user password authentication in the same user session.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Create a TLS Profile element that defines the trusted certificate authority for the users' certificates.
   Only users whose certificates are signed by the trusted certificate authority can successfully authenticate using client certificates.
   a) Select **Administration**.
   b) Browse to **Certificates** > **Other Elements** > **TLS Profiles**
   c) Right-click **TLS Profiles**, then select **New TLS Profile**.
   d) In the **Trusted Certificate Authorities** section, select **Trust Selected**, then click **Add** to specify the trusted certificate authorities that sign the users' client certificates.
   e) Configure the other settings as needed, then click **OK**.
2) Configure client certificate authentication on the Engine.
   a) Right-click a Security Engine, then select **Edit <element type>**.
   b) Browse to **Add-Ons** > **User Authentication**.

c) Next to the **TLS Profile** field, click **Select**, then select the TLS Profile element that you created.

d) Browse to **Advanced** > **Authentication**.

e) From the **Client Certificate Identity Field for TLS** drop-down list, select the client certificate field that is used to look up the user entry from the user domain.

f) Configure the other advanced options as needed, then click **Save**.

3) If you selected **Distinguished Name** as the **Client Certificate Identity Field for TLS**, configure the Active Directory Server or LDAP Server element for client certificate authentication.

a) Select **User Authentication**.

b) Browse to **Servers**.

c) Right-click the LDAP Server or Active Directory Server element, then select **Properties**.

d) On the **Client Certificate** tab, enter the name of the value in the distinguished name that is checked to verify the client identity.

The supported values are CN, email, and UID.

e) Make sure that the value of the **UserId** field on the **Attributes** tab matches the attribute that contains the specified user information.

- CN — The value of the **UserId** field on the **Attributes** tab must be `CN`.
- email — The value of the **E-mail** field on the **Attributes** tab must match the attribute that contains the user's email address.
- UID — The value of the **UserId** field on the **Attributes** tab match the attribute that contains the user's UID.

f) Click **OK**.

4) Enable client certificate authentication on the Engine.

a) Right-click a Security Engine, then select **Edit <element type>**.

b) Browse to **Add-Ons** > **User Authentication**.

c) Next to the **TLS Profile** field, click **Select**, then select the TLS Profile element that you created.

d) Click **Save and Refresh**.

Authenticate to the Security Engine

End users can authenticate and re-authenticate using a compatible VPN client or a web browser.

Before you begin

To use smart cards for authentication, you must have smart card reader hardware and software.

To use certificate files for authentication, you must save the certificates in a location that is accessible from your web browser.

If the users are authenticating for VPN access, they must authenticate using a compatible VPN client.

CAUTION

If users authenticate over an unsecured connection, use a one-time password scheme to reduce the risk of unauthorized access.

Steps

1) Access the authentication prompt in one of the following ways:
   - Follow the instructions for the VPN client about connecting and authenticating.
   - Enter the IP address and port of the Engine to open an authentication page in a web browser.

2) To authenticate using a user name and password, enter the user credentials.
   If you enter your user name without specifying the LDAP domain, the default LDAP Domain is used. If your user account does not belong to the default LDAP Domain, add the LDAP Domain to the user name with the `@` character as a separator.
   For example, enter `fred@mobileusers` for the user `fred` in the LDAP Domain `mobileusers`.

3) To authenticate using a client certificate, do the following.
   a) If you have a smart card, insert the smart card into the smart card reader.
   b) If there is more than one certificate on the smart card or on your computer, select the certificate to use for authentication.
   c) (Smart card only) Enter the PIN for the smart card if you are prompted to do so.

4) To re-authenticate the active session, do one of the following.
   - Follow the instructions for the VPN client about re-authenticating.
   - For browser-based re-authentication, click **Re-Authenticate** on the status page. Depending on the authentication method, enter your password or enter the response to the challenge.

Customize the User Authentication Pages for browser-based user authentication

You can customize the logon, challenge, re-authentication, and status pages that are shown to users when they authenticate using a web browser.

The pages you define are the same for all user groups. The default HTML pages that you can base your customized pages on are included in dynamic update packages.

Export the default HTML pages for browser-based user authentication

You can export the HTML pages from the default User Authentication Page element, then edit the HTML pages using any HTML editor.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Browse to **Other Elements** > **Engine Properties** > **User Authentication Pages**.
3) Right-click the default User Authentication Page element, then select **Properties**.
4) Click **Export**, then save the .zip file under a different name.
5) Click **Cancel** to close the dialog box.

Customize the default HTML pages for browser-based user authentication

You can use an HTML editor to edit the HTML pages used for browser-based user authentication. For example, you can add images and adjust the CSS files for the default pages.

Steps

1) Decompress the exported .zip file to your computer.

2) Open and edit the HTML files for each page in an HTML editor.

Important

Do not remove the following fields and buttons from the default HTML files.

HTML file	Required fields
Logon page	The **Username** and **Password** fields, and the **Log On** button.
Challenge page	The **Challenge response** field and the **Log On** button.
Status page	The **Logged on**, **Connection status**, and **Authentication message** fields.
Status Page With Automatic Redirection	The **Logged on**, **Connection status**, and **Authentication message** fields, and the link for users to continue to their original destination.
Status Page With Manual Redirection	The **Logged on**, **Connection status**, and **Authentication message** fields, and the link for users to continue to their original destination.
Re-Authentication page	The **Username** and **Password** fields, and the **Log On** button.

3) Compress the custom files into a .zip file.

Import the custom HTML pages for browser-based user authentication

To use the customized HTML pages, you must create a custom User Authentication Page element and import the custom files.

The HTML files are validated during import to make sure that the HTML pages are valid HTML.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Browse to **Other Elements** > **Engine Properties** > **User Authentication Pages**.

3) Select **New** > **User Authentication Page**.

4) Next to **Zip File**, click **Browse**, then select the .zip file that contains the custom HTML files.
The name is added automatically based on the name of the .zip file.

5) (Optional) To preview the HTML pages, click the corresponding preview option.

6) Click **OK**.

Monitoring and testing user authentication

You can find information about successful and failed user authentication attempts, as well as the system's own connections to external authentication servers in the logs.

You can also create reports based on this data. There is a separate view for monitoring the currently active authenticated users.

If there are problems with integration with external components, you can activate more detailed diagnostics logging for authentication.

Related concepts

Monitoring connections, block lists, VPN SAs, users, routing, SSL VPNs, and neighbors on page 236

Related tasks

Enable or disable diagnostics on page 356

Examples of user authentication

These examples illustrate common uses for User Authentication and general steps on how each scenario is configured.

Example: authenticating VPN client users

This scenario shows an example of restricting VPN access so that only specific users can access the secure network.

Company A's employees include several consultants who frequently work at customer locations, but also remotely access Company A's secure network. All users are stored in the Management Server's internal directory, and there is a separate User Group called Consultants for accounts belonging to the consultants. The administrators have set up a mobile VPN for remote access. They want to allow all users to establish a VPN tunnel to the office, but allow only users in the Consultants group to access the secure network.

The administrators:

1) Create a rule that establishes a VPN tunnel and allows users in the Consultants group to access the Secure Network after successful authentication:

Source	Destination	Service	Action	Authentication
DHCP address range for VPN clients Internal Networks	Secure Network	HTTP SSH FTP	Enforce VPN	Consultants User Group User Password Authentication

- This rule allows any users in any directory that is defined in the SMC to authenticate to a VPN client if their allowed authentication methods include User Password.
- This rule allows any user whose account is stored in the internal directory to use a VPN client to establish a VPN tunnel to the office.

2) Create a rule to allow users who have established VPN tunnels to access the company's internal networks from the DHCP-assigned IP addresses for VPN clients:

Source	Destination	Service	Action	Authentication
DHCP address range for VPN clients	Internal Networks	ANY	Allow	

3) Transfer the policy to the engine.

Example: integrating Microsoft Active Directory Servers

A general overview of integrating Active Directory servers.

For more information about configuring the NPS, see Microsoft's documentation at https://technet.microsoft.com.

Company B has an existing Microsoft Active Directory server that stores user information. They decide to use this existing information for user authentication.

The administrators:

1) Define an Active Directory Server element.

2) Add the SMC-specific classes and attributes into the Active Directory server's configuration to be able to fully manage the user accounts through the SMC Client.

3) Define the Management Server as an LDAP client for the Active Directory server.

4) Define the Engine as an authentication client for the NPS.

5) Add an LDAP Domain element for the Active Directory server in the SMC Client.

6) Add an Access rule with authentication defined as shown here.

Example Access rule for NPS authentication

Source	Destination	Authentication
IP addresses of authenticated hosts.	IP addresses of network services that require authentication.	Some User or User Group elements from the AD's LDAP Domain. Require authentication with "Network Policy Server" Authentication Method.

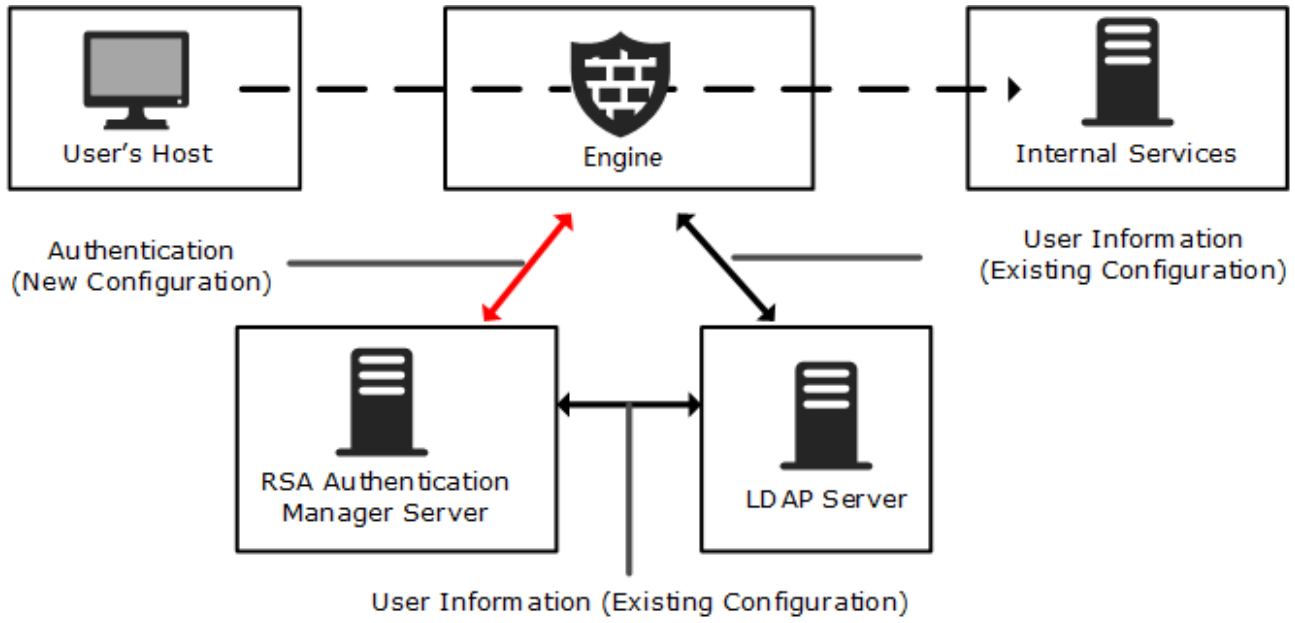
Example: using SecurID authentication with the Forcepoint VPN Client

This example shows a general overview of using SecurID authentication for the Forcepoint VPN Client.

For more information about using SecurID authentication, see the RSA documentation at https://www.rsa.com.

Company C is about to introduce remote Forcepoint VPN Client access to their network. The administrators decide to add one-time passwords with SecurID cards with their existing RSA Authentication Manager server that already shares the user information with the company's LDAP server.

Company C's authentication scheme



The administrators:

1) Create an Agent Host record for the Engine in the RSA Authentication Manager server.

2) Configure a mobile VPN in the SMC Client with the default Hybrid Authentication selected as the authentication method for connecting clients.
   - Hybrid authentication is available for the Forcepoint VPN Client. Hybrid authentication requires the client to authenticate the VPN Gateway (the engine) by using a certificate. The users must provide the correct User Name/Password combination (validated by the RSA Authentication Manager server in this case).

3) Create a RADIUS Authentication Server element.

4) Create a custom Authentication Method element for the server, named "SecurID".

5) Open the Active Directory Server Properties dialog box, and do the following:
   a) Click the **Authentication** tab.

b) Add the new authentication method under **External Authentication Methods** to indicate that it is supported by the server.

6) To configure the SecurID authentication method to always be used with the domain by default, do the following:

a) Open the LDAP Domain element properties dialog box.

b) Click the **Default Authentication** tab.

c) Select the **SecurID** authentication method.

Note

If SecurID is not the default authentication method, then users must login with syntax "*username;SecurID*".

7) Add Access rules with both an authentication and a VPN requirement defined as shown here:

Example Access rule for SecurID authentication

Source	Destination	Authentication	Action
The virtual IP address range used on the virtual adapters of the Forcepoint VPN Client.	IP addresses of network services that require authentication.	User or User Group elements. Require authentication with "SecurID" Authentication Method.	**Allow**, with the **VPN Action** option set to **Enforce VPN**.

Part XI

Virtual private networks

Contents

- VPNs in Forcepoint Network Security Platform on page 1157
- Configuring VPNs on page 1169
- Example VPN configurations on page 1223
- Managing VPN certificates on page 1251
- Reconfiguring existing VPNs on page 1269
- VPN client settings on page 1279
- Configuring the SSL VPN Portal on page 1285

Forcepoint Network Security Platform supports both policy-based and Route-based Tunnels between VPN gateways. For full remote access, Forcepoint Network Security Platform supports both IPsec and SSL VPN tunnels for VPN clients.

Forcepoint Network Security Platform also supports an SSL VPN Portal that provides access to internal HTTP and HTTPS services through a standard web browser or through a client application that allows direct network access.

Chapter 72

VPNs in Forcepoint Network Security Platform

Contents

- Types of VPNs in Forcepoint Network Security Platform on page 1157
- Pre-shared key (PSK) authentication in VPNs on page 1160
- Certificate-based authentication in VPNs on page 1161
- Configuring VPNs with external gateway devices on page 1162
- Adaptive Forward Erasure Correction (FEC) for VPN tunnels on page 1163
- Logs related to VPNs on page 1164
- Clustering and VPNs on page 1164
- VPNs and Multi-Link for VPN on page 1165
- VPN Broker on page 1167

A VPN extends a secured private network over public networks by encrypting connections so that they can be transported over insecure links without compromising confidential data.

Types of VPNs in Forcepoint Network Security Platform

Forcepoint Network Security Platform provides two types of VPNs. The main difference between the two is how traffic is selected to use the VPN.

- *Policy-based VPNs* are configured using Policy-Based VPN elements. The Engine Access rules define which traffic is sent to the VPN and which traffic is allowed out of the VPN.
- *Route-based VPNs* are configured using the Route-based Tunnels elements. Any traffic that is routed to engine interfaces that are designated as endpoints for a VPN tunnel is sent into the VPN tunnel. If Access rules allow the traffic, it is automatically sent through the tunnel to the peer endpoint.

Policy-based VPNs are recommended for the following uses:

- To create mobile VPNs with IPsec tunnels, SSL VPN tunnels, or both IPsec and SSL VPN tunnels.
- To easily create VPN topologies with multiple connections between multiple gateways, such as full mesh, partial mesh, star, and hub topologies.

Route-based Tunnels are recommended for the following uses:

- To use VPN tunnels as paths in dynamic routing.
- To protect the integrity of dynamic routing communications that are sent through the Internet.
- To protect and route multicast streams through the Internet.
- To configure GRE, IP-IP, or SIT tunnels that encapsulate traffic but do provide encryption.

Limitations of VPNs in Forcepoint Network Security Platform

These limitations apply to VPNs in Forcepoint Network Security Platform.

- You cannot use the same pair of endpoints for VPN tunnels in several configurations for a single Security Engine. For example:
  - You cannot use the same pair of endpoints l in two policy-based VPNs.
  - You cannot create two Route-based Tunnels elements that use the same pair of endpoints.
  - You cannot create a Route-based Tunnels element that uses the same pair of endpoints that is used in a VPN tunnel in a policy-based VPN.
- VPNs are not supported on layer 2 physical interfaces on Engines.
- VPNs are not supported on Layer 2 Engines.
- If your Forcepoint Network Security Platform installation is configured in a restricted operating mode to comply with regulatory requirements, some VPN options are not available to you.
- Version-specific limitations in supported features for different Forcepoint Network Security Platform versions are listed in the Release Notes for the versions you are using. The SMC automatically prevents the use of unsupported settings based on engine version.

Policy-based VPNs in Forcepoint Network Security Platform

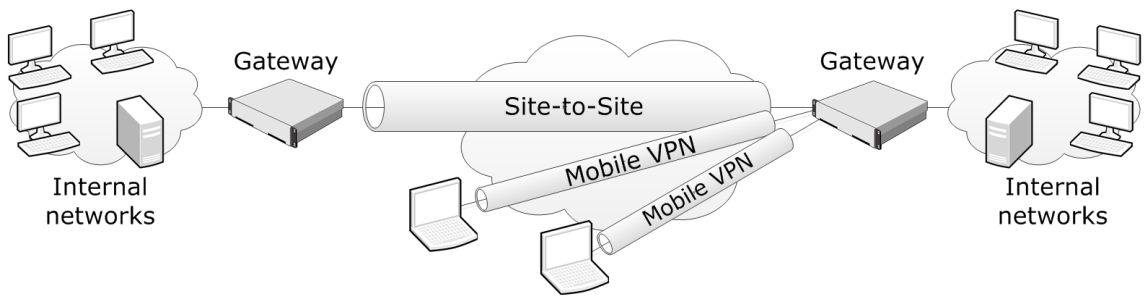
In Policy-based VPNs, the Access rules determine which traffic is sent into the VPN tunnels.

Site-to-site and mobile VPNs in Forcepoint Network Security Platform

You can create VPNs between VPN gateway devices or between a VPN client and a VPN gateway device.

- A *site-to-site VPN* is created between two or more gateway devices that provide VPN access to several hosts in their internal networks. Site-to-site VPNs are supported for IPv4 and IPv6 traffic.
- A *mobile VPN* is created between a VPN client running on an individual computer and a gateway device.

Site-to-site and mobile VPNs



For mobile VPNs, we recommend using the Forcepoint VPN Client solution. Forcepoint VPN Client is available for the following platforms:

- Android (SSL VPN only)
- Mac OS (SSL VPN only)
- Windows (IPsec or SSL VPN)

In mobile VPNs with IPsec tunnels, you can alternatively use a third-party IPsec-compatible VPN client. However, third-party clients do not support all features offered by Forcepoint Network Security Platform.



Note

Most VPN clients that are a part of a vendor-specific VPN gateway solution are incompatible with gateways from other vendors.

The following limitations apply to mobile VPNs:

- All mobile VPNs that you configure in Forcepoint Network Security Platform must be valid for Forcepoint VPN Client even if you use only third-party VPN client software.
- VPN clients cannot connect directly to engines that have a dynamic IP address.
  Instead, VPN clients connect through a central gateway that forwards the connections to the non-compatible gateways using a site-to-site VPN.

Types of encryption for tunnels in policy-based VPNs

Tunnels in policy-based VPNs can use IPsec or SSL VPN encryption.

- *IPsec* — The IPsec protocol allows any IP traffic to be transported in the VPN regardless of which higher-level protocol the traffic uses on top of the IP protocol. Hosts can communicate through the VPN as if it was a normal link without the need for application-specific configurations on the gateway device. IPsec is part of both the IPv4 and IPv6 standards. IPsec is defined in RFC 4301.
  You can use IPsec VPN tunnels in both site-to-site and mobile VPNs.
- *SSL VPN* — SSL VPN tunnels use secure sockets layer (SSL) encryption to provide secure remote access. You can use SSL VPN tunnels in mobile VPNs.

You can use SSL VPN tunnels alone, IPsec tunnels alone, or both SSL VPN and IPsec tunnels together in the same policy-based VPN.

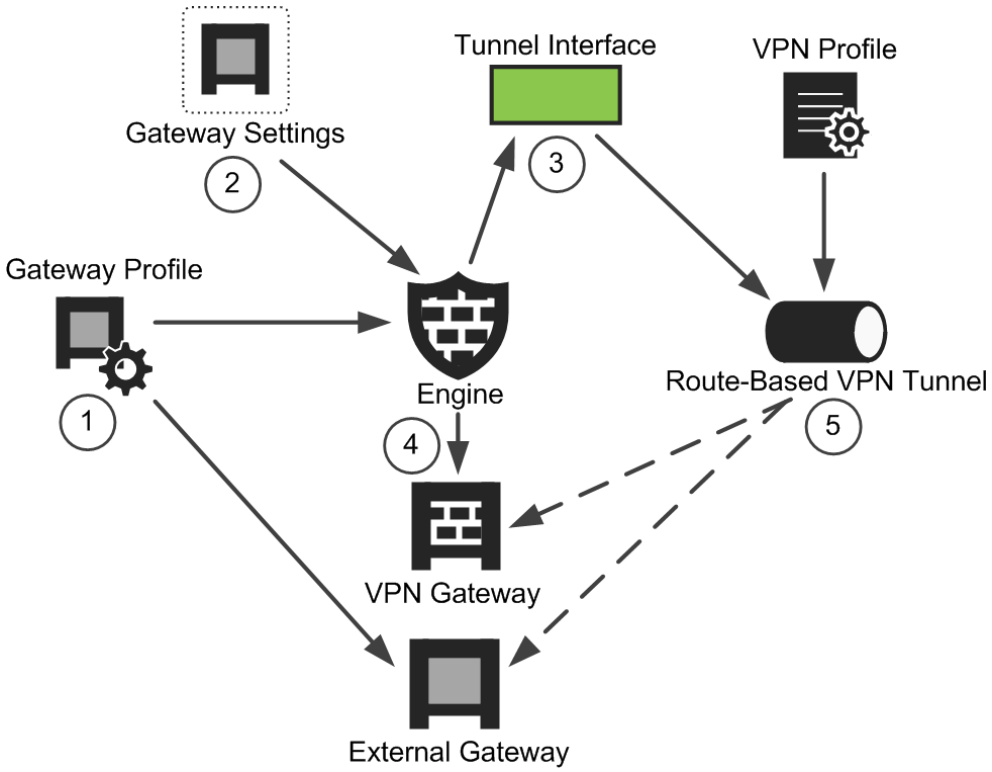
How route-based VPNs work

In route-based VPNs, any traffic that is routed to a tunnel interface and allowed by Access rules is automatically sent through the tunnel to the peer endpoint.

Devices that provide VPN access are called *VPN gateways*. With route-based VPNs, you can create only *site-to-site* VPN tunnels between gateway devices. It is not possible to create a mobile VPN tunnel using Route-based Tunnels.

Due to the various authentication and encryption methods that the IPsec protocol supports, the number of settings is rather high. To reduce configuration work, you can use reusable profiles for storing different types of settings. These and other elements related to Route-based Tunnels configurations are pictured in this illustration.

Route-based Tunnels configuration example



1. The Security Engine and External Gateway refer to a Gateway Profile element that contains information about the capabilities of different types of gateways.
2. The Security Engine optionally refers to a Gateway Settings element that defines settings for advanced VPN performance tuning.
3. The Security Engine has a Tunnel Interface element defined.

   The Tunnel Interface refers to a Route-based Tunnels element.
4. The Security Engine has a VPN Gateway element defined.

   One VPN Gateway is automatically created for each Security Engine in the Engine/VPN role. You can optionally add more VPN Gateways to the Security Engine.
5. The Route-based Tunnels element refers to both ends of the VPN tunnel: the VPN Gateway and External Gateway.

   The Route-based Tunnels also refers to a VPN Profile, which contains the IPsec authentication and encryption settings (IKE settings)

Pre-shared key (PSK) authentication in VPNs

A pre-shared key is a string of characters that is used as an authentication key. You can use pre-shared keys for site-to-site VPN authentication and with third-party VPN clients.

Both gateways create a hash value based on the pre-shared key and other information. The hash values are then exchanged and verified to authenticate the other party. As its name suggests, the pre-shared key has to be

distributed beforehand to all devices that use it. Pre-shared keys must be transferred confidentially because their security benefit is immediately lost if the key is exposed to unauthorized parties.

The pre-shared keys must also be long and random to be secure. Short or predictable pre-shared keys can be easily broken in brute-force attacks. Administrators must also remember to renew the pre-shared keys periodically to maintain a high level of security. Forcepoint Network Security Platform includes tools for generating sufficiently long, random pre-shared keys for VPN components. The keys are automatically transferred to any Security Engines that need them using the secure management communications channel.

Related tasks

Create VPN Profile elements on page 1191
Replace pre-shared keys for VPNs on page 1274

Certificate-based authentication in VPNs

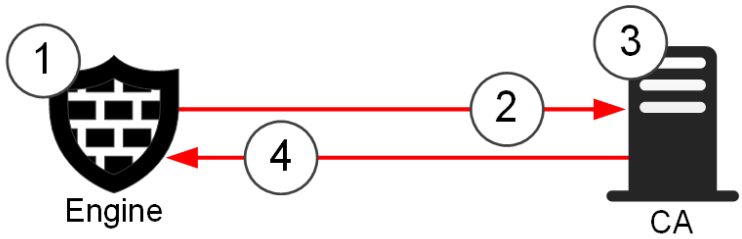
In all site-to-site VPNs and in mobile VPNs with third-party VPN clients, you can use certificates for authentication.

With the Forcepoint VPN Client, the following types of authentication are available:

- *Hybrid authentication* requires the presence of a valid certificate on the gateway and some other form of authentication from the VPN client user.
- *Certificate exchange authentication* requires a certificate from both the gateway and the VPN client.

Certificates often provide a higher level of security than pre-shared keys. Certificates only have to be renewed at an interval of a few years, and have an automatic expiration mechanism that makes sure the certificate is renewed. Certificate files cannot be compromised in transit, because they cannot be used without a private key. This illustration outlines the basics of how a certificate is generated.

VPN certificate creation



1. When a certificate request process is started, a private key is generated and stored.
2. The certificate requester uses the private key to generate a certificate request that is transferred to the certificate authority (CA).
3. The CA signs the certificate request, which validates the certificate.
4. The signed certificate is transferred to the original certificate requester.

For VPN gateways, all steps can be automatic if the default internal CA for gateways is used for signing the certificate. If another certificate authority is used, the certificate request is exported from the SMC and the signed certificate is imported back into the SMC.

For VPN clients, certificate-based authentication typically uses an existing certificate that is stores on an external token or in an operating system certificate store. The signed certificate is imported into the SMC. Alternatively,

the certificate request file can be created manually in the VPN client and transferred manually to be signed by an internal certificate authority in the SMC or another certificate authority. The signed certificate is then transferred manually into the VPN client computer.

Private keys are always generated automatically. If the private key is lost, such as due to a hardware failure, any associated certificate becomes unusable and a new certificate must be created. The private key is securely and automatically synchronized between clustered engine nodes to allow all nodes to use the same certificate.

Unlike pre-shared keys, certificates do not need to be distributed to all gateways in the VPN. Instead, the other gateways are configured to trust the CA that signed the certificate, after which they trust all certificates from that issuer. This trust relationship also allows renewing or re-creating the certificate on one gateway without having to reconfigure the other gateways. Only certificates from trusted CAs are accepted for authentication. For this reason, VPN gateways must be configured to trust the CAs that sign the certificates that the other gateways use for authentication.

Related concepts

VPN certificates and how they work on page 1251

Related tasks

Create VPN Profile elements on page 1191

Configuring VPNs with external gateway devices

An External VPN Gateway is any VPN gateway that is not controlled by the same Management Server (and the same administrative Domain) on which you are configuring the gateway element.

Often, external gateway devices are at a partner organization, not under your control, and not Forcepoint Network Security Platform devices. Because IPsec is a networking standard, you can create a VPN between gateways of different brands by selecting the settings you want identically for both gateways. Any option that both gateways support is a valid option for the VPN.

The settings that must match are:

- The IKE SA settings.
- The IPsec SA settings.
- The site definitions (IP addresses) defined for both gateways at both ends (possibly translated using NAT).
- The endpoint identity type and value. The endpoint identity value is often the IP address of each gateway, but other options are also possible.

When the listed settings are identical, the VPN works. However, there are some things that you must consider when you configure VPNs with external gateway devices:

- Every setting must match to produce a fully functional VPN, but the supported options might be partly different on the different gateways.
- Because there is not a single common standard for naming the different options, the two gateways might use a different name for the same authentication or encryption method.
- If Forcepoint Network Security Platform devices are used as External VPN Gateways, you can export and import some settings between the two Management Servers (or between administrative Domains). However, you must still construct many of the configurations manually.
- The IP addresses accessible through each gateway must match.

In VPN Gateways controlled by the Management Server on which the VPN is configured, the IP addresses included in the policy-based VPN are defined as separate Site elements. The security association (SA) granularity setting defines whether a new VPN tunnel is established for each communicating host or for each network. In most gateways, there is an option for the SA setting. However, some gateways might select the SA automatically based on the type of IP address definition or even have a fixed setting.

Note

Site definitions are always defined for the VPN Gateway or External VPN Gateway element and are used in all policy-based VPNs where the same gateway is used. If you add a site to a gateway in one policy-based VPN, disable it in other policy-based VPNs where you do not want the site to be included.

Note

In route-based VPNs, the site information is ignored. In Route-based Tunnels the site definition is always 0.0.0.0/0 for IPv4 and ::/0 for IPv6 (any network).

Related concepts

Supported encryption methods for VPNs on page 1186

Adaptive Forward Erasure Correction (FEC) for VPN tunnels

The FEC feature helps to control errors in data transmission over an unreliable or noisy communication channel.

Requirements:

- The FEC feature is supported on engine version 7.2 and later.
- When FEC is configured for a VPN tunnel, the gateways at both ends of the tunnel must support the FEC implementation. If this condition is not met the FEC configuration is ignored.
- The FEC feature should not be used if packet loss is caused due to traffic congestion.
- The FEC feature is not suitable for bulk data transmissions, and should not be used with applications which already implement end to end FEC like feature.

When FEC is enabled, the engine sends a combination of M data packets, N correction packets, and metadata information in a data set through a link to the destination. This allows recovering up to N missing data packets within one set of data and correction packets.

Important

It is recommended to only use this feature as a last resort option for the critical application traffic that is highly sensitive to packet loss, when no better links are available. Note that when the feature is enabled, it increases the network bandwidth usage for the traffic selected for FEC and can cause traffic congestion making things worse if packet loss is due to congestion.

The FEC feature can be enabled by:

1) Selecting what QoS classes are enabled for FEC. For more details, refer to the Create Link Usage Profile elements topic.

2) (Optional) Controlling what link types must be used for FEC. For more details, refer to the Create Link Usage Profile elements topic.

3) Using the configured Link Usage Profile in the engine configuration. For more details, refer to the Select a Link Usage Profile element for an Secure SD-WAN Engine topic.

4) Selecting the QoS class in the access rule to match the traffic for FEC. For more details, refer to the Use QoS Class elements to apply custom link selection options to traffic topic.

Note

- The maximum transmission unit (MTU) size for the connections subjected to FEC is slightly smaller than the MTU size with the normal VPN tunnel. The ICMP fragmentation needed message is generated by the engine when needed.
- When engine configuration includes FEC, but engine version does not support FEC, the FEC settings are ignored, and warning is displayed during the policy installation.
- The configuration done must match in both the tunnel endpoint engines to achieve symmetric FEC for both sent and received packets of the critical connections.
- The FEC configuration is interpreted on the sender side and recipient side adjusts itself to the configuration.
- The FEC correction packets have the same QoS class as the data packets in the set. When it comes to interface used to send out ESP packets, potential DSCP marking is applied to correction packets. Also, the correction packets are included in the QoS processing and in the interface statistics.
- The correction packets are counted as traffic inside VPN tunnel.
- The tunnel statistics still reflect the tunnel itself and FEC does not change the tunnel statistics. However, the application health is reported as the application determines it.

Logs related to VPNs

VPN negotiations and VPN traffic are logged as informational messages and can be viewed in the **Logs** view like any other logs.

New connections that are allowed through the policy-based VPN are logged like any other traffic according to the logging options in the Access rules.

If there are VPN related problems, you can activate IPsec diagnostics for the engine to get more detailed information about the VPN negotiations. The Troubleshooting topics contain further information about possible problems, including explanations for the most common messages you might see in the logs.

Clustering and VPNs

A Forcepoint Network Security Platform cluster can be used as a gateway in policy-based and route-based VPNs. There are no additional configuration steps compared to a Single Security Engine.

Clustering provides high availability and load balancing at the VPN gateway with multiple nodes in a cluster. If one of the nodes is commanded offline or fails, the remaining nodes in the cluster take over the VPN traffic that was handled by that node. To allow the nodes to use the same certificate, the associated private encryption key

is exchanged securely through the heartbeat channel. To external VPN gateways, the cluster presents itself as a single device with a single endpoint (CVI IP address) to contact.

VPNs and Multi-Link for VPN

Using Multi-Link enhances the reliability of the VPN communications by ensuring the availability of network connections.

Forcepoint Network Security Platform can balance the VPN traffic load between multiple network connections and redistribute traffic when a connections becomes unavailable. Using Multi-Link reduces the possibility of traffic congestion or ISP network connectivity failures. Multi-Link is not a part of the IPsec standards.



Note

Multi-Link is only supported with Forcepoint Network Security Platform gateways at both ends of the VPN tunnel. If an external gateway device allows configuring multiple VPN tunnels between two devices, you might still be able to use some Multi-Link features. Not all Multi-Link features are available with external gateway devices.

In a Multi-Link VPN configuration, the traffic can use one or several alternative tunnels to reach the same destination. Multi-Link guarantees that even if one or more tunnels fail, the VPN service continues as long a tunnel is available.

You can use Multi-Link between two Forcepoint Network Security Platform gateways when one or both gateways use multiple network connections. VPN traffic is balanced between the tunnels based on availability and performance checks on each VPN tunnel. If one of the links fails or becomes congested, the VPN traffic is routed through the other tunnels.

The Forcepoint VPN Client can also use Multi-Link. If the ISP connection for one of the gateway endpoints fails, the client automatically connects to the NetLink of the next available endpoint.

The VPN links can be in three different modes: active, aggregate, or standby. If there are multiple links in active mode, traffic is dynamically balanced across the links. The balancing decision can be based on a performance measurement or based on the links' relative bandwidths. In active mode, a single connection uses one of the active links at a time. With multiple connections, all links are used. If there are multiple links in aggregate mode, each connection is balanced on a packet-by-packet basis between all aggregate links in round robin fashion. Standby tunnels are used only if all active or aggregate tunnels become unavailable. Individual tunnels can also be disabled so that they are never used in the VPN.

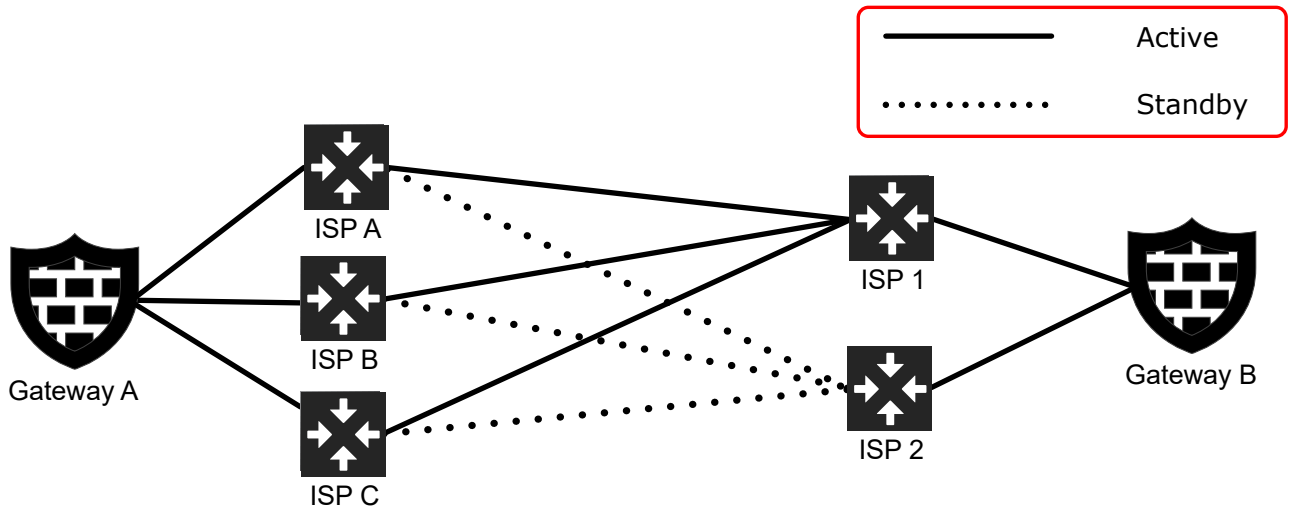


Note

Aggregate mode in a Multi-Link VPN is likely to cause packet reordering due to different latencies of different links. Packet reordering can decrease performance if the TCP stacks of the connection endpoints do not handle reordering well. Use Active mode instead.

This illustration shows a Multi-Link VPN between two VPN Gateways that both have multiple ISP connections. In this configuration, ISP 2 at Gateway B acts as a backup connection for VPN traffic. The three tunnels (one from each ISP at Gateway A) with their endpoints in the ISP 2 network have been set to standby. They are only used if ISP 1 fails. The standby setting is not tied to a particular ISP (NetLink). It is possible to set, for example, only the ISP A to ISP 2 tunnel to active mode while leaving the other tunnels in standby mode.

Example of a Multi-Link VPN with standby tunnels



Related tasks

Change VPN link modes in Multi-Link VPNs on page 1201

Multi-Link packet duplication

The multi-link packet duplication feature enables duplication of traffic packets over multiple links that are sent to the same destination. This eliminates all packet loss due to link failure or delay in packet loss detection

Requirements:

- There must be at least two available links in active mode to the destination gateway.
- The destination gateway must run the engine version that is supported by the packet duplication feature.
- The packet duplication feature is only supported on engine version 7.2 and later.

Important

It is recommended to use this feature only for critical application traffic that is highly sensitive to packet loss and when normal Multi-Link VPN setup packet loss detection is not fast enough to select alternative link on a link failure or on sudden increase in packet loss. Note that when the feature is enabled, it doubles the network bandwidth usage for the traffic selected for packet duplication.

The multi-link packet duplication feature can be enabled by:

1) Selecting what QoS classes are enabled for packet duplication. For more details, refer to the *Create Link Usage Profile elements* topic.

2) (Optional) Controlling what link types should not be used for packet duplication. For more details, refer to the *Create Link Usage Profile elements* topic.

3) Using the configured Link Usage Profile in the engine configuration. For more details, refer to the *Select a Link Usage Profile element for an Engine* topic.

4) Selecting the QoS class in the access rule to match the traffic for multi-link packet duplication. For more details, refer to the *Use QoS Class elements to apply custom link selection options to traffic* topic.

Note

- The receiving gateway forwards only the first received packet. The duplicate packet received later is dropped.
- The link selection for packet duplication is done automatically.
- The maximum transmission unit (MTU) size for the connections subjected to packet duplication is slightly smaller than the MTU size with the normal VPN.
- The configuration done must match in both the tunnel endpoint engines to achieve symmetric packet duplication for both sent and received packets of the critical connections.

Related tasks

Create Link Usage Profile elements on page 789
Select a Link Usage Profile element for an Security Engine on page 790
Use QoS Class elements to apply custom link selection options to traffic on page 787

VPN Broker

The VPN Broker creates highly-scalable, full-mesh VPN environments. VPN tunnels are automatically created between Security Engines when they communicate with each other. The VPN tunnels are automatically removed when they are no longer needed.

You can configure the VPN Broker in the Security Management Center on a dedicated Forcepoint Network Security Platform appliance. The VPN Broker is a component of Forcepoint Network Security Platform.

For more information about VPN Broker, see the *Forcepoint Security Engine Manager and VPN Broker Product Guide*.

Chapter 73

Configuring VPNs

Contents

- VPN configuration overview on page 1169
- Define a custom Gateway Profile element on page 1173
- Defining VPN gateways on page 1173
- Defining Site elements for VPN gateways on page 1180
- Define VPN Traffic Selector elements on page 1184
- Defining VPN profiles on page 1185
- Defining Policy-Based VPN elements on page 1193
- Configuring route-based VPNs on page 1209
- Examples of policy-based VPNs on page 1217
- Examples of route-based VPNs on page 1219

VPNs allow creating secure, private connections through networks that are not otherwise secure.

VPN configuration overview

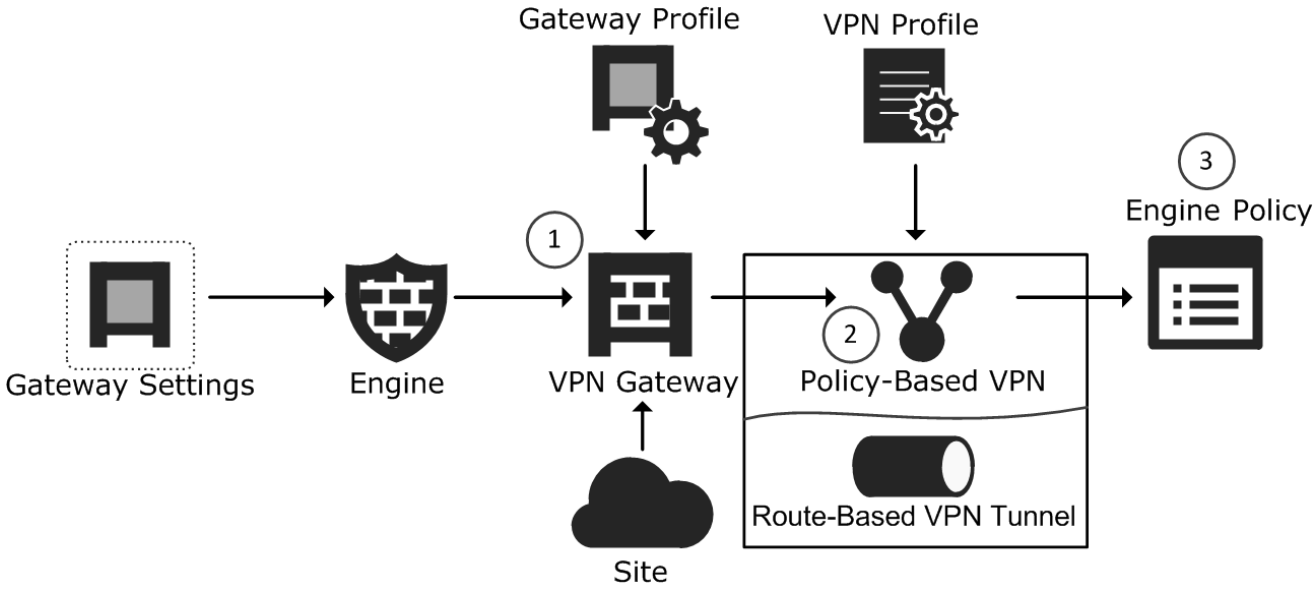
Many steps might be required to configure a VPN, depending on the complexity of the configuration.

Devices that provide VPN access to other computers are called VPN gateways. There are two general types of VPN gateways in the Engine/VPN role:

- *VPN Gateway* elements represent Security Engines that are managed by the Management Server (and administrative Domain) you are currently connected to with your SMC Client.
- All other gateway devices are represented by *External VPN Gateway* elements. Engines that are managed by a different Management Server (or administrative Domain) are also External VPN Gateways.

Due to the various authentication and encryption methods that are supported in VPNs, there are many settings in policy-based VPNs. To prevent repeated configuration work, reusable profiles are used for storing different types of settings. These profiles and other elements related to the configuration of policy-based VPNs are shown in the following illustration, excluding the elements that are related to managing certificates.

Elements in the VPN configuration (excluding certificate-related elements)



1. The VPN Gateway element represents a Engine/VPN role device in VPNs. One VPN Gateway element is automatically created for each Security Engine in the Engine/VPN role. You can optionally add more VPN Gateways to the Security Engine. Each VPN Gateway element can be used in several VPNs. The Gateway element refers to the following other elements:
   - The Security Engine element contains the VPN settings for the VPN Gateway. The Security Engine element refers to a Gateway Settings element that defines settings for advanced VPN performance tuning. The default settings are usually recommended.
   - *Gateway Profile* elements contain information about the capabilities of different gateways, so that the system can disable unsupported options and find incompatible combinations of settings automatically. Gateway Profiles can be created and selected for External VPN Gateways. The Gateway Profiles of VPN Gateways are selected based on the installed software version.
   - *Site* elements define real or translated IP addresses that are routable through the policy-based VPNs. The system can add the IP addresses automatically from routing or you can adjust the sites yourself.
2. The Policy-Based VPN element combines other elements together to define the settings used in one particular policy-based VPN and defines the topology for the VPN.

   Route-based Tunnels elements define endpoints for tunnels in route-based VPNs.

   The VPN elements refer to a VPN Profile, which contains the IPsec authentication and encryption settings (IKE settings) for establishing a VPN.
3. The *Engine Policy* controls policy-based VPN traffic in the same way as any other traffic.
   - The Access rules determine which connections are directed out through each VPN and which traffic is allowed in from each VPN.
   - The NAT rules define how address translation is done for VPN connections. The VPN communications between the gateway devices are always subject to NAT as usual. The traffic that uses the tunnels is subject to NAT only if address translation is enabled for the policy-based VPN.

The same elements used in the configuration of policy-based VPNs can also be used when configuring route-based VPNs.

VPN configuration work flow

VPN configuration requires several high-level steps.

This workflow contains steps for all kinds of VPN configurations. Alternative next steps are included as necessary to achieve a particular type of configuration.

1) (Optional) If you are configuring a VPN with an external device, you might want to create a custom Gateway Profile specific to the device.

2) (External VPN gateways only) Add the necessary number of External VPN Gateway elements to represent the VPN devices. External VPN Gateway elements define the VPN endpoints (gateway IP addresses) and the sites (see the next point).



Note

One VPN Gateway elements is automatically created for each Security Engine that is managed by the Management Server and administrative Domain that you are currently connected to with your SMC Client.

3) (Policy-based VPNs only) Configure the sites. Sites define the IP addresses that can be made routable through VPNs. The sites can be adjusted in different VPNs that the gateway establishes.

4) (Optional) If the existing VPN Profiles do not have suitable settings for your new VPN, create a custom VPN Profile element. The custom VPN Profile element defines the IPsec settings (authentication, encryption, and integrity checking).

5) Define the VPN in one of the following ways:
   - Create a Policy-Based VPN element. The Policy-Based VPN element defines the topology (which gateways create tunnels with each other).
   - Create Route-based Tunnels elements to define endpoints for tunnels in route-based VPNs.

6) Create certificates, if necessary.

7) Add the necessary Access rules according to the type of VPN:
   - (Policy-based VPNs) Add the Access rules that allow traffic and select the policy-based VPN to be used. If necessary, the NAT rules for VPN traffic. Adding rules for policy-based VPNs also activates the VPN on the engines.
   - (Route-based Tunnels) Add Access rules to allow traffic between the internal network and the networks that are reachable through the Route-based Tunnels.

Default elements for policy-based VPNs

There are several default elements for policy-based VPN configuration.

Default elements for policy-based VPN configuration

Element type	Default elements
Certificates	The Internal RSA CA for Gateways VPN Certificate Authority element represents the Management Server's internal RSA certificate authority. You can use the element to define certificate trust relationships if you configure other CAs in the SMC.
Connection Types	The default Connection Type elements represent the Active, Aggregate, and Standby modes for endpoints in a Multi-Link configuration.
Gateways	The predefined VPN Client gateway element that represents VPN clients, including the Forcepoint VPN Client and third-party VPN clients. You can change the Gateway Profile associated with this default element.
Gateway Profiles	Several different Gateway Profiles are included for different Engine/VPN role and Forcepoint VPN Client versions. With third-party VPN devices, you can use the Default (All Capabilities) profile, which enables all options. You can also create a more restrictive profile yourself for better automatic configuration validation.
Gateway Settings	Gateway Default Settings is a predefined Gateway Settings element that contains the default recommended settings for most environments. Each engine has settings that are common to all VPNs the engine establishes, set in the Gateway Settings element. These settings are mostly for performance tuning. Usually there is no need to change them at all. If there is some particular need to change the settings, you must create a custom Gateway Setting element. You cannot edit the Gateway Default Settings system element. There are some advanced properties on the **General** tab, meant for advanced users only. The default values are the recommended values. These options affect the VPN directly.
VPN Profiles	The predefined VPN Profiles are provided to allow you to quickly try out VPNs without creating a VPN Profile yourself. ■ The Forcepoint Cloud Connection VPN profile contains VPN settings forconnecting VPN with Forcepoint Cloud services. ■ The iOS Suite VPN Profile contains only iOS-compatible encryption algorithms and protocols. For example, iOS VPN clients only support IKEv1 key exchange, which must be enabled in the profile. ■ The SOHO Suite VPN profile contains VPN settings for SOHO firewalls. ■ The Suite-B-GCM-128 and Suite-B-GCM-256 VPN Profiles contain the VPN settings specified for the respective cryptographic suites in RFC 6379. ■ The CNSA-GCM-256-ECDH-384, CNSA-GCM-256-DH-3072 and CNSA-GCM-256-DH-4096 Profiles contain the VPN settings specified for the respective cryptographic suites in RFC 9206. The predefined VPN Profiles also allow you to change settings that are not specified in RFC 4308, RFC 6379, and RFC 9206. You might need to adjust the settings to achieve a valid VPN in some configurations.

Define a custom Gateway Profile element

The Gateway Profile element introduces information about the features and options available so that the VPN configuration can be automatically validated.

The general settings directly affect the settings used in VPNs. The authentication and encryption settings defined in the Gateway Profile do not directly influence which of the displayed settings are used for any VPNs. The settings in the Gateway Profile help you make sure that the settings defined for the VPNs correspond to the options supported by the gateway devices involved.

For VPN Gateways that represent engines, the Gateway Profiles are automatically selected according to the software version, and you cannot change the selection. If you use an Security Engine managed by a different Management Server or administrative Domain as an External VPN Gateway, select the Gateway Profile according to the software version. If you use a third-party device as an External VPN Gateway, you have the following options:

- You can use the **Default (all capabilities)** profile, which allows any of the options to be selected for the External VPN Gateway.
- You can define a custom Gateway Profile to restrict the options to a supported set to prevent configuration errors.

For the Forcepoint VPN Client, there are predefined Gateway Profiles.

The **IKE Capabilities** and **IPsec Capabilities** are not directly used in a VPN. The settings are selected for use in the VPN Profile element. The settings define a set of options that the gateway supports, so that the SMC can automatically check for misconfigured settings.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.
2) Browse to **Other Elements** > **Profiles**.
3) Right-click **Gateway Profiles**, then select **New Gateway Profile**.
4) Configure the settings.
5) Click **OK**.

Defining VPN gateways

VPN Gateway and External VPN Gateway elements represent the physical devices that establish the VPN in the configuration.

- *VPN Gateway* elements represent Security Engines that are managed by the Management Server and administrative Domain that you are currently connected to with your SMC Client. One VPN Gateway element is automatically created for each Forcepoint Network Security Platform in the Engine/VPN role. You can optionally add more VPN Gateways to the Engine.

Each VPN Gateway can have multiple VPN endpoints but each endpoint can belong to only one VPN Gateway. For example, using multiple endpoints for a VPN Gateway is required for a Multi-Link VPN configuration.

- *External VPN Gateway* elements represent all other gateway devices. Security Engines that are managed by a different Management Server or administrative Domain are also External VPN Gateway elements. External VPN Gateway elements define settings for the external gateway devices in their role as VPN gateways.

Only one VPN Gateway or External VPN Gateway element is required for each device, even if there are many VPNs. You can use the same Gateway in several different VPNs, possibly overriding some of the Gateway's settings as necessary. You can create several Gateway elements to represent the same Engine. However, each Gateway element reserves a VPN endpoint (IP address) that other Gateway elements cannot use. You cannot use the same pair of endpoints for VPN tunnels in several configurations for a single Security Engine.

The predefined VPN Client element represents all instances of the Forcepoint VPN Client and third-party IPsec VPN clients in mobile VPNs. When you set up a mobile VPN with the Forcepoint VPN Client, the VPN Client element must always be used. Usually, we recommend using the element with third-party VPN clients as well. However, it is possible to configure an individual third-party VPN client using an External VPN Gateway element if there is a specific need to do so. In this configuration, only one client at a time can connect to each gateway.

Using a dynamic IP address for a VPN endpoint

The following restriction applies when a VPN endpoint has a dynamic IP address that has been assigned using DHCP, PPPoA, or PPPoE.

IKEv1 main mode with pre-shared key authentication is not supported. Aggressive mode allows the use of pre-shared keys, but for security reasons certificate-based authentication is also recommended when IKEv1 is set in aggressive mode. Always use IKEv2 if both VPN endpoints support it.

Using a NAT address for a VPN endpoint

VPN traffic is protected against modifications, so there are some restrictions when NAT is applied to the encrypted traffic.

If a gateway does not have a public IP address as a VPN endpoint, you might need to configure NAT traversal.

You might also need to configure the VPN with contact addresses so that the gateways are aware of the NAT operation:

- Engines that are used as VPN Gateways in a NAT environment must have Locations and Contact Addresses defined for the endpoint interfaces involved. On Engine Clusters, CVIs must have Locations and Contact Addresses defined. If Contact Addresses have already been configured for non-VPN use, the same general configuration applies to VPN communications as well. The Forcepoint VPN Client downloads its configuration from the gateway, including any contact address configuration as needed.
- Usually, External VPN Gateways must be defined using their private IP addresses. The public IP address must be added as the Contact Address for the Location of the contacting Forcepoint Network Security Platform in the Engine/VPN role.

Define endpoints for VPN Gateway elements

Each endpoint is dedicated for one VPN Gateway element.

Any IP address that is already an endpoint for another VPN Gateway element is not shown on the **Endpoints** list for other Gateways that you create for the same Security Engine. Each VPN Gateway element can be used in

several policy-based VPNs or Route-based Tunnels. However, you cannot use the same pair of local and remote endpoints in different VPN configurations for the same Security Engine.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click the Engine, then select **Edit Single Engine** or **Edit Engine Cluster**.

3) Browse to **VPN** > **Endpoints**.

4) (Optional) Change the selection of IP addresses that you want to use as endpoints in VPNs.
   - Typically, these are IP addresses that belong to interfaces toward the Internet, which are automatically selected based on the engine's default routing table.
   - If loopback IP addresses are defined for the Security Engine, you can select a loopback IP address as the endpoint IP address. On clustered engines, the IP addresses are CVIs.
   - (Optional) If you have more than one Internet connection, select an IP address from each ISP.

5) Double-click the endpoint, then configure the following optional settings according to your environment.
   a) (Optional) In the **Name** field, enter a descriptive name for the endpoint.

   b) (Multi-Link tunnels only) From the **Connection Type** drop-down list, select the Connection Type element that defines how the endpoint is used in a Multi-Link configuration.
      You can override these settings in each individual VPN.

   c) (Optional) From the **Use NAT-T** drop-down list, select an option to activate encapsulation for NAT traversal in site-to-site VPNs.
      You might need NAT traversal to traverse a NAT device at the local or at the remote gateway end. The gateway always allows VPN clients to use NAT-T regardless of these settings. NAT-T always uses the standard UDP port 4500.

      Note

      If a private external IP address is translated to a public IP address by an external NAT device, make sure that Contact Addresses and Locations are defined for the Engine.

6) In the **Phase-1 ID** settings, select an option from the **ID Type** drop-down list according to your environment.
   The ID identifies the Gateways during the IKE SA negotiations.

7) In the **ID Value** field, enter an ID value according to the selected ID type.
   - If you selected **DNS Name**, enter a DNS name.
   - If you selected **Email**, enter an email address.
   - If you selected **Distinguished Name**, enter the distinguished name that is used in the gateway certificate.
   - If you selected **IP Address** and the endpoint has a static IP address, the value is filled in automatically. If the endpoint has a dynamic IP address, you must manually enter an IP address.

8) (Optional) If the endpoint must use different Phase-1 ID settings in individual policy-based VPNs, add VPN-specific exceptions.
   a) Click **Exceptions**.
   b) Click **Add**, then select the type of ID from the drop-down list.
   c) Select a **Policy-Based VPN** element, then click **Select**.
   d) In the **ID Value** cell, enter the value of the ID.
   e) Click **OK**.
9) (Optional) In the **VPN Type** settings, restrict the types of VPNs that the endpoint can be used in.
   a) Select **Selected types only**.
   b) Select one or more types of VPNs.
10) Click **OK** to save your changes to the endpoint.
11) Save the changes.
   - To save the changes, click **Save**.
   - To save the changes and refresh the security policy on the engine, click **Save and Refresh**.

Create an External VPN Gateway element

External VPN Gateway elements represent third-party VPN devices or Forcepoint Network Security Platform devices managed by a different Management Server. To use third-party VPN devices or Forcepoint Network Security Platform devices managed by a different Management Server in VPNs, create an External VPN Gateway element.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.
2) Right-click **VPN Gateways**, then select **New External VPN Gateway**.
3) Configure the settings.
4) Click **OK**.

Define endpoints for External VPN Gateways

Each endpoint is dedicated for one External VPN Gateway element.

Before you begin

You must have an External VPN Gateway element.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.
2) Browser to **External VPN Gateway**.
3) Right-click the **External VPN Gateway** element, then select **Properties**.
4) On the **Endpoints** tab, click **Add**.
5) Configure the following optional settings according to your environment if needed.
   a) (Optional) In the **Name** field, enter a descriptive name for the endpoint.
   b) (Policy-Based VPNs only) From the **Connection Type** drop-down list, select an option to define how the endpoint is used in a Multi-Link configuration.
      You can override these settings in each individual VPN.
   c) (Optional) From the **Use NAT-T** drop-down list, select an option to activate encapsulation for NAT traversal in site-to-site VPNs.
      You might need NAT traversal to traverse a NAT device at the local or at the remote gateway end. The gateway always allows VPN clients to use NAT-T regardless of these settings. NAT-T always uses the standard UDP port 4500.

      Note

      If a private external IP address is translated to a public IP address by an external NAT device, make sure that Contact Addresses and Locations are defined for the Engine.

   d) If necessary, change the default Contact Address or add Exceptions for the Locations of other gateways involved in the VPN.
      The Contact Address must be defined if the IP address for contacting this gateway is different from the IP address that the gateway actually has on its interface (for example, because of NAT).
      Example: An external gateway is behind a NAT device. The real address is defined as the endpoint address, because the IP address is also used as the Phase 1 ID inside the encrypted traffic. Contact must be made using the translated address, so it is defined as a Contact Address.
6) In the **Phase-1** settings, select an option from the **ID Type** drop-down list to according to your environment.
   The ID identifies the Gateways during the IKE SA negotiations. The **IP Address** might not work as an ID if the address is translated using NAT.

7) In the **ID Value** field, enter an ID value according to the selected ID type.

Note

Make sure that the ID value matches the identity configured on the external gateway device.

- If you selected **DNS Name**, enter the DNS name that is configured on the external gateway device.
- If you selected **Email**, enter the email address that is configured on the external gateway device.
- If you selected **Distinguished Name**, enter the distinguished name that is used in the gateway certificate.
- If the endpoint has a dynamic IP address, enter a specific IP address as the value for the **IP Address** type.

Note

If the endpoint has a static IP address, the value for the **IP Address** type is filled in automatically.

8) Click **OK** to save your changes to the endpoint.

Restrict the trusted CAs for a VPN gateway

Certificate Authorities (CA) verify certificate authenticity with their signatures. By default, the gateways trust all VPN CAs, but you can restrict the trusted CAs.

Before you begin

You must have more than one VPN Certificate Authority element.

When you restrict the trusted CAs for a VPN gateway, the VPN gateways accept certificates only from the trusted CAs that you select. When you restrict the trusted CAs for an external VPN gateway, the system uses the trusted CA definition in the External VPN Gateway element to check that all gateways have the necessary certificates.

Tip

You can also restrict trusted CAs in VPN Profiles.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Access the **Trusted VPN Certificate Authorities** settings in one of the following ways:
   - Right-click a Engine element, select **Edit <element type>**, then browse to **VPN** > **Certificates**.
   - Right-click an **External VPN Gateway** element, select **Properties**, then click the **Trusted CAs** tab.

2) Select **Trust only selected**, then select one or more CAs.

3) Save the changes in one of the following ways:
   - In the Engine Editor, click **Save**.

- In the **External VPN Gateway Properties** dialog box, click **OK**.

Related tasks
Define additional VPN certificate authorities on page 1254

Define VPN client settings for Security Engines

VPN client settings in the Engine Editor define the settings that are used when the Security Engine acts as a VPN Gateway in a mobile VPN.

If you use Forcepoint VPN Client, configure the Virtual Adapter. The alternative NAT Pool method does not allow the Forcepoint VPN Client computers to use your organization's internal DNS servers. Virtual IP addresses work with all Forcepoint VPN Client versions and with third-party VPN clients that support this feature.

If you use Forcepoint VPN Client, the policy-based VPN configuration defined in the SMC Client is also used for creating the configuration for Forcepoint VPN Client. Forcepoint VPN Client downloads the settings from the VPN gateway the first time that the Forcepoint VPN Client connects to the VPN Gateway, and automatically whenever there are relevant changes. All IPsec and address management settings are included in the download. For example, the download includes information about which encryption methods are used, which VPN endpoints are available, and which internal networks clients can access through the gateway. The decision whether a VPN tunnel is used is based on the IP addresses you have defined for the Sites of the gateway.

For third-party VPN clients and external VPN gateways, you must duplicate the VPN Gateway settings in the configuration of the VPN client or gateway. You must also duplicate the VPN Gateway settings for engines under a different administrative Domain. The settings that you must duplicate include the following:

- All IPsec-related settings, such as the authentication, encryption, and integrity checking options.
- The encryption domain (the IP addresses that are allowed in the VPN as a source or destination IP address).

If a VPN Gateway that contains VPN Client settings is used in a Route-based Tunnels, the VPN Client settings are ignored.

Note

The Virtual Adapter IP addresses must be assigned by a DHCP server. It is not possible to define the IP addresses in the VPN client or in the VPN gateway configuration. When you use a Single Engine's internal DHCP server, use the IP address of the interface on which the internal DHCP server is enabled as the IP address of the DHCP Server element.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click a Engine element, then select **Edit <element type>**.

3) Browse to **VPN** > **VPN Client**.

4) Configure the settings.
   If you selected a VPN Mode that includes SSL VPN, configure the settings in the Virtual Address section.

5) Click **Save**.

Related concepts
Defining IP addresses for VPN clients on page 1279

Defining Site elements for VPN gateways

The Site element defines the internal IP addresses that can send or receive traffic through the VPN.

Note

In route-based VPNs, the site information is ignored. In Route-based Tunnels the site definition is always 0.0.0.0/0 for IPv4 and ::/0 for IPv6 (any network).

The IP addresses work like routing definitions when the gateway selects which VPN tunnel a packet is sent through. The Site elements must contain the IP addresses of all protected hosts that potentially send or receive VPN traffic through any site-to-site or mobile VPN. IP addresses that are not included in the Site elements are not allowed as source or destination addresses in policy-based VPNs.

Note

An IP address must be included in a Site to be valid in the VPN. The Access rules define which connections are allowed to enter and exit a VPN tunnel.

By default, each site is included in all VPNs where the gateway is used. You can manually disable individual sites in individual VPNs without affecting other VPNs. It is not possible to partially disable sites. If the IP address space must be different in different VPNs, you need several sites. You can define as many Site elements as you need.

Automatic VPN Site management

The VPN settings for Security Engines include a Site that is automatically populated and updated according to the routing definitions.

All interfaces and networks are included in the automatic Site, except interfaces with the Any Network element. If loopback IP addresses are defined for the engine, you can use a loopback IP address as an endpoint IP address.

You can change this automatic Site in the following ways:

- You can disable individual interfaces through their right-click menu. This way, you can exclude some of the internal interfaces from VPNs.
- You can add addresses to the automatic Site at the top level (at the same level with the Interface elements, not inside them) by dragging and dropping the correct Networks or other elements.
- You can add more Sites alongside the automatic Site.
- You can define the automatic Site as Private in some VPNs.

Disable or enable automatic VPN Site management

Automatic Site management is active by default in the VPN settings for Security Engines.

If you prefer not to update the information automatically for any interface, you can disable automatic site management.

When you disable automatic site management, the automatic Site is removed. There must be another Site configured for the gateway for it to be valid in a policy-based VPN.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click a Engine element, then select **Edit <element type>**.
3) Browse to **VPN** > **Sites**.
4) Deselect or select the **Add and Update Addresses Based on Routing** option.
   - When the option is not selected, you must manually define the addresses that you want to be routable through the VPN.
   - When the option is selected, the Site content updates automatically according to changes made in the routing configuration for the engine (for interfaces that are not disabled).
5) Click **Save**.

Define a VPN site

You must define Site elements for all Security Engines and External VPN Gateways that are used in policy-based VPNs. You must also define sites for Security Engines and External VPN Gateways that are used in Route-based Tunnels in which the value of the **Encryption** option is **Tunnel Mode**.

The Site elements must always contain the actual IP addresses that are used inside the VPN tunnel. If traffic in the tunnel is subject to NAT, you must add the NAT addresses to the site. For Security Engines, you must add both the NAT addresses and any untranslated IP addresses that are not automatically added to the site. Sites for External VPN Gateways only require the translated address space that the Security Engine actually contacts.

The local and remote site definitions must match the same information about the other gateways involved in the VPN because the gateways verify this information during IKE negotiation. When creating VPNs with external Gateways, make sure that the IP address spaces of both gateways are defined identically in the SMC and on the external device. Otherwise, the VPN establishment can fail in one or both directions. Make sure to update the policies of any engines that are involved in the VPN when there are changes in the Site elements at either end.

If you want to use a central gateway as a hub that forwards traffic from one VPN tunnel to another, include all IP addresses that are accessible through the central gateway in the central gateway's Site elements.



Note

You cannot add or change Site elements under the VPN Client Gateway element.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.
2) Browse to **VPN Gateways**.
3) Right-click a **VPN Gateway** or an **External VPN Gateway**, then select **New** > **Site**.
4) Select the elements that represent the protected IP addresses behind the Gateway, then click **Add** to include them in this site.
   - Do not include IP addresses outside the Gateway's local networks in the site. There is no need to include the Gateways' own IP addresses in the sites. However, there is usually no need to exclude those addresses if they are in the networks you add to the site.
   - IP address ranges might be interpreted differently from lists of IP addresses and networks depending on the VPN device. The system converts Group or Expression elements into address ranges, networks, or individual IP addresses depending on the IP addresses included. Other VPN devices might treat the same types of values differently.
   - VPN Traffic Selector elements allow you to define the IP addresses, protocols, and ports used by a specific host in a VPN site.
5) Click **OK**.

Next steps

If you edited a previously configured VPN, make sure that the configuration of any external VPN gateway device involved contains the same IP address information. Refresh the policy on all affected gateways to transfer the changes.

Adjust VPN-specific Site settings for VPN Gateways or External VPN Gateways

Site elements allow you to adjust how the Site is used in each VPN.

Before you begin

Manually added Site elements to VPN Gateways.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Open the list of sites for the gateway in one of the following ways:
   - Right-click an Security Engine, select **Edit <element type>**, then browse to **VPN** > **Sites**.

- Browse to **Secure SD-WAN Configuration** > **VPN Gateways**, double-click the **External VPN Gateway** element, then click the **Sites** tab.

3) Right-click a manually added Site, then select **Properties**.

4) Configure the settings.

5) Save the changes in one of the following ways:
   - In the Engine Editor, click **Save and Refresh**.
   - In the **External VPN Gateway Properties** dialog box, click **OK**.

Temporarily disable a VPN site in all VPNs

You can disable a site that has been manually added to the Gateway. The site is disabled globally in all VPNs.

To remove the automatic site from an Security Engine that acts as a VPN Gateway, disable automatic VPN site management. There must be at least one enabled site.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Open the list of sites for the gateway in one of the following ways:
   - Right-click an Security Engine, select **Edit <element type>**, then browse to **VPN** > **Sites**.
   - Browse to **Secure SD-WAN Configuration** > **VPN Gateways**, double-click the **External VPN Gateway** element, then click the **Sites** tab.

3) Right-click the site, then select **Properties**.

4) On the **VPN References** tab, deselect the **Enable** cell.

5) Click **OK**.

6) Save the changes in one of the following ways:
   - In the Engine Editor, click **Save and Refresh**.
   - In the **External VPN Gateway Properties** dialog box, click **OK**.

Next steps

If you edited a previously configured VPN, refresh the policy on all affected gateways to transfer the changes. The configurations of external gateways might also require an update.

Remove a VPN site from all VPNs

You can remove a site that has been manually added to the VPN gateway. The site is removed from all VPNs where the VPN gateway is used.

Note

To remove the automatic site from an Security Engine that acts as a VPN Gateway, disable automatic VPN site management. There must be at least one enabled site.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Open the list of sites for the gateway in one of the following ways:
   - Right-click an Security Engine, select **Edit <element type>**, then browse to **VPN** > **Sites**.
   - Browse to **Secure SD-WAN Configuration** > **VPN Gateways**, double-click the **External VPN Gateway** element, then click the **Sites** tab.

3) Right-click the site, then select **Remove**.

4) Save the changes in one of the following ways:
   - In the Engine Editor, click **Save and Refresh**.
   - In the **External VPN Gateway Properties** dialog box, click **OK**.

Next steps

If you edited a previously configured VPN, refresh the policy on all affected gateways to transfer the changes. The configurations of external gateways might also require an update.

Related tasks

Disable or enable automatic VPN Site management on page 1181

Define VPN Traffic Selector elements

VPN Traffic Selector elements allow you to define the IP addresses, protocols, and ports used by a specific host in a VPN site.

Before you begin

You must have sites for VPN Gateways.

You can use VPN Traffic Selector elements in policy-based VPNs. VPN Traffic Selector elements are intended to be used in specific advanced VPN configurations.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Open the list of sites for the gateway in one of the following ways:
   - Right-click a Engine element, select **Edit <element type>**, then browse to **VPN** > **Sites**.
   - Open the properties of the **External VPN Gateway** element, then click the **Sites** tab.
2) Select ⋮ **More actions** > **New** > **VPN Traffic Selector**.
3) Configure the settings.
4) Click **OK**.

Defining VPN profiles

VPN Profile elements contain settings related to authentication, integrity checking, and encryption.

The VPN Profile element is the main point of configuration for IKE and IPsec settings. These settings are used or agreed on during IKE SA and IPsec SA negotiations. You can select any combination of settings as long as all gateways and VPN clients involved in the VPN support those settings and are configured to accept them.

The authentication methods for VPN clients are selected separately in the VPN Profile. A certificate-based method is always included in the VPN, but you can optionally add other authentication methods.

If you want to use certificates signed by a particular certificate authority (CA), you must define the CA as an element. By default, all VPN CAs are considered trusted, but you can restrict the trusted CAs for particular VPNs.

Each VPN refers to a VPN Profile. You can use the same VPN Profile in several VPNs if the settings are compatible. You can use the same VPN Profile in both policy-based and route-based VPNs. You can also easily copy the element to create custom versions of the same basic settings. There are predefined VPN Profile elements, which are mostly useful for site-to-site VPNs between engines that act as VPN Gateways.

Mobile VPNs usually require a custom profile. However, there is a predefined VPN Profile element that simplifies configuration for VPNs between iOS devices and VPN gateways. The iOS Suite VPN profile contains only iOS-compatible encryption algorithms and protocols.

Before editing a VPN Profile that is used in active VPNs, we recommend backing up the settings. You can back up the settings by duplicating the element, exporting it, or creating a Management Server backup. After editing a VPN profile that is used in active VPNs, check all VPNs that use the profile for issues that the changes might have caused.

Security associations (SA) in IPsec VPNs

The settings that are used for a tunnel are stored in *Security Associations* (SA). There are two SAs for each IPsec VPN tunnel: one for outgoing traffic, and another one for incoming traffic.

For any communications to be able to use the VPN, the gateways must construct and maintain the VPN tunnels. The gateways negotiate which settings to use between each other. The gateways store this information so that it can be used for handling the traffic throughout the lifetime of the VPN tunnel.

The term *SPI* (security parameter index) is sometimes used with SAs in IPsec VPNs. SPIs are used to identify the SAs.

For security reasons, each SA has an expiration time. After the expiration time, the gateways discard the old SAs and agree on new ones if there is still traffic going through the VPN.

Internet key exchange (IKE) in IPsec VPNs

SAs for IPsec VPNs are created in a process called the *Internet key exchange* (IKE) negotiations.

During the IKE negotiations, the VPN gateways negotiate the parameters to use, such as the encryption keys and the authentication methods. This information is then stored in the SAs. Both IKEv1 and IKEv2 are supported with Forcepoint Network Security Platform.

The IKE negotiations consist of two phases:

- Phase 1 — During the *IKE SA* negotiations, the gateways authenticate themselves to each other and establish a secure (encrypted) channel for the IPsec SA negotiations. Authentication in IKE SA negotiations can be done with signatures, or with pre-shared keys. These parameters are then stored in IKE SAs.
- Phase 2 — During the *IPsec SA* negotiations, the gateways select parameters for encrypting the traffic going through the VPN tunnels. These parameters are then stored in IPsec SAs.

The IPsec SA negotiations are much faster than the IKE SA negotiations. Because IKE SA negotiations involve heavy computation, it is common to configure the IKE SAs to expire less frequently than the IPsec SAs.

IKEv2 also provides support for *IKEv2 Mobility and Multihoming Protocol* (MOBIKE) protocol. MOBIKE enables transparent recovery for VPN clients when the VPN clients change their IP addresses. For example, the IP address can change when a laptop is connected to a different network while a VPN connection is open. MOBIKE also allows the IP addresses associated with IKE SAs and IPsec SAs to change in a VPN Multi-Link configuration. When a VPN client fails to connect to a gateway, it checks if another gateway address is available. If the VPN client can connect using the new gateway address, the gateway's IP address is updated in the IKE SAs and the IPsec SAs. VPN traffic can continue uninterrupted. There is no need to renegotiate the SAs.

Supported encryption methods for VPNs

Select encryption settings in your VPN according to the guidelines in your organization's security policy.

The message digest algorithms (for integrity checking) and encryption methods that are available in VPNs are listed. The IPsec standards mandate support for some options, but also allow other options to be provided by IPsec-compatible products. RFC 8221 lists the IPsec standard requirements that all IPsec-compliant products must follow.

Estimates of how common support for the various algorithms is in IPsec-compatible products are listed. This information can be helpful when deciding which methods to use when establishing a VPN with a third-party VPN device.

If your organization is required to follow FIPS encryption standards, some of the options presented are not available in your system. See the *Common Criteria Certification User's Guide* for more information.

Supported message digest algorithms for IPsec VPNs

Message digest algorithms are used to guarantee the integrity of data (that the packets have not been changed in transit). These algorithms are often also referred to using the MAC or HMAC abbreviations (keyed-hash message authentication code).

Supported message digest algorithms

Algorithm	Description
AES-XCBC-MAC	128-bit hash algorithm. Available only for checking the integrity of IPsec traffic. Many IPsec-compatible VPN devices do not support this algorithm, but support is becoming increasingly common. Reference: RFC 3566.
MD5	Message-Digest algorithm 5 A 128-bit hash algorithm (also referred to as HMACMD5). Available for checking the integrity of the IKE negotiations and IPsec traffic. Most IPsec-compliant VPN devices still support this algorithm, but it is not considered secure and should not be used unless required for compatibility with legacy devices. Reference: RFC 2403.
SHA-1	Secure Hash Algorithm Has a 160-bit hash (sometimes referred to as HMAC-SHA-1). Available for checking the integrity of the IKE negotiations and IPsec traffic. All VPN devices must support this algorithm to be fully IPsec-compliant. Reference: RFC 2404.
SHA-2	Secure Hash Algorithm Has 256-bit, 384-bit, and 512-bit hashes (includes SHA- 224, SHA-256, SHA-384, and SHA-512). Available for checking the integrity of the IKE negotiations and IPsec traffic. All VPN devices must support this algorithm to be fully IPsec-compliant. Reference: RFC 4868.

Supported encryption algorithms for IPsec VPNs

Encryption algorithms scramble data, so that it is not viewable while in transit.

Supported encryption methods

Method	Description
AES-128 AES-256	Advanced Encryption Standard (also referred to as Rijndael) with a 128-bit/192-bit/256- bit encryption key. The AES-128 option uses 128-bit keys by default, but accepts stronger 192-bit or 256-bit keys if requested by the other gateway. All VPN devices must support this algorithm to be fully IPsec-compliant. Reference: RFC 3602.
AES-GCM-128 AES-GCM-256	Advanced Encryption Standard (also referred to as Rijndael) in GCM (galois/counter mode), uses a 16-octet ICV (integrity check value). AES-GCM-128 uses a 128-bit encryption key, and AES-GCM-256 uses a 256-bit encryption key. Provides both authentication and encryption. These methods replace the selected message digest algorithm. In high-performance networks, these encryption methods are recommended. All VPN devices must support this algorithm to be fully IPsec-compliant. Reference: RFC 4106.
DES	Data Encryption Standard (also referred to as the Data Encryption Algorithm, DEA) uses a 56-bit encryption key. Do not use DES if you can avoid it. DES has been largely abandoned because the short key makes it vulnerable to attacks. If you must use DES, make sure that PFS is enabled and that the encryption keys are frequently renegotiated. Many IPsec-compliant VPN devices still support this method, but it is not considered secure and should not be used unless required for compatibility with legacy devices. Reference: RFC 2405.
Blowfish	Uses up to 448-bit keys. Policy-based VPNs use 128-bit keys by default, but accept up to 448-bit keys if requested by the other gateway. Many IPsec-compatible VPN devices do not support this method. Reference: RFC 2451.
3DES	Triple-DES (also referred to as TDES or TDEA, Triple Data Encryption Algorithm), uses 168-bit encryption achieved with three different 56-bit encryption keys. 3DES is processor-intensive compared to the level of protection it offers and is therefore not the most efficient method. It might not be optimal for VPNs that handle large traffic volumes or systems that otherwise have a high load. Most IPsec-compliant VPN devices support this method. Reference: RFC 2451.

Method	Description
Null	No encryption. Traffic is sent as cleartext just like any non-VPN traffic. Anyone who intercepts the VPN traffic in transit can view the traffic. Null encryption is useful only in special cases. Do not select it for any VPN that requires protected data transfer. Most IPsec-compliant VPN devices support this method. Reference: RFC 2410.

Configuring Post-quantum Pre-Shared Key

The Post-quantum Pre-Shared Key (PPK) is used in addition to the Diffie-Hellman (DH) key exchange in the IKEv2 protocol to prevent threats from quantum computers.

It is believed that a quantum computer can in the future break the DH algorithm and find the secret key used to protect communication. PPK adds a shared secret to the DH result mitigating the threat. Hence, when PPK is used with the IKEv2 protocol it prevents traffic decryption even if the encrypted traffic is stored and later a quantum computer is used to decrypt the traffic.

A PPK consists of a Key ID and a Pre-shared secret. For more details on PPK, refer to RFC 8784.

Note

- Devices in communication share the Key ID during the IKEv2 handshake, but the pre-shared secret is shared only as part of the configuration.
- When the same SMC is managing two gateways and only one of the two gateways has the PPK configured, the SMC automatically creates a new PPK element and then selects it for the gateway that does not have PPK configured. Also, this new PPK element has an auto-generated unique primary PPK ID.
- The PPK feature is not supported for VPN broker tunnels.

Requirements:

- To establish a VPN tunnel between devices using PPK with the IKEv2 protocol, both devices involved in the communication must support the PPK feature.
- The PPK feature is only supported on Engine versions 7.3 and later.
- The PPK feature is only supported for route-based and policy-based tunnels.

Steps

1) Create a PPK element. For more details, refer to the *Create Post-quantum Pre-shared Key element* topic.

2) Configure PPK settings in the VPN profile. For more details, refer to the *Create VPN Profile elements* topic.

3) Select the VPN profile that is configured with the PPK settings and the PPK element for the VPN tunnel. For more details:
   - On Policy-based tunnel, refer to the *Define VPN tunnel settings for policy-based VPNs* topic.
   - On Route-based tunnel, refer to the *Create Route-based Tunnels elements* topic.

Create Post-quantum Pre-shared Key element

A Post-quantum Pre-shared Key (PPK) element can be used along with the IKEv2 protocol to prevent threats from quantum computers. A PPK consists of a Key ID and a Pre-shared secret.

Steps

1) Select **Secure SD-WAN Configuration**.
2) Browse to **Other Elements** > **Post-quantum Preshared Keys**.
3) Right-click **Post-quantum Preshared Keys**, then select **New Post-quantum Preshared Key**.
4) Configure the settings.
5) 5. Click **OK**.

Update the Post-quantum Pre-shared Key

The pre-shared secret key is associated with the primary Post-quantum Pre-shared Key (PPK) ID. Follow the steps below to update the secret key.

Important

To change the primary pre-shared secret on both gateways, you must define the new secret in the **Secondary PPK** section and then swap the secrets.

Steps

1) Select **Secure SD-WAN Configuration**.
2) Browse to **Other Elements** > **Post-quantum Preshared Keys**.
3) Right-click **Post-quantum Preshared Keys** element, then select the **Properties**.

4) Do one of the following:

   a) If you want to update the **Primary PPK** section with a new pre-shared secret:

      i) Update the **Secondary PPK** section with the details of the new pre-shared secret.

      ii) Click the **OK** button.

      iii) Refresh the current policy. For more details, refer to the *Refresh the currently installed policy* topic.

      iv) Re-open the PPK element properties.

      v) Click the **Swap** button.

      vi) Click the **OK** button to save the changes.

   b) If both the **Primary PPK** section and the **Secondary PPK** sections are already populated and you only want to update the **Primary PPK** section with the existing **Secondary PPK** section details:

      i) Click the **Swap** button.

      ii) Click the **OK** button to save the changes.

5) Refresh the current policy. For more details, refer to the *Refresh the currently installed policy* topic.

Related tasks

Refresh the currently installed policy on page 353

Create VPN Profile elements

If the default VPN Profile elements do not meet your VPN configuration needs, create a custom VPN Profile element.

The options you select are a balance between performance and security. A higher level of security generally requires more processing power.

If External VPN Gateways are involved, you must make sure that all settings match between the gateways.

Limitations:

- The FIPS mode engine allows only the following algorithms for VPN:
  - AES with 256 bits key.
  - If AES-GCM-256 is selected for IKE, then the digest must be SHA_512.
  - SHA-384 or larger.
- The FIPS mode engine allows only the following authentication methods for VPN:
  - Certificate using ECDSA-384
  - Certificate using RSA-3072 or larger

Note

If VPN peer requests authentication method that is not allowed, Engine responds with AUTHENTICATION_FAILED.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.
2) Browse to **Other Elements** > **Profiles** > **VPN Profiles**.
3) Right-click **VPN Profiles**, then select **New VPN Profile**.
4) Configure the settings.
5) Click **OK**.

Export iOS VPN configuration profiles

Exporting iOS VPN configuration profiles simplifies the VPN client configuration for users who have iOS or iPadOS devices.

Before you begin

- The mobile VPN must be fully configured.
- The mobile VPN must use the iOS Suite VPN profile.

Note

The iOS VPN configuration profile is for use with the native iOS or iPadOS VPN client only. Do not use iOS VPN configuration profiles with the Forcepoint VPN Client for macOS.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.
2) Browse to **Policy-Based VPNs**.
3) Right-click the **VPN** to which iOS VPN client users connect, then select **More actions** > **Export iOS VPN configuration Profile**.
4) Select where to save the configuration file.

5) Select the **VPN gateway** and endpoint to which iOS VPN client users connect.

6) (Optional) To require users to enter a password to open the configuration file, enter a password in the **Password** field.

7) Click **Export**.

Next steps

Make the configuration file available to the iOS VPN client users. The VPN client setup procedure automatically starts when the users open the configuration file on the iOS device.

Defining Policy-Based VPN elements

The Policy-Based VPN element collects together the gateways and the VPN Profile, and provides the settings for defining the topology and the tunnels of the policy-based VPN.

The configuration of a Policy-Based VPN element has two stages: first you define some basic properties for the element, then you can add gateways and adjust the tunnels.

The main configuration for the VPN consists of defining which gateways are in the VPN and which of the gateways form tunnels with each other. You can also enter and renew pre-shared keys if you use them for authentication in this VPN.

The Sites and networks for each gateway element can be adjusted in the policy-based VPN, but most of the settings are not specific to the Policy-based VPN. The only change that is specific to the policy-based VPN is to disable a Site element in the Policy-based VPN. Disabling a Site excludes the IP addresses from that policy-based VPN only. Any other adjustments to the Sites and networks affect all other VPNs where the same gateway element is used.

You can also change some of the properties for tunnels between two particular gateways or endpoints, such as the VPN Profile used. You can also define Multi-Link VPN settings that allow you to select standby and active tunnels and to set tunnels to aggregate mode. In aggregate mode, each connection is automatically balanced between the aggregate tunnels.



Note

Although the VPN endpoints usually correspond to the NetLink interfaces in a Multi-Link configuration, the VPN endpoint settings are separate from the NetLink and Outbound Multi-Link definitions. For example, a NetLink that is set to standby for non-VPN traffic in an Outbound Multi-Link element can still be used as an active endpoint for VPN traffic.

Topologies for policy-based VPNs

The topology of Policy-based VPNs is determined by selecting whether individual VPN gateways are central or satellite gateways in each particular policy-based VPN.

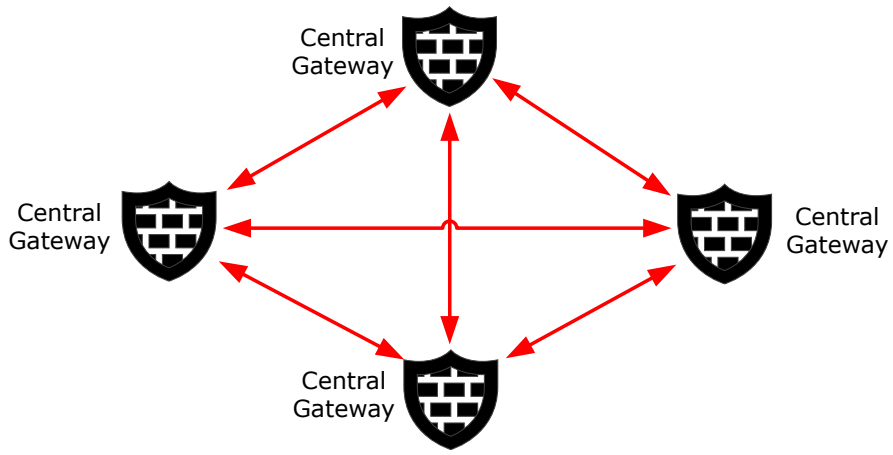
- A *central gateway* establishes VPN tunnels with any other central or satellite gateway in the VPN, unless you specifically disable the tunnels.
- A *satellite gateway* establishes VPN tunnels only with central gateways.

- You can also create a VPN hub by adding a gateway so that it is listed under some other (central or satellite) gateway in the topology. Other gateways connect to the higher-level gateway, which forwards the connections to the lower-level gateway.

Tunnels are generated from each central gateway to all other gateways based on the overall topology. You can adjust the tunnels to limit which gateways and endpoints form tunnels with each other.

You can define policy-based VPN tunnels using different topologies:

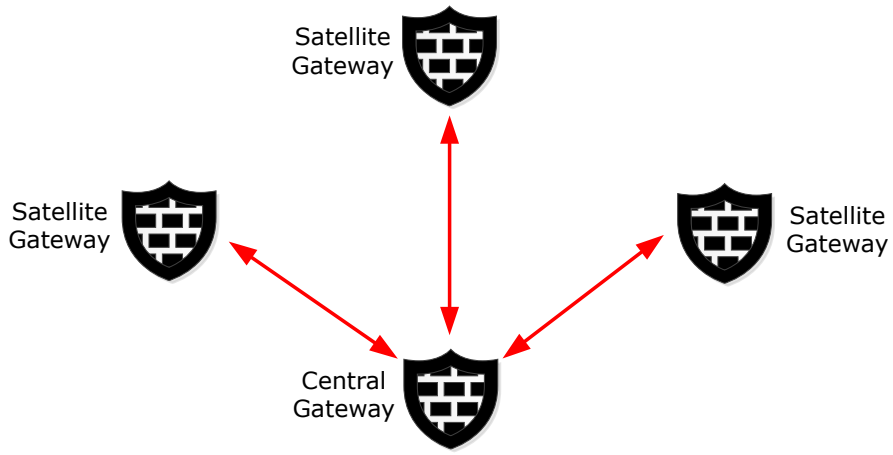
Full mesh VPN topology



Full-mesh topology connects each site to every other site in the same VPN. All gateways are central gateways, which means that all gateways can establish tunnels with all other gateways in the VPN.

The full mesh topology is formed between sites that must all be able to connect to any other site.

Star VPN topology



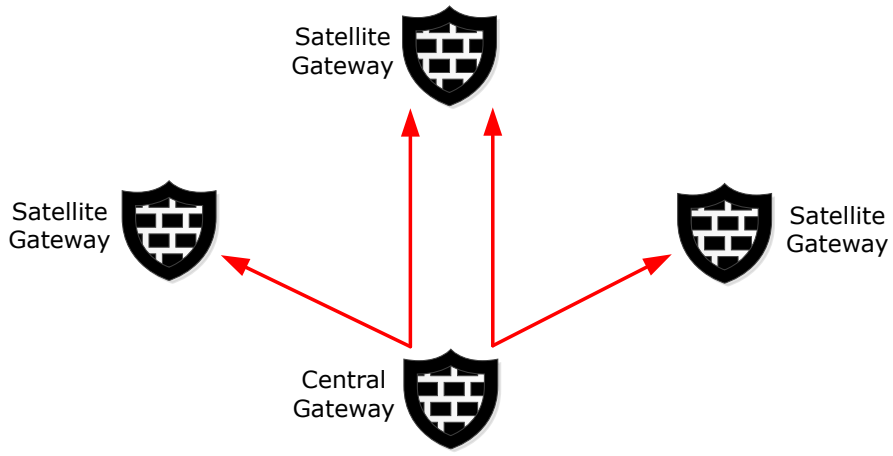
Star topology connects sites behind satellite gateways to the sites behind central gateways. No VPN tunnels are established between the satellite gateways.

In VPNs with partner organizations or remote offices, VPN connectivity is often needed between remote sites and a main site, but not from one remote site to another. This topology is a star topology.

The star topology is defined with satellite gateways that connect only to the central gateway. There is no VPN between the satellite gateways. This topology reduces the number of VPN tunnels that the gateways maintain compared to full-mesh topology. Having fewer tunnels can save resources on the remote gateways.

Sometimes the star topology is preferred even if VPN connectivity is needed between the remote offices. In this case, the central gateway can be used as a hub that relays traffic from one VPN tunnel to another. Traffic can be forwarded from either a site-to-site tunnel or a mobile VPN tunnel.

Hub VPN topology

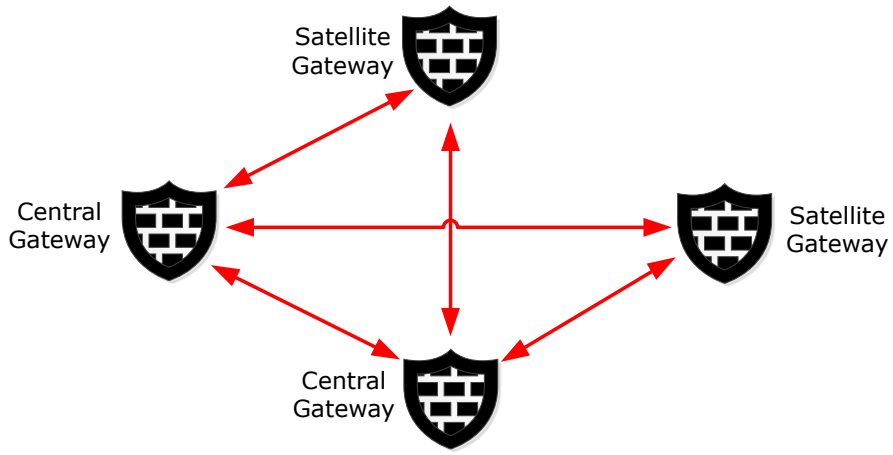


VPN hub topology routes site-to-site or mobile VPN connections to other sites through a central (hub) gateway using other site-to-site VPNs. The hub is usually a central gateway, but it can also be a satellite gateway.

The hub topology simplifies VPN client use if the clients connect to several gateways. It can also make setting up site-to-site VPNs easier, especially if the satellite gateways are third-party devices. VPN encryption and decryption require heavy computing. Consider hardware performance before high volumes of traffic are concentrated at a hub gateway.

Because the connectivity requirements vary from location to location, the VPN configuration can be a mix of the different topologies. This illustration shows an example of a mixed topology:

Combination of different topologies



Replacing two of the central gateways from the full mesh example with satellite gateways results in a VPN where all but two gateways still have a VPN with each other.

Considerations for creating new policy-based VPNs

There are some things to consider when you create a new policy-based VPN.

- Check whether you can use an existing policy-based VPN instead. Most settings can be set individually for each site-to-site tunnel even within a single policy-based VPN. The VPN Profile, pre-shared key, and Multi-Link settings can all be selected separately for each VPN tunnel. Site definitions are the only major exception to this rule.
- There must not be duplicate tunnels (two tunnels between the same two endpoints) in the configuration of any Engine. Duplicate tunnels cause a policy installation failure. The easiest way to avoid duplicate tunnels is to define all VPNs between your Engines in the same Policy-based VPN.
- If you are creating VPNs with partner organizations, you might only want to include a subset of the internal IP address space in the Site definitions. Limiting the IP address space allows you to avoid revealing all internal addresses to your partner. Any cases where Site definitions must be different for different VPN tunnels requires creating separate policy-based VPNs.
- IPsec tunnel between two Virtual Engines running on same Master Engine cluster is not supported. This limitation also applies to Master Engine clusters with only one node. However, to allow communication between two Virtual Engines, the inter-engine traffic must either be routed through an external router or by using a Shared Interface.

Create a Policy-Based VPN element

Create a Policy-Based VPN element and define the properties for the element.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.
2) Browse to **Policy-Based VPNs**.
3) Right-click **Policy-Based VPNs**, then select **New Policy-Based VPN**.
4) Configure the settings, then click **OK**.
   The **Policy-Based VPN** opens for editing.

Next steps

Define the VPN topology.

Edit a Policy-based VPN

The Policy-based VPN element can be configured in two ways: the basic properties are defined in the Policy-Based VPN element's properties. All other settings, including the included gateways, Sites, and tunnels are configured in the Policy-Based VPN editing view.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.
2) Browse to **Policy-Based VPN**.
3) Open the correct view for the settings that you want to edit:
   - To edit the basic properties, right-click the **Policy-Based VPN**element, then select **Properties**.
   - To adjust the other settings, right-click the **Policy-Based VPN** element, then select **Edit <element name>**.

Define VPN topology for policy-based VPNs

For a valid policy-based VPN, you must have at least two gateways in the VPN. At least one of the gateways must be listed as a central gateway. The satellite gateways list can be left empty (for a full-mesh topology).

The **Policy-Based VPN** editing view has three tabs. The gateway selection on the **Site-to-Site VPN** tab determines the following:

- Which gateways are included in the VPN.

- Which gateways form tunnels with each other.
- Which gateways contact each other through a hub gateway instead of contacting each other directly.

You define general VPN topology by classifying gateways as Central Gateways or Satellite Gateways. This classification defines which tunnels are generated on the **Tunnels** tab, and which gateways can be selected for mobile VPN access on the **Mobile VPN** tab.

IPv4 Access rules control which connections use the VPN tunnels. Always check the Access rules after you add or remove tunnels.

Note

Each endpoint-to-endpoint tunnel can only exist in one active VPN. If you use the same two gateway elements in more than one VPN, make sure that the topology does not create duplicate tunnels. You can also disable any duplicates of existing tunnels on the **Tunnels** tab.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.

2) Browse to **Policy-Based VPNs**.

3) Right-click the **Policy-Based VPN** element, then select **Edit**.

4) On the **Site-to-Site VPN** tab, drag and drop the Gateways you want to include in this VPN into either of the two panes for the VPN topology.
   - If you add a gateway under **Central Gateways**, the gateway can establish a VPN with any other gateway in the VPN. The **Tunnels** tab is populated with tunnels between the endpoints of the gateway you add and the endpoints of all other gateways in the VPN.
   - If you add a gateway under **Satellite Gateways**, the gateway can establish a VPN only with central gateways in this VPN. The **Tunnels** tab is populated with tunnels between the endpoints of the gateway you add and the endpoints of the central gateways.
   - The **Issues** pane alerts you to any incompatible or missing settings that you must correct.

   Note

   Be careful to not unintentionally drop gateways on top of other gateways. Dropping gateways on top of other gateways creates a forwarding relationship on a hub gateway.

5) (Optional) If you want to forward connections from one VPN tunnel into another through a hub gateway, drag and drop a gateway on top of another gateway. The gateway is added under the other gateway at the same level as the Sites.

   The Gateway used as a hub requires a special Site configuration.

6) (Optional) If you want to exclude a gateway's Site (some IP addresses) from this VPN, right-click the Site element under the gateway, then select **Disable.**

7) (Optional) Define which VPN Gateways provide Mobile VPN access.
   a) On the **Mobile VPN** tab, select one of the following options:
      - **Only central Gateways from overall topology** — Only the VPN Gateways in the **Central Gateways** listed on the **Site-to-Site VPN** tab provide mobile VPN access.
      - **All Gateways from overall topology** — All VPN Gateways included in the VPN provide mobile VPN access.
      - **Selected Gateways below** — Only the VPN Gateways that you add to the **Mobile VPN Gateways** tree provide mobile VPN access. Drag and drop the VPN Gateways from the **Resources** pane.

8) Click **Save**.

Define VPN tunnel settings for policy-based VPNs

The **Tunnels** tab in the Policy-Based VPN editing view allows you to define settings particular to individual tunnels or disable some tunnels altogether.

The topology of the policy-based VPN (defined on the **Site-to-Site VPN** tab) determines which tunnels are shown on the **Tunnels** tab. If you have set up connection forwarding between the gateways on the **Site-to-Site VPN** tab, the number of generated tunnels is reduced according to the relationships and the capabilities of the gateway that forwards the traffic. The forwarding relationships are shown under Define VPN tunnel settings for policy-based VPNs on page 1199**Forwarding Gateways**.

There are two types of tunnels:

- The **Gateway<->Gateway** list shows connections between pairs of gateways.
- The **Endpoint<->Endpoint** list shows the individual connections that form the tunnels in the **Gateway<->Gateway** list. There can be several connections at this level for any Gateway pair if one or both of the Gateways have multiple endpoints (Multi-Link). If both Gateways have only one endpoint, there is only one tunnel also at this level for the Gateway pair.

If a VPN Gateway has a Multi-Link VPN configuration, you can select whether to use tunnels as backups or actively balance traffic between them. Multi-Link is specific to Forcepoint Network Security Platform, and is not part of the IPsec standard. You might not be able to use Multi-Link with third-party gateways. Satisfactory results can be achieved if the third-party gateway allows ICMP probes, RTT ICMP probes, and supports DPD. You can disable redundant tunnels to the third-party gateway on the **Tunnels** tab if necessary.

This tab is also where you can view the *link summary*. The link summary is a summary of addresses and settings that have been configured for individual tunnels. You might want to check the link summary when there are complex setups involving external components (such as a VPN hub configuration).

Before editing a policy-based VPN that is used in active VPNs, we recommend making a backup of the Management Server.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.

2) Browse to **Policy-Based VPNs**.

3) Right-click the **Policy-Based VPN** element, then select **Edit**.

4) Click the **Tunnels** tab.

5) (Optional) If there are tunnels listed that are not needed, right-click the tunnel, then select **Disable**.

Duplicate tunnels are not allowed between VPNs. If another VPN already defines a tunnel between the same endpoints, disable the duplicate tunnel in one of the VPNs.

6) If you use pre-shared keys for authentication with external gateways, either set a key with your partner or export the keys that have been generated for your partner.

To view, change, or export the pre-shared key for a particular tunnel, double-click in the **Key** column in the **Gateway<->Gateway** list. This pre-shared key is used only with gateway devices. Set pre-shared keys for third-party VPN clients in the User elements. The Forcepoint VPN Client does not allow pre-shared key authentication.

CAUTION

The pre-shared key must be long and random to provide a secure VPN. Change the pre-shared key periodically (for example, monthly). Make sure that it is not possible for outsiders to obtain the key while you transfer it to other devices.

7) (Optional) Change the **VPN Profile** used at the tunnel level to override the profile selected for the VPN element:

- If you change a profile for a tunnel on the **Gateway<->Gateway** list, both IKE SA and IPsec SA settings are overridden from the default for the VPN.
- If you change a profile for a tunnel on the **Endpoint<->Endpoint** list, only the IPsec SA settings are overridden from the selection for the main tunnel on the gateway level.

8) In the **PPK** column, select a PPK element for the tunnel from the available option.

Important

You must select the VPN profile that is configured with the PPK settings in the **VPN Profile** column to be able to select a PPK element in the **PPK** column.

9) (Optional, Multi-Link only) Select the **Mode** in which **Endpoint<->Endpoint** links are used.

a) Select a tunnel on the **Gateway<->Gateway** list.

b) Right-click the **Mode** column for a link on the **Endpoint<->Endpoint** list, then select the mode from the right-click menu.

The Mode that you select for a link overrides the Mode setting in the endpoint properties. You can also configure the link's Mode to be automatically calculated based on the Mode defined for the endpoints. You can also define QoS Exceptions to select the Mode based on the QoS class of the traffic that is directed to the link.

10) (Optional) To review the IP addresses and settings used in the individual tunnels, right-click the tunnels on the **Endpoint<->Endpoint** list, then select **View Link Summary**.

11) After making all changes, check the **Validity** column for all tunnels.

a) If a tunnel has a warning icon in the **Validity** column, right-click the tunnel, then select **View Issues**.

b) Resolve all problems indicated in the messages shown.

If all tunnels are shown as valid, the policy-based VPN is correctly configured, although the Management Server cannot check all possible problems in this view. More issues might be shown at policy installation. Any validation and issues that are shown for external gateways are based only on the definitions that have been entered manually into the related elements.

12) Click **Save**.

Next steps

Add Access rules and possibly also NAT rules to direct outgoing traffic to the VPN and allow incoming traffic from the VPN.

Related concepts

Defining VPN gateways on page 1173
Defining Policy-Based VPN elements on page 1193
VPN certificates and how they work on page 1251

Related tasks

Back up system configurations on page 1297

Change VPN link modes in Multi-Link VPNs

The mode of a VPN link determines how the link is used in a Multi-Link VPN for VPN.

You can select the Mode in which **End-Point<->End-Point** tunnels are used if there are multiple links between two Gateways (Multi-Link configuration). The Mode you select is the default mode for the link.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.

2) Browse to **Policy-Based VPNs**.

3) Right-click the **Policy-Based VPN** element, then select **Edit <element name>**.

4) On the **Tunnels** tab, select a tunnel on the **Gateway<->Gateway** list.
The links between the gateways are displayed in the **End-Point<->End-Point** list.

5) Right-click the **Mode** column for a link on the **End-Point<->End-Point** list, then select **Edit Mode**.

6) Configure the settings.

7) Click **OK**.

Related tasks
Create QoS Class elements on page 981

Access rules for policy-based VPNs

The Access rules define which traffic is sent to the policy-based VPN and which traffic is allowed out of the policy-based VPN.

No traffic is sent out through the policy-based VPN until you direct traffic to the VPN in the Access rules. The Policy-Based VPN element must be referenced in at least one Access rule. The IKE and IPsec packets required to establish the VPN are allowed automatically based on the VPN definitions for the VPN Gateways. If there are intermediate engines between the VPN endpoints, make sure that the policies of those engines allow the required IKE and IPsec traffic.

You can set the VPN options in the Action options of the following Actions: Allow, Continue, or Jump. The **VPN Action** setting has three options, which have different effects depending on the source and destination of the traffic.

- **Apply VPN** — Directs traffic from protected local networks into the policy-based VPN tunnel. It allows traffic that arrives through a policy-based VPN to proceed. The rule does not match non-VPN traffic from outside networks into the protected networks regardless of whether the other cells in the rule match. This action allows handling special cases in which VPN and clear text traffic that match the same rule must be passed through the engine.
- **Enforce VPN** — Directs traffic from protected local networks into the policy-based VPN tunnel. It allows traffic that arrives through a policy-based VPN to proceed. The rule drops non-VPN connections from outside networks into the protected networks if the other cells in the rule match the connection.
- **Forward** — Directs traffic from protected local networks or from a policy-based VPN tunnel into another policy-based VPN tunnel. This action is useful for forwarding connections from one policy-based VPN tunnel into another (VPN hub configuration), or from local networks to VPN client computers that are currently connected.

When traffic is sent out through a policy-based VPN, the correct tunnel is selected based on the Sites of the gateway elements. If a VPN Access rule matches a connection with a source or destination IP address that is not included in the Sites, tunnel selection fails and the connection is dropped.

Incoming connections that arrive through the policy-based VPN are matched just like connections that do not use a VPN. Incoming connections do not have to match a VPN Access rule to be allowed in through the policy-based VPN. Any Access rule can match a policy-based VPN connection.

Create Access rules for policy-based VPN traffic

No traffic is sent out through the policy-based VPN until you direct traffic to the VPN in the Access rules. The Policy-Based VPN element must be referenced in at least one Access rule.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Create rules for incoming site-to-site VPN traffic.

   a) To allow traffic from a single policy-based VPN with an **Apply** or **Enforce** action, insert the following type of rule:

   Basic rule for allowing incoming VPN traffic from a single policy-based VPN

Source	Destination	Service	Action
Remote networks.	Local networks.	Set as needed.	Select **Allow**, then open the Action options. Set **VPN Action** to **Apply VPN** or **Enforce VPN**, then select a **Policy-Based VPN**.

   b) (Optional) To match the rule based on whether traffic is using a policy-based VPN, insert the following type of rule:

   Rule for allowing incoming policy-based VPN traffic from any number of different policy-based VPNs

Source	Destination	Service	Action	Source VPN
Remote networks.	Local networks.	Set as needed.	Allow.	To ignore this rule for non-VPN traffic, select **Match traffic based on source VPN**. Add one or more **Policy-Based VPN** elements according to where the traffic is coming from. This rule does not match traffic from other sources.

2) To create rules for outgoing policy-based VPN traffic, insert the following type of rule:

   Basic rule for outgoing VPN traffic

Source	Destination	Service	Action
Local networks.	Remote networks.	Set as needed.	Select **Allow**, then open the Action options. Set **VPN Action** to **Apply VPN**, **Enforce VPN**, or **Forward**, then select a **Policy-Based VPN**.

Note

If Access rules send traffic into a policy-based VPN, but the source or destination IP addresses are not included in the Site definitions, the traffic is dropped. This configuration error is shown as the message “tunnel selection failed” in the logs.

Create Access rules for VPN client connections

The Security Engine automatically allows policy-based VPN traffic to form and maintain the tunnels. VPN client user authentication is also allowed as part of this VPN connection establishment process.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) To allow incoming connections from VPN clients, insert the following type of rule:

Rule for allowing incoming traffic from VPN clients

Source	Destination	Service	Action	Authentication
VPN clients' Virtual Adapter address space. If Virtual Adapters are not used, select ANY.	Local networks.	Set as needed.	Select **Allow**, then open the Action options. Set **VPN Action** to **Apply VPN** or **Enforce VPN**, then select a **Policy-Based VPN**.	Add User or User Group elements and allowed Authentication Methods.

- When a policy-based VPN and Authentication Methods are specified in the installed policy, the corresponding configurations are activated on the engine. Connections from VPN client users are also matched against all other rules.
- Any users who can authenticate using the specified authentication method can connect with a VPN client. Any such connected users can access resources if there is a matching rule that allows connections without specific Users defined. You can also use the **Source VPN** cell to prevent unwanted matches in Access rules.
- When filled in, the **User** and **Authentication** cells are equal to Source, Destination, and Service as rule matching criteria. Matching continues from the next rule if the defined User and Authentication Method do not match the connection that is being examined. You can, for example, create rules that give the same user access to different resources depending on the authentication method used.

2) (Optional) To allow internal hosts to open connections to the VPN client computers when the VPN is active, insert the following type of rule:

Rule for sending traffic to VPN clients

Source	Destination	Service	Action
Local networks.	VPN clients' Virtual Adapter address space.	Set as needed.	Select **Allow**, then open the Action options. Set **VPN Action** to **Forward**. Select a specific **Policy-Based VPN** element, or select **Any Mobile VPN** to match any VPN client connection.

To use the policy-based VPN, the connecting hosts' IP addresses must be included in the gateway's Site definition.

Related concepts

Defining VPN gateways on page 1173
Defining Policy-Based VPN elements on page 1193
Defining Site elements for VPN gateways on page 1180

Create Access rules to forward traffic on hub gateways

Create access rules for forwarding policy-based VPN traffic from one tunnel to another and for forwarding tunneled traffic to the Internet.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) To forward policy-based VPN traffic from one tunnel to another, insert the following type of rule:

Basic rule for forwarding policy-based VPN traffic

Source	Destination	Service	Action	Source VPN
Addresses in remote (spoke) networks as needed.	Addresses in remote (spoke) networks as needed.	Set as needed.	Select **Allow**, then open the Action options. Set **VPN Action** to **Forward**, then select a **Policy-Based VPN**.	Select **Match traffic based on source VPN**, then add one or more **Policy-Based VPN** elements according to where the traffic is coming from.

2) To forward tunneled traffic to the Internet, insert the following type of rule:

Rule for allowing traffic except if it arrives through Policy-based VPNs

Source	Destination	Service	Action	Source VPN
Set as needed.	Set as needed.	Set as needed.	Allow.	Select **Match traffic based on source VPN**, then add one or more **Policy-Based VPN** elements according to where the traffic is coming from. This rule does not match traffic from other sources.

In most cases, the source IP addresses are from a private address space. You must add a NAT rule to translate them to publicly routable IP addresses. Make sure that NAT is enabled in the properties of the Policy-Based VPN element. Add a NAT rule for the VPN traffic if a suitable NAT rule does not exist already.

3) Create a VPN rule that directs all traffic to the policy-based VPN with the hub gateway.



Note

For the traffic to be allowed into the VPN, the destination IP address must be part of the Site definition of the hub gateway. When you forward Internet traffic, the hub's Site must usually include the Any Network element. This Site can interfere with other VPN configurations. We recommend disabling it in other VPNs.

Related concepts

Topologies for policy-based VPNs on page 1193

Prevent other Access rules from matching policy-based VPN traffic

Access rules that do not have any Source VPN definition can match any traffic, including traffic that is received through a VPN. You can optionally use the Source VPN cell to match traffic based on whether the traffic is coming from a policy-based VPN tunnel.

When the Source VPN cell is set to match policy-based VPNs, the rule only matches traffic from the selected policy-based VPNs.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Insert the following type of rule:

Rule for allowing traffic except if it arrives through VPNs

Source	Destination	Service	Action	Source VPN
Set as needed.	Set as needed.	Set as needed.	Set as needed.	Select **Match traffic based on source VPN**, then select **Rule does not match traffic from any VPN**.

Using NAT for policy-based VPN traffic

NAT configurations only apply to the encrypted packets in the VPN tunnel by default. To translate the addresses of the packets going through the policy-based VPN tunnel, you must specifically enable NAT for the policy-based VPN.

Note

NAT is needed for the NAT Pool feature in VPN client communications and for the Server Pool feature in inbound traffic management. To use these features in a policy-based VPN, NAT must be enabled in the properties of the Policy-Based VPN element.

Observe the following guidelines:

- Define Sites (encryption domains) that contain the translated IP addresses that the packets use when they are inside the policy-based VPN tunnel. Set the Sites that contain the real IP addresses to Private mode in the policy-based VPN.
  For example, if you translate IP addresses of traffic going into the policy-based VPN, add a Site that includes the translated IP addresses to your VPN Gateway element. The Sites that contain the internal addresses are set to Private mode.
- If address translation for VPN clients is enabled for the engine in the Engine Editor, NAT Pool translation is applied before the NAT rules. NAT rules cannot match traffic to which NAT pool translation is applied. NAT Pool is the preferred method for translating VPN client addresses.
- If you want to forward traffic originating from VPN clients to the Internet, you must typically have at least two NAT rules. The first rule is for connections to internal resources to prevent NAT from being applied or to translate to an internal IP address as necessary. The second rule translates internal IP addresses to an external IP address for the Internet connections.

The order of processing for traffic going into a policy-based VPN tunnel is:

1) Access Rules
2) NAT Rules
3) VPN tunnel

The order of processing for traffic coming out of a VPN tunnel is:

1) Access Rules
2) VPN client NAT Pool
3) NAT Rules
4) Internal Network

Other than these guidelines, there are no other VPN-specific issues with NAT rules. The first matching NAT rule is applied to those connections that are matched against the NAT rules and the rest of the NAT rules are ignored.

Related concepts
Getting started with NAT rules on page 851

NAT traversal in VPNs

NAT traversal (NAT-T) is an optional IKE standard mechanism to detect when an IPsec VPN tunnel goes through a NAT device. NAT-T allows IPsec VPNs to work reliably through networks where NAT is applied to connections.

If NAT-T is enabled and NAT is detected, the gateway automatically uses UDP port 4500 for IKE negotiation messages, and encapsulates ESP packets in UDP packets that use port 4500.

NAT-T is always enabled for mobile VPNs.

NAT-T encapsulation is not always necessary even if static NAT is applied to a site-to-site VPN. You can define Contact Addresses so that the VPN works even when NAT is applied. The NAT-T option is activated in the endpoint properties in the Engine Editor or in the External VPN Gateway element.

NAT for VPN gateway communications

You can apply NAT to the communications between VPN Gateways.

The communications that establish and maintain VPN tunnels between VPN Gateways are always translated according to the NAT rules. Create a matching NAT rule to translate addresses in these communications and make sure that the communications do not match the wrong NAT rule unintentionally.

When you add NAT, you might need to change VPN settings. Add contact addresses for the engines. Contact Addresses can be used with both internal and external gateways.

There is nothing VPN-specific about creating the actual NAT rules.

NAT for traffic in VPN tunnels

You can configure NAT for traffic in VPN tunnels in the properties of the Policy-Based VPN element.

By default, IP addresses in traffic that enters or leaves a VPN tunnel are not translated. An option in the properties of the Policy-Based VPN element, accessible through the right-click menu for the policy-based VPN, defines whether NAT is applied to traffic in VPN tunnels.

If the option to translate the IP addresses is enabled, the IP addresses in traffic that uses site-to-site VPN tunnels are translated according to the NAT rules. There is nothing VPN-specific in creating these NAT rules. However, the VPN configuration is affected if local protected addresses are translated using NAT:

- Set the Site element that contains the private local addresses (before translation) in the Private mode in VPNs in which those addresses are translated using NAT.
- Add the translated addresses as a new Site for the gateway (disable the Site in other VPNs). This Site is in the default Normal mode.

VPN client traffic is translated according to the NAT Pool settings defined for the Engine in the Engine Editor, or as defined in the NAT rules.

Monitoring policy-based VPNs

You can monitor the status of VPNs in the Dashboard view. The overall status of the VPNs and the tunnels they contain is shown in the tree of monitored elements.

Logging for policy-based VPNs is separate for the tunnels and the traffic that uses the tunnels:

- VPNs negotiations are always logged (regardless of the logging options in Access rules) as informational messages.
- More detailed logging is available when you activate IPsec diagnostic logging for the Engine/VPN role for troubleshooting purposes.
- The traffic using the VPNs tunnels is logged according to the logging options in the rule that allows the traffic in or out of the VPNs.
- The Dashboard view provides shortcuts to logs filtered for the specific policy-based VPNs or VPNs Gateway element referenced in the log event.
  - Right-click a **policy-based VPN** in the Status tree, then select **Monitoring** > **Logs by VPN Gateway**.
  - Right-click a **VPN Gateway** in the Status tree or connectivity diagram, then select **Monitoring** > **Logs by VPN Gateway**.
  - Right-click the connection between two **VPN Gateways** in the connectivity diagram, then select **Monitoring** > **Logs by VPN Gateways** to view logs of traffic between the two VPN Gateways.

Log pruning filters can delete some (or even all) of the generated messages.

Related concepts

How component statuses are indicated on page 224
Information about IPsec tunnels in logs on page 1422

Related tasks

Enable or disable diagnostics on page 356
Define logging options for Access rules on page 901

Define Mobile VPNs

In a mobile VPN, a VPN client on a user's device connects to a VPN gateway.

You can use both SSL VPN and IPsec tunnels together in the mobile VPN configuration in the same policy-based VPN.

Note

Route-based Tunnels do not support mobile VPNs.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Configure VPN Client settings in the Engine Editor.
   a) Select **Engine Configuration**.
   b) Right-click a Engine element, then select **Edit <element type>**.
   c) Browse to **VPN** > **VPN Client**.
   d) Configure the settings.
   e) (Optional) Configure the settings on the **Advanced** branch.
   f) Click **Save and Refresh** to save the changes to the configuration and refresh the policy on the engine.
2) Create a **policy-based VPN** or edit an existing **policy-based VPN**.
3) On the **Mobile VPN** tab of the policy-based VPN, specify which VPN Gateways provide mobile VPN access.
4) Click **Save**.

Configuring route-based VPNs

In route-based VPNs, the routing defines which traffic is sent through the VPN tunnel.

Route-based Tunnels elements represent the endpoints of the tunnel. Tunnel Interfaces allow routing information to be used to determine the correct VPN tunnel to use.

The routing configuration also determines the physical network interfaces on the engine to which the tunnel interfaces are automatically mapped. You can statically define which networks are reachable through each tunnel interface. You can also use dynamic routing to create the routes for traffic to be sent through the VPN tunnels.

When Route-based Tunnels are in transport mode, the packets are not encapsulated into new IPsec packets. Instead, the original headers of the packet are left intact, and the IP payload of the packet is encrypted. IPsec transport mode is used to encrypt the packets. Other encapsulation, such as *generic routing encapsulation* (GRE) or *IP in IP* (IP-IP), must be used to add the tunnel endpoint IP addresses in front of the original packet header.

When Route-based Tunnels are in tunnel mode, the encryption is provided by a policy-based VPN.

Configuring route-based VPNs consists of the following guidelines:

1) Create a tunnel interface for one end of the VPN.

2) Create a tunnel interface for the other end of the VPN.

3) Create a Route-based Tunnels element that references both ends of the VPN.
In the Route-based Tunnels element, you can select the tunnel type and a VPN Profile to use.

4) (Optional) Create Tunnel Groups to group Route-based Tunnels elements.
Each tunnel can be added to a Tunnel Group element. The groups allow you to organize the tunnels, and you can view the groups in the VPN section in the Dashboard view.

5) Create Access rules to allow traffic between the internal network and the networks that are reachable through the route-based VPNs.

Using route-based VPNs for dynamic routing

Route-based VPNs can protect and route dynamic routing communications between sites to protect the confidentiality and integrity of the dynamic routing communications.

Routing protocols, such as RIP, OSPF, and BGP, send non-routable multicast packets between routing devices, such as routers and engines. Because IPsec accepts only unicast traffic, these packets cannot be directly sent into IPsec tunnels. Instead, dynamic routing communications are forwarded to tunnel interfaces, which encapsulate the traffic and send it into the Route-based Tunnels tunnel. The following configuration considerations apply when route-based VPNs are used to protect dynamic routing protocols:

- The TTL value for the tunnel must be high enough to allow the packets to be routed through each hop in the route.
- If IP addresses are defined for tunnel interfaces, the netmask must be defined according to the functionality of the interface. For OSPF, the peers must belong to the same subnet.

Add tunnel interfaces for Security Engines

Tunnel Interfaces allow routing information to be used to determine the correct VPN tunnel to use in route-based VPNs.

Any traffic that is routed to a tunnel interface and allowed by Access rules is automatically sent through the tunnel to the peer endpoint defined in the Route-based Tunnels element. Tunnel interfaces are only used in route-based VPNs.

You can optionally add IPv4 or IPv6 addresses to a tunnel interface. Tunnel interfaces can only have static IP addresses. Any IP address can be added to a tunnel interface, even if the same IP address is used on another interface or as a loopback IP address. Adding an IP address to a tunnel interface allows you to define the source IP address of traffic sent from the Security Engine itself. For example, an IP address is recommended to provide a source IP address for dynamic routing daemons, for IGMP proxy, and for Protocol Independent Multicast - Sparse-Mode (PIM-SM) configuration. If no IP address is added to the tunnel interface, the source IP address for traffic sent from the Security Engine is automatically selected. The selection is done according to the **Bypass Default IP Address** setting in the loopback interface configuration for the Security Engine.

The mapping of tunnel interfaces to physical network interfaces on the Security Engine is done automatically based on the routing configuration.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click an Security Engine, then select **Edit <element type>**.
3) Browse to **Interfaces**.
4) Select **Add** > **Tunnel Interface**.
5) Configure the settings.
6) Click **OK**.
7) If you want to add a source IP address for traffic sent from the engine node, add IPv4 addresses or IPv6 addresses to the tunnels.
8) If you do not want to add IP addresses, select system communication roles for engine interfaces to define how the source IP address for traffic sent from the engine node is selected.
9) Click **Save and Refresh**.

Create Route-based Tunnels elements

Route-based Tunnels elements represent the endpoints of the tunnel.

Before you begin

Add tunnel interfaces for Security Engines.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.
2) Browse to **Route-Based Tunnels**.
3) Right-click **Route-Based Tunnels**, then select **New Route-Based Tunnel**.

4) Configure the settings.

Note

Specifying tunnel options for individual Route-based Tunnels overrides the default settings defined for the tunnel interface on the Security Engine.

5) Click **OK**.

6) Click **Save**.

Note

IPsec tunnel between two Virtual Engines running on same Master Engine cluster is not supported. To allow communication between two virtual engines, the inter-engine traffic must either be routed through an external router or by using a Shared Interface.

Using GRE keep alive to check the status of Route-based Tunnels

You can optionally use GRE keep alive to check that Route-based Tunnels of the GRE tunnel type are still functioning.

When GRE keep alive is enabled, the Security Engine sends keep alive packets at the specified interval. If no reply is received after the specified number of packets, the GRE tunnel is considered to be down.

You can enable and configure GRE keepalive in the properties of tunnel interfaces on Security Engines and in the properties of Route-Based Tunnel elements. When you enable GRE keepalive for a tunnel interface on an Security Engine, GRE keepalive is used in all GRE Route-based Tunnels where the tunnel interface is an endpoint. Enabling GRE keepalive for individual Route-based Tunnels overrides the default settings defined for the tunnel interface on the Security Engine.

To use GRE keepalive, your environment must meet these requirements:

- The router to which the Security Engine is connected must support GRE keepalive.
- **No Encryption** must be selected for the **Encryption** option in the properties of the tunnel interface or Route-based Tunnels element.

Enable GRE keepalive for a tunnel interface on an Security Engine

When you enable GRE keepalive for a tunnel interface on an Security Engine, GRE keepalive is used in all GRE Route-based Tunnels where the tunnel interface is an endpoint.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click an Security Engine, then select **Edit <element type>**.

3) Browse to **Interfaces**.

4) Right-click a tunnel interface, then select **Properties.**

5) On the **Tunnel** tab, enable GRE keepalive.
   a) Make sure that **No Encryption** is selected from the **Encryption** drop-down list.
   b) In the **Tunnel options** section, select **Use GRE Keepalive**.
   c) Configure the settings.

6) Click **OK**.

7) Click **Save and Refresh**.

Enable GRE keepalive for a Route-Based Tunnel element

You can enable GRE keepalive for individual route-based tunnels to override the default settings defined for tunnel interfaces on the Security Engine.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.

2) Browse to **Route-Based Tunnels**.

3) Right-click a **Route-Based Tunnel**element, then select **Properties**.

4) Enable GRE keepalive.
   a) Make sure that **No Encryption** is selected from the **Encryption** drop-down list.
   b) In the **Tunnel options** section, select **Use GRE Keepalive**.
   c) Configure the settings.

5) Click **OK**.

Create Tunnel Group elements

You can group Route-based Tunnels elements, then monitor them in the Dashboard view.

Note

To add multiple existing Route-based Tunnels elements to a Tunnel Group, select the elements, right-click, then select **More actions** > **Select Tunnel Group**.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.
2) Browse to **Other Elements** > **Tunnel Groups**.
3) Right-click **Tunnel Groups**, then select **New Tunnel Group**.
4) Configure the settings, then click **OK**.

Use a policy-based VPN to encrypt tunnels in route-based VPNs

You can use a policy-based VPN to provide encryption for Route-based Tunnels.

Before you begin

Define the policy-based VPN that provides the encryption.

Using a policy-based VPN to encrypt tunnels in a Route-based Tunnels allows you to do the following:

- Encrypt multiple tunnels in the same VPN tunnel. This configuration improves compatibility with third-party devices and cloud-based services that do not support multiple, separately encrypted tunnels.
- Create multiple tunnels between remote and local sites when only one public IP address is available.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Create a Host element.
   a) Select **Network Elements**.
   b) Right-click **Hosts**, then select **New Host**.

c) In the **IPv4 Address** or **IPv6 Address field**, enter the same IP address as the endpoint you use in the Route-based Tunnels.



Note

You might receive a warning that the IP address of the Host element is not unique. Ignore the warning and save the element.

d) Configure the other settings according to your needs.

e) Click **OK**.

2) Configure the VPN settings for the engine that acts as the VPN gateway.

a) Select **Engine Configuration**.

b) Right-click the Security Engine, then select **Edit <element type>**.

c) Browse to **VPN** > **Endpoints**, then define at least two endpoints: one for the policy-based VPN and one for the Route-based Tunnels.

d) Browse to **Sites**, then add the Host element to the site for the VPN Gateway.

e) Click **Save**.

3) Configure the policy-based VPN that provides the encryption.

a) Select **Secure SD-WAN Configuration**.

b) Browse to **Policy-Based VPNs**.

c) Open the policy-based VPN for editing.

d) On the **Site-to-Site VPN** tab, add the VPN Gateway that represents the engine to the **Central Gateways** or **Satellite Gateways** list.

e) Click **Save**.

4) Create the Route-Based Tunnel element.

a) Select **Secure SD-WAN Configuration**.

b) Browse to **Route-Based Tunnels**.

c) Right-click **Route-Based Tunnels**, then select **New Route-Based Tunnel**.

d) Use the following settings:

Setting	Configuration
Tunnel type	**GRE**, **IP-IP**, or **SIT**.
Encryption	**Tunnel Mode**.
VPN	Select the policy-based VPN that provides the encryption.
Local engine	Select the same VPN Gateway that is used in the policy-based VPN.
CVI	Select the CVI that has the same IP address as the endpoint that is used in the policy-based VPN.

Configure the other settings according to your needs.

e) Click **OK**.

5) Add Access rules to allow traffic between the internal network and the networks that are reachable through the Route-based Tunnels.

Note

The Access rules that direct the Route-based Tunnels traffic into the policy-based VPN are automatically generated for the Engines associated with the VPN Gateway elements. The rules are not visible in the Engine policy, and cannot be edited. If a policy that contains the automatically generated rules is installed on a Engine that is not involved in the VPN, the rules are ignored.

a) Open the Engine policy for editing.

b) Add IPv4 Access rules or IPv6 Access rules that have the following settings:

Source	Destination	Service	Action
Elements that represent the internal network	Elements that represent the networks that are reachable through the Route-based Tunnels.	Select a service, or set to **ANY**.	Allow

Configure the other settings for the rules according to your needs.

c) Click **Save**.

d) Install the policy on all Engines that are involved in the VPNs.

Examples of policy-based VPNs

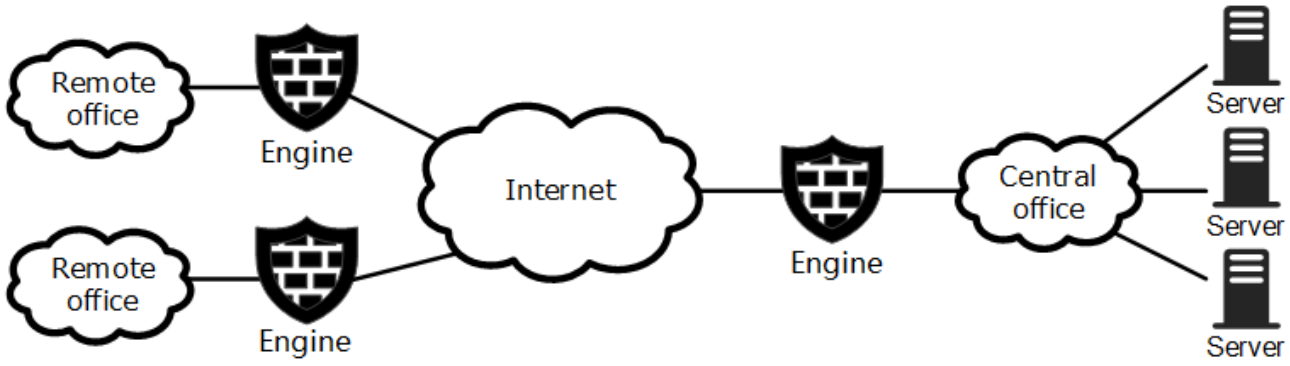
These examples illustrate some common uses for policy-based VPNs and general steps for how each example is configured.

Example: creating a policy-based VPN between three offices

This example shows how to create a VPN between two or more locations.

Company A has a central office and two remote offices, each with their own Engine/VPN role device. The company needs secured communications links between the remote offices and the central office for access to various services, such as file servers, at the central office.

Company A's networks



All shared servers are at the central office, and internal emails and other communications are also handled by the central servers. There is no need for secure connectivity between the remote offices.

All Engines have a public IP address toward the Internet. The internal networks at each site use private IP addresses. There is no need to translate the VPN traffic because all offices use their own distinct address space.

The security policy of the company requires certificate-based authentication. The administrators decide to use the Management Server's Internal RSA CA for Gateways for issuing the VPN certificates.

The administrators:

1) Select each engine's public IP address as the VPN endpoint, then activate automatic certificate management.

2) Add Site elements for all gateways, then add the entire local internal network as the content for each Site.

3) Create a VPN Profile, and select RSA Signatures as the authentication method. RSA certificates are automatically generated for each gateway.

4) Create a Policy-Based VPN element called "Inter-Office VPN" that includes the central office gateway as a central gateway and the two remote site gateways as satellite gateways.

5) Add the following types of Access rules in the policy of the central office engine:

Source	Destination	Action
Network elements for remote office 1 and remote office 2 internal IP addresses	Network elements for central office's internal networks	Select **Allow**, then open the Action options. Set **VPN Action** to **Enforce VPN**, then select the "Inter-office VPN" **Policy-Based VPN** element.
Network elements for central office's internal networks	Network elements for remote office 1 and remote office 2 internal IP addresses	Select **Allow**, then open the Action options. Set **VPN Action** to **Enforce VPN**, then select the "Inter-office VPN" **Policy-Based VPN** element.

6) Add the following types of Access rules in the policies of both remote office engines:

Source	Destination	Action
Network element for each remote office's internal IP addresses	Network elements for central office's internal networks	Select **Allow**, then open the Action options. Set **VPN Action** to **Enforce VPN**, then select the "Inter-office VPN" **Policy-Based VPN** element.
Network elements for central office's internal networks	Network element for each remote office's internal IP addresses	Select **Allow**, then open the Action options. Set **VPN Action** to **Enforce VPN**, then select the "Inter-office VPN" **Policy-Based VPN** element.

Example: creating a policy-based VPN for mobile users

An example of a policy-based VPN that allows mobile users to authenticate and connect to internal networks.

Company A has service technicians and salespeople who must be able to connect to their office networks to access information when they are on customer visits. The administrators need to add VPN client access to the existing VPN infrastructure. The administrators decide to use Forcepoint VPN Client. As the authentication method, the administrators decide to use passwords stored in the Management Server's internal database.

The administrators also want to provide only one point of access so that the users do not have to select which gateway to connect to. The central office has site-to-site VPN tunnels to both remote offices that can be used for forwarding traffic to those sites as needed. The existing DHCP server at the central office can be used for assigning IP addresses to the VPN clients' Virtual Adapter. A Virtual Adapter is required for this type of forwarding.

The administrators:

1) Edit the central office engine element, then activate the Virtual Adapter method for VPN client address management.

2) Edit the VPN Profile to use Hybrid Authentication for authenticating the VPN client users.

3) Create a **Policy-Based VPN** element called "Remote User VPN" that includes the central office gateway as a Central Gateway.

4) Select the **Only central Gateways from overall topology** option on the **Mobile VPN** tab.

5) Create a "Forward Addresses" Site element under the central office gateway.

6) Populate the site with the remote office networks to route those IP addresses through the "Remote User VPN".

7) Disable the “Forward Addresses” Site in the existing “Inter-Office VPN” between the central office and the remote offices. Sites are global for all policy-based VPNs, so this Site must be disabled to avoid a misconfiguration in the Inter-Office VPN.

8) Create User Group and User elements to define the user names and passwords for the VPN client users.

9) Add the following Access rules in the policy of the central office engine:

Source	Destination	Action	Authentication	Source VPN
ANY	Central office internal networks	Select **Allow**, then open the Action options. Set **VPN Action** to **Enforce VPN**, then select the “Remote User VPN” Policy-Based VPN element.	**Users** tab: “VPN Client Users” User Group **Authentication Methods** tab: “User Password” Authentication Service	
VPN Client DHCP addresses	Remote offices’ internal IP addresses	Select **Allow**, then open the Action options. Set **VPN Action** to **Forward**, then select the “Inter-office VPN” **Policy-Based VPN** element.		Rule matches traffic from any VPN client

10) Create a customized Forcepoint VPN Client installation package for Windows. A customized installation package allows users of Forcepoint VPN Client for Windows to install using a silent installation package that does not require their input. The administrators include the gateway contact information in the package so that the users do not need to enter it manually even when they use the Forcepoint VPN Client for the first time.

Examples of route-based VPNs

These examples illustrate some common uses of route-based VPNs and general steps for how each example is configured.

Example: protecting dynamic routing communications with a Route-based Tunnels

This scenario shows an example of protecting communications when public Internet networks are used for backup connectivity.

Company A is a large company with enterprise networks at multiple sites. The networks are currently connected with a private backbone network that is built with dynamic routing using OSPF. The administrators want to use public Internet networks for backup connectivity in case the private backbone fails. To route the traffic and to protect the confidentiality and integrity of the dynamic routing communications, the administrators decide to send dynamic routing communications through tunnels in a Route-based Tunnels.

The administrators:

1) Define tunnel interfaces on the engines that act as VPN Gateways at each site.

Note

One tunnel interface is required for each remote VPN Gateway endpoint.

2) Add IP addresses to each tunnel interface.

3) Create a **Route-Based Tunnel** element that specifies the gateways, endpoints, and tunnel interfaces, and select the appropriate tunnel type and VPN Profile. The following options are used:
   - **TTL**: Default.
   - **MTU**: Default.
   - **PMTU Discovery**: Enabled.

4) Create Access rules that allow traffic between the internal networks and the networks that are reachable through the Route-based Tunnels.

5) Refresh the policy on the engines that act as VPN Gateways.

6) Configure dynamic routing on the engines.

Example: configuring route-based VPNs with external gateways

This scenario shows an example of creating a Route-based Tunnels between an internal and external network.

The administrators at Company B want to create a Route-based Tunnels between their own network and a partner's network. The administrators:

1) Create a Network element to represent the partner's network.

2) Define a Tunnel Interface on the Company B engine that acts as the VPN Gateway.

3) Configure routing to define a route to the partner's network through the Tunnel Interface.

4) Define an External VPN Gateway element to represent the partner company's gateway device.

5) Add a Route-Based Tunnel element with the following settings:

Local Gateway	Remote Gateway
■ **Gateway** — VPN Gateway element that represents the engine ■ **Endpoint** — Endpoint IP address in the Internal Network ■ **Interface** — Tunnel Interface defined on the engine	■ **Gateway** — External VPN Gateway element ■ **Endpoint** — Endpoint IP address in the Partner Network

6) Select an IPsec Profile and an encapsulation Mode that is compatible with the External VPN Gateway.

7) Create an Access rule that allows traffic from the internal network to the partner network that is reachable through the Route-based Tunnels.

8) Refresh the policy on the engine that acts as a VPN Gateway.

Chapter 74

Example VPN configurations

Contents

- Getting started with example VPN configurations on page 1223
- Example VPN configuration 1: Basic VPN between Security Engines on page 1224
- Example VPN configuration 2: Basic VPN with a partner gateway on page 1228
- Example VPN configuration 3: Basic VPN for remote clients on page 1236
- Example VPN configuration 4: Basic VPN hub on page 1244

The following example configurations outline common VPN use cases.

Getting started with example VPN configurations

You can follow these examples when you set up your own VPNs and add other features after the basic scenario is configured and working.



Note

VPNs are not supported on Layer 2 Engines or on layer 2 physical interfaces on Engines.

The following example configurations are explained:

- Example 1 shows a site-to-site VPN between two or more Security Engines that are managed through the same Management Server. A default set of VPN settings is used to simplify the configuration.
- Example 2 shows a site-to-site VPN between an Security Engine and an IPsec-compatible VPN gateway that is not managed through the same Management Server. A customized set of VPN settings is created. Customized settings are typically mandatory for this configuration. A pre-shared key is used for authentication.
- Example 3 shows a mobile VPN between an Security Engine and the Forcepoint VPN Client installed on individual computers. A default set of VPN settings is used to simplify the configuration.
- Example 4 shows a site-to-site VPN in which several remote gateway connect to a hub gateway. The hub gateway forwards connections to the other remote gateways as necessary. A default set of VPN settings is used to simplify the configuration.

Example VPN configuration 1: Basic VPN between Security Engines

This scenario shows an example of how to create a policy-based VPN between two or more Security Engines managed through the same SMC.

This example VPN requires all engines to have a static IP address (not assigned using DHCP or PPPoE).

The address spaces protected by the different Security Engines that act as gateways must not overlap within any single VPN. If you use the same IP addresses at the different locations, you must apply NAT to the communications. You must also define the sites using the translated IP addresses that are used inside the VPN tunnels.

This scenario uses the default **Suite-B-GCM-128** VPN profile that contains the VPN settings specified for the Suite-B-GCM-128 cryptographic suite in RFC 6379. The profile uses pre-shared keys for authentication.

The configuration consists of the following general steps:

1) Configure VPN settings for the Security Engines.
2) Create a Policy-Based VPN element.
3) Create Access rules.

Begin by configuring VPN settings for the Security Engines.

Example VPN configuration 1: configure VPN settings for the Security Engines

Follow these steps for each Security Engine that is used as a VPN gateway.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click the Engine element, then select **Edit Single Engine** or **Edit Engine Cluster**.
3) Browse to **VPN** > **Endpoints**.
4) (Optional) Change the selection of IP addresses that you want to use as endpoints in VPNs.
   - Typically, these are IP addresses that belong to interfaces toward the Internet, which are automatically selected based on the engine's default routing table.
   - If loopback IP addresses are defined for the Security Engine, you can select a loopback IP address as the endpoint IP address. On clustered engines, the IP addresses are CVIs.
   - (Optional) If you have more than one Internet connection, select an IP address from each ISP.

5) In the navigation pane on the left, browse to **VPN** > **Sites**.

   The Sites represent the addresses that are routable through the VPN. Sites do not grant any host access directly. The Access rules define the allowed connections.

6) (Optional) Select the internal networks that you want to exclude from the VPN by disabling the interface they are under in the automatic site.

   Disabled interfaces are grayed-out.

   - If you want to include some individual network that is under an otherwise disabled interface, drag and drop it from under the disabled interface onto the Site element. The element is copied to the higher level. The copied definition is not updated automatically.
   - The Sites must include only internal networks. Do not add interfaces with the **Any Network** element in this type of VPN.

7) Click **Save**.

Next steps

Create a Policy-Based VPN element.

Related concepts
Defining VPN gateways on page 1173

Example VPN configuration 1: create a Policy-Based VPN element

In this configuration, you must create a Policy-Based VPN element.

Note

This configuration scenario does not explain all settings related to Policy-Based VPN elements.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.

2) Right-click **Policy-Based VPNs** in the element tree and select **New Policy-Based VPN**.

3) In the **Name** field, enter a unique name.

4) In the **Default VPN Profile** drop-down list, make sure that **Suite-B-GCM-128** is selected.

   Note

   The VPN Profile element defines most of the IPsec settings. You can optionally create a custom VPN Profile element.

5) If you want to apply NAT rules to the communications that go through the VPN, select **Apply NAT to traffic that uses this VPN**.

   This setting does not affect the communications that the two gateways have with each other to set up and maintain the VPN. Communications between the gateways are always matched to the automatic rules or the NAT rules.

6) Click **OK**.

   The **VPN Editing** view opens on the **Site-to-Site VPN** tab.

7) To define which gateways can create a VPN with each other, drag and drop two or more VPN Gateway elements from the **Resources** pane to the **Central Gateways** or **Satellite Gateways** lists.

   - If you add a VPN Gateway to the **Central Gateways**, the VPN Gateway can establish a VPN with any other VPN Gateway in this VPN (both Central and Satellite). Add at least one of the VPN Gateways under **Central Gateways**.
   - If you add a VPN Gateway to the **Satellite Gateways**, the VPN Gateway can establish a VPN only with VPN Gateways defined as Central in this VPN. You do not have to add any VPN Gateways to the **Satellite Gateways** (all gateways can be Central).

   Note

   Be careful that you do not accidentally drop VPN Gateway elements on top of other VPN Gateway elements. This configuration creates a hub topology where the top-level VPN Gateway forwards connections from other components to the lower-level VPN Gateway.

8) On the **Tunnels** tab, make sure that the **Validity** column in the **Gateway<->Gateway** and the **End-Point<->End-Point** tables has a green check mark to indicate that there are no problems.

   a) If the **Validity** column of a tunnel has a warning icon, see the **Issues** pane to check what the problem is. If the pane is not shown, select **Menu** > **View** > **Panels** > **Issues**.

   b) If issues are shown, correct them as indicated. Long issues are easiest to read by hovering the cursor over the issue text so that the text is shown as a tooltip.

9) Click **Save** to save the Policy-Based VPN.

Next steps

Create Access rules

Related concepts

Defining Policy-Based VPN elements on page 1193

Example VPN configuration 1: create Access rules

The rules in this example allow protected hosts to open connections both ways. Two rules are created here to allow the different directions of traffic separately.

VPN rules are matched based on source, destination, and service like any other rules.

Note

This configuration scenario does not explain all settings related to Access rules for VPNs.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Browse to **Engine** > **Policies** > **Engine Policies**.

3) Right-click the Engine policy that is used by the Security Engines involved in the VPN, then select **Edit Engine Policy**.

4) Add two IPv4 Access rules in a suitable location in the policy.
   - Make sure that these rules are above other rules that match the same traffic with **Allow**, **Discard**, or **Refuse** as their action.
   - Traffic that you do not want to send through the VPN must not match these rules. Traffic that is not routable through the VPN is dropped if it matches these rules.

5) Fill in the rules.
   If NAT is enabled in the VPN, remember that the Access rules are checked before the NAT rules are applied.

   Example VPN rules

Source	Destination	Service	Action
Local internal networks	Remote internal networks	Set as needed.	Select **Allow**, then open the Action options. Set **VPN Action** to **Enforce VPN**, then select a **Policy-Based VPN**.
Remote internal networks	Local internal networks	Set as needed.	Select **Allow**, then open the Action options. Set **VPN Action** to **Enforce VPN**, then select a **Policy-Based VPN**.

6) Save the policy.

7) Add the same rules in the policies of all engines involved in the VPN.

CAUTION

If you continue to use this VPN, change the pre-shared key periodically to guarantee continued confidentiality of your data. Alternatively, you can change to certificate-based authentication by creating a custom VPN profile.

8) Refresh the policies of all engines involved in the VPN to activate the new configuration.

Result

The VPN is established when traffic matches the Access rules. Example VPN configuration 1 is now complete.

Related concepts
Access rules for policy-based VPNs on page 1202
Monitoring policy-based VPNs on page 1208

Example VPN configuration 2: Basic VPN with a partner gateway

This scenario walks you through creating a site-to-site VPN between one Security Engine and one external VPN gateway that is not managed through the same SMC.

This example VPN requires the local engine to have a static IP address (not assigned using DHCP or PPPoE).

The address spaces protected by the different VPN Gateways must not overlap within any single VPN. If you use the same IP addresses at the different locations, you must apply NAT to the communications and define the Sites using the translated IP addresses. The translated addresses are the addresses that are used inside the VPN tunnels.

You can create VPN with IPsec-compliant gateway devices from many different manufacturers. You can create VPN with partner organizations that use a third-party VPN solution. The authentication and encryption options to use must be decided beforehand in co-operation with the administrator of the other gateway.

The configuration consists of the following general steps:

1) Configure VPN settings for the Security Engine.
2) Create an External VPN Gateway element.
3) Define a site for the external VPN gateway.
4) Create a VPN Profile element.
   The VPN Profile must contain VPN settings that match the settings defined on the external VPN gateway.
5) Create a Policy-Based VPN element.
6) Create Access rules.

Begin by configuring VPN settings for the Security Engine.

Example VPN configuration 2: configure VPN settings for the Security Engine

If you have already configured VPN settings for the Security Engine, there is no need to change any of the settings.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click the Engine element, then select **Edit Single Engine** or **Edit Engine Cluster**.

3) In the navigation pane on the left, browse to **VPN** > **Sites**.
   The Sites represent the addresses that are routable through the VPN. Sites do not grant any host access directly. The Access rules define the allowed connections.

4) (Optional) Select the internal networks that you want to exclude from the VPN by disabling the interface they are under in the automatic site.
   Disabled interfaces are grayed-out.
   - If you want to include some individual network that is under an otherwise disabled interface, drag and drop it from under the disabled interface onto the Site element. The element is copied to the higher level. The copied definition is not updated automatically.
   - The Sites must include only internal networks. Do not add interfaces with the **Any Network** element in this type of VPN.

5) To use NAT to translate the IP addresses of the hosts that make connections through this VPN, drag and drop the networks for the translated addresses on top of the (top-level) automatic Site element on the right.

6) Click **Save**.

Next steps

Create an External VPN Gateway element.

Related concepts
Defining VPN gateways on page 1173

Example VPN configuration 2: create an External VPN Gateway element

You need an External VPN Gateway element for this configuration.

Note

This configuration scenario does not explain all settings related to External VPN Gateway elements.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.

2) Right-click **VPN Gateways** in the element tree, then select **New External VPN Gateway**.

3) In the **Name** field, enter a unique name.

4) Click **Select** for **Gateway Profile**, then select one of the following profiles:
   - For-third party gateways, select the **Default (All Capabilities)** profile for third-party gateways.
   - For Forcepoint Network Security Platforms managed by a different Management Server, select the appropriate version-specific profile .

5) On the **Endpoints** tab, click **Add**, then define the IP address for the endpoint:
   - If the endpoint has a static (manually defined) IP address, enter in it the **IPv4 Address** field.
   - If the endpoint has a dynamic (DHCP-assigned) IP address, select **Dynamic.**.

6) If the external gateway has a dynamic IP address:
   a) In the **Phase 1 ID** section at the bottom of the dialog box, change the **ID Type** to **E-mail**.

   b) Enter an email address in the **ID value** field.
   This email address can be any address that is not used as an ID in any of your other endpoints. The address entered here is used only as an identification, not for actually sending email.

   Note

   Make sure that the **ID Type** and **ID Value** match the identity configured on the external gateway device. If the device has a static IP address, make sure that the device uses it as its identity in this VPN or change the External VPN Gateway element configuration.

   c) Click **OK**.

7) Leave the properties dialog box open.

Next steps

Define a site for the external VPN gateway.

Related concepts

Defining VPN gateways on page 1173

Example VPN configuration 2: define a site for the external VPN gateway

The External VPN Gateway element needs a Site element.

Before you begin

You must have created an External VPN Gateway element for configuration 2. The **External VPN Gateway Properties** dialog box should still be open.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) On the **Sites** tab of the **External VPN Gateway Properties** dialog box, double-click the new Site element on the right.

2) (Optional) In the **Name** field, enter a descriptive name for the site.

3) Select or create the elements that represent the protected IP addresses behind the Gateway in the left pane, then click **Add** to include them.

   The internal IP address used as the source or destination address must be included in the site of the Gateway. Other traffic cannot use the VPN.

4) Click **OK** in both open dialog boxes.

Next steps

Create a VPN Profile element.

Example VPN configuration 2: create a VPN Profile element

The VPN Profile must contain VPN settings that match the settings defined on the external VPN gateway.

Before you begin

You must have defined a site for the external VPN gateway in configuration 2.



Note

This configuration scenario does not explain all settings related to VPN Profiles.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.

2) Browse to **Other Elements** > **Profiles** > **VPN Profiles**.

3) Right-click **VPN Profiles**, then select **New VPN Profile**.

4) In the **Name** field, enter a unique name.

5) On the **IKE SA** tab, configure the IKE SA settings.

   a) Select the **Version.**

   You can select **IKEv1**, **IKEv2**, or both. If both versions are selected, IKEv2 is tried first in the negotiations, and IKEv1 is only used if the remote gateway does not support IKEv2.

   b) In the **Cipher Algorithms** section, select one or more encryption methods that match the settings of the external gateway device.

   We recommend that you limit the selection to as few choices as possible. Do not select **DES** unless you are required to do so. DES is no longer secure because it is relatively easy to break DES encryption with modern computers. **3DES** (Triple-DES) has a relatively high overhead compared to other protocols with a comparable level of security. For this reason, 3DES is not a good choice when high throughput is required.

   Note

   The restricted (-R) product version has no strong encryption algorithms.

   c) Select the **Message Digest Algorithm** that matches the settings of the external gateway device.

   - In IKE, the message digest algorithm is used for integrity checking and key derivation.
   - If you select **SHA-2**, define the **Minimum Length** for the digest: **256**, **384**, or **512** bits. Set the length so that it is in line with the overall security strength.

   d) Select the Diffie-Hellman group or groups (used for key exchange) to be allowed to be used with the external gateway device.

   We recommend that you select from groups 14-21 according to the security requirements for the VPN. Groups 1, 2, and 5 are not sufficiently secure in all cases, although they might be required for interoperability with legacy systems.

   e) Select the **Authentication Method**.

   f) If IKEv1 is selected as the Version, adjust the **SA Lifetime in Minutes** to match the settings of the external gateway device.

   In IKEv2, lifetime is set locally, so it does not have to match the lifetime settings of the external gateway.

   g) If one of the Gateways has a dynamic IP address, change the **IKEv1 Negotiation Mode** to **Aggressive.**

6) On the **IPsec SA** tab, configure the IPsec SA settings.

   a) Select the **IPsec Type**:

   - The recommended setting is **ESP** (the communications are encrypted).
   - Usually, **AH** is not a valid option. The AH setting disables encryption for the VPN, fully exposing all traffic that uses the VPN to anyone who intercepts it in transit. You can use AH to authenticate and check the integrity of communications without encrypting them.

   b) In the **Cipher Algorithms** section, select one or more encryption methods that match the settings of the external gateway device

   - Do not select **Null**. This option disables encryption and allows anyone to view the data in transit.
   - Do not select **DES** unless you are required to do so. DES is no longer secure, as it is relatively easy to break DES encryption with modern computers.
   - **3DES** (Triple-DES) has a relatively high overhead compared to other protocols with a comparable level of security. It is not a good choice when high throughput is required.
   - **AES-GCM-128** or **AES-GCM-256** are recommended for high-speed networks.

   c) Select the **Message Digest Algorithm** that matches the settings of the external gateway device.

   - In IPsec, the message digest algorithm is used for integrity checking (except when authenticated encryption such as AES-GCM is used).
   - If you select **SHA-2**, define the **Minimum Length** for the digest: **256**, **384**, or **512** bits. Set the length so that it is in line with the overall security strength.

   d) Make sure that **Compression Algorithm** is set to **None**.

   The external gateway must not use compression.

   e) Adjust the **IPsec Tunnel Lifetime** to match the settings of the external gateway device.

   f) Select the **Security Association Granularity for Tunnel Mode** that matches the settings of the external gateway device.

   g) (Recommended) Select **Use PFS with Diffie-Hellman Group** if the external gateway device can use perfect forward secrecy (PFS), and select the Diffie-Hellman group to use with PFS.

   We recommend that you select from groups 14-21 according to the security requirements for the VPN. Groups 1, 2, and 5 are not considered sufficiently secure in all cases, although they might be required for interoperability with legacy systems.

7) Click **OK**.

Next steps

Create a Policy-Based VPN element.

Related concepts

Defining VPN profiles on page 1185

Example VPN configuration 2: create a Policy-Based VPN element

You must add a Policy-Based VPN element for this configuration.

Before you begin

You must have created a VPN Profile for configuration 2.

Note

This configuration scenario does not explain all settings related to Policy-Based VPN elements.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click **Policy-Based VPNs** in the element tree, then select **New Policy-Based VPN**.

2) In the **Name** field, enter a unique name.

3) From the **Default VPN Profile** drop-down list, select the VPN profile.

4) If you want to apply NAT rules to the communications that go through the VPN, select **Apply NAT to traffic that uses this VPN**.
   This setting does not affect the communications that the two gateways have with each other to set up and maintain the VPN. Communications between the gateways are always matched to the automatic rules or the NAT rules.

5) Click **OK**.
   The **VPN Editing** view opens on the **Site-to-Site VPN** tab.

6) Drag and drop the VPN Gateway element that represents the engine to **Central Gateways**.

7) Drag and drop the External VPN Gateway element to **Satellite Gateways**.

8) On the **Tunnels** tab, double-click the **Key** cell for the tunnel shown in the **Gateway<->Gateway** pane.

9) To match the pre-shared key between the two gateways:
   - To use the key that is automatically generated on the Management Server, click **Export**, then transfer the key securely to the external gateway.
   - To use a different key, replace the shown key with the one that you have agreed on with the administrator of the external gateway device.

CAUTION

The pre-shared key must be long and random to provide a secure VPN. Change the pre-shared key periodically (for example, monthly).

10) Click **OK** to close the **Pre-Shared Key** dialog box.

11) Make sure that the **Validity** column in the **Gateway<->Gateway** and the **End-Point<->End-Point** tables has a green check mark to indicate that there are no problems.

   a) If the **Validity** column of a tunnel has a warning icon, see the **Issues** pane to check what the problem is. If the pane is not shown, select **Menu** > **View** > **Panels** > **Issues**.

   b) If issues are shown, correct them as indicated. Long issues are easiest to read by hovering over the issue text so that the text is shown as a tooltip.

12) Click **Save**.

Next steps

Create Access rules.

Related concepts
Defining Policy-Based VPN elements on page 1193

Example VPN configuration 2: create Access rules

The rules in this example allow protected hosts to open connections both ways. Two rules are created here to allow the different directions of traffic separately.

VPN rules are matched based on source, destination, and service like any other rules.

Note

This configuration scenario does not explain all settings related to VPN Access rules.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Browse to **Engine** > **Policies** > **Engine Policies**.

3) Right-click the Engine policy that is used by the Security Engines involved in the VPN, then select **Edit Engine Policy**.

4) Add two IPv4 Access rules in a suitable location in the policy.

   - Make sure that rules for sending traffic through the VPN are above other rules that match the same traffic the **Allow**, **Discard**, or **Refuse** action.
   - Traffic that you do not want to send through the VPN must not match these rules. Traffic that is not routable through the VPN is dropped if it matches these rules.

5) Fill in the rules as outlined here. If NAT is enabled in the VPN, remember that the Access rules are checked before the NAT rules are applied.

Example VPN rules

Source	Destination	Service	Action
Local internal networks	Remote internal networks	Set as needed.	Select **Allow**, then open the Action options. Set **VPN Action** to **Enforce VPN**, then select a **Policy-Based VPN**.
Remote internal networks	Local internal networks	Set as needed.	Select **Allow**, then open the Action options. Set **VPN Action** to **Enforce VPN**, then select a **Policy-Based VPN**.

6) Save the policy.

CAUTION

If you continue to use this VPN, change the pre-shared key periodically (for example, monthly) to guarantee continued confidentiality of your data. Alternatively, you can switch to certificate-based authentication by creating a custom VPN profile.

7) Refresh the policies of all engines involved in the VPN to activate the new configuration.

Result

The VPN is established when traffic matches the Access rules created here. Example VPN configuration 2 is now complete.

Related concepts

Access rules for policy-based VPNs on page 1202

Example VPN configuration 3: Basic VPN for remote clients

This configuration scenario walks you through creating a mobile VPN between an Security Engine and more than one Forcepoint VPN Client.

To be able to configure a mobile VPN, the engine must have a static IP address (not assigned using DHCP or PPPoE).

Depending on the configuration that you want to use, you can add VPN client access to an existing site-to-site VPN as well. However, in this example scenario, a separate policy-based VPN is created for VPN clients.

This scenario assumes that automatic Site management is used, and that the Sites do not need to be changed.

In this scenario, the VPN settings are defined in a copy of the default **Suite-B-GCM-128** VPN Profile. The **Suite-B-GCM-128** VPN Profile contains the VPN settings specified for the VPN-A cryptographic suite in RFC 6379.

The configuration consists of the following general steps:

1) Configure virtual IP addresses for VPN clients.

2) Configure VPN settings for the Security Engine.

3) Create a VPN Profile element.

4) Create a Policy-Based VPN element.

5) Create User elements.

6) Create Access rules.

Begin by configuring virtual IP addresses for VPN clients.

Related tasks
Define virtual IP addresses for VPN clients on page 1281

Example VPN configuration 3: configure virtual IP addresses for VPN clients

VPN clients cannot use their local IP address in the internal corporate network. In this scenario, virtual IP addresses are used to solve this problem.

Before you begin

To use virtual IP addresses for VPN clients:

- You must use an external DHCP server to assign the IP addresses.
- The users must use a VPN client that has a Virtual Adapter feature. The Forcepoint VPN Client always has this feature installed and active.

You can use a Virtual Adapter to assign the VPN client an IP address in the VPN, independent of the address the VPN client computer uses in its local network. The virtual IP address is only used in communications through the VPN tunnels. The VPN gateway gets the IP address and network settings of the Forcepoint VPN Client from the an external DHCP server and forwards the information to the Forcepoint VPN Client.

This configuration scenario does not explain all settings related to VPN client address management.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click the engine element, then select **Edit Single Engine** or **Edit Engine Cluster**.

3) In the navigation pane on the left, browse to **VPN** > **VPN Client**.

4) From the **DHCP Mode** drop-down list, select how DHCP requests from VPN clients are sent.

5) From the **Interface** or **Interface for DHCP Relay** drop-down list, select the source address for the DHCP packets when querying the DHCP server (the interface toward the DHCP server).

6) Click **Add**, then select the DHCP server element that assigns IP addresses for the VPN clients.

7) (Optional) From the **Add Information** drop-down list, select what VPN Client user information is added to the Remote ID option field in the DHCP Request packets.
   - **Add User information** — VPN Client user information (in the form user@domain) is automatically added to the Remote ID option field in the DHCP Request packets.
   - **Add Group information** — VPN Client user information (in the form group@domain) is automatically added to the Remote ID option field in the DHCP Request packets.

   Your DHCP server must support the **DHCP Relay Agent Information** option to use this information. Depending on your DHCP server configuration, this information can be used as a basis for IP address selection.

8) (Optional) Select **Restrict Virtual Address Ranges**, then enter the IP address range in the field on the right.
   With this option, you can restrict the VPN clients' addresses to a set range, even if the DHCP server tries to assign another IP address. If an incorrect address is assigned, the user might not be able to access resources.

9) Click **Save**.

Next steps

Configure VPN settings for the Security Engine.

Related concepts
VPN client settings and how they work on page 1279

Related tasks
Define virtual IP addresses for VPN clients on page 1281

Example VPN configuration 3: configure VPN settings for the Security Engine

Configure certificates and sites for the mobile VPN.

Before you begin

The Security Engine must have a certificate for a mobile VPN. You can check the gateway certificates in the **Secure SD-WAN Configuration** > **Other Elements** > **VPN Certificates** > **Gateway Certificates** branch.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click the engine element, then select **Edit Single Engine** or **Edit Engine Cluster**.

3) In the navigation pane on the left, browse to **VPN** > **Certificates**.

4) Make sure that **Automated RSA Certificate Management** is selected.

5) In the navigation pane on the left, browse to **VPN** > **Sites**.
   The Sites represent the internal addresses that VPN clients can reach through the VPN. Sites do not grant any host access directly. The Access rules define the allowed connections.

6) (Optional) Leave **Add and update addresses based on routing** selected.
   This option automatically updates this information based on routing changes. You can exclude some interfaces while keeping the others automatically updated.

7) (Optional) Select the internal networks that you want to exclude from the VPN by disabling the interface they are under in the automatic site.
   Disabled interfaces are grayed-out.
   - If you want to include some individual network that is under an otherwise disabled interface, drag and drop it from under the disabled interface onto the Site element. The element is copied to the higher level. The copied definition is not updated automatically.
   - The Sites must include only internal networks. Do not add interfaces with the **Any Network** element in this type of VPN.

8) Click **Save**.

Next steps

Create a VPN Profile element.

Related concepts
Defining VPN gateways on page 1173
Defining Site elements for VPN gateways on page 1180

Example VPN configuration 3: create a VPN Profile element

You must create a custom VPN Profile element to define the settings for VPN clients.

Note
This configuration scenario does not explain all settings related to authenticating VPN client users.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.

2) In the element tree, browse to **Other Elements** > **Profiles** > **VPN Profiles**.

3) Right-click **Suite-B-GCM-128** and select **New** > **Duplicate**.
The settings from the default profile are copied into the **VPN Profile Properties** dialog box that opens.

4) In the **Name** field, enter a unique name.

5) On the **IKE SA** tab, configure the IKE SA settings.
   a) In the **Version** drop-down list, select the IKE version.
   You can select **IKEv1**, **IKEv2**, or both. If both versions are selected, IKEv2 is tried first in the negotiations, and IKEv1 is only used if the remote gateway does not support IKEv2.

   b) (Only if IKEv1 is selected) Make sure **IKEv1 Negotiation Mode** is set to **Main**.
   Using Main mode helps guarantee that the user names and passwords of the VPN client users remain confidential.

6) On the **IPsec Client** tab, configure the VPN client settings.
   a) Make sure that the **Authentication Method** is set to **RSA Signatures**.

   b) Select **Allow Hybrid/EAP Authentication**.
   Hybrid authentication is used with IKEv1. EAP (Extensible Authentication Protocol) is used with IKEv2.

   c) Make sure **IPsec Security Association Granularity for Tunnel Mode** is set to **SA Per Net**.

7) Click **OK**.

Next steps

Create a Policy-Based VPN element.

Example VPN configuration 3: create a Policy-Based VPN element

You must add a Policy-Based VPN element for this configuration.

Before you begin

You must have a custom VPN Profile element for configuration 3.



Note

This configuration scenario does not explain all settings related to Policy-Based VPN elements.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.

2) Right-click **Policy-Based VPNs** in the element tree, then select **New Policy-Based VPN**.

3) In the **Name** field, enter a unique name.

4) From the **Default VPN Profile** drop-down list, select the custom VPN Profile that you created.

5) Select **Apply NAT Rules to Traffic That Uses This VPN**.
   This option applies the NAT rules in the policy and the global NAT definition for the Engine.

6) Click **OK**.
   The **VPN Editing** view opens on the **Site-to-Site VPN** tab.

7) Drag and drop the VPN Gateway element that represents the engine to **Central Gateways**.

8) On the **Mobile VPN** tab, select **Only central Gateways from overall topology** to define which VPN Gateways provide Mobile VPN access.

9) On the **Tunnels** tab, make sure that the **Validity** column in the **Gateway<->Gateway** and the **End-Point<->End-Point** tables has a green check mark to indicate that there are no problems.
   a) If the **Validity** column of a tunnel has a warning icon, see the **Issues** pane to check what the problem is. If the pane is not shown, select **Menu** > **View** > **Panels** > **Issues**.
   b) If issues are shown, correct them as indicated. Long issues are easiest to read by hovering over the issue text so that the text is shown as a tooltip.

10) Click **Save**.

Next steps

Create User elements.

Related concepts
Defining Policy-Based VPN elements on page 1193

Example VPN configuration 3: create User elements

User authentication is configured in the same way for VPN client connections and normal, unencrypted connections. The same User elements (user accounts) can be used for both.

Note

This configuration scenario does not explain all settings related to user authentication.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Make sure that user database replication is enabled for the Security Engine.
   a) Right-click the Security Engine, then select **Options**.
   b) Make sure that **User DB Replication** is selected.
2) Select **User Authentication**.
3) Browse to **Users** > **InternalDomain**.
4) Right-click the **stonegate** Internal User Group, then select **New** > **Internal User**.
5) In the **Name** field, enter the user name that the end user uses to authenticate to the VPN.
6) On the **Authentication** tab, click **Add** in the **Authentication Method** section.
7) Select **User Password** and click **Select**.
   This default element allows user password authentication for the internal LDAP database.
8) In the **Password** and **Confirm Password** fields, enter and confirm the password.
   Make a note of the password so that you can communicate it to the user. The passwords entered in the VPN client are encrypted so that they remain confidential as they are transferred over the Internet.
9) Click **OK**.

Result

The information is added to the Management Server's internal LDAP user database.

Next steps

Create Access rules.

Related concepts

Getting started with user authentication on page 1127

Example VPN configuration 3: create Access rules

Create a rule to allow specific users access to internal networks after having authenticated.

The authentication connection from VPN clients is allowed in the Firewall Template. Authentication is always required to establish a VPN tunnel. VPN client connections are matched based on Source, Destination, and Service like any other traffic. The example rule matches only specific users and only after the users have already successfully authenticated. We recommend always adding the authentication requirement to rules that are specific to VPN clients.

After the VPN tunnel is established, any connection from the VPN clients to the internal network is matched against the Access rules as usual. The example rule that is created here allows these connections.

Note

This configuration scenario does not explain all settings related to VPN Access rules.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Browse to **Engine** > **Policies** > **Engine Policies**.

3) Right-click the Engine policy that is used by the Security Engines involved in the VPN, then select **Edit Engine Policy**.

4) Add an IPv4 Access rule in a suitable location in the policy and configure the rule as outlined here:

Example VPN rule

Source	Destination	Service	Action	Authentication
Network element that represents the virtual IP address range for the VPN Client	Local internal networks	Set as needed.	Select **Allow**, then open the Action options. Set **VPN Action** to **Enforce VPN**, then select a **Policy-Based VPN**.	**Users** tab: stonegate Internal User Group (under InternalDomain). **Authentication Methods** tab: ANY or a specific method.

5) Save the policy.

6) Refresh the policies of all engines involved in the VPN to activate the new configuration.

Result

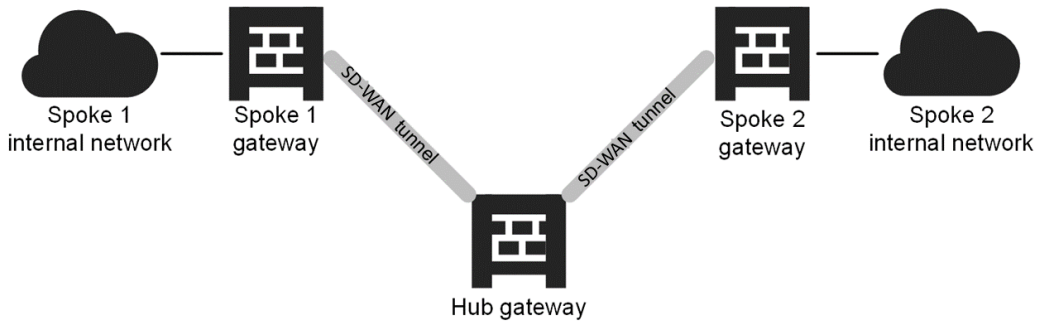
The VPN is established when traffic matches the created Access rules. Example VPN configuration 3 is now complete.

Related concepts
Access rules for policy-based VPNs on page 1202

Example VPN configuration 4: Basic VPN hub

In a VPN hub configuration, a gateway is configured to forward VPN traffic between different VPN tunnels.

The gateway that does this forwarding is called a hub gateway. The gateways that contact each other through a hub are called spoke gateways.



All traffic from the Spoke 1 internal network to the Spoke 2 internal network, and from the Spoke 2 internal network to the Spoke 1 internal network is sent through a VPN tunnel to the hub gateway. The hub gateway forwards the traffic through another VPN tunnel to its destination.

The hub gateway must be set up specifically as a hub. The hub configuration is reflected in the topology, the Site definitions, and the VPN Access rules. The spoke gateways do not require any hub-specific configuration. In this example configuration, VPN tunnels are established from all spoke gateways to the hub gateway. All networks of all gateways are configured as reachable through the hub. Connections are allowed only as defined in the Engine Access rules.

Note

There must not be duplicate endpoint-to-endpoint tunnels in different VPN. If there are existing tunnels between the hub gateway and the other gateways in other active VPN, you must remove the overlapping configurations.

This scenario explains a configuration in which all connections are defined within the same Policy-Based VPN element. A single Policy-Based VPN element is simpler to set up and maintain than forwarding traffic between VPN tunnels defined in different Policy-Based VPN elements. In this scenario, all gateways are Engines controlled by the same Management Server. You can add External VPN Gateways to this configuration even though their creation is not covered in detail in this workflow.

The configuration consists of the following general steps:

1) Create a Policy-Based VPN element.

2) Create a Site element for the hub gateway.

3) Create Access rules.

Begin by creating a Policy-Based VPN element.

Example VPN configuration 4: create a Policy-Based VPN element

You must add a VPN element for this configuration.



Note

This configuration scenario does not explain all settings related to Policy-Based VPN elements.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.

2) Right-click **Policy-Based VPNs** in the element tree, then select **New Policy-Based VPN**.

3) In the **Name** field, enter a unique name.

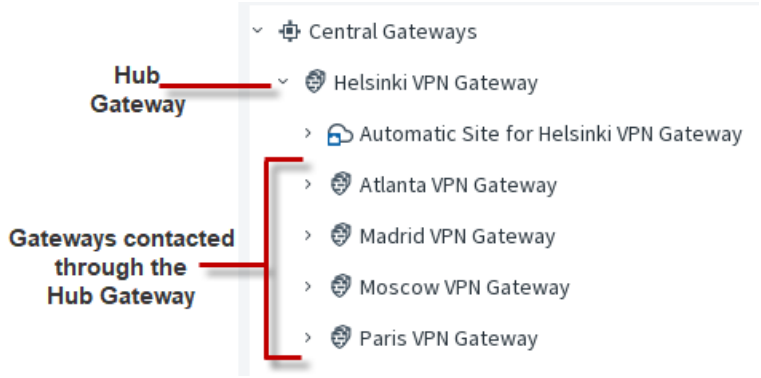
4) From the **Default VPN Profile** drop-down list, select **Suite-B-GCM-128**.

5) If you want to apply NAT rules to the communications that go through the VPN, select **Apply NAT to traffic that uses this VPN**.
   This setting does not affect the communications that the two gateways have with each other to set up and maintain the VPN. Communications between the gateways are always matched to the automatic rules or the NAT rules.

6) Click **OK**.
   The **VPN Editing** view opens on the **Site-to-Site VPN** tab.

7) Drag and drop the engine that acts as the hub gateway to **Central Gateways**.

VPN editing view - Site-to-Site VPN tab



8) Drag and drop the other engines on top of the hub gateway so that the engines are added as branches (spokes) under the hub gateway.
Spoke gateways can be any other engines or External VPN Gateways.

9) Click **Save**, but do not close the **VPN Editing** view.

Next steps

Define a Site element for the hub gateway.

Related concepts
Defining Policy-Based VPN elements on page 1193

Example VPN configuration 4: define a Site element for the hub gateway

The VPN Gateway that acts as the hub gateway needs a Site element.

Note
This configuration scenario does not explain all settings related to Site elements.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click the **VPN Gateway** that acts as the hub gateway, then select **New** > **Site**.

2) In the **Name** field, enter a unique name.

3) Add all networks protected by the spoke gateways to the site contents on the right.
   After you add the protected networks, the site contains all remote IP addresses that are used in spoke-to-hub traffic that is forwarded from the hub to other spokes. The site should not contain the hub gateway's local networks. These are defined using the automatic site management features in this example.

4) On the **VPN References** tab, select **Enable** for this VPN element, then deselect it for all other VPNs.
   The site is still shown in all VPNs, but is grayed-out (disabled) and not included in the configuration.

5) In the **Mode** cell, select **Hub** to activate VPN hub-related features for the VPN Gateway.

6) Click **OK** to close the dialog box.
   You return to the main **VPN editing** view.

7) Click the **Tunnels** tab.

8) Check that the **Validity** column in the **Gateway<->Gateway** and the **End-Point<->End-Point** tables has a green checkmark to indicate that there are no problems.
   a) If the **Validity** column of a tunnel has a warning icon, see the **Issues** pane to check what the problem is. If the pane is not shown, select **Menu** > **View** > **Panels** > **Issues**.

   b) If issues are shown, correct them as indicated. Long issues are easiest to read by hovering over the issue text so that the text is shown as a tooltip.

9) Click **Save**.

Next steps

Create Access rules.

Related concepts
Defining Site elements for VPN gateways on page 1180

Example VPN configuration 4: create Access rules

The rules in this example allow connections between hosts in protected networks of all gateways to connect to all other protected networks.

Note

This configuration scenario does not explain all settings related to VPN Access rules.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Browse to **Engine** > **Policies** > **Engine Policies**.

3) Add rules to the policy that is used by the Security Engine that acts as a hub.

   a) Right-click the Engine policy, then select **Edit Engine Policy**.

   b) Add the following rules in a suitable location in the policy:

   Make sure that rules for sending traffic through the VPN are above other rules that match the same traffic with the **Allow**, **Discard**, or **Refuse** action. Traffic that you do not want to send through the VPN must not match this rule. Traffic that is not routable through the VPN is dropped if it matches this rule. If NAT is enabled in the VPN, remember that the Access rules are checked before the NAT rules are applied.

   Example VPN rules in the hub policy

Source	Destination	Service	Action	Source VPN
Rules for traffic between the hub spoke 1, and between spoke 1 and spoke 2				
Hub gateway	Spoke 1 internal network	Set as needed.	Select **Allow**, then open the Action options. Set **VPN Action** to **Enforce VPN**, then select the **Policy-Based VPN** element that you created.	Your **Policy-Based VPN** element
Spoke 1 internal network	Hub gateway	Set as needed.	Select **Allow**, then open the Action options. Set **VPN Action** to **Enforce VPN**, then select the **Policy-Based VPN** element that you created.	Your **Policy-Based VPN** element
Spoke 1 internal network	Spoke 2 internal network	Set as needed.	Select **Allow**, then open the Action options. Set **VPN Action** to **Forward**, then select your **Policy-Based VPN**.	Your **Policy-Based VPN** element
Rules for traffic between the hub spoke 2, and between spoke 2 and spoke 1				
Hub gateway	Spoke 2 internal network	Set as needed.	Select **Allow**, then open the Action options. Set **VPN Action** to **Enforce VPN**, then select the **Policy-Based VPN** element that you created.	Your **Policy-Based VPN** element
Spoke 2 internal network	Hub gateway	Set as needed.	Select **Allow**, then open the Action options. Set **VPN Action** to **Enforce VPN**, then select the **Policy-Based VPN** element that you created.	Your **Policy-Based VPN** element
Spoke 2 internal network	Spoke 1 internal network	Set as needed.	Select **Allow**, then open the Action options. Set **VPN Action** to **Forward**, then select your **Policy-Based VPN**.	Your **Policy-Based VPN** element

   c) Save the policy.

4) Add rules to the policy that is used by the Security Engine that acts as spoke 1.

   a) Right-click the Engine policy, then select **Edit Engine Policy**.

   b) Add the following rules in a suitable location in the policy:

   Make sure that rules for sending traffic through the VPN are above other rules that match the same traffic with the **Allow**, **Discard**, or **Refuse** action. Traffic that you do not want to send through the VPN must not match this rule. Traffic that is not routable through the VPN is dropped if it matches this rule. If NAT is enabled in the VPN, remember that the Access rules are checked before the NAT rules are applied.

   Example VPN rules in the spoke 1 policy

Source	Destination	Service	Action	Source VPN
Spoke 1 internal network	Hub gateway	Set as needed.	Select **Allow**, then open the Action options. Set **VPN Action** to **Enforce VPN**, then select the **Policy-Based VPN** element that you created.	Your **Policy-Based VPN** element
Hub gateway	Spoke 1 internal network	Set as needed.	Select **Allow**, then open the Action options. Set **VPN Action** to **Enforce VPN**, then select the **Policy-Based VPN** element that you created.	Your **Policy-Based VPN** element
Spoke 1 internal network	Spoke 2 internal network	Set as needed.	Select **Allow**, then open the Action options. Set **VPN Action** to **Forward**, then select your **Policy-Based VPN**.	Your **Policy-Based VPN** element
Spoke 2 internal network	Spoke 1 internal network	Set as needed.	Select **Allow**, then open the Action options. Set **VPN Action** to **Forward**, then select your **Policy-Based VPN**.	Your **Policy-Based VPN** element

   c) Save the policy.

5) Add rules to the policy that is used by the Security Engine that acts as spoke 2.

   a) Right-click the Engine policy, then select **Edit Engine Policy**.

b) Add the following rules in a suitable location in the policy:

Make sure that rules for sending traffic through the VPN are above other rules that match the same traffic with the **Allow**, **Discard**, or **Refuse** action. Traffic that you do not want to send through the VPN must not match this rule. Traffic that is not routable through the VPN is dropped if it matches this rule. If NAT is enabled in the VPN, remember that the Access rules are checked before the NAT rules are applied.

Example VPN rules in the spoke 2 policy

Source	Destination	Service	Action	Source VPN
Spoke 2 internal network	Hub gateway	Set as needed.	Select **Allow**, then open the Action options. Set **VPN Action** to **Enforce VPN**, then select the **Policy-Based VPN** element that you created.	Your **Policy-Based VPN** element
Hub gateway	Spoke 2 internal network	Set as needed.	Select **Allow**, then open the Action options. Set **VPN Action** to **Enforce VPN**, then select the **Policy-Based VPN**element that you created.	Your **Policy-Based VPN** element
Spoke 2 internal network	Spoke 1 internal network	Set as needed.	Select **Allow**, then open the Action options. Set **VPN Action** to **Forward**, then select your **Policy-Based VPN**.	Your **Policy-Based VPN** element
Spoke 1 internal network	Spoke 2 internal network	Set as needed.	Select **Allow**, then open the Action options. Set **VPN Action** to **Forward**, then select your **Policy-Based VPN**.	Your **Policy-Based VPN** element

c) Save the policy.

6) Refresh the policies of all engines involved in the VPN to activate the new configuration.

CAUTION

If you continue to use this VPN, change the pre-shared key periodically (for example, monthly) to guarantee continued confidentiality of your data. Alternatively, you can switch to certificate-based authentication by creating a custom VPN profile.

Result

The VPN is established when traffic matches the created Access rules. Example VPN configuration 4 is now complete.

Related concepts

Access rules for policy-based VPNs on page 1202

Chapter 75

Managing VPN certificates

Contents

- VPN certificates and how they work on page 1251
- Define additional VPN certificate authorities on page 1254
- Create an internal ECDSA certificate authority for VPN gateways on page 1256
- Select the default internal certificate authority on page 1257
- Create a VPN certificate or certificate request for a VPN Gateway element on page 1257
- Import an externally signed VPN gateway certificate on page 1259
- Sign external VPN certificate requests with an internal certificate authority on page 1260
- Select which internal certificate authority signs each certificate on page 1261
- Replacing expired VPN certificates on page 1262
- Export signed VPN gateway certificates or VPN certificate authority certificates on page 1264
- Check when VPN gateway certificates expire on page 1266
- Check when VPN certificate authorities expire on page 1266

A digital certificate is a proof of identity. Forcepoint Network Security Platform in the Engine/VPN role supports using certificates for authenticating gateways and the Forcepoint VPN Client.

VPN certificates and how they work

Certificates can be used for authenticating VPN gateways and the Forcepoint VPN Client.

In site-to-site VPNs, you can use both pre-shared keys and certificates as the authentication method. In mobile VPNs, certificates are always needed when the Forcepoint VPN Client is involved. However, if you use the hybrid authentication method with the Forcepoint VPN Client, only the gateway needs a certificate.

Certificates do not contain information that is specific to a particular VPN. You can use certificates for authentication in both the policy-based and route-based VPNs. A certificate authority (CA) issues certificates as proof of identity. Gateways that form a VPN tunnel are configured to trust the CA that signed the other gateway's certificate. All certificates issued by a trusted CA are accepted as valid, so certificates can be added, renewed, and changed without affecting the VPN as long as the actual identity information is correct. The same certificate can be used for any number of VPNs with any number of gateways and VPN clients.

Certificates are always required for gateways to which the Forcepoint VPN Client connects. Certificates can optionally be used to identify VPN clients, but are not mandatory.

Certificates reduce the required maintenance work, because they do not have to be changed as frequently as pre-shared keys. All certificates are created with an expiration date, after which the certificate is no longer valid. Certificates signed by an Internal RSA CA for Gateways or an Internal ECDSA CA for Gateways are valid for three years from their creation. When a certificate expires, a new certificate is needed.

Certificate management in VPNs

Certificate management tasks in the SMC mostly involve VPN Gateways that represent engines.

VPN certificates can be generated by any internal or external certificate authorities that both gateways are configured to trust. There are several options for signing VPN Gateway certificates:

- The Management Server includes a dedicated Internal RSA CA for Gateways and optionally an Internal ECDSA CA for Gateways for signing VPN certificates. You use these certificate authorities through the SMC Client.
- One Internal CA for Gateways can be selected as the default CA. Certificate management can be automatic if the certificate is signed using the Management Server's internal default CA.
- You can create certificate requests in the SMC Client, export them, sign them using an external CA, and then import the signed certificate back into the SMC.

RSA certificates can be created and renewed automatically using the default CA. Some manual steps are required in the following cases:

- You have both an Internal RSA CA for Gateways and an Internal ECDSA CA for Gateways. Only one Internal CA for Gateways can be selected as the default certificate authority. You must manually create and renew any certificates that are not signed by the default CA.
- You use DSA certificates.
- You want to use an external CA to sign certificates.

The Internal RSA CA for Gateways or Internal ECDSA CA for Gateways can also sign certificate requests created by external components. This feature is meant to support VPN client deployments. If you have used the Internal RSA CA for Gateways or Internal ECDSA CA for Gateways to sign certificate requests, you cannot cancel the issued certificates. Consider how widely you can use them for signing external certificate requests within your organization.

Limitations of certificates in VPNs

Certificates in VPNs have these limitations.

- All gateways in the same VPN must support the same CA algorithm. Otherwise, VPN communication fails. For example, if you use an Internal ECDSA CA for Gateways as the default CA, all other gateways used in the same VPN must support ECDSA.
- Certificates created for VPN gateways for establishing the VPN are stored on the VPN gateway devices (Engines). These certificates are not included in the Management Server backup, and are not changed in any way when a Management Server backup is restored.
- Certificates can become unusable if the private key for that certificate is lost. The key can be lost, for example, if the Security Engine hardware fails and must be replaced. Engine Clusters share each VPN certificate and can synchronize the private key from node-to-node as needed. If the private key is erased from a Single Engine or from all the nodes of a Engine Cluster, a new certificate must be created.
- Externally issued VPN certificates can be revoked by the certificate authority that issued them. This safety measure is used when the certificate is suspected to be compromised.
- A single VPN Gateway can use only one type of certificate (for example, RSA or ECDSA). However, if different VPNs require different certificate types, then different VPN Gateway elements must be created for those VPNs.

Validity of VPN certificates

VPN certificates are always valid starting from a specific date and time and expire at a specific date and time in the future.

All components that use (or sign) certificates must have the correct time settings to avoid unexpected certificate rejections. The **Internal RSA CA for Gateways** and the **Internal ECDSA CA for Gateways** of the Management Server generate certificates that are valid starting immediately until three years from their creation.

A Certificate Revocation List (CRL) or online Certificate Status Protocol (OCSP) server can be used to cancel a certificate before it reaches its expiration. For example, a certificate might be revoked if unauthorized parties have obtained a copy of both the certificate and the associated private key. The Internal RSA CA for Gateways and the Internal ECDSA CA for Gateways do not support certificate revocation lists. If you want to use certificate validation, you must use an external certificate authority (either one you maintain yourself or a commercial service). The Security Engine contacts the certificate validation servers using HTTP. If all defined certificate validation servers are unreachable, the certificates are treated as invalid until the validity of the certificate can be checked.

Internal VPN certificate authorities

The Management Server includes a dedicated Internal RSA CA for Gateways and optionally an Internal ECDSA CA for Gateways for signing VPN certificates.

You can use both an Internal ECDSA CA for Gateways and an Internal RSA CA for Gateways at the same time.

The internal certificate authorities run on the same computer as the Management Server. If you have both types of internal certificate authorities, only one certificate authority can be selected as the default certificate authority. Only the default CA is used in automated RSA certificate management. You must manually create and renew any certificates that are not signed by the default CA.

If you want to use the internal certificate authorities to sign other certificates, you must export, transfer, and import certificate requests and signed certificates manually. The Internal RSA CA for Gateways and the Internal ECDSA CA for Gateways do not support certificate revocation lists. We do not recommend using the internal certificate authorities to sign certificates for components that are outside the control of your organization.

The Internal RSA CA for Gateways and the Internal ECDSA CA for Gateways are each valid for 10 years. A new Internal RSA CA for Gateways or Internal ECDSA CA for Gateways is automatically created to replace the default certificate authority 6 months before the expiration date. The internal certificate authority that is not selected as the default certificate authority is not automatically renewed.

If automatic RSA certificate management is activated for an Security Engine, RSA certificates issued by the default certificate authority are renewed automatically. You must manually renew certificates if the certificate-related files, including the private key stored on the engines, are damaged or lost. You must also manually create and renew any certificates that are not signed by the default certificate authority. If certificates for authenticating VPN client users were signed by the expiring Internal CA for Gateways, you must manually create new certificates for the VPN clients. You must also create new certificates manually for external components that have certificates signed by the Internal RSA CA for Gateways or the Internal ECDSA CA for Gateways.

External certificate authorities for VPNs

External certificate authorities can create certificates for VPN Gateways, External VPN Gateways, or VPN clients.

All IPsec certificates follow the ITU-T X.509 standard, which is also used in protocols such as TLS/SSL and HTTPS. External certificate authorities are especially useful when creating VPNs with partner organizations. Using external certificate authorities allows both organizations to use their preferred certificate authority. Different gateways in a VPN can have certificates signed by different certificate authorities.

To make Security Engines accept externally signed certificates of external components, you must import the public key of the external certificate authority into the SMC and add the CA to the list of trusted certificate authorities for the VPN Profile and for the Security Engine.

To create a certificate for Security Engines or the Forcepoint VPN Client, you must generate a certificate request and have it signed by the external certificate authority. The external certificate authority must support PKCS#10 certificate requests in PEM format and the signed certificates must also be in the PEM format. Furthermore, the certificate authority must be able to copy all attributes from the certificate request into the certificate. Especially, the X.509 extension Subject Alternative Name must be copied into the certificate because its value is used for identification.

VPN certificate configuration overview

Configuring VPN certificates involves several main steps.

1) (Optional) If you want to use certificates that are signed by some external certificate authority (CA), define the CA in the SMC Client.

2) (Optional) If you want to use an Internal ECDSA CA for Gateways to sign certificates, create an Internal ECDSA CA for Gateways.

3) (Optional) If you have both an Internal RSA CA for Gateways and an Internal ECDSA CA for Gateways, select which CA is the default.

4) Start by creating a VPN certificate or certificate request for a VPN Gateway in the following cases:
   - To use an externally signed certificate.
   - To use a ECDSA certificate.
   - If automated RSA certificate management is disabled for gateways.

5) (For externally signed certificates) When the certificate is signed, import the certificate.

6) Select a certificate-based **Authentication Method** on the **IKE SA** tab of the VPN Profile.

Define additional VPN certificate authorities

If you want to use certificates that are signed by an external CA, define an additional VPN CA.

Before you begin

You must have the root certificate (or a valid certificate) from the certificate authority.

You must define additional VPN CAs in the following cases:

- In a VPN with an external gateway where you do not want to use the Internal RSA CA for Gateways or the Internal ECDSA CA for Gateways to create a certificate for the external gateway. The external gateway must also be configured to trust the issuer of the certificate.
- If you want to use a certificate signed by an external CA for a VPN Gateway or for a VPN client.

Note

Only the Internal RSA CA for Gateways and Internal ECDSA CA for Gateways of your SMC are configured as trusted CAs for gateways in VPNs by default. The Internal RSA CA for Gateways is automatically created when you install the SMC.

You can configure the CA as trusted by importing its certificate to VPN Certificate Authorities. The certificates must be X.509 certificates in PEM format (Base64 encoding). It might be possible to convert between formats using, for example, OpenSSL or the certificate tools included in Windows.

The CAs you use can be either private (for self-signed certificates) or public (commercial certificate issuers). When you define a CA as trusted, all certificates signed by that CA are valid until their expiration date (or until the CA's certificate expires). Optionally, you can also enable certificate revocation checking by the VPN gateways by enabling validity check options in the properties of the VPN Certificate Authority element. Certificate Revocation Lists (CRLs) and Online Certificate Status Protocol (OCSP) are supported when the issued certificates contain pointers for the revocation information. The CA can cancel a certificate, for example, because it is compromised.

By default, all CAs you have defined are trusted by all gateways and in all VPNs. If necessary, you can limit trust to a subset of the defined CAs when you configure the VPN Gateway and VPN Profile elements. The trust relationships can be changed at the gateway level and in the VPN Profiles.

To obtain a certificate from an external certificate authority, first create element for the certificate authority.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.
2) Browse to **Other Elements** > **VPN Certificates** > **VPN Certificate Authorities**.
3) Right-click **VPN Certificate Authorities**, then select **New VPN Certificate Authority**.
4) On the **General** tab, configure the settings.

   Note

   All fields but the **Name** on the **General** tab are grayed out. The grayed out fields are always filled in automatically based on information contained in the certificate you import. You cannot change the information in the grayed out fields. The information is shown when you close and reopen the VPN Certificate Authority element after importing the information.

   CAUTION

   When certificate checking is defined, all certificates signed by the CA are treated as invalid if the validity check cannot be performed. For example, the validity check might not be performed due to incorrectly entered addresses or connectivity problems.

5) On the **Certificate** tab, import the certificate in one of the following ways:
   - Click **Import**, then import a certificate file.

- Copy and paste the information into the field. Include the “Begin Certificate” header and “End Certificate” footer in the information that you copy and paste.



Tip

You can copy and paste the certificate information for many public certificate authorities from the default Trusted Certificate Authority elements. The default Trusted Certificate Authority elements are in the **Configuration** view under **Administration** > > **Certificates** > **Certificate Authorities** > **Trusted Certificate Authorities**.

6) Click **OK**.

Next steps

If you see an invalid certificate error, the certificate you imported might be in an unsupported format. Try converting the certificate to an X.509 certificate in PEM format (Base64 encoding) using OpenSSL or the certificate tools included in Windows.

If your Engine Policy is based on the Firewall Template, both LDAP (port 389) and HTTP (port 80) connections from the Engine are allowed. If your engine or server configuration differs from these standard definitions, edit the Engine Policy to allow the necessary connections from the Engines.

Create an internal ECDSA certificate authority for VPN gateways

The Elliptic Curve Digital Signature Algorithm (ECDSA) is a digital signature algorithm that uses elliptic curve cryptography. If you want to use the ECDSA signature algorithm for signing VPN certificates, create an Internal ECDSA CA for Gateways.

You can create one Internal ECDSA CA for Gateways. You can use both an Internal ECDSA CA for Gateways and an Internal RSA CA for Gateways at the same time. When there is more than one valid CA, you can select which CA signs each certificate.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.

2) Browse to **Other Elements** > **VPN Certificates** > **VPN Certificate Authorities**.

3) Right-click **VPN Certificate Authorities**, then select **Create New VPN ECDSA Certificate Authority**.
A new Internal ECDSA CA for Gateways is created.

Result

The ECDSA CA for Gateways is ready to use for signing certificates.

Select the default internal certificate authority

If you have both an Internal RSA CA for Gateways and an Internal ECDSA CA for Gateways, only one certificate authority can be selected as the default certificate authority.

Only the default certificate authority is used in automated RSA certificate management. You must manually create and renew any certificates that are not signed by the default CA.

CAUTION

All gateways in the same VPN must support the CA algorithm used by the default certificate authority. Otherwise, VPN communication fails.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.
2) Browse to **Other Elements** > **VPN Certificates** > **VPN Certificate Authorities**.
3) Right-click the Internal CA for Gateways that is not currently the default certificate authority, then select **More actions** > **Set Default Certificate Authority**.

Create a VPN certificate or certificate request for a VPN Gateway element

You can create a certificate request and sign it either using an Internal CA for Gateways or an external certificate authority (CA).

If automated RSA certificate management is active for the VPN Gateway, these steps are necessary only in the following cases:

- You have both an Internal RSA CA for Gateways and an Internal ECDSA CA for Gateways. Only the default CA is used in automated RSA certificate management. You must manually create and renew any certificates that are not signed by the default CA.
- You want to use DSA certificates.
- You want to create a certificate request to be signed by an external CA.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.
2) Select **VPN Gateways**.
   The gateways are displayed.

3) Right-click the VPN Gateway element and select **More Actions** > **Generate Certificate**.

4) In the **Generate Certificate** dialog box, enter the certificate information.

5) Select the **Public Key Algorithm** according to the requirements of your organization.



Note

The Public Key Algorithm can be different from the internal CA type. For example, you can use RSA key algorithm with an Internal ECDSA CA for Gateways.

6) Select how you want to Sign the certificate.

7) (Optional) Select the **Signature Algorithm** used to sign the certificate signing request and for an internal CA to sign the certificate.

- If you selected an Internal CA for Gateways, you can define the Signature Algorithm if the selected Public Key Algorithm is compatible with the algorithm used by the Internal CA. In other cases, the default algorithm for the Internal CA is used (for example, RSA / SHA-256 for Internal RSA CA for Gateways).
- If you selected an external certificate authority, you can define a Signature Algorithm that is compatible with the selected Public Key Algorithm type.

8) (Optional, if supported by the Public Key Algorithm) Enter the **Key Length** for the generated public-private key pair.

- The default Key Length depends on the Public Key Algorithm.
- The Key Length cannot be changed for some Public Key Algorithms.

9) Click **OK**.

There might be a slight delay while the certificate request is generated. If you signed the certificate using an Internal CA for Gateways, the certificate is automatically transferred to the Engine and no further action is needed.

The signed certificate or unsigned certificate request is added under the gateway in the gateway list.

10) (With external certificate authorities only) Right-click the certificate request, select **Export Certificate Request**, and save it.

- To generate certificates for a VPN Gateway element, the CA must support PKCS#10 certificate requests in PEM format (Base64 encoding). The signed certificates must also be in the PEM format. It might be possible to convert between formats using, for example, OpenSSL or the certificate tools included in Windows.
- The CA must be able to copy all attributes from the certificate request into the certificate. In particularly, the X.509 Subject Alternative Name extension must be copied as it is in the request when the value is used for identification in VPN negotiation.

When you receive the signed certificate, import it.

Related tasks

Import an externally signed VPN gateway certificate on page 1259

Import an externally signed VPN gateway certificate

You can import a certificate signed by an external certificate issuer for a VPN Gateway element when the certificate request has been created in the SMC.

For security reasons, it is not possible to import externally generated private keys.

Note

Prior to software versions 6.10.8 all CAs that issues certificates for your VPNs must be configured in the SMC and be included as trusted both at the gateway and VPN Profile levels. In later versions only trust anchor certificates must be configured as trusted. Possible intermediate CAs must be included in the certificate bundle that are being imported as the VPN gateway certificate.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.
2) Open the **VPN Gateways** branch and expand the tree under the VPN Gateway element.
3) Right-click the certificate request and select **Import Certificate**.
4) Select the certificate authority that signed the certificate.
5) Browse to the certificate file on your local workstation or copy and paste the content of the certificate into the dialog box.

   If you copy and paste the certificate, include the“Begin Certificate Request” header and the “End Certificate Request” footer.

   Note

   Starting with version 6.10.8: If VPN gateway certificate is signed by an intermediate CA, imported certificate must be a certificate bundle that contains all the VPN gateway certificate and all intermediate certificates. SMC verifies if the complete certificate chain is present when importing. Certificate bundle is concatenation of the certificates in PEM format.

   The signed certificate is imported and transferred to the engine automatically.

Related tasks

Create a VPN certificate or certificate request for a VPN Gateway element on page 1257

Sign external VPN certificate requests with an internal certificate authority

You can use an internal certificate authority to sign VPN certificate requests for VPN clients and internal VPN gateways.

Before you begin

For VPN clients, you must have a PKCS#10 certificate request file in PEM format. For internal VPN gateways, you must have already generated a certificate request.

You can use the SMC's Internal RSA CA for Gateways and Internal ECDSA CA for Gateways to sign external certificate requests. You can also use an internal certificate authority to sign any certificate request that is in the supported format (PKCS#10 certificate requests in PEM format). An alternative is to configure the Internal Gateway to accept an externally signed certificate by defining the external certificate issuer as trusted.

If more than one valid internal certificate authority is available, you can select which internal CA signs the certificate request. There can be multiple valid Internal CAs for Gateways in the following cases:

- There is both an Internal RSA CA for Gateways and an Internal ECDSA CA for Gateways.
- The Internal CA for Gateways is in the process of being renewed and both the previous CA and the new CA are temporarily available.

Make sure that the date, time, and time zone are all set correctly on the Management Server and on the external component that uses the certificate. Certificates are valid for three years starting from the date and time they are created. The validity start and end date and time are written in the certificate and are enforced in the authentication.

Note

The Internal RSA CA for Gateways and the Internal ECDSA CA for Gateways do not support certificate revocation lists. It is not possible to cancel an internally signed certificate before it expires.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.

2) Browse to **Other Elements** > **VPN Certificates**

3) Sign VPN Client certificates.
   a) Select **More actions** > **Sign VPN Client Certificate**.

Tip

You can sign any X.509 certificate requests in this dialog box (not only VPN client certificate requests).

b) If more than one valid internal certificate authority is available, select which internal CA signs the certificate request.



Note

If the Internal CA for Gateways is in the process of being renewed and both the previous CA and the new CA are temporarily available, select the new CA.

c) Browse to the certificate request file on your local workstation or copy and paste the content of the certificate request into the dialog box.

If you copy and paste the certificate request, include the Begin Certificate Request header and the End Certificate Request footer.

d) Click **Sign**.

The certificate is signed and the **Export Certificate** dialog box opens. Click the **Certificate** tab to view the validity information for the certificate.

e) Click the **General** tab, then click **Export** to save the certificate for transfer to the device that needs it.

f) Click **OK**.

4) Sign certificate requests for internal VPN gateways.

a) Click **Gateways**, then expand the VPN Gateway element for which you generated a certificate request.

b) Right-click the certificate request, then select **Sign Internally**.

c) If more than one valid internal certificate authority is available, select which internal CA signs the certificate request.



Note

If the Internal CA for Gateways is in the process of being renewed and both the previous CA and the new CA are temporarily available, select the new CA.

d) Click **Sign**.

Select which internal certificate authority signs each certificate

When there is more than one valid CA, you can select which CA signs each certificate.

The Management Server includes a dedicated Internal RSA CA for Gateways for signing VPN certificates. You can optionally also create an Internal ECDSA CA for Gateways. If you have both an **Internal RSA CA for Gateways** and an **Internal ECDSA CA for Gateways**, only one certificate authority can be selected as the default certificate authority. If you want to sign a certificate with the certificate authority that is not the default CA, you must select which Internal CA for Gateways you want to use.

The Internal RSA CA for Gateways and the Internal ECDSA CA for Gateways are each valid for 10 years. A new Internal RSA CA for Gateways or Internal ECDSA CA for Gateways is automatically created to replace the default certificate authority six months before the expiration date. The certificate authority that is not selected as the default certificate authority is not automatically renewed. You must manually renew the certificate authority.

If the default certificate authority is in the process of being renewed, there is temporarily an extra valid Internal CA for Gateways. In this case, select the new Internal CA for Gateways to sign the certificate.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the **Certificate Properties** dialog box, select **Other** in the **Internally with** field.

2) (Optional) To make sure that you are selecting the correct Internal CA for Gateways, right-click the Internal CA for Gateways, select **Properties**, then check the following information:
   - Validity information in the **Valid from** and **Valid to** fields
   - Status information

3) Select the CA you want to use and click **Select**.

Replacing expired VPN certificates

For security reasons, VPN certificates have an expiration date, after which the certificates must be replaced with new ones.

The VPN certificates issued by the Internal RSA CA for Gateways and the Internal ECDSA CA for Gateways are valid for three years.

Internal certificate authorities also have an expiration date. The system automatically generates a new internal certificate authority and a new internal VPN certificate authority six months before their expiration dates. Each component that uses certificates signed by the internal certificate authority or the internal VPN certificate authority requires a new certificate that is signed by the new internal certificate authority or internal VPN certificate authority.

If certificates signed by the expiring Internal CA for Gateways are used to authenticate VPN client users, you must manually create new certificates for the VPN clients. You must also create new certificates manually for any other external components that have certificates signed by the expiring Internal RSA CA for Gateways or Internal ECDSA CA for Gateways.

Note

When you renew the VPN certificate, Forcepoint VPN Client users receive a notification about the certificate fingerprint change. We recommend that you notify users before you renew the certificate if possible.

Related tasks

Select the default internal certificate authority on page 1257
Create a VPN certificate or certificate request for a VPN Gateway element on page 1257
Import an externally signed VPN gateway certificate on page 1259
Sign external VPN certificate requests with an internal certificate authority on page 1260

Renew an externally signed certificate for a VPN Gateway element

Steps

1) Create a certificate request.
2) Sign the certificate with the external CA.
3) Import the signed certificate.

Renew an internally signed certificate for an external component

Steps

1) Create a certificate request in the external component.
2) Sign the certificate with the internal CA.
3) Export the signed certificate and import it to the external component.
4) If an external gateway trusts the internal VPN CA and the internal VPN CA has been renewed, create a certificate for the external gateway and sign it with the new internal VPN CA.

   You must also set the new VPN CA as a trusted CA in the External Gateway's properties and also in the properties of the VPN Profile element that is used in the VPN configuration.

Renew an internally signed certificate for a VPN Gateway element

New certificates signed by the new default certificate authority are automatically created for VPN Gateway elements. You must manually create and renew any certificates that are not signed by the default certificate authority.

If you have both an Internal RSA CA for Gateways and an Internal ECDSA CA for Gateways, only one certificate authority can be selected as the default certificate authority. If automatic RSA certificate management is activated for an Security Engine, RSA certificates issued by the default certificate authority are renewed automatically as long as the certificate-related files, including the private key stored on the engines, are intact. You must manually create and renew any other certificates that are not signed by the default certificate authority.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.

2) Browse to **Other Elements** > **VPN Certificates** > **Gateway Certificates**.
The certificates are shown with their expiration dates and signer information.

3) Right-click the certificate you want to renew and select **Renew Certificate**.
You are prompted to confirm that you want to renew the certificate.

4) Click **Yes**.
There is a delay while the certificate is renewed, after which you are notified that the certificate was renewed. The certificate is transferred to the engine automatically.

5) Refresh the policy of the Engine to activate the new certificate.
This procedure renews the certificate when the certificate-related information is intact on the engine and on the Management Server. If the certificate has not expired but has other problems, delete the existing certificate element in the SMC Client and create a new one.

Renew an external certificate authority used in VPN configurations

Steps

1) Configure a new certificate authority and make sure that it is a trusted certificate authority in the VPN configurations.

2) Create new certificates for the components involved in the VPN configuration, signed by the new certificate authority.

Export signed VPN gateway certificates or VPN certificate authority certificates

You can export signed gateway certificates, the certificates of the Internal RSA CA for Gateways, and the certificates of the Internal ECDSA CA for Gateways.

In most cases, it is not necessary to export signed VPN gateway certificates or VPN certificate authority certificates, but can be done as needed.

If the SMC has created a new Internal RSA CA for Gateways or Internal ECDSA CA for Gateways to replace an expiring default certificate authority, you must export the certificate of the new default certificate authority. You must import the certificate on external gateways that use certificates signed by the default certificate authority or communicate with gateways that use certificates signed by the default certificate authority. If the external gateway

itself uses a certificate signed by the default certificate authority, you must also create a new certificate for the external gateway.

You must export certificates that are created when an internal certificate authority signs an external certificate request at the time of signing the certificate request. They are not stored for exporting later.

Export a signed VPN gateway certificate

You can export signed VPN gateway certificates for external VPN gateways.

Before you begin

You must have a signed VPN gateway certificate.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.
2) Browse to **Other Elements** > **VPN Certificates** > **Gateway Certificates**.
3) Right-click a certificate and select **Export Certificate**.
4) Browse to the location where you want to save the file on your local workstation and click **Save**.

Export the certificate of an internal CA for gateways

You can export the certificates of the Internal RSA CA for Gateways and the Internal ECDSA CA for Gateways.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.
2) Browse to **Other Elements** > **VPN Certificates** > **VPN Certificate Authorities**.
3) Right-click a **VPN Certificate Authority** element, then select **More actions** > **Export Certificate**.
4) Browse to the location where you want to save the file on your local workstation and click **Save**.

Check when VPN gateway certificates expire

Certificates expire according to the information written in the certificate when it was generated

Before you begin

A signed certificate must be present.

By default, RSA certificates issued by the default certificate authority for VPN Gateway elements are renewed automatically. VPN Gateways never accept expired certificates.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.

2) Browse to **Other Elements** > **VPN Certificates** > **Gateway Certificates**.
   The existing certificates are shown.

3) See the **Expiration Date** column for information about the certificate's expiration date.
   - You can renew internally signed certificates through their right-click menu. You must manually renew certificates if automated RSA certificate management is not active, or if the certificate was not signed by the default certificate authority. VPN client users might be prompted to accept the change of certificate.
   - Elements with no expiration date are certificate requests (Status: To be signed).

Check when VPN certificate authorities expire

The Internal RSA CA for Gateways and the Internal ECDSA CA for Gateways are valid for 10 years.

The Management Server includes a dedicated Internal RSA CA for Gateways for signing VPN certificates. You can optionally also create an Internal ECDSA CA for Gateways. If you have both an Internal RSA CA for Gateways and an Internal ECDSA CA for Gateways, only one certificate authority can be selected as the default certificate authority.

A new Internal RSA CA for Gateways or Internal ECDSA CA for Gateways is automatically created to replace the default certificate authority six months before the expiration date. The certificate authority that is not selected as the default certificate authority is not automatically renewed.

When a new internal VPN CA has been created, the VPN gateways that trust the old VPN CA must be made to trust the new VPN CA. VPN clients that use certificates for user authentication also require new certificates signed by the new VPN CA.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.

2) Browse to **Other Elements** > **VPN Certificates** > **VPN Certificate Authorities**.

3) See the **Expiration Date** column for information about the CA's expiration date.

4) To view detailed information, right-click an Internal RSA CA for Gateways or an Internal ECDSA CA for Gateways, then select **Properties**. Check the following information in the Properties dialog box:
   - **Validity** information in the Valid from and Valid to fields.
   - **Status** information:
     - **Active**: You can use this Internal CA for Gateways to sign certificates.
     - **Renewal Started**: This certificate authority is a new Internal CA for Gateways that the SMC has created automatically. The process of renewing VPN certificates has begun.
     - **Expires Soon**: A new Internal CA for Gateways has been created but some components might still use certificates signed by this Internal CA for Gateways.
     - **Inactive**: This Internal CA for Gateways has expired or no SMC components use a certificate signed by this internal VPN CA.

Chapter 76

Reconfiguring existing VPNs

Contents

- Changing tunnels in a VPN on page 1269
- Add gateways to an existing VPN on page 1270
- Changing gateway IP addresses in an existing VPN on page 1270
- Give VPN access to more hosts in policy-based VPNs on page 1271
- Route all Internet traffic through policy-based VPNs on page 1272
- Redirect traffic between VPN tunnels on page 1273
- Replace pre-shared keys for VPNs on page 1274
- Adjusting gateway settings for Security Engines in existing VPNs on page 1275

You can reconfigure and tune existing VPNs.

Changing tunnels in a VPN

You can add or remove tunnels in a VPN.



Note

Before changing the tunnels that are used in active VPNs, we recommend that you back up the Management Server.

You must add or remove Route-based Tunnels elements manually.

In a policy-based VPN, the gateway topology and the number of active endpoints in each gateway element determine the number of tunnels generated for a VPN. After changing the topology of a policy-based VPN, always check that all new or changed tunnels are valid on the **Tunnels** tab.

- Each central gateway forms a tunnel with each central and satellite gateway in the VPN. No other Gateway<->Gateway tunnels are created. Tunnels are not generated between endpoints that cannot connect to each other. For example, tunnels are not generated between two endpoints if they both have a dynamic IP address.
- Adding a gateway under another gateway instead of directly at the main level in the central gateways list can prevent tunnel generation. This configuration implies that the gateway at the main level forwards connections to the gateways below it in the hierarchy. For the forwarding to work, it must be explicitly configured in the central gateway's Access rules with the **VPN Action** setting in the Action options set to **Forward**.
- Endpoint<->endpoint tunnels are created between all endpoints defined in the properties of any two gateways that form a Gateway<->Gateway tunnel.

Related tasks

Back up system configurations on page 1297
Create Route-based Tunnels elements on page 1211
Define VPN topology for policy-based VPNs on page 1197

Add gateways to an existing VPN

To create connectivity to different gateways, add new gateways to Route-based Tunnels and policy-based VPNs.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Create a gateway element to represent the physical gateway device in VPNs if the element does not exist already.

   VPN Gateway elements are automatically created for Forcepoint Network Security Platform in the Engine/VPN role. The same element can be used in many VPNs.

2) If the VPN uses certificates for authentication, you might need to create a VPN certificate for the gateway.

   The same certificate can be used in many VPNs, providing it fulfills the following criteria:

   - The certificate must match the type of certificate selected for the VPN in the VPN Profile.
   - The certificate must be issued by a certificate authority that the other Gateways trust.

3) Add the gateway to a policy-based VPN or to a Route-based Tunnels element.

   - Edit the **Policy-Based VPN** element and add the gateway on the **Site-to-Site VPN** tab.
   - Edit the **Route-based Tunnels** Tunnel element and select the gateway.

4) Check and adjust the tunnels between the new gateway and the existing gateways.

5) Refresh the policies of all Security Engines that are involved in the tunnels.

Changing gateway IP addresses in an existing VPN

If the IP address of a device that you use as a VPN gateway changes, change the gateway IP addresses in NPNs where the gateways are used.

There are special considerations depending on whether you change the IP address of a VPN Gateway element or an External VPN Gateway element.

For VPN Gateway elements, the IP addresses you have defined for the engine's interfaces determine the VPN endpoint addresses. On Engine Clusters, only CVI addresses are used as VPN endpoints.

Note

If the gateway's identity in the VPN is based on its IP address, you must update the configurations of all gateways in the VPN. You must update the configurations even if the IP address is NATed and not directly used for contact. For VPN Gateway elements, you update the configuration by refreshing the engine's policy after you change the IP addresses. For External VPN Gateways, change the information in the configuration of the gateway device.

- If you change the IP address for a engine interface, the corresponding VPN endpoint IP address also changes automatically. The existing tunnels in the Policy-Based VPN element and the Route-based Tunnels elements are preserved.
- If continuous connectivity is required, define the new address as a second endpoint before you change the IP address. The Multi-Link VPN automatically selects the IP address that works before and after the change.
- If you add or remove interfaces, you might need to select or deselect endpoints manually and then check the tunnel configuration in the Policy-Based VPN element or the Route-based Tunnels elements.



Note

You cannot use the same endpoint in a policy-based VPN and a Route-based Tunnels.

For External VPN Gateways, you always enter the VPN endpoint addresses manually. Change the IP address configured in the SMC Client, then refresh the policies of all affected engines.

Give VPN access to more hosts in policy-based VPNs

If you want to give access to hosts with IP addresses that are not already configured for your policy-based VPN, you must follow several general steps.



Note

In route-based VPNs, it is not necessary to change the VPN configuration to allow access through the VPN for more hosts. Any traffic that is routed to a tunnel interface and allowed by the Access rules automatically uses the Route-based Tunnels.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Make sure that the IP addresses are included in one of the Sites of the correct VPN gateway. If the IP addresses must not be included in other VPNs where the same gateway element is used, add them to a separate Site. Disable the Site in other VPNs.

2) (VPN with external gateways) Add the new IP addresses to the configuration of the external gateway device, so that it routes the traffic through the VPN.

3) Check that the Access rules of all gateways involved specify that this traffic is sent or allowed through the policy-based VPN. If NAT is enabled in the policy-based VPN, also check the NAT rules.

Route all Internet traffic through policy-based VPNs

You can force all traffic from VPN clients or clients in protected networks to be routed through a policy-based VPN.

Before you begin

You must have a working VPN between all gateways.

Routing all traffic from VPN clients or clients in protected networks through a policy-based VPN allows the traffic to be inspected centrally.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Enable NAT for tunneled traffic in the Policy-Based VPN element's properties.
2) Change the mode of the central gateway's sites in this policy-based VPN to **Private** and replace them with a Site element that contains the **Any Network** element. Disable the **Any Network Site** in other VPNs.
3) Add Access rules to redirect the traffic from VPN clients or clients in protected networks to the central gateway as necessary.
4) Redirect the traffic from external components to the central gateway as necessary.
   For VPN Gateway elements, add an Access rule that sends the allowed traffic to the VPN.
5) (VPN Clients only) Configure the Virtual Adapter.

Related concepts

Using NAT for policy-based VPN traffic on page 1206

Related tasks

Edit a Policy-based VPN on page 1197
Create Access rules for policy-based VPN traffic on page 1203

Redirect traffic between VPN tunnels

In policy-based VPNs, you can redirect traffic from one VPN tunnel to another VPN tunnel through a hub gateway.

Before you begin

You must have a policy-based VPN configured. If VPN client traffic is forwarded, you must configure virtual IP addressing for VPN clients.

Redirecting traffic between VPN tunnels is especially useful for VPN client users who need access to resources at different locations. When you redirect traffic between VPN tunnels, users do not have to separately connect to different gateways.

Steps

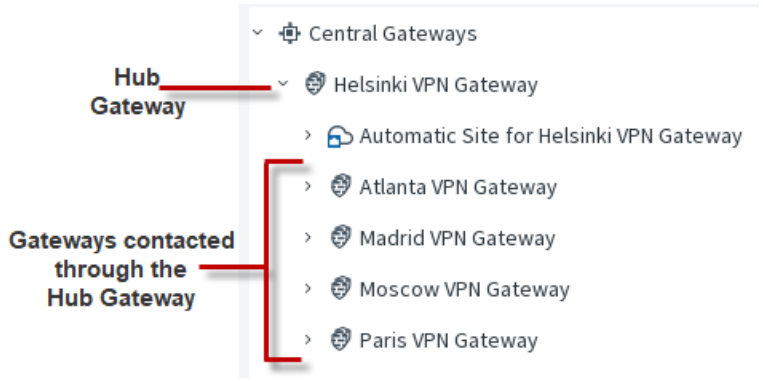
For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.

2) Browse to **Policy-Based VPNs**.

3) Right-click the **Policy-Based VPN** element, then select **Edit <element name>**.

4) On the **Site-to-Site VPN** tab of the **Policy-Based VPN** editing view, place the forwarding hub gateway at the top level of the **Central Gateways** list.



5) Place the gateways that are contacted through the hub under the hub gateway.

Note

Duplicate tunnels are not allowed. There must not be site-to-site connections between the hub and the other gateways in other active VPNs.

6) Add a Site element that contains the IP addresses behind the spoke gateways under the hub gateway.
   a) Set the **Mode** of the Site element to **Hub**.

b) Disable the Site element in any other VPNs where it is used.

The protected IP addresses are behind the spoke gateways.

7) To forward VPN client traffic, add a Site element that contains the virtual IP address space used for the VPN clients under the hub gateway.

8) Add Access rules that forward the traffic between tunnels.

9) Refresh the policies of all engines involved in the VPN, starting from the engine that acts as the hub gateway.

Optionally, all traffic (including Internet traffic) can be routed through the hub gateway.

Related tasks

Create Access rules to forward traffic on hub gateways on page 1205
Define virtual IP addresses for VPN clients on page 1281

Replace pre-shared keys for VPNs

As a security precaution, we recommend that you periodically change the pre-shared key.

Before you begin

Pre-shared key authentication must be selected in the VPN Profile and allowed in the Gateway Profiles

You can renew or generate pre-shared keys automatically or manually.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.

2) For policy-based VPNs, right-click the **Policy-Based VPN** element, select **Edit <element type>**, then follow these steps.

a) Click the **Tunnels** tab.

b) To automatically generate pre-shared keys for multiple tunnels, select the tunnels.

c) Right-click the selected tunnels, then select **Delete Pre-Shared Key**.

d) Right-click the selected tunnels again, then select **Generate Missing Pre-Shared Key**.

A new pre-shared key is generated for each tunnel.

e) To manually enter the key for a single tunnel, double-click the **Key** column, then enter or paste the key.

f) To transfer the key for a tunnel to external components, double-click the **Key** column, then copy the key, or click **Export**.



Note

Make sure that outsiders cannot obtain the key while you transfer it to other devices. The key must remain secret to be an effective security measure.

3) For route-based VPNs, right-click a **Route-Based Tunnel** element, select **Properties**, then follow these steps.

a) Next to **Pre-Shared Key**, select **Edit**.

- To automatically generate a key, click **Generate**.
- To manually enter the key, enter or paste the key.
- To transfer the key to external components, copy the key, or click **Export**.

b) Click **OK**.

4) Click **Save**.

Adjusting gateway settings for Security Engines in existing VPNs

The Gateway Settings element defines performance-related VPN options for the Security Engines.

The gateway settings are used internally and there is no need to match them exactly with settings of other gateways in VPNs.

Gateway setting	
MOBIKE Return Routablility Check	MOBIKE (mobile IKE) return routablility checks (RRC) can be used with IKEv2 to verify the validity of VPN client or gateway IP addresses if the IP address changes in the middle of an open VPN connection. The IP address is updated in the negotiated SAs when the new IP address has been verified. If the new IP address cannot be verified, the VPN connection is closed. By default, no return routablility checks are done.
Negotiation Retry	If a negotiation for a VPN does not complete successfully, the VPN establishment is retried according to settings in the Negotiation Retry section in Gateway Settings properties. The default settings are the recommended values. VPN establishment might fail because you have frequent intermittent problems with network connectivity or because your network connection is too slow. In these cases, increasing Negotiation Retry values might be a work-around solution for getting the VPN to establish.

Gateway setting	
Certificate Cache	The **CRL Validity** setting in the Certificate Cache section in Gateway Settings properties has an effect only if you use certificates to authenticate VPN gateways in IKE negotiations. The default setting is the recommended value. We do not recommend adjusting this setting.

By default, all Security Engines use the **Gateway Default Settings** Gateway Settings element. To customize the gateway settings, define a custom Gateway Settings element.

Define a custom Gateway Settings element

The Gateway Settings element defines performance-related VPN options for Security Engines.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.
2) Browse to **Other Elements** > **Profiles** > **Gateway Settings**.
3) Right-click **Gateway Settings** branch in the element, then select **New Gateway Settings**.
4) Configure the settings, then click **OK**.

Assign gateway settings to an Security Engine

To use custom gateway settings, select the custom Gateway Settings element.

Before you begin

Define a custom Gateway Settings element.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.
2) Right-click a Engine, then select **Edit <element type>**.
3) Browse to **VPN** > **Advanced**.
4) Select an Gateway Settings element from the Gateway Settings list.
   To browse for an element, click **Select**. To create an element, select **New**.

5) Click **Save and Refresh** to transfer the changes.

Chapter 77

VPN client settings

Contents

- VPN client settings and how they work on page 1279
- Defining IP addresses for VPN clients on page 1279

Forcepoint VPN Client does not have controls for many settings that are needed for establishing a VPN. These settings are defined in the SMC. Forcepoint VPN Client downloads the settings from the gateways it connects to. VPN clients are only supported in policy-based VPNs.

VPN client settings and how they work

Forcepoint VPN Client settings are configured centrally in the SMC. The settings are automatically updated to the Forcepoint VPN Client from the engines when the clients connect.

The following settings are transferred from the gateway to the client:

- Routing information (VPN Site definitions). Generally, if an IP address that the client wants to contact is included in the Site definition, the traffic is routed into the VPN.
- Authentication settings
- Encryption settings
- Information about the gateway's endpoints
- Settings for NAT traversal methods allowed
- Settings for local security checks on the client computer
- Secondary IPsec VPN gateways to contact in case there is a disruption at the IPsec VPN gateway end

VPN client settings have the following limitations:

- When the Forcepoint VPN Client is first installed, it has no configuration. Either the user or the administrator must add the basic information about gateways, such as the IP address to use for connecting.
- There are version-specific dependencies between the Forcepoint VPN Client and Engine/VPN role software. See the Release Notes of the Forcepoint VPN Client version you intend to use for information about compatibility with your Engine/VPN role gateway's software version.
- The SMC does not create configurations for third-party VPN clients. You must create the configuration through the controls and tools of the third-party VPN client product.

Defining IP addresses for VPN clients

There are two different methods to define the IP addresses that VPN clients use in the internal network.

You must always configure one of the following methods for the mobile VPN to be valid:

1) You can use NAT to translate the IP addresses in communications. Using NAT gives the VPN clients an 'internal' IP address in the internal network without the need for a DHCP server. This method is called a *NAT Pool*.
   - This method is not recommended for the Forcepoint VPN Client. It does not allow the clients to make queries to internal DNS servers without manual configuration.
   - NAT rules are not applied to communications from clients that receive their address through the NAT Pool feature. The NAT Pool translation is applied before the NAT rules.
   - The NAT Pool method does not require any other client-side features.

2) (Recommended for the Forcepoint VPN Client) You can use a DHCP server to assign a virtual IP address that VPN clients use in communications through the VPN tunnel. The IP address is attached to a *Virtual Adapter*. Using this method provides the following benefits over the NAT Pool:
   - Centrally configure the DNS settings for VPN clients when connected (using the DHCP server).
   - Control how the IP address each VPN client is assigned (depending on the DHCP server).
   - Forward mobile VPN traffic to a site-to-site VPN or route the Internet traffic from the client computer through the gateway for inspection.
   - Open new connections from the internal network to the VPN client computers through the VPN.

To use the Virtual Adapter, the VPN client software must support this feature. Not all third-party VPN clients have a Virtual Adapter feature. The Virtual Adapter is required when there is a need to open connections from the internal network to the VPN client. Activating both the NAT Pool and the Virtual Adapter is technically possible. However, the NAT Pool address translation is applied to all VPN client traffic when activated, including connections from hosts that use a Virtual Adapter.



Note

For a detailed technical discussion on using a virtual IP address, see RFC 3456.

Translate VPN client IP addresses using NAT Pool

The NAT pool defines a range of IP addresses that the engine can use to translate the source address of connections from VPN clients.

The NAT pool translates the addresses in the same way as NAT rules do. Connections that use the NAT Pool must not match any NAT rules.



Note

Make sure that NAT is enabled for this VPN. The **Apply NAT to traffic that uses this VPN** option in the properties of the VPN element must be selected. Otherwise, the NAT pool options have no effect.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click the engine element, then select **Edit <element type>**.

3) Browse to **VPN** > **Advanced**.

4) Select the **Translate IP Addresses Using NAT Pool** option.



Note

If the NAT Pool is active, it is also used for translating connections from VPN clients that have a virtual IP address. It is not possible to exclude hosts with a virtual IP address from being subject to the NAT Pool address translation.

5) In the **IP Address Range** and **Port Range** fields, enter the IP addresses and ports you want to use for translating VPN client traffic.



CAUTION

Make sure the addresses that you define here do not overlap with addresses that are in use in your networks. Also, the addresses must not overlap with any translated address space in your NAT rules.

Define virtual IP addresses for VPN clients

You can use a Virtual Adapter to assign the VPN client an IP address in the VPN, independent of the address the VPN client computer uses in its local network.

Before you begin

To use virtual IP addresses for VPN clients:

- You must use an external DHCP server to assign the IP addresses.
- The users must use a VPN client that has a Virtual Adapter feature. The Forcepoint VPN Client always has this feature installed and active.

The virtual IP address is only used in communications through the VPN tunnels. The VPN gateway gets the IP address and network settings of the Forcepoint VPN Client from the an external DHCP server and forwards the information to the Forcepoint VPN Client. For one-way access without DNS resolving, the VPN gateway can alternatively be set up to apply NAT to translate the Forcepoint VPN Client connections. This method is meant for testing purposes.

The VPN gateway specifies the destination IP addresses for traffic that the Forcepoint VPN Client sends into the VPN tunnel. The IP addresses are configured as Site elements for each gateway in the SMC Client. When the Sites contain specific internal networks, the Forcepoint VPN Client receives a configuration for *split tunneling*. Split tunneling means that only the specified portion of traffic uses the VPN tunnel, and other connections use the local network as usual.

Most DHCP servers allow a configuration in which a particular client computer is always assigned a particular IP address. For example, the DHCP server might assign the IP address based on the MAC address if VPN clients have fixed MAC addresses for their Virtual Adapters. By default, when the Forcepoint VPN Client virtual adapter requests an IP address, it uses the MAC address of the physical interface used in the VPN connection.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click the engine, then select **Edit <element type>**.

3) Browse to **VPN** > **VPN Client**.

4) From the **DHCP Mode** drop-down list, select how DHCP requests from VPN clients are sent.

> **Note**
>
> If **SSL VPN** or **Both IPsec & SSL VPN** is selected from the **VPN Type** drop-down list, only **Direct** and **DHCP Relay** are shown.

5) From the **Interface** or **Interface for DHCP Relay** drop-down list, select the source address for the DHCP packets when querying the DHCP server (the interface toward the DHCP server).

6) Click **Add**, then select the DHCP server element that assigns IP addresses for the VPN clients.

7) (Optional) From the **Add Information** drop-down list, select what VPN Client user information is added to the Remote ID option field in the DHCP Request packets.

- **Add User information** — VPN Client user information (in the form user@domain) is automatically added to the Remote ID option field in the DHCP Request packets.
- **Add Group information** — VPN Client user information (in the form group@domain) is automatically added to the Remote ID option field in the DHCP Request packets.

Your DHCP server must support the **DHCP Relay Agent Information** option to use this information. Depending on your DHCP server configuration, this information can be used as a basis for IP address selection.

8) (Optional) Select **Restrict Virtual Address Ranges**, then enter the IP address range in the field on the right.

With this option, you can restrict the VPN clients' addresses to a set range, even if the DHCP server tries to assign another IP address. If an incorrect address is assigned, the user might not be able to access resources. These address ranges must not overlap with the NAT Pool.

> **Note**
>
> If the NAT Pool is active, it is also used for translating connections from VPN clients that have a virtual IP address. It is not possible to exclude hosts with a virtual IP address from being subject to the NAT Pool address translation.

9) (Optional) Configure the Engine to act as a proxy for the VPN client's ARP requests.

a) Select **Proxy ARP**.

b) In the field on the right, enter the IP address range for proxy ARP.



Note

The **Proxy ARP** option might be required for a working VPN depending on your network configuration.

10) Click **Save and Refresh**.

Chapter 78

Configuring the SSL VPN Portal

Contents

- Getting started with the SSL VPN Portal on page 1285
- Make services available in the SSL VPN Portal on page 1286
- Allow access to services using the SSL VPN Portal on page 1286
- Define an SSL VPN Portal element on page 1287
- Enable the SSL VPN Portal for an Security Engine on page 1288

The SSL VPN Portal uses secure sockets layer (SSL) encryption to allow authenticated users to establish secure connections to internal HTTP and HTTPS services through a standard web browser or through a client application that allows direct network access.

Getting started with the SSL VPN Portal

The SSL VPN Portal provides secure browser-based access to services in the protected network.

The SSL VPN Portal is an integrated feature of Forcepoint Network Security Platform. It provides remote access to applications and information in the protected network from standard web browsers. End users must authenticate to access the SSL VPN Portal webpage. You can configure single sign-on (SSO), to allow users to access different services under the same Domain without logging on to each service separately. The SSL VPN Portal proxies end-user connections to HTTP-based services in the protected network. The end user is never directly connected to the back-end services.

The configuration consists of the following general steps:

1) Define SSL VPN Portal Service elements to make services available in the SSL VPN Portal.

2) Add rules to the SSL VPN Portal Policy to allow access to services using the SSL VPN Portal.

3) Create an SSL VPN Portal element to define the settings for connecting to the SSL VPN Portal and the look and feel of the SSL VPN Portal.

4) Enable the SSL VPN Portal for each Security Engine that provides SSL VPN Portal access.

Make services available in the SSL VPN Portal

To make services in the protected network available in the SSL VPN Portal, define SSL VPN Portal Service elements.

SSL VPN Portal Service elements map external URLs to HTTP-based services in the protected network. SSL VPN Portal Service elements contain settings that define how the internal URLs of the HTTP-based services are translated to external URLs. URL translation makes sure that all traffic to registered web resource hosts is routed through the SSL VPN Portal. End users can access the SSL VPN Portal Services through the SSL VPN Portal, or directly through web browser bookmarks.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.
2) Browse to **SSL VPN Portal** > **SSL VPN Portal Services**.
3) Right-click **SSL VPN Portal Services**, then select **New SSL VPN Portal Service**.
4) Configure the settings, then click **OK**.

Next steps

You are now ready to define which users are allowed to access the services.

Allow access to services using the SSL VPN Portal

The SSL VPN Portal Policy defines which services are available in the SSL VPN Portal and which users can access the services.

Before you begin

You must have one or more SSL VPN Portal Service elements.

The SSL VPN Portal Policy contains rules that define which users can use each SSL VPN Portal Service, and the authentication requirements for accessing the SSL VPN Portal Services.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.

2) Create an SSL VPN Portal Policy.
   a) Expand the **SSL VPN Portal** branch.
   b) Right-click **SSL VPN Portal Policies** and select **New SSL VPN Portal Policy**.
   c) Configure the settings, then click **OK**.

   The SSL VPN Portal Policy opens for editing in a new tab.

3) Add rules in one of the following ways:
   - Right-click the last row of an empty policy and select **Add Rule**.
   - Right-click the **ID** cell of an existing rule and select **Add Rule Before** or **Add Rule After**.

4) Drag and drop one or more SSL VPN Portal Service elements from the **Resources** pane to the **SSL VPN Portal Service** cell.

5) Drag and drop one or more User or User Group elements from the **Resources** pane to the **Authentication** cell.

6) Save the SSL VPN Portal Policy.

Next steps

You are now ready to select the SSL VPN Portal Policy for an SSL VPN Portal element.

Define an SSL VPN Portal element

Create an SSL VPN Portal element to define the settings for connecting to the SSL VPN Portal and the look and feel of the SSL VPN Portal.

Before you begin

You must have an SSL VPN Portal Policy.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.

2) Browse to **SSL VPN Portal** > **SSL VPN Portals**.

3) Right-click **SSL VPN Portals**, then select **New SSL VPN Portal**.

4) Configure the settings, then click **OK**.

5) Click **OK**.

Next steps

Enable the SSL VPN Portal for each Security Engine that provides SSL VPN Portal access.

Enable the SSL VPN Portal for an Security Engine

In the Engine Editor, enable the SSL VPN Portal for each Security Engine that provides SSL VPN Portal access.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Engine Configuration**.

2) Right-click an engine element, then select **Edit <element type>**.

3) Browse to **VPN** > **SSL VPN Portal**.
The **SSL VPN Properties** pane opens on the right.

4) Click **Select** to select the SSL VPN Portal you want to use.

5) Click **Save and Refresh** to transfer the new configuration to the engines.

Part XII

Maintenance and upgrades

Contents

- Configuration of automatic updates and upgrades on page 1291
- Backing up and restoring system configurations on page 1295
- Managing log data on page 1307
- Managing and scheduling Tasks on page 1319
- Managing licenses on page 1331
- Upgrading the SMC on page 1345
- Upgrading Security Engines on page 1351
- Manual dynamic updates on page 1357
- SMC Appliance maintenance on page 1361

Maintenance includes procedures that you do not typically need to do frequently.

Chapter 79

Configuration of automatic updates and upgrades

Contents

- Getting started with automatic updates and upgrades on page 1291
- Configure automatic updates and upgrades on page 1292

You can configure the Management Server to automatically download and install dynamic update packages, remote upgrades for engines, and licenses.

Getting started with automatic updates and upgrades

Automatic updates are available for dynamic update packages, remote upgrades for engines, and licenses.

Before dynamic updates and engine upgrades are downloaded and activated or installed, they are verified. The Management Server and the Security Engines check the digital signature of each dynamic update or engine upgrade using a valid Trusted Update Certificate. Only updates and upgrades with a valid signature can be downloaded and installed. If there is a verification failure, the administrator receives an error. Verification failure can result from an out-of-date SMC version or an invalid or missing signature.

What automatic updates and engine upgrades do

The Management Server can automatically perform the following tasks:

- Check for new dynamic update packages and automatically download and install them according to your selection.
- Check for new engine upgrade packages. Engine upgrades can also be automatically downloaded, but they must always be installed manually.
- Upgrade the licenses.

When automatic updates and engine upgrades are active, you can also view information regarding the maintenance contract and support level of your licenses in the SMC Client.

Limitations

- (Multiple Management Servers only) Dynamic update packages are downloaded and activated on the active Management Server (the Management Server that controls all Domains). The settings for automatic updates and upgrades are configured in the properties of the active Management Server.
- There are no automatic updates for the SMC software.

- New engine software versions might require an upgraded version of the SMC. Check the Release Notes for compatibility information before upgrading the engines.
- Upgrades and updates (both automatic and manual) require an active maintenance or support contract.
- If you select the **Notify When Updates Become Available** setting, you must manually download the updates and engine upgrades.

Related concepts
Checking maintenance contract information on page 250

Related tasks
Configure automatic updates and upgrades on page 1292

Configure automatic updates and upgrades

There are several options for handling automatic updates and engine upgrades.

Before you begin

- You must have a valid maintenance or support contract.
- Automatic dynamic updates and engine upgrades require the Management Server to be able to connect to the license server at https://smc-pool.stonesoft.com and to the dynamic update service at https://autoupdate.ngfw.forcepoint.com using HTTPS on port 443.

The Management Server can periodically check for new dynamic update packages, engine upgrades, and licenses. This feature is active by default. In an environment with multiple Management Servers, automatic updates and upgrades must be enabled on the active Management Server (the Management Server that controls all Domains).

Update Service elements define sets of URLs for automatic dynamic updates and engine upgrades. You can optionally change which Update Service element is used for automatic dynamic updates and engine upgrades.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Settings** > **Global System Properties**.

2) On the **Updates** tab, select **Allow Sending License and Installation Telemetry Data to Forcepoint Servers**.
   Selecting this option allows you to select settings for dynamic updates and for engine and license upgrades.

3) Configure the **Dynamic Updates** settings.



Note

Because update packages can change system elements, the policies might require editing after update activation.

4) Select one of the **Remote Upgrades for Engines** settings.

5) (Optional) Select **Generate and Install New Licenses Automatically** to automatically regenerate and install the licenses required for upgrading SMC components to a major new release.

6) (Optional) Select the **Update Check Interval** to define how often the SMC checks for new updates.

7) Click **OK**.

Related concepts

Checking maintenance contract information on page 250
Getting started with licenses on page 1331
Getting started with upgrading the SMC on page 1345
Getting started with upgrading Security Engines on page 1351
Getting started with manual dynamic updates on page 1357

Related tasks

Schedule Tasks on page 1327

Chapter 80

Backing up and restoring system configurations

Contents

- Backups and how they work on page 1295
- Back up system configurations on page 1297
- Copy backup files to a storage location on page 1298
- Import backup files into the SMC on page 1298
- Restoring backups on page 1299
- Restore system configurations after a hardware failure on page 1302
- Managing SMC Appliance backups on page 1303

Backups contain the necessary configuration information to restore the SMC to the state it was in when the backup was taken.

Backups and how they work

Backups are needed to recover from the loss of the system configurations, for example, due to hardware failure. Backups also allow you to relocate the SMC servers onto different hardware.

The Management Server is the only component that contains usable, complete configuration information for any individual engine component. The engines contain a working copy of the configuration details that allows them to carry out traffic inspection independently. It is not possible to extract this information from the engines if the Management Server is lost. For this reason, regular Management Server backups are essential and must be stored in a safe storage location outside of the computer where the SMC servers are installed.

Always take the backups using the proprietary backup tools in the SMC Client, on the Management Server command line, or on the SMC Appliance command line. Third-party backup applications that back up the host system might not produce usable backups of your SMC servers, especially if the SMC servers are running when you take the backup.

Restoring backups allows you to restore the configurations to the state they were when the backup was taken, even if you restore the backup in a different SMC.

Different types of backups contain different information:

- The Management Server backup contains the policies, elements, and other configuration details for all Security Engines that they manage. The Management Server backup also contains the configuration information of the Web Access Server and of the Management Server itself.
- The Log Server backup contains the Log Server's local configuration and optionally the logs.
- On the SMC Appliance, the Management Server and Log Server backups also contain the SMC Appliance configuration files.

Note

If your configuration includes TLS Credentials and Client Protection Certificate Authority elements, the private keys and certificates used in these elements are included as plain text in the Management Server backup. Use the encryption option for the backups when the configuration contains these kinds of elements.

The backup files are compressed to .zip files (unencrypted backups) or .enc files (encrypted backups) and they can also be decompressed manually if needed. If necessary, the backups are split into several files to fit the maximum file size. Each backup has its own subdirectory.

The following limitations apply:

- In FIPS-compatible operating mode, you can only restore backups that were created for an SMC in FIPS-compatible operating mode.
- You cannot restore backups that were created in an SMC in FIPS-compatible operating mode on an SMC that is not in FIPS-compatible operating mode.
- The private keys of engine certificates are stored locally on the engines and are not backed up.
- If you restore an SMC Appliance backup onto third-party hardware, SMC Appliance configuration information is ignored. Only the Management Server and Log Server backups are applied.

Related concepts

TLS inspection and how it works on page 1063

Backup configuration overview

Managing backups involves several main steps.

1) Back up the Management Servers and Log Servers regularly or schedule backup tasks to run at regular intervals.

2) Store the backup files in a safe location.
When you create backup files, you can save the files on the Management Server or on your local workstation. We recommend copying the backup files to a safe location, such as removable media or another host, for long-term storage.

3) When necessary, restore a backup.

Back up system configurations

Backups are needed to recover from the loss of the system configurations, for example, due to hardware failure. A backup also allows you to relocate the SMC servers onto different hardware.

Before you begin

To back up a Management Server, there must be enough free disk space on the server. Twice the size of the management database is required. If there is not enough available disk space, the backup process does not start.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click the Management Server or Log Server you want to back up, then select **Backup**.
2) (Optional) To back up other servers, select the servers from the list on the left, then click **Add**.
3) (Optional) To encrypt the backup, select **Encrypted**, then enter and confirm a password.
   We recommend this option if the configuration contains TLS Credentials and Client Protection Certificate Authority elements.
4) (Optional) If you are creating a backup of Log Servers and you want to back up the log files, select **Back up Log Files**.
5) Click **OK**.
   The backup starts and the progress is shown on a new tab.

Next steps

Copy the backup files to a storage location.

Related concepts

Log data management and how it works on page 1307

Related tasks

Create Backup Tasks on page 1323
Schedule Tasks on page 1327

Copy backup files to a storage location

Copying backup files to a storage location ensures that the backup files are available if the data on the server on which they were created is lost.

When you create backup files, you can save the files on the Management Server or on your local workstation. We recommend copying the backup files to a safe location, such as removable media or another host, for long-term storage.

When you save backup files on the Management Server, the backup files are saved in the `<installation directory>/backups/` directory of the Management Server on which they were created.

Note

If you installed the Management Server in the `C:\Program Files\Forcepoint\SMC` directory in Windows, some program data might be stored in the `C:\ProgramData\Forcepoint\SMC` directory.

On the SMC Appliance, backup files can be automatically stored in a remote location through a CIFS share.

Steps

1) On the Management Server on which the backup files were created or on your local workstation, copy the backup files to a safe storage location.

   Note

   Handle the backup files securely. They contain all configuration information for the system.

Related tasks

Mount a CIFS share on page 1304

Import backup files into the SMC

To restore a backup of the Management Server, Log Server, or SMC Appliance, import the backup file into the SMC.

When you import backup files, the backup files are added to the `<installation directory>/backups/` directory of the Management Server, Log Server, or SMC Appliance.

Note

Handle the backup files securely. They contain all configuration information for the system.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.

2) Browse to **Backups** > **Management Server** or **Backups** > **Log Server**.

3) Right-click **Management Server** or **Log Server**, then select **New** > **Import Backup**.

4) Browse to the backup file, then click **Open**.
The backup file is imported into the SMC.

5) When the import finishes, click **OK**.

Next steps

You can now restore the backup on a Management Server, Log Server, or the SMC Appliance.

Restoring backups

Restoring backups allows you to recover from the loss of the system configurations, or to relocate the SMC servers onto different hardware.

You can restore backups that were created in one operating system to an installation running on another operating system. You can also restore Management Server and Log Server backups that were created on a non-appliance installation to an SMC Appliance.



Note

In FIPS-compatible operating mode, you can only restore backups that were created for an SMC in FIPS-compatible operating mode. You cannot restore backups that were created in an SMC in FIPS-compatible operating mode on an SMC that is not in FIPS-compatible operating mode.

You can restore backups from the previous major version of the SMC in the current major version of the SMC. You might not be able to restore Backups taken from older versions. Generally, you can restore backups between versions that support direct upgrades between the versions. See the upgrade instructions in the Release Notes. If an intermediate upgrade is required between your current version and the newest version, upgrade the existing installation to the intermediate version to create a working backup.

When you restore a backup, the backup restoration process checks that there is enough disk space on the destination drive. Twice the size of the backup file is required. If there is not enough available disk space, the restoration fails.

Restore Management Server backups

Restoring Management Server backups allows you to recover from the loss of the system configurations, or to relocate the SMC servers onto different hardware.

Before you begin

Import the backup files into the SMC or copy the backup files to the <installation directory>/backups/ directory.

Note

Handle the backup files securely. They contain all configuration information for the system.

Note

If the restore operation fails, the original configuration remains unchanged.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) To restore a backup on a Management Server installed in Windows or Linux, stop the Management Server service through the operating system's service management features or using the command-line script, then start the backup restoration script.
   - In Windows, run <installation directory>/bin/sgRestoreMgtBackup.bat
   - In Linux, run <installation directory>/bin/sgRestoreMgtBackup.sh

2) To restore a backup on the SMC Appliance, run the following command to stop the Management Server service and start the backup restoration script:

   ```
   sudo smca-restore
   ```

   Note

   To exclude SMC Appliance configuration files from being restored, use the `-nosmca` option.

   To restore only the SMC Appliance configuration files, use the -smcaonly option.

3) Select the backup file to be restored.

   The default backup file names have the following structure: sgm_vVERSION.[BUILD]_YYYYMMDD_HHMMSS[comment].

4) Type `y`, then press **Enter** to confirm the restoration.

5) If the backup is encrypted, enter the password.

6) If you are restoring the backup on a system that uses a different IP address, change the IP address of the Management Server.

7) If components have certificates from a different certificate authority (CA) than the one contained in the backup, regenerate the certificates.
   The backup contains the internal CAs. If components have certificates from a different CA than the one contained in the backup, the certificates are not accepted as valid after restoring the backup.

8) In Windows or Linux, start the Management Server.
   The SMC Appliance automatically starts the Management Server after the backup has been restored.

9) (SMC Appliance only) Restart the SMC Appliance if you are prompted to do so.

Related concepts
Renewing certificates on page 162

Related tasks
Change the Management Server IP address on page 491

Restore Log Server backups

The Log Server backup contains the Log Server's local configuration and optionally the logs.

Before you begin

Import the backup files into the SMC or copy the backup files to the <installation directory>/backups/ directory.

Note
Handle the backup files securely. They contain all configuration information for the system.

For Log Servers installed in Windows or Linux, if it is not possible to transfer the logs through a backup, you can copy log files to the Log Server through the operating system.

Note
If the restore operation fails, the original configuration remains unchanged.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) To restore a backup on a Log Server installed in Windows or Linux, stop the Management Server service through the operating system's service management features or using the command-line script, then start the backup restoration script.
   - In Windows, run <installation directory>/bin/sgRestoreLogBackup.bat
   - In Linux, run <installation directory>/bin/sgRestoreLogBackup.sh

2) To restore a backup on the SMC Appliance, run the following command to stop the Log Server service and start the backup restoration script:

```
sudo smca-restore
```



Note

To exclude SMC Appliance configuration files from being restored, use the `-nosmca` option.

To restore only the SMC Appliance configuration files, use the -smcaonly option.

3) Select the backup file to be restored.

The default backup file names have the following structure: sgl_vVERSION.[BUILD]_YYYYMMDD_HHMMSS[comment].

4) Type `y`, then press **Enter** to confirm the restoration.

5) If the backup is encrypted, enter the password.

6) If you are restoring the backup on a system that uses a different IP address, change the IP address of the Log Server.

7) In Windows or Linux, start the Log Server.

The SMC Appliance automatically starts the Log Server after the backup has been restored.

8) (SMC Appliance only) Restart the SMC Appliance if you are prompted to do so.

Related tasks

Change the Log Server IP address on page 493

Restore system configurations after a hardware failure

Restoring backups on replacement hardware restores the SMC configurations to the state they were in when you took the backup.

Steps

1) Restore the Management Server configurations on replacement hardware.

   a) Install the Management Server software.

   The same exact version is not required for recovery, but all SMC components must run the same version to work together.

   b) Restore the Management Server backup.

2) Restore Log Server configurations on replacement hardware.
   a) Install the Log Server software, if not installed together with the Management Server software.
      The same exact version is not required for recovery, but all SMC components must run the same version to work together.
   b) Restore the Log Server backup.
3) Restore engine configurations on replacement hardware.
   a) Generate an initial configuration for the engine in the SMC.
   b) Add the hardware to the network and configure it in the same way as a new installation.
   c) When contact with the Management Server is established, install the policy.
      The full working configuration is transferred to the engine.



Note

In some cases, the IPsec VPN certificate information can be lost and the policy installation fails. If VPN certificate information is lost, delete the old VPN certificates in the SMC Client and create new VPN certificates for the engine. When you use the same CA and certificate details, the other components accept the new certificates. Policy installation is also possible if you disable the invalid configurations (for example, by disabling all VPN-specific Access rules in the policy).

Related concepts
Connect Security Engines to the SMC on page 631

Managing SMC Appliance backups

You can create and manage Management Server and Log Server backups for SMC Appliance using the SMC Client. You can also use the smca-backup script on the command line of the SMC Appliance.

When you create a Management Server or Log Server backup for the SMC Appliance using the SMC Client, SMC Appliance configuration files are included in a directory in the backup .zip file.

There are two ways to restore backups:

- You can use the smca-restore command to restore Management Server and Log Server backups, and SMC Appliance configuration files on the SMC Appliance.
- You can use the sgRestoreMgtBackup and sgRestoreLogBackup commands to restore Management Server and Log Server backups on another platform. When you restore a backup that includes SMC Appliance configuration information on another platform, the SMC Appliance configuration information is ignored.

If you are restoring a backup from a different version of the SMC Appliance, read the version-specific Release Notes before restoring the backup.

To manage backups on the SMC Appliance command line, you must have SMC Appliance Superuser administrator permissions. The backup commands must be run with elevated permissions using sudo. A list of available sudo commands can be found by running `sudo -l` at the command line.

The SMC Appliance has a dedicated partition for storing logs and backups. There are several files included in a full backup of the system.

- Log Server backup
- Management Server backup
- SMC Appliance OS configuration backup

Note

If password protection is enabled for backups, the files are saved with AES 128-bit encryption.

Backup files can be automatically stored in a remote location through a CIFS share.

Related reference
Forcepoint Security Management Center commands on page 1429

Mount a CIFS share

Setting up a CIFS share allows the SMC Appliance to export backups and other files to a remote storage location. A CIFS share is persistent, even if the appliance is restarted.

Before you begin

If there is a engine between the SMC Appliance and the CIFS share location, make sure that CIFS traffic is allowed.

Note

CIFS shares are not supported in FIPS mode.

Steps

1) Log on to the appliance from the command line.

2) Enter the CIFS command.

```
sudo smca-cifs add -n <name> -s //<server>/<share> -u <username> -p <password> -d <domain>
```

Where *<name>* is the locally used name, *<server>/<share>* is the server or IP address and share name on that server, and *<domain>* is the domain of the share.

3) Press **Enter**.

Result

The CIFS share is mounted in the /mnt/cifs/<name> directory. The CIFS credentials are stored in the /etc/smca/cifs/<name> file.

Create an SMC Appliance backup on the command line

You can create a backup of the SMC Appliance configuration from the command line.

Steps

1) Use the command line to log on to the appliance.
2) Run the backup command.

   ```
   sudo smca-backup
   ```

   You can append various options to the command to customize the backup process.
3) Press **Enter**.

Move backups to remote storage

Regularly storing appliance backups in a remote location allows for a complete appliance recovery in the event of an unexpected failure.

Before you begin

Mount a CIFS share.

Steps

1) Log on to the appliance from the command line.
2) (Optional) Create a backup of the appliance.
3) Enter the rsync command.

   ```
   sudo smca-rsync add -i <source directory> -o <destination directory>
   ```

   <source directory> is where the backups are stored when they are created and *<destination directory>* is the remote location to store the backups. If you do not include the `-i`, the appliance defaults to the SMC backup directory file.
4) Press **Enter**.

Result

The rsync command copies the backup files to the destination hourly. The command also runs a check to verify that there are no duplicate files in the destination directory.

Restore SMC Appliance backups

You can use the smca-restore command to restore Management Server backups, Log Server backups, and SMC Appliance configuration files on the SMC Appliance.

Before you begin

If you are restoring a backup from a different version of the SMC Appliance, read the version-specific Release Notes before restoring the backup.

Steps

1) Log on to the SMC Appliance command line.

2) Run the restore command.

```
sudo smca-restore [options]
```

You can append various options to the command to customize the backup process.

- To specify the backup file to restore, use the -backup option.
- To restore the Management Server or Log Server backup without restoring the SMC Appliance configuration, use the `-nosmca` option.
- To restore the SMC Appliance configuration without restoring the Management Server or Log Server backup, use the `-smcaonly` option.

3) If you did not specify the backup file, select the backup file to be restored.
The default backup file names have the following structure:
   - sgm_vVERSION.[BUILD]_YYYYMMDD_HHMMSS[comment]
   - sgl_vVERSION.[BUILD]_YYYYMMDD_HHMMSS[comment]

4) If the backup is encrypted, enter the password.

5) When the backup restoration has finished, restart the SMC Appliance if you are prompted to do so.

Chapter 81

Managing log data

Contents

- Log data management and how it works on page 1307
- Reducing unnecessary log generation on page 1310
- Archive log data on page 1311
- Delete log data on page 1312
- Export log data on page 1315
- View history of completed Log Tasks on page 1316
- Overwrite old log or audit entries when log storage is full on page 1317
- Examples of log management on page 1317

Log management consists of configuring when log data produced, which log entries are stored, and when stored log entries are deleted or archived. To prevent the Log Server storage from filling up, log data management tools help you manage log entries automatically.

Log data management and how it works

Log data management keeps the number of logs at a reasonable level and prevents log files from filling the storage space.

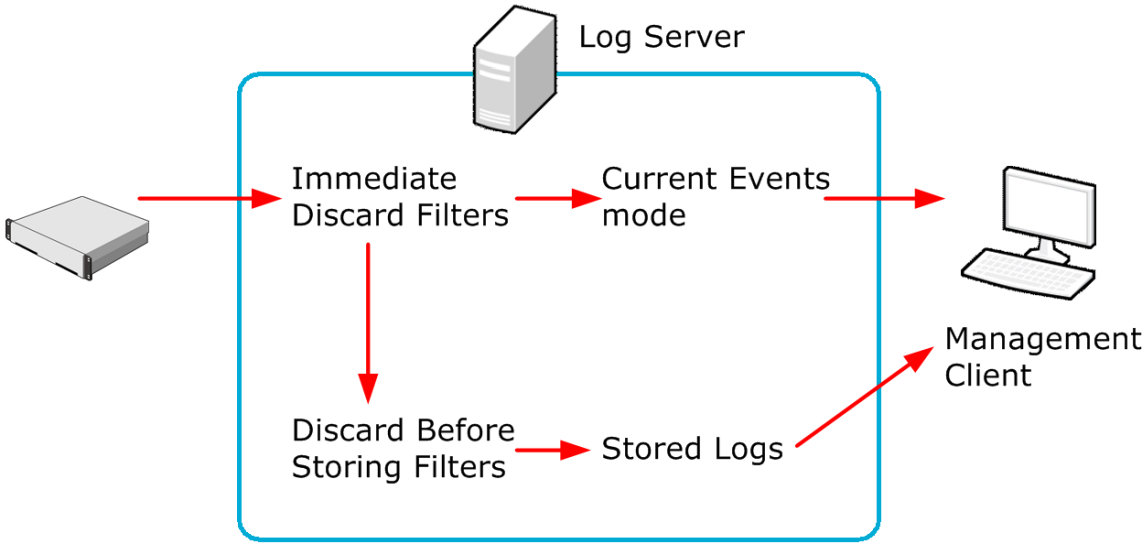
Log entries are stored in log files on the Log Server. Audit entries are stored on the Log Server or Management Server that originally created them. If these files are never removed, they eventually fill up the storage space on the Log Server or Management Server. In the properties of the Log Server and the Management Server, you can optionally specify what happens when the log storage is full. For the Log Server, an alert is automatically sent when the amount of log and audit data on the Log Server exceeds 75% of the total storage capacity.

You can manage the log data in the following ways:

- Configure logging options in rules to prevent unnecessary log entries from being created.
- Export log data so that it can be used elsewhere.
- Copy log data to an archive location.
- Delete old or unnecessary log data.
- Set up automatic log management tasks to run automatically at regular intervals for exporting, copying, and deleting selected data.
- Discard irrelevant log entries by pruning some of the log entries before they are stored on the Log Server.

This illustration demonstrates how log pruning filters are used in log data management.

Log pruning



The engines send their logs to their configured Log Server. The Log Server either stores the log entries or just relays them to be viewed immediately in the Current Events mode in the Logs view. Some logs might be discarded through pruning before these operations. When you view logs, the information is fetched directly from the Log Servers. Some other tasks, such as processing data for statistical reports, are also partially carried out by the Log Server.

You can prune log entries in two phases using Immediate Discard filters and Discard Before Storing filters. Immediate Discard filters delete log entries as they arrive to the Log Server. The Discard Before Storing filters delete log entries before the log entries are stored on the Log Server.

Note

Alert entries and audit entries cannot be pruned.

Limitations

Only the logs in the active storage are used in reporting. If you archive logs, you can still view them in the Logs view, but they are no longer available when you generate reports.

Alert and audit logs cannot be pruned.

Related concepts

Getting started with reports on page 311

What log entries are

Most often, Access rules trigger *Log entries*.

Other types of rules can be set to create log entries as well. However, in most recommended configurations, the volumes are much smaller than in Access rules. The system can also produce detailed *Diagnostic logs* and always produces some other internal log entries (such as entries related to policy installation).

What Alert entries are

Alert entries are notifications of important events that require the administrator's attention.

The difference between alerts and normal log entries is that alerts are highlighted in the SMC Client and they can be escalated through external notification channels. In addition to rule matches, alerts can be produced when an automatic test fails, a monitored element becomes unreachable, or if there is a system error.

Related concepts

Alert escalation and how it works on page 411

What Audit entries are

Audit entries provide a record of administrator actions and some internal events in the SMC.

Internal events in the SMC that produce audit entries include creation, modification and deletion of elements, administrator logons, and the execution of scheduled tasks. This data is useful, for example, when troubleshooting the cause of malfunctions caused by undocumented configuration changes.

Only the Management Server and the Log Server send audit entries. The audit entries are stored on the Management Server or the Log Server that originally created them.

Where log data files are stored

Logs are stored as a file on the Log Server.

A separate folder is created for the logged events each hour. The log files are stored by default in the <installation directory>/data/ storage/ directory on the Log Server.

The log files have the following naming: YYYYMMDD_hh_C<ORIGINATOR>_MMDDhhmmssss.arch.

- The date YYYYMMDD_hh refers to the date and hour of the logged events.
- The rest of the file name starting with "_C" refers to the file creation date and the <ORIGINATOR> refers to the originator of the logged events.

The time and date in the file name always use the UTC time zone, which is the system's internal time zone.

Archive directory options for the Log Server

You can change the Log Server's default archive directory and define up to 32 alternative or additional directories.

For example, you can directly archive some or all logs on a network drive to free resources on the Log Server. The Log Server's default archive directory is <installation directory>data/archive.

Domain boundaries in viewing and managing log data

If administrative Domains are configured, all log, alert, and audit entries are Domain-specific.

When you log on to a Domain, only the entries related to that specific Domain can be viewed or managed. However, Audit entries from all Domains are displayed to administrators who are logged on to the Shared Domain.

Log data management configuration overview

Log data management involves several main steps.

1) Configure logging to generate only the log entries you need.

2) (Optional) Set up log archiving to store older important logs for possible later use and free up the space on the Log Server.

3) (Recommended) Set up scheduled log data tasks for deleting logs that are not needed in the long term.

4) (Optional) Configure log pruning to prune out any unnecessary logs if any are generated.

Reducing unnecessary log generation

The primary way to manage logging is to set up the system to create all necessary logs and alerts and a minimum of unnecessary log entries.

The main generator of logs that you can configure are the rules in traffic handling policies. Another major point of configuration is the automatic tester, which you can set up to create alerts on various events. Some other features also generate logs and alerts, but it is not always possible to reduce the generation of logs from these features.

Normal and Alert logs are generated both based on internal conditions in the operation of a component and based on traffic that the engines handle.

Internal conditions that trigger logs or alerts:

- There is a system error or warning.
- An engine test fails. You can configure the engine tester in detail and select whether test failures trigger an alert.
- The status of an engine changes (not active by default).
- When the values of a monitored item exceed a threshold limit in an Overview (not active by default).
- Diagnostics are active on a Engine (not active by default).

Traffic conditions that trigger logs and alerts:

- An IPS engine's or a Layer 2 Engine's limit for the number of times tunneled traffic is rematched has been reached (not active by default).
- Traffic matches a rule in your policy.
- Diagnostics are active on an engine (not active by default).

You can also set up Log Servers to receive logs from any devices that can be set up to send syslog.

In addition to activating and deactivating logging and the listed features, you can optimize the number of generated logs on the engines in the following ways:

- You can configure log compression for Discard logs for Engines, IPS engines, and Layer 2 Engines.
- On Engines, you can configure log compression also for antispoofing logs.

Log Tasks can export, archive, and delete logs. It is possible to schedule these tasks to run automatically. The greater the volume of log data, the more frequently cleanup operations must run. For example, if the number of stored log entries is constantly high, you might need to export and delete logs daily. The schedules are defined in the SMC Client's local time. The Log Server might have a different time zone.

If administrative Domains are used, Log Tasks are always Domain-specific. You must define and run the Log Tasks in a specific Domain to apply them to the log data in that Domain.

Related concepts

Getting started with the Security Engine tester on page 637
Getting started with editing policies on page 885
Getting started with monitoring third-party devices on page 261

Related tasks

Enable or disable status monitoring on page 356
Set thresholds for monitored items in Overview elements on page 232
Configure inspection of tunneled traffic on page 683
Enable or disable diagnostics on page 356
Configure log handling settings on page 685

Archive log data

You can set up an Archive Log Task to copy log data from the active storage on the Log or Management Server to some other location.

The same task can also delete the log data from the active storage. You do not have to set up a separate task for freeing up the space.

By default, the log archive location is on the same disk drive as the active storage.



Note

The Archive Log Task copies the existing log files without compression. Copying the files without compression enables you to view the archived logs in the Logs view but they are not used in the Reports view when reports are generated.

Create an Archive Log Task

You can set up an Archive Log Task to copy log data from the active storage on the Log or Management Server to some other location.

You can define where the archived data is stored, and how the source data is handled after it has been copied to the target location.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.
2) Browse to **Tasks**.
3) Right-click **Tasks**, then select **New** > **Archive Log Task**.
4) Configure the settings.
5) Select the server from which the logs are archived and click **Add**.
6) Click **OK**.

Result

The task appears under **Task Definitions** in the **Tasks** branch of the **Administration** tree. You can run the task either manually or according to a fixed schedule.

Delete log data

To permanently delete generated log data, you can delete it from the active storage or delete it as it arrives to the Log Server using pruning filters.

Create a Delete Log Task

The recommended way to delete logs is to set up a Delete Log Task.

You can use Filter elements to select log data to be deleted.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.
2) Browse to **Tasks**.
3) Right-click **Tasks**, then select **New** > **Delete Log Task**.
4) Configure the settings.
5) Select the server from which the logs are deleted, then click **Add**.
6) Click **OK**.

Result

The task appears under Task Definitions in the Tasks branch of the Administration tree. You can run the task either manually or according to a fixed schedule.

CAUTION

When this task starts (either manually or as according to the schedule), all logs matching the selected filter and time range are permanently deleted from the active storage. Make sure that the data you want to keep is exported or copied to a safe location before the operation is started.

Create a Delete Elasticsearch Data Task

The Delete Elasticsearch Data Task deletes log data that has been forwarded to an Elasticsearch cluster.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.
2) Browse to **Tasks**.
3) Right-click **Tasks**, then select **New** > **Delete Elasticsearch Data Task**.
4) Configure the settings, then click **OK**.

Result

The task appears under Task Definitions in the Tasks branch of the Administration tree. You can run the task either manually or according to a fixed schedule.

CAUTION

When this task runs, all logs matching the selected filter and time range are permanently deleted from the active storage. Make sure that the data you want to keep is exported or copied to a safe location before the operation is started.

Discard unnecessary logs

Log data pruning allows you to discard some of the generated logs according to detailed filtering criteria you set.

Before you begin

If Domains are in use, you must log on to the Shared Domain.

You can manage the amount of log data by defining log pruning filters. Log pruning is needed, for example, when a rule generates both useful and unnecessary logs. Log pruning gives you the ability to discard newly generated irrelevant logs entries on the Log Server. Only logs can be pruned: alerts and audit entries are never pruned.

It is better to adjust log generation options instead of letting log entries be generated and then pruning them. Pruning log entries after they are generated wastes system resources by creating and transferring the unnecessary logs.

You can define two types of log pruning filters:

- *Immediate Discard* filters delete log entries immediately as they arrive to the Log Server. The deleted log entries are not displayed in the SMC Client.
- *Discard Before Storing* filters deletes log entries before they are saved. Log entries are shown the **Current Events** mode in the **Logs** view before they are deleted. This option converts Essential or Stored type log entries to Transient log entries.

If Domain elements have been configured, log pruning filters can only be defined in the Shared Domain. The same log pruning filters are used in all Domains. Administrators who have the right to manage log pruning can view the log pruning filters when they are logged on to other Domains.

CAUTION

Be careful when defining the pruning filters. The matching log events are irreversibly deleted at the Log Server.

You can prune log entries as soon as they arrive on the Log Server or before they are stored. When you prune log entries before they are stored, you can still view them in the **Current logs** view).

Use pruning with caution. The data is deleted without leaving any traces and there is no way to recover incorrectly pruned entries. Pruning also wastes resources compared to preventing the entries from being generated. Pruned entries still have to be created and transferred to the Log Server, and the Log Server still has to process them.

You can prune normal log entries; alert and audit entries cannot be pruned. The logs are pruned using log filters.

Note

Pruning has no effect on logs that have already been stored.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Menu** > **File** > **New Tab**, then select **Log Data Pruning**.

2) Select the tab of the log data type you want to prune.

3) Select a log filter for pruning log data from the tree in the Resources pane. If you add several Log Filters, they are combined with a logical OR. Each filter is matched individually, so all logs that match any of the selected filters are deleted.

   To create a Filter, select **New** > **Filter**.

   CAUTION

   Never select the Match All filter for log pruning. Pruning with the Match All filter irreversibly destroys all new logs that are created.

4) Activate the pruning for the correct stage:
   - Click **Add** below the **Immediate Discard** field to prune logs before they are even shown in the **Current Logs** view.

- Click **Add** below the **Discard Before Storing** field to show the log entries in the **Current Logs** view and then delete them before they are permanently stored.

CAUTION

Any log entry that matches the filter you have selected is irrevocably deleted. The changes you make to pruning filters are applied immediately.

5) A warning message is displayed. Click **Yes** to prune the selected log entries. This change is applied immediately without any further action.

Related concepts

Getting started with filtering data on page 333

Cancel log pruning

You can disable Pruning filters by removing them from the **Log Data Pruning** panes.

An empty pane means that no logs are pruned.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Menu** > **File** > **New Tab**, then select **Log Data Pruning**.

2) Select the correct tab according to log data type.

3) Select the filter you want to remove from pruning in the **Immediate Discard** or **Discard Before Storing** field.

4) Click **Remove** below the field that contains the filter.

Note

Log Filters that are removed from pruning remain available for other use until you delete them separately.

5) Click **OK** in the dialog box that is displayed to affirm this action. This change is applied immediately without any further action.

Export log data

Exporting log data allows you to copy log data to a location outside of the SMC Manager without archiving the log data.

You can use Log Export tasks to export log data from the Log Server or from the Management Server. You can either run the Log Export Task manually or schedule the task to run automatically. If administrative Domains have been configured, you must run the Log Export Task in the Domain to which the Log Server belongs.

You can also export extracts of log data while browsing logs.

Create an Export Log Task

You can use Log Export tasks to export log data from the Log Server or from the Management Server.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.
2) Browse to **Tasks**.
3) Right-click **Tasks**, then select **New** > **Export Log Task**.
4) Configure the settings.
5) Select the server from which the log data is exported and click **Add**.
6) Click **OK**.

Result

The task appears under **Task Definitions** in the **Tasks** branch of the **Administration** tree. You can run the task either manually or according to a fixed schedule.

View history of completed Log Tasks

You can view all previously executed tasks related to logs from a history file.

The history file includes information related to Export Logs Tasks, Archive Log Tasks, and Delete Log Tasks. The SMC never erases this file.

Steps

1) Open the `<installation directory>/data/logfile.txt` file in a text editor to view the previously executed tasks.

Related concepts

Log data management and how it works on page 1307

Overwrite old log or audit entries when log storage is full

By default, Log Servers stop receiving log entries when the log storage is full, and Management Servers shut down when the audit storage is full. You can optionally overwrite old log entries when the log storage is full.

When you configure Log Servers or Management Servers to overwrite old log entries or audit entries when the log or audit storage is full, the Log Server or Management Server writes new log entries or audit entries over the existing entries, starting with the oldest entries, until more disk space is available.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Network Elements**.

2) Browse to **Servers**, right-click a Log Server or Management Server element, then select **Properties**.

3) Configure the option to overwrite old log or audit entries depending on the type of server.
   - Log Server — From the **Log Storage Full** drop-down list, select **Overwrite Oldest**.
   - Management Server — From the **Audit Storage Full** drop-down list, select **Overwrite Oldest Audit Entries**.

4) Click **OK**.

Examples of log management

These examples illustrate some common uses for log management and general steps on how each scenario is configured.

Example: archiving old log data

This scenario shows an example of archiving and deleting logs to free up disk space on the Log Server.

Old logs are taking up too much disk space on the Log Servers at Company A, but some of the logs are still needed for the company's records. The administrators decide to archive the needed logs on another server and to delete last month's log data from the Log Servers. Because not all old logs need to be archived, they delete the unnecessary logs. They want to repeat the same archiving operation once a month. The administrators do the following:

1) Create an Archive Log Task for archiving the data with the following settings:

Archive Log Task for Company A

Option	Setting
Time Range	Last Full Month

Option	Setting
Filter for Copying	A custom filter that matches the important log data that the administrators want to archive.
Delete Other Data	**Match All** to delete all last month's logs from the Log Server.
Archive Target Directory	A network drive.

2) Save the Archive Log Task.

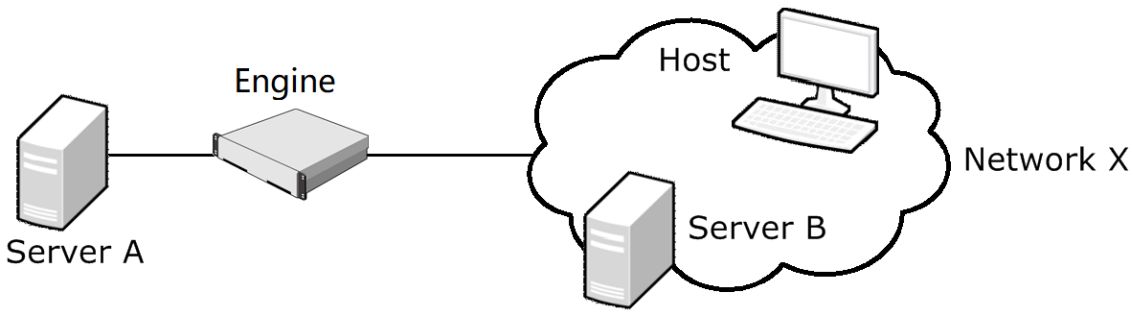
3) Create a Scheduled Task for running the Archive Log Task and set it to be repeated monthly.

4) Save the Scheduled Task.

Example: filtering out irrelevant logs

This scenario shows an example of using Log Data Pruning to temporarily discard unnecessary logs.

Company B's network



Server A provides services to users in network X, as well as to Server B. The administrators are interested in tracking how many of the users in network X actually use Server A. Server B also connects frequently to Server A, and generates a large amount of traffic.

The administrators are only interested in connections from other hosts in network X to Server A. The administrators decide to temporarily eliminate logs related to Server B's connections. All hosts in network X including Server B are currently logged according to a single rule. Creating a separate rule to handle Server B's connections with logging set to "None" would create unnecessary clutter in the policy. The administrators decide to set up log pruning to filter the logs so that only the relevant ones are stored on the Log Server. The administrators do the following:

1) Select one of the irrelevant log entries in the Logs view.

2) Create a temporary filter based on the log entry, and save the filter as a permanent filter.

3) Add the new filter to the **Discard Before Storing** list in the **Log Data Pruning** view.

Chapter 82

Managing and scheduling Tasks

Contents

- Getting started with Tasks on page 1319
- Task configuration overview on page 1320
- Task types on page 1321
- Creating Task Definitions on page 1322
- Schedule Tasks on page 1327
- Start Tasks manually on page 1328
- Pause the scheduled execution of Tasks on page 1329
- Remove Tasks from schedules on page 1329
- Stop running Tasks on page 1330

Tasks define parameters of system maintenance operations. You can run maintenance operations manually or automatically according to a schedule you set.

Getting started with Tasks

You can use Task elements for manual or automatic maintenance operations in the SMC.

What Tasks do

With Task elements, you can start maintenance operations either manually or according to a schedule. You can do the following with Tasks:

- Back up the Management Servers and Log Servers.
- Refresh policies.
- Upload policies.
- Validate policies.
- Upgrade engine software remotely.
- Export, archive, and delete logs.

There are also certain predefined system tasks.

Scheduling the Tasks allows you to run regular or one-time maintenance operations automatically, for example, during a regular maintenance window.

Limitations

Tasks are Domain-specific. The elements that are the target of the Task must belong to the Domain in which the Task is run. For example, to export log data from a Log Server you must run the Log Export Task in the Domain to which the Log Server belongs.

What do I need to know before I begin?

When scheduling automatic backups, you might want the data to be moved to a safe place automatically. This can be achieved through operating system scripts, which Tasks can start automatically upon completion. With Log Servers, you can change the backup and log archive locations in the Log Server's local configuration file.

Related reference

Task types on page 1321

Task configuration overview

Before you can run a Task, you must define the parameters for the Task element. If you want the task to be run automatically, you can also set up the Task execution schedule.

Follow these general steps to configure and use Tasks:

1) Define the Task parameters.

2) (Optional) Set up automatic Task execution.

3) Run Tasks when necessary.

Related concepts

Creating Task Definitions on page 1322

Related tasks

Schedule Tasks on page 1327
Start Tasks manually on page 1328
Pause the scheduled execution of Tasks on page 1329
Remove Tasks from schedules on page 1329
Stop running Tasks on page 1330

Task types

There are several types of predefined system Task Definitions. In addition, you can create custom Task Definitions.

Tasks are based on Task Definitions. There are two kinds of Task Definitions: custom Task Definitions and predefined System Task Definitions. To view Task Definitions or to create new Task Definitions, browse to **Administration** > **Tasks** > **Definition** in the **Configuration** view.

The following table explains the types of custom Tasks that you can create.

Custom Task Definitions

Task Definition	Explanation
Backup Task	Creates backup files for the selected Management Servers and Log Servers.
sgInfo Task	Creates a .zip file that contains copies of configuration files and system trace files for the selected components for Forcepoint Customer Hub.
Refresh Policy Task	Refreshes the currently installed policy on the selected engines and Master Engines, the currently installed Alert Policy on the selected Domains.
Refresh Policy on Master Engines and Virtual Security Engines Task	Refreshes the currently installed policy on the selected Master Engines and the Virtual Security Engines associated with the Master Engines. If a Virtual Engine belongs to another administrative Domain, the policy is refreshed in that Domain.
Upload Policy Task	Uploads the selected policy to the selected engines.
Validate Policy Task	Validates the selected policy on the selected engines.
Remote Upgrade Task	Remotely upgrades the software on the selected engines.
Export Log Task	Copies log data from the active storage or archive to the selected location.
Archive Log Task	Copies log data from the active storage to the selected location.
Delete Elasticsearch Data Task	Deletes log data that has been forwarded to an Elasticsearch cluster.
Delete Log Task	Deletes log data from the active storage.

In addition to Task Definitions that you create and customize, there are predefined Task Definitions for several system tasks. You can run the System Tasks manually or reschedule them, but you cannot change the options in System Task Definitions.

System Task Definitions

Task Definition	Explanation
Create Snapshot of All System Elements	Automatically creates a snapshot of all system elements after an update package has been activated. The snapshot information is used when administrators compare policy snapshots.
Delete Old Executed Tasks	Deletes information about previously executed tasks if there are more than 1000 executed tasks.
Delete Old or Unused Web Access Data	Deletes Web Access files that are older than 30 days or folders associated with a deleted administrator account.
Delete Old Snapshots	Deletes Policy Snapshots.

Task Definition	Explanation
Disable Unused Administrator Accounts	Disables the accounts of administrators who have not been active within the time period defined in the SGConfiguration.txt file. The accounts are disabled if the Enforce Password Settings option is enabled. An Administrator with unrestricted permissions (superuser) can re-enable the disabled accounts.
Fetch Certificate Revocation Lists	Downloads Certificate Revocation Lists (CRLs) from a CRL server.
Renew Gateway Certificates	Generates new certificates for Security Engines that act as VPN Gateways if automatic certificate renewal is enabled for the Security Engines.
Renew Internal Certificate Authorities	Checks the status of Internal Certificate Authorities for automatic renewal. To make sure that the automatic certificate authority renewal works correctly, do not change the schedule of this Task. This Task can only be run in the Shared Domain. If administrative Domains have been configured, it renews the internal certificate authority in all Domains.
Renew Internal Certificates	Checks the status of internal certificates for automatic renewal. To make sure that the automatic certificate renewal works correctly, do not change the schedule of this Task.

Related tasks

Create Backup Tasks on page 1323
Create sgInfo Tasks on page 1327
Create Refresh Policy Tasks on page 1323
Create Refresh Policy on Master Engines and Virtual Security Engines Tasks on page 1324
Create Upload Policy Tasks on page 1325
Create Remote Upgrade Tasks on page 1326
Enable and define password policy settings on page 386

Creating Task Definitions

You can create new Tasks to help you maintain your system. All Tasks are Domain-specific.

Related concepts

Log data management and how it works on page 1307

Related reference

Task types on page 1321

Create Backup Tasks

Create a Task for backing up servers.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.
2) Right-click **Tasks**, then select **New** > **Backup Task**.
3) Configure the settings, then click **OK**.

Create Refresh Policy Tasks

Create a Task for refreshing a policy on selected engines or Domains.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.
2) Right-click **Tasks**, then select **New** > **Refresh Policy Task**.
3) Give the Task a descriptive **Name** and optionally a free-form **Comment**.
4) Select the engines on which you want to refresh the policy, the Domains on which you want to refresh the Alert Policy from the list on the left. Click **Add** to add the policy to the list.

   Note

   If engines are the Target of the Task and you want already established connections to continue using the same configuration information (such as NAT rules), verify that **Keep Previous Configuration Definitions** is selected.

   The selected elements are added to the list on the right.
5) (Optional, only if Domains or a single engine is the Target.) To validate the rules when the Task is started manually, verify that **Validate Policy before Upload** is selected and select the related settings.
6) (Optional) Add an **Upload Comment** that is shown in Policy Snapshots or Alert Snapshots created from the policy installations.
7) Click **OK**.
   The new Refresh Policy Task is added to the list of Task Definitions.

Related tasks

Validate rules automatically on page 912
Create Refresh Policy on Master Engines and Virtual Security Engines Tasks on page 1324
Create Upload Policy Tasks on page 1325

Create Validate Policy Tasks

Create a Task for validating a policy on selected engines.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.
2) Right-click **Tasks**, then select **New** > **Validate Policy Task**.
3) Configure the settings.
4) Click **OK**.

Result

The new task is added to the list of Task Definitions.

Create Refresh Policy on Master Engines and Virtual Security Engines Tasks

The Refresh Policy on Master Engines and Virtual Security Engines Task is used for refreshing the currently installed policy on the selected Master Engines and all Virtual Engines that are associated with the Master Engines.

The policies of the Virtual Engines are installed in the administrative Domains to which the Virtual Engines belong.

This Task is primarily meant for situations in which it is necessary to refresh the policy on many Master Engines and Virtual Engines. This might be necessary, for example, after replacing the hardware on which the Master Engines run). In most cases, use the Refresh Policy Task if you want to refresh the policy on the Master Engines.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.
2) Right-click **Tasks**, then select **New** > **Refresh Policy on Master Engines and Virtual Security Engines Task**.

3) Give the Task a descriptive **Name** and optionally a free-form **Comment**.

4) Select the Master Engines on which you want to refresh the policy from the list on the left, then click **Add**.



Note

If you want already established connections to continue using the same configuration information (such as NAT rules), verify that **Keep Previous Configuration Definitions** is selected.

The selected Master Engines are added to the list on the right.

5) (Optional) Add an **Upload Comment** that is shown in Policy Snapshots created from the policy installations.

6) Click **OK**.
The new Refresh Policy on Master Engines and Virtual Security Engines Task is added to the list of Task Definitions.

Related tasks
Create Refresh Policy Tasks on page 1323
Create Upload Policy Tasks on page 1325

Create Upload Policy Tasks

Create a Task for uploading a policy on selected engines.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.

2) Right-click **Tasks**, then select **New** > **Upload Policy Task**.

3) Give the Task a descriptive **Name** and optionally a free-form **Comment**.

4) Select the engines on which you want to upload the policy, the Domains on which you want to upload the Alert Policy from the list on the left. Click **Add**.
The selected elements are added to the list on the right.

5) (Only if engines or Domains are selected as the Target). Click **Select** next to the **Policy** field, then select the policy you want to upload.



Note

If engines are the Target of the Task and you want already established connections to continue using the same configuration information (such as NAT rules), verify that **Keep Previous Configuration Definitions** is selected.

6) (Optional, only if Domains or a single engine is the Target). To validate the rules when the Task is started manually, verify that **Validate Policy before Upload** is selected, then select the related settings.

7) (Optional) Add an **Upload Comment** that is shown in Policy Snapshots or Alert Policy Snapshots created from the policy installations.

8) Click **OK**.
The new Upload Policy Task is added to the list of Task Definitions.

Related tasks
Validate rules automatically on page 912
Create Refresh Policy Tasks on page 1323
Create Refresh Policy on Master Engines and Virtual Security Engines Tasks on page 1324

Create Remote Upgrade Tasks

Create a Task for remotely upgrading the software on selected engines.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.

2) Right-click **Tasks**, then select **New** > **Remote Upgrade Task**.

3) Give the Task a descriptive **Name** and optionally a free-form **Comment**.

4) Select the upgrade **Operation**:
- **Remote Upgrade (transfer + activate)** — Loads the new configuration and reboots the node.
- **Remote Upgrade (transfer)** — Loads the new configuration without rebooting the node.
- **Remote Upgrade (activate)** — Reboots the node to activate a previously loaded configuration.

CAUTION
Do not activate the new configuration simultaneously on all nodes of a cluster. If you want to schedule a Remote Upgrade Task for several nodes, create two separate Remote Upgrade Tasks: one to transfer the new configuration and another to activate it. Schedule the Activate Task to run only after the Transfer Task is complete.

5) Select the engines that you want to upgrade from the list on the left, then click **Add**.
The selected engines are added to the list on the right.

6) Select the correct previously imported **Engine Upgrade** file for the upgrade.

7) Click **OK**.
The new Remote Upgrade Task is added to the list of Task Definitions.

Related tasks
Schedule Tasks on page 1327

Create sgInfo Tasks

Create a sgInfo Task to collect information about your system for Forcepoint Customer Hub.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.
2) Right-click **Tasks**, then select **New** > **sgInfo Task**.
3) Give the Task a descriptive **Name** and optionally a free-form **Comment**.
4) Select the engines whose configuration files and system trace files you want to get from the list on the left, then click **Add**.
   The selected engines are added to the list on the right.
5) If instructed to do so by Forcepoint Customer Hub, select **Include core files** to include the core files for troubleshooting.
6) Click **OK**.
   The new sgInfo Task is added to the list of Task Definitions.

Schedule Tasks

After creating Task Definition elements, you can schedule tasks to run at a convenient time.

If necessary, you can also schedule system tasks.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.
2) Select **Tasks** > **Definition**.
   A list of defined tasks opens.
3) Right-click the task, then select **Schedule**.

4) Set the schedule properties.



Tip

The date and time can be entered manually in the format YYYY-MM-DD HH:MM:SS. You can also right-click the Up or Down arrows next to the date field to select a date from the calendar.

5) From the **Final Action** drop-down list, select the action when the task finishes running.

6) Click **OK**.
The Task schedule is added under the Task Definition.

Related concepts
Creating Task Definitions on page 1322

Related tasks
Start Tasks manually on page 1328
Pause the scheduled execution of Tasks on page 1329
Remove Tasks from schedules on page 1329

Related reference
Task types on page 1321

Start Tasks manually

To run a Task immediately, start it from the Task Definitions list.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.

2) Select **Tasks** > **Definition**.
A list of Task Definitions opens.

3) Right-click the Task you want to start, then select **Start**.
The Task starts.

4) (Optional) Click the History branch to view the progress of the Task.
To view the Task, verify that **More actions** > **Show Executed Tasks Show** is enabled in the History branch.

Related tasks

Schedule Tasks on page 1327
Stop running Tasks on page 1330

Related reference

Task types on page 1321

Pause the scheduled execution of Tasks

If you want to temporarily stop a Scheduled Task from running at the scheduled time, you can suspend the Scheduled Task.

When a Scheduled Task is suspended, the schedule remains under the Task Definition, but the Task does not run at the scheduled time.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.

2) Select **Tasks** > **Definition**.
   The Task Definitions list opens.

3) Expand the Task Definition you want to suspend.
   The schedule information for the Task is displayed below the Task Definition.

4) Right-click the schedule information, then select **Suspend**.
   The Task is suspended.

5) To restart a suspended Task, right-click the schedule information, then select **Continue**.
   The Task resumes and runs at the next scheduled time.

Remove Tasks from schedules

You can remove Tasks from the schedule by moving the schedule information (Task Schedule) from the Task Definition to the Trash.

Moving the schedule information to the Trash does not delete the Task Definition: the same Task can be scheduled again.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.

2) Select **Tasks** > **Definition**.
The Task Definitions list opens.

3) Expand the Task you want to remove from the schedule.
The schedule information for the Task is displayed below the Task Definition.

4) Right-click the schedule information, then select **Delete**.
The schedules of default System Tasks cannot be deleted.
A confirmation dialog box opens.

5) To confirm that you want to move the selected Task Schedule to the Trash, click **Yes**.
The Task Schedule is removed from the Scheduled Tasks list.

Stop running Tasks

Abort a Task that is no longer needed.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.

2) Select **Tasks** > **Task History**.

3) Right-click the Task you want to abort, then select **Abort**.
A confirmation dialog box opens.

4) To stop the running Task, click **Yes**.

Chapter 83

Managing licenses

Contents

- Getting started with licenses on page 1331
- Generate licenses on page 1333
- Upgrading licenses manually on page 1335
- Changing license binding details on page 1336
- Install licenses for unlicensed components on page 1339
- Replacing licenses of previously licensed components on page 1340
- Check that all components are licensed on page 1342
- Check validity and status of licenses on page 1343

All SMC components and Security Engines must be licensed as a proof of purchase. In addition, some additional features can be activated by installing a feature license.

Getting started with licenses

Licenses prove that your organization has legally purchased the SMC components and Security Engines. You can upgrade them to new supported versions as part of each component's support and maintenance contract.

Licenses are issued when you purchase a product. They are generated at https://stonesoftlicenses.forcepoint.com and installed in the SMC as files. Licenses are shown as elements in the **Administration Configuration** view.

What do the different license types mean?

Licenses can be bound to a component in several different ways. The possible binding methods depend on the licensed component and the software version.

License binding methods

License binding	Description
IP address binding	The license is statically bound to the IP address of the licensed component. **Note** Only licenses for SMC servers can be bound to an IP address. Existing IP-address-bound licenses for other components continue to work and can be upgraded. Any new licenses for other components must be bound to the Management Server's proof-of-license (POL) code.

License binding	Description
UIID binding	The license is statically bound to the unique installation identifier (UIID) for the SMC. The UIID is automatically generated when you install the SMC. The UIID is also shown in the properties of the Management Server or Log Server elements.  **Note** Only licenses for SMC servers can be bound to the UIID for the SMC.
Management Server proof-of-license (POL) code binding	Licenses are dynamically bound to the Management Server's proof-of-license (POL) code. You must manually bind Management Server POL-bound licenses to the correct element. Licenses are valid only for components that are managed by the Management Server that has the same POL code.
Appliance proof-of-serial (POS) code binding	The license is bound to the unique POS code of a pre-installed Forcepoint Network Security Platform appliance. The appliance identifies itself when contacting the Management Server. The Management Server allows the use of the appliance if the license POS code matches the reported code. The Management Server automatically binds the correct license to the Security Engine element based on the POS code. For the Management Server and pre-installed appliances, the Management Server can use this licensing method automatically with new appliances.

The license types that are available depend on the SMC server or type of Security Engine.

License binding	Description
Management Servers	A static IP-address-bound license or a static UIID-bound license.
Log Servers	A static IP-address-bound license, a static UIID-bound license, or a dynamic license bound to the Management Server's POL code
Pre-installed Forcepoint Network Security Platform appliances	A license bound to the POS code of the appliance (all current models) or a dynamic license bound to the Management Server's POL code (older models)
Security Engines installed on your own hardware	Always a dynamic license bound to the Management Server's POL code
Security Engines installed on a virtualization platform	Always a dynamic license bound to the Management Server's POL code
Feature-specific licenses	A dynamic license bound to the Management Server's POL code or a license bound to the POS code of the appliance depending on the feature

Management Server license limits

Some Management Server licenses impose a restriction on the number of managed or monitored elements.

If your Management Server license allows an unlimited number of managed components, this restriction does not apply to you. To check this information, see the properties of the Management Server's License.

Each Security Engine is counted as one managed unit in the limitation, with some exceptions:

- Any number of clustered Security Engine nodes counts as a single managed unit, regardless of the number of nodes in the cluster. (One cluster equals one managed unit.)

- Third-party components that are monitored through the SMC count as one-fifth of a unit. (Five components equal one managed unit.)
- Virtual Security Engines are not counted against the license limit.

You cannot combine licenses in the SMC. For example, two Management Server licenses that each contain a restriction for five managed components only allow you to manage five components even if you bind both licenses to a single Management Server.

Check the appliance data sheets at https://www.forcepoint.com/resources/datasheets/next-generation-firewall to see the Management Count for your appliance.

Master Engine license limits

Master Engines use the same Security Engine licenses as other Security Engines. The Security Engine licenses enable the Virtual Engine features.

Virtual Engines do not require their own licenses. However, the Security Engine license limits the number of Virtual Resources that can be created. The limit for the number of Virtual Resources limits how many Virtual Engines can be created. You can optionally increase the allowed number of Virtual Resources by purchasing and installing a feature-specific license.

Generate licenses

Generally, each SMC component and Security Engine must have a separate license. Some additional features might also require a separate license.

There are some exceptions:

- In a high availability environment where there are multiple Management Servers, all Management Servers in the same SMC share a single license.
- All currently available Forcepoint Network Security Platform appliance models can fetch a license automatically through the Management Server if automatic updates are enabled. If automatic licensing fails, the appliances have a 30-day temporary initial license to allow time for manual licensing.
- Forcepoint Network Security Platforms deployed in the AWS cloud with the Bring Your Own License image must have a license in the SMC. Forcepoint Network Security Platforms deployed in the AWS cloud with the Hourly (pay as you go) license image do not require a separate license in the SMC.

Licenses always indicate the newest software version that you are entitled to, but they are valid for licensing any older software versions as well.



Note

Your SMC might be able to automatically generate licenses for new Forcepoint Network Security Platform appliances. For automatic licensing to work, install a license for the SMC components and make sure that automatic updates are enabled on the Management Server. The temporary initial license is automatically replaced with a permanent POS-bound license after the policy is first installed on the appliance.

Steps

1) Go to https://stonesoftlicenses.forcepoint.com.

2) In the **License Identification** field, enter the POL or POS code, as follows.

- Proof-of-license (POL) code — Identifies the license. For previously licensed components, the POL code is shown in the Licenses tree in the **Administration Configuration** view.
- Proof-of-serial (POS) number — The Forcepoint Network Security Platform appliances additionally have a proof-of-serial number that you can find on a label attached to the appliance hardware.

3) Click **Submit**.

4) Check which components are listed as included in this license, then click **Register**.

5) Read the instructions on the page, then fill in the required fields for all included components.

6) Enter the details that bind each license to a component, as follows:

Component	Details for license binding
Management Servers	■ IP-address-bound license — Enter the IP address that you plan to use on the server. If your license allows several Management Servers in the same SMC for high availability, enter a comma-separated list of the IP addresses of all Management Servers. ■ UIID-bound license — Enter the UIID of the SMC. The UIID is automatically generated when you install the SMC. The UIID is also shown in the properties of the Management Server or Log Server elements.
Other SMC servers	■ IP-address-bound license — Enter the IP address that you plan to use on the server. ■ Management Server POL-bound license — Enter the Management Server's POL code. ■ UIID-bound license — Enter the UIID of the SMC. The UIID is automatically generated when you install the SMC. The UIID is also shown in the properties of the Management Server or Log Server elements.
Master Security Engines	Enter the POS code of a Forcepoint Network Security Platform appliance (see the label attached to the appliance).
Security Engines	■ For Forcepoint Network Security Platform appliances, enter the POS code of the appliance (see the label attached to the appliance). ■ For Forcepoint Network Security Platform software installed on your own hardware or on a virtualization platform, enter the POL code of the Management Server that you use to manage the Security Engine. **Note** POS binding is always recommended when the option is available.

Note

If the binding information is incorrect, the license is unusable. If you accidentally generated a license with the wrong binding information, request a license change through the License Center.

7) Click **Submit Request**.
The license file available for download at the License Center.

Upgrading licenses manually

Licenses are valid for any older software versions in addition to the version indicated on the license. You can upgrade the licenses at any time without affecting the system's operation.



Note

IP-address-bound licenses have been previously available for Security Engines. You can use and update a previously generated IP-address-bound Security Engine license, but you must change the license binding to the Management Server's POL code if the Security Engine's control IP address changes.

You can view, change, and download your current licenses at https://stonesoftlicenses.forcepoint.com by logging on with your personal account (to view all licenses linked to that account) or by entering a proof-of-license (POL) or proof-of-serial (POS) code (to view information related to a particular license).

If automatic license upgrades have been enabled in the Management Server properties, your licenses are kept up to date automatically. Otherwise, you can upgrade licenses manually in the following ways:

- When you log on to the online License Center, you can upgrade the license for the displayed components through the link provided. Save the license as a file that you can install in the SMC.
- You can export information about licenses through the SMC Client and use the resulting text file to upgrade the licenses.

When do I have to upgrade licenses?

Components do not run without a valid license. Always make sure that you have an updated license before you make any change that is not supported by the current license.

Licenses must be updated for new software versions and if the binding detail in the license changes:

- Software upgrades — No action is required if automatic license upgrades have been enabled on the Management Server. In an environment with multiple Management Servers, no action is required if the automatic license upgrade feature has been enabled on the active Management Server that controls the Shared Domain. Otherwise, upgrade the licenses manually. License upgrades are available shortly before a new version is released.
- Changes in license binding — Change licenses manually if:
  - You change a control IP address used for license binding in a static IP-address-bound license.
  - You move a dynamic Management Server POL-bound license to a different Management Server.
  - You replace an IP-address-bound license for an SMC server to a license bound to the unique installation identifier (UIID) for the SMC.

  If you have IP-address-bound licenses for Security Engine components, you must change the license to a Management Server POL-bound license if you change the Security Engine's control IP address.

Each license indicates the maximum version for which the license is valid, but the license is also valid for all previous software versions. You must update the license if you upgrade a component to a new major release indicated by a change in the first two digits of the version number (for example, an upgrade from 1.2.3 to 1.3.0

or an upgrade from 1.2.3 to 2.0.0). If only the last number changes, the existing license is valid also for the later software version.

Appliance licenses do not allow upgrading an appliance beyond the last supported software version on which the appliance can run. See Knowledge Base article 9743 to check which appliance models have a last supported software version. With third-party hardware, be careful not to upgrade the software to a version that exceeds the hardware's capabilities.

Upgrade licenses manually

If you have not enabled automatic license upgrades, upgrade licenses manually through the SMC Client.

Licenses are valid for any older software versions in addition to the version indicated on the license. You can upgrade the licenses at any time without affecting the system's operation.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.

2) Select **Licenses**, then browse to the type of licenses that you want to upgrade.

3) Select the license that you want to upgrade.
   Details about the selected license open in the **Info** pane.

4) In the **Info** pane, copy the license information to the clipboard in one of the following ways:
   - From the **Proof of License** field, copy the POL code.
   - From the **Proof of Serial** field, copy the POS code.

5) Go to https://stonesoftlicenses.forcepoint.com.

6) In the **License Identification** field, paste the POL or POS code, then click **Submit**.

7) Under the license information, click **Update**.

8) Enter any information needed for the upgrade request, then select the license files to update.

9) To send the license request, click **Submit**.
   A confirmation page opens, showing the details of your request. The licenses are available for download on the license page.

Changing license binding details

You must manually update license binding details for licenses that are bound to IP addresses or POL codes.

Licenses that are bound to an IP address must be changed if the IP address of the component changes.

Note

Only licenses for SMC servers can be bound to an IP address. If you have IP-address-bound licenses for Security Engine components, you must change the license to a Management Server POL-bound license if you change the Security Engine's control IP address.

You can optionally replace IP-address-bound licenses for SMC servers with UIID-bound licenses.

Note

Only licenses for SMC servers can be bound to the UIID for the SMC.

Licenses that are bound to the POL code of the Management Server must be changed if:

- You want to transfer the licenses to a different Management Server.
- You replace the Management Server's license with a license that has a different POL code.

Forcepoint Network Security Platform appliances use POS-based licenses. The licenses are bound to the serial number of the appliance hardware and are automatically bound to the correct element.

You must change IP address binding and POL-based binding manually. To view, change, and download your current licenses at https://stonesoftlicenses.forcepoint.com, log on with your personal account (for all licenses that your account is authorized to view). Alternatively, you can enter a POL or POS code (to view information related to a particular license).

Note

If automatic license updates have been enabled in the Management Server properties, changed licenses are automatically downloaded and bound to the correct element as long as the license's identification code remains the same.

Change the binding method for licenses for SMC servers

You can optionally replace IP-address-bound licenses for SMC servers with UIID-bound licenses.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Change the license binding in the License Center.
   a) Go to https://stonesoftlicenses.forcepoint.com.
   b) In the **License Identification** field, enter the POL code.
      The proof-of-license (POL) code identifies the license. For previously licensed components, the POL code is shown in the Licenses tree in the **Administration Configuration** view.
   c) Click **Submit**.
   d) Below the list of licenses, click **Change binding**.

e) In the **Binding Type** drop-down list for the license that you want to change, select **Bind to Unique Installation ID (UIID)**.

f) In the **Binding** field, enter the UIID of the SMC.

The UIID is automatically generated when you install the SMC. The UIID is also shown in the properties of the Management Server or Log Server elements.

g) Click **Submit Request**.

The license file is available for download at the License Center.

2) In the SMC Client, delete the IP-address-bound licenses that you want to replace with UIID-bound licenses.

a) Select **Engine Configuration**, then browse to **Administration**.

b) Expand the **Licenses** branch, then browse to the correct type of license.

c) Right-click the license, then select **Delete**.

3) Install the UIID-bound licenses.

a) Select **Settings** > **Install Licenses**.

b) Select the UIID-bound license file or files, then click **Install**.

Each UIID-bound license is automatically bound to the SMC server. The **Binding** column shows the UIID of the SMC server.

4) Check the license information that is shown and verify that all components you meant to license have the correct new license.

Related tasks

Check validity and status of licenses on page 1343
Check that all components are licensed on page 1342

Install licenses for unlicensed components

All components must have a valid license. You must install licenses for any unlicensed components.

CAUTION

- If you have configured Domains, the licenses are always shown in the shared domain. Also, the licenses are visible in the domain where they are allocated to. Licenses that are not allocated to any Domain can be manually allocated to a particular Domain through their right-click menu.
- If you use Management Server POL code-based licenses for Forcepoint Network Security Platform appliances, verify that the licenses are bound to the correct Security Engine elements. Remember to change the license if you change the appliance hardware.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Settings** > **Install Licenses**.

2) Select one or more license files, then click **Install**.
   The licenses are installed. Most types of licenses are also automatically bound to the correct components.

3) Select **Administration**.

4) Expand the **Licenses** branch, then browse to the types of licenses you installed.

5) If you installed a dynamic Security Engine license (bound to the Management Server's POL code), right-click the newly installed license, then select **Bind** to select the correct Security Engine.
   If the license is bound to an incorrect element, right-click the license, then select **Unbind**.

   CAUTION

   The license is permanently bound to the Security Engine and cannot be bound to another Security Engine after you install or refresh the policy on the Security Engine. Permanently bound licenses cannot be rebound unless you relicense or delete the Security Engine element that the license is bound to. When unbound, a permanently bound license is shown as **Retained**.

6) Check that components are now licensed as intended.
   - The list of Unlicensed Components might show Security Engines that have a license bound to the POS code of an appliance. The reason for this is that POS-bound licenses are bound to the correct Security Engines automatically when the Security Engine is installed and makes initial contact with the Management Server. The Security Engine is moved to the correct type of Licenses branch after initial contact with the Management Server.
   - Engine licenses are applied when you upload or refresh the policy on the Security Engines.
   - If any Security Engines are not correctly licensed, you might need to upgrade or generate the licenses again.

Related tasks

Check validity and status of licenses on page 1343
Check that all components are licensed on page 1342

Replacing licenses of previously licensed components

All components must have a valid license. You must replace any invalid or missing licenses with new ones.

If you have configured Domains, licenses can be viewed only in the Shared Domain. When the license is bound to an element (automatically or manually), the license is shown only in the Shared Domain, regardless of the Domain to which the element belongs. Licenses that are not allocated to any Domain can be manually allocated to a particular Domain through their right-click menu.

CAUTION

If you use Management Server POL code-based licenses for Forcepoint Network Security Platform appliances, verify that the licenses are bound to the correct Security Engine elements. Remember to change the license if you change the appliance hardware.

Invalid or missing licenses can prevent components from working. If you are manually replacing working licenses with new ones, we recommend that you take a backup of the Management Server before you make changes. Then you can easily roll back the changes, if needed.

Related tasks

Back up system configurations on page 1297

Replace a Management Server POL-bound license with an IP-address-bound license or a POS-bound license

If you bind an IP-address-bound license or a POS-bound license to a component that already has a Management Server POL-bound license, the component has two licenses. To avoid duplicate licenses, replace the existing Management Server POL-bound license with a new license.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Settings** > **Install Licenses**.

2) Select one or more license files, then click **Install**.

3) Select **Administration**.

4) Expand the **Licenses** branch, then browse to the correct type of license.

5) Right-click the Management Server POL-bound license that is bound to the element, then select **Replace by Static License**.

6) Select the license to bind to the Security Engine.

7) Check the license information that is shown and verify that all components you meant to license have the correct new license.

8) Right-click any old licenses that might still be shown, then select **Delete**.

9) Refresh or install the policy to transfer the license changes to the Security Engine.

CAUTION

After you install or refresh the policy on the Security Engine, the license is permanently bound to the Security Engine and cannot be bound to any other Security Engine. Permanently bound licenses cannot be rebound unless you relicense or delete the Security Engine element that the license is currently bound to. When unbound, a permanently bound license is shown as **Retained**.

Related tasks

Check validity and status of licenses on page 1343
Check that all components are licensed on page 1342

Replace a Management Server POL-bound license with a different Management Server POL-bound license

If a component has a Management Server POL-bound license that bound to another Management Server, replace the license with a POL-bound license that is bound to the Management Server that manages the component.

For example, you might need to replace a Management Server POL-bound license with a different Management Server POL-bound license if you move an element from one Management Server to a different Management Server.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Settings** > **Install Licenses**.

2) Select one or more license files, then click **Install**.

3) Select **Administration**.

4) Expand the **Licenses** branch, then browse to the correct type of license.

5) If there is a previous Management Server POL-bound license bound to the element, right-click the old license, then select **Unbind**.

6) Right-click the new license, then select **Bind**.

7) Select the Security Engine to bind to the license.

8) Check the license information that is shown and verify that all components you meant to license have the correct new license.

9) Right-click any old licenses that might still be shown, then select **Delete**.

10) Refresh or install the policy to transfer the license changes to the Security Engine.

CAUTION

After you install or refresh the policy on the Security Engine, the license is permanently bound to the Security Engine and cannot be bound to any other Security Engine. Permanently bound licenses cannot be rebound unless you relicense or delete the Security Engine element that the license is currently bound to. When unbound, a permanently bound license is shown as **Retained**.

Related tasks

Check validity and status of licenses on page 1343
Check that all components are licensed on page 1342

Check that all components are licensed

Each SMC Server, Security Engine, and Master Engine must have its own license. There is no difference between licenses for nodes in a cluster and licenses for single Security Engines.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.

2) Select **Licenses** > **Unlicensed Components** > **All**.
This view displays all elements that require a license, but do not currently have a valid license.

3) If any components are shown as unlicensed, check the following:
   - Engine licenses generated based on a Forcepoint Network Security Platform appliance POS code are bound when the Security Engine makes initial contact with the Management Server. It is normal to see the corresponding elements as unlicensed until initial contact is made.

- If you have already generated and installed licenses, check that the binding details are correct (POS code, POL code, or IP address).

4) If needed, generate and install new licenses for the components listed.

Related tasks
Generate licenses on page 1333

Check validity and status of licenses

You can check license validity and binding information in the SMC Client.

The view displays each license and the component it is bound to, with the newest software version that the license allows you to install. Licenses are valid for all minor releases within the displayed major version (for example, a 5.1 license allows installing 5.1.0, 5.1.1, and 5.1.2) and for any previous software version.

It is not possible to create new IP-address-bound licenses for Security Engine components. You can upgrade previously created IP-address-bound Security Engine licenses to new versions. However, if a license is changed in any other way, the binding must also be changed.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.

2) Select **Licenses** > **All Licenses**.

Related tasks
Troubleshoot licenses that are shown as retained on page 1389

Chapter 84

Upgrading the SMC

Contents

- Getting started with upgrading the SMC on page 1345
- Upgrading the SMC configuration overview on page 1346
- Obtain SMC installation files on page 1347
- Upgrade SMC servers on page 1348
- Default SMC installation directories on page 1349

You can upgrade the Management Servers, SMC Clients, Log Servers, and Web Access Servers in your SMC.

Getting started with upgrading the SMC

Upgrading the SMC updates the system software version, activates the newest dynamic update package, and can include changes to elements. Backing up the system before an upgrade is recommended.

What SMC upgrades do

In addition to updating the SMC software, the upgrade makes other changes in your SMC:

- New system elements and policies can be added and obsolete system elements can be removed. Elements that are used are not deleted, but instead converted from system elements to regular elements when they have no default role anymore.
- Any element can be updated with new types of options (related to new or changed features), and occasionally obsolete options can be removed or changed.
- A new dynamic update package is activated, unless you have already installed the same or a newer update package before the installation. The previous installation can be the cause of some, but not necessarily all, of the preceding changes listed.
- The online help in the SMC Client is updated.

A summary of changes to elements is created during each upgrade; a link to these HTML reports is displayed when the Management Server upgrade is finished.

Limitations

All SMC components (Management Server, SMC Client, Log Server, and the optional Web Access Server) must have the same software version. All other components try to connect to the Management Server when they start. They do not start if their software version does not match with the Management Server software version. If you have multiple Management Servers or Log Servers, you must upgrade each server separately so that they have the same software version.

What do I need to know before I begin?

Although the need to do so is unlikely, the upgrade can be easily reversed if you take the correct precautions. If the SMC upgrade fails, you can automatically revert to the previous installation if you select the option in the installer before the upgrade starts. In any case, we recommend that you take a backup of the Management Server using the SMC's internal backup tool before you upgrade.

The backup contains all necessary information to restore the configurations (including the engine configurations). The backup does not contain software version or operating system specific information. It can always be restored from an older version of the SMC to a newer version of the SMC if a direct upgrade is supported between the software versions involved.

The SMC is offline during the upgrade. The engines continue to operate normally and store their generated log data in their local spool while the Management Servers and Log Servers are offline. Once connectivity is restored, the spooled log data is transferred from the engines to the Log Servers.

To check which version of the SMC you are currently using, select **Help** > **About** in the SMC Client. Also, the SMC Client's version is displayed in the SMC Client's logon dialog box.

Related tasks

Back up system configurations on page 1297

Upgrading the SMC configuration overview

All SMC components must be upgraded at the same time. Before upgrading, make sure that SMC licenses are up to date.

Follow these general steps to upgrade the SMC.

1) Obtain the installation files and check the installation file integrity.

2) (If automatic license upgrades have been disabled) Upgrade the licenses.

3) Upgrade all components that work as parts of the same SMC.

4) (Multiple Management Servers only) Synchronize the management database between the Management Servers.

5) Upgrade the SMC Clients that are installed locally on workstations.
If you are using the SMC Client in a web browser through Web Access, there is no need to upgrade.

Related concepts

Upgrading licenses manually on page 1335

Related tasks

Synchronize databases between the active Management Server and additional Management Servers on page 478
Obtain SMC installation files on page 1347
Upgrade SMC servers on page 1348

Obtain SMC installation files

Before running the installer, you must download the correct installation file from the download page and make sure that the checksums are correct.

You can check the installation file integrity using the file checksums. The checksums are on the installation DVD and in the release notes. In Windows environments, you can use Windows PowerShell to generate checksums. Several third-party programs are also available.

We provide only recent versions of the software for download. We recommend that you store the upgrade files yourself to make sure you can install the exact same version later (for example, if there is a hardware failure). This is recommended especially if your organization's policies mandate lengthy testing periods that limit the speed of adopting new versions.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Download the installation file from https://support.forcepoint.com/s/download.

   Separate .zip packages are available for downloading installation files for all supported platforms or just one supported platform.

2) Browse to the directory that contains the files to be checked.

3) Generate a checksum of the file.

   Examples in a Linux environment where `filename` is the installation file:

   ```
   sha1sum filename
   sha256sum filename
   sha512sum filename
   ```

   Example result:

   ```
   sha1sum smc_1.0.1.1000.zip
   79785edab5d2a1191a3065510756f72883952455 smc_1.0.1.1000.zip
   ```

4) Compare the displayed output to the checksum on the website.

   CAUTION

   Do not use files that have invalid checksums. If downloading the files again does not help, contact Forcepoint Customer Hub to resolve the issue.

5) Unzip all folders and files in the archive to the server you want to upgrade.

Upgrade SMC servers

You can upgrade SMC servers without uninstalling the previous version. A change in the Management platform, such as a new operating system or different hardware, requires reinstalling the SMC.

CAUTION

All SMC components (Management Server, SMC Client, Log Server, and the optional Web Access Server) must use the same SMC software version to work together. If you have multiple Management Servers or Log Servers, you must upgrade each server separately.

The same installer works with all SMC components, including locally-installed SMC Clients.

If you have multiple Management Servers or Log Servers, you can upgrade them in any order. Management Servers are automatically isolated from database replication during the upgrade. There is no need to explicitly isolate the Management Servers before upgrading.

If you are upgrading from a very old version of the SMC, you might have to upgrade to an intermediate version first before upgrading to the latest version. See the Release Notes.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Log on to the operating system with administrator rights in Windows or as the root user in Linux.

2) Start the Installation Wizard from a `.zip` file or the Installation DVD.
Decompress the `.zip` file.
   - On Windows, the executable is `\Forcepoint_SMC_Installer\Windows-x64\setup.exe`
   - On Linux, the executable is `/Forcepoint_SMC_Installer/Linux-x64/setup.sh`

   If the DVD is not automatically mounted in Linux, use the following command:

   ```
   mount /dev/cdrom /mnt/cdrom
   ```

3) Select the language for the installation, then click **OK**.
The language that you select is also set as the default language of the SMC Client.

4) Read the information on the **Introduction** page, then click **Next**.

   Tip

   Click **Previous** to go back to the previous page, or click **Cancel** to close the wizard.

5) Select **I accept the terms of the License Agreement**, then click **Next**.

6) To accept the installation directory that was automatically detected, click **Next**.
The Installation Wizard displays the components to be upgraded.

7) (Management Server only, optional) To save a copy of the current installation that you can revert to after the upgrade, select **Save Current Installation**, then click **Next**.

8) (Management Server only) Select whether to back up the server, then click **Next**.

- To create a backup that can be used and viewed without a password, select **Yes**.
- To create a password-protected backup, select **Yes, encrypt the backup**. You are prompted for the password as you confirm the selection.
- If you already have a recent backup of the Management Server, select **No**.

9) Check that the information in the **Pre-Installation Summary** is correct, then click **Install**.

10) (Optional) When the upgrade is complete, click the links in the notification to view the reports of changes the installer has made.
The report opens in your web browser.

11) When the installation has completed, click **Done**.

Next steps

1) Upgrade any SMC components that run on other computers (for example, additional Management Servers or Log Servers).

2) (Multiple Management Servers only) Synchronize the management database between the Management Servers.

Related tasks
Change the Management Server or Log Server platform on page 490
Synchronize databases between the active Management Server and additional Management Servers on page 478

Default SMC installation directories

The location of the SMC installation directory depends on the operating system.

- Default installation directory on Windows — `C:\Program Files\Forcepoint\SMC`

Note
If you installed the Management Server in the `C:\Program Files\Forcepoint\SMC` directory in Windows, some program data might be stored in the `C:\ProgramData\Forcepoint\SMC` directory.

- Default installation directory on Linux — `/usr/local/forcepoint/smc`

Under the installation directory are the following folders:

- `/backups/` — Stores Management Server (sgm_) and Log Server (sgl_) backups. The backups must be in this directory to be listed in the SMC Client and when running scripts without specifying a backup file.
- `/bin/` — Contains the SMC command-line tools as well as some additional scripts that are used by the Installation Wizard.

Related reference
Forcepoint Security Management Center commands on page 1429

Chapter 85

Upgrading Security Engines

Contents

- Getting started with upgrading Security Engines on page 1351
- Upgrading the Security Engines configuration overview on page 1352
- Obtain and import Security Engine upgrade files on page 1353
- Upgrade Security Engines remotely on page 1354

You can upgrade Engines, IPS engines, Layer 2 Engines, and Master Engines.

Getting started with upgrading Security Engines

You can remotely upgrade engines using the SMC Client or locally on the engine command line.

Remote upgrade is recommended in most cases. See the *Forcepoint Network Security Platform Installation Guide* for detailed instructions if you want to upgrade engines locally.

How engine upgrades work

The upgrade package is imported to the Management Server manually or automatically. Before the import, the Management Server verifies the digital signature of the upgrade package using a valid Trusted Update Certificate. The signature must be valid for the import to succeed. Verification failure can result from an out-of-date SMC version, in which case the SMC must be upgraded, or an invalid or missing signature, in which case the administrator must obtain an official upgrade package.

After the upgrade package has been imported, you can apply it to selected engines through the SMC Client. Before the upgrade is installed on the engines, the Management Server verifies the digital signature of the upgrade package. Also the engines verify the digital signature of the upgrade package before the upgrade is installed. Upgrade package digests are calculated using an SHA-512 hash and signed with an ECDSA key.

The engines have two alternative partitions for the software. When you install a new software version, it is installed on the inactive partition and the current version is preserved. This allows rollback to the previous version in case the installation is interrupted or other problems arise. If the engine is not able to return to operation after the upgrade, it automatically switches back to the previous software version at the next restart. You can also switch the active partition manually.

You can upload and activate the new software separately. For example, you can upload the upgrade during office hours but activate it during a service window.

The currently installed working configuration (routing, policies) is stored separately and is not changed in an upgrade or a rollback. Although parts of the configuration can be version-specific (for example, if system communications ports are changed), the new software version can use the existing configuration. Possible version-specific adjustments are made when you refresh the policy after the upgrade.

Limitations

You cannot upgrade Virtual Engines directly. To upgrade Virtual Engines, you must upgrade the Master Engine that hosts the Virtual Engines.

What do I need to know before I begin?

The SMC must be up to date before you upgrade the engines. An old SMC version might not be able to recognize the new version engines and can generate an invalid configuration for them. The Management Server can control several older versions of engines. See the Release Notes for version-specific compatibility information.

During a cluster upgrade, it is possible to have the upgraded nodes online and operational side by side with the older version nodes. This way, you can upgrade the nodes one by one while the other nodes handle the traffic. However, you must upgrade all nodes to the same version as soon as possible, as prolonged use with mismatched versions is not supported.

The current engine version is displayed on the **General** tab in the **Info** pane when you select the engine. If the **Info** pane is not shown, select **Menu** > **View** > **Panels** > **Info**.

Upgrading the Security Engines configuration overview

You can upgrade Security Engines remotely using the SMC Client.

Follow these general steps to upgrade engines:

1) (Manual download of Security Engine upgrade files) Prepare the installation files.
2) (Manual license updates) Update the licenses.
3) Upgrade the Security Engines.

Related concepts

Upgrading licenses manually on page 1335

Related tasks

Manually update the anti-malware database on page 991
Upgrade Security Engines remotely on page 1354

Obtain and import Security Engine upgrade files

If the Management Server is not set up to download engine upgrades automatically, download the installation files manually.

You must also check the installation file integrity using the file checksums. In Windows environments, you can use Windows PowerShell to generate checksums. Several third-party programs are also available.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Go to https://support.forcepoint.com.

2) Enter your license code or log on using an existing user account.

3) Select **Downloads**.

4) Under **Network Security**, click the version of the Forcepoint Network Security Platform software that you want to download, then download the .zip file installation file.

5) On your local computer, change to the directory that contains the files to be checked.

6) (Linux only) Generate a checksum of the file using one of the following commands, where `filename` is the name of the installation file:

   - `sha1sum filename`
   - `sha256sum filename`
   - `sha512sum filename`

   For Windows, see the documentation for the third-party checksum program.

   Example:

   ```
   $ sha1sum sg_engine_1.0.0.1000.iso
   869aecd7dc39321aa2e0cfaf7fafdb8f sg_engine_1.0.0.1000.iso
   ```

7) Compare the displayed output to the checksum on the website.

   CAUTION

   Do not use files that have invalid checksums. If downloading the files again does not help, contact Forcepoint Customer Hub to resolve the issue.

8) Log on to the SMC Client, then select **Menu** > **File** > **Import** > **Import Engine Upgrades**.

9) Select the engine upgrade (sg_engine_version_platform.zip file), then click **Import**.

The import takes a while. You can see the related messages in the status bar at the bottom of the SMC Client window.

Note

The Management Server verifies the digital signature of the .zip file before importing it. The signature must be valid for the import to succeed. If the verification fails, an error message is shown. Verification failure can result from an out-of-date SMC version or an invalid or missing signature.

Related concepts

Getting started with automatic updates and upgrades on page 1291

Upgrade Security Engines remotely

The Management Server can remotely upgrade Security Engine components that it manages.

Before you begin

Read the Release Notes for the new version, especially the required SMC version and any other version-specific upgrade issues that might be listed. To access the release notes, select **Administration** > **Other Elements** > **Engine Upgrades**. Select the type of Security Engine you are upgrading. A link to the release notes is included in the upgrade file's information. If the Management Server has no Internet connectivity, you can find the release notes at https://support.forcepoint.com/s/article/Documentation-Featured-Article.

CAUTION

If McAfee Endpoint Intelligence Agent (McAfee EIA) is configured on the Engine when you upgrade to version 6.3 or later, the Engine node is returned to the initial configuration state and stops processing traffic. You must remove the McAfee Endpoint Intelligence Agent (McAfee EIA) configuration and refresh the policy before you upgrade to version 6.3 or later. For more information, see Knowledge Base article 14093.

You can upgrade several Security Engines of the same type in the same operation. However, we recommend that you upgrade clusters one node at a time and wait until an upgraded node is back online before you upgrade the other nodes. Clusters operate normally throughout the upgrade when the upgrade is done in stages. However, it is recommended to upgrade all nodes in the cluster to the same version as soon as possible. Prolonged use with mismatched versions is not supported. It is not possible to have 32-bit and 64-bit Security Engines online in the cluster at the same time.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Dashboard** > **Engines Dashboard**.

2) Expand the nodes of the Security Engine that you want to upgrade.

3) Right-click the node that you want to upgrade, then select **Commands** > **Go Offline**.

4) (Optional) Enter an **Audit Comment** to be shown in the audit log entry that is generated when you send the command to the Security Engine.

5) When prompted to confirm that you want to set the node offline, click **Yes**.
The node goes offline shortly.

6) When the node is offline, right-click the node, then select **Upgrade Software** or **Configuration** > **Upgrade Software** depending on your selection.

Note

You cannot upgrade Virtual Engines directly. To upgrade Virtual Engines, you must upgrade the Master Engine that hosts the Virtual Engines.

7) From the **Operation** drop-down list, select the type of operation that you want to perform:
   - Select **Remote Upgrade (transfer + activate)** to install the new software and reboot the node with the new version of the software.
   - Select **Remote Upgrade (transfer)** to install the new software on the node without an immediate reboot and activation. The node continues to operate with the currently installed version until you choose to activate the new version.
   - Select **Remote Upgrade (activate)** to reboot the node and activate the new version of the software that was installed earlier.

CAUTION

To avoid an outage, do not activate the new configuration simultaneously on all nodes of a cluster. Activate the new configuration one node at a time, and proceed to the next node only after the previous node is back online.

8) If necessary, add or remove Security Engines in the **Target** list.
All Security Engines in the same Upgrade Task must be of the same type.

9) Click **Select** next to the **Engine Upgrade** field, select the upgrade file, then click **OK**.
If you choose to activate the new configuration, you are prompted to acknowledge a warning that the node will be rebooted. A new tab opens showing the progress of the upgrade. The time the upgrade takes varies depending on the performance of your system and the network environment. The Security Engine is automatically rebooted and brought back online.

The upgrade overwrites the inactive partition and then changes the active partition. To undo the upgrade, use the `sg-toggle-active` command or the Security Engine's boot menu to change back to the previous software version on the other partition. This change can also happen automatically at the next reboot if the Security Engine is not able to successfully return to operation when it boots up after the upgrade.

Note

The Management Server verifies the digital signature of the upgrade package before installing it. The signature must be valid for the upgrade to succeed. If the verification fails, an error message is shown. Verification failure can result from an out-of-date SMC version or an invalid or missing signature.

Related tasks

Access the Security Engine command line on page 364
Create Remote Upgrade Tasks on page 1326

Related reference

Security Engine commands on page 1445

Chapter 86

Manual dynamic updates

Contents

- Getting started with manual dynamic updates on page 1357
- Dynamic update configuration overview on page 1358
- Import dynamic update packages on page 1358
- Activate dynamic update packages on page 1359

Dynamic Update packages include changes and additions to the system Policies, Situations, and other elements of the SMC.

Getting started with manual dynamic updates

It is important to keep the system policies and situations up to date so that newly discovered vulnerabilities can be detected. Changes and additions are provided in dynamic update packages.

Dynamic updates are available at https://autoupdate.ngfw.forcepoint.com.

Dynamic update packages are imported to the Management Server manually or automatically. Before the import, the Management Server verifies the digital signature of the dynamic update package using a valid Trusted Update Certificate. The signature must be valid for the import to succeed. Verification failure can result from an out-of-date SMC version, in which case the SMC must be upgraded, or an invalid or missing signature, in which case the administrator must obtain an official dynamic update package.

What dynamic updates do

Dynamic update packages provide updates for Security Engines, especially for deep inspection features. For example, new threat patterns and changes in the system Templates and Policies are introduced in dynamic updates for up-to-date detection. They can also revise the default elements you use to configure the system.

Limitations

Some limitations apply to installing dynamic updates:

- You might need to upgrade first before you can use a certain dynamic update package. For more information about the update packages, see the *Release Notes*.
- If there are several Domains defined in the SMC, manual dynamic updates can only be installed in the Shared Domain.

What do I need to know before I begin?

As an alternative to downloading the updates manually as explained here, you can configure the dynamic updates to be downloaded and optionally activated automatically.

Malware database updates are always done automatically and directly by the engines. Updates are always active when the anti-malware feature is active.

Related concepts

Getting started with automatic updates and upgrades on page 1291

Dynamic update configuration overview

Download and activate dynamic update packages using the SMC Client.

Follow these general steps to import and activate dynamic updates:

1) Download the latest dynamic update package and import it in the SMC Client.

2) Activate the dynamic update package in the SMC Client.

Related tasks

Import dynamic update packages on page 1358
Activate dynamic update packages on page 1359

Import dynamic update packages

Download the latest dynamic update package and import it in the SMC Client.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Go to https://autoupdate.ngfw.forcepoint.com.

2) On the **Dynamic Updates** tab, download the latest dynamic update package .jar file.
For details about the dynamic update package, click **Release Notes** under the .jar file.

3) Save the update package file to a location accessible from the computer you use to run the SMC Client.



Note

Make sure that the checksums for the original files and the files that you have downloaded match.

4) In the SMC Client, select **Menu** > **File** > **Import** > **Import Update Packages**.

5) Browse to the file, select it, then click **Import**.
The import takes some time, and the completion of the import is displayed in the status bar of the SMC Client window.

Note

The Management Server verifies the digital signature of the dynamic update package before importing it. The signature must be valid for the import to succeed. If the verification fails, an error message is shown. Verification failure can result from an out-of-date SMC version or an invalid or missing signature.

Activate dynamic update packages

To introduce the changes from an imported dynamic update package into your SMC, activate the dynamic update package in the SMC Client.

Note

The Management Server verifies the digital signature of the dynamic update package before activating it. The signature must be valid for the activation to succeed. If the verification fails, an error message is shown. Verification failure can result from an out-of-date SMC version or an invalid or missing signature.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.

2) Browse to **Other Elements** > **Updates**.

3) Right-click the imported dynamic update package, then select **Activate**.
You are prompted to create a Management Server backup. We recommend that you create a Management Server backup before you activate the dynamic update package.
The progress of the import and the items included are shown on a new tab.

4) When the activation is finished, click **Close**.

5) Refresh the policy on all Security Engines.
If your policy uses a custom template, you might need to edit the policy.

Chapter 87

SMC Appliance maintenance

Contents

- Getting started with SMC Appliance maintenance on page 1361
- Patching and upgrading the SMC Appliance on page 1362
- Roll back the SMC Appliance to the previous version on the command line on page 1365

The SMC Appliance has a specific patching process that keeps the SMC software, operating system, and appliance firmware up-to-date.

Getting started with SMC Appliance maintenance

SMC Appliance patches can include improvements, enhancements, and upgrades for the SMC software, the operating system, and the appliance firmware.

The SMC Appliance patch (SAP) format is specific to the SMC Appliance. The SAP numbering is appended to the version number. Patch digests are calculated using an SHA-512 hash and signed with an ECDSA key.

There are two kinds of SMC Appliance patches:

- Hotfix patches include improvements and enhancements for the current SMC Appliance version.
  Hotfix patch files use the letter P as a separator between the version number and the patch number. Example: 7.3.0.1P01
- Upgrade patches upgrade the SMC Appliance to a new version.
  Upgrade patch files use the letter U as a separator between the version number and the patch number. Example: 7.3.0.1U01

When you install a patch, a configuration backup and a file system snapshot are automatically created for the SMC Appliance. The backup and snapshot allow you to roll back the SMC Appliance to its previous configuration if needed. If the patch activation fails, the appliance reverts to the snapshot automatically. The file system of the SMC Appliance has two partitions: an active partition and an alternative partition. Some patches update the alternative partition. You can toggle between the partitions to roll back the SMC Appliance upgrade.



Note

SMC Appliance patches apply only to the SMC Appliance hardware or to SMC Appliance software installed on a virtualization platform. SMC components installed on third-party platforms do not offer a patching and rollback feature that includes the SMC software, the operating system, and the appliance firmware.

You can patch and upgrade the SMC Appliance remotely using the SMC Client or using the appliance maintenance and bug remediation (AMBR) patching utility on the command line.

Configuration overview

1) Check for new SMC Appliance patches.
   There is no automatic notification when new SMC Appliance patches are available. We recommend checking for new SMC Appliance patches once a month.

2) Obtain the patch files.
   - You can use the SMC Client or the AMBR utility to automatically download patch files directly into the SMC Client or onto the SMC Appliance.
   - In environments without Internet connectivity, you must manually download patch files, then import them into the SMC Client or transfer them to the SMC Appliance.

3) (If automatic license upgrades have been disabled) Upgrade the licenses.

4) Upgrade the SMC Clients that are installed locally on workstations.
   If you are using the SMC Client in a web browser through Web Access, there is no need to upgrade.

Patching and upgrading the SMC Appliance

To introduce improvements and enhancements for the current SMC Appliance version, install a hotfix patch. To upgrade the SMC Appliance, install an SMC Appliance upgrade patch.

Before upgrading, read the Release Notes.

It is important to upgrade the SMC Appliance before upgrading the engines. An old SMC version might not be able to recognize the new version engines and can generate an invalid configuration for them. The Management Server can control several older versions of engines. See the release notes for version-specific compatibility information.

Patch or upgrade the SMC Appliance in the SMC Client

You can use the SMC Client to patch or upgrade the SMC Appliance. In some certified environments, you must use the SMC Client to install SMC Appliance patches.

Before you begin

In environments without Internet connectivity, you must download the SMC Appliance patch file from https://update.stonesoft.com/download/appliance/patches/, then transfer the file to a location that is accessible from the SMC Client.

Installing some SMC Appliance patches might restart the SMC Appliance. The Restart Required column in the SMC Client indicates whether the appliance must restart as part of the installation. You can use the SMC Client to install SMC Appliance patches regardless of whether the installation requires restarting the SMC Appliance.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Administration**.

2) Browse to **SMC Appliance Patches**.

3) Download or import SMC Appliance patches.
   - To download an individual patch, right-click a patch for which the **State** column shows Available, then select **Download SMC Appliance Patch**.
   - To automatically download all available SMC Appliance patches, right-click **SMC Appliance Patches**, then select **Download SMC Appliance Patches**.
   - To import SMC Appliance patches that you manually downloaded, right-click **SMC Appliance Patches**, select **Import SMC Appliance Patches**, browse to the SMC Appliance patch file, then click **Import**.

4) Install a hotfix patch or an upgrade patch.
   - To patch the current SMC Appliance version, right-click a hotfix patch file, then select **Activate**.
   - To upgrade the SMC Appliance to a new version, right-click an upgrade patch file, then select **Activate**.

Result

The SMC Appliance patch is installed on the SMC Appliance.

If you installed an SMC Appliance upgrade patch, the installation process restarts the appliance and installs the patch. When the upgrade is finished, the appliance restarts again.

Patch or upgrade the SMC Appliance on the command line

You can use the appliance maintenance and bug remediation (AMBR) patching utility to patch or upgrade the SMC Appliance on the command line.

Before you begin

In environments without Internet connectivity, you must download the SMC Appliance patch file from https://update.stonesoft.com/download/appliance/patches/, then transfer the files to the SMC Appliance.

If you do not have physical access to the SMC Appliance, use SSH to access the SMC Appliance remotely.

Note

In FIPS mode, SSH access to the SMC Appliance command line is not supported.

You must have SMC Appliance Superuser permissions to log on to the SMC Appliance command line. Administrators with unrestricted permissions (superusers) are allowed to log on to the SMC Appliance command line only if there are no administrators with Console Superuser permissions.

Use sudo if you need elevated privileges. For a list of available sudo commands, enter the following command:

```
sudo -l
```

Steps

1) From the command line, log on to the SMC Appliance.

2) To update the list of available remote patches from the download server, enter the following command:

```
sudo ambr-query -u
```

3) To show all local and remote patches, enter the following command:

```
sudo ambr-query -a
```

4) To automatically download a patch, or to load a patch that you manually downloaded, enter the following command:

```
sudo ambr-load <patch>
```



Note

If you manually downloaded the patch and transferred it to the SMC Appliance, append the command with the `-f` option and specify the full path to the patch file.

Example:

```
sudo ambr-load –f /var/tmp/7.3.0.0P001.sap
```

5) To activate the patch, enter the following command:

```
sudo ambr-install <patch>
```

Result

The SMC Appliance patch is installed on the SMC Appliance.

If you installed an SMC Appliance upgrade patch, the installation process restarts the appliance and installs the patch. When the upgrade is finished, the appliance restarts again.

Related reference

Forcepoint Security Management Center commands on page 1429

Unload SMC Appliance patches on the command line

When you load patches, they are copied to the patch storage on the SMC Appliance. Use the unload command to remove patches that have been loaded but not installed from the patch storage on the SMC Appliance.

Note

To revert changes made by a patch that has been installed, roll back the SMC Appliance to the previous version.

Steps

1) From the command line, log on to the SMC Appliance.

2) To unload the patch, enter the following command:

```
sudo ambr-unload <patch>
```

Result

The patch is removed from the patch storage on the SMC Appliance. To verify that it has been removed, enter the following command:

```
ambr-query
```

Note

Patch files that you manually transferred to the SMC Appliance are removed from the patch storage, but are not deleted from the file system of the SMC Appliance. You must manually delete these patch files.

Roll back the SMC Appliance to the previous version on the command line

To revert changes made by a patch that has been installed, you can roll back the SMC Appliance to the previous configuration if needed.

Note

When you roll back the SMC Appliance, configuration changes made after the current version was installed are lost. We recommend rolling back as a recovery option only for a short time after an upgrade.

Steps

1) From the command line, log on to the SMC Appliance.

2) Enter the following command:

```
sudo smca-system toggle
```

3) Restart the SMC Appliance.

Result

When the SMC Appliance restarts, the previous SMC Appliance configuration is in use.

Part XIII
Troubleshooting

Contents

- General troubleshooting tips on page 1369
- Troubleshooting Administrator accounts and passwords on page 1371
- Messages for troubleshooting on page 1375
- Troubleshooting Security Engine operation on page 1385
- Troubleshooting licenses on page 1389
- Troubleshooting logging on page 1393
- Troubleshooting the SMC Client on page 1397
- Troubleshooting NAT on page 1403
- Troubleshooting policies on page 1407
- Troubleshooting reporting on page 1415
- Troubleshooting upgrades on page 1419
- Troubleshooting VPNs on page 1421

Troubleshooting helps you resolve common problems in the Forcepoint Network Security Platform and SMC.

Chapter 88

General troubleshooting tips

Contents

- If your problem is not listed on page 1369
- Tools for further troubleshooting on page 1369

General troubleshooting tips help you troubleshoot situations that are not covered by more specific troubleshooting topics.

If your problem is not listed

There are several possible causes and solutions for problems that are not listed in the troubleshooting topics.

When having problems with your system, first make sure that you have followed the relevant instructions.

Some problems you are having can be related to known issues, which you can view at:

https://support.forcepoint.com

If your organization is entitled to technical support, contact Forcepoint Customer Hub.

Tools for further troubleshooting

Information in logs and alerts, and networking tools on engine can be useful for troubleshooting.

Logs and alerts provide useful information about what the components do and what is happening in your system. You can increase the detail level of logs on certain areas of operation. To gain useful information from the logs produced, you must be able to filter the logs efficiently.

There are Forcepoint Network Security Platform-specific as well as standard networking tools available on the engines.

Related concepts

Benefits of filtering log entries on page 289

Considerations for working on the Security Engine command line on page 363

Related tasks

Enable or disable diagnostics on page 356

Related reference

Security Engine commands on page 1445

Chapter 89

Troubleshooting Administrator accounts and passwords

Contents

- Replace forgotten passwords on page 1371
- Troubleshoot user accounts on page 1372
- Create an emergency administrator account on page 1372

There are several common problems and solutions related to Administrator accounts and passwords.

Replace forgotten passwords

There are several ways to replace forgotten passwords depending on the type of account.

Problem description: You or someone else in your organization forgets one of the passwords related to the SMC.

Solution: You can regain access by changing the password. If none of the administrators can log on due to account issues, you can create an emergency administrator account. An administrator who has unrestricted permissions (superuser) can change any password in the SMC. The password recovery procedures for the different passwords are as follows:

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Change SMC Client logon passwords in the Administrator elements. Administrator elements can be found in the **Administration** branch of the **Configuration** view under **Access Rights** > **Administrators**.

2) Change Web Portal logon passwords in the Web Portal User elements. Web Portal elements can be found in the **Administration** branch of the **Configuration** view under **Access Rights** > **Web Portal Users**.

3) Change the Engine Root account password (for command-line access) by right-clicking the individual engine node and selecting **Commands** > **Change Password**.

   If the engine is not connected to the Management Server (because, for example, it is a spare appliance), you can reset all the engine's settings. You can reset settings through a boot menu option in the local console accessible through a serial connection or through a directly connected monitor and keyboard.

   CAUTION

   Resetting the engine through the boot menu stops the engine from processing traffic because all configurations are cleared.

4) Change user passwords used for end-user authentication in the User element. User elements are stored in the **User Authentication** branch of the **Configuration** view under **Users** (if the user is stored in the internal LDAP database or an external LDAP database that the SMC is configured to use).

5) To change the Management Server Database password, select **Settings** > **Password** > **Change Database Password**.

   The default Management Server Database user account is *dba*, and the password is created automatically.

Related tasks

Add administrator accounts on page 379
Create an emergency administrator account on page 1372

Troubleshoot user accounts

Resolve problems when end-user passwords are not accepted for authentication.

Problem description: You add a User element (end user account for authentication) or change the password in a User element, but the new user account or new password are not accepted when the end user tries to authenticate. Previously created and unmodified user accounts work as expected. If you changed the password, the previous password is still accepted.

Reason: There might be a replication problem that prevents synchronizing the user database information from the Management Server to the local database on the Engines.

Steps

1) Reset the user database by right-clicking an individual Engine node (not the upper-level Single Engine/Engine Cluster element) and selecting **Commands** > **Reset User Database**. This action copies all user information from the Management Server to the engine.

2) Make sure **User DB Replication** (automatic user database replication) is active under **Options** in the right-click menu for the Single Engine/Engine Cluster (top-level) element.

Create an emergency administrator account

If none of the administrators can log to the SMC Client because of account-related issues, you can create an emergency administrator account.

Steps

1) On the command line of the Management Server computer, stop the Management Server service.

2) Create an account with unrestricted permissions (superuser) using the `sgCreateAdmin` script (located in the <installation directory>/bin folder).



Note

If you installed the Management Server in the `C:\Program Files\Forcepoint\SMC` directory in Windows, some program data might be stored in the `C:\ProgramData\Forcepoint\SMC` directory.

Result

You can now log on to the SMC Client using this account, change passwords, and create new accounts normally.

Chapter 90

Messages for troubleshooting

Contents

- Alert log messages for troubleshooting on page 1375
- Log messages for troubleshooting on page 1377
- Error messages for troubleshooting on page 1382

Some common alert and log messages that you might see in the Logs view are useful for troubleshooting.

Alert log messages for troubleshooting

Alert log messages provide useful information for troubleshooting.

Respond to Log spool filling alerts

The “Log spool filling” alert indicates that logs are not being transferred from the engine at all, or the engine is generating logs more quickly than they can be transferred to the Log Server.

Follow the instructions in the topics about troubleshooting log storage.

Related tasks

Troubleshoot log storage on page 1394

Respond to alerts about inoperative Security Engines

The Inoperative Security Engines alert is triggered when the Management Server does not receive the expected status updates from an Security Engine. The Status Surveillance option must also be selected for the engine.

If you see these alerts, one of the problems listed here might exist or might have existed temporarily:

- The connection between the engine and the Management Server might have been lost due to network connectivity problems or due to a technical issue on the Management Server. Problems that affect only management communications do not interfere with the operation of the engines - the engines continue processing traffic.
- The engine might be experiencing technical problems.

A console connection to the affected engine is recommended, if possible, when you suspect that the engine might not be operating properly. Connecting to the engine allows you to see any possible error messages printed out to the console before you take corrective actions, such as rebooting the node.

Steps

1) Check if the status of the engine or the system connections (shown in the info view when the engine is selected) still shows problems.

2) Check if there is a steady log stream from the affected engine. Also check if there are any further alerts or logs from the engine that could explain the reason for the message.

3) Check if the engine is actually processing traffic normally even if the Management Server is not able to monitor the engine and show the log stream.

4) If you suspect technical problems on the engine, run the `sginfo` script on the engine before rebooting it (if possible) and contact Forcepoint Customer Hub.

Related reference
Forcepoint Security Management Center commands on page 1429

Respond to System Alerts

"System Alert" is a general category for alert messages that are generated because something in the operation of the SMC components or Security Engines requires your attention.

Steps

1) Select the alert entry in the Logs view and click **Details** in the toolbar to view the alert entry information.

Respond to Tester alerts

Tester alerts indicate that the automatic tester running on the engines has detected a failure in one of the tests that the tester is configured to run.

The Tester is configured to run some tests by default. You can add more tests in the Engine Editor for each engine.

Steps

1) Make sure the condition that caused the test to fail is resolved.

Related concepts
Getting started with the Security Engine tester on page 637

Respond to throughput-based license exceeded alerts

Throughput-based license exceeded alerts are triggered traffic exceeds the throughput limits allowed by your license.

Some licenses limit the throughput of the engine to a certain fixed value. If the throughput limit is reached at any particular moment, the exceeding traffic is dropped and an alert is created to notify you how many packets have been dropped. The throughput limit is counted as the total throughput of all traffic handled by the engine at any one moment. All traffic is taken into consideration, regardless of type, direction, or the links used. Usually, temporary spikes in traffic trigger the messages, and the messages do not cause major problems. If you see these messages often, you should take action.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Make sure that your appliance licenses are matched to the correct elements according to the type of appliance. If the license generated with the POS code of a lower-throughput appliance is applied to a higher-throughput appliance, the throughput is needlessly limited.

2) If the hardware can handle a higher throughput than what it is licensed for, you can switch to a higher-throughput license (contact your reseller).

3) If the license throughput corresponds to the maximum throughput achievable with your hardware, you might be able to install an additional cluster node. Alternatively, you can switch to hardware with a higher maximum throughput (contact your reseller).

4) You can restrict the traffic in a more controlled way, for example, using the traffic management features of a Engine. See *Getting Started with QoS* for more information.

Related concepts

Quality of Service (QoS) and how it works on page 973

Log messages for troubleshooting

Log messages provide useful information for troubleshooting.

Respond to Connection closed abnormally log messages

Logs might contain the message “connection closed abnormally” if the connection closing does not occur in the expected order of a normal TCP connection.

Connection closed by client/by server or connection reset by client/by server messages in logs inform you about an event that the Engine has detected in the network. Connection closing is an expected event at the end of

each standard TCP connection. The logging settings in the Access rules determine whether connection closing is logged. Frequent abnormal connection closing or resets might indicate problems in the network, such as an overloaded server.

Steps

1) If connection problems affect services, perform standard network troubleshooting steps along the whole communication path.

Related concepts
Traffic captures and how they work on page 249

Respond to Connection removed during connection setup log messages

The “Connection removed during connection setup” message in logs notifies you that a connection was abnormally cut during the TCP connection setup phase because of an RST (reset) sent by one of the communicating parties.

Steps

1) Check why the connection is torn down, for example, if the server is overloaded or if the service is down or using a non-standard port.

Respond to Connection state might be too large log messages

Logs that contain “connection state might be too large” messages indicate problems with synchronizing state information between nodes in a Engine Cluster.

Problem description: You see “error when serializing for state sync” messages with a “connection state might be too large” clarification log entries for a Engine Cluster. You might also experience intermittent or continuous problems with clustering and traffic flow, which are typically alleviated for some time by rebooting all clustered nodes.

Reason: The Engine keeps a record of all connections that are handled statefully to be able to track the connection. When the Engine is clustered, this connection table must be synchronized between the nodes to allow connections to continue if a node goes down. When the state table grows excessively large, the Engines can no longer effectively use it.

A misconfiguration usually causes this message. Typical configuration problems include:

- Using the Oracle Protocol Agent on the actual database connections between the client and the server. The Oracle Protocol Agent is meant for cases where TCP port 1521 is used only for negotiating the port number for Oracle database connections. The port number for the actual connection is assigned dynamically. The Oracle Protocol Agent must not be used in any other cases.
- Excessive idle timeouts defined in Access Rules. All TCP connections are normally explicitly closed by the communicating parties and can therefore be cleared from the state table based on actual connection state. Non-TCP protocols do not establish connections. The communications are still handled as virtual connections

on the Engine to allow all Engine features to be used on the traffic. Because the communicating parties do not have a closing mechanism, these virtual connections are never cleared from the Engines' connection records before the communications are left idle (unused) for the duration of the defined timeout. If Access rules define excessively long timeouts for such traffic between many different hosts, the connection state table can grow very large.

Steps

1) If you use the Oracle Protocol Agent, make sure that it is not applied incorrectly.

   If necessary, replace the default service that has a Protocol Agent attached with a custom service that matches the correct port without a Protocol Agent.

2) Check the Access Rules to see if there are rules that override the default idle timeout value for non-TCP traffic.
   - Make sure that the override is not applied to any traffic that does not absolutely need a longer timeout (make the rule as specific as possible)
   - Try reducing the timeout (generally, the idle timeout should not be more than a few minutes).
   - In some cases, allowing both communications directions separately might remove the need for long timeouts

Respond to Connection timeout log messages

Connection timeout log messages are generated for inactive connections that the Engine clears out from its connection tracking table.

Connections are inactive when the hosts involved in the connection stop transmitting packets between each other.

Most connection timeouts are normal and necessary to ensure that the Engine cleans up inactive connections from its records, freeing up the resources. However, sometimes the timeout can prevent communications from continuing.

Steps

1) If some application in your network leaves connections inactive for long periods of time before continuing again, you can increase the timeout for those connections. You can change the timeout in the Action options for the Access rule that allows the connection. The rule-specific timeouts override the global timeouts that are set per connection state in the Engine element's properties (Advanced Settings).

   CAUTION

   Setting long timeouts for a high number of connections considerably increases the resource consumption of the Engine and can even lead to performance problems. This issue applies especially to non-TCP protocols that do not include connection closing messages, because such virtual connections are never closed before the timeout is reached.

2) If the protocol is not connection-oriented (for example, if the protocol is SNMP), you can disable connection tracking for the traffic in the Access rule's Action options. Disabling connection tracking requires that you explicitly allow both directions of the communications in the rule because without connection tracking, reply packets cannot be automatically allowed. NAT rules are not applied to connections that are not tracked. We recommend that you deactivate logging in rules that have connection tracking off because these rules create a separate log entry for each packet transmitted. The number of log entries generated greatly increases, and can potentially lead to an unmanageable level of logging traffic.

Respond to Incomplete connection closed log messages

Logs that contain “incomplete connection closed” messages indicate that a Engine determined that a connection was unsuccessful and removed it from its records.

“Incomplete connection closed” messages are shown in logs when the Engine allows a connection and passes the first packet of a connection (the SYN packet), but the reply packet (SYN/ACK) from the destination host does not arrive at the Engine.

One of the following situations can cause the connection to be incomplete:

- The SYN packet did not reach the destination.
- The SYN packet reached its destination, but the destination host did not send any reply.
- The SYN packet reached its destination and the destination host replied, but the reply packet did not reach the Engine.

It is normal to see a few of these messages in the log from time to time. However, a higher number of these messages can indicate problems in your network or the communicating applications.

Steps

1) If this message appears in the logs often for legitimate traffic, there is a networking problem that you must address. Use normal network troubleshooting tools to find out where the packets are lost. You can generate a tcpdump file by taking a Traffic Capture from the SMC.

2) In some cases, SYN packets can be sent maliciously to random hosts as an attempt to find out your network structure. These attempts can sometimes be seen as SYN packets to hosts that do not exist. If access to those addresses is allowed and routable, this process can trigger the Incomplete Connection Closed messages. The possibility of successful scans can be reduced by using dynamic NAT on the Firew all.

Related concepts

Traffic captures and how they work on page 249
Getting started with NAT rules on page 851

Respond to NAT balance log messages

Logs that contain NAT balance messages indicate that connections were dropped when the Engine tried to forward the connections after applying NAT.

NAT balance messages are shown in the logs when a connection has been allowed, the Engine has applied a NAT rule that defines source and/or destination translation, and the traffic has been forwarded according to the Engine’s routing configuration, but a reply is never received.

Steps

1) If NAT is applied to the connection in error, adjust your NAT rules accordingly. It is also possible to create a NAT rule that defines no translation to disable NAT for any matching connection.

2) Make sure that the Engine routes the traffic correctly. The routing decision is made based on the translated destination IP address.

3) Make sure that the destination host is up and providing the requested service, and that any intermediary Engine allows the connection.

4) Try to trace the path that the communications take and use traffic captures as necessary to find the point of failure.

Respond to Not a Valid SYN Packet log messages

Logs that contain “Not a Valid SYN Packet” messages indicate that packets were discarded due to connection tracking.

Problem description: The “Not a Valid SYN Packet” message appears in logs with entries on discarded packets.

Reason: “Not a Valid SYN Packet” is a TCP packet that is not the first packet of a TCP connection (the packet does not have the SYN flag set), but is not part of an existing connection either (there is no connection tracking entry on the Engine matching this packet). The policy would allow this packet if the packet was part of an existing tracked connection.

The message usually also contains a code inside square brackets that indicates the flags set in the discarded packet (A=Ack, F=FIN, R=RST, P=Push, S=SYN).

Some examples of situations, where “Not a Valid SYN packet” messages can be seen:

- Asymmetric routing, which means that the opening packet does not go through the Engine, but the reply (the SYN/ACK) does. Asymmetric routing can indicate that there is a configuration error in the routing of the surrounding network that must be fixed.
- Connections that are idle for more than the defined connection timeout (connection has been erased from the Engine records). If necessary, you can increase the timeout.
- Connections that have been made to look like TCP connections even though they are not. If necessary, you can allow these connections as individual packets without connection tracking.
- Network scans or attacks that use ACK packets.
- Heavily loaded server or client that sends a packet after the host at the other end of the connection has already timed out and closed the connection.

It is normal to see some messages like this in the logs. If a certain type of communication that you want to allow is always prevented because of connection tracking, check these troubleshooting steps.

Steps

1) If there are connections that are left idle for a long time, you can change the idle timeout value for the Access rule that allows that specific traffic. There are also default values that you can set globally for different TCP connection states in the Engine element’s properties (Advanced settings). This solution does not usually apply to non-TCP connections, so take care that the rule only matches the specific connections involved.



CAUTION

Setting long idle timeouts for a high number of connections considerably increases the resource consumption of the Engine and can even lead to performance problems. Especially, non-TCP protocols do not include connection closing messages, so such virtual connections are never closed before the timeout is reached.

2) You might have to disable connection tracking in the rule allowing the connection. We recommend that you deactivate logging in rules that have connection tracking off. These rules create a separate log entry for each packet transmitted, which increases the number of log entries generated. NAT cannot be applied to traffic that is allowed without connection tracking, and both communication directions must be explicitly allowed in the Access rules (replies are not automatically allowed).

3) For some types of connections, problems can be solved by using a Service that includes a special Protocol Agent for that traffic.

Related concepts
Protocol elements and how they work on page 931

Respond to Requested NAT cannot be done log messages

Logs that contain “Requested NAT cannot be done” error messages can indicate problems with dynamic NAT or Server Pools.

Steps

1) A Dynamic NAT operation might be applied to the wrong type of traffic. Dynamic (many-to-one) NAT is done by assigning different hosts the same IP address, but different ports. For this reason, dynamic NAT does not work when the protocol in question does not use ports. Only the TCP and UDP transport protocols use ports. See the **TCP** and **UDP** branches in the **Services** tree in the SMC Client to check which protocols are transported over TCP or UDP.

2) Dynamic NAT can run out of ports if there are too many simultaneous connections in relation to the IP addresses and the port range you have configured for dynamic NAT. You can increase the available ports for translation by adding a new IP address for your dynamic NAT rule. Alternatively, you can expand the port range, if the rule does not currently use the whole range of high ports.

3) If the Server Pool element is used, check the NAT rules. Because the Server Pool element always does NAT, errors can occur when the Server Pool element is used and the same connection matches an overlapping NAT rule.

4) Check if the information message in the log states that dynamic NAT is denied due to excessive number of connections. This can happen when a single host is opening connections at an excessive rate to a single destination IP address and port through dynamic source NAT. This message indicates the triggering of a self-protection mechanism, which prevents excessive use of processing resources to dynamic NAT operations. Set up a static NAT rule to allow these types of connections if it is not possible to adjust the connection settings of the application.

Error messages for troubleshooting

Error messages provide useful information for troubleshooting.

Respond to Command Failed or Connect Timed out errors

There are several possible causes and solutions when to command an engine using the SMC Client and receive an error.

Steps

1) Make sure that there is connectivity between the Management Server and the engine.

   The most common connectivity problems include traffic filtering by an intermediate engine and incorrectly configured NAT configuration. When NAT is configured incorrectly, a required NAT rule for the connection is missing, or the engine's contact address is missing. The engine sends its status reports through the Log Server, so a green operating status does not guarantee that the Management Server can reach the engine.

2) Try to refresh the engine's policy. Read all messages displayed and make sure that none of the nodes perform a rollback to the previous policy. If policy installation fails, see *Troubleshooting Policy Installation*.

Related tasks

Troubleshoot policy installation on page 1407

Respond to Unexpected Error messages

In some cases, the SMC Client might display a pop-up message stating that an "unexpected error" has occurred.

Steps

1) Exit and restart the SMC Client.

2) If the condition persists, contact Forcepoint Customer Hub.

Chapter 91

Troubleshooting Security Engine operation

Contents

- Troubleshoot Security Engines that do not go or stay online on page 1385
- Troubleshoot errors when commanding Security Engines on page 1386
- Troubleshoot heartbeat and synchronization errors on page 1387
- Troubleshoot contact between Security Engines and the Management Server on page 1387

There are several common errors and problems that are directly related to the operation of Engines, IPS engines, and Layer 2 Engines.

Troubleshoot Security Engines that do not go or stay online

There are several possible causes and solutions when you command a node online, but it does not go online or turns itself offline shortly after going online.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) If you have just updated the engines or you are using an evaluation license, open the **Administration** view and browse to **Licenses** to check that your licenses are valid for the version of engines you are using.

2) If the nodes are in a cluster, and only one node at a time stays online, try the following troubleshooting steps:
   - Check whether the cluster is in Standby mode. Standby mode keeps one node online at a time and uses the other nodes as backups in case the online node fails.
   - Refresh the policy of the cluster and check that the installation is successful so that no nodes roll back to the previous configuration. All nodes in the cluster must have the same configuration that has been installed in the same policy installation operation. You might have to adjust the rollback timeout in the cluster's properties if policy rollback on some node is the problem.
   - Check for alerts in the Logs view about tests failing. Check if any of the failed tests are configured to turn the node offline when they fail. The tester leaves one node in a cluster online even if the test fails on all nodes. If you see a test failure, it might indicate a genuine problem that you need to solve or the test might be misconfigured and might have to be disabled or reconfigured.

3) See the Logs view for any alerts or logs regarding the functioning of the nodes.
   Certain internal error conditions (for example, heartbeat connection failures or missing certificates) can cause nodes to go offline. These events are shown as logs and alerts.

Related concepts
Getting started with the Security Engine tester on page 637
Adjusting Engine clustering options on page 676

Related tasks
Adjust IPS clustering options on page 679
Generate licenses on page 1333
Troubleshoot policy installation on page 1407

Troubleshoot errors when commanding Security Engines

There are several possible causes and solutions when you try to command an engine using the SMC Client and receive an error.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) If there are disabled nodes in the same cluster, enable them before refreshing the policy.
2) Refresh the engine's policy.
3) Read all messages displayed and make sure none of the nodes roll back to the previous policy.
4) If policy installation fails, see *Troubleshooting Policy Installation*.
   If policy installation fails and the status and statistics information displayed in the SMC Client is missing (status is displayed as unknown) even though the engine is online and processing traffic, the problem might be due to an expired certificate.

Related concepts
How certificates work on page 149

Related tasks
Troubleshoot policy installation on page 1407

Troubleshoot heartbeat and synchronization errors

There are several possible causes and solutions when you receive alerts regarding the heartbeat connection between the nodes in a cluster.

Steps

1) Apply normal network troubleshooting (for example, check speed and duplex settings and cabling) to make sure that the heartbeat link works reliably.

2) Use a primary and a backup heartbeat connection using separate physical links.

   It is highly recommended that you use a dedicated link for both the primary and the backup heartbeat. The heartbeat and state synchronization are time-critical communications. The heartbeat connection is critical to the operation of a cluster. The cluster cannot work without a reliable heartbeat connection.

3) If you have installed two or more clusters with a single LAN as a shared heartbeat, and you see extra log entries about unauthenticated heartbeat messages, change the Heartbeat IP and the Synchronization IP so that each cluster uses a different address.

Related concepts

Traffic captures and how they work on page 249
Adjusting Engine clustering options on page 676

Related tasks

Adjust IPS clustering options on page 679
Adjust Layer 2 Engine clustering options on page 680

Troubleshoot contact between Security Engines and the Management Server

Sometimes, the engine cannot establish initial contact to the Management Server, or all subsequent attempts to command the engine through the SMC Client fail. There are several possible causes and solutions for these failures.

For a full list of all system communications in all configurations, see *Default Communication Ports*.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Apply normal network troubleshooting (for example, check speed and duplex settings and cabling).

2) If there is a local Engine between a remote site Engine and the Management Server, make sure that the local Engine does not block the communication.
A Engine with reversed management connections (for example, because it has a dynamic IP address) contacts the Management Server on port 8906. Create an Access rule in the policy of the main site Engine to allow the connection:
   - **Source**: Remote site Engine
   - **Destination**: Contact address of the Management Server
   - **Service**: SG-dynamic-control
   - **Action**: Allow

Related concepts
Traffic captures and how they work on page 249

Related reference
Forcepoint Security Management Center ports on page 1457

Chapter 92

Troubleshooting licenses

Contents

- Problems with licenses on page 1389
- Troubleshoot licenses that are shown as retained on page 1389
- Troubleshoot licenses that are shown as unassigned on page 1390

Licenses are a proof of purchase used for ensuring that your organization is a legal license holder of the software.

Problems with licenses

Components that do not have a license do not work.



Note

Do not confuse Licenses with Certificates. Licenses are a proof of purchase used for ensuring that your organization is a legal license holder of the software. Certificates are proof of identity that components use to authenticate themselves in system communications.

Licenses are introduced to the SMC as elements that you install in your Management Server. You must install licenses to set up the components your organization has purchased. However, on current Forcepoint Network Security Platform appliance models, the licenses can be generated automatically.

If the Management Server does not have a valid license, you see a license-related dialog box each time you log on using the SMC Client. You cannot create a working configuration without a Management Server license because most controls are disabled. If an engine component is missing a license, you can create a configuration for it, but you cannot transfer that configuration to the engine.

Related concepts

Getting started with licenses on page 1331

Troubleshoot licenses that are shown as retained

There are several possible solutions when the State of a license is shown as “retained” in the SMC Client.

Licenses can be generated based on the Management Server's proof-of-license (POL) code, the appliance’s proof-of-serial (POS) code, or a fixed IP address. When you start using a Management Server POL-based license, you bind it to the correct component, and the binding is fixed when you install the component's policy. Because of this fixed binding, if you unbind a Management Server POL-based license, it is shown as “Retained.”

In this case, it is not possible to reuse the Management Server POL-based license on another component, or delete the Management Server POL-based license.

You can resolve issues with licenses shown as "retained" in the following ways:

- You can rebind the license to the component it was previously bound to.
- You can replace the Management Server POL-based license with a new Management Server POL-based license.
- You can replace the Management Server POL-based license with a new IP-address-bound license or POS-bound license.
- You can delete the element to which the Management Server POL-based license is bound. After the element is deleted, the state of the Management Server POL-based license changes to **Unassigned** and it can be rebound to some other component or deleted.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) To rebind the license to the component it was previously bound to, right-click the license and select **Cancel Unbind**.

2) To replace the Management Server POL-based license with a new Management Server POL-based license, do the following:
   a) Install a new Management Server POL-based license.
   b) Bind the license to the component that the retained license is bound to.
      The state of the previous Management Server POL-based license changes to **Unassigned** and it can be rebound to some other component or deleted.

3) To replace the Management Server POL-based license with a new IP-address-bound license or POS-bound license, do the following:
   a) Install a new IP-address-bound license or POS-bound license.
   b) Right-click the Management Server POL-based license and select **Replace by Static License**.

Troubleshoot licenses that are shown as unassigned

When you install a new license, it does not bind itself to any component and is shown as “unassigned.”

The license does not contain identifying information that the Management Server could use to attach the license to a component. This is normal for Management Server POL-based licenses.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Right-click the license and select **Bind** to select which component you want to license.

Related tasks

Check validity and status of licenses on page 1343

Chapter 93

Troubleshooting logging

Contents

- Troubleshoot the Logs view on page 1393
- Troubleshoot log storage on page 1394
- Troubleshoot Log Server operation on page 1395

There are some common problems you might encounter when viewing logs or performing tasks related to the log files.

Troubleshoot the Logs view

There are some common problems and solutions related to viewing logs.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) If the Logs view is unable to contact some Log Server, do one of the following:
   a) Check that the Log Server is running, and that the Log Server is reachable from the computer used for running the SMC Client. The status shown for a Log Server in the SMC Client is based on information that the Management Server has. However, logged data is not routed through the Management Server, so a green status is not an indication of whether the Log Server is reachable for log browsing or not.
   b) If there is a NAT device between some SMC Client and a Log Server, administrators must select the correct Location for the SMC Client in the status bar at the bottom right corner of the SMC Client window.
2) If some or all logs are not visible in the Logs view, do one of the following:
   a) Check that the filtering, time range, and source settings in the **Logs** view are correct. Also make sure that you have clicked **Apply** after making the latest changes in these settings.
   b) Check the logging options in your policies. Not all connections automatically create log entries. The Alert, Stored, and Essential options create permanent log entries. The Transient option means that logs are not stored, and they can only be viewed in the Current Events mode in the Logs view when they reach the Log Server.
   c) Check that logs you want to keep are not being pruned. The Log Server deletes selected logs according to how pruning is configured.
   d) Check that the logs are being transferred from the engines to the Log Server. The log entries are spooled on the engines if a connection to the Log Server is unavailable. Connections between engines and the Log Server should be shown as green in the **Dashboard** view.

e) The logging process might slow down due to a lack of resources on the engine, in the network, or on the Log Server. Your logging settings and the number of engines that send data to the same Log Server affect the speed of the logging process.

Related concepts
Considerations for setting up system communications on page 125

Related tasks
Discard unnecessary logs on page 1313

Troubleshoot log storage

There are several possible causes and solutions when alerts indicate that the disk of an engine or a Log Server is filling up with log files.

The “Log spool filling” alert indicates that logs are not being transferred from the engine at all, or the engine is generating logs more quickly than they can be transferred to the Log Server.

Steps

1) If an engine is filling up with logs, do one of the following:
   a) Check that the Log Server is running. If it is not running, try to start it. If the Log Server is running, check for network problems between the engine and the Log Server. The log entries are spooled on the engines if they cannot be sent to the Log Server. Stopping and restarting the Log Server process can help in resetting the connection.

   b) If the volume of logs is high, they might not be transferred quick enough, and logs must be spooled even though they are being transferred. If you suspect this is the case, turn off all diagnostics logs for all engines that you are not actively troubleshooting. Also turn off logging for all rules that have connection tracking set to off (because these rules log each packet individually). Finally, check if logs that are currently pruned could be prevented from being generated in the first place.

2) If the Log Server is filling up with logs, do one of the following:
   a) Set up log management tasks that archive and remove the oldest Log entries from the Log Server hard disk. To avoid problems in the future, set up tasks to run automatically at regular intervals.

   b) In an emergency, you can also move or delete old log entries manually. The logs are stored on the Log Server machine under a folder structure based on dates and times (default location is <installation directory>/data/storage). You should always avoid manual handling of the newest entries.

Related concepts
Log data management and how it works on page 1307

Related tasks
Enable or disable diagnostics on page 356

Troubleshoot Log Server operation

There are several possible causes and solutions when the Log Server is not running and does not stay running when you try to start it.

Steps

1) On the Log Server command line, try to start the Log Server using the script <installation directory>/bin/sgStartLogSrv[.bat|.sh] to get more information.

   Startup messages are shown on the command line.

2) Check for the following possible problems and solutions:

   - The software version might be incorrect. The Log Server must have the exact same software version as the Management Server. Upgrade components as necessary.
   - The Log Server's certificate for system communications might have expired, been deleted, or become otherwise invalid.
   - The Log Server might not have a license, the license might be bound to the wrong IP address, or the IP address the license is bound to might not be active on the server.
   - There might not be enough space for logs on the hard disk.
   - Sometimes files that are necessary for the Log Server to run might be moved or deleted by external processes, such as a malware scanner, or lost due to hard disk errors. In these types of cases, Java shows an error message stating "Could not find the main class." Reinstall the same software version and check the configuration of the host computer to prevent the same from occurring in the future.

Related concepts
Renewing certificates on page 162
Problems with licenses on page 1389

Related tasks
Troubleshoot log storage on page 1394

Chapter 94

Troubleshooting the SMC Client

Contents

- Troubleshoot disabled options in the SMC Client on page 1397
- Troubleshoot slow SMC Client startup and use on page 1398
- Troubleshoot logging on to the SMC Client on page 1398
- Troubleshoot SMC Client layout and views on page 1399
- Troubleshoot missing or incomplete statistics on page 1400
- Troubleshoot status monitoring on page 1400
- Troubleshoot Management Server commands on page 1401

There are several general problems that you might encounter when using the SMC Client.

Troubleshoot disabled options in the SMC Client

There are several possible causes and solutions when some SMC Client options are disabled.

When Management options are disabled, you can see them but they are grayed out and you cannot change them.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Make sure that the IP address in the license is correct and that the network interface with that IP address is active on the Management Server.

   If nearly all controls are disabled, the Management Server license might not be installed or valid.

2) Make sure that your administrator permissions allow you to use the SMC Client options.

   Your administrator permissions might restrict your actions. For example, Operator administrators are only allowed to view the contents of the elements granted to them. They are not allowed to change the elements' properties. If you want to edit an element that is locked, first unlock the element.

3) Check whether the disabled options are part of an element or object created by the installation. Select the element in the **Configuration** view and check the **History** tab in the Info pane.

   System elements display "System" as the Creator. System objects cannot be edited, but usually you can make copies of them and edit and use the copies instead.

4) Check whether the options are part of a feature that is not supported by the engine version you are using.

   If so, upgrade the engines before you can use the feature. See the Release Notes for information about new and enhanced features in different versions.

5) Check whether you must select some other option before you can change the disabled options. Look for an option to activate the feature or override the default values, or a way to change the default values instead.

Related concepts
Getting started with administrator accounts on page 373

Related tasks
Lock elements on page 197
Unlock elements on page 198
Generate licenses on page 1333

Troubleshoot slow SMC Client startup and use

There are several possible causes and solutions when SMC Client logon and other operations are slow.

For example, refreshing a list of network elements might take a long time.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Make sure the Management Server and Log Server host names can be resolved on the computer running the SMC Client (even if the SMC Client is running on the same computer as the SMC server components). Add the IP address and host name pairs into the local hosts file on the client computer:
   - In Linux: /etc/hosts
   - In Windows: \WINNT\system32\drivers\etc\hosts

2) Exclude temporarily unavailable Log Servers from Sources. If there is a NAT device in between your client and a Log Server, also make sure that the Location is set correctly for your SMC Client (in the status bar).

   Log browsing is slowed down if some of the selected Sources (Log Servers) are not available (data is displayed only after queries time out).

Troubleshoot logging on to the SMC Client

If you have trouble logging on to the SMC Client, check the connectivity, certificate, and password.

Steps

1) If the SMC Client reports a connection problem:
   a) Make sure you entered the address correctly on the logon screen.
      If there is a NAT device between the SMC Client and the Management Server, the IP address used must be the translated address. Locations and Contact Addresses are not used for selecting the correct address when you are just logging on.
   b) Make sure that the Management Server is running.
   c) Make sure that the network between the SMC Client and the Management Server is working and routing the traffic correctly. Make sure that the correct ports and protocols are used for the communications.
      These communications are not allowed in the predefined Firewall Template for Engines.
2) Make sure that the Management Server certificate has not expired.
3) Make sure that your password is correct.

Related tasks
Replace forgotten passwords on page 1371
Recertify SMC servers on page 162

Related reference
Forcepoint Security Management Center ports on page 1457

Troubleshoot SMC Client layout and views

There are some common problems and solutions related to the SMC Client layout.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) If a view that was previously visible is now missing:
   a) Select the appropriate view from the **Menu** > **View** menu. If the view is not listed, it is not available in this window.
   b) If the view is still not visible, select **Menu** > **View** > **Layout** > **Reset Layout**.
2) If all views that you want to be displayed on the screen are visible, but you want to change the layout:
   - You can drag and drop the views by their titles to different predefined positions on the screen and resize them by dragging the borders.

- You can reset the layout to the default positions, select **Menu** > **View** > **Layout** > **Reset Layout**.

Troubleshoot missing or incomplete statistics

There are some common causes and solutions when some or all types of statistics are not visible.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) If you can see some types of statistics for an element, but not other types:
   - This is normal. Not all statistics are compatible with all element types.

2) If an element does not show any statistics at all:
   a) If it is a server element, this is normal. Servers do not produce statistical information. Only engine elements do.

   b) Check that the Log Server is selected correctly in the Management Server properties and that the Log Server is running. Also check that the network connections are up between the engines, the Log Server, and the Management Server. Statistical information is sent from the engines to the Log Server, which relays the information to the Management Server.

Troubleshoot status monitoring

There are some common problems and solutions related to viewing the status of elements in the SMC Client.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) If a status icon for an element is always white, turn on status monitoring for the element.
   A white status icon indicates that status monitoring for the element is turned off.

2) If you do not know what the status icon means, place the cursor over the status icon and wait for a tooltip to show the element's status as text.

3) If the status of the Management Server indicates problems, but the engine is processing traffic normally, check network connectivity.
   Status information is normally sent from the engines to the Log Server, which then relays the information to the Management Server. There is also a backup channel directly from the engines to the Management Server.

4) If no information is available on the **Status** tab in the **Info** pane, click **Refresh View**.
By default, the appliance status information is not displayed automatically.

Related concepts
Getting started with monitoring the system on page 211

Related tasks
Monitor tasks on page 248

Related reference
Forcepoint Security Management Center ports on page 1457

Troubleshoot Management Server commands

Commands sent to the Management Server might fail, or the status of the Management Server might indicate a replication error.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) If commands sent to the Management Servers using the **SMC Client HA Administration** dialog box fail, an error message indicating that the SMC Client is unable to connect to a Management Server might appear.
   a) Make sure that the SMC Client Location is selected correctly so that the SMC Client connects to the Management Server at the correct Contact Address.
   b) Make sure that the Location and Contact Address are correctly configured for the Management Server.
2) If the Status of the Management Server on the General tab in the Info pane reads "Replication error":
   a) Switch to the Replication tab in the Info pane for more information.

Related tasks
Define Management Server or Log Server contact addresses on page 128
Select the Location for the SMC Client on page 142

Chapter 95

Troubleshooting NAT

Contents

- Problems with NAT and possible causes on page 1403
- Troubleshoot NAT that is not applied correctly on page 1404
- Troubleshoot NAT that is applied when it should not be on page 1405

There are some common problems you might encounter with NAT.

Problems with NAT and possible causes

There are several possible causes and solutions for problems with NAT.

Consider the following when troubleshooting NAT issues:

- A Dynamic NAT operation can be applied to most types of traffic. For TCP and UDP connections, dynamic (many-to-one) NAT is done by assigning different hosts the same IP address but different ports, so that the subsequent replies can be recognized and forwarded correctly. For ICMP connections, different hosts are also assigned the same IP address. The ICMP ID in the packets is used to recognize the replies and to forward them to the correct recipients. Only the TCP and UDP transport protocols use ports. See the **TCP** and **UDP** branches in the **Services** tree in the SMC Client to check which protocols are transported over TCP or UDP. Dynamic NAT is suitable when different internal hosts communicate with different external hosts. If a single internal host communicates with a single external host, we recommend the use of static NAT, especially if the number of simultaneous connections is high.
- Dynamic NAT can run out of ports if there are too many simultaneous connections in relation to the IP addresses and the port range you have configured for dynamic NAT. You can increase the available ports for translation by adding a new IP address for your dynamic NAT rule. If the rule does not currently use the whole range of high ports, you can also expand the port range. The number of simultaneous NATed connections equals the number of IP addresses multiplied by the number of ports.
- Check the NAT rules for configurations that overlap with the following NAT configurations:
  - Address translation configured in an Outbound Multi-Link or Server Pool element
  - A NAT pool defined for VPN clients in the Engine element's properties
  - Element-based NAT

  Errors can occur when one of the listed elements is used and the same connection matches an overlapping NAT rule, because the elements also use NAT. Only one address translation operation can be done for each packet and overlapping configurations can cause conflicts. Overlap within the NAT rules is allowed because the rules are resolved based on their order (first matching rule is applied). If you use element-based NAT, a more specific manually created NAT rule can prevent traffic from matching the automatically generated NAT rules.
- Check that the NAT configurations do not overlap with an IP address that is used by some physical host in the network. This configuration error is most common with source address translation for a DMZ or external IP address. Overlapping NAT configurations can create conflicts between IP addresses and other hosts in the network.

Troubleshoot NAT that is not applied correctly

Resolve problems when NAT is not applied at all or is applied incorrectly.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Make sure that any connection that you want to NAT is allowed by an Access rule that has Connection Tracking enabled.

2) If the target of the translation is traffic that is entering or exiting a VPN tunnel, enable address translation for traffic transmitted over that VPN in the properties of the VPN element.

   The default setting is to disable all address translation for tunneled VPN traffic. The setting affects only traffic inside the VPN tunnel, not the tunnel itself (the encrypted packets).

3) If traffic is not translated at all or the wrong translation is applied, check the NAT rules:

   a) Search the rules using the original (before translation) source and destination addresses and check if the traffic matches the wrong NAT rule higher up in the rule table. Only the first matching rule is considered. Note that NAT rules with an empty NAT cell are valid and specify that addresses are not translated for matching traffic.

   b) In addition to NAT rules, NAT is also used in NetLink or Server Pool elements, and as a NAT pool defined for VPN clients in the Engine element's properties. There must not be overlapping NAT rules that match the same connections.

   c) NAT rules are automatically generated from NAT definitions that are added to an element's properties. The NAT rules that are generated from NAT definitions do not override the rules that have been manually added to the Engine policy. However, a more specific manually created NAT rule can prevent traffic from matching automatically generated NAT rules.

Related concepts

Element-based NAT and how it works on page 633

Troubleshoot NAT that is applied when it should not be

Resolve problems when the Engine translates an IP address to some other IP address even though it should not.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Check the order of the NAT rules.

   The Engine reads the NAT rules from top to bottom. Only the first rule that matches is considered, so you can make exceptions to rules by placing a different, partially overlapping rule above. Leaving the **NAT** cell empty tells the Engine that addresses in any connections that match the rule should not be translated.

2) Check for other configurations that apply NAT.

   For VPN traffic, you can also enable and disable address translation for all traffic transmitted over a VPN in the properties of the VPN element. The default setting is to disable all address translation for tunneled VPN traffic. The setting affects only traffic wrapped inside the VPN tunnel, not the tunnel itself (the encrypted packets).

   In addition to NAT rules, NAT is also used in NetLink or Server Pool elements, and as a NAT pool defined for VPN clients in the Engine element's properties. There must not be overlapping NAT rules that match the same connections. For NetLinks, NAT rules are used to select traffic for balancing, and only the actual IP addresses used for the translation are defined in the NetLink elements. NAT is required for the operation of these features and you must exclude the connections in question from the scope of these features to disable NAT.

Related concepts

Getting started with NAT rules on page 851

Chapter 96

Troubleshooting policies

Contents

- Troubleshoot policy installation on page 1407
- Troubleshoot rules on page 1410
- Troubleshooting packets incorrectly dropped by antispoofing on page 1413
- Troubleshooting unsupported definitions in IPv6 Access rules on page 1414

There are some common problems you might encounter when working with policies and the rules that they contain.

Troubleshoot policy installation

There are several possible causes and solutions when policy installation fails, or there is a warning message indicating problems in the policy installation window.

Troubleshoot unintended policy rollback

Resolve problems when policy installation results in a rollback to the previously installed policy version.

Problem description: The policy installation reports that the Management Server can contact the engines and installs the new policy successfully. However, the policy installation results in a rollback to the previously installed policy version.

Reason: The rollback is a safety mechanism that prevents changing the engines' policy in ways that cut the connectivity between the engines and the Management Server. After each policy installation, the engine contacts the SMC using its new configuration and automatically reverts its policy if the contact does not succeed within a time-out period.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Make sure the policy and the configuration changes that you have made do not prevent communications between the Management Server and the engine.
   a) Check the IPv4 Access Rules and NAT rules (as applicable). You can also validate the policy to see if there are issues in it that prevent the policy installation. The rule search is useful for finding the first rule that matches the connections.
   b) Check the Routing. You can use the Route Query tool to check where the packets will be routed after a policy installation.

c) Check the Locations and Contact Addresses of the SMC components, which are required if NAT is applied to these system communications.

2) The rollback occurs after a timeout set in the engine element's advanced properties. If you are sure that there are no configuration or policy design issues, you can increase the timeout to allow for longer delays in contact. Increasing the timeout can help if the timeout is caused by poor network reliability or delays caused by processing a policy that is very large considering the engine's available resources.

Related concepts

Considerations for setting up system communications on page 125
Searching in rules on page 894

Related tasks

Check routes using the Route Query tool on page 709
Validate rules automatically on page 912

Troubleshoot policy installation failure due to connection timeouts

Resolve problems when policy installation fails because the connection between the engine and the Management Server times out.

The engine is up and running, but policy installation fails when the Management Server is contacting the nodes. When node-initiated contact is active, the Management Server might also wait for contact from a node, but the contact never happens.

The connection might time out for the following reasons:

- There is no network-level connectivity.
- The engine or the Management Server uses the wrong IP address.
- The engine and the Management Server reject each others' certificates.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Check for network problems, such as faulty or loose cables, mismatching speed/duplex settings, IP addresses, and routing.

2) Check the Locations and Contact Addresses of the SMC components, which are required if NAT is applied to these system communications.

3) In a cluster, all nodes that the Management Server tries to contact must be reachable and operational to install a policy on any of the clustered engines. If you have taken down an engine for maintenance, temporarily disable it to install the policy on the other cluster members.

4) If the problem seems to be related to certificates, you can recertify the engine to re-establish contact between the engine and the SMC.

5) Check the engine software version (shown in the Info pane when you select the element in the SMC Client). See the Release Notes for information regarding version compatibility between the engine and SMC software versions.

Related concepts

Considerations for setting up system communications on page 125
Management connections for Security Engines and how they work on page 629

Related tasks

Disable cluster nodes temporarily on page 360

Troubleshoot policy installation failure for other reasons

Resolve problems with policy installation that are not related to the connection between the engine and the Management Server.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) If the policy installation ends with an error, read the messages both in the main window and in the Issues tab and correct as necessary. Warnings do not prevent policy installation; you can still press **Continue** in the policy installation even if a warning is triggered. You can also validate the policy separately to see if there are issues that prevent you from installing it.

2) Try installing the policy with the **Keep Previous Configuration Definitions** option deselected. Normally, the option should be selected. Under certain conditions, the old configuration definitions might not be compatible with the new policy, so the engine cannot fulfill this request and policy installation fails.

CAUTION

When the policy is installed with the **Keep Previous Configuration Definitions** option deselected, even some currently active connections that are allowed in the new policy can be cut. The applications must then reopen the connections.

3) Make sure that a current dynamic update package is imported and activated on the Management Server.

Related tasks

Validate rules automatically on page 912

Policy installation warnings about ignored Automatic Proxy ARP options

Misconfigurations can result in warnings about ignored Automatic Proxy ARP options.

When installing a Engine policy, the “Automatic Proxy ARP option in NAT rule <rule tag> is ignored: none of the CVI interfaces are directly connected to the network in question” warning is shown when proxy ARP has been defined, but there is no matching CVI network configured in the Engine element. Automatic proxy ARP is used in NAT to handle ARP requests to the translated IP address for hosts in networks that are directly connected to the Engine. This warning can be due to an incorrect IP address or netmask setting, or the (not directly connected) Network in question missing from the Routing tree. It can also result from selecting the option for a NAT rule that involves an IP address for which the Engine cannot act as an ARP proxy.

Related settings can be configured in NAT rules, in a Server Pool element, and in the Engine Editor for the Engine element.

Related concepts

Configuring interfaces for Engines on page 547
Getting started with outbound traffic management on page 731
Getting started with NAT rules on page 851

Troubleshoot rules

There are several possible causes and solutions for problems with rules in policies.

Policy validation for troubleshooting rules

You can automatically validate a policy and check the rules for invalid configurations, for example, if policy installation fails.

Related tasks

Validate rules automatically on page 912

Troubleshoot traffic that is blocked even though rules allow ANY Service

The possible causes and solutions depend on the engine role when a connection that you want to allow is stopped.

In IPS policies, Access rules allow all connections by default. If a connection you want to allow is stopped because of an IPS Access rule, your Access rules contain a specific rule that stops these connections.

In Engine and Layer 2 Engine Access rules, even if you set the Source, Destination, and Service to ANY and set the rule to allow the traffic, certain connections might still be discarded.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Use a Protocol Agent to allow connections with a protocol that assigns ports dynamically.
   The Protocol Agent enables the Engine to track the assigned port.

2) Make sure that there is a matching rule with Continue as the action further up in the rule table with a Service in which the correct Protocol Agent is used if you want to use a Protocol Agent in a rule with ANY as the Service.

3) Add your own rules as necessary.
   The Firewall Template contains a rule that does uses a Protocol Agent for some, but not all protocols that use a dynamic port assignment.

4) If you must allow connections in your network for some application that implements TCP incorrectly, you might need to adjust or even disable connection tracking in the Access rules for those connections.
   Connections that violate the standard TCP connection sequence are dropped due to connection tracking. We recommend that you disable logging for rules that have connection tracking set to off, because such rules create a log entry for each packet.

Troubleshoot false positives in the Inspection Policy

There are several possible causes and solutions when the Inspection Policy produces alerts or terminates traffic that you consider to be normal.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) If a Situation is not valid in your environment under any conditions, change the action for the Situation to **Permit** on the **Rules** tab.

2) If a Situation is not valid between some hosts, add an Exception for the Source, Destination, and Situation that produce false positives. Then set the action to **Permit**. This editing can be done manually or based on a log entry through its right-click menu.

3) If a custom Situation produces false positives, adjust the parameters in the Situation to better match the traffic pattern that you want to detect.

Enable passthrough for PPTP traffic

You might need to specifically allow Point-to-Point Tunneling Protocol (PPTP) traffic if you use PPTP tunneling.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) To allow PPTP passthrough, add matching Access rules with the following two services:
   - The TCP Service for PPTP. The default Service element for PPTP uses the standard destination port 1723. Check the actual port used and create a Service with a different port, if necessary.
   - The IP Service for GRE (IP protocol 47).

2) Make sure that the GRE traffic is not matched against any dynamic NAT rule, including the dynamic NAT rule required to load-balance connections between NetLinks in a Multi-Link configuration.

   Use static NAT instead if IP address translation is required or configure the communicating applications to encapsulate the traffic in TCP or UDP (NAT traversal mode).

   Dynamic NAT cannot be applied because it uses ports to track connections using the same IP address. GRE works directly on top of IP and does not have the concept of ports, so it is not possible to do the same with GRE. It requires a static translation that forms a fixed one-to-one relationship between an original and translated IP address. Use a static IP address to IP address or network to same-size network mapping in the NAT rules.

   Even with static NAT, some PPTP implementations require extra setup (for example, encapsulation of the packets) to work correctly when IP addresses are translated.

Troubleshoot traffic that is incorrectly stopped by the engine

The Engine might incorrectly stop traffic that you want to allow. To check possible configuration errors, use log entries to locate the correct policy.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) In the Logs view, check whether the connection is logged.
   - Add a quick filter for both the source and destination IP address of the traffic you want to allow in the **Query** pane and click **Apply**.
   - If the logs show that the connection is discarded or refused by a rule, click the **Rule Tag** link in the log entry to check the rule.

2) Check the Access rules and NAT rules of the active policy for rules that match the same source, destination, and service.
   a) Open the **Search Rules** pane through the policy view's toolbar, then drag and drop the corresponding elements to the search fields at the bottom of the rule table.

   b) Select **Show Only Matching Rules** from **Options** in the search pane.

c) Deselect **Do Not Match ANY** from **Options** in the search pane.

d) If several rules are shown, the topmost rule is the one that is applied, unless the **Source VPN** cell (in IPv4 Access rules) has a definition that does not match. The other cells are not used for matching, but define options for what happens when traffic does match.

3) If the first matching Access rule is set to allow the traffic, check that other parts of the rule are correct:
   - Some protocols require the correct Protocol Agent, which is set by including the correct Service with the correct Protocol attached. In some cases, you might need to change the options of the Protocol Agent.
   - ANY rules do not use most Protocol Agents by default.
   - You can create new Services for any source or destination port or port range as needed.
   - The Connection Tracking Action options define if stateful inspection is used and how strict the checks are. Connection tracking allows NAT rules to be applied to the connection and a rule table where reply packets do not need to be separately allowed. The Engine checks that the communications follow the standards of the protocol used and discards invalid communications. If invalid communications must be allowed, you might need to adjust connection tracking options.

4) If there is a matching NAT rule, make sure that they are applied correctly. Particularly, dynamic NAT must only be used for protocols that work on top of TCP or UDP because dynamic NAT uses ports to track the translated connections.

5) Check your routing configuration. If Routing is incorrectly configured on the Engine, packets can be dropped because the Engine has no route where to send them.

Related concepts

Protocol elements and how they work on page 931
Alert log messages for troubleshooting on page 1375

Related tasks

Troubleshoot traffic that is blocked even though rules allow ANY Service on page 1410

Troubleshooting packets incorrectly dropped by antispoofing

You can add exceptions to antispoofing if traffic that should be allowed is incorrectly dropped as spoofed.

The antispoofing rules are automatically generated based on your routing configuration. Generally, traffic is only allowed if the IP address seen in the communications corresponds to the IP address space that is defined for routing through that interface in the Routing tree. Normally, communications require this routing information in any case for any reply packets to be correctly routed. In cases where communications are one way, however, you can make exceptions to the antispoofing in the Antispoofing tree.

By default, the antispoofing tree is read by selecting the most specific entry defined. For example, a definition of a single IP address is selected over a definition of a whole network. If some IP address must be allowed access

through two or more different interfaces, the definition for each interface must be at the same level of detail for the IP address in question.

If Interface A contains a Host element for 192.168.10.101 and Interface B contains a Network element for 192.168.10.0/24, connections from 192.168.10.101 are considered spoofed if they enter through Interface B. If this behavior is not wanted, the Host element must be added also under Interface B (in addition to the Network element already included).

The preceding behavior can be changed by setting the more general setting (network) as **Absolute** through its right-click menu in the antispoofing tree. This setting allows the address through the interface even if there is a more specific definition attached to some other interface.

Antispoofing also discards packets that are in a routing loop: if the Engine accepts a packet, but then receives the exact same packet again through a different interface, the Engine drops it. This behavior does not affect communications, but saves the Engine and other equipment in your network from handling the same packet over and over again until it finally expires. If so, you must correct the routing in your network. Often, routing loops are indicated by “NIC index changed” information in logs that discard the connection. The same packet enters the Engine a second time, but through a different interface - usually because the device where the Engine is configured to send the packet routes the packet right back to the Engine.

Troubleshooting unsupported definitions in IPv6 Access rules

Using elements that do not contain an IPv6 address in the Source and Destination fields of the IPv6 Access rules causes an error message indicating the rule is invalid.

Elements that contain only an IPv4 address cannot be used. Make sure that all elements used in the Source and Destination field of the IPv6 Access Rules contain an IPv6 address.

Chapter 97

Troubleshooting reporting

Contents

- Error messages for reports on page 1415
- Troubleshoot report generation on page 1415
- Troubleshoot reports with empty sections or incomplete data on page 1416

There are some common problems that you might encounter when generating reports from raw statistical and log data stored on the Log Server.

Error messages for reports

Error messages for reports are shown in the **Comment** column of the **Stored Reports** view. Check the status of the report task there before you proceed with the troubleshooting.

Related tasks

Troubleshoot report generation on page 1415

Troubleshoot reports with empty sections or incomplete data on page 1416

Troubleshoot report generation

There are some common causes and solutions for problems with generating reports.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) If the report generation seems to work, but you cannot find it in the **Stored Reports** view, you might have created the report without selecting the **Stored** option. Try creating the report again and make sure that you have selected all outputs you want to produce.

2) If the report generation does not begin, and there is no error indicated, you might have defined a start time that is in the future. The delay in the **Start Earliest** field can affect when report generation begins. Check the start time for the report task. The time is interpreted according to the clock of the computer you are using to run the SMC Client. Make sure that the time and time zone settings are correct.

3) If the “out of memory” error appears, check if you have placed one or more IP address-based top rate items under progress sections in the Report Design:
   - Generating reports with such a design consumes large amounts of memory. The report collects the full progress information for every IP address that appears in the logs over the chosen period. Any unnecessary data is discarded only after the top items are selected at the completion of the report task.
   - To reduce the memory load, use a Drill-down top rate design. The Drill-down top rate design first finds the top IP addresses and then gets the progress information about those IP addresses.
   - Memory consumption can also be reduced by restricting the amount of data included in the report (for example, set a shorter time range).

4) If the “unreachable server” error is shown, a Log Server is not running or is not reachable from the Management Server. If there are Log Server elements that do not represent any physical, running Log Server, those Log Servers might cause this error. Select the **Exclude This Log Server from Statistics and Reporting** option in the properties of the Log Server elements.

CAUTION

Be careful when excluding Log Servers from reporting. If you select this setting for a Log Server that is actively used, there is no warning that the reports are missing parts of the log data.

Troubleshoot reports with empty sections or incomplete data

There are several possible causes and solutions when reports are generated, but the reports contain empty sections or incorrect data.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) If sections show a “No Data” message, check that the log data that the section requires is in the active log data directory of your Log Servers. Archived logs are not included in reporting.
   a) Open the **Logs** view.
   b) Make sure only **Active Log Data** is selected for all Log Servers on the **Sources** tab.
   c) Select the same time period and filter that you used in the report, and click **Apply**.

   The logs shown correspond to the log data available for generating the logs. Also take note of the following point.

2) If you are only missing information about traffic volumes (for example, **Traffic by Source IP**) check your Access rules. Traffic information is available only for connections that match rules for which you collect accounting information.

   - The collected accounting information is shown in the **Logs** view for log entries that have “Connection Closed” as their Event (see the **Details** pane for those entries).

     Note

     Connection closing might not be logged at all, depending on the logging options of your Access rules.

   - If accounting data has not been collected at all, traffic volumes are unknown and the report sections on traffic volumes return “No Data”.
   - If only some rules collect accounting data, only the traffic that matches those rules is included in the report sections on traffic volumes.
   - Alternatively, the report items listed under **Counters** can be used to generate reports on traffic volumes. The data for these items comes from stored summaries of the statistical data that you can view as live statistics in an Overview. This data is always stored and includes information about all traffic.
   - To start collecting information about traffic volumes from now on.

3) If you are missing all data from one or more Log Servers, make sure that the **Exclude This Log Server from Statistics and Reporting** option is not selected. Log Servers with this option selected are ignored in all reporting.

4) The time range you enter is interpreted according to the clock and time zone setting of the computer you are using to run the SMC Client. If the clock and time zone are not correctly set in the operating system, the report period might be different from what you intend.

Related tasks

Define logging options for Ethernet rules on page 898

Chapter 98

Troubleshooting upgrades

Contents

- Troubleshoot upgrade failure because of running services on page 1419
- Respond to Security Management Center installation failed messages on page 1419

There are some common problems that you might encounter when upgrading SMC components.

Troubleshoot upgrade failure because of running services

Troubleshoot when you are unable to upgrade because the upgrade reports that some services are still running.

Steps

1) In the Windows environment, check the **Services** window and stop any SMC services that are still running (Management Server, Log Server, or Web Access Server).

2) If all Services are stopped in the Windows **Services** window, but the upgrade still reports that services are still running, follow these steps:

   a) Set the services to **Manual startup**.

   b) Restart the computer.

Respond to Security Management Center installation failed messages

If any anomaly is detected during the installation or upgrade, you might see the “Security Management Center installation failed” message.

Steps

1) To find the cause of the problem:
   a) Check the installation log in the installation directory for messages regarding the upgrade.



Note

If you installed the Management Server in the `C:\Program Files\Forcepoint\SMC` directory in Windows, some program data might be stored in the `C:\ProgramData\Forcepoint\SMC` directory.

   b) If you are skipping versions, check the Release Notes for the version you are installing, and make sure the upgrade that you attempted is possible. The Release Notes also list any known issues that can cause errors during the upgrade.

2) To solve problems indicated:
   a) Missing files are one of the most common errors. Files might be missing if you copy installation source files manually. Make sure that you have copied all necessary installation source files and folders, and run the installation again. If you have not checked the integrity of the installation files, compare the checksum of your local files to the correct checksum.

   b) Start the component in question. The message is shown for many types of errors, and the component might still be able to run without problems.

   c) If you are unable to get the component running, uninstall the existing installation and install the component as a new installation. To upgrade, you can restore elements from a Management Server backup. You can restore a backup taken with the previous version on the upgraded SMC. See the Release Notes for any version-specific limitations or exceptions.

Related tasks

Obtain SMC installation files on page 1347

Chapter 99

Troubleshooting VPNs

Contents

- Find VPN issues on page 1421
- Information about IPsec tunnels in logs on page 1422
- VPN certificate issues on page 1423
- Problems with VPNs with external gateways on page 1423
- Forcepoint VPN Client connection issues on page 1424
- If traffic is not sent into route-based VPNs on page 1425

There are some common problems that you might encounter when creating and managing VPNs.

Find VPN issues

The SMC Client has automatic VPN validation that checks the settings you have selected are compatible with each other.

When you configure the policy-based VPN, the validity of the VPN is automatically checked based on settings in VPN Profile elements, VPN Gateway Profile elements assigned to gateways in the VPN, endpoint identity information, and gateways certificates. If problems are found, they are shown in the **Issues** view. While useful in many cases, the automatic check does not detect all problems, especially regarding external gateways or interference between several separate policy-based VPNs.

Steps

For more details about the product and how to configure features, click **Help** or press **F1**.

1) Select **Secure SD-WAN Configuration**.

2) Browse to **Policy-Based VPNs**

3) Right-click the **Policy-Based VPN** element, then select **Preview VPN** or **Edit VPN**.

4) Switch to the **Issues** tab.
   If the tab is not visible, select **Menu** > **View** > **Panels** > **Issues**.

5) If issues are shown, read the descriptions, then fix the problems that are described.

Information about IPsec tunnels in logs

The Engine logs contain information about IPsec tunnel negotiations.

Log messages related to IPsec tunnel negotiations contain the value IPsec in the **Facility** field. You can use this value to filter logs for viewing IPsec messages. The **IKE Cookie** and **IPsec SPI** fields contain identifiers related to each particular VPN instance, which helps further in reading and filtering the logs. The **Situation** and **Information Message** fields include the actual VPN-related events. If possible, examine logs from the devices at both ends of the IPsec tunnel for more information.



Tip

Right-click a VPN log entry and select **Search Related Events** to see logs related to the same IPsec VPN negotiation.

You can collect more detailed information by enabling the IPsec diagnostics. For VPN clients, also enable authentication and DHCP relay diagnostics.

Log messages generated by Access rules might also contain relevant information. These logs contain information about the connections that the gateway processes, and whether policy-based VPN traffic is directed correctly to VPN tunnels by the policy. Log messages generated by Access rules are not included if you are filtering the logs to only show IPsec logs.

A normal IPsec tunnel negotiation proceeds as follows:

1) The negotiations start when a connection matches a rule in the Engine Policy that triggers the VPN negotiation (or a similar mechanism at the other end).

2) The gateway at the source end of the connection or the VPN client (the *initiator*, I) contacts the gateway at the other end (the *responder*, R). The gateways establish trust and exchange keys in the IKE Phase 1 negotiations.

3) If Phase 1 negotiation succeeds, IKE Phase 2 negotiations begin. At this stage, the gateways agree on further settings used for handling the connection.

4) If Phase 2 negotiations succeed, the VPN tunnel is ready and ESP or AH packets (the actual traffic) can be seen in the logs. New connections that are opened through the VPN are logged using a VPN-specific log message “New Connection Through VPN.”

Related concepts

IPsec VPN log messages

Monitoring policy-based VPNs on page 1208

Related tasks

Enable or disable diagnostics on page 356

VPN certificate issues

Your VPN Gateway always needs a certificate if VPN clients connect to it. In a site-to-site VPN, certificates are required when the VPN Profile used includes a certificate-based authentication method (ECDSA signatures, RSA signatures, or DSS (DSA) signatures).

Certificate acceptance

By default, the gateways only accept certificates signed by your Management Server. To accept certificates from other sources, you must define the certificate authority (CA) that signed the certificate as trusted. By default, all Gateways and all VPNs accept any valid CA that you have configured. You can configure the trusted CAs at the Gateway level and at the VPN level. A CA must be trusted on both levels to be accepted as a trusted CA for a VPN.

Creating, signing, renewing, transferring to gateways

Internally signed certificates are created, uploaded to the engines, and renewed automatically if automatic certificate management is enabled for the Security Engine.

You can manually create certificate requests, import certificates, and sign certificate requests in the **Administration** > **Certificates** branch of the Configuration view. Any certificate request you create is, by default, also immediately signed using the internal CA and uploaded to the engine. To disable this action (for example, to sign the certificate using an external CA), you must deactivate this option in the new certificate request you create.

To sign or upload a certificate, display the certificates, then select ⋮ **More actions** and the corresponding option.

Related concepts

VPN certificates and how they work on page 1251

Related tasks

Define additional VPN certificate authorities on page 1254
Restrict the trusted CAs for a VPN gateway on page 1178

Problems with VPNs with external gateways

There are some common problems and solutions when you create a VPN with an external gateway.

Both policy-based and route-based VPNs can form IPsec VPN tunnels with any other fully IPsec compliant device.



Note

Make sure that you have successfully installed or refreshed the policy on all affected Engines after you have changed any part of the VPN configuration.

When creating a VPN with an external gateway:

- There are no settings that always work with a device of a certain brand and model. Most IPsec settings depend on user preference and there are many alternative settings that you can use, regardless of the type of gateway.
- Make sure that all VPN settings are the same at both ends (for both gateways at both ends: typically, four definitions in all).
- Make sure that matching networks and netmasks are defined at both ends. In the SMC, all networks you want to be accessible through the VPN must be placed in a Site element attached to the correct Gateway element. The networks defined must be identical at the other end.
- One commonly missed setting is the SA (Security Association) setting, which can be per net or per host. Some gateways might not have an explicit setting for the SA setting. Find out the setting used.
- For third-party devices, check for parameters that are set in the VPN configuration in the SMC Client but not on the other device. Find out the default settings used.
- The problem might be due to an overlapping, but mismatching lifetime or encryption domain in the SMC, or the IP address definitions in Site elements under the following conditions:
  - The VPN works when the connection is initiated from one end, but not when initiated from the other.
  - The Engine's policy has rules for both ways.

Forcepoint VPN Client connection issues

If you experience connection issues, there are several troubleshooting processes to follow.

Mobile VPNs are only supported in policy-based VPNs.

If NAT is used and the configuration download succeeds, but the Forcepoint VPN Client cannot connect to the VPN Gateway, follow these troubleshooting steps:

1) If NAT is done between the Forcepoint VPN Client and the Engine, set the Contact Address for interfaces that are used as a VPN endpoint. The Contact Address tells the VPN clients the external NATed address they must contact.

2) Refresh the Engine Policy and make sure the Forcepoint VPN Client downloads a new configuration from the engine.

If NAT is configured to translate the Forcepoint VPN Client address, but NAT is not done, check the following:

- In the VPN properties, make sure that the **Apply NAT to traffic that uses this VPN** option is selected. NAT is only done if the option is selected.
- Make sure that the NAT rules are correct. Usually, NAT is performed using the NAT Pool address range defined for the Engine element in the **VPN** > **Advanced** branch in the Engine Editor. The same traffic must not match any of the NAT rules in the Engine's policy. The only exception is a rule that specifically defines that no NAT is performed on this traffic to prevent subsequent NAT rules from matching.

For any general problems:

- Make sure that the Forcepoint VPN Client version is up to date. Older clients might have known issues that prevent correct operation and might not support all features configured for the gateway.
- Check for any VPN-capable devices between the Engine/VPN role device and the Forcepoint VPN Client. These devices can sometimes attempt to take part in the VPN negotiations.
- Check the Engine logs for information about mobile VPN connections.

Related concepts
Define contact IP addresses on page 127

If traffic is not sent into route-based VPNs

There are some common causes and solutions when traffic is not sent into route-based VPNs.

It is possible to create a half-configured Route-based Tunnels by configuring only one tunnel interface and the routing. This configuration creates a black-hole routing situation in which traffic routed to the tunnel interface is silently discarded. No warnings are given when you install the Engine policy, as the configuration is treated as valid. Traffic is only sent into the a Route-based Tunnels after you fully define the Route-based Tunnels elements.

Appendices

Contents

- Command line tools on page 1429
- Default communication ports on page 1457
- Working with expressions on page 1465
- Predefined Aliases on page 1471
- Situation Context parameters on page 1473
- Regular expression syntax on page 1477
- Schema updates for external LDAP servers on page 1489
- Log fields on page 1491
- Keyboard shortcuts on page 1493
- Multicasting on page 1495

Appendix A

Command line tools

Contents

- Forcepoint Security Management Center commands on page 1429
- Security Engine commands on page 1445
- Server Pool Monitoring Agent commands on page 1454

There are command line tools for the SMC and the Security Engines.

Forcepoint Security Management Center commands

SMC commands include commands for the Management Server, Log Server, and Web Access Server.

In Windows, the command line tools are *.bat script files. In Linux, the files are *.sh scripts. Commands are found in the following locations:

- For SMC installations on Linux or Windows, commands are found in the `<installation directory>/bin` directory.
- For the SMC Appliance, general SMC commands are found in the `/usr/local/forcepoint/smc/bin` directory.
- Commands that are specific to the SMC Appliance are found in the `/usr/bin` directory.

If you enabled the restricted shell when you installed the SMC Appliance, only a limited set of commands is available. These commands include patching utilities, appliance maintenance, service handling, and other basic functionality. To show the list of allowed commands, enter `?`.

Note

When the restricted shell is enabled, all administrator accounts that you create in the SMC automatically use the restricted shell.

On the SMC Appliance, some commands must be run with elevated permissions using sudo. Commands in the restricted shell automatically prompt you to enter the password when required. A list of available sudo commands can be found by running `sudo -l` at the command line.

Note

Only administrators who have SMC Appliance Superuser administrator permissions can log on to the SMC Appliance command line.

Commands that require parameters must be run through the command line (cmd.exe in Windows). Commands that do not require parameters can alternatively be run through a graphical user interface, and can be added as shortcuts during installation.

CAUTION

`login` and `password` parameters are optional. Giving them as command-line parameters can pose a security vulnerability. Do not enter logon and password information unless explicitly prompted to do so by a command line tool.

Forcepoint Security Management Center commands

Command	Description

`ambr-crl`
(SMC Appliance only)
`[-a ADD|--add=ADD]`
`[-d DELETE|--delete=DELETE]`
`[-q|--query]`
`[-i IMPORT_CRL|--import=IMPORT_CRL]`
`[-v]`
`[-l <log file path>]`
`[-h|--help]`
	Fetches the certificate revocation lists (CRLs) for the CA certificates used by the appliance maintenance and bug remediation (AMBR) utilities.
	`-a ADD, --add=ADD` adds a CRL distribution point URL in the form of `http://<url>`.
	`-d DELETE, --delete=DELETE` deletes a CRL distribution point URL.
	`-q, --query` lists CRL distribution points.
	`-i IMPORT_CRL, --import=IMPORT_CRL` imports a CRL from a file.
	`-v` increases the verbosity of the command. You can repeat this command up to two times (`-vv` or `-v -v`) to further increase the verbosity.
	`-l <log file path>` specifies the path to a log file.
	`-h, --help` shows information about the command.

`ambr-decrypt` (SMC Appliance only)	Decrypts an ambr patch; not normally used by administrators. `ambr-install` automatically decrypts patches.

`ambr-install <patch>`
(SMC Appliance only)
`[-F|--force]`
`[-r|--skip-revocation]`
`[--no-backup]`
`[--no-snapshot]`
`[--no-prompt]`
`[-v]`
`[-l <log file path>]`
`[-h|--help]`
	Installs an ambr patch that has been loaded on the system.
	You can install multiple patches with a space between each patch name.
	`-F, --force` forces the reinstallation of the patch or patches.
	`-r, --skip-revocation` skips the certificate revocation checks.
	`--no-backup` does not create a configuration backup.
	`--no-snapshot` does not create a recovery snapshot.
	`--no-prompt` does not prompt before restarting.
	`-v` increases the verbosity of the command. You can repeat this command up to two times to further increase the verbosity.
	`-l <log file path>` specifies the path to a log file.
	`-h, --help` shows information about the command.

`ambr-load <patch>`
(SMC Appliance only)
`[-f IN_FILES|--file=IN_FILES]`
`[-r|--skip-revocation]`
`[-v]`
`[-l <log file path>]`
`[-h|--help]`
	Loads an ambr patch onto the system from either the patch server or from the local file system. A loaded patch means that the file is copied to the local file system, but not installed.
	You can load multiple patches with a space between each patch name.
	`-f IN_FILES, --file=IN_FILES` specifies the local file to load.
	`-r, --skip-revocation` skips the certificate revocation checks.
	`-v` increases the verbosity of the command. You can repeat this command up to two times to further increase the verbosity.
	`-l <log file path>` specifies the path to a log file.
	`-h, --help` shows information about the command.

Command	Description

`ambr-query`
(SMC Appliance only)
`[-c|--clean]`
`[-u|--update]`
`[-a|--all]`
`[-j|--json]`
`[-i INFO|--info=INFO <patch>]`
`[-L <log file path>]`
`[-v]`
`[-h|--help]`
	Shows patch information including:
	■ What is loaded or installed on the system
	■ A list of available updates from the patch server
	■ Detailed information about a specific patch
	`-u , --update` updates the remote patch list from a web server .
	`-c, --clean` cleans the remote patch cache.
	`-a, --all` shows all local and remote patches.
	`-j, --json` formats output as JSON.
	`-i INFO, --info=INFO <patch>` shows detailed information about the patch. You can get information about multiple patches in one command by separating the patch names with a space.
	`-v` increases the verbosity of the command. You can repeat this command up to two times to further increase the verbosity.
	`-L <log file path>` specifies the path to the file where log messages are written.
	`-h, --help` shows information about the command.

`ambr-unload <patch>`
(SMC Appliance only)
`[-a|--all]`
`[-v]`
`[-l <log file path>]`
`[-h|--help]`
	Unloads an ambr patch from the system. The command deletes the patch file if it has not been installed, but it does not uninstall the patch.
	You can unload multiple patches with a space between each patch name.
	`-a, --all` unloads all loaded patches.
	`-v` increases the verbosity of the command. You can repeat this command up to two times to further increase the verbosity.
	`-l <log file path>` specifies the path to a log file.
	`-h, --help` shows information about the command.

`ambr-verify` (SMC Appliance only)	Verifies the signature of a patch file; not normally used by administrators. `ambr-install` automatically verifies patches.
`daemon-ctl {start | stop |restart | enable | disable| status} <daemon>` (SMC Appliance only)	Manages and controls SMC and system services on the SMC appliance. This command allows you to start, stop, restart, enable, disable or check the status of various services. `start <daemon>`, starts the <daemon> service. `stop <daemon>`, stops the <daemon> service. `restart <daemon>`, restarts the <daemon> service. `enable <daemon>`, enables the <daemon> service. `disable <daemon>`, disables the <daemon> service. `status <daemon>`, displays the status of the <daemon>service. **Note** Daemons are chronyd, iptables, ip6tables, network, rsyslog, sgLogServer, sgMgtServer, nmpd, snmptrapd, and sshd.

Command	Description
`revert`	Reverts to the previous installation saved during the upgrade process. The previous installation can be restored at any time, even after a successful upgrade.  **Note** This script is located in `<installation directory>/bin/uninstall`.

`sgActivateWebswing`
`[host=<Management Server Address[\Domain>]`
`login=<login name>`
`pass=<password>`
`port=<port number>`
`mgtserver=<name>`
`enable=<true|false>`
`hostname=<host name>`
`listening_address=<IP address>`
`https=<true|false>`
`generate_logs=<true|false>`
`use_ssl=<true|false>`
`https_validity=<number of days>`
`public_key_output=<path>`
	Configures Web Access to run the SMC Client in a web browser.
	`Host` specifies the address of the Management Server. If the parameter is not defined, the loopback address (localhost) is used.
	`login` defines the user name for the account that is used for this operation. If this parameter is not defined, the user name root is used.
	`pass` defines the password for the user account.
	`port` specifies the port number of the Web Access service on the Management Server. The default is 8085.
	`mgtserver` specifies the name of the Management Server element. The default is Management Server.
	`enable` specifies whether Web Access is enabled (true) or disabled (false). The default is true.
	`hostname` specifies the host name of the Web Access service.
	`listening_address` specifies the listening IP address of the Web Access service if the server has several addresses. If not specified, requests to any of this server's IP addresses are allowed.
	`https` specifies whether HTTPS is enabled for the Web Access service. If true, the public key is returned in the output. The default is true.
	`generate_logs` specifies whether to log all file load events in Combined Log format in a file on the server for further analysis with external web statistics software. The default is false.
	`use_ssl` specifies whether SSL is used to track sessions in your web application. If SSL connections are managed by a proxy or a hardware accelerator they must populate the SSL request headers. The default is false.
	`https_validity` specifies the number of days for which the self-signed certificate for HTTPS is valid. The default is 365.
	`public_key_output` specifies the path for the HTTPS public key.

Command	Description

`sgArchiveExport`
`[host=<Management Server Address[\Domain>]`
`[login=<login name>]`
`[pass=<password>]`
`[format=<CSV|XML|JSON>]`
`i=<input files and/or directories>`
`[o=<output file name>]`
`[f=<filter file name>]`
`[e=<filter expression>]`
`[-h|-help|-?]`
`[-v]`
	Shows and exports logs from archive. Supports CEF, LEEF, and ESM formats in addition to CSV, XML, and JSON.
	This command is only available on the Log Server. The operation checks permissions for the supplied administrator account from the Management Server to prevent unauthorized access to the logs.
	Enclose details in double quotes if they contain spaces.
	`Host` specifies the address of the Management Server. If the parameter is not defined, the loopback address (localhost) is used.
	`login` defines the user name for the account that is used for this operation. If this parameter is not defined, the user name root is used.
	`pass` defines the password for the user account.
	`format` defines the file format for the output file. If this parameter is not defined, the XML format is used.
	`i` defines the source from which the logs are exported. Can be a folder or a file. The processing recurses into subfolders.
	`o` defines the destination file where the logs are exported. If this parameter is not defined, the output is shown on screen.
	`f` defines a file that contains the filtering criteria you want to use for filtering the log data. You can export log filters individually in the SMC Client through **More actions** > **Save for Command Line Tools** in the filter's right-click menu.
	`e` allows you to enter a filter expression manually (using the same syntax as exported filter files).
	`-h`, `-help`, or `-?` shows information about using the script.
	`-v` shows verbose output on the command execution.
	Example (exports logs from one full day to a file using a filter): `sgArchiveExport login=admin pass=abc123 i= C:\Program Files\Forcepoint\SMC \data \archive\engine\year2011\month12\.\sgB.day01\ f= C:\Program Files\Forcepoint\SMC \export\MyExportedFilter.flp format=CSV o=MyExportedLogs.csv`

Command	Description

`sgBackupLogSrv`
`[-pwd=<password>]`
`[-path=<destpath>]`
`[-nodiskcheck]`
`[-comment=<comment>]`
`[-nofsstorage]`
`[-h|--help]`
	Note
	For the SMC Appliance, use the `smca-backup` command.
	
	Creates a backup of Log Server configuration data.
	The backup file is stored in the `<installation directory>/backups/` directory.
	Twice the size of the log database is required on the destination drive. Otherwise, the operation fails.
	`pwd` enables encryption.
	`path` defines the destination path.
	`nodiskcheck` ignores the free disk check before creating the backup.
	`comment` allows you to enter a comment for the backup. The maximum length of a comment is 60 characters.
	`nofsstorage` creates a backup only of the Log Server configuration without the log data.
	`-h` or `--help` shows information about using the script.
	Also see `sgRestoreLogBackup`.

`sgBackupMgtSrv`
`[pwd=<password>]`
`[path=<destpath>]`
`[nodiskcheck]`
`[comment=<comment>]`
`[-h|--help]`
	Note
	For the SMC Appliance, use the `smca-backup` command.
	
	Creates a complete backup of the Management Server (including both the local configuration and the stored information in the configuration database). The backup file is stored in the `<installation directory>/backups/` directory.
	Twice the size of the Management Server database is required on the destination drive. Otherwise, the operation fails.
	`pwd` enables encryption.
	`path` defines the destination path.
	`nodiskcheck` ignores the free disk check before creating the backup.
	`comment` allows you to enter a comment for the backup. The maximum length of a comment is 60 characters.
	`-h` or `--help` shows information about using the script.
	Also see `sgRestoreMgtBackup` and `sgRecoverMgtDatabase`.

`sgCertifyLogSrv` `[host=<Management Server Address[\Domain]>`	Contacts the Management Server and creates a certificate for the Log Server to allow secure communications with other SMC components. Renewing an existing certificate does not require changing the configuration of any other SMC components. `host` specifies the address of the Management Server. If the parameter is not defined, the loopback address (localhost) is used. `Domain` specifies the administrative Domain the Log Server belongs to if the system is divided into administrative Domains. If the Domain is not specified, the Shared Domain is used. Stop the Log Server before running this command. Restart the server after running this command.

Command	Description
`sgCertifyMgtSrv [login=<login name>]` `[pass=<password>]` `[standby-server=<name of additional Management Server>]` `[active-server=<IP address of active Management Server>]` `[mode=ext-pki-init` `[dn=<Subject DN>` `dns=<SubjectAltName DNS>` `key-size=<256|384|521>` `csr-out=<path>` `crt-in=<path>` `ca-file=<path>]]` `[-nodisplay]` `[-h|-help|-?]`	Creates a certificate for the Management Server to allow secure communications between the SMC components. Renewing an existing certificate does not require changes on any other SMC components. In an environment with only one Management Server, or to certify the active Management Server, stop the Management Server before running the `sgCertifyMgtSrv` command. Run the command without parameters. Restart the Management Server after running this command. To certify an additional Management Server, stop the additional Management Server before running the `sgCertifyMgtSrv` command. The active Management Server must be running when you run this command. The management database is replicated to the additional Management Server during the certification. The additional Management Server must have a connection to the active Management Server when you run this command. `[login=<login name>]` defines the user name for the account that is used for this operation. If this parameter is not defined, the user name root is used. `[pass=<password>]` defines the password for the user account. `[standby-server]` specifies the name of the additional Management Server to be certified. `[active-server]` specifies the IP address of the active Management Server. `[mode=ext-pki-init]` enables commands for external certificate management. `[dn]` specifies the Subject DN to use in the certificate request for the Management Server. `[dns]` specifies the SubjectAltName DNS value to use in the certificate request for the Management Server. `[key-size]` specifies the key size to use in the certificate request for the Management Server. `[csr-out]` specifies the output path where the certificate request is saved. `[crt-in]` specifies the input path for importing a certificate in PEM format. `[ca-file]` specifies the input path for importing a CA file in PEM format. `-nodisplay` sets a text-only console. `-h`, `-help`, or `-?` shows information about using the script.
`sgCertifyWebAccessSrv` `[host=<Management Server Address[\Domain]>]`	Contacts the Management Server and creates a certificate for the Web Access Server to allow secure communications with other SMC components. Renewing an existing certificate does not require changing the configuration of any other SMC components. `host` specifies the address of the Management Server. If the parameter is not defined, the loopback address (localhost) is used. `Domain` specifies the administrative Domain the Web Access Server belongs to if the system is divided into administrative Domains. If the Domain is not specified, the Shared Domain is used. Stop the Web Access Server before running this command. Restart the server after running this command.
`sgChangeMgtIPOnLogSrv <IP address>`	Changes the Management Server's IP address in the Log Server's local configuration to the IP address you give as a parameter. Use this command if you change the Management Server's IP address. Restart the Log Server service after running this command.

Command	Description
`sgChangeMgtIPOnMgtSrv <IP address>`	Changes the Management Server's IP address in the local configuration to the IP address you give as a parameter. Use this command if you change the Management Server's IP address. Restart the Management Server service after running this command.
`sgClient`	Starts a locally installed SMC Client.
`sgCreateAdmin`	Creates an unrestricted (superuser) administrator account. The Management Server must be stopped before running this command.

`sgExport [host=<Management Server Address[\Domain]>]`
`[login=<login name>]`
`[pass=password]`
`file=<file path and name>`
`[type=<all|nw|ips|sv|rb|al|vpn>`
`[name=<element name 1, element name 2, ...>]`
`[recursion]`
`[-system]`
`[-h|-help|-?]`
	Exports elements stored on the Management Server to an XML file.
	Enclose details in double quotes if they contain spaces.
	`host` specifies the address of the Management Server. If the parameter is not defined, the loopback address (localhost) is used.
	`Domain` specifies the administrative Domain for this operation if the system is divided into administrative Domains. If the Domain is not specified, the Shared Domain is used.
	`login` defines the user name for the account that is used for this operation. If this parameter is not defined, the user name root is used.
	`pass` defines the password for the user account.
	`file` defines the name and location of the export `.zip` file.
	`type` specifies which types of elements are included in the export file:
	■ `all` for all exportable elements
	■ `nw` for network elements
	■ `ips` for IPS elements
	■ `sv` for services
	■ `rb` for security policies
	■ `al` for alerts
	■ `vpn` for VPN elements.
	`name` allows you to specify by name the elements that you want to export.
	`recursion` includes referenced elements in the export, for example, the network elements used in a policy that you export.
	`-system` includes any system elements that are referenced by the other elements in the export.
	`-h`, `-help`, or `-?` shows information about using the script.

Command	Description

`sgHA`
`[host=<Management Server Address[\Domain]>]`
`[login=<login name>]`
`[pass=<password>]`
`[master=<Management Server used as master server for the operation>]`
`[-set-active]`
`[-set-standby]`
`[-check]`
`[-retry]`
`[-force]`
`[-restart]`
`[-h|-help|-?]`
	Controls active and standby Management Servers.
	If you want to perform a full database synchronization, use the `sgOnlineReplication` command.
	`host` specifies the address of the Management Server. If the parameter is not defined, the loopback address (localhost) is used.
	`Domain` specifies the administrative Domain for this operation if the system is divided into administrative Domains. If the Domain is not specified, the Shared Domain is used.
	`login` defines the user name for the account that is used for this operation. If this parameter is not defined, the user name root is used.
	`pass` defines the password for the user account.
	`master` defines the Management Server used as a master Management Server for the operation.
	`-set-active` activates and locks all administrative Domains.
	`-set-standby` deactivates and unlocks all administrative Domains.
	`-check` checks that the Management Server's database is in sync with the master Management Server.
	`-retry` retries replication if this has been stopped due to a recoverable error.
	`-force` enforces the operation even if all Management Servers are not in sync.
	Note
	This option can cause instability if used carelessly.
	`-restart` restarts the specified Management Server.
	`-h`, `-help`, or `-?` shows information about using the script.

`sgImport`
`[host=<Management Server Address[\Domain]>]`
`[login=<login name>]`
`[pass=<password>]`
`file=<file path and name>`
`[-replace_all]`
`[-h|-help|-?]`
	Imports Management Server database elements from an XML file.
	When importing, existing (non-default) elements are overwritten if both the name and type match.
	`host` specifies the address of the Management Server. If the parameter is not defined, the loopback address (localhost) is used.
	`Domain` specifies the administrative Domain for this operation if the system is divided into administrative Domains. If the Domain is not specified, the Shared Domain is used.
	`login` defines the user name for the account that is used for this operation. If this parameter is not defined, the user name root is used.
	`pass` defines the password for the user account.
	`file` defines the `.zip` file whose contents you want to import.
	`-replace_all` ignores all conflicts by replacing all existing elements with new ones.
	`-h`, `-help`, or `-?` shows information about using the script.

Command	Description

`sgImportExportUser`
`[host=<<Management Server Address[\Domain]>>]`
`[login=<login name>]`
`[pass=password]`
`action=<import|export>`
`file=<file path and name>`
`[-h|-help|-?]`
	Imports and exports a list of Users and User Groups in an LDIF file from or to a Management Server's internal LDAP database.
	To import User Groups, all User Groups in the LDIF file must be directly under the stonegate top-level group (dc=stonegate).
	CAUTION
	The user information in the export file is stored as plaintext. Handle the file securely.
	`host` specifies the address of the Management Server. If the parameter is not defined, the loopback address (localhost) is used.
	`Domain` specifies the administrative Domain for this operation if the system is divided into administrative Domains. If the Domain is not specified, the Shared Domain is used.
	`login` defines the user name for the account that is used for this operation. If this parameter is not defined, the user name `root` is used.
	`pass` defines the password for the user account.
	`action` defines whether users are imported or exported.
	`file` defines the file that is used for the operation.
	Example: `sgImportExportUser login=admin pass=abc123 action=export file=c:\temp\exportedusers.ldif`
	`-h`, `-help`, or `-?` shows information about using the script.

Command	Description

`sgInfo`
`SG_ROOT_DIR`
`FILENAME`
`[fast=<timestamp>]`
`[list]`
`[hprof=none|limited|all]`
`[-nolog]`
`[-client]`
`[-h|-help|-?]`
	Creates a `.zip` file that contains copies of configuration files and the system trace files.
	The resulting `.zip` file is stored in the logged on user's home directory. The file location is shown on the last line of screen output. Provide the generated file to support for troubleshooting purposes.
	Note
	On the SMC Appliance, you must always specify the path to the directory in which the `.zip` file is stored. The directory must be accessible from the account that you use to log on to the command line of the SMC Appliance.
	`SG_ROOT_DIR` SMC installation directory.
	`FILENAME` name of output file.
	`fast` collects only traces that changed after the specified time stamp. Enter the time stamp in milliseconds or in the format yyyy-MM-dd HH:mm:ss. No other information is collected, except for threaddumps.
	`[list]` only lists files. It does not create a `.zip` file or generate threaddumps.
	`hprof` defines whether hprof memory dump files are included.
	■ `none` does not include hprof memory dump files.
	■ `limited` includes only hprof memory dump files that are created with makeheap.
	■ `all` includes memory dump files that are created with makeheap and java_pid.
	`-nolog` extended Log Server information is not collected.
	`-client` collects traces only from the SMC Client.
	`-h`, `-help`, or `-?` shows information about using the script.

`sgOnlineReplication`
`[active-server=<name of active Management Server>]`
`[-nodisplay]`
`[-h|-help|-?]`
	Replicates the Management Server's database from the active Management Server to an additional Management Server.
	Stop the Management Server to which the database is replicated before running this command. Restart the Management Server after running this command.
	Use this script to replicate the database only in the following cases:
	■ The additional Management Server's configuration has been corrupted.
	■ In new SMC installations if the automatic database replication between the Management Servers has not succeeded.
	Otherwise, synchronize the database through the SMC Client.
	CAUTION
	This script also has parameters that are for the internal use of the Management Server only. Do not use this script with any parameters other than the ones listed here.
	`active-server` specifies the IP address of the active Management Server from which the Management database is replicated.
	`-nodisplay` sets a text-only console.
	`-h`, `-help`, or `-?` shows information about using the script.

Command	Description
`sgReinitializeLogServer`	Creates a Log Server configuration if the configuration file has been lost.  **Note** This script is located in `<installation directory>/bin/install`.
`sgRestoreArchive <ARCHIVE_DIR>`	Restores logs from archive files to the Log Server. This command is available only on the Log Server. `ARCHIVE_DIR` is the number of the archive directory (0–31) from where the logs will be restored. By default, only archive directory 0 is defined. The archive directories can be defined in the `<installation directory>/data/LogServerConfiguration.txt` file: `ARCHIVE_DIR_ xx=PATH`.

`sgRestoreLogBackup`
`[-pwd=<password>]`
`[-backup=<backup file name>]`
`[-nodiskcheck]`
`[-overwrite-syslog-template]`
`[-h|-help]`
	Restores the Log Server (logs or configuration files) from a backup file in the `<installation directory>/backups/` directory.
	`-pwd` defines a password for encrypted backup.
	`-backup` defines a name for the backup file.
	`-nodiskcheck` ignores the free disk check before backup restoration.
	`-overwrite-syslog-template` overwrites a syslog template file if found in the backup.
	`-h` or `-help` shows information about using the script.

`sgRestoreMgtBackup`
`[-pwd=<password>]`
`[-backup=<backup file name>]`
`[-import-license <license file name>]`
`[-nodiskcheck]`
`[-h|-help]`
	Restores the Management Server (database or configuration files) from a backup file in the `<installation directory>/backups/` directory.
	`-pwd` defines a password for encrypted backup.
	`-backup` defines a name for the backup file.
	`-import-license` specifies a license file to import during the backup restoration.
	`-nodiskcheck` ignores the free disk check before backup restoration.
	`-h` or `-help` shows information about using the script.

`sgShowFingerPrint` `[-server]`	Shows the CA certificate's fingerprint on the Management Server. `-server` displays Management Server Certificate Fingerprint.
`sgStartLogSrv`	Starts the Log Server and its database.
`sgStartMgtDatabase`	Starts the Management Server's database. There is usually no need to use this script.
`sgStartMgtSrv`	Starts the Management Server and its database.
`sgStartWebAccessSrv`	Starts the Web Access Server.
`sgStopLogSrv`	Stops the Log Server.
`sgStopMgtSrv`	Stops the Management Server and its database.
`sgStopMgtDatabase`	Stops the Management Server's database. There is usually no need to use this script.
`sgStopWebAccessSrv`	Stops the Web Access Server.

Command	Description

`sgStopRemoteMgtSrv`
`[host=<Management Server address[\Domain]>]`
`[login=<login name>]`
`[pass=<password>]`
`[-h|-help|-?]`
	Stops the Management Server service when run without arguments.
	To stop a remote Management Server service, provide the arguments to connect to the Management Server.
	`host` is the Management Server's host name if not localhost.
	`login` is an SMC administrator account for the logon.
	`pass` is the password for the administrator account.
	`-h`, `-help`, or `-?` shows information about using the script.

`sgTextBrowser`
`[host=<Management Server address[\Domain]>]`
`[login=<login name>]`
`[pass=<password>]`
`[format=<CSV|XML|JSON>]`
`[o=<output file>]`
`[f=<filter file>]`
`[e=<filter expression>]`
`[m=<current|stored>]`
`[limit=<maximum number of unique records to fetch>]`
`[-h|-help|-?]`
	Shows or exports current or stored logs.
	This command is available on the Log Server.
	Enclose the file and filter names in double quotes if they contain spaces.
	`host` defines the address of the Management Server used for checking the logon information. If this parameter is not defined, Management Server is expected to be on the same host where the script is run. If Domains are in use, you can specify the Domain the Log Server belongs to. If `domain` is not specified, the Shared Domain is used.
	`login` defines the user name for the account that is used for this export. If this parameter is not defined, the user name root is used.
	`pass` defines the password for the user account used for this operation.
	`format` defines the file format for the output file. If this parameter is not defined, the XML format is used.
	`o` defines the destination output file where the logs will be exported. If this parameter is not defined, the output is shown on screen.
	`f` defines the exported filter file that you want to use for filtering the log data.
	`e` defines the filter that you want to use for filtering the log data. Type the name as shown in the SMC Client.
	`m` defines whether you want to view or export logs as they arrive on the Log Server (current) or logs stored in the active storage directory (stored). If this option is not defined, the current logs are used.
	`limit` defines the maximum number of unique records to be fetched. The default value is unlimited.
	`-h`, `-help`, or `-?` shows information about using the script.

`smca-agent` (SMC Appliance only)	SMC uses it to exchange configuration data between SMC and the operating system; not normally used by administrators. The agent configures the NTP and SNMP daemons and sets the logon and SSH banners.

Command	Description

`smca-backup`
(SMC Appliance only)
`[-pwd <password>]`
`[-comment <comment>]`
`[-nodiskcheck]`
`[-nofsstorage]`
`[-path <destination>]`
`[-log]`
`[-mgt]`
`[-h|--help]`
	Creates a configuration backup of the SMC Appliance operating system and includes an SMC backup.
	`-pwd <password>` enables the encryption of the backup file and sets the password.
	`-comment <comment>` adds a comment to the backup file name.
	`-nodiskcheck` turns off the available disk space check.
	`-nofsstorage` excludes the log files for the Log Server from the backup.
	`-path <destination>` specifies a path for backup file storage. The default directory for backups is `/usr/local/forcepoint/smc/backups`.
	`-log` creates a Log Server backup.
	`-mgt` creates a Management Server backup.
	`-h, --help` shows information about the command.
	Also see `sgRestoreLogBackup` and `sgRestoreMgtBackup`.

`smca-backup-remove`
(SMC Appliance only)
`[-f|--file]`
`[--force]`
`[--age <days>]`
`[--log]`
`[--mgt]`
`[-h|--help]`
	Removes old SMC Appliance backup files.
	`-f,--file` specifies the backup file to be removed.
	`--force` forces backup file delete without confirmation.
	`[--age <days>]` remove any backups older than the specified number of days. The default is 30 days.
	`[--log]` removes Log Server backups.
	`[--mgt]` removes Management Server backups.
	`-h, --help` shows information about the command.

`smca-cifs`
(SMC Appliance only)
`[add]`
`[remove]`
`[-n <name>]`
`[-s //<server>/<share>]`
`[-u <username>]`
`[-p <password>]`
`[-d <domain>]`
	Configures the mounting of remote CIFS file shares on the SMC Appliance.
	`add` adds the CIFS share.
	`remove` removes the CIFS share. Use with the name option.
	`-n <name>` specifies the name of the share.
	`-s //<server>/<share>` specifies the server or IP address of the share.
	`-u <username>` specifies the user name to authenticate with the CIFS server to get access to the share.
	`-p <password>` specifies the password on remote system.
	`-d <domain>` specifies the domain of the share.

Command	Description

`smca-restore`
(SMC Appliance only)
`[-pwd <password>]`
`[-nodiskcheck]`
`[-backup <filename>]`
`[-nosmca]`
`[-smcaonly]`
`[-overwrite-syslog-template]`
`[-h|-help]`
	Restores a backup on the SMC Appliance.
	`-pwd <password>` specifies the password for decrypting an encrypted backup file.
	`-nodiskcheck` turns off the available disk space check.
	`-backup <filename>` specifies the backup file name. If you do not specify the backup file name, you are prompted to select the backup file.
	`[-nosmca]` restores the Management Server or Log Server backup without restoring the SMC Appliance configuration
	`[-smcaonly]` restores the SMC Appliance configuration without restoring the Management Server or Log Server backup.
	`-overwrite-syslog-template` overwrites any existing syslog templates in the log backup file.
	`-h, --help` shows information about the command.

`smca-rsync`
(SMC Appliance only)
`[add]`
`[modify]`
`[remove]`
`[enable]`
`[disable]`
`[list]`
`[run]`
`[-t task_id]`
`[-i <source directory>]`
`[-o <destination directory>]`
`[-m <mode>]`
`[-h|-help]`
	Configures automated backup tasks. Typically used with the `smca-cifs` command to move backups off the appliance.
	`add` adds a backup task. You can specify an existing source and destination directories. If not specified, the default is `/usr/local/forcepoint/smc/backups/`.
	`modify` changes an existing backup task by its task ID. All attributes can be changed, except for the task ID. To change an attribute, use the appropriate option with a new value.
	`remove` removes an existing backup task by its task ID.
	`enable` enables an existing backup task by its task ID.
	`disable` disables an existing backup task by its task ID.
	`list` provides a list of all configured backup tasks.
	`run` runs all enabled backup tasks.
	`-t task_id` specifies the task ID. Use the list command to view the task IDs.
	`-i <source directory>` specifies the directory where the backups are stored when they are created. If omitted, the source directory defaults to the SMC backups directory `/usr/local/forcepoint/smc/backups/`.
	`-o <destination directory>` specifies the remote location to store the backups.
	`-m <mode>` specifies the rsync mode. You can indicate whether rsync appends or mirrors the source directory to the destination directory. Appending the directory means that existing files in the destination directory, that are not in the source directory or are newer than those files in the source directory, are not changed. If omitted, the mode defaults to append.
	`-h, --help` shows information about the command.

Command	Description

`smca-system`
(SMC Appliance only)
`[toggle]`
`[toggle-vcdrom]`
`[mirror [-n <name>]]`
`[snapshot [-C|--create] [-R|--restore] [-D, --delete] [-n <name>]]`
`[serial-number]`
`[fingerprint]`
`[toggle-console]`
`[bootloader-password [-s|--set] [-r|--remove]]`
`[netconfig]`
`[log-view [<filename>]]`
`[fips-config]`
`[file-remove [filename] [-h] [-l] [--autoremove] [-a <age>] [--no-prompt]]`
`[-f]`
`[-h|-help]`
	Manages recovery snapshots, alternate partition mirroring, and changing system partition boot preference.
	`toggle` restarts the appliance to the alternate partition.
	`toggle-vcdrom` sets the appliance's default boot option to the vcdrom.
	`mirror` mirrors the active system to the alternate system. `-n <name>` specifies the name of the snapshot used for mirror operations.
	`snapshot` manages recovery snapshots.
	■ `-C, --create` creates a snapshot.
	■ `-D, --delete` deletes the snapshot.
	■ `-R, --restore` restores the snapshot.
	■ `-n <name>` specifies the name of the snapshot used for snapshot operations.
	`[serial-number]` shows the hardware serial number for the SMC Appliance.
	`[fingerprint]` shows the CA certificate fingerprint. If an external CA is configured, shows the Management Server certificate fingerprint instead.
	`toggle-console` enables or disables the serial console on the SMC Appliance.
	`bootloader-password` manages the bootloader password for the SMC Appliance.
	■ `-s, --set` sets or changes the bootloader password.
	■ `-r, --remove` removes the bootloader password.
	`netconfig` sets up network-related configuration, such as IPv6 configuration.
	`log-view <filename>` shows the contents of the specified log file in the SMC Appliance log data directory `/var/log` or in any of the subdirectories of `/var/log`. `log-view -l` shows a list of all available log files.
	`fips-config` modifies the SMC Appliance configuration to support FIPS certification.
	`file-remove` deletes the specified SMC files from the SMC Appliance.
	■ `filename` specifies the file to remove. This command can remove files in the following directories:
	  ■ `<SMC_DATA_DIR>/storage`
	  ■ `<SMC_DATA_DIR>/mgtserver`
	  ■ `<SMC_DATA_DIR>/SGInfo`
	  ■ `<SMC_DATA_DIR>/TrafficCapture`
	  ■ `<SMC_DATA_DIR>/datamgtserver/webswing/users`
	■ `-h, --help` shows information about the remove command.
	■ `-l, --list` lists the files that can be deleted.
	■ `--autoremove` removes Web Swing files that are older than the number of days that are specified in the `-a <age>, --age <age>` option.
	■ `-a <age>, --age <age>` specifies the age of Web Swing files to remove with the `--autoremove` option. The default is 30 days.
	■ `--no-prompt` deletes the selected files without prompting for confirmation.
	`-f` forces the procedure, does not prompt for any confirmation.
	`-h, --help` shows information about the command.

`smca-user` (SMC Appliance only)	This utility is used by the SMC Appliance to keep user accounts in sync between the SMC and the operating system; not normally used by administrators.

Security Engine commands

There are commands that can be run on the command line on Engine, Layer 2 Engine, IPS engines, or Master Engines.



Note

Using the SMC Client is the recommended configuration method, as most of the same tasks can be done through it.



Note

All command line tools that are available for single Security Engines are also available for Virtual Engines that have the same role. However, there is no direct access to the command line of Virtual Engines. Commands to Virtual Engines must be sent from the command line of the Master Engine using the `se-virtual-engine` command.

Some commands are only available when the Security Engine is in the Engine (FW), Layer 2 Engine (L2FW), or IPS engine (IPS) role.

Security Engine command line tools

Command	Role	Description
`activate-alternative-policy`	FW L2FW IPS	Not supported on Master Engine or Virtual Engine. Activate alternative policies. `-s <number>|--slot=<number>` Use given alternative policy slot instead of prompting interactively. `-p <name>|--policy-name=<name>` Use policy with given name instead of prompting interactively. `-h|--help` Show this help and exit.
`remove-alternative-policy [options]`	FW L2FW IPS	Not supported on Master Engine or Virtual Engine. Remove selected local alternative policy(s). `-s <number>|--slot=<number>` Remove given alternative policy slot instead of prompting interactively. `-p <name>|--policy-name=<name>` Use policy with given name instead of prompting interactively. `-f|--force` No error if specified policy does not exist. `-a|--all` Remove all alternative policies. `--h|--help` Show this help and exit. **Note** You can remove just one policy at a time. User must select the local alternative policy to remove.
`restore-alternative-policy [options]`	FW L2FW IPS	Not supported on Master Security Engine or Virtual Security Engine. Restore selected local alternative policy(s). You can remove just one policy at a time. User must select the local alternative policy to remove.

Command	Role	Description

`sg-blocklist`
`show [-v] [-f FILENAME ] |`
`add [`
`[-i FILENAME]|`
`[src IP_ADDRESS/MASK]`
`[src6 IPv6_ADDRESS/PREFIX]`
`[dst IP_ADDRESS/MASK]`
`[dst6 IPv6_ADDRESS/PREFIX]`
`[proto {tcp|udp|icmp|NUM}]`
`[srcport PORT {-PORT}]`
`[dstport PORT {-PORT}]`
`[duration NUM]`
`[ve VIRTUAL_ENGINE_ID]`
`] |`
`del [`
`[-i FILENAME]|`
`[src IP_ADDRESS/MASK]`
`[src6 IPv6_ADDRESS/PREFIX]`
`[dst IP_ADDRESS/MASK]`
`[dst6 IPv6_ADDRESS/PREFIX]`
`[proto {tcp|udp|icmp|NUM}]`
`[srcport PORT{-PORT}]`
`[dstport PORT{-PORT}]`
`[duration NUM]`
`[ve VIRTUAL_ENGINE_ID]`
`] |`
`iddel NODE_ID ID |`
`flush`
	FW
	L2FW
	IPS
		Used to view, add, or delete active block list entries.
		The block list is applied as defined in Access Rules.
		`show` shows the current active block list entries in format: `engine node ID | block list entry ID | (internal) | entry creation time | (internal) | address and port match | originally set duration | (internal) | (internal)`. Use the `-f` option to specify a storage file to view (`/data/block list/db_<number>`). The `-v` option adds operation's details to the output.
		`add` creates a block list entry. Enter the parameters or use the `-i` option to import parameters from a file.
		`del` deletes the first matching block list entry. Enter the parameters or use the `-i` option to import parameters from a file.
		`iddel` removes one specific block list entry on one specific Security Engine. `NODE_ID` is the ID of the Security Engine, `ID` is the block list entry's ID (as shown by the show command).
		`flush` deletes all block list entries.
		Add/Del Parameters:
		Enter at least one parameter. The default value is used for the parameters that you omit. You can also save parameters in a text file; each line in the file is read as one block list entry.
		`src` defines the source IP address and netmask to match. Matches any IP address by default.
		`src6` defines the source IPv6 and prefix length to match. Matches any IPv6 address by default.
		`dst` defines the destination IP address and netmask to match. Matches any IP address by default.
		`dst6` defines the destination IPv6 address and prefix length to match. Matches any IPv6 address by default.
		`proto` defines the protocol to match by name or protocol number. Matches all IP traffic by default.
		`srcport` defines the TCP/UDP source port or range to match. Matches any port by default.
		`dstport` defines the TCP/UDP destination port or range to match. Matches any port by default.
		`duration` defines in seconds how long the entry is kept. Default is 0, which cuts current connections, but is not kept.
		`ve` specifies the Virtual Engine on which the block list entry is created or deleted.
		Examples:
		`sg-blocklist add src 192.168.0.2/32 proto tcp dstport 80 duration 60`
		`sg-blocklist add -i myblock list.txt`
		`sg-blocklist del dst 192.168.1.0/24 proto 47`

Command	Role	Description

`sg-bootconfig`
`[--primary-console=tty0|ttyS PORT,SPEED]`
`[--secondary-console=[tty0|ttyS PORT,SPEED]]`
`[--flavor=up|smp]`
`[--initrd=yes|no]`
`[--crashdump=yes|no|Y@X]`
`[--append=kernel options]`
`[--help]`
`apply`
	FW
	L2FW
	IPS
		Used to edit boot command parameters for future bootups.
		`--primary-console` defines the terminal settings for the primary console.
		`--secondary-console` defines the terminal settings for the secondary console.
		`--flavor` defines whether the kernel is uniprocessor or multiprocessor.
		`--initrd` defines whether Ramdisk is enabled or disabled.
		`--crashdump` defines whether kernel crashdump is enabled or disabled, and how much memory is allocated to the crash dump kernel (Y). The default is 24M. X must always be 16M.
		`--append` defines any other boot options to add to the configuration.
		`--help` shows usage information.
		`apply` applies the specified configuration options.

`sg-clear-all`
`[--help]`
`[--flash-defaults]`
`[--dry-run]`
`[--on-boot]`
`[--reboot | --shutdown]`
`[--fast] | --wipe <number>]`
`[--debug | --verbose]`
	FW
	L2FW
	IPS
		This command restores the factory default settings on the Security Engine.
		`[--help]` shows usage information.
		`[--flash-defaults]` assumes that the Security Engine has a flash data partition and a RAM spool partition.
		`[--dry-run]` exits without shutting down or restarting when command execution finishes.
		`[--on-boot]` indicates that Security Engine is starting up. This option is not intended to be used in normal command line usage.
		`[--reboot]` the Security Engine always restarts when command execution finishes.
		`[--shutdown]` the Security Engine always shuts down when command execution finishes.
		`[--fast]` runs a minimal, non-interactive clear for testing purposes.
		`[--wipe <number>]` globally specifies the number of times to wipe partitions.
		`[--debug]` shows full debug messages during command execution.
		`[--verbose]` shows additional informational messages during command execution.
		 **Note**
		If you run the command without specifying any options, the Security Engine requests confirmation before restarting. When the Security Engine restarts, you are prompted to select the system restore options.
		After using this command, you can reconfigure the Security Engine using the `sg-reconfigure` command.

Command	Role	Description

`sg-cluster`
`[-v <Virtual Engine ID>]`
`[status [-c SECONDS]]`
`[versions]`
`[online]`
`[lock-online]`
`[offline]`
`[lock-offline]`
`[standby]`
`[safe-offline]`
`[force-online]`
`[move]`
	FW
	L2FW
	IPS
		Shows or changes the status of the node.
		`-v` (Master Engine only) specifies the ID of the Virtual Engine on which to execute the command.
		`status` shows cluster status. When `-c SECONDS` is used, the status is shown continuously with the specified number of seconds between updates.
		`version` shows the Security Engine software versions of the nodes in the cluster.
		`online` sends the node online.
		`lock-online` sends the node online and keeps it online, even if another process tries to change its state.
		`offline` sends the node offline.
		`lock-offline` sends the node offline and keeps it offline, even if another process tries to change its state.
		`standby` sets an active node to standby.
		`safe-offline` sets the node to offline only if there is another online node.
		`force-online` sets the node online regardless of state or any limitations. Also sets all other nodes offline.
		`[move]` (Master Engine only) moves the specified Virtual Security Engine to this node.

`sg-contact-mgmt`	FW L2FW IPS	Used for establishing a trust relationship with the Management Server as part of Security Engine installation or reconfiguration (see `sg-reconfigure`). The Security Engine contacts the Management Server using the one-time password created when the Security Engine's initial configuration is saved.
`sg-diagnostics [-s|-u] -f <facility_number>`	FW L2FW IPS	Enables or disables diagnostics for the specified facility. When enabled, diagnostic information for the specified facility is included in the log data. `-f <facility_number>` specifies the facility for which to enable diagnostics. Use the `sg-logger -s` command to get a list of facility numbers. `-s` enables diagnostics. `-u` disables diagnostics. When you run the command without `-s` or `-u`, the output shows the current value for the specified facility.

Command	Role	Description

`sg-dynamic-routing`
`[start]`
`[stop]`
`[restart]`
`[force-reload]`
`[backup <file>]`
`[restore <file>]`
`[sample-config]`
`[route-table]`
`[info]`
	FW	`start` starts the Free Range Routing (FRR) routing suite.
		`stop` stops the FRR routing suite and flushes all routes made by zebra.
		`restart` restarts the FRR routing suite.
		`force-reload` forces reload of the saved configuration.
		`backup` backs up the current configuration to a compressed file.
		`restore` restores the configuration from the specified file.
		`sample-config` creates a basic configuration for FRR.
		`route-table` prints the current routing table.
		`info` shows the help information for the `sg-dynamic-routing` command, and detailed information about FRR suite configuration with vtysh.

`sg-ipsec -d`
`[-u <username[@domain]> |`
`-si <session id>|`
`-ck <ike cookie> |`
`-tri <transform id> |`
`-ri <remote ip> |`
`-ci <connection id>]`
	FW	Deletes VPN-related information (use the `vpntool` command to view the information). Option `-d` (for delete) is mandatory.
		`-u` deletes the VPN session of the named VPN client user. You can enter the user account in the form `<user_name@domain>` if there are several user storage locations (LDAP domains).
		`-si` deletes the VPN session of a VPN client user based on session identifier.
		`-ck` deletes the IKE SA (Phase one security association) based on IKE cookie.
		`-tri` deletes the IPSEC SAs (Phase two security associations) for both communication directions based on transform identifier.
		`-ri` deletes all SAs related to a remote IP address in site-to-site VPNs.
		`-ci` deletes all SAs related to a connection identifier in site-to-site VPNs.

Command	Role	Description

`sg-log-view`
`-h | --help`
`-c CONFIGURATION_FILE`
`-C | --show-configuration`
`-N | --show-field-names`
`-A | --alerts`
`-o {list|table|json|json-pretty} | --output-format {list|table|json|json-pretty}`
`-f | --follow`
`-t TABLE_FIELDS [TABLE_FIELDS ...] | --table-fields TABLE_FIELDS [TABLE_FIELDS ...]`
`-a ADD_TABLE_FIELDS [ADD_TABLE_FIELDS ...] | --add-table-fields ADD_TABLE_FIELDS [ADD_TABLE_FIELDS ...]`
`-r REMOVE_TABLE_FIELDS [REMOVE_TABLE_FIELDS ...] | --remove-table-fields REMOVE_TABLE_FIELDS [REMOVE_TABLE_FIELDS ...]`
`-I | --add-event-id-table-field`
`-i EVENT_IDS [EVENT_IDS ...] | --event-ids EVENT_IDS [EVENT_IDS ...]`
`-s START_DATE | --start-date START_DATE`
`-e END_DATE | --end-date END_DATE`
`-F FILTERS [FILTERS ...] | --filters FILTERS [FILTERS ...]`
`--input-file-format {binary|json}`
`--log-files [LOG_FILES [LOG_FILES ...]]`
`--timestamp-type {date|integer}`
`-S | --show-log-counter`
	FW
	L2FW
	IPS
		If you have saved copies of the most recent log and alert entries locally on the Security Engine, allows you to browse log and alert entries on the command line of the Security Engine.
		`-h | --help` shows usage information.
		`-c` specifies a configuration file for viewing stored log entries. If you do not specify a configuration file in this command, the `LOG_VIEW_CONF` environment variable specifies the configuration file. If no configuration file is specified in the `LOG_VIEW_CONF` variable, the default configuration is used.
		`-C | --show-configuration` shows the active configuration.
		`-N | --show-field-names` shows all available log field names.
		`-A | --alerts` shows alert entries instead of log entries.
		`-o | --output-format` specifies the output format for log entries. The default is table.
		`-f | --follow` shows log entries in real time as they are generated.
		`-t | --table-fields` shows the specified fields in a table view. You can specify the width and position of the field in the table using numbers and semicolons. For example, `situation:40:3`.
		`-a | --add-table-fields` adds the specified fields to the table view. You can specify the width and position of the field in the table using numbers and semicolons. For example, `situation:40:3`.
		`-r | --remove-table-fields` removes the specified fields from the table view.
		`-I | --add-event-id-table-field` adds `event_id` as the first log field in the table view.
		`-i | --event-ids` shows details about the specified events (event ids) in a list view.
		`-s | --start-date` shows log entries starting from the specified date.
		`-e | --end-date` shows log entries ending on the specified date.
		`-F | --filters` specifies log filters as either a simple filter string or a complete JSON filter string.
		`--input-file-format` specifies the input log file format. The default is binary.
		`--log-files` specifies the log files to show. If you do not specify a log file, all available log files found in the specified log directories are shown.
		`--timestamp-type` shows timestamp values as dates or integers. The default is date.
		`-S | --show-log-counter` shows log counters in table and list views.

Command	Role	Description

`sg-logger`
`-f FACILITY_NUMBER`
`-t TYPE_NUMBER`
`[-e EVENT_NUMBER]`
`[-i "INFO_STRING"]`
`[-s]`
`[-h]`
	FW
	L2FW
	IPS
		Used in scripts to create log messages with the specified properties.
		`-f` defines the facility for the log message.
		`-t` defines the type for the log message.
		`-e` defines the log event for the log message. The default is `0 (H2A_LOG_EVENT_UNDEFINED)`.
		`-i` defines the information string for the log message.
		`-s` dumps information about option numbers to stdout
		`-h` shows usage information.

`sg-raid`
`[-status] [-add] [-re-add]`
`[-force] [-help]`
	FW
	L2FW
	IPS
		Configures a new hard drive.
		This command is only for Security Engine appliances that support RAID (Redundant Array of Independent Disks) and have two hard drives.
		`-status` shows the status of the hard drive.
		`-add` adds a new empty hard drive. Use `-add -force` if you want to add a hard drive that already contains data and you want to overwrite it.
		`-re-add` adds a hard drive that is already partitioned. This command prompts for the drive and partition for each degraded array. Use `-re-add -force` if you want to check all arrays.
		`-help` shows usage information.

`sg-reconfigure`
`[--maybe-contact]`
`[--no-shutdown]`
`[--stop-autocontact]`
	FW
	L2FW
	IPS
		Starts the Security Engine Configuration Wizard. Used for reconfiguring the node manually.
		CAUTION
		This script also has parameters that are for the internal use of the Security Engine only. Do not use this script with any parameters other than the ones listed here.
		`--maybe-contact` contacts the Management Server if requested. This option is only available on Engines.
		`--no-shutdown` allows you to make limited configuration changes on the node without shutting it down. Some changes might not be applied until the node is rebooted.
		`--stop-autocontact` (unconfigured Security Engine appliances with valid POS codes only) prevents the Security Engine from contacting the installation server for plug-and-play configuration when it reboots.

`sg-selftest [-d] [-h]`	FW	Runs cryptography tests on the Security Engine. `-d` runs the tests in debug mode. `-h` shows usage information.
`sg-status [-l] [-h]`	FW L2FW IPS	Shows information about the Security Engine status. `-l` shows all available information about Security Engine status. `-h` shows usage information.

Command	Role	Description

`sg-toggle-active`
`SHA1 SIZE |`
`--force [--debug ]`
	FW
	L2FW
	IPS
		Switches the Security Engine between the active and the inactive partition.
		This change takes effect when you reboot the Security Engine.
		You can use this command, for example, if you have upgraded an Security Engine and want to switch back to the earlier Security Engine version. When you upgrade the Security Engine, the active partition is switched. The earlier configuration remains on the inactive partition. To see the currently active (and inactive) partition, see the directory listing of `/var/run/stonegate` (`ls -l /var/run/stonegate`).
		The `SHA1` option is used to verify the signature of the inactive partition before changing it to active. If you downgrade the Security Engine, check the checksum and the size of the earlier upgrade package by extracting the signature and size files from the `sg_engine_[version.build]_i386.zip` file.
		`--debug` reboots the Security Engine with the debug kernel.
		`--force` switches the active configuration without first verifying the signature of the inactive partition.

`sg-upgrade`	FW	Upgrades the node by rebooting from the installation DVD. Alternatively, the node can be upgraded remotely using the SMC Client.
`sg-version`	FW L2FW IPS	Shows the software version and build number for the node.

`se-virtual-engine`
`-l | --list`
`-v <Virtual Engine ID>`
`-e | --enter`
`-E "<command [options]>"`
`-h | --help`
	FW	 **Note**
		Master Engine only.
		Used to send commands to Virtual Engines from the command line of the Master Engine.
		All commands that can be used for the Engine role can also be used for Virtual Engines.
		`-l` or `--list` list the active Virtual Engines.
		`-v` specifies the ID of the Virtual Engine on which to execute the command.
		`-e` or `--enter` enters the command shell for the Virtual Engine specified with the `-v` option. To exit the command shell, type `exit`.
		`-E` executes the specified command on the Virtual Engine specified with the `-v` option.
		`-h` or `--help` shows usage information.

Command	Role	Description

`sginfo`
`[-f]`
`[-d]`
`[-s]`
`[-p]`
`[--]`
`[--help]`
	FW
	L2FW
	IPS
		Gathers system information you can send to Forcepoint Customer Hub.
		Use this command only when instructed to do so by Forcepoint Customer Hub.
		`-f` forces sgInfo even if the configuration is encrypted.
		`-d` includes core dumps in the sgInfo file.
		`-s` includes slapcat output in the sgInfo file.
		`-p` includes passwords in the sgInfo file (by default passwords are erased from the output).
		`--` creates the sgInfo file without showing the progress.
		`--help` shows usage information.

The following table lists some general Linux operating system commands that can be useful in running your Security Engines. Some commands can be stopped by pressing **Ctrl+C**.

General command line tools on Security Engines

Command	Description
`dmesg`	Shows system logs and other information. Use the `-h` option to see usage.
`halt`	Shuts down the system.
`ip`	Shows IP address information. Type the command without options to see usage. Example: type `ip addr` for basic information about all interfaces.
`ping`	Tests connectivity with ICMP echo requests. Type the command without options to see usage.
`ps`	Reports the status of running processes.
`reboot`	Reboots the system.
`scp`	Secure copy. Type the command without options to see usage.
`sftp`	Secure FTP. Type the command without options to see usage.
`ssh`	SSH client (for opening a terminal connection to other hosts). Type the command without options to see usage.
`tcpdump`	Gives information about network traffic. Use the `-h` option to see usage. You can also analyze network traffic by creating tcpdump files from the SMC Client with the Traffic Capture feature.
`top`	Shows the top CPU processes taking most processor time. Use the `-h` option to see usage.

Command	Description
`traceroute`	Traces the route packets take to the specified destination. Type the command without options to see usage.
`vpntool`	Shows VPN information and allows you to issue some basic commands. Type the command without options to see usage.

Server Pool Monitoring Agent commands

You can test and monitor the Server Pool Monitoring Agents on the command line.

Server Pool Monitoring Agent commands

Command	Description

`agent`
`[-v level]`
`[-c path]`
`[test [files]]`
`[syntax [files]]`
	(Windows only) Allows you to test different configurations before activating them.
	`-v` sets the verbosity level. The default level is 5. Levels 6–8 are for debugging where available.
	`-c` uses the specified path as the first search directory for the configuration.
	`test` runs in the test mode - status queries do not receive a response. If you specify the files, they are used for reading the configuration instead of the default files.
	`syntax` checks the syntax in the configuration file. If no files are specified, the default configuration files are checked.

`sgagentd [-d]`
`[-v level]`
`[-c path]`
`[test [files]]`
`[syntax [files]]`
	(Linux only) Allows you to test different configurations before activating them.
	`-d` means Don't Fork as a daemon. All log messages are printed to stdout or stderr only.
	`-v` sets the verbosity level. The default level is 5. Levels 6–8 are for debugging where available.
	`-c` uses the specified path as the first search directory for the configuration.
	`test` runs in the test mode - status queries do not receive a response. If you specify the files, they are used for reading the configuration instead of the default files. The output is directed to syslog or eventlog instead of the console where the command was run unless you use the `-d` option.
	`syntax` checks the syntax in the configuration file. If no files are specified, the default configuration files are checked. The output is directed to syslog or eventlog instead of the console where the command was run unless you use the `-d` option.

Command	Description

`sgmon`
`[status|info|proto]`
`[-p port]`
`[-t timeout]`
`[-a id]`
`host`
	Sends a UDP query to the specified host and waits for a response until received, or until the timeout limit is reached. The request type can be defined as a parameter. If no parameter is given, status is requested. The commands are:
	`status` queries the status.
	`info` queries the agent version.
	`proto` queries the highest supported protocol version.
	`-p` connects to the specified port instead of the default port.
	`-t` sets the timeout (in seconds) to wait for a response.
	`-a` acknowledge the received log messages up to the specified id. Each response message has an id, and you can acknowledge more than one message at a given time by using the id parameter. Messages acknowledged by sgmon will no longer appear in the engine logs.
	`host` is the IP address of the host to connect to. To get the status locally, you can give localhost as the host argument. This parameter is mandatory.

Appendix B

Default communication ports

Contents

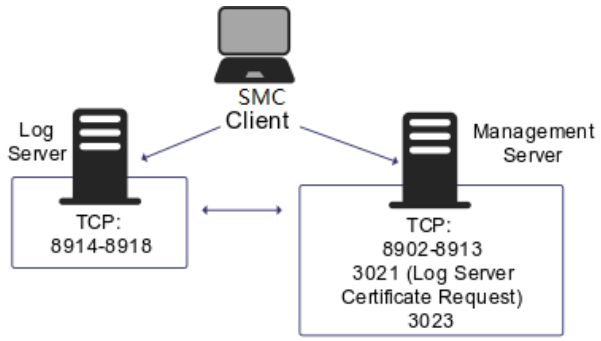
- Forcepoint Security Management Center ports on page 1457
- Security Engine ports on page 1460

There are default ports used in connections between SMC components and default ports that SMC components use with external components.

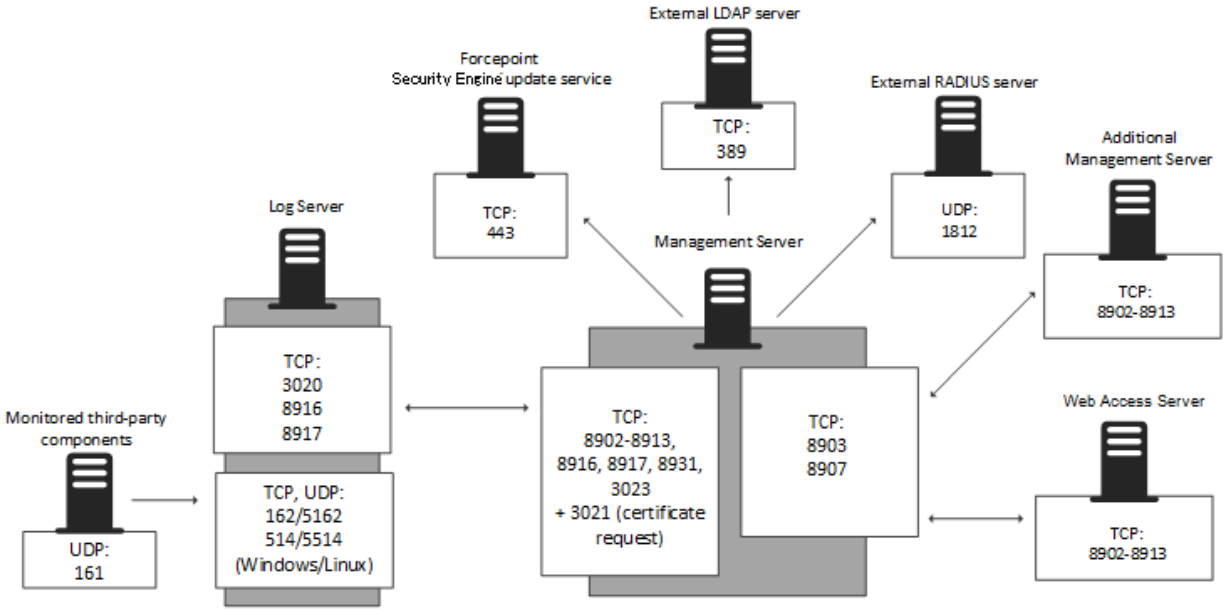
Forcepoint Security Management Center ports

The most important default ports used in communications to and from SMC components are presented in the following illustrations.

Destination ports for basic communications within the SMC



Default destination ports for optional SMC components and features



This table lists the default ports SMC uses internally and with external components. Many of these ports can be changed. The names of corresponding default Service elements are also included for your reference.

SMC default ports

Listening host	Port/protocol	Contacting hosts	Service description	Service element name
Additional Management Servers	8902- 8913/TCP	Management Server	Database replication (push) to the additional Management Server.	SG Control
DNS server	53/UDP, 53/TCP	SMC Client, Management Server, Log Server	DNS queries.	DNS (UDP)
LDAP server	389/TCP	Management Server	External LDAP queries for display/ editing in the SMC Client.	LDAP (TCP)
Log Server	162/UDP, 5162/ UDP	Monitored third-party components	SNMPv1 trap reception from third-party components. Port 162 is used if installed on Windows, port 5162 if installed on Linux.	SNMP (UDP)
Log Server	514/TCP, 514/ UDP, 5514/TCP, 5514/UDP	Monitored third-party components	Syslog reception from third-party components. Port 514 is used if installed on Windows, port 5514 if installed on Linux.	Syslog (UDP) [Partial match]
Log Server	2055/UDP	Monitored third-party components	NetFlow or IPFIX reception from third-party components. Port 2055 is used in both Windows and Linux.	NetFlow (UDP)

Listening host	Port/protocol	Contacting hosts	Service description	Service element name
Log Server	3020/TCP	Log Server, Web Access Server, Security Engines	Alert sending from the Log Server and Web Access Server. Log and alert messages; monitoring of block lists, connections, status, and statistics from Security Engines.	SG Log
Log Server	8914-8918/TCP	SMC Client	Log browsing.	SG Data Browsing
Log Server	8916-8917/TCP	Log Server, Web Access Server	Database replication (push) to the Log Server; Log browsing on the Web Access Server.	SG Data Browsing (Web Access Server)
Management Server	3021/TCP	Log Server, Web Access Server	System communications certificate request/renewal.	SG Log Initial Contact
Management Server	8902-8913/TCP	SMC Client, Log Server, Web Access Server	Monitoring and control connections.	SG Control
Management Server	3023/TCP	Additional Management Servers, Log Server, Web Access Server	Log Server and Web Access Server status monitoring. Status information from an additional Management Server to the active Management Server.	SG Status Monitoring
Management Server	8903, 8907/TCP	Additional Management Servers	Database replication (pull) to the additional Management Server.	SG Control
Management Server	8085/TCP	Web Access clients	Communication for using Web Access.	HTTPS
Monitored third-party components	161/UDP	Log Server	SNMP status probing to external IP addresses.	SNMP (UDP)
NTP server	123/TCP or UDP	SMC Appliance	Receiving NTP information.	NTP
RADIUS server	1812/UDP	Management Server	RADIUS authentication requests for administrator logon. The default ports can be edited in the properties of the RADIUS Server element.	RADIUS (Authentication)
Security Engine update service	443/TCP	SMC servers	Update packages, engine upgrades, and licenses.	HTTPS
SMC Appliance	161/UDP	Third-party components	Requesting health and other information about the SMC Appliance.	SNMP
Update servers	443/TCP	SMC Appliance	Receiving appliance patches and updates.	HTTPS
SMC Appliance	22/TCP	Terminal clients	SSH connections to the command line of the SMC Appliance.  **Note** Do not use SSH in FIPS mode.	SSH

Listening host	Port/protocol	Contacting hosts	Service description	Service element name
Syslog server	514/UDP, 5514/UDP	Log Server	Log data forwarding to syslog servers. The default ports can be edited in the LogServerConfiguration.txt file.	Syslog (UDP) [Partial match]
Terminal Client Engine, Layer 2 Engine, IPS, Master Engine	22/TCP	SMC Appliance	Contacting engines and moving SMC Appliance backups off the appliance. **Note** Do not use SSH in FIPS mode.	SSH
Third-party components	2055/UDP	Log Server	NetFlow or IPFIX forwarding to third-party components. Port 2055 is used in both Windows and Linux.	NetFlow (UDP)
Third-party components	162/UDP	SMC Appliance	Sending SNMP status probing to external devices.	SNMP
Third-party components	445/TCP	SMC Appliance	Moving SMC Appliance backups off the appliance. **Note** You cannot use CIFS in FIPS mode.	CIFS
Web Access Server	8083/TCP	Web Access clients	Communication for using Web Access.	HTTPS

Security Engine ports

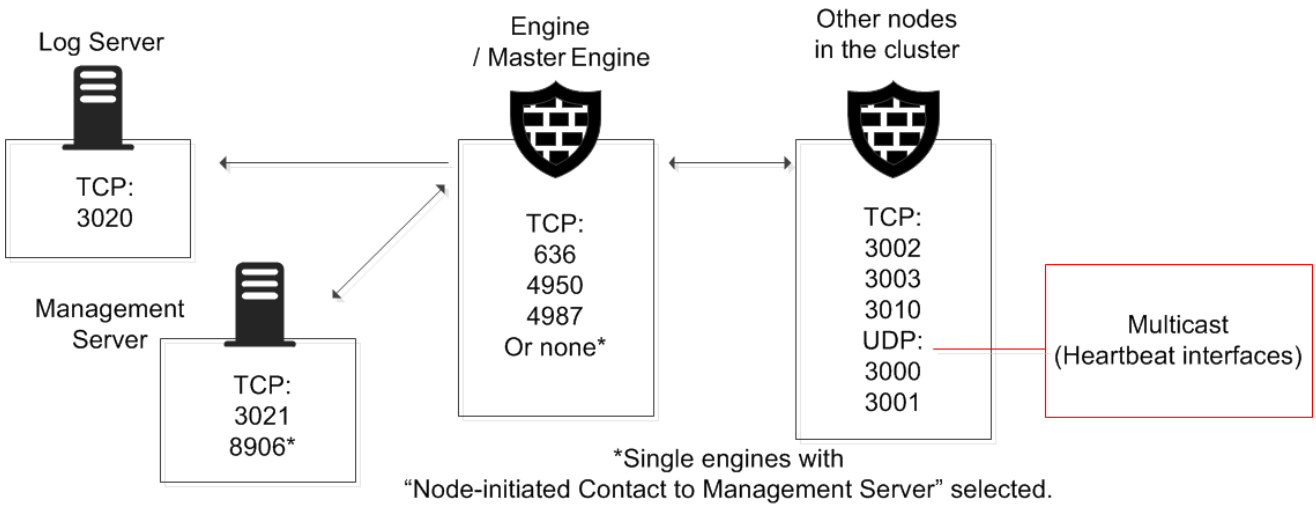
The most important default ports used in communications to and from Security Engines and Master Engines are presented in the following illustrations.

See the table for a complete list of default ports for the engines.

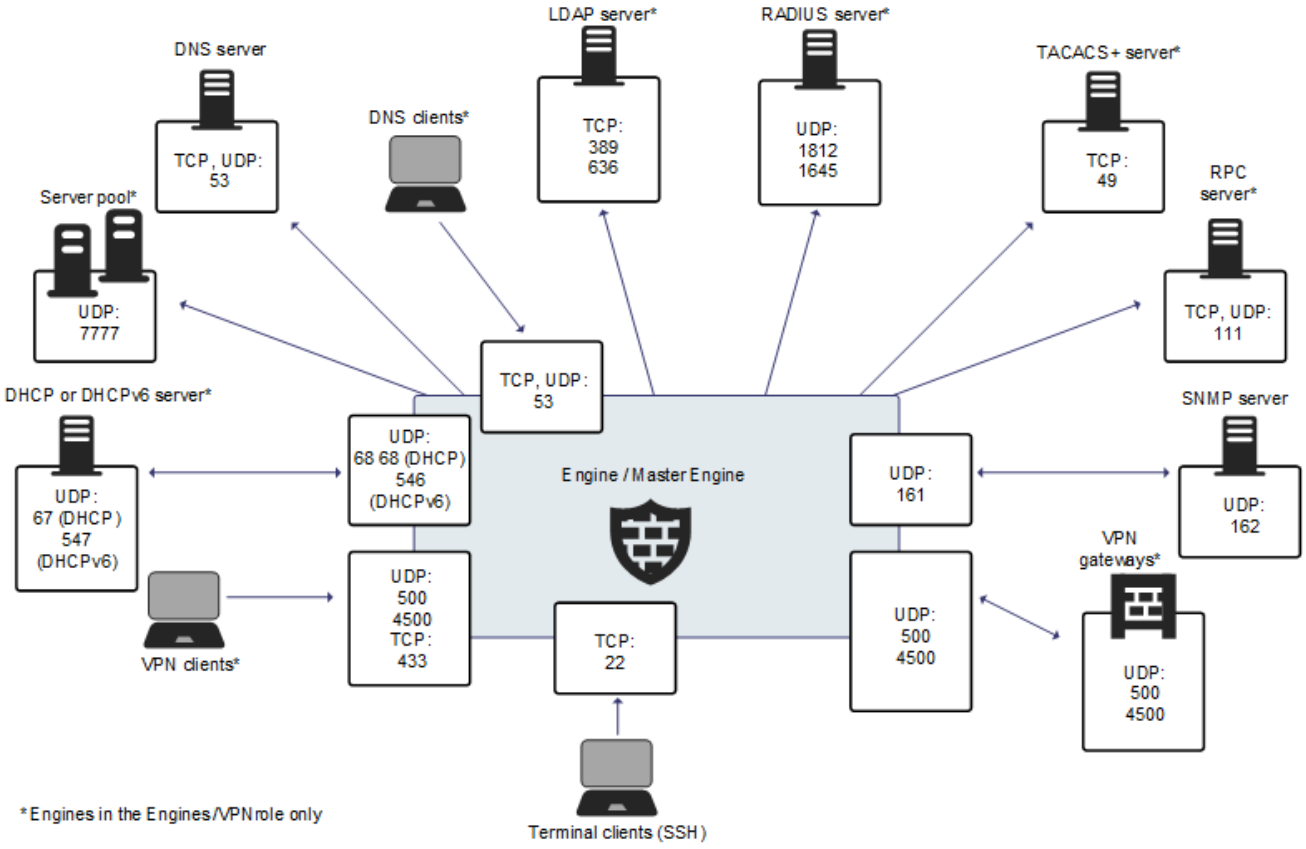
Note

Master Engines use the same default ports as clustered Security Engines. Virtual Engines do not communicate directly with other system components.

Destination ports for basic Security Engine communications



Default destination ports for Security Engine service communications



This table lists the default ports for Security Engines and Master Engines. Many of these ports can be changed. The names of corresponding default Service elements are also included for your reference.

Security Engine and Master Engine default ports

Listening host	Port/protocol	Contacting hosts	Service description	Service element name
Certificate Revocation List (CRL) server	80/TCP	Engine	Online certificate status protocol (OCSP) queries and fetching CRLs.	HTTP
DHCP server	67/UDP	Engine	Relayed DHCP requests and requests from a engine that uses dynamic IP address.	BOOTPS (UDP)
DHCPv6 server	547/UDP	Engine	Requests from a engine that uses dynamic IPv6 address.	N/A
External DNS server	53/UDP, 53/TCP	Engine, Master Engine	DNS resolution and dynamic DNS updates.	DNS (TCP), DNS (UDP)
File reputation server	443/TCP	Engine, Layer 2 Engine, IPS, Master Engine	GTI File Reputation Server	HTTPS
Engine	67/UDP	Any	DHCP relay on engine.	BOOTPS (UDP)
Engine	68/UDP	DHCP server	Replies to DHCP requests.	BOOTPC (UDP)
Engine	80/TCP	Clients that need to authenticate to the Engine	Browser Based User Authentication	HTTP
Engine	443/TCP	Clients that need to authenticate to the Engine	Browser Based User Authentication	HTTPS
Engine	443/TCP	VPN clients using SSL tunneling	VPN client SSL tunneling	TLS
Engine	443/TCP	SSL Portal users	SSL VPN Portal	HTTPS
Engine	546/UDP	DHCPv6 server	Replies to DHCPv6 requests.	N/A
Engine, Master Engine	53/UDP, 53/TCP	Clients in the internal network	DNS relay	DNS (TCP), DNS (UDP)
Engine, Master Engine	500/UDP	VPN clients, VPN gateways	VPN negotiations, VPN traffic.	ISAKMP (UDP)
Engine, Master Engine	636/TCP	Management Server	Internal user database replication.	LDAPS (TCP)
Engine, Master Engine	4500/UDP	VPN client, VPN gateways	VPN traffic using NAT-traversal.	NAT-T
Engine Cluster Node, Master Engine cluster node	3000-3001/UDP, 3002–3003, 3010/TCP	Engine Cluster Node, Master Engine cluster node	Heartbeat and state synchronization between clustered Engines.	SG State Sync (Multicast), SG State Sync (Unicast), SG Data Sync
Engine, Layer 2 Engine, IPS, Master Engine	22/TCP	Terminal clients	SSH connections to the engine command line. **Note** Do not use SSH in FIPS mode.	SSH
Engine, Layer 2 Engine, IPS, Master Engine	4950/TCP	Management Server	Remote upgrade.	SG Remote Upgrade

Listening host	Port/protocol	Contacting hosts	Service description	Service element name
Engine, Layer 2 Engine, IPS, Master Engine	4987/TCP	Management Server	Management Server commands and policy upload.	SG Commands
Engine, Layer 2 Engine, IPS, Master Engine	15000/TCP	Management Server, Log Server	block list entries.	SG block listing
Engine, Layer 2 Engine, IPS, Master Engine	161/UDP	SNMP server	SNMP monitoring.	SNMP (UDP)
Engine, Layer 2 Engine, IPS	9111/TCP	Forcepoint One Endpoint client	Endpoint information from the Forcepoint One Endpoint client.	N/A
Forcepoint User ID Service server	5000/TCP	Engine, Layer 2 Engine, IPS	Information about user name and IP address mappings.	N/A
IPS Cluster Node	3000-3001/UDP, 3002–3003, 3010/TCP	IPS Cluster Node	Heartbeat and state synchronization between clustered IPS engines.	SG State Sync (Multicast), SG State Sync (Unicast), SG Data Sync
LDAP server	389/TCP	Engine, Master Engine	External LDAP queries, including StartTLS connections.	LDAP (TCP)
Layer 2 Engine Cluster Node	3000-3001/UDP, 3002–3003, 3010/TCP	Layer 2 Engine Cluster Node	Heartbeat and state synchronization between clustered Layer 2 Engines.	SG State Sync (Multicast), SG State Sync (Unicast), SG Data Sync
Log Server	3020/TCP	Engine, Layer 2 Engine, IPS, Master Engine	Log and alert messages; monitoring of block lists, connections, status, and statistics.	SG Log
Malware signature server	80/TCP	Engine, Layer 2 Engine, IPS, Master Engine	Malware signature update service.	HTTP
Management Server	3021/TCP	Engine, Layer 2 Engine, IPS, Master Engine	System communications certificate request/renewal (initial contact).	SG Initial Contact
Management Server	8906/TCP	Engine, Layer 2 Engine, IPS	Management connection for engines with "Node-Initiated Contact to Management Server" selected.	SG Dynamic Control
RADIUS server	1812, 1645/UDP	Engine, Master Engine	RADIUS authentication requests.	RADIUS (Authentication), RADIUS (Old)
RPC server	111/UDP, 111/TCP	Engine, Master Engine	RPC number resolve.	SUNRPC (UDP), Sun RPC (TCP)
Server Pool Monitoring Agents	7777/UDP	Engine, Master Engine	Polls to the servers' Server Pool Monitoring Agents for availability and load information.	SG Server Pool Monitoring
SNMP server	162/UDP	Engine, Layer 2 Engine, IPS, Master Engine	SNMP traps from the engine.	SNMP Trap (UDP)
TACACS+ server	49/TCP	Engine, Master Engine	TACACS+ authentication requests.	TACACS (TCP)
ThreatSeeker Intelligence Cloud server	443/TCP	Engine, Layer 2 Engine, IPS, Master Engine	ThreatSeeker Intelligence Cloud URL categorization service.	HTTPS

Listening host	Port/protocol	Contacting hosts	Service description	Service element name
VPN gateways	500, 4500/UDP	Engine, Master Engine	VPN traffic. Ports 443/TCP (or custom port) can also be used, depending on encapsulation options.	ISAKMP (UDP)

Appendix C

Working with expressions

Contents

- Introduction to expressions on page 1465
- Using operands on page 1466
- Expression processing order on page 1468
- Grouping operands using parentheses on page 1468
- Nesting expressions on page 1469

Expressions are elements that allow you to create simple definitions for representing complex sets of IP addresses by using logical operands.

Introduction to expressions

An *expression* is an element that combines other network elements (IP addresses) with logical operands.

Expressions make it easier to define complex sets of network resources, even though you can arrive at the same definitions without expressions. For example, a single, simple expression can include a whole network except for a few individual IP addresses scattered throughout the address space. Otherwise, several Address Range elements might be needed for defining the same set of IP addresses.

The expressions consist of the following parts:

- *Parentheses* group sets of elements and define the processing order in the same way as they do in mathematical equations. The parentheses in expressions are always the basic curved type “(” and “)”.
- *Negation* operators take a set and form a new set that includes every possible element except the ones in the original set. Negations are expressed with “~”.
- *Intersection* operators take two sets and forms a new set that includes only those IP addresses that are found in both sets. Intersections are expressed with “#”.
- *Union* operators combine two sets and form a new set that includes every IP address in both sets. Unions are expressed with “#”.

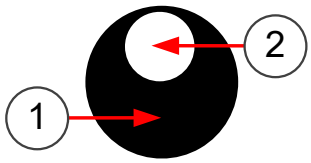
Using operands

When creating expressions, you can use the negation, intersection, and union operands.

Negation

The negation operand can be understood based on common language use: it corresponds to the word “NOT”.

Graphical representation of a negation



1. Address space to which the negation applies
2. The negation excludes a subset

Example

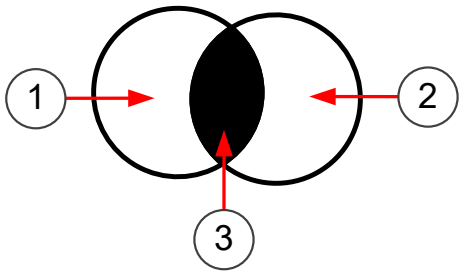
~1.2.3.4 (negation of IP address 1.2.3.4) includes all other possible (IPv4) addresses except the IP address 1.2.3.4. As you see, negations are a good way to create a simple element that includes large IP address spaces with some exceptions. Usually, the negation appears in constructions like the following: 192.168.10.0/24 # ~192.168.10.200. This example basically means “include all addresses in network 192.168.10.0/24, but do not include address 192.168.10.200”.

This definition uses the intersection operand, which is explained next. We return to this same example to explain the intersection part of the equation. Also, the section explaining the union operand returns to this example once more to explain why a union operand is not appropriate here.

Intersection

Intersection means “include only those IP addresses that are a part of both sets”.

Graphical representation of an intersection



1. Address Space A
2. Address Space B
3. IP addresses that are in common. This is the intersection.

Example

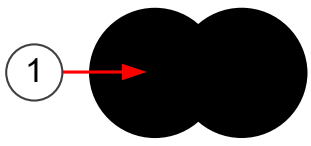
We could intersect two address ranges, A (192.168.10.200 – 192.168.10.300) and B (192.168.10.250 – 192.168.10.350). The expression reads A # B and it resolves to the following IP addresses: 192.168.10.250 – 192.168.10.300 (the IP addresses that appear in both ranges).

We now return to the previous example on the negation operand, where an intersection was also used: 192.168.10.0/24 # ~192.168.10.200. On the left side, there is a specific network that we intersect with the right side that contains all possible IP addresses except one IP address. The intersection resolves to the IP addresses that the left side and the right side have in common. These addresses include the IP addresses in network 192.168.10.0/24 except the one IP address that is excluded on the right side of the equation. As shown here, intersections allow us to make expressions more specific.

Union

The common language equivalent for the union operand is the word “AND”. The union operand’s role is to widen the scope of the expression.

Graphical representation of a union



1. The union includes all IP addresses in both sets combined

Example

The expression 1.2.3.0/24 # 2.3.4.0/24 includes all IP addresses in the two networks. As is evident from this example, using unions is the same as including the elements in the same Group element. For this reason, unions are not the only operand in an expression. Perhaps a better example is: ~192.168.1.1 # ~192.168.1.255 (includes

all IP addresses except the two IP addresses mentioned). However, this example is wide in scope and might require further restriction to be practical. For example, adding an intersection with the network 192.168.1.0/24 restricts the result to addresses in that network.

Unions do have the potential to become too wide in scope if you are not careful. In the preceding example, we used the example expression: 192.168.10.0/24 # ~192.168.10.200. If we replace the intersection (“#”) with a union (“#”), the expression then includes all addresses from the left side (network 192.168.10.0/24) and the right side (all IPv4 addresses except for one). The expression includes even the single IP address that is excluded on the right side because it is part of the network on the left side. The result corresponds to the default “Any Network” element that matches all possible IP addresses. The processing order of the operands is also a factor in this result.

Expression processing order

The processing order of expressions is fixed.

As in mathematical equations, items inside parentheses are always resolved before other comparisons. Next, the operands are processed by type: first the negations, then intersections, and last the unions.

For example, the expression A # ~(B # C) # D is processed as follows:

1) The formula between parentheses is solved first (the union of B and C). If we replace this result with X, the expression becomes A # ~X # D.

2) The negation is processed next, inverting the value of X. If we replace this result with Y, the expression becomes A # Y # D.

3) Next, the intersection between Y and D is resolved. If we replace this result with Z, the expression becomes A # Z.

4) Finally, the union of A and Z yields the actual value that the expression represents (the full contents of both A and Z).

As shown here, the order in which the operand-value combinations appear in the expression have no significance to the order of processing. The only way to change the processing order is by using parentheses as explained next.

Grouping operands using parentheses

Parentheses allow grouping the expression so that the operands you add are processed in a non-standard order.

Operands inside parentheses are always processed before other operands. Parentheses can also be placed inside parentheses, in which case the operands are processed starting from the innermost parentheses.

We can change the earlier example by adding a set of parentheses like this: (A # ~(B # C)) # D. With the two sets of parentheses, the inner parentheses are processed first (B # C as before), and the negation is processed next (~X as before). However, the outer parentheses are processed next instead of processing the intersection, changing the result. If ~(B # C) results in Y (as before), the expression becomes (A # Y) # D. The order of processing is then different than without the parentheses: instead of intersecting Y and D, the expression performs a union of A and Y, and the intersection is then the last operand to be processed.

Complicated expressions with extensive use of parentheses can become difficult to read and edit. In these situations, nested expressions might sometimes be a better option.

Nesting expressions

You can nest expressions by placing other expressions inside an expression. Nesting is a good way to simplify the creation of complex expressions.

When you construct complex expressions from smaller incremental units, you can more easily find and fix problems. You can also reuse the smaller units in other expressions or policies as appropriate. Arguably, the expressions are also easier to read and edit when broken down into smaller units.

For example, if we want to create an expression that includes all IP addresses in three networks, except for one IP address in each, we have a long expression: (192.168.1.0/24 # ~192.168.1.1) # (192.168.2.0/24 # ~192.168.2.1) # (192.168.3.0/24 # ~192.168.3.1).

Instead of creating just one expression, it might make more sense to create several expressions: one for each set of parentheses (for example, Expression A: 192.168.1.0/24 # ~192.168.1.1) and then add an expression that collects those three expressions together as follows: Expression A # Expression B # Expression C, or simply create a Group element that contains the three expressions.

All three subexpressions can be used individually or easily combined in other ways as needed, for example, Expression A # Expression C. Naturally, when changes are made to an expression used inside some other expression, the definitions are updated in both places.

You can also create expressions that you use as templates for creating new expressions: when you insert an expression into another expression, you have the choice of extracting the contents from the expression instead of just inserting the expression. Extracting the contents allows you to make further changes and additions to the expression you insert. Extracting the contents also removes the link between the expressions, so changes are not propagated if the inserted expression is later changed.

Appendix D

Predefined Aliases

Contents

- Predefined User Aliases on page 1471
- System Aliases on page 1471

Predefined Aliases are used in the default policies. Some of them might be useful when you create your own rules.

Predefined User Aliases

User Aliases are usually created by administrators, but there are also some predefined User Aliases in the SMC.

User Aliases are preceded with one $ character. The following table lists all editable automatically created User Aliases.

System-defined User Aliases

Predefined User Alias	Description
$ DHCP address pools	Addresses that can be allocated by DHCP servers.
$ DHCP address pools for Mobile VPN clients	Address pools for assigning virtual IP addresses to mobile VPN clients.
$ DHCP servers	All DHCP servers defined for the Engine.
$ DHCP servers for Mobile VPN clients	The DHCP servers defined for assigning virtual IP addresses to mobile VPN clients.

System Aliases

System Aliases are automatically created non-editable Aliases.

The System Aliases are preceded with two $$ characters. The following table lists the definitions of all System Aliases. These Aliases are used in the Engine's default policies.

System Aliases

System Alias	Description
$$ DHCP Enabled Interface Addresses	IP addresses (of CVIs on clusters) which have DHCP relay enabled.
$$ DHCP Enabled interface addresses for Mobile VPN clients	IP addresses (of NDIs on clusters) which have DHCP relay enabled for mobile VPN clients.
$$ DHCP Interface X.dns	IP address of the DHCP-assigned DNS server for the interface with DHCP index number X.

System Alias	Description
$$ DHCP Interface X.gateways	IP address of the DHCP-assigned default router for the interface with DHCP index number X.
$$ DHCP Interface X.ip	DHCP-assigned IP address for the interface with DHCP index number X.
$$ DHCP Interface X.net	Directly connected network behind the interface with DHCP index number X.
$$ Interface ID X.ip	First IP address (CVI) of Physical Interface ID X.
$$ Interface ID X.net	Directly connected networks behind Physical Interface ID X.
$$ Local Cluster	All addresses of the cluster.
$$ Local Cluster (CVI addresses only)	All CVI addresses of the cluster.
$$ Local Cluster (DHCP Interface Addresses)	All DHCP-assigned IP addresses of the engine.
$$ Local Cluster (NDI addresses only)	All NDI addresses of all nodes in the cluster.
$$ Local Cluster (NDI for heartbeat addresses only)	Heartbeat NDI addresses of all nodes in the cluster.
$$ Local Cluster (NDI for management addresses only)	Management NDI addresses of all nodes in the cluster.
$$ Log Servers	IP addresses of all Log Servers.
$$ Management Servers	IP addresses of all Management Servers.
$$ Valid DHCP Address Pools for Mobile VPN clients	Address pools defined for assigning virtual IP addresses to mobile VPN clients.
$$ Valid DHCP Servers	All DHCP servers defined for the Engine.
$$ Valid DHCP Servers for Mobile VPN clients	The DHCP servers defined for assigning virtual IP addresses to mobile VPN clients.

Appendix E

Situation Context parameters

Contents

- Correlation Context parameters on page 1473
- Other Context parameters on page 1476

There are parameters you can define for Situation Contexts.

Note

The details related to the Contexts in your system might be different from what is described here. The Contexts might have been updated through dynamic update packages after this guide was published. Read the Release Notes of each update package you import to see which elements are affected.

Correlation Context parameters

You can configure Correlation Context parameters to find patterns in event data.

Event Compress

Event Compress combines repeated similar events into the same log entry, reducing clutter in the **Logs** view.

Event Compress parameters

Field	Option (if any)	Explanation
Correlated Situations		Situations you want to compress.
Time Window		All matches to the Situations selected are combined to a common log entry when they are triggered within the defined time from each other.
Log Fields Enabled	Select	Events triggered by the selected Situations are regarded as the same when the values those entries have in the Log Fields you place in Event Binding are identical.
	Ignore	Events triggered by the selected Situations are regarded as the same, except when the values those entries have in the Log Fields you place in Event Binding are identical.
Event Binding		The selected log fields are used by the matching option you selected in the previous step.

Field	Option (if any)	Explanation
Location	Very Early	The execution order of the Compress operation in relation to other operations. Compress operations that share the Location are executed in parallel; each compress operation receives the same events as the other compress operations in the same Location. “Very Early” and “Early” locations can affect the operation of other Correlations.
	Early	
	Late	
	Very Late	
Compress Filter		Filters in data for the compression.

Event Count

Event Count finds recurring patterns in traffic by counting the times certain Situations occur within the defined period, so that action can be taken if the threshold values you set are exceeded.

Event Count parameters

Field	Option (if any)	Explanation
Correlated Situations		Situations you want to count.
Time Window		The period of time within which the matches to the Situation must occur the specified number of times.
Alarm Threshold		The number of times that the selected Situations must occur for the Correlation Situation to match.
Log Fields Enabled	Select	Events triggered by the selected Situations are regarded as the same when the values those entries have in the Log Fields you place in Event Binding are identical.
	Ignore	Events triggered by the selected Situations are regarded as the same, except when the values those entries have in the Log Fields you place in Event Binding are identical.
Event Binding		The selected log fields are used by the matching option you selected in the previous step.

Event Group

Event Group finds event patterns in traffic by following if all events in the defined set of Situations match at least once in any order within the defined time period.

Event Group parameters

Field	Option (if any)	Explanation
Member (column)	Event Match	Filter for grouping.
	Needed Number	How many occurrences of the Event selected for this Member are required for them to be included in the grouping.
	Binding	Log field used for the grouping.

Field	Option (if any)	Explanation
Correlated Situations		Situations you want to group.
Keep and Forward Events	Yes	Makes the Correlation Situation examine the events and trigger the response defined in the Inspection Policy but does not actually group the matching events into one. All individual events are still available for further inspection, even though they have already triggered a response.
	No	Makes the Correlation Situation group the matching events together. Only the response defined in the Inspection Policy is triggered, and no further processing is done on the individual events.
Time Window Size		The period of time within which the Situation must occur for them to be grouped.
Continuous Responses	Yes	Makes the Security Engine or Log Server respond as defined in the Inspection Policy to each occurrence of the defined event within the selected Time Window.
	No	Makes the Security Engine or Log Server respond only to the first occurrence of the defined event within the selected Time Window.

Event Match

Event Match allows filtering event data produced by specific Situations using Filter expressions.

Event Match parameters

Field	Explanation
Correlated Situations	Situations you want the Correlation Situation to match.
Filter	Filter for finding a pattern in the event data.

Event Sequence

Event Sequence finds event patterns in traffic by following if all events in the defined set of Situations match in a specific order within the defined time period.

Event Sequence parameters

Field	Option (if any)	Explanation
Entry to/Exit from (columns)	Event Match	Filter for selecting data for the sequencing.
	Binding	Log field that the Correlation Situation traces to find a sequence.
Correlated Situations		Situations from which you want to find sequences.

Field	Option (if any)	Explanation
Keep and Forward Events	Yes	Makes the Correlation Situation examine the events and trigger the response defined in the Inspection Policy but does not actually group the matching events into one. All individual events are still available for further inspection, even though they have already triggered a response.
	No	Makes the Correlation Situation group the matching events together. Only the response defined in the Inspection Policy is triggered, and no further processing is done on the individual events.
Time Window Size		The period of time within which the Situation must occur for them to be reagarded as a sequence.

Other Context parameters

See the properties dialog box of the Context in question.

The Contexts are shown as branches/sub-branches in the **Other Elements** > **Situations** > **By Context** tree in the **Engine Configurations** view.

Appendix F

Regular expression syntax

Contents

- SMC regular expression syntax on page 1477
- Special character sequences on page 1479
- Pattern-matching modifiers on page 1480
- Variable expression evaluation on page 1482
- Stream operations on page 1484
- System variables on page 1485
- Independent subexpressions on page 1487
- Parallel matching groups on page 1487
- Tips for working with regular expressions on page 1488

The SMC has its own regular expression syntax. Regular expressions are used in Situations for matching network traffic. Situations are used in the Inspection rules on Security Engines.

SMC regular expression syntax

A regular expression is a sequence of characters that defines a matching pattern. These patterns are used for matching byte sequences in network traffic.

The expression matching always starts from the beginning of the traffic stream, defined by the associated Situation Context. Depending on the context, this can mean:

- The beginning of a TCP stream.
- The beginning of a UDP packet.
- A protocol-specific field or header, such as the beginning of an HTTP request header or the beginning of an HTTP Request URI.

A regular expression consists of one or more branches that are separated by a logical OR symbol “|”. A Situation match occurs if any of the branches matches the traffic stream.

Regular expression matching

```
# This regular expression matches
# if any of the following patterns are seen
# at the beginning of the traffic stream: "aaa", "bbb", "ccc"
aaa|bbb|ccc
```

The basic sequences that can be used in an SMC regular expression are listed in the following table:

SMC regular expression syntax

Sequence	Description	Example
`<char>`	Matches only the defined characters.	"2", "A", "foo" match exactly to the defined characters: "2", "A", and "foo" respectively.
`. (dot)`	Matches any character, including the null character `\x00` and a missing character. Matches also other than printable characters, such as the linefeed. A missing character is a special character used by the engine to represent characters missing from a TCP connection. For example, in capture mode, the engine might not see all traffic of a TCP connection.	"`.`" matches any single character or byte.
`\x<hex>`	Matches the hexadecimal byte value ranging from `\x00` to `\xFF`.	"`\x4d`" matches hexadecimal value "4d" which represents the decimal value 77 and the ASCII character "M".
`[<char>]`	Matches any single character in the list.	"`[15aB]`" matches when any of the characters " 1 ", "5", "a", or "B" are in the matching location of the inspected string.
`[^<char>]`	Matches any single character that is not on the list.	"`[^aBc]`" matches if none of the characters "a ", "B", or "c" is present in the matching location of the inspected string.
`[<char1>-<char2>]`	Matches all characters ranging from `<char1>` to `<char2>`, these two characters included.	"`[a-f]`" matches any character within the range from "a" to "f ", with "a" and "f" included.
`\<char>`	Used for escaping special metacharacters to be interpreted as normal characters. The metacharacters are: `|)(][^-*+?.#`	"`\[`" matches the "`[`" character instead of interpreting it as the regular expression class metacharacter.
`#<text>`	Anything starting with "`#`" up to the linefeed (`\x0a`) or the carriage return character (`\x0d`) is regarded as a comment and not used in the matching process.	"`# my comment.`" is not used in the matching process.
`(<expr1>|<expr2>)`	Matches if either expression `<expr1>` or `<expr2>` matches.	"`a(bc|de)`" matches "abc" and "ade".

Example regular expressions

```
# This regular expression matches any of the following strings:
# "login.php", "login1.php", "login2.php", "login_internal.php"
# Note: to match the "." character, the character must be escaped in the
# regular expression by prefixing the character with "\"
login\.php|login[12]\.php|login_internal\.php

# Alternatively, the branches of the above regular expression can be
# combined into one single branch as shown below
login([123]|_internal)?\.php
```

It is also possible to indicate repeated, consecutive characters, or regular expressions using quantifiers. The quantifiers available in SMC regular expression syntax are listed in the following table.

SMC regular expression quantifiers

Sequence	Description	Example
`<expr>*`	Matches if there are zero or more consecutive `<expr>` strings.	"`a*`" matches "`<empty>`", "`a`", "`aa`" and so on.
`<expr>+`	Matches if there are one or more consecutive `<expr>` strings.	"`a+`" matches "`a`", "`aa`", "`aaa`" and so on, but not the empty string.
`<expr>?`	Matches if there is zero or one `<expr>` string.	"`a?`" matches "`<empty>`" and "`a`".
`<expr>{n,m}`	`{num}` matches exactly `num` times the expression. `{num,}` matches `num` or more times the expression. `{num,max}` matches at least `num` and no more than `max` times the expression.	" `a{5,}`" matches five or more consecutive "`a`" characters. "`a{5,7}`" matches 5, 6, or 7 consecutive "`a`" characters.

The quantifiers always apply only to the single previous character (or special character sequence), unless otherwise indicated by parentheses. For example, the regular expression “login*” matches “logi”, “login” or “loginnnn”, whereas the regular expression “(login)*“ matches the empty string “”, “login” or “loginloginlogin”.

As the matching of a regular expression is always started from the beginning of the traffic stream, “`.*`” (any character zero or more times) is often needed when writing SMC regular expressions. For example, the regular expression “`.*/etc/passwd`” searches for the string “`/etc/passwd`” anywhere in the traffic stream.

Note

Use the wildcard characters '`*`' and '`+`', as well as '`<expr>{n,m}`' (where m has a large value) with care. If used in the middle of a regular expression, they can result in an expression that has a very large number of matching states, and that is too complex for efficient use. It is recommended to use these wildcards only in the beginning of a branch.

Special character sequences

Printable characters, such as “a” or “b”, are defined by simply typing them into a regular expression. In addition, there are some shorthands for common non-printable characters and character classes.

Special character sequences are listed in the following table:

Special character sequences

Sequence	Description
`\a`	Bell (BEL) = `\x07`
`\t`	Horizontal tab (HT) = `\x09`
`\n`	Linefeed (LF) = `\x0A`
`\f`	Formfeed (FF) = `\x0C`
`\r`	Carriage return (CR) = `\x0D`
`\e`	Escape (ESC) = `\x1B`
`\OOO`	Octal code `OOO` of the character.

Sequence	Description
`\xHH`	Hexadecimal code `HH` of the character. Case-insensitive. For example, "`\xaa`" is regarded to be the same as "`\xAA`".
`\c<char>`	Control character that corresponds to `Ctrl+<char >`, where `<char>` is an uppercase letter.
`\w`	"word" class character = `[A-Za-z0-9_]`
`\W`	Non-"word" class character = `[^A-Za-z0-9_]`
`\s`	Whitespace character = `[ \t\r\n\f]`
`\S`	Non-whitespace character = `[^ \t\r\n\f]`
`\d`	Digit character = `[0-9]`
`\D`	Non-digit character = `[^0-9]`
`\b`	Backspace (BS) = `\x08` **Note** Allowed only in bracket expressions.

`\Q`
`<expr>`
`\E`
	Quotes all metacharacters between `\Q` and `\E`. Backslashes are regarded as normal characters. For example, "`\QC:\file.exe\E`" matches the "`C:\file.exe`" string, not the "`C:\x0Cile.exe`" string, where `\x0C` is the formfeed "`\f`".

Example of using special character sequences

```
# This fingerprint matches HTTP content
# for which the length is >= 10000
# The situation context for this regular expression could be either
# "HTTP Request Header Line" or "HTTP Reply Header Line"
Content-Length: \d\d\d\d\d

# The regular expression could be also written as shown below
Content-Length: \d{5}
```

Pattern-matching modifiers

The regular expression syntax has Perl-like extensions. The pattern-matching modifiers are extensions that can be used to control the matching process in more detail.

The modifiers are enabled with (`?<modifiers>`) and disabled with a minus (`?-<modifiers>`), where `<modifiers>` is a list of one or more modifiers.

Example of pattern-matching modifiers

```
# This fingerprint is identical to the special character sequence example,
# except for the (?i) modifier.
# HTTP Header names are case-insensitive. For this reason,
# case-insensitivity is enabled in this fingerprint.
(?i)Content-Length: \d\d\d\d\d
```

The modifiers (`?C`), (`?L`), and (`?s`)are enabled by default.

The pattern-matching modifiers are listed in the following table:

Pattern-matching modifiers

Sequence	Description
`(?i)`	Case insensitive mode When enabled, case insensitive matching is used for the uppercase and lowercase letters. Thus, a letter matches regardless of its capitalization. When disabled, the letters are matched case-sensitively so that capitalization is taken into account in the matching process.
`(?s)`	Single line mode When enabled, the dot character "." matches also the null character `\x00` and a missing character in addition to matching any character (including linefeed and other non-printable characters). When disabled, the linefeed or a missing character are not matched. This modifier is enabled by default. Use `(?-s)` to disable it.
`(?x)`	Extended readability mode When enabled, equals to enabling `(?C)`, `(?L)`, and `(?S)`. Comments, linefeeds and spaces are not used in the matching process, allowing to use them for readability of the expression. When disabled, equals to disabling `(?C)`, `(?L)`, and `(?S)`. Comments, linefeeds, and spaces are used in the matching process.
`(?C)`	Allow comments mode When enabled, anything after the hash character "# " is regarded as a comment and not included in the matching process. When disabled, the hash character "# " and anything following are used in the matching process. This modifier is enabled by default. Use `(?-C)` to disable it.
`(?L)`	Ignore linefeeds mode When enabled, linefeed and carriage return characters are not included in the matching process unless defined (`\x0A` or `\n` for linefeed and `\x0D` or `\r` for carriage return). When disabled, linefeeds and carriage returns are used in the matching process. This modifier is enabled by default. Use `(?-L)` to disable it.
`(?S)`	Ignore spaces mode When enabled, the space and horizontal tab characters are not used in the matching process unless defined (`\x20` for space and `\x09` or `\t` for horizontal tab). When disabled, the space and horizontal tab characters are used in the matching process.
`(?<modifiers>:<expr>)`	Applies the `<modifiers>` modifiers only to the expression `<expr>`. These modifiers are not used in other parts of the regular expression.

Variable expression evaluation

Variable expression evaluation is an extension to regular expression syntax that provides the ability to use variables, parse values from the traffic stream and perform arithmetic operations.

Variable expression syntax

Sequence	Description
`(?[<expression>])`	`<expression>` is one or more comma-separated expressions

Example of setting a variable in a variable expression

```
# This regular expression searches for "aaa" anywhere in the traffic stream,
# and then sets the value of "parameter1" to 1
.*aaa(?[parameter1=1])
```

The default variable size is one bit. Variable size can be changed by appending "`@<size>`" to the variable name. For example, "`parameter1@8`" is an 8-bit variable. Possible variable sizes, in addition to 1, are 8, 16, 32 and 64 bits. By default variables are visible within a situation context. For example, a variable used in a situation with context "HTTP Request URI" is visible to all other situations in that context. Prefixing the variable name with a dollar sign "$" makes it a connection variable. A connection variable is visible in all situations contexts for a single TCP connection, for example in both client and server stream contexts.

By default, no situation match is created when the end of a variable expression in reached. To create a match when a variable expression is used, the "`sid()`" function must be called.

Variable expression syntax

Sequence	Description
`<varexpr_a> -> <varexpr_b>`	`varexpr_b` is executed only if `varexpr_a` is true

The following example shows a typical case where we want to search one string followed by another, for example "aaa" followed by "bbb". An expression such as "`.*aaa.*bbb`" breaks the guideline of not using "`.*`" in the middle of a regular expression. You can circumvent this issue using variable expressions.

Example of setting and checking a variable value in a variable expression

```
# This regular expression searches for "aaa" anywhere in the traffic stream,
# and then sets the value of 'my_var' to 1.
# It also searches for "bbb", and checks whether "aaa" has already been
# seen earlier (i.e. the value of 'my_var' is one). If "aaa" has been seen
# already, a match is created using the "sid()" function.

# The following traffic matches this regular expression: "aaabbb",
# "xxaaaxxxxxbbbxx", "aaaxbbb"
# The following traffic does not match this regular expression:
#"bbbaaa", "aabbbxxaaa"
(?x)
.*aaa(?[my_var=1]) |
.*bbb(?[my_var==1 -> sid()])
```

Example of setting and checking a variable

```
# This regular expression matches when "login.php" is seen in the traffic
# stream before "user=admin" Situation Context, e.g. "HTTP Request URI"
(?x)
.*login\.php(?[login_page_seen=1]) |
.*user=admin(?[login_page_seen==1 -> sid()])
```

All the arithmetic operations that are available in SMC regular expressions are listed in the table below. Operator precedence is the same as in the C programming language, except that '->' is the lowest in precedence. Statements inside parentheses '()' are always evaluated first, so the order of operations can be overridden with parentheses.

Operations on expression results

Sequence	Description
`false`	Always evaluates to a false.
`true`	Always evaluates to a true.
`<number>`	A literal number in decimal, octal, and hexadecimal format, for example "32" or "0x20".
`<var> = <expr>`	Sets a value returned by expression `<expr>` to a variable `<var>`. See variable syntax below.
`<var> += <expr>`	Adds the value of variable `<var>` with the value returned by expression `<expr>` and sets the result to variable `<var>`.
`<var> -= <expr>`	Subtracts the value from variable `<var>` by the value returned by expression `<expr>` and sets the result to variable `<var>`.
`<var> *= <expr>`	Multiplies the value of `<var>` by the value returned by expression `<expr>` and sets the result to variable `<var>`.
`<var> /= <expr>`	Divides the value of `<var>` with the value returned by expression `<expr>` and sets the result to variable `<var>`.
`<var> %= <expr>`	Divides the value of `<var>` with the value returned by expression `<expr>` and sets the modulo of the result to variable `<var>`.
`<var> <<= <expr>`	Shifts the value of `<var>` to left by number of steps returned by expression `<expr>` and sets the result to variable `<var>`.
`<var> >>= <expr>`	Shifts the value of `<var>` to right by number of steps returned by expression `<expr>` and sets the result to variable `<var>`.
`<var> &= <expr>`	Performs bitwise AND with the value of variable `<var>` and the value returned by expression `<expr>` and sets the result to variable `<var>`.
`<var> |= <expr>`	Performs bitwise OR with the value of variable `<var>` and the value returned by expression `<expr>` and sets the result to variable `<var>`.
`<var> ^= <expr>`	Performs bitwise XOR with the value of variable `<var>` and the value returned by expression `<expr>` and sets the result to variable `<var>`.
`<expr_a> -> <expr_b>`	Expression `<expr b>` is evaluated only if `<expr a>` is true.
`<expr_a> ? <expr_b> : <expr_c>`	Expression `<expr b>` is evaluated only if `<expr b>` is true and expression `<expr c>` is evaluated if `<expr a>` is false.
`<expr_a> == <expr_b>`	Test if expressions `<expr a>` and `<expr b>` return an equal value.
`<expr_a> != <expr_b>`	Test if expressions `<expr a>` and `<expr b>` do not return an equal value.

Sequence	Description
`<expr_a> < <expr_b>`	Test if expression `<expr b>` returns higher value than expression `<expr a>`.
`<expr_a> <= <expr_b>`	Test if expression `<expr b>` returns higher or equal value than expression `<expr a>`.
`<expr_a> > <expr_b>`	Test if expression `<expr a>` returns higher value than expression `<expr b>`.
`<expr_a> >= <expr_b>`	Test if expression `<expr a>` returns higher or equal value than expression `<expr b>`.
`<expr_a> & <expr_b>`	Performs bitwise AND with expressions <expr_a> and `<expr b>` and returns the result.
`<expr_a> | <expr_b>`	Performs bitwise OR with expressions `<expr a>` and `<expr b>` and returns the result.
`<expr_a> ^ <expr_b>`	Performs bitwise XOR with expressions `<expr a>` and `<expr b>` and returns the result.
`<expr_a> && <expr_b>`	Performs AND with expressions `<expr a>` and `<expr b>` and returns the result.
`<expr_a> || <expr_b>`	Performs OR with if expressions `<expr a>` and `<expr b>` and returns the result.
`<var>++, ++<var>`	Increase value of variable `<var>` by one.
`<var>--, --<var>`	Decrease value of variable `<var>` by one.
`-<expr>`	Negate the result of the expression `<expr>`.
`~<expr>`	Bitwise invert the result of the expression `<expr>`.
`!<expr>`	Perform NOT operation with the expression `<expr>`.

Note

In a regular expression such as "`.*aaa(?[var1=1])`", the starting of the variable expression "`(?[var1=1])`" is the most time-consuming operation, whereas setting or checking a variable value is a relatively fast operation. For example, the regular expression "`.*/(?[parameter1=1])`" in an HTTP context would cause the starting of a variable expression after every "/" character in the traffic stream. As this character is common in HTTP protocol, the regular expression might degrade the system performance.

Stream operations

Stream operations can be used to read data from the traffic stream.

The value returned by stream operations can either be written to a variable or used directly in an arithmetic operation. The stream operations are listed in the tables below.

ASCII data variable expressions

Sequence	Description
`parse_dec(<length>)`	Parse ASCII decimal value. `<length>` is the maximum number of the characters to parse. The actual number of parsed digits is available in the variable `$parse_length@32`. If no characters could be parsed, then the variable is set to zero.
`parse_hex(<length>)`	Parse ASCII hexadecimal value. `<length>` is the maximum number of the characters to parse. The actual number of parsed digits is available in the variable `$parse_length@32`. If no characters could be parsed, then the variable is set to zero.

Sequence	Description
`parse_int(<length>)`	Parse ASCII value; parses hexadecimal if the string starts with "`0x`", octal if the string starts with zero ("`0`") and decimal otherwise. `<length>` is the maximum number of the characters to parse. The actual number of parsed digits is available in the variable `$parse_length@32`. If no characters could be parsed, then the variable is set to zero.
`parse_oct(<length>)`	Parse ASCII octal value. `<length>` is the maximum number of the characters to parse. The actual number of parsed digits is available in the variable `$parse_length@32`. If no characters could be parsed, then the variable is set to zero.

Miscellaneous input stream operations

Sequence	Description
`CRC(<length>)`	Calculates a 32-bit CRC value starting from the current byte up to number of bytes specified by the `<length>` parameter. This function can be used as a space optimizer for probabilistically matching against a specific large binary block by its CRC. The CRC used is the 32-bit CRC with polynomial 0x104C11DB7 (used for example in Ethernet).
`skip(<length>)`	Skip `<length>` number of bytes.
`regex(<regexp>)`	Launch an independent subexpression.

The binary data from the input stream can be read into variables with the following expressions.

Binary data variable expressions

Sequence	Description
`parse_be@<size>`	Parse big endian value. `<size>` is the size of the value to be read in bits, and it can be one of the following: 8, 16, 24, 32, 40, 48, 56 or 64.
`parse_le@<size>`	Parse little endian value. `<size>` is the size of the value to be read in bits, and it can be one of the following: 8, 16, 24, 32, 40, 48, 56 or 64.

Example of parsing a value from the traffic stream

```
# This regular expression finds the string "&parameter1=", parses the
# following three bytes as an ASCII decimal number, and writes the values
# to the "var1@8" variable
# The regular expression matches only if the number is greater than 100
(?x)
.*&parameter1=(?[var1@8=parse_dec(3), var1@8>100 -> sid()])
```

System variables

System variables are connection variables whose values are set by the Security Engine.

A regular expression can only read the value of these variables. The two most commonly used variables are `$dport` and `$offset`. The `$dport` variable contains the destination port of the connection/datagram, and it is useful especially in:

- “Any Application Protocols” contexts, which receive all traffic (any TCP/UDP port).
- Unknown Application Protocols” contexts, which receive traffic that does not have a dedicated, protocol-specific context (mostly high TCP/UDP ports).

The `$offset` variable contains the number of bytes that have been matched since the beginning of the traffic stream. The following table lists all system variables.

System variables

Sequence	Description
`$major`	The major version number of the Security Engine.
`$minor`	The minor version number of the Security Engine.
`$patch`	The patch level number of the Security Engine.
`$build`	The build number of the Security Engine.
`$dir`	32-bit integer type expression that is evaluated to the current direction of the current fingerprinted data. The expression returns the following values: ■ 0 — Client direction ■ 1 — Server direction **Note** The value is relative to direction from which the connection was established. If a connection was detected while it was in progress, the value might not correspond to the actual client or server direction.
`$dport`	The current destination port of the connection. For TCP, `$dport` is the destination port of the SYN packet. For UDP, `$dport` is the destination port of the first UDP packet sent between two hosts.
`$icmp_code`	A 32-bit integer type expression that is evaluated to the current ICMP code number of the connection.
`$icmp_type`	A 32-bit integer type expression that is evaluated to the current ICMP type number of the connection.
`$ipproto`	A a 32-bit integer type expression that is evaluated to the current IP-protocol number of the connection.
`$offset`	The byte that is under inspection when counted from the beginning of the traffic stream. For implementation-specific reasons, the value is increased only after the first byte of a traffic stream (after the first byte, the value is still `0`). For this reason, the value of `$offset` is actually the real offset minus one.
`$parse_length@32`	Number of digits parsed by last `parse_dec()`, `parse_hex()`, `parse_oct()`, or `parse_in()` expression.
`$sport`	A 32-bit integer type expression that is evaluated to the current source port of the connection.

Example of system variable use

```
# This regular expression matches
# if hexadecimal bytes "0x01", "0x02",
# and "0x03" are seen in port 5000
.*\x01\x02\x03(?[$dport==5000 -> sid()])
```

Related reference
Stream operations on page 1484

Independent subexpressions

Independent subexpressions allow starting another regular expression from inside a variable expression.

The function used for starting the subexpression is "`regex()`". The "`cancel`" function must always be called after a match in a subexpression. This function stops the execution of the subexpression and frees resources. The "`cancel`" function is always called without parentheses "()" unlike other functions.

Subexpressions are useful for splitting a single complex regular expression into two. For example, "`.*&filename=[^&]{256}`" breaks the guideline of not using "`.*`" or "`<expr>{n,m}`" with a large m in the middle of a regular expression. The following illustration shows how to circumvent this limitation by using an independent subexpression.

Example of independent subexpression use

```
# This fingerprint detects an HTTP parameter file name with value longer than # 256 bytes
(?x)
.*&filename=(?[
    regex(
        [^&]{256}(?[sid(),cancel])
    )
])
```

Parallel matching groups

With complex regular expressions, you might need to set up different regular expressions for matching in parallel groups.

You can set different regular expressions to be matched in parallel groups within one Situation Context. Normally, manual Situation group definitions are not needed and the engine automatically compiles all your custom Situations in the same group (group 0). Manual group definitions are needed if the policy upload fails due to fingerprint/DFA compilation problems that can occur with complex regular expressions.

To use grouping, add a new preprocessing tag to the beginning of the regular expression.

Preprocessing tag for setting a group for matching

Syntax	Description

`#!!GROUP(X)`
`Comment`
`#!!#`
	'X' is the group number from 0 to 7. The comment is optional. If you do not specify the group with this tag, the Situation is processed in group zero.

Tips for working with regular expressions

Before you work with regular expressions, review the following tips.

- For more examples of regular expressions, you can view the **Context** tab of the **Situation Properties** dialog box.
- When adding a Situation to an Inspection rule, it is often useful to select the “Excerpt” logging option. This option includes an excerpt of the traffic that the regular expression matches and also the matching position (“Excerpt position”) in the log entry. This helps in verifying that the regular expression works as expected.
- Freely available tools, such as wget, can be used for generating traffic for testing regular expressions.
- If a policy upload fails with an error message such as “Fingerprint compilation failed”, it indicates that a regular expression is too complex. In this case, the regular expression must be edited. For example, use a variable expression or an independent subexpression. If it is not possible to edit the regular expression, the regular expression can be moved to a parallel matching group.

Related reference

Parallel matching groups on page 1487

Appendix G

Schema updates for external LDAP servers

Contents

- Schema updates for external LDAP servers on page 1489

There are SMC-specific LDAP classes and attributes that you add to the schema of external LDAP servers.

Schema updates for external LDAP servers

When adding SMC-specific LDAP classes and attributes to the schema of external LDAP servers, see the following tables.

The SMC-specific attribute and class names start with “sg”. The classes are listed in the following table.

SMC-specific LDAP classes

Class	Description
sggroup	SMC user group
sguser	SMC user account

The SMC-specific attributes are listed in the following table.

SMC-specific LDAP attributes

Attribute	Related classes	Description
sgactivation	sguser	Activation date for the user account.
sgauth	sggroup, sguser	Authentication service for the user or group.
sgdelay	sggroup, sguser	Number of days the user account is valid after the activation.
sgexpiration	sguser	Last day when the user account is valid and the user can log in.
sggrouptype	sggroup	Indicates the type of the group: a subtree or discrete group.
sgmember	sggroup	The Distinguished Name (DN) for the user member of this group.
sgpassword	sguser	MD5 message digest hash of the user password.
sgpresharedkey	sguser	IPsec PreSharedKey for the user account.
sgsubjectaltnames	sguser	IPsec certificate SubjectAltNames for the user account.

Attribute	Related classes	Description
sgvirtualip	sggroup, sguser	Virtual IP allocation allowed for the user.

In addition to updating the directory schema, there can be some server-specific requirements. For Netscape and OpenLDAP version 1.2.11 servers, you must configure the following lines to the LDAP server's slapd.conf configuration file after stopping the LDAP service.

Additional configuration for OpenLDAP v1.2.11 and Netscape server

```
include /etc/openldap/slapd.at.conf
include /etc/openldap/slapd.oc.conf
include /etc/openldap/sg-schema.conf
schemacheck on
```

For OpenLDAP server versions 2.0 and later, you must configure the following lines to the LDAP server's slapd.conf configuration file after stopping the LDAP service.

Additional configuration for OpenLDAP version 2.0 or later

```
include /etc/openldap/schema/core.schema
include /etc/openldap/schema/cosine.schema
include /etc/openldap/schema/inetorgperson.schema
include /etc/openldap/schema/sg-v3.schema
```

Appendix H

Log fields

For descriptions of all log fields, see Knowledge Base article 38581.

Appendix I

Keyboard shortcuts

For a list of available shortcut keys in the SMC Client, see Knowledge Base article 38538.

Appendix J

Multicasting

Contents

- Multicasting vs. unicasting or broadcasting on page 1495
- Overview of IP Multicasting on page 1496
- Internet Group Management Protocol and how it works on page 1497
- Ethernet multicasting on page 1497
- Multicasting and Forcepoint Network Security Platforms on page 1498

The multicasting reference describes the general principles of multicasting and how it can be used with CVIs (cluster virtual IP addresses) in Engine Clusters.



Note

Use the Packet Dispatch CVI mode instead of multicast CVIs as it uses unicast and requires no additional switch or router configuration. The other CVI modes are provided mainly for backward compatibility.

Multicasting vs. unicasting or broadcasting

Multicasting differs in certain important respects from unicasting and broadcasting as a transmission technique.

A distinction can be made between multicasting traffic at the network layer (based on special class D IP addresses) and at the data link layer (based on multicast MAC addresses). The general differences how multicasting can be distinguished from unicasting and broadcasting are highlighted in the following sections.

Multicasting vs. unicasting

In unicasting, the transmitted datagrams are intended only for a single host having a unique address. In multicasting, the data is transmitted likewise to a single address (that is, the multicast group address), but the actual data reaches all hosts that belong to the group identified by the multicast address. This way the data needs only to be sent once, and not separately to each host. This naturally saves bandwidth.

Multicasting vs. broadcasting

In broadcasting, the data is sent from a host to other hosts within a given network, so they must all use their resources to process the data. In contrast, in multicasting, the hosts that do not belong to a multicast group do not have to use their resources for multicast data. Moreover, multicasting is not restricted to a single network. Hosts on remote networks can receive IP multicast datagrams if they belong to a specific host group, and that there are multicast routers forwarding the traffic. Thus, IP multicasting can in principle be used globally whereas broadcasting is limited to a single network.

Overview of IP Multicasting

In the RFC 1112, IP multicasting is defined as the transmission of an IP datagram to a group of hosts identified by a single IP destination address.

In addition to this common multicast group address, the hosts in the group all have separate and unique unicast addresses. The actual multicast host group can consist of any number of hosts, possibly even located in different networks. The number can vary over time, as hosts can join in and leave from a group at any time. Moreover, a particular host can belong to several groups simultaneously.

The multicast group addresses are class D addresses. They are identified by the high-order initial four-bit sequence *1110*. In the dotted decimal notation, the multicast group address range runs from 224.0.0.0 to 239.255.255.255. There are certain special addresses:

- 224.0.0.0 is never assigned.
- 224.0.0.1 is assigned to the permanent group of all hosts, including gateways, in the local network.
- 224.0.0.2 is assigned to all local multicast routers.

Multicast IP addresses are not allowed to be used as source addresses. A multicast source address implies forging of an IP address.

The multicast groups are either permanent or transient. Permanent groups have administratively assigned IP addresses, while the addresses of the transient multicast groups can be assigned dynamically from the pool of multicast addresses not reserved for permanent groups. The IP address of an established permanent group persists even if the group would not have any members at a given time. The transient groups cease to exist as soon as they no longer have member hosts, and the assigned multicast address is released.

See, for example, https://www.iana.org/assignments/multicast-addresses for a list of addresses registered with IANA.

How Multicasting can be used with applications

Multicasting is a viable option for many types of transmissions.

Multicasting is widely used in local area networks for various purposes. Moreover, multicasting can be used both for receiving a publicly transmitted session on an intranet, or for transmitting an internal communication to a public network (for example, for announcing a product launch). Multicasting is particularly important solution for bandwidth-intensive applications, such as multimedia. The most typical protocol for multicast traffic is UDP.

Multicasting can be a suitable solution, for example, for the following applications:

- Work groups, electronic whiteboards.
- Video/voice-over-IP conferences.
- Real-time streaming media (for example, Internet radio).
- File transfer.
- Spreading of any information to certain selected destinations.

Internet Group Management Protocol and how it works

Internet Group Management Protocol (IGMP) is an integral part of Internet Protocol.

The IGMP messages are encapsulated in IP datagrams. IGMP is used both between hosts and multicast routers, and between multicast routers. It keeps multicast routers informed of the multicast group memberships on a given local network. Each host supporting multicasting must join the multicast group with the address 224.0.0.1 on each network interface at the initialization time. They shall remain members of this group as long as they are active. With IGMP, the hosts on a LAN can inform the routers that they want to be able to receive multicast messages from external networks.

How membership messages are used with IGMP

Multicast routers use IGMP for enquiring periodically which multicast groups have members in the connected local networks.

This inquiry is carried out by sending *Host Membership Query* messages to the all-hosts address 224.0.0.1. The hosts receiving the query respond by sending *Host Membership Reports* to all neighboring multicast routers.

A host joining a new group immediately transmits a report, instead of waiting for a query. When a host wants to stop receiving a multicast transmission, it sends a *Leave Report* message with the destination address 224.0.0.2 to all subnet routers. A router receiving a Leave Report message sends in response a *Group Specific Query* to the multicast address to check whether there still are hosts in that group. In case no response is received, multicasting to that address is stopped.

Ethernet multicasting

When multicasting is implemented at the data link layer, stations are identified by their Media Access Control (MAC) addresses as well as their network level IP addresses.

So far we have seen how multicasting is implemented at the network layer and how multicast IP addresses differ from other types of IP addresses. In addition, we must also distinguish multicasting at the data link layer where stations are identified, not only by their network level IP addresses, but also by their MAC addresses. As opposed to unicast and broadcast addresses, the relation of multicast addressing to IP addressing applies also at this level.

Most local area network (LAN) topologies allow for multicasting by using a group addressing scheme. Some topologies offer better support for multicasting than others. In Ethernet (as defined in IEEE 802.3), all MAC addresses that have the least significant bit of the most significant byte as “1” are multicast addresses. Thus, for example, 01:00:00:00:00:00 and 49:aa:bb:cc:dd:ee are both multicast MAC addresses; while 02:00:00:00:00:00 and fe:fe:fe:fe:fe:fe are not. The devices with a given multicast MAC defined are able to listen to all traffic sent to that particular MAC address.

A specific subset of MAC addresses is reserved for mapping the IP multicasting addresses to data link layer addresses. In Ethernet, the multicast MAC addresses that correspond to multicast IP addresses range from 01:00:5e:00:00:00 to 01:00:5e:7f:ff:ff.

Multicasting and Forcepoint Network Security Platforms

After distinguishing between network layer multicasting and data link layer multicasting, we can now have a look at how the engine uses multicasting and unicasting.



Note

Use the Packet Dispatch CVI mode instead of multicast CVIs as it uses unicast and requires no additional switch or router configuration. The other CVI modes are provided mainly for backward compatibility.

When using clustering technology, the clustered engine nodes share a common *unicast* IP address, which is called a *CVI* (cluster virtual IP address). This shared IP address is assigned to the node that receives traffic that arrives from the network for distributing and load-balancing between all nodes. Any traffic that has a specific node in the cluster as its final destination (such as management connections) is sent to *NDIs* (node dedicated IP addresses).

CVIs allow the cluster to appear as a single virtual entity to other network devices, rather than a group of individual nodes. Traffic addressed to CVIs is load-balanced between the nodes according to the cluster's load-balancing filters. The load-balancing filters determine which traffic is distributed to which individual nodes. This way, a specific node in a cluster handles all packets in the connection as long as the node stays online.

In addition to the shared unicast IP address, each node must also share a data link layer address (MAC) at the CVI. Only this way will each of the nodes be provided with the exact same traffic. There are different options for the cluster-wide MAC address, and the selection depends on the features of the other connected networking devices, such as switches and hubs. This document is not a definitive reference for different types of switch configurations, but it gives an overview of possible considerations when implementing engine clusters in different types of network environments.

The method can be selected based on the surrounding network devices. Unicast MAC configuration can be used with hubs and with switches that support sending a specified unicast MAC address to several ports at the same time. When a layer 2 network is not able to do this, multicast MAC can be used instead. Because multicast MAC sends all packets to all ports, unicast MAC mode gives better performance with hubs. However, in large networks with large amounts of traffic, the action of sending packets to all ports can create extra load. In that situation, static MAC address forwarding tables can be used to limit traffic to Cluster multicast MAC to cluster ports only. With switches that do not support static MAC address forwarding tables, IGMP snooping can be used for the same task. With switches, Packet Dispatch mode creates less load to switches than unicast MAC or multicast MAC modes.

The different configuration options are presented in the following sections.

Unicast MAC

A common unicast MAC can be defined at the CVIs if the cluster is connected to hubs or switches that can forward frames with a unicast destination to multiple ports.

This way the network devices forward the same packets to each of the connected engine nodes sharing this combination of unicast IP and MAC addresses. This mode is recommended whenever the networking devices support sending packets to a specified unicast MAC address to a predefined set of ports at the same time (as opposed to one port, which is typically the default). Hubs by default support this; however, with switches this is not as frequent, and they usually need additional configuration. With unicast MAC, only the switches directly connected to the cluster need special configuration.

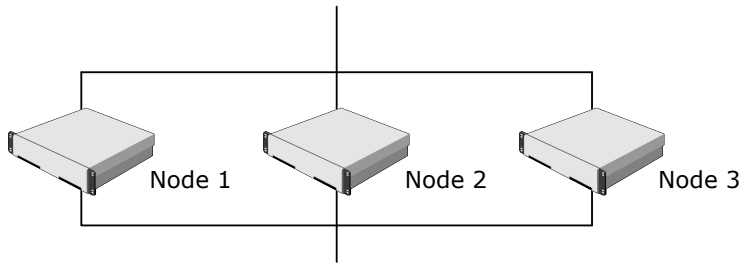
Note

Unlike multicast MAC addresses, there can be only one unicast MAC address defined per Interface ID. Thus, all NDIs and the unicast CVIs on the same physical interface use the same MAC address.

In addition to the common CVI IP address, each node can optionally have unique unicast IP addresses defined at the same physical interface as the CVI. These unicast IP addresses are assigned to NDIs (node dedicated IP addresses), and used when an individual node is the endpoint of a connection. Because there can only be one unicast MAC address at a given interface, also the node-specific NDI IP addresses are mapped to the common unicast MAC.

The following illustration exemplifies the IP and MAC address configuration of a cluster's interfaces that are connected to an external network. By default, the CVI of each node share one unicast IP address. The CVI is mapped to a common unicast MAC address. In addition, for each node, an NDI is defined at the same physical interface as the CVI. The NDI IP addresses are unique, but they all are mapped to the same unicast MAC as the CVI IP address, as there can be only one unicast MAC defined for a physical interface. Traffic directed from the Internet to the cluster's external CVI IP address is sent by the connected switch or hub to all nodes because they all are identified by the same unicast MAC.

CVI with unicast MAC



Interface (external)	Node 1	Node 2	Node 3
CVI IP Address	203.0.113.254	203.0.113.254	203.0.113.254
CVI Unicast MAC	08:08:08:08:08:08	08:08:08:08:08:08	08:08:08:08:08:08
NDI IP Address	203.0.113.21	203.0.113.22	203.0.113.23
NDI Unicast MAC	08:08:08:08:08:08	08:08:08:08:08:08	08:08:08:08:08:08

Multicast MAC

In case it is not feasible to use a switch that works in unicast mode with clusters, a shared multicast MAC can be defined for the cluster nodes.

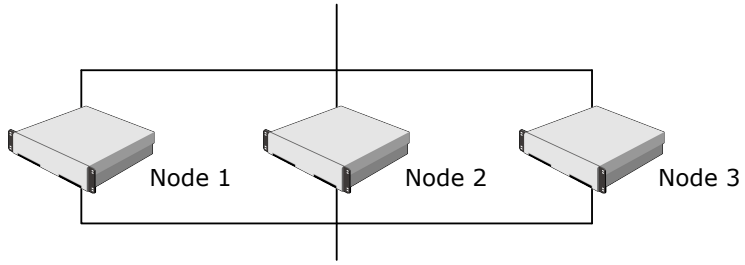
Most switches support this mode, however, not all switches in the same virtual LAN (VLAN) need to be configured. By default, most switches send packets with a multicast MAC address to all ports connected to the same VLAN. If the size of the VLAN is small, this type of flooding is acceptable. However, with larger VLANs performance problems can occur as the device needs to send each packet to each port connected to the same VLAN. In some switches, it's possible to prevent this type of flooding by statically restricting multicast traffic with a given MAC address to some predefined ports only.

Note

Some networking devices discard ARP replies specifying a multicast MAC. In this case, static ARP entries must be used.

The following illustration presents an example where a common multicast MAC is configured for all cluster nodes. For instance, if a switch is not able to send packets with the same unicast MAC to multiple ports, this type of configuration might be used. Each node has also a unique unicast MAC address mapped to the corresponding IP addresses defined at the NDIs.

CVI with multicast MAC



Interface (external)	Node 1	Node 2	Node 3
CVI IP Address	203.0.113.254	203.0.113.254	203.0.113.254
CVI Unicast MAC	09:08:08:08:08:08	09:08:08:08:08:08	09:08:08:08:08:08
NDI IP Address	203.0.113.21	203.0.113.22	203.0.113.23
NDI Unicast MAC	04:08:08:08:08:08	06:08:08:08:08:08	08:08:08:08:08:08

Multicast MAC with IGMP

Internet Group Management Protocol (IGMP) can be used in combination with multicast MAC addresses to avoid flooding with switches that do not support statically defined destinations for multicast.

In this mode, switches are configured to send multicast traffic only to the ports from which they have received IGMP *Host Membership Report* messages corresponding to the MAC address in question. Multicast with IGMP must be selected as the mode for the cluster, and IGMP snooping enabled on the switch. For the IGMP messaging, a common multicast *IP address* for the cluster nodes should be specified. The multicast *MAC address* is then computed automatically based on it. Do note, however, that the CVIs are still identified solely by the common *unicast IP address*; the multicast IP address is only used as the source address for the IGMP messages sent to the switch.

Note

Some routers that use router redundancy protocols such as HSRP or VRRP listen to all multicast traffic in addition to the routing-related traffic. Thus, multicast packets are rerouted to the network. To prevent the rerouting, you can either configure the router to send this traffic only to the cluster ports or define the router's access control list (ACL) to drop all incoming packets with the cluster's multicast MAC.

© 2025 Forcepoint