

FortiOS - Administration Guide

Version 6.4.0

Netskope Exhibit 1017

FORTINET DOCUMENT LIBRARY
https://docs.fortinet.com

FORTINET VIDEO GUIDE
https://video.fortinet.com

FORTINET BLOG
https://blog.fortinet.com

CUSTOMER SERVICE & SUPPORT
https://support.fortinet.com

FORTINET TRAINING & CERTIFICATION PROGRAM
https://www.fortinet.com/training-certification

NSE INSTITUTE
https://training.fortinet.com

FORTIGUARD CENTER
https://www.fortiguard.com

END USER LICENSE AGREEMENT
https://www.fortinet.com/doc/legal/EULA.pdf

FEEDBACK
Email: techdoc@fortinet.com



June 3, 2021
FortiOS 6.4.0 Administration Guide
01-640-607590-20210603

Netskope Exhibit 1017

TABLE OF CONTENTS

Change Log **16**
Getting started **19**
Differences between models 19
Using the GUI 19
Connecting using a web browser 19
Menus 20
Tables 21
Entering values 23
Using the CLI 24
Connecting to the CLI 24
CLI basics 27
Command syntax 33
Subcommands 36
Permissions 38
FortiExplorer for iOS 38
Getting started with FortiExplorer 39
Connecting FortiExplorer to a FortiGate via WiFi 42
Running a security rating 43
Upgrading to FortiExplorer Pro 44
Basic administration 44
Registration 45
FortiCare and FortiGate Cloud login 46
Troubleshooting your installation 49
Zero touch provisioning 51
Zero touch provisioning with FortiDeploy 51
Zero touch provisioning with FortiManager 53
Dashboards and widgets **56**
Using dashboards 56
Viewing device dashboards in the security fabric 58
Creating a fabric system and license dashboard 59
Using widgets 60
Monitor dashboards and widgets 62
Static & Dynamic Routing Monitor 63
DHCP monitor 64
IPSEC monitor 66
SSL VPN monitor 67
Firewall Users Monitor 68
Device inventory 69
FortiView 73
FortiView dashboards and widgets 74
Adding top FortiView widgets by category 75
VDOMs and dashboards 78
FortiView interface 79
FortiView from disk 86

Netskope Exhibit 1017

FortiView from FortiAnalyzer ... 88
FortiView from FortiGate Cloud ... 89
FortiView top sources ... 91
Viewing top websites and sources by category ... 92
Cloud application view ... 96
Configuration backups ... 108
Fortinet Security Fabric ... 112
Security Fabric settings and usage ... 112
Components ... 113
Configuring the root FortiGate and downstream FortiGates ... 115
Configuring FortiAnalyzer ... 120
Configuring other Security Fabric devices ... 120
Using the Security Fabric ... 144
Deploying the Security Fabric ... 158
Synchronizing objects across the Security Fabric ... 166
Security Fabric over IPsec VPN ... 170
Leveraging LLDP to simplify security fabric negotiation ... 176
Configuring the Security Fabric with SAML ... 179
Configuring single-sign-on in the Security Fabric ... 179
CLI commands for SAML SSO ... 184
SAML SSO with pre-authorized FortiGates ... 185
Navigating between Security Fabric members with SSO ... 185
Integrating FortiAnalyzer management using SAML SSO ... 188
Integrating FortiManager management using SAML SSO ... 190
Advanced option - FortiGate SP changes ... 191
Advanced option - unique SAML attribute types ... 192
Security rating ... 195
Security Fabric score ... 198
External connectors ... 199
SDN connectors ... 200
Kubernetes (K8s) SDN connectors ... 245
Endpoint/Identity connectors ... 258
Threat feeds ... 270
Automation stitches ... 284
Creating automation stitches ... 284
Triggers ... 297
Actions ... 301
Execute a CLI script based on CPU and memory thresholds ... 328
Troubleshooting ... 333
Viewing a summary of all connected FortiGates in a Security Fabric ... 333
Diagnosing automation stitches ... 336
Network ... 340
Interfaces ... 340
Interface settings ... 341
Aggregation and redundancy ... 344
VLANs ... 346
Enhanced MAC VLANs ... 352
Inter-VDOM routing ... 355

Netskope Exhibit 1017

Software switch ... 360
Zone ... 361
Virtual Wire Pair ... 363
Virtual switch support for FortiGate 300E series ... 365
Failure detection for aggregate and redundant interfaces ... 367
VLAN inside VXLAN ... 367
Virtual Wire Pair with VXLAN ... 370
DNS ... 372
Important DNS CLI commands ... 372
DNS domain list ... 373
FortiGate DNS server ... 375
DDNS ... 377
DNS latency information ... 380
DNS over TLS ... 382
DNS troubleshooting ... 382
Explicit and transparent proxies ... 383
Explicit web proxy ... 383
FTP proxy ... 386
Transparent proxy ... 388
Proxy policy addresses ... 391
Proxy policy security profiles ... 399
Explicit proxy authentication ... 405
Transparent web proxy forwarding ... 410
Upstream proxy authentication in transparent proxy mode ... 411
Multiple dynamic header count ... 413
Restricted SaaS access (Office 365, G Suite, Dropbox) ... 416
Explicit proxy and FortiSandbox Cloud ... 418
Proxy chaining (web proxy forwarding servers) ... 421
Agentless NTLM authentication for web proxy ... 425
DHCP server ... 428
Configure DHCP on the FortiGate ... 428
DHCP options ... 429
IP address assignment with relay agent information option ... 430
DHCP client options ... 432
Static routes ... 433
Policy routes ... 434
RIP ... 434
OSPF ... 434
BGP ... 434
Direct IP support for LTE/4G ... 435
LLDP reception ... 438
Route leaking between VRFs ... 441
SD-WAN ... 444
SD-WAN quick start ... 444
Configuring the SD-WAN interface ... 445
Adding a static route ... 446
Selecting the implicit SD-WAN algorithm ... 446

Netskope Exhibit 1017

Configuring firewall policies for SD-WAN 447
Link monitoring and failover 448
Results 449
Configuring SD-WAN in the CLI 453
WAN path control 455
Performance SLA - link monitoring 455
Performance SLA - SLA targets 457
Factory default health checks 458
Implicit rule 461
SD-WAN rules - best quality 464
SD-WAN rules - lowest cost (SLA) 467
SD-WAN rules - maximize bandwidth (SLA) 470
Application steering using SD-WAN rules 472
SD-WAN traffic shaping and QoS 484
Advanced configuration 488
Self-originating traffic 489
SDN dynamic connector addresses in SD-WAN rules 489
Forward error correction on VPN overlay networks 492
Using BGP tags with SD-WAN rules 495
BGP multiple path support 498
Controlling traffic with BGP route mapping and service rules 501
Applying BGP route-map to multiple BGP neighbors 507
IBGP and EBGP support in VRF 513
ADVPN and shortcut paths 516
SD-WAN monitor on ADVPN shortcuts 528
SD-WAN integration with OCVPN 529
DSCP matching (shaping) 537
SD-WAN health check packet DSCP marker support 541
Dual VPN tunnel wizard 541
Internet service customization 542
SD-WAN with FGCP HA 543
Hub and spoke SD-WAN deployment example 549
Datacenter configuration 550
Branch configuration 554
Validation 558
Dynamic definition of SD-WAN routes 559
Adding another datacenter 560
Dynamic connector addresses in SD-WAN policies 561
SD-WAN configuration portability 562
Interface speedtest 564
Configuring SD-WAN in an HA cluster using internal hardware switches 566
Troubleshooting SD-WAN 570
Understanding SD-WAN related logs 570
SD-WAN related diagnose commands 573
SLA logging 577
SLA monitoring using the REST API 579
SD-WAN bandwidth monitoring service 581

Netskope Exhibit 1017

System ... **585**
Basic system settings ... 585
Advanced system settings ... 585
Operating modes ... 586
Administrators ... 587
Administrator profiles ... 587
Add a local administrator ... 589
Remote authentication for administrators ... 589
Password policy ... 591
Admin profile option for diagnose access ... 592
Firmware ... 593
Firmware upgrade notifications ... 594
Downloading a firmware image ... 595
Testing a firmware version ... 597
Upgrading the firmware ... 598
Downgrading to a previous firmware version ... 599
Installing firmware from system reboot ... 600
Restoring from a USB drive ... 601
Controlled upgrade ... 602
Settings ... 602
Default administrator password ... 602
Changing the host name ... 603
Setting the system time ... 603
Configuring ports ... 607
Setting the idle timeout time ... 608
Setting the password policy ... 608
Changing the view settings ... 608
Setting the administrator password retries and lockout time ... 609
Virtual Domains ... 610
Split-task VDOM mode ... 611
Multi VDOM mode ... 615
Configure VDOM-A ... 618
Configure VDOM-B ... 620
Configure the VDOM link ... 622
Configure VDOM-A ... 627
Configure VDOM-B ... 629
High Availability ... 631
Introduction to the FGCP cluster ... 631
Failover protection ... 632
FGSP (session synchronization) peer setup ... 634
UTM inspection on asymmetric traffic in FGSP ... 635
UTM inspection on asymmetric traffic on L3 ... 637
Encryption for L3 on asymmetric traffic in FGSP ... 639
Synchronizing sessions between FGCP clusters ... 639
Using standalone configuration synchronization ... 641
Troubleshoot an HA formation ... 643
Check HA sync status ... 643
Disabling stateful SCTP inspection ... 644

Netskope Exhibit 1017

Upgrading FortiGates in an HA cluster 645
HA cluster setup examples 646
Routing data over the HA management interface 655
Override FortiAnalyzer and syslog server settings 657
Force HA failover for testing and demonstrations 661
Querying autoscale clusters for FortiGate VM 664
SNMP 665
Interface access 665
MIB files 666
SNMP agent 666
SNMP v1/v2c communities 667
SNMP v3 users 669
Important SNMP traps 670
SNMP traps and query for monitoring DHCP pool 672
Replacement messages 673
Replacement message images 673
Modifying replacement messages 674
Replacement message groups 676
Example 677
FortiGuard 679
IPv6 FortiGuard connections 680
Configuring antivirus and IPS options 681
Manual updates 681
Automatic updates 682
Sending malware statistics to FortiGuard 684
Update server location 685
Filtering 685
Override FortiGuard servers 686
Online security tools 687
FortiGuard anycast and third-party SSL validation 687
Using FortiManager as a local FortiGuard server 690
Cloud service communication statistics 691
IoT detection service 692
Feature visibility 694
Security feature presets 694
Certificates 695
Microsoft CA deep packet inspection 695
Purchase and import a signed SSL certificate 700
Configuration scripts 703
Workspace mode 704
Policy and Objects 706
Policies 706
Firewall policy parameters 707
Profile-based NGFW vs policy-based NGFW 708
NGFW policy mode application default service 712
Policy views and policy lookup 714
Policy with source NAT 715
Policy with destination NAT 728

Netskope Exhibit 1017

Policy with Internet Service 741
NAT64 policy and DNS64 (DNS proxy) 758
NAT46 policy 762
Multicast processing and basic Multicast policy 765
Local-in policies 766
IPv4/IPv6 access control lists 768
Mirroring SSL traffic in policies 769
Inspection mode per policy 772
OSPFv3 neighbor authentication 774
Firewall anti-replay option per policy 776
Enabling advanced policy options in the GUI 777
Recognize anycast addresses in geo-IP blocking 778
Matching GeoIP by registered and physical location 779
Authentication policy extensions 780
HTTP to HTTPS redirect for load balancing 781
GTPv2 in policies 783
Use active directory objects directly in policies 784
FortiGate Cloud / FDN communication through an explicit proxy 788
No session timeout 789
Objects 790
Address group exclusions 791
MAC addressed-based policies 792
ISDB well-known MAC address list 795
Dynamic policy — fabric devices 796
FSSO dynamic address subtype 798
ClearPass integration for dynamic address objects 802
Group address objects synchronized from FortiManager 806
Using wildcard FQDN addresses in firewall policies 808
IPv6 geography-based addresses 810
Array structure for address objects 812
Traffic shaping 814
Determining your QoS requirements 815
Packet rates 816
Changing traffic shaper bandwidth unit of measurement 818
Shared traffic shaper 818
Per-IP traffic shaper 822
Type of Service-based prioritization and policy-based traffic shaping 825
Interface-based traffic shaping profile 828
Interface-based traffic shaping with NP acceleration 837
Classifying traffic by source interface 838
Configuring traffic class IDs 839
Traffic shaping schedules 842
QoS assignment and rate limiting for quarantined VLANs 843
Weighted random early detection queuing 844
Security Profiles 851
Inspection modes 851
Flow mode inspection (default mode) 852
Proxy mode inspection 852
Inspection mode feature comparison 854

Netskope Exhibit 1017

Antivirus 856
Protocol comparison between antivirus inspection modes 857
Other antivirus differences between inspection modes 857
Databases 858
Content disarm and reconstruction for antivirus 859
FortiGuard outbreak prevention 862
External malware block list for antivirus 864
Checking flow antivirus statistics 867
CIFS support 870
Using FortiSandbox appliance with antivirus 876
Using FortiSandbox Cloud with antivirus 886
Web filter 894
URL filter 895
FortiGuard filter 901
Credential phishing prevention 908
Usage quota 910
Web content filter 913
File filter 916
Advanced filters 1 922
Advanced filters 2 926
External resources for web filter 931
Reliable web filter statistics 937
Flow-based web filtering 940
URL certificate blocklist 941
DNS filter 942
How to configure and apply a DNS filter profile 943
FortiGuard category-based DNS domain filtering 945
Botnet C&C domain blocking 948
DNS safe search 952
Local domain filter 953
DNS translation 956
Using a FortiGate as a DNS server 959
Troubleshooting for DNS filter 961
Application control 963
Basic category filters and overrides 964
Port enforcement check 967
Protocol enforcement 968
SSL-based application detection over decrypted traffic in a sandwich topology 970
Matching multiple parameters on application control signatures 971
Intrusion prevention 974
Botnet C&C IP blocking 974
Detecting IEC 61850 MMS protocol in IPS 978
Email filter 979
Protocol comparison between Email filter inspection modes 980
Local-based filters 980
FortiGuard-based filters 985
File type-based filters 986
Protocols and actions 991
Configuring webmail filtering 992

Netskope Exhibit 1017

Data leak prevention ... 993
Protocol comparison between DLP inspection modes ... 994
Basic DLP filter types ... 994
DLP fingerprinting ... 997
DLP watermarking ... 1001
VoIP solutions ... 1004
General use cases ... 1004
SIP message inspection and filtering ... 1008
SIP pinholes ... 1010
SIP over TLS ... 1011
Custom SIP RTP port range support ... 1012
Voice VLAN auto-assignment ... 1014
ICAP ... 1015
ICAP configuration example ... 1016
ICAP response filtering ... 1018
Web application firewall ... 1021
Protecting a server running web applications ... 1021
SSL & SSH Inspection ... 1024
Certificate inspection ... 1025
Deep inspection ... 1026
Protecting an SSL server ... 1029
Ignoring the AUTH TLS command ... 1030
SSH traffic file scanning ... 1031
Redirect to WAD after handshake completion ... 1034
Custom signatures ... 1035
Application groups in policies ... 1035
Overrides ... 1038
Web rating override ... 1038
Web profile override ... 1041
VPN ... 1046
IPsec VPNs ... 1046
General IPsec VPN configuration ... 1046
Site-to-site VPN ... 1071
Remote access ... 1116
Aggregate and redundant VPN ... 1150
Overlay Controller VPN (OCVPN) ... 1185
ADVPN ... 1216
Other VPN topics ... 1249
VPN IPsec troubleshooting ... 1275
SSL VPN ... 1282
SSL VPN best practices ... 1283
SSL VPN quick start ... 1285
SSL VPN tunnel mode ... 1292
SSL VPN web mode for remote user ... 1299
SSL VPN authentication ... 1302
SSL VPN to IPsec VPN ... 1377
SSL VPN protocols ... 1387
SSL VPN troubleshooting ... 1388

Netskope Exhibit 1017

User & Authentication **1392**
Endpoint control and compliance 1392
Per-policy disclaimer messages 1392
Compliance 1395
FortiGuard distribution of updated Apple certificates 1401
User Definition 1402
User types 1402
Removing a user 1402
User Groups 1403
Configuring POP3 authentication 1403
Guest Management 1404
Configuring guest access 1404
Retail environment guest access 1406
LDAP Servers 1407
Configuring an LDAP server 1407
FSSO polling connector agent installation 1409
Enabling Active Directory recursive search 1414
Configuring LDAP dial-in using a member attribute 1415
Configuring wildcard admin accounts 1416
Exchange Server connector with Kerberos KDC auto-discovery 1418
Configuring least privileges for LDAP admin account authentication in Active Directory 1419
RADIUS Servers 1419
Configuring RADIUS SSO authentication 1420
RSA ACE (SecurID) servers 1426
Support for Okta RADIUS attributes filter-Id and class 1430
TACACS+ servers 1432
SAML 1433
SAML SP for VPN authentication 1434
Authentication Settings 1436
FortiTokens 1438
Configuring FortiTokens 1440
FortiToken Cloud 1443
Configuring FortiToken Mobile 1443
FortiToken maintenance 1444
FortiToken Mobile Push 1445
Configuring the maximum log in attempts and lockout period 1445
PKI 1446
Creating a PKI/peer user 1446
Configuring firewall authentication 1447
Creating a locally authenticated user account 1447
Creating a RADIUS-authenticated user account 1448
Creating an FSSO user group 1449
Creating a firewall user group 1451
Defining policy addresses 1451
Creating security policies 1452

Netskope Exhibit 1017

Wireless configuration **1454**
Switch Controller **1455**
FortiLink setup 1455
Viewing managed FortiSwitches 1456
VLAN interface templates for FortiSwitches 1458
Create VLAN interface templates 1459
Example 1460
FortiLink auto network configuration policy 1462
FortiLink network sniffer extension 1463
FortiLink MCLAG configuration 1465
Standalone FortiGate as switch controller 1466
Standalone FortiGate as switch controller 1466
Multiple FortiSwitches managed via hardware/software switch 1470
Multiple FortiSwitches in tiers via aggregate interface with redundant link enabled 1474
Multiple FortiSwitches in tiers via aggregate interface with MCLAG enabled only on distribution 1478
HA (A-P) mode FortiGate pairs as switch controller 1482
Multiple FortiSwitches managed via hardware/software switch 1482
Multiple FortiSwitches in tiers via aggregate interface with redundant link enabled 1486
Multiple FortiSwitches in tiers via aggregate interface with MCLAG enabled only on distribution 1490
Multiple FortiSwitches in tiers via aggregate interface with MCLAG enabled on all tiers 1494
Authentication and security 1499
MAC-based 802.1X authentication 1499
Port-based 802.1X authentication 1502
MAC layer control - Sticky MAC and MAC Learning-limit 1505
Quarantine 1506
Flow and device detection 1511
Data statistic 1511
Security Fabric showing 1512
NAC policies on switch ports 1513
Use FortiSwitch to query FortiGuard IoT service for device details 1517
Voice device detection 1519
FortiSwitch multi-tenant support 1523
Persistent MAC learning 1526
Split port mode (for QSFP / QSFP28) 1527
Dynamic VLAN name assignment from RADIUS attribute 1529
MSTI support 1529
Netflow and IPFIX support 1530
Log and Report **1533**
Viewing event logs 1533
Sample logs by log type 1534
Log buffer on FortiGates with an SSD disk 1554
Checking the email filter log 1557
Supported log types to FortiAnalyzer, syslog, and FortiAnalyzer Cloud 1557

Netskope Exhibit 1017

Configuring multiple FortiAnalyzers on a multi-VDOM FortiGate 1558
Checking FortiAnalyzer connectivity 1559
Configuring multiple FortiAnalyzers (or syslog servers) per VDOM 1561
Source and destination UUID logging 1562
Troubleshooting 1564
Log-related diagnose commands 1564
Backing up log files or dumping log messages 1570
SNMP OID for logs that failed to send 1572
VM 1576
Amazon Web Services 1576
Microsoft Azure 1576
Google Cloud Platform 1576
Oracle OCI 1576
AliCloud 1576
Private cloud 1576
VM license 1576
Uploading a license file 1577
Types of VM licenses 1578
CLI troubleshooting 1579
FortiGate multiple connector support 1580
Adding VDOMs with FortiGate v-series 1583
Terraform: FortiOS as a provider 1585
Troubleshooting 1590
PF SR-IOV driver support 1590
Troubleshooting 1592
Troubleshooting methodologies 1592
Verify user permissions 1593
Establish a baseline 1593
Create a troubleshooting plan 1595
Troubleshooting scenarios 1596
Checking the system date and time 1597
Checking the hardware connections 1598
Checking FortiOS network settings 1599
Troubleshooting CPU and network resources 1602
Troubleshooting high CPU usage 1603
Checking the modem status 1607
Running ping and traceroute 1608
Checking the logs 1612
Verifying routing table contents in NAT mode 1612
Verifying the correct route is being used 1613
Verifying the correct firewall policy is being used 1613
Checking the bridging information in transparent mode 1614
Checking wireless information 1615
Performing a sniffer trace (CLI and packet capture) 1616
Debugging the packet flow 1619
Testing a proxy operation 1622
Displaying detail Hardware NIC information 1622

Netskope Exhibit 1017

Performing a traffic trace ... 1625
Using a session table ... 1626
Finding object dependencies ... 1629
Diagnosing NPU-based interfaces ... 1630
Identifying the XAUI link used for a specific traffic stream ... 1631
Running the TAC report ... 1631
Other commands ... 1632
FortiGuard troubleshooting ... 1633
Additional resources ... 1636
Technical documentation ... 1636
Fortinet video library ... 1636
Release notes ... 1637
Knowledge base ... 1637
Fortinet technical discussion forums ... 1637
Fortinet training services online campus ... 1637
Fortinet Support ... 1637

Netskope Exhibit 1017

Change Log

Date	Change Description
2020-03-31	Initial release.
2020-04-08	Added SD-WAN integration with OCVPN on page 529 and Allow FortiClient to join OCVPN on page 1199.
2020-04-13	Updated Device inventory on page 69.
2020-04-14	Added Synchronizing objects across the Security Fabric on page 166.
2020-04-15	Added Viewing device dashboards in the security fabric on page 58 and Creating a fabric system and license dashboard on page 59. Updated Topology on page 147.
2020-04-16	Added NAC policies on switch ports on page 1513, Use FortiSwitch to query FortiGuard IoT service for device details on page 1517, UTM inspection on asymmetric traffic on L3 on page 637, Admin profile option for diagnose access on page 592, Interface speedtest on page 564, FortiView on page 73, Adding top FortiView widgets by category on page 75 and VDOMs and dashboards on page 78.
2020-04-17	Added Microsoft Teams integration webhook on page 326, SD-WAN monitor on ADVPN shortcuts on page 528, SD-WAN health check packet DSCP marker support on page 541, SD-WAN configuration portability on page 562, Weighted round robin for IPsec aggregate tunnels on page 1183, Matching multiple parameters on application control signatures on page 971, and UDP hole punching for spokes behind NAT on page 1246.
2020-04-21	Added Support for Okta RADIUS attributes filter-Id and class on page 1430, SNMP traps and query for monitoring DHCP pool on page 672, Interface-based traffic shaping with NP acceleration on page 837, Querying autoscale clusters for FortiGate VM on page 664, Detecting IEC 61850 MMS protocol in IPS on page 978, NAS-IP support per SSL-VPN realm on page 1375, DHCP client options on page 432, Defining gateway IP addresses in IPsec with mode-config and DHCP on page 1271, Group address objects synchronized from FortiManager on page 806, and Voice device detection on page 1519.
2020-04-22	Added VLAN interface templates for FortiSwitches on page 1458, Firmware upgrade notifications on page 594, and Identifying the XAUI link used for a specific traffic stream on page 1631.
2020-04-23	Added IoT detection service on page 692.
2020-04-24	Added Route leaking between VRFs on page 441, No session timeout on page 789, and Redirect to WAD after handshake completion on page 1034.

Date	Change Description
2020-04-27	Added SD-WAN quick start on page 444, FortiView from disk on page 86, FortiView from FortiAnalyzer on page 88, Cloud application view on page 96, and FortiView Top Source and Top Destination Firewall Objects widgets on page 106.
2020-04-28	Added FortiView top sources on page 91, Cloud application view on page 96, and Top application: YouTube example on page 102. Updated Synchronizing sessions between FGCP clusters on page 639.
2020-05-01	Added Application steering using SD-WAN rules on page 472.
2020-05-04	Added Viewing top websites and sources by category on page 92. Updated Additional devices on page 143.
2020-05-07	Added Applying BGP route-map to multiple BGP neighbors on page 507 and SD-WAN with FGCP HA on page 543.
2020-05-08	Updated Private Cloud K8s SDN connector on page 245.
2020-05-13	Updated Credential phishing prevention on page 908.
2020-05-19	Added Array structure for address objects on page 812.
2020-05-22	Added SSL VPN with RADIUS on Windows NPS on page 1331.
2020-05-27	Added SSL VPN with multiple RADIUS servers on page 1336.
2020-06-01	Updated Using wildcard FQDN addresses in firewall policies on page 808 and VXLAN over IPsec tunnel on page 1267.
2020-06-05	Added Integrating FortiManager management using SAML SSO on page 190.
2020-06-12	Updated VLAN inside VXLAN on page 367. Added Configuring least privileges for LDAP admin account authentication in Active Directory on page 1419.
2020-06-18	Updated DHCP options on page 429.
2020-06-22	Updated SSL VPN multi-realm on page 1370.
2020-06-23	Added Proxy chaining (web proxy forwarding servers) on page 421.
2020-07-03	Updated FortiClient EMS on page 132.
2020-08-14	Added Self-originating traffic on page 489.
2020-08-20	Updated Cisco ACI SDN connector on page 232 and Nuage SDN connector on page 236.
2020-08-20	Updated Reliable web filter statistics on page 937.
2020-08-25	Updated FortiSwitch multi-tenant support on page 1523.
2020-08-31	Updated Add FortiToken multi-factor authentication on page 1127, SSL VPN best practices on page 1283, Set up FortiToken multi-factor authentication on page 1290, and Configuring FortiToken Mobile on page 1443.

Netskope Exhibit 1017

Date	Change Description
2020-09-02	Updated Encryption algorithms on page 1259.
2020-09-04	Added General IPsec VPN configuration on page 1046.
2020-09-25	Updated SSL VPN best practices on page 1283 and SSL VPN with certificate authentication on page 1350.
2020-10-02	Added VM license on page 1576.
2020-10-05	Updated Threat feeds on page 270 and External resources for DNS filter on page 279.
2020-10-22	Updated Agentless NTLM authentication for web proxy on page 425 and LLDP reception on page 438.
2021-01-08	Updated FortiGuard outbreak prevention on page 862.
2021-01-21	Updated Matching multiple parameters on application control signatures on page 971.
2021-02-17	Updated Antivirus on page 856.
2021-05-31	Updated IPsec VPN to an Azure with virtual WAN on page 1111.
2021-06-03	Updated Adding top FortiView widgets by category on page 75.

Netskope Exhibit 1017

Getting started

This section explains how to get started with a FortiGate.

Differences between models

Not all FortiGates have the same features, particularly entry-level models (models 30 to 90). A number of features on these models are only available in the CLI.



Consult your model's QuickStart Guide, hardware manual, or the Feature / Platform Matrix for further information about features that vary by model.

FortiGate models differ principally by the names used and the features available:

- Naming conventions may vary between FortiGate models. For example, on some models the hardware switch interface used for the local area network is called *lan*, while on other units it is called *internal*.
- Certain features are not available on all models. Additionally, a particular feature may be available only through the CLI on some models, while that same feature may be viewed in the GUI on other models.
  If you believe your FortiGate model supports a feature that does not appear in the GUI, go to *System* > *Feature Visibility* and confirm that the feature is enabled. For more information, see Feature visibility on page 694.

Using the GUI

This section presents an introduction to the graphical user interface (GUI) on your FortiGate.

The following topics are included in this section:

- Connecting using a web browser
- Menus
- Tables
- Entering values

For information about using the dashboards, see Dashboards and widgets on page 56.

Connecting using a web browser

In order to connect to the GUI using a web browser, an interface must be configured to allow administrative access over HTTPS or over both HTTPS and HTTP. By default, an interface has already been set up that allows HTTPS access with the IP address 192.168.1.99.

Browse to https://192.168.1.99 and enter your username and password. If you have not changed the admin account's password, use the default user name, `admin`, and leave the password field blank.

The GUI will now be displayed in your browser.

To use a different interface to access the GUI:

1. Go to *Network* > *Interfaces* and edit the interface you wish to use for access. Take note of its assigned IP address.
2. In *Administrative Access*, select *HTTPS*, and any other protocol you require. You can also select *HTTP*, although this is not recommended as the connection will be less secure.
3. Click *OK*.
4. Browse to the IP address using your chosen protocol.
   The GUI will now be displayed in your browser.

Menus



If you believe your FortiGate model supports a menu that does not appear in the GUI, go to *System* > *Feature Visibility* and ensure the feature is enabled. For more information, see Feature visibility on page 694.

The GUI contains the following main menus, which provide access to configuration options for most FortiOS features:

Dashboard	The dashboard displays various widgets that display important system information and allow you to configure some system options. For more information, see Dashboards and widgets on page 56.
Security Fabric	Access the physical topology, logical topology, automation, and settings of the Fortinet Security Fabric. For more information, see Fortinet Security Fabric on page 112.
Network	Options for networking, including configuring system interfaces and routing options. For more information, see Network on page 340.
System	Configure system settings, such as administrators, FortiGuard, and certificates. For more information, see System on page 585.
Policy & Objects	Configure firewall policies, protocol options, and supporting content for policies, including schedules, firewall addresses, and traffic shapers. For more information, see Policy and Objects on page 706.
Security Profiles	Configure your FortiGate's security features, including Antivirus, Web Filter, and Application Control. For more information, see Security Profiles on page 851.
VPN	Configure options for IPsec and SSL virtual private networks (VPNs). For more information, see IPsec VPNs on page 1046 and SSL VPN on page 1282.
User & Authetication	Configure user accounts, groups, and authentication methods, including external authentication and single sign-on (SSO).

WiFi & Switch Controller	Configure the unit to act as a wireless network controller, managing the wireless Access Point (AP) functionality of FortiWiFi and FortiAP units. On certain FortiGate models, this menu has additional features allowing for FortiSwitch units to be managed by the FortiGate. For more information, see Wireless configuration on page 1454 and Switch Controller on page 1455.
Log & Report	Configure logging and alert email as well as reports. For more information, see Log and Report on page 1533.

Tables

Many GUI pages contain tables of information that can be filtered and customized to display specific information in a specific way. Some tables allow content to be edited directly on that table, or rows to be copied and pasted.

Navigation

Some tables contain information and lists that span multiple pages. Navigation controls will be available at the bottom of the page.

Filters

Filters are used to locate a specific set of information or content in a table. They can be particularly useful for locating specific log entries. The filtering options vary, depending on the type of information in the log.

Depending on the table content, filters can be applied using the filter bar, using a column filter, or based on a cell's content. Some tables allow filtering based on regular expressions.

Administrators with read and write access can define filters. Multiple filters can be applied at one time.

To manually create a filter:

1. Click *Add Filter* at the top of the table. A list of the fields available for filtering is shown.
2. Select the field to filter by.
3. Enter the value to filter by, adding modifiers as needed.
4. Press *Enter* to apply the filter.

To create a column filter:

1. Click the filter icon on the right side of the column header
2. Choose a filter type from the available options.
3. Enter the filter text, or select from the available values.
4. Click *Apply*.

To create a filter based on a cell's content:

1. Right click on a cell in the table.
2. Select a filtering option from the menu.

Netskope Exhibit 1017

Column settings

Columns can be rearranged, resized, and added or removed from tables.

To add or remove columns:

1. Right a column header, or click the gear icon on the left side of the header row that appears when hovering the cursor over the headers.
2. Select columns to add or remove.
3. Click *Apply*.

To rearrange the columns in a table:

1. Click and drag the column header.

To resize a column:

1. Click and drag the right border of the column header.

To resize a column to fit its contents:

1. Click the dots or filter icon on the right side of the column header and select *Resize to Contents*.

To resize all of the columns in a table to fit their content:

1. Right a column header, or click the gear icon on the left side of the header row that appears when hovering the cursor over the headers.
2. Click *Best Fit All Columns*.

To reset a table to its default view:

1. Right a column header, or click the gear icon on the left side of the header row that appears when hovering the cursor over the headers.
2. Click *Reset Table*.
   Resetting a table does not remove filters.

Editing objects

In some tables, parts of a configuration can be edited directly in the table. For example, security profiles can be added to an existing firewall policy by clicking the edit icon in a cell in the *Security Profiles* column.

Copying rows

In some tables, rows can be copied and pasted using the right-click menu. For example, a policy can be duplicated by copying and pasting it.

Netskope Exhibit 1017

Entering values

Numerous fields in the GUI and CLI require text strings or numbers to be entered when configuring the FortiGate. When entering values in the GUI, you will be prevented from entering invalid characters, and a warning message will be shown explaining what values are not allowed. If invalid values are entered in a CLI command, the setting will be rejected when you apply it.

- Text strings on page 23
- Numbers on page 24

Text strings

Text strings are used to name entities in the FortiGate configuration. For example, the name of a firewall address, administrator, or interface are all text strings.

The following characters cannot be used in text strings, as they present cross-site scripting (XSS) vulnerabilities:

- " - double quotes
- ' - single quote
- > - greater than
- < - less than

Most GUI text fields prevent XSS vulnerable characters from being added.



VDOM names and hostnames can only use numbers (0-9), letters (a-z and A-Z), dashes, and underscores.

The `tree` CLI command can be used to view the number of characters allowed in a name field. For example, entering the following commands show that a firewall address name can contain up to 80 characters, while its FQDN can contain 256 characters:

```
config fire address
(address) # tree
    -- [address] --*name     (80)
              |- uuid
              |- subnet
              |- type
              |- start-mac
              |- end-mac
              |- start-ip
              |- end-ip
              |- fqdn     (256)
              |- country     (3)
              |- wildcard-fqdn     (256)
              |- cache-ttl     (0,86400)
              |- wildcard
              |- sdn     (36)
              |- interface     (36)
              |- tenant     (36)
              |- organization     (36)
              |- epg-name     (256)
              |- subnet-name     (256)
```


Netskope Exhibit 1017

```
|- sdn-tag    (16)
|- policy-group    (16)
|- comment
|- visibility
|- associated-interface    (36)
|- color    (0,32)
|- filter
|- sdn-addr-type
|- obj-id
|- [list] --*ip    (36)
          |- obj-id    (128)
          +- net-id    (128)
|- [tagging] --*name    (64)
             |- category    (64)
             +- [tags] --*name    (80)
+- allow-routing
```

Numbers

Numbers are used to set sizes, rated, addresses, port numbers, priorities, and other such numeric values. They can be entered as a series of digits (without commas or spaces), in a dotted decimal format (such as IP addresses), or separated by colons (such as MAC addresses). Most numeric values use base 10 numbers, while some use hexadecimal values.

Most GUI and CLI fields prevent invalid numbers from being entered. The CLI help text includes information about the range of values allowed for applicable settings.

Using the CLI

The Command Line Interface (CLI) can be used in lieu of the GUI to configure the FortiGate. Some settings are not available in the GUI, and can only be accessed using the CLI.

This section briefly explains basic CLI usage. For more information about the CLI, see the FortiOS CLI Reference.

- Connecting to the CLI on page 24
- CLI basics on page 27
- Command syntax on page 33
- Subcommands on page 36
- Permissions on page 38

Connecting to the CLI

You can connect to the CLI using a direct console connection, SSH, the CLI console in the GUI, or the FortiExplorer app on your iOS device.

You can access the CLI in three ways:

- Console connection: Connect your computer directly to the console port of your FortiGate.
- SSH access: Connect your computer through any network interface attached to one of the network ports on your FortiGate.

Netskope Exhibit 1017

- FortiExplorer: Connect your device to the FortiExplorer app on your iOS device to configure, manage, and monitor your FortiGate. See FortiExplorer for iOS on page 38 for details.

Console connection

A direct console connections to the CLI is created by directly connecting your management computer or console to the FortiGate unit, using its DB-9 or RJ-45 console port.

Direct console access to the FortiGate may be required if:

- You are installing the FortiGate for the first time and it is not configured to connect to your network.
- You are restoring the firmware using a boot interrupt. Network access to the CLI will not be available until after the boot process has completed, making direct console access the only option.

To connect to the FortiGate console, you need:

- A computer with an available communications port
- A console cable to connect the console port on the FortiGate to a communications port on the computer (a USB adapter may also be required)
- Terminal emulation software

To connect to the CLI using a direct console connection:

1. Using the null modem or RJ-45 to USB (or DB-9) cable, connect the FortiGate unit's console port to the serial communications (COM) port on your management computer. A DB-9-to-USB adapter may be required.
2. Start a terminal emulation program on the management computer, select the COM port, and use the following settings:

Bits per second	9600
Data bits	8
Parity	None
Stop bits	1
Flow control	None

3. Press Enter on the keyboard to connect to the CLI.
4. Log in to the CLI using your username and password (default: *admin* and no password).
5. You can now enter CLI commands, including configuring access to the CLI through SSH.

SSH access

SSH access to the CLI is accomplished by connecting your computer to the FortiGate unit using one of its network ports. You can either connect directly, using a peer connection between the two, or through any intermediary network.



If you do not want to use an SSH client and you have access to the GUI, you can access the CLI through the network using the CLI console in the GUI.

The CLI console can be accessed from the upper-right hand corner of the screen and appears as a slide-out window. For policies and objects, the CLI can be also be accessed by right clicking on the element and selecting *Edit in CLI*.

Netskope Exhibit 1017

SSH must be enabled on the network interface that is associated with the physical network port that is used.

If your computer is not connected either directly or through a switch to the FortiGate, you must also configure the FortiGate with a static route to a router that can forward packets from the FortiGate to the computer. This can be done using a local console connection, or in the GUI.

To connect to the FortiGate CLI using SSH, you need:

- A computer with an available serial communications (COM) port and RJ-45 port
- The RJ-45 to USB (or DB-9) or null modem cable included in your FortiGate package
- Terminal emulation software
- A network cable
- Prior configuration of the operating mode, network interface, and static route.

To enable SSH access to the CLI using a local console connection:

1. Using the network cable, connect the FortiGate unit's port either directly to your computer's network port, or to a network through which your computer can reach the FortiGate unit.
2. Note the number of the physical network port.
3. Using direct console connection, connect and log into the CLI.
4. Enter the following command:

```
config system interface
    edit <interface_str>
        append allowaccess ssh
    next
end
```

   Where `<interface_str>` is the name of the network interface associated with the physical network port, such as `port1`.
5. Confirm the configuration using the following command to show the interface's settings:

```
show system interface <interface_str>
```

   For example:

```
show system interface port1
    config system interface
        edit "port1"
            set vdom "root"
            set ip 192.168.1.99 255.255.255.0
            set allowaccess ping https ssh
            set type hard-switch
            set stp enable
            set role lan
            set snmp-index 6
        next
    end
```

Connecting using SSH

Once the FortiGate unit is configured to accept SSH connections, use an SSH client on your management computer to connect to the CLI.

The following instructions use PuTTy. The steps may vary in other terminal emulators.

Netskope Exhibit 1017

To connect to the CLI using SSH:

1. On your management computer, start PuTTy.
2. In the *Host Name (or IP address)* field, enter the IP address of the network interface that you are connected to and that has SSH access enabled.
3. Set the port number to 22, if it is not set automatically.
4. Select *SSH* for the *Connection type*.
5. Click *Open*. The SSH client connect to the FortiGate.
   The SSH client may display a warning if this is the first time that you are connecting to the FortiGate and its SSH key is not yet recognized by the SSH client, or if you previously connected to the FortiGate using a different IP address or SSH key. This is normal if the management computer is connected directly to the FortiGate with no network hosts in between.
6. Click *Yes* to accept the FortiGate unit's SSH key.
   The CLI displays the log in prompt.
7. Enter a valid administrator account name, such as `admin`, then press *Enter*.
8. Enter the administrator account password, then press *Enter*.
   The CLI console shows the command prompt (FortiGate hostname followed by a `#`). You can now enter CLI commands.



If three incorrect log in or password attempts occur in a row, you will be disconnected. If this occurs, wait for one minute, then reconnect and attempt to log in again.

CLI basics

Basic features and characteristics of the CLI environment provide support and ease of use for many CLI tasks.

Help

Press the question mark (?) key to display command help and complete commands.

- Press the question mark (?) key at the command prompt to display a list of the commands available and a description of each command.
- Enter a command followed by a space and press the question mark (?) key to display a list of the options available for that command and a description of each option.
- Enter a command followed by an option and press the question mark (?) key to display a list of additional options available for that command option combination and a description of each option.
- Enter a question mark after entering a portion of a command to see a list of valid complete commands and their descriptions. If there is only one valid command, it will be automatically filled in.

Shortcuts and key commands

Shortcut key	Action
?	List valid complete or subsequent commands.

Shortcut key	Action
	If multiple commands can complete the command, they are listed with their descriptions.
Tab	Complete the word with the next available match. Press multiple times to cycle through available matches.
Up arrow or Ctrl + P	Recall the previous command. Command memory is limited to the current session.
Down arrow, or Ctrl + N	Recall the next command.
Left or Right arrow	Move the cursor left or right within the command line.
Ctrl + A	Move the cursor to the beginning of the command line.
Ctrl + E	Move the cursor to the end of the command line.
Ctrl + B	Move the cursor backwards one word.
Ctrl + F	Move the cursor forwards one word.
Ctrl + D	Delete the current character.
Ctrl + C	Abort current interactive commands, such as when entering multiple lines. If you are not currently within an interactive command such as `config` or `edit`, this closes the CLI connection.
\ then Enter	Continue typing a command on the next line for a multiline command. For each line that you want to continue, terminate it with a backslash (\ ). To complete the command, enter a space instead of a backslash, and then press *Enter*.

Command tree

Enter `tree` to display the CLI command tree. To capture the full output, connect to your device using a terminal emulation program and capture the output to a log file. For some commands, use the `tree` command to view all available variables and subcommands.

Command abbreviation

You can abbreviate words in the command line to their smallest number of non-ambiguous characters.

For example, the command `get system status` could be abbreviated to `g sy stat`.

Adding and removing options from lists

When configuring a list, the `set` command will remove the previous configuration.

For example, if a user group currently includes members A, B, and C, the command `set member D` will remove members A, B, and C. To avoid removing the existing members from the group, the command `set members A B C D` must be used.

Netskope Exhibit 1017

To avoid this issue, the following commands are available:

append	Add an option to an existing list. For example, `append member D` adds user D to the user group without removing any of the existing members.
select	Clear all of the options except for those specified. For example, `select member B` removes all member from the group except for member B.
unselect	Remove an option from an existing list. For example, `unselect member C` removes only member C from the group, without affecting the other members.

Environment variables

The following environment variables are support by the CLI. Variable names are case-sensitive.

$USERFROM	The management access type (`ssh`, `jsconsole`, and so on) and the IPv4 address of the administrator that configured the item.
$USERNAME	The account name of the administrator that configured the item.
$SerialNum	The serial number of the FortiGate unit.

For example, to set a FortiGate device's host name to its serial number, use the following CLI command:

```
config system global
   set hostname $SerialNum
end
```

Special characters

The following characters cannot be used in most CLI commands: `<`, `>`, `(`, `)`, `#`, `'`, and `"`

If one of those characters, or a space, needs to be entered as part of a string, it can be entered by using a special command, enclosing the entire string in quotes, or preceding it with an escape character (backslash, \).

To enter a question mark (`?`) or a tab, Ctrl + V must be entered first.



Question marks and tabs cannot be typed or copied into the CLI Console or some SSH clients.

Character	Keys
?	Ctrl + V then ?
Tab	Ctrl + V then Tab
Space (as part of a string value, not to end the string)	Enclose the string in single or double quotation marks: `"Security Administrator"`. or `'Security Administrator'`.

Character	Keys
	Precede the space with a backslash: `Security\ Administrator`.
' (as part of a string value, not to begin or end the string)	\'
" (as part of a string value, not to begin or end the string)	\"
\	\\

Using grep to filter command output

The `get`, `show`, and `diagnose` commands can produce large amounts of output. The `grep` command can be used to filter the output so that it only shows the required information.

The grep command is based on the standard UNIX grep, used for searching text output based on regular expressions.

For example, the following command displays the MAC address of the internal interface:

```
get hardware nic internal | grep Current_HWaddr
    Current_HWaddr        00:09:0f:cb:c2:75
```

The following command will display all TCP sessions that are in the session list, including the session list line number in the output:

```
get system session list | grep -n tcp
```

The following command will display all of the lines in the HTTP replacement message that contain `URL` or `url`:

```
show system replacemsg http | grep -i url
```

The following options can also be used:

```
-A <num> After
-B <num> Before
-C <num> Context
```

The `-f` option is available to support contextual output, in order to show the complete configuration. The following example shows the difference in the output when `-f` is used versus when it is not used:

Without `-f`:

```
show | grep ldap-group1
    edit "ldap-group1"
        set groups "ldap-group1"
```

With `-f`:

```
show | grep -f ldap-group1
    config user group
        edit "ldap-group1"
            set member "pc40-LDAP"
        next
    end
    config firewall policy
        edit 2
            set srcintf "port31"
            set dstintf "port32"
            set srcaddr "all"
```

```
            set action accept
            set identity-based enable
            set nat enable
            config identity-based-policy
                edit 1
                    set schedule "always"
                    set groups "ldap-group1"
                    set dstaddr "all"
                    set service "ALL"
                next
            end
        next
    end
```

Language support and regular expressions

Characters such as ñ and é, symbols, and ideographs are sometimes acceptable input. Support varies depending on the type of item that is being configured. CLI commands, objects, field names, and options must use their exact ASCII characters, but some items with arbitrary names or values can be input using your language of choice. To use other languages in those cases, the correct encoding must be used.

Input is stored using Unicode UTF-8 encoding, but is not normalized from other encodings into UTF-8 before it is stored. If your input method encodes some characters differently than in UTF-8, configured items may not display or operate as expected.

Regular expressions are especially impacted. Matching uses the UTF-8 character values. If you enter a regular expression using a different encoding, or if an HTTP client sends a request in a different encoding, matches may not be what is expected.

For example, with Shift-JIS, backslashes could be inadvertently interpreted as the symbol for the Japanese yen (¥), and vice versa. A regular expression intended to match HTTP requests containing monetary values with a yen symbol may not work it if the symbol is entered using the wrong encoding.

For best results:

- use UTF-8 encoding, or
- use only characters whose numerically encoded values are the same in UTF-8, such as the US-ASCII characters that are encoded using the same values in ISO 8859-1, Windows code page 1252, Shift-JIS, and other encoding methods, or
- for regular expressions that must match HTTP requests, use the same encoding as your HTTP clients.



HTTP clients may send requests in encodings other than UTF-8. Encodings usually vary based on the client's operating system or input language. If the client's encoding method cannot be predicted, you might only be able to match the parts of the request that are in English, as the values for English characters tend to be encoded identically, regardless of the encoding method.

If the FortiGate is configured to use an encoding method other than UTF-8, the management computer's language may need to be changed, including the web browse and terminal emulator. If the FortiGate is configured using non-ASCII characters, all the systems that interact with the FortiGate must also support the same encoding method. If possible, the

Netskope Exhibit 1017

same encoding method should be used throughout the configuration to avoid needing to change the language settings on the management computer.

The GUI and CLI client normally interpret output as encoded using UTF-8. If they do not, configured items may not display correctly. Exceptions include items such as regular expression that may be configured using other encodings to match the encoding of HTTP requests that the FortiGate receives.

To enter non-ASCII characters in a terminal emulator:

1. On the management computer, start the terminal client.
2. Configure the client to send and receive characters using UTF-8 encoding.
   Support for sending and receiving international characters varies by terminal client.
3. Log in to the FortiGate unit.
4. At the command prompt, type your command and press *Enter*.
   Words that use encoded characters may need to be enclosed in single quotes (').
   Depending on your terminal client's language support, you may need to interpret the characters into character codes before pressing *Enter*. For example, you might need to enter: `edit '\743\601\613\743\601\652'`
5. The CLI displays the command and its output.

Screen paging

By default, the CLI will pause after displaying each page worth of text when a command has multiple pages of output. this can be useful when viewing lengthy outputs that might exceed the buffer of terminal emulator.

When the display pauses and shows `--More--`, you can:

- Press *Enter* to show the next line,
- Press *Q* to stop showing results and return to the command prompt,
- Press an arrow key, *Insert*, *Home*, *Delete*, *End*, *Page Up*, or *Page Down* to show the next few pages,
- Press any other key to show the next page, or
- Wait for about 30 seconds for the console to truncate the output and return to the command prompt.

When pausing the screen is disable, press *Ctrl* + *C* to stop the output and log out of the FortiGate.

To disable pausing the CLI output:

```
config system console
   set output standard
end
```

To enable pausing the CLI output:

```
config system console
   set output more
end
```

Changing the baud rate

The baud rate of the local console connection can be changed from its default value of 9600.

Netskope Exhibit 1017

To change the baud rate:

```
config system console
   set baudrate {9600 | 19200 | 38400 | 57600 | 115200}
end
```

Editing the configuration file

The FortiGate configuration file can be edited on an external host by backing up the configuration, editing the configuration file, and then restoring the configuration to the FortiGate.

Editing the configuration file can save time is many changes need to be made, particularly if the plain text editor that you are using provides features such as batch changes.

To edit the configuration file:

1. Backup the configuration. See Configuration backups on page 108 for details.
2. Open the configuration file in a plain text editor that supports UNIX-style line endings.
3. Edit the file as needed.



Do not edit the first line of the configuration file.

This line contains information about the firmware version and FortiGate model. If you change the model number, the FortiGate unit will reject the configuration when you attempt to restore it.

4. Restore the modified configuration to the FortiGate. SeeConfiguration backups on page 108 for details.
The FortiGate downloads the configuration file and checks that the model information is correct. If it is correct, the configuration file is loaded and each line is checked for errors. If a command is invalid, that command is ignored. If the configuration file is valid, the FortiGate restarts and loads the downloaded configuration.

Command syntax

When entering a command, the CLI console requires that you use valid syntax and conform to expected input constraints. It rejects invalid commands. Indentation is used to indicate the levels of nested commands.

Each command line consists of a command word, usually followed by configuration data or a specific item that the command uses or affects.

Notation

Brackets, vertical bars, and spaces are used to denote valid syntax. Constraint notations, such as `<address_ipv4>`, indicate which data types or string patterns are acceptable value input.

All syntax uses the following conventions:

Angle brackets < >	Indicate a variable of the specified data type.
Curly brackets { }	Indicate that a variable or variables are mandatory.
Square brackets [ ]	Indicate that the variable or variables are optional.

Netskope Exhibit 1017

	For example: `show system interface [<name_str>]` To show the settings for all interfaces, you can enter `show system interface` To show the settings for the Port1 interface, you can enter `show system interface port1`.
Vertical bar |	A vertical bar separates alternative, mutually exclusive options. For example: `set protocol {ftp | sftp}` You can enter either `set protocol ftp` or `set protocol sftp`.

Space	A space separates non-mutually exclusive options.
	For example:
	`set allowaccess {ping https ssh snmp http fgfm radius-acct probe-response capwap ftm}`
	You can enter any of the following:
	`set allowaccess ping`
	`set allowaccess https ping ssh`
	`set allowaccess http https snmp ssh ping`
	In most cases, to make changes to lists that contain options separated by spaces, you need to retype the entire list, including all the options that you want to apply and excluding all the options that you want to remove.

Optional values and ranges

Any field that is optional will use square-brackets. The overall config command will still be valid whether or not the option is configured.

Square-brackets can be used is to show that multiple options can be set, even intermixed with ranges. The following example shows a field that can be set to either a specific value or range, or multiple instances:

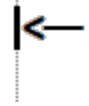
```
config firewall service custom
   set iprange <range1> [<range2> <range3> ...]
end
```

next

The `next` command is used to maintain a hierarchy and flow to CLI commands. It is at the same indentation level as the preceding `edit` command, to mark where a table entry finishes.

The following example shows the next command used in the subcommand `entries`:

```
config dlp filepattern
    edit <1>
        set name <name>
        set comment [comment]
        config entries
            edit <2>
                set filter-type {pattern | type}
            next
```



After configuring table entry <2> then entering `next`, the <2> table entry is saved and the console returns to the `entries` prompt:

```
FGT60E1Q23456789 (entries) #
```

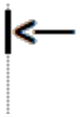
You can now create more table entries as needed, or enter `end` to save the table and return to the `filepattern` table element prompt.

end

The `end` command is used to maintain a hierarchy and flow to CLI commands.

The following example shows the same command and subcommand as the `next` command example, except `end` has been entered instead of `next` after the subcommand:

```
config dlp filepattern
    edit <1>
        set name <name>
        set comment [comment]
        config entries
            edit <2>
                set filter-type {pattern | type}
    end
```



Entering `end` will save the <2> table entry and the table, and exit the `entries` subcommand entirely. The console returns to the filepattern table element prompt:

```
FGT60E1Q23456789 (1) #
```

Subcommands

Subcommands are available from within the scope of some commands. When you enter a subcommand level, the command prompt changes to indicate the name of the current command scope. For example, after entering:

```
config system admin
```

the command prompt becomes:

```
(admin)#
```

Applicable subcommands are available until you exit the command, or descend an additional level into another subcommand. Subcommand scope is indicated by indentation.

For example, the `edit` subcommand is only available in commands that affects tables, and the `next` subcommand is available only in the `edit` subcommand:

```
config system interface
   edit port1
      set status up
   next
end
```

The available subcommands vary by command. From a command prompt under the `config` command, subcommands that affect tables and fields could be available.

Table subcommands

edit <table_row>	Create or edit a table value.
	In objects such as security policies, `<table_row>` is a sequence number. To create a new table entry without accidentally editing an existing entry, enter `edit 0`. The CLI will confirm that creation of entry 0, but will assign the next unused number when the entry is saved after entering `end` or `next`.
	For example, to create a new firewall policy, enter the following commands:
	`config firewall policy`
	`edit 0`
	`....`
	`next`
	`end`
	To edit an existing policy, enter the following commands:
	`config firewall policy`
	`edit 27`
	`....`
	`next`
	`end`
	The `edit` subcommand changes the command prompt to the name of the table value that is being edited.

delete <table_row>	Delete a table value.
	For example, to delete firewall policy 30, enter the following commands:
	`config firewall policy`
	`delete 30`
	`end`

purge	Clear all table values. The `purge` command cannot be undone. To restore purged table values, the configuration must be restored from a backup.

move	Move an ordered table value.
	In the firewall policy table, this equivalent to dragging a policy into a new position. It does not change the policy's ID number.
	For example, to move policy 27 to policy 30, enter the following commands:
	`config firewall policy`
	`move 27 to 30`
	`end`
	The `move` subcommand is only available in tables where the order of the table entries matters.

clone <table_row> to <table_row>	Make a clone of a table entry.
	For example, to create firewall policy 30 as a clone of policy 27, enter the following commands:
	`config firewall policy`
	`clone 27 to 30`
	`end`
	The `clone` subcommand may not be available for all tables.

rename <table_row> to <table_row>	Rename a table entry.
	For example to rename an administrator from Flank to Frank, enter the following commands:
	`config system admin`
	`rename Flank to Frank`
	`end`
	The `rename` subcommand is only available in tables where the entries can be renamed.

get	List the current table entries.
	For example, to view the existing firewall policy table entries, enter the following commands:
	`config firewall policy`
	`get`

show	Show the configuration. Only table entries that are not set to default values are shown.
end	Save the configuration and exit the current `config` command.



Purging the `system interface` or `system admin` tables does not reset default table values. This can result in being unable to connect to or log in to the FortiGate, requiring the FortiGate to be formatted and restored.

Field subcommands

set <field> <value>	Modify the value of a field.

	For example, the command `set fsso enable` sets the `fsso` field to the value `enable`.
unset	Set the field to its default value.
select	Clear all of the options except for those specified. For example, if a group contains members A, B, C, and D, to remove all members except for B, use the command `select member B`.
unselect	Remove an option from an existing list. For example, if a group contains members A, B, C, and D, to remove only member B, use the command `unselect member B`.
append	Add an option to an existing multi-option table value.
clear	Clear all the options from a multi-option table value.
get	List the configuration of the current table entry, including default and customized values.
show	Show the configuration. Only values that are not set to default values are shown.
next	Save changes to the table entry and exit the `edit` command so that you can configure the next table entry.
abort	Exit the command without saving.
end	Save the configuration and exit the current `config` command.

Permissions

Administrator, or access, profiles control what CLI commands an administrator can access by assigning read, write, or no access to each are of FortiOS. For information, see Administrator profiles on page 587.

Read access is required to view configurations. Write access is required to make configuration changes. Depending on your account's profile, you may not have access to all CLI commands. To have access to all CLI commands, an administrator account with the *super_admin* profile must be used, such as the *admin* account.

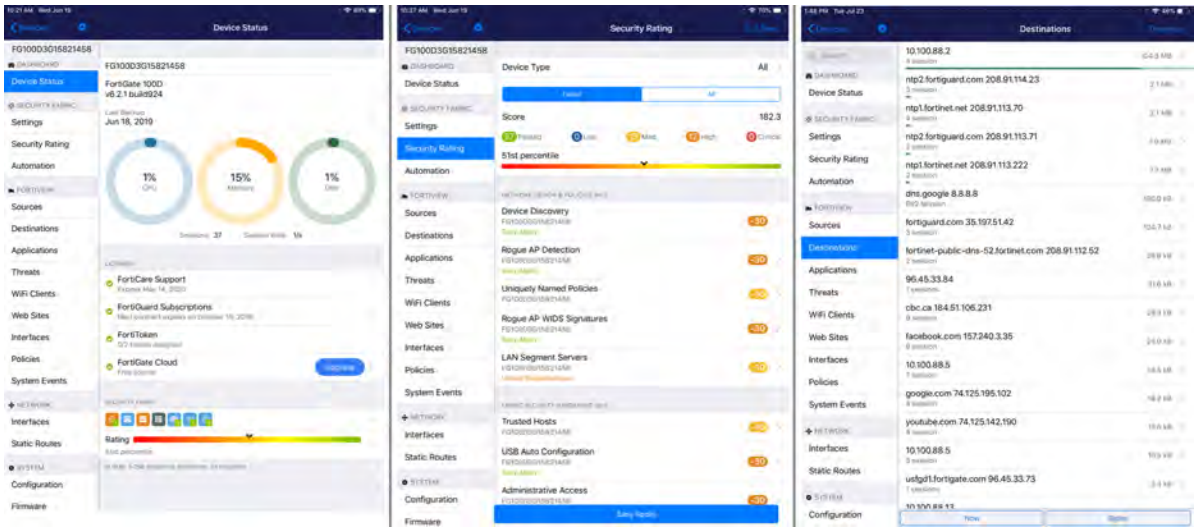
Accounts assigned the *super_admin* profile are similar to the root administrator account. They have full permission to view and change all FortiGate configuration options, including viewing and changing other administrator accounts.

To increase account security, set strong passwords for all administrator accounts, and change the passwords regularly.

FortiExplorer for iOS

FortiExplorer for iOS is a user-friendly application that helps you to rapidly provision, deploy, and monitor Security Fabric components from your iOS device.

Netskope Exhibit 1017



FortiExplorer for iOS requires iOS 10.0 or later and is compatible with iPhone, iPad, and Apple TV. It is supported by FortiOS 5.6 and later, and is only available on the App Store for iOS devices.

Advanced features are available with the purchase of FortiExplorer Pro. Paid features include the ability to add more than two devices and the ability to download firmware images from FortiCare.

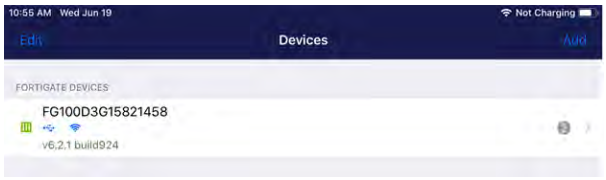
Up to six members can use this app with 'Family Sharing' enabled in the App Store.

Getting started with FortiExplorer

If your FortiGate is accessible on a wireless network, you can connect to it using FortiExplorer provided that your iOS device is on the same network (see Connecting FortiExplorer to a FortiGate via WiFi). Otherwise, you will need to physically connect your iOS device to the FortiGate using a USB cable.

To connect and configure a FortiGate with FortiExplorer using a USB connection:

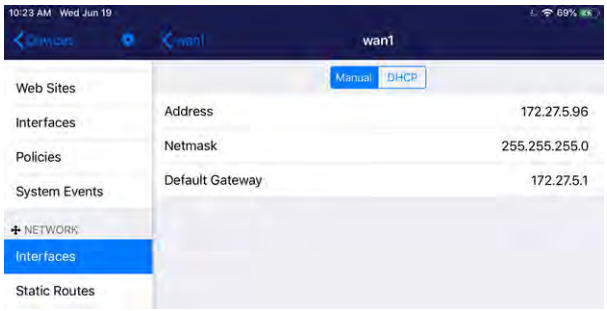
1. Connect your iOS device to your FortiGate USB A port. If prompted on your iOS device, *Trust* this computer.
2. Open FortiExplorer and select your FortiGate from the *FortiGate Devices* list . A blue USB icon will indicate that you are connected over a USB connection.



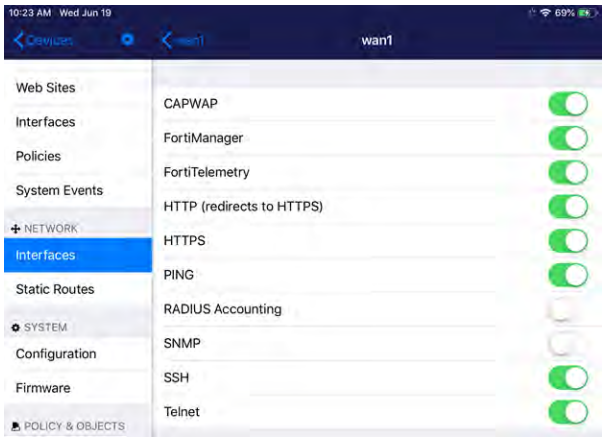
3. On the *Login* screen, select *USB*.
4. Enter the default *Username* (`admin`) and leave the *Password* field blank.
5. Optionally, select *Remember Password*.
6. Tap *Done* when you are ready.
   FortiExplorer opens the FortiGate management interface to the *Device Status* page:



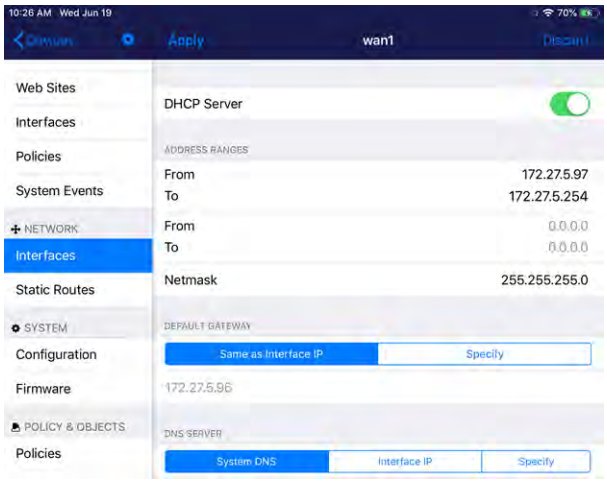
7. Go to *Network* > *Interfaces* and configure the WAN interface or interfaces.
8. The *wan1* interface *Address* mode is set to *DHCP* by default. Set it to *Manual* and enter its *Address*, *Netmask*, and *Default Gateway*, and then *Apply* your changes.



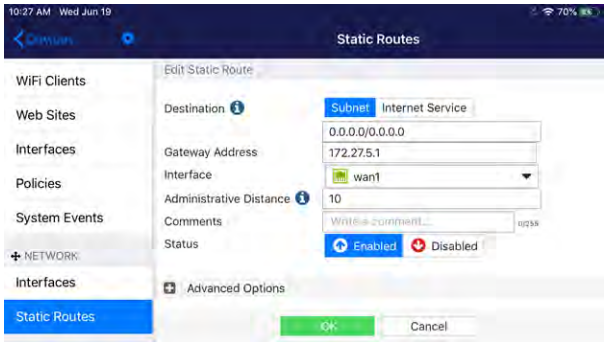
9. Optionally, configure *Administrative Access* to allow *HTTPS* access. This will allow administrators to access the FortiGate GUI using a web browser.



10. Go to *Network* > *Interfaces* and configure the local network (internal) interface.
11. Set the *Address* mode as before and configure *Administrative Access* if required.
12. Configure a *DHCP Server* for the internal network subnet.

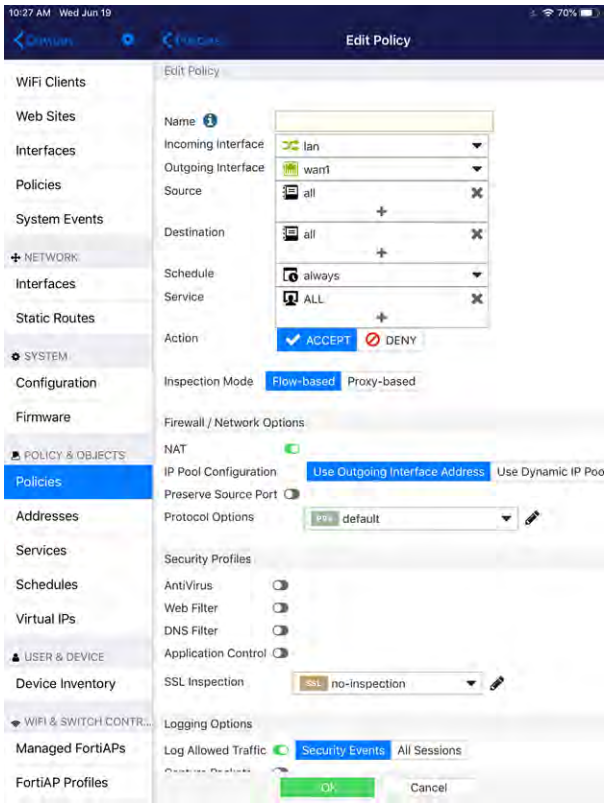


13. Return to the internal interface using the < button at the top of the screen.
14. Go to *Network* > *Static Routes* and configure the static route to the gateway.



15. Go to *Policy & Objects* > *Firewall Policy* and edit the Internet access policy. Enter a *Name* for the policy, enable the required *Security Profiles*, configure *Logging Options*, then tap *OK*.

Netskope Exhibit 1017

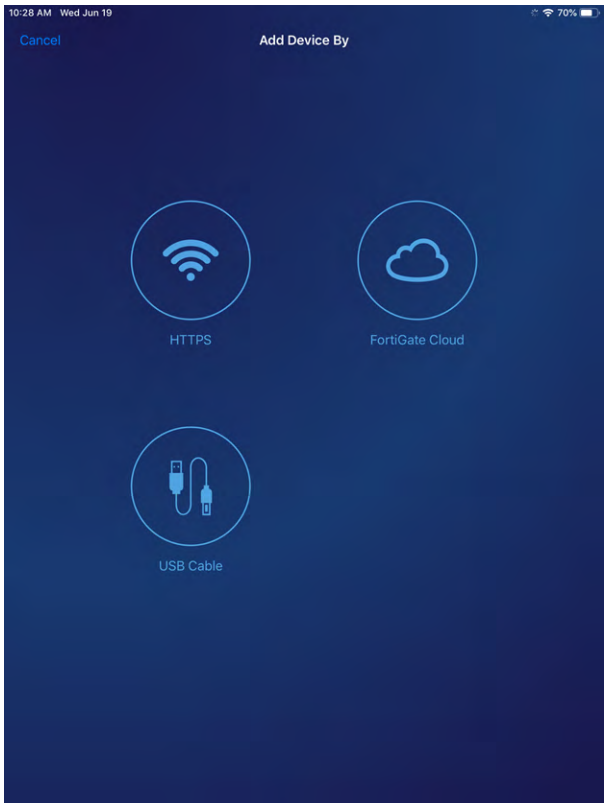


Connecting FortiExplorer to a FortiGate via WiFi

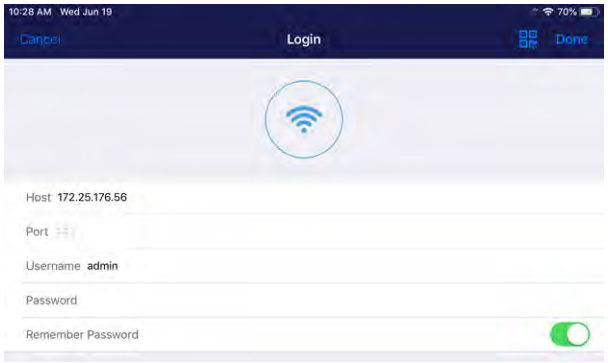
You can wirelessly connect to the FortiGate if your iOS device and the FortiGate are both connected to the same wireless network.

To connect and configure a FortiGate with FortiExplorer wirelessly:

1. Open the FortiExplorer app and tap *Add* on the *Devices* page.
2. On the *Add Device By* page, tap *HTTPS*.



3. Enter the *Host* information, *Username*, and *Password*.
4. If required, change the default *Port* number, and optionally enable *Remember Password*.

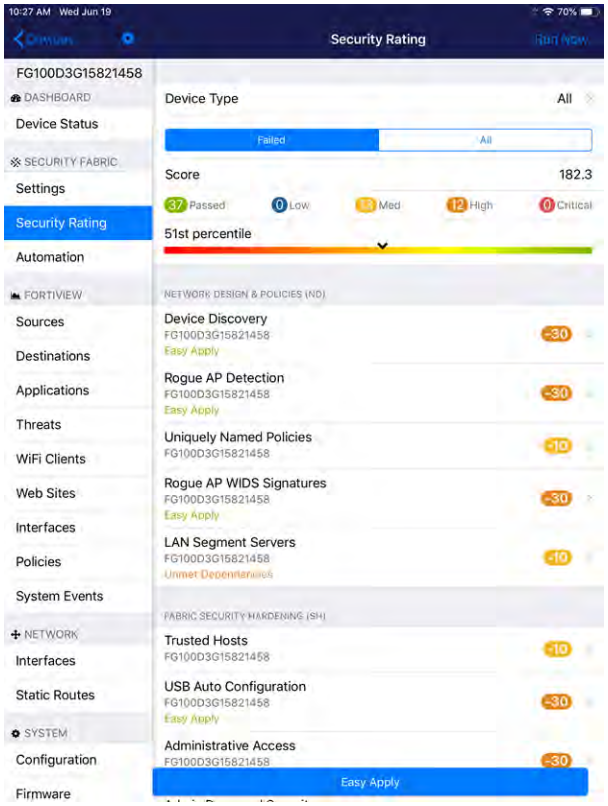


5. Tap *Done*.
6. If the FortiGate device identity cannot be verified, tap *Connect* at the prompt.
   FortiExplorer opens the FortiGate management interface to the *Device Status* page.

Running a security rating

After configuring your network, run a security rating check to identify vulnerabilities and highlight best practices that could improve your network's security and performance.

Go to *Security Fabric* > *Security Rating* and follow the steps to determine the score. See Security rating on page 195 for more information.



Upgrading to FortiExplorer Pro

FortiExplorer Pro includes the option to add more than two devices and download firmware images from FortiCare.

To upgrade to FortiExplorer Pro:

1. In FortiExplorer, go to *Settings*.
2. Tap *Upgrade to FortiExplorer Pro*.
3. Follow the on-screen prompts.

Basic administration

This section contains information about basic FortiGate administration that you can do after you installing the unit in your network.

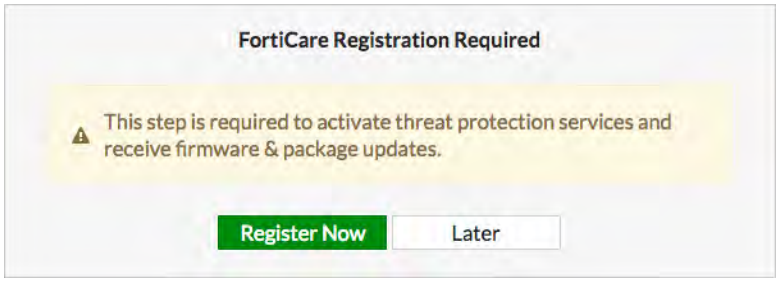
- Registration on page 45
- FortiCare and FortiGate Cloud login on page 46

Registration

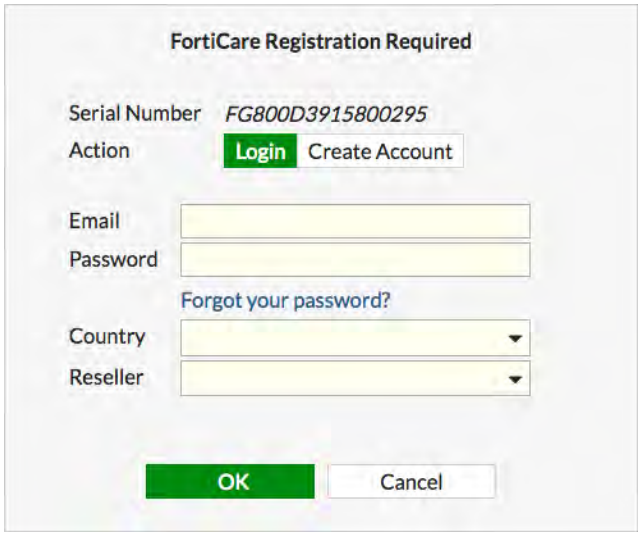
The FortiGate must be registered to have full access to Fortinet Customer Service and Support, and FortiGuard services.

To register your FortiGate:

1. Connect to the FortiGate GUI. A message is shown stating that FortiCare registration is required.

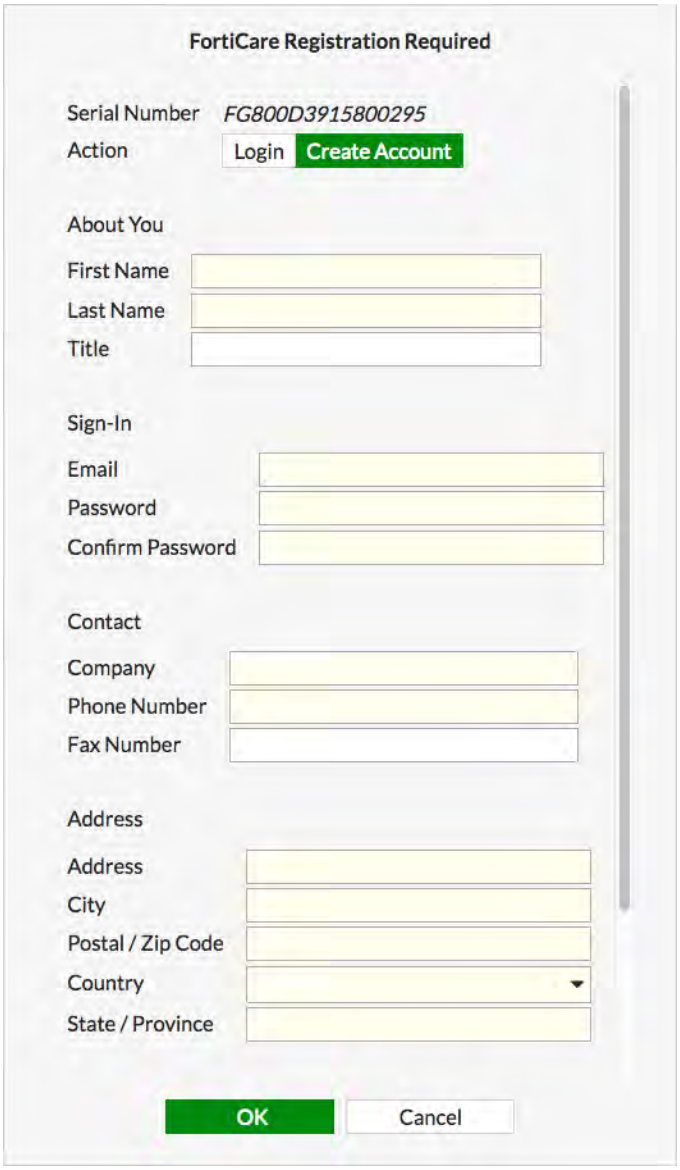


2. Click *Register Now*.
3. If you already have a support account, set *Action* to *Login*, and enter the required information.



If you need to create an account, set *Action* to *Create Account*, and enter the required information.

Netskope Exhibit 1017



4. Click *OK*.
5. Go to *System* > *FortiGuard*.
6. In the *License Information* table, the *FortiCare Support* status is *Registered*. There may be a delay before the status is updated on your FortiGate.

FortiCare and FortiGate Cloud login

With FortiCloud, FortiGate supports a unified login to FortiCare and FortiGate Cloud. The FortiGate Cloud setup is a subset of the FortiCare setup.

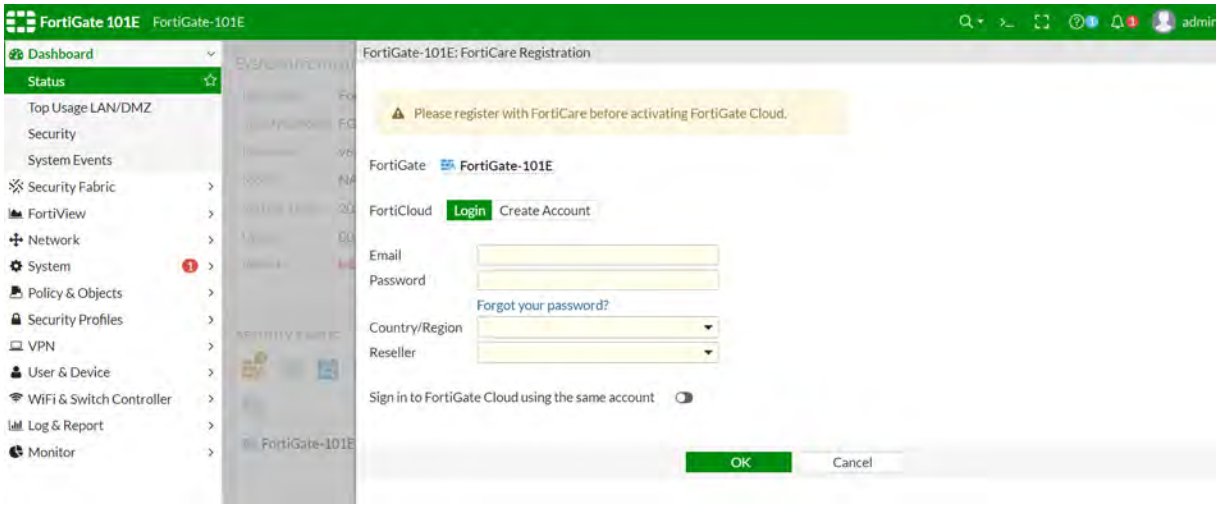
- If the FortiGate is not registered, activating FortiGate Cloud will force you to register with FortiCare.
- If a FortiGate is registered in FortiCare using a FortiCloud account, then only that FortiCloud account can be used to activate FortiGate Cloud.
- If a different FortiCloud account was already used to activate FortiGate Cloud, then a notification asking you to migrate to FortiCloud is shown in the GUI after upgrading FortiOS.

Netskope Exhibit 1017

The CLI can be used to activate FortiGate Cloud without registration, or with a different FortiCloud account.

To activate FortiGate Cloud and register with FortiCare at the same time:

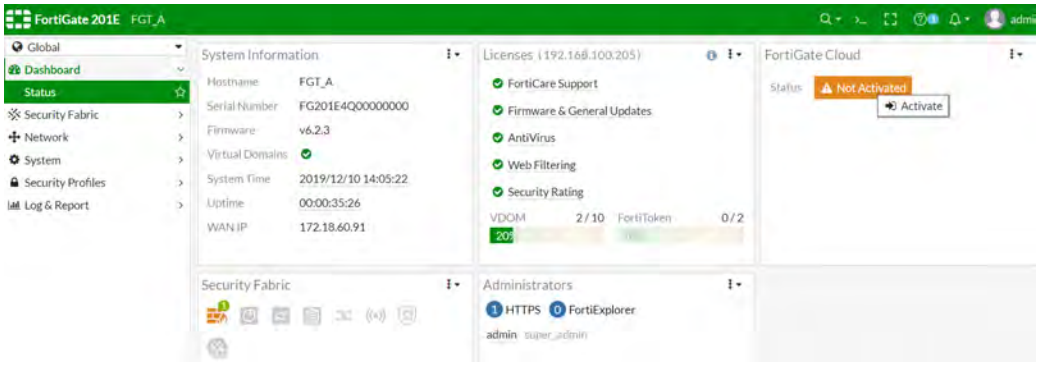
1. Go to *Dashboard* > *Status*.
2. In the FortiGate Cloud widget, click *Not Activated* > *Activate*.
   You must register with FortiCare before activating FortiGate Cloud.



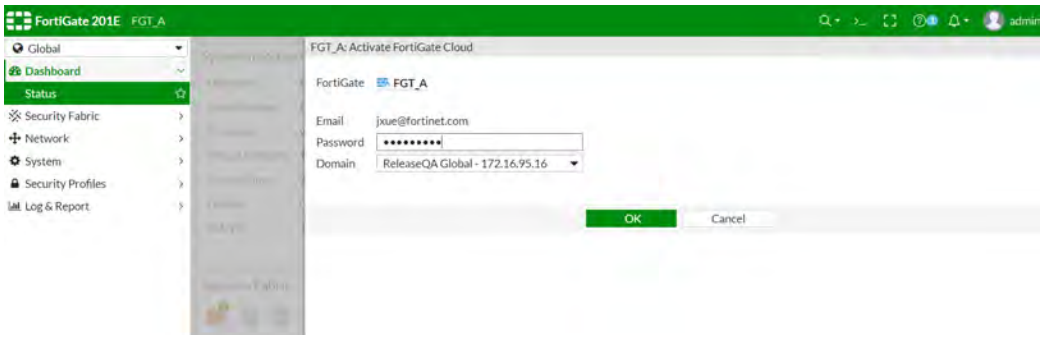
3. Enter your FortiCare *Email* address and *Password*.
4. Select your *Country/Region* and *Reseller*.
5. Enable *Sign in to FortiGate Cloud using the same account*.
6. Click *OK*.

To activate FortiGate Cloud on an already registered FortiGate:

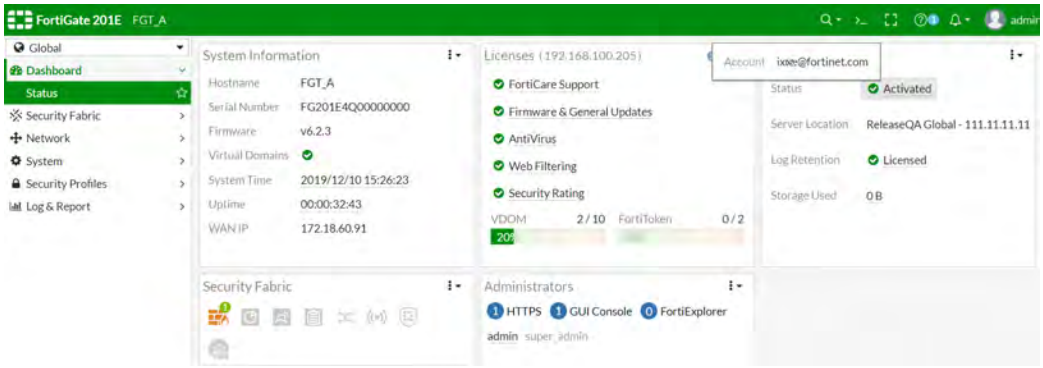
1. Go to *Dashboard* > *Status*.
2. In the FortiGate Cloud widget, click *Not Activated* > *Activate*.



3. Enter the password for the account that was used to register the FortiGate.

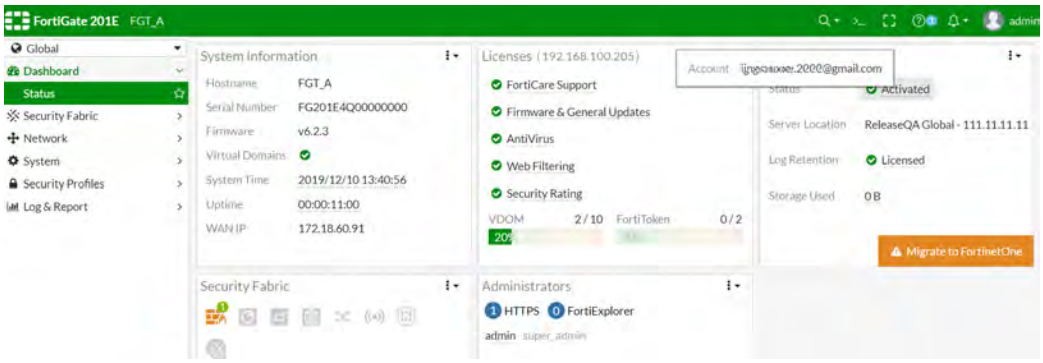


4. Click *OK*.
   The FortiGate Cloud widget now shows the FortiCloud account.

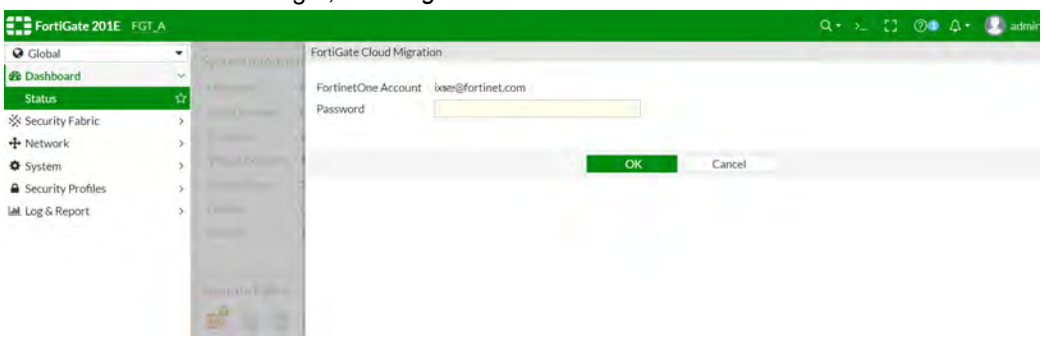


To migrate from the activated FortiGate Cloud account to the registered FortiCloud account:

1. Go to *Dashboard* > *Status*.

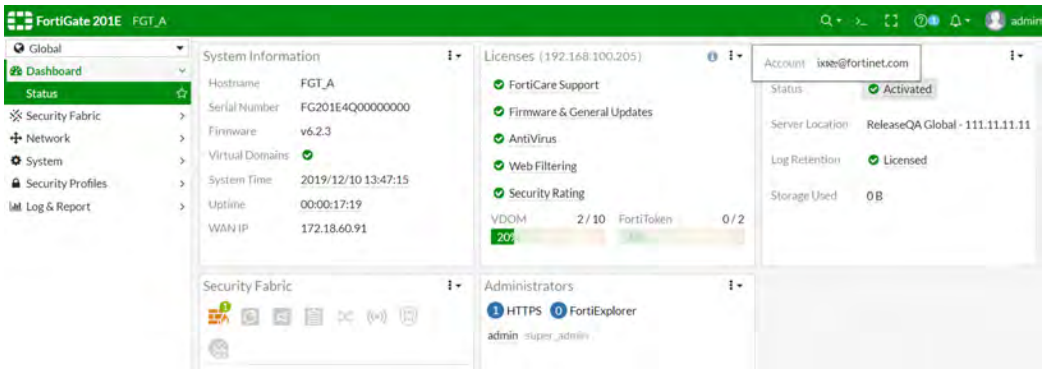


2. In the FortiGate Cloud widget, click *Migrate to FortiCloud*.



Netskope Exhibit 1017

3. Enter the password for the account that was used to register the FortiGate, then click *OK*.
   The FortiGate Cloud widget now shows the FortiCloud account.



To activate FortiGate Cloud using an account that is not used for registration:

1. In the CLI, enter the following command:

   ```
   execute fortiguard-log login <account_id> <password>
   ```

   Where the <account_id> and <password> are the credentials for the account that you are using to activate FortiGate Cloud.
2. Check the account type with following command:

   ```
   # diagnose fdsm contract-controller-update
   Protocol=2.0|Response=202|Firmware=FAZ-4K-FW-2.50-
   100|SerialNumber=FAMS000000000000|Persistent=false|ResponseItem=HomeServer:172.16.95.151
   :443*AlterServer:172.16.95.151:443*Contract:20200408*NextRequest:86400*UploadConfig:Fals
   e*ManagementMode:Local*ManagementID:737941253*AccountType:multitenancy

   Result=Success
   ```



A FortiCloud account that is not used for the support portal account cannot be used to register FortiGate. Attempting to activate FortiGate Cloud with this type of account will fail.

Troubleshooting your installation

If your FortiGate does not function as desired after installation, try the following troubleshooting tips:

1. **Check for equipment issues**
   Verify that all network equipment is powered on and operating as expected. Refer to the QuickStart Guide for information about connecting your FortiGate to the network.
2. **Check the physical network connections**
   Check the cables used for all physical connections to ensure that they are fully connected and do not appear damaged, and make sure that each cable connects to the correct device and the correct Ethernet port on that device.

3. **Verify that you can connect to the internal IP address of the FortiGate**
   Connect to the GUI from the FortiGate's internal interface by browsing to its IP address. From the PC, try to ping the internal interface IP address; for example, `ping 192.168.1.99`. If you cannot connect to the internal interface, verify the IP configuration of the PC. If you can ping the interface but can't connect to the GUI, check the settings for administrative access on that interface. Alternatively, use SSH to connect to the CLI, and then confirm that HTTPS has been enabled for Administrative Access on the interface.
4. **Check the FortiGate interface configurations**
   Check the configuration of the FortiGate interface connected to the internal network (under *Network* > *Interfaces*) and check that *Addressing mode* is set to the correct mode.
5. **Verify the security policy configuration**
   Go to *Policy & Objects* > *Firewall Policy* and verify that the internal interface to Internet-facing interface security policy has been added and is located near the top of the policy list. Check the *Active Sessions* column to ensure that traffic has been processed (if this column does not appear, right-click on the table header and select *Active Sessions*). If you are using NAT mode, check the configuration of the policy to make sure that *NAT* is enabled and that *Use Outgoing Interface Address* is selected.
6. **Verify the static routing configuration**
   Go to *Network* > *Static Routes* and verify that the default route is correct. Go to *Monitor* > *Routing Monitor* and verify that the default route appears in the list as a static route. Along with the default route, you should see two routes shown as *Connected*, one for each connected FortiGate interface.
7. **Verify that you can connect to the Internet-facing interface's IP address**
   Ping the IP address of the Internet-facing interface of your FortiGate. If you cannot connect to the interface, the FortiGate is not allowing sessions from the internal interface to Internet-facing interface. Verify that PING has been enabled for *Administrative Access* on the interface.
8. **Verify that you can connect to the gateway provided by your ISP**
   Ping the default gateway IP address from a PC on the internal network. If you cannot reach the gateway, contact your ISP to verify that you are using the correct gateway.
9. **Verify that you can communicate from the FortiGate to the Internet**
   Access the FortiGate CLI and use the command `execute ping 8.8.8.8`. You can also use the `execute traceroute 8.8.8.8` command to troubleshoot connectivity to the Internet.
10. **Verify the DNS configurations of the FortiGate and the PCs**
    Check for DNS errors by pinging or using traceroute to connect to a domain name; for example: `ping www.fortinet.com`.

    If the name cannot be resolved, the FortiGate or PC cannot connect to a DNS server and you should confirm that the DNS server IP addresses are present and correct.
11. **Confirm that the FortiGate can connect to the FortiGuard network**
    Once the FortiGate is on your network, you should confirm that it can reach the FortiGuard network. First, check the *License Information* widget to make sure that the status of all FortiGuard services matches the services that you have purchased. Go to *System* > *FortiGuard*. Scroll down to *Filtering Services Availability* and select *Check Again*. After a minute, the GUI should indicate a successful connection.Verify that your FortiGate can resolve and reach FortiGuard at `service.fortiguard.net` by pinging the domain name. If you can reach this service, you can then verify the connection to FortiGuard servers by running the command `diagnose debug rating`. This displays a list of FortiGuard IP gateways you can connect to, as well as the following information:
    - `Weight`: Based on the difference in time zone between the FortiGate and this server
    - `RTT`: Return trip time
    - `Flags`: D (IP returned from DNS), I (Contract server contacted), T (being timed), F (failed)
    - `TZ`: Server time zone
    - `Curr Lost`: Current number of consecutive lost packets
    - `Total Lost`: Total number of lost packets

Netskope Exhibit 1017

12. **Consider changing the MAC address of your external interface**
    Some ISPs do not want the MAC address of the device connecting to their network cable to change. If you have added a FortiGate to your network, you may have to change the MAC address of the Internet-facing interface using the following CLI command:
    ```
    config system interface
       edit <interface>
          set macaddr <xx:xx:xx:xx:xx:xx>
       end
    end
    ```
13. **Check the FortiGate bridge table (transparent mode)**
    When a FortiGate is in transparent mode, the unit acts like a bridge sending all incoming traffic out on the other interfaces. The bridge is between interfaces on the FortiGate unit. Each bridge listed is a link between interfaces. Where traffic is flowing between interfaces, you expect to find bridges listed. If you are having connectivity issues and there are no bridges listed, that is a likely cause. Check for the MAC address of the interface or device in question.To list the existing bridge instances on the FortiGate, use the following CLI command:
    ```
    diagnose netlink brctl name host root.b
    show bridge control interface root.b host.
    fdb: size=2048, used=25, num=25, depth=1
    Bridge root.b host table
    port no device devname mac addr ttl attributes
    3 4 wan1 00:09:0f:cb:c2:77 88
    3 4 wan1 00:26:2d:24:b7:d3 0
    3 4 wan1 00:13:72:38:72:21 98
    4 3 internal 00:1a:a0:2f:bc:c6 6
    1 6 dmz 00:09:0f:dc:90:69 0 Local Static
    3 4 wan1 c4:2c:03:0d:3a:38 81
    3 4 wan1 00:09:0f:15:05:46 89
    3 4 wan1 c4:2c:03:1d:1b:10 0
    2 5 wan2 00:09:0f:dc:90:68 0 Local Static
    ```
14. **Use FortiExplorer if you can't connect to the FortiGate over Ethernet**
    If you can't connect to the FortiGate GUI or CLI, you may be able to connect using FortiExplorer. Refer to the QuickStart Guide or see the section on FortiExplorer for more details.
15. **Either reset the FortiGate to factory defaults or contact Fortinet Support for assistance**
    To reset the FortiGate to factory defaults, use the CLI command `execute factoryreset`. When prompted, type `y` to confirm the reset.

If you require further assistance, visit the Fortinet Support website.

Zero touch provisioning

This section contains instructions for configuring zero touch provisioning:

- Zero touch provisioning with FortiDeploy on page 51
- Zero touch provisioning with FortiManager on page 53

Zero touch provisioning with FortiDeploy

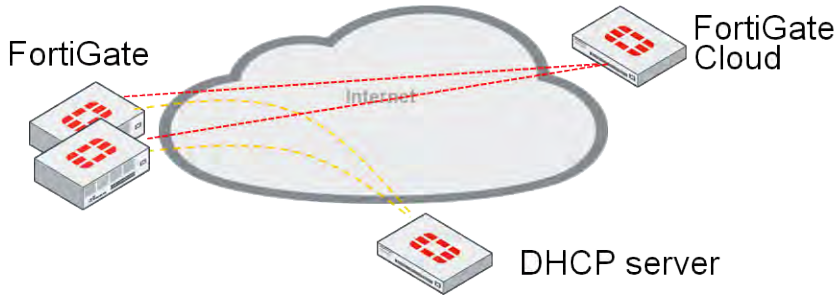
You can use this feature only when the FortiGate boots up from factory reset.

Netskope Exhibit 1017

Topology

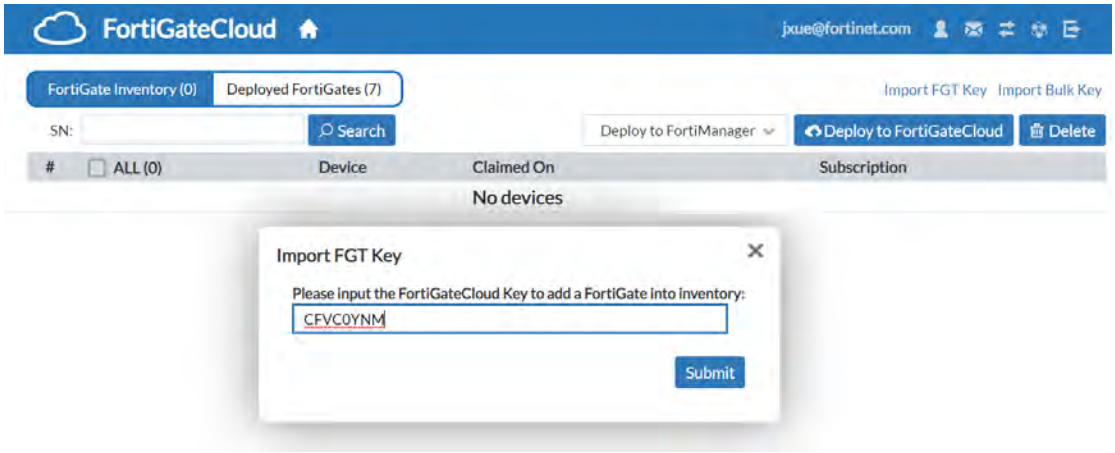
DHCP assignment

Management tunnel to FortiGate Cloud

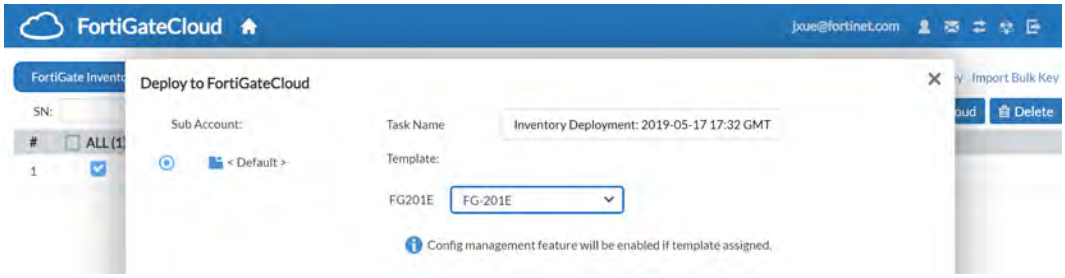


FortiGate zero touch provisioning workflow

1. Add the FortiGate Cloud product key to the FortiGate Cloud portal so that the FortiGate serial number appears in the portal.



2. Set up a configuration template with the basic configuration in the FortiGate Cloud portal.
3. Deploy the FortiGate to FortiGate Cloud with that template.

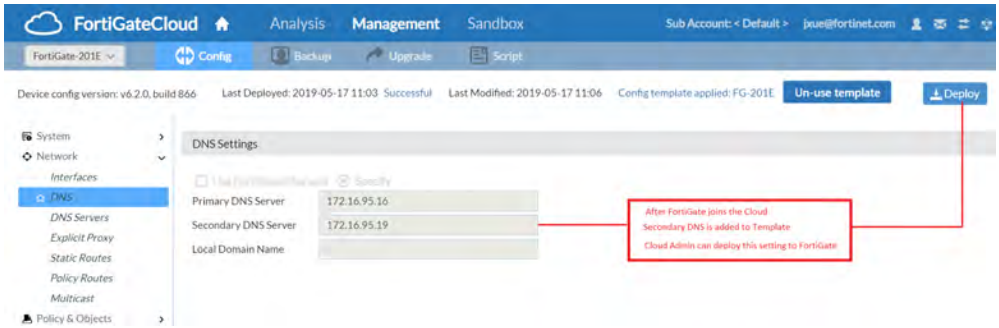


4. Ensure the FortiGate has an interface in default DHCP client mode and is connected to the ISP outlet.
5. Boot the FortiGate in factory reset. The FortiGate gets the DHCP lease so that it can access FortiGate Cloud in the Internet and join FortiGate Cloud.
   The FortiGate Cloud server checks that the FortiGate key is valid and then deploys the FortiGate to FortiGate Cloud.
   To prevent spoofing, FortiGate Cloud invalidates that key after a successful join.

6. Complete zero touch provisioning by obtaining configuration from platform template in the Cloud.

```
0:     set admintimeout 50
0: end
0: config system interface
0:     edit "wan1"
0:         set allowaccess ping ssh fgfm
0:     next
0:     edit "port1"
0:         set allowaccess ping
0:         set ip 1.1.1.1 255.255.255.0
0:     next
0:     edit "port2"
0:         set allowaccess ping
0:         set ip 2.2.2.2 255.255.255.0
0:     next
0: end
```

7. The FortiGate Cloud admin can change the template for different configuration requirements and then deploy the updated template to the FortiGate.
   For example, you can add a secondary DNS to the template and deploy it to FortiGate.



Zero touch provisioning with FortiManager

You can use this feature only when the FortiGate boots up from factory reset. This feature is for FortiGate devices that cannot access the Internet.

A DHCP server includes option 240 and 241 which records FortiManager IP and domain name. FortiGate has an interface with the default DHCP client mode that is connected to the DHCP server in the intranet.

The FortiManager admin can authorize the FortiGate the specific ADOMs and install specific configurations on the FortiGate.

In the whole operation, you do not need to do any manual configuration on the FortiGate except connect to the DHCP server. This is called zero touch deployment.

To prevent spoofing, if a different FortiManager IP comes from the DHCP server later, FortiGate does not change the central management configuration.

Example of configuring DHCP server with option 240

```
config system dhcp server
    edit 2
        set dns-service default
```


Netskope Exhibit 1017

```
        set default-gateway 172.16.200.254
        set netmask 255.255.255.0
        set interface "wan1"
        config ip-range
            edit 2
                set start-ip 172.16.200.201
                set end-ip 172.16.200.209
            next
        end
        set timezone-option default
        config options
            edit 1
                set code 240
                set type ip
                set ip "172.18.60.115"
            next
        end
    next
end
```

FortiGate zero touch provisioning workflow

1. Boot the FortiGate in factory reset.

   ```
   G201E4Q17901047 # diagnose fdsm fmg-auto-discovery-status
   dhcp: fmg-ip=0.0.0.0, fmg-domain-name='', config-touched=0
   ```

   `config-touched=0` means no configuration change from the default.
2. When FortiGate boots in factory reset, it gets the DHCP lease including IP, gateway, DNS, and the FortiManager IP/URL. Central management is automatically configured by using FortiManager IP in option 240.

   ```
   FG201E4Q17901047 # show system central-management
   config system central-management
       set type fortimanager
       set fmg "172.18.60.115"
   end
   ```
3. If FortiGate changes from factory reset, you can see it in central management in `config-touched=1`.

   ```
   FG201E4Q17901047 # diagnose fdsm fmg-auto-discovery-status
   dhcp: fmg-ip=172.18.60.115, fmg-domain-name='', config-touched=1(/bin/dhcpcd)
   ```

Example of a spoofing DHCP server with a fake FortiManager IP

```
config options
    edit 1
        set code 240
        set type ip
        set ip "172.18.60.117"
    end
```

After FortiGate reboots and gets DHCP renew, central management will not use the fake FortiManager IP because `config-touched=1` shows that the FortiGate is not in factory reset.

```
FG201E4Q17901047 # diagnose fdsm fmg-auto-discovery-status
dhcp: fmg-ip=0.0.0.0, fmg-domain-name='', config-touched=1(/bin/dhcpcd)

FG201E4Q17901047 # show system central-management
```


Netskope Exhibit 1017

```
config system central-management
    set type fortimanager
    set fmg "172.18.60.115"
end
```


Netskope Exhibit 1017

Dashboards and widgets

FortiOS includes predefined dashboards so administrators can easily monitor device inventory, security threats, traffic, and network health. You can customize the appearance of a default dashboard to display data pertinent to your security fabric, or combine widgets to create custom dashboards. Many dashboards also allow you to switch views between fabric devices.

Each dashboard contains a set of widgets and monitors that allow you to view drill down data and take actions to prevent threats. Use widgets to perform tasks such as viewing device inventory, creating and deleting DHCP reservations, and disconnecting dial-up users. You can add or remove widgets to a dashboard, or save a widget as a standalone dashboard.

This section contains the following topics:

- Using dashboards on page 56
- Using widgets on page 60
- Monitor dashboards and widgets on page 62
- FortiView on page 73
- Configuration backups on page 108

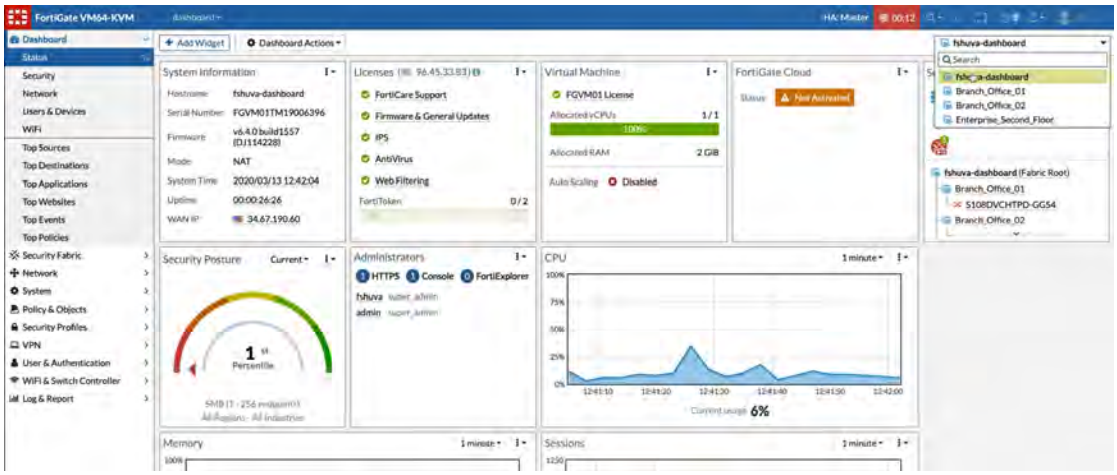
Using dashboards

You can use the dashboard GUI to view fabric devices in the security fabric. You can also combine widgets to create custom dashboards.

Netskope Exhibit 1017

To view downstream fabric devices in the GUI:

1. At the right side of the dashboard, click the device dropdown and select a device.





The device dropdown is available in the *Status*, *Security*, *Network*, *Users & Devices*, and *WiFi* dashboards. You can also enable the dropdown when you create a dashboard.

To create a new dashboard in the GUI:

1. In the banner, click *Dashboard Actions* > *Create New*. The *Add Dashboard* window opens.
2. Configure the dashboard options:
   a. In the *Name* field, enter the dashboard name.
   b. In the *Layout Type* area, select a dashboard layout.

Option	Description
Responsive	The number of columns the dashboard has is determined by screen size. Widgets can only be re-sized horizontally, but the dashboard will look good on all screen sizes.
NOC	The number of columns the dashboard has is explicitly set. Widgets can be re-sized both vertically and horizontally, but the dashboard will look best on the screen size it's originally set up on. This option will work best for a *Network Operation Center* dashboard.
Standalone	Only one widget can be added as full-size.

   c. In the *FortiGate Selection* area, enable the device dropdown.

Option	Description
Dashboard	Enables the device dropdown allowing you to view downstream devices.
Widget	Displays data for the device selected at the time you create the dashboard.

   d. Click *OK*.

You can also use a widget to create a new dashboard. See Using widgets on page 60.

To add a widget to a dashboard in the GUI:

1. In the tree menu, select a dashboard.
2. (Optional) Click the device dropdown, and select a device in the network.
3. In the banner, click *Add Widget*. The *Add Dashboard Widget* window opens.
4. Click the widgets you want to add to the dashboard.
5. Click *Close*.

To delete a dashboard in the GUI:

1. In the tree menu, select a widget.
2. Click *Dashboard Actions* > *Delete*. The *Confirm* window opens.
3. *Click OK.*

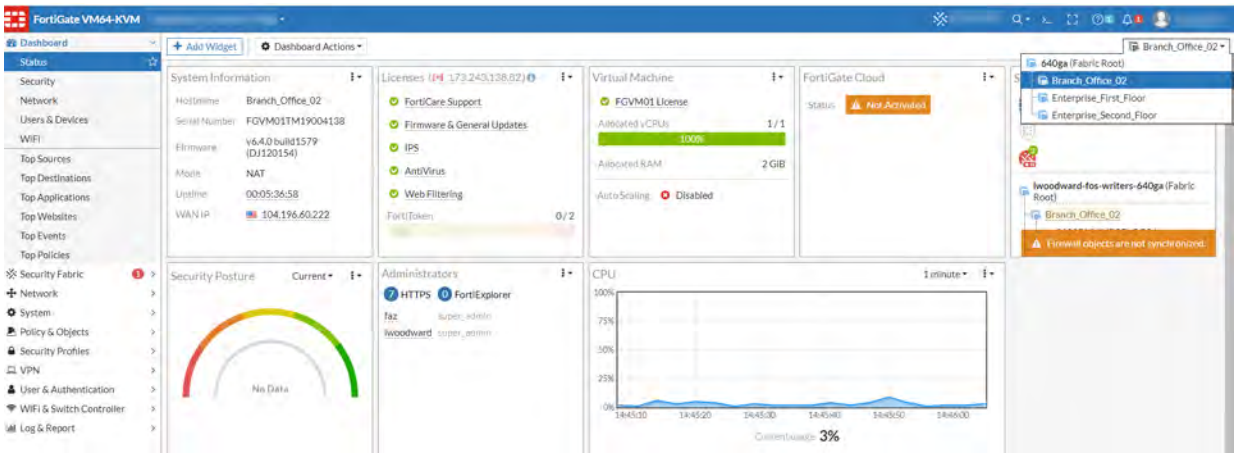
To reset the all the dashboards to the default widgets:

1. In the tree menu, select a widget.
2. Click *Dashboard Actions* > *Reset All*. The *Confirm* window opens.
3. Click *OK*.

Viewing device dashboards in the security fabric

Use the device dropdown in the built-in dashboards to quickly navigate between downstream fabric devices. You can also create dedicated device dashboards devices or log in and configure fabric devices.

To view fabric devices, click the device dropdown at the right side of the page, and select a device from the list.





The device dropdown is available in the *Status*, *Security*, *Network*, *Users & Devices*, and *WiFi* dashboards. You can also enable the dropdown when you create a dashboard.

To create device dashboard in the GUI:

1. In the banner, click *Dashboard Actions* > *Create New*.
2. Configure the dashboard settings.
   a. In the *Name* field, enter a name for the dashboard such as *Branch Office*.
   b. Select the *Layout Type*.

Option	Description
Responsive	The number of columns the dashboard has is determined by screen size. Widgets can only be re-sized horizontally, but the dashboard will look good on all screen sizes.
NOC	The number of columns the dashboard has is explicitly set. Widgets can be re-sized both vertically and horizontally, but the dashboard will look best on the screen size it's originally set up on. This option will work best for a *Network Operation Center* dashboard.
Standalone	Only one widget can be added as full-size.

   c. In the *FortiGate Selection* area, click *Dashboard*.
   d. Click *OK*. The dashboard is added to the tree menu.
3. Open the dashboard in the tree menu.
4. In the banner click, *Add Widget*. The *Add Dashboard* window opens.
5. Click the *Add* icon to add a widget to the dashboard.
6. From the *FortiGate* dropdown, select a fabric device, and click *Add Widget*. The widget is added to the dashboard.

To log into a device from the device dropdown:

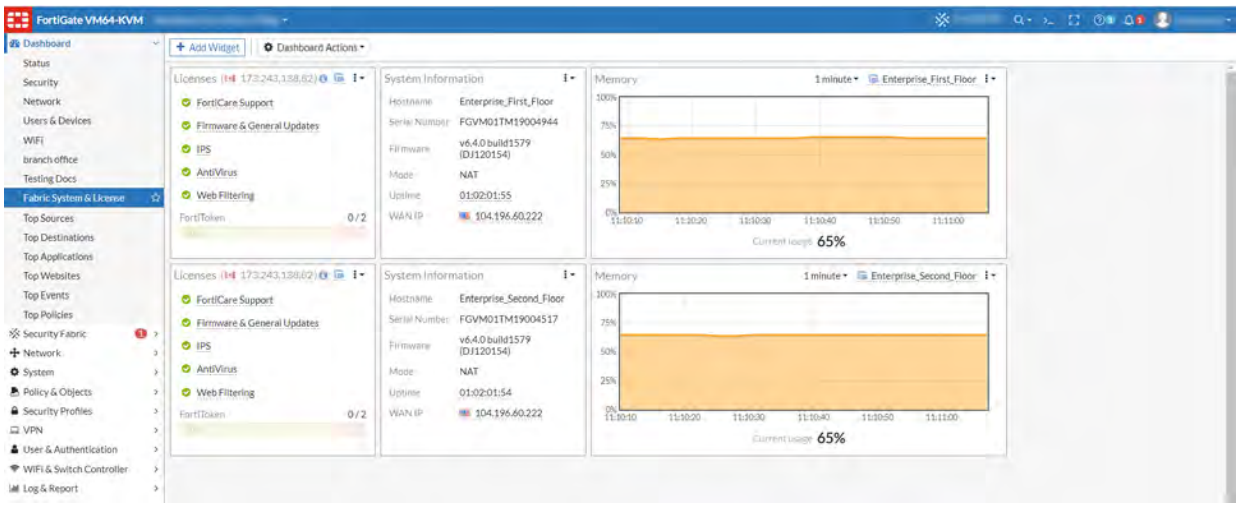
1. Hover over the device in the dropdown, and click *Login* You are redirected to the device login page or System dashboard if you are already logged in.

To configure a device from the device dropdown:

1. Hover over the device in the dropdown, and click *Configure*. The *Configure* page opens.

Creating a fabric system and license dashboard

Create a dashboard summary page to monitor all the fabric devices in a single view. You can use the dashboard to monitor aspects of the devices such as system information, VPN, and routing.



To create a system dashboard in the GUI:

1. In the dashboard banner, click *Dashboard Actions* > *Create New*. The *Add Dashboard* window opens.
2. Configure the dashboard settings.
   a. In the *Name* field, enter a name for the dashboard. such as *Fabric System & License*.
   b. Set the *Layout Type*:

Option	Description
Responsive	The number of columns the dashboard has is determined by screen size. Widgets can only be re-sized horizontally, but the dashboard will look good on all screen sizes.
NOC	The number of columns the dashboard has is explicitly set. Widgets can be re-sized both vertically and horizontally, but the dashboard will look best on the screen size it's originally set up on. This option will work best for a *Network Operation Center* dashboard.
Standalone	Only one widget can be added as full-size.

   c. In the *FortiGate Selection* area, select *Widget*.
   d. Click *OK*. The new dashboard appears in the tree menu.
3. In the banner, click *Add Widget*. The *Add Dashboard Widget* window opens.
4. Click the *Add* icon next to widget. The *Add Dashboard Widget* window opens.
5. From the *FortiGate* dropdown, select a fabric device, and then click *Add Widget*. The widget is added to the dashboard.
   Repeat this step for all the devices you want to view in the dashboard.
6. (Optional) Arrange the widgets in the dashboard by fabric device.

Using widgets

You can save a widget as a standalone dashboard, change the view type, as well as configure tables and filter data.

To save a widget as a dashboard:

1. Hover over a widget in the dashboard, and click *Expand to Full Screen*.
2. In the banner, next to the widget name, click *Save* > *Save as Standalone*. The *Add Dashboard* window opens.
3. Select a *Layout Type*.

Option	Description
Responsive	The number of columns the dashboard has is determined by screen size. Widgets can only be re-sized horizontally, but the dashboard will look good on all screen sizes.
NOC	The number of columns the dashboard has is explicitly set. Widgets can be re-sized both vertically and horizontally, but the dashboard will look best on the screen size it's originally set up on. This option will work best for a *Network Operation Center* dashboard.
Standalone	Only one widget can be added as full-size.

4. Click *OK*.

To save options specific to a widget:

In the widget banner, next to the widget name, click *Save* > *Save View Settings*.



Only radio options or dropdowns specific to the widget can be saved. For example, if you use the *Device Inventory* widget, you can change the view to be based on *Software OS* instead of *Hardware*. and save it.

To view the widget settings:

1. At the top-right of the widget, click the *Settings* icon. The *Edit Dashboard Widget* window opens.
2. Configure the widget settings, and click *OK*.



The settings will vary depending on the widget.

To configure a table in the widget:

1. Hover over the left side of the table header, and click the *Configure Table* icon.
2. Configure the table options.

Option	Description
Best Fit All Columns	Resizes all of the columns in a table to fit their content.
Reset Table	Resets the table to the default view.
Select Columns	Adds or removes columns from the view.

3. Click *Apply*.

Netskope Exhibit 1017

To filter or configure a column in a table:

1. Hover over a column heading, and click the *Filter/Configure Column* icon.
2. Configure the column options, and click *Apply*.
   table

Option	Description
Resize to Contents	Resizes the column to fit the content.
Group by this Column	Groups the table rows by the contents in the selected column.

3. To filter a column, enter a value in the *Filter* field, and click *Apply*.



Filtering is not supported in all the widgets.

Monitor dashboards and widgets

Monitor dashboards and widgets allows you to view various states of your FortiGate pertaining to routing, VPN, DHCP, devices, users, quarantine, and wireless connections.

The following default monitor dashboards are built into FortiOS:

- Network
- Users & Devices
- WiFi

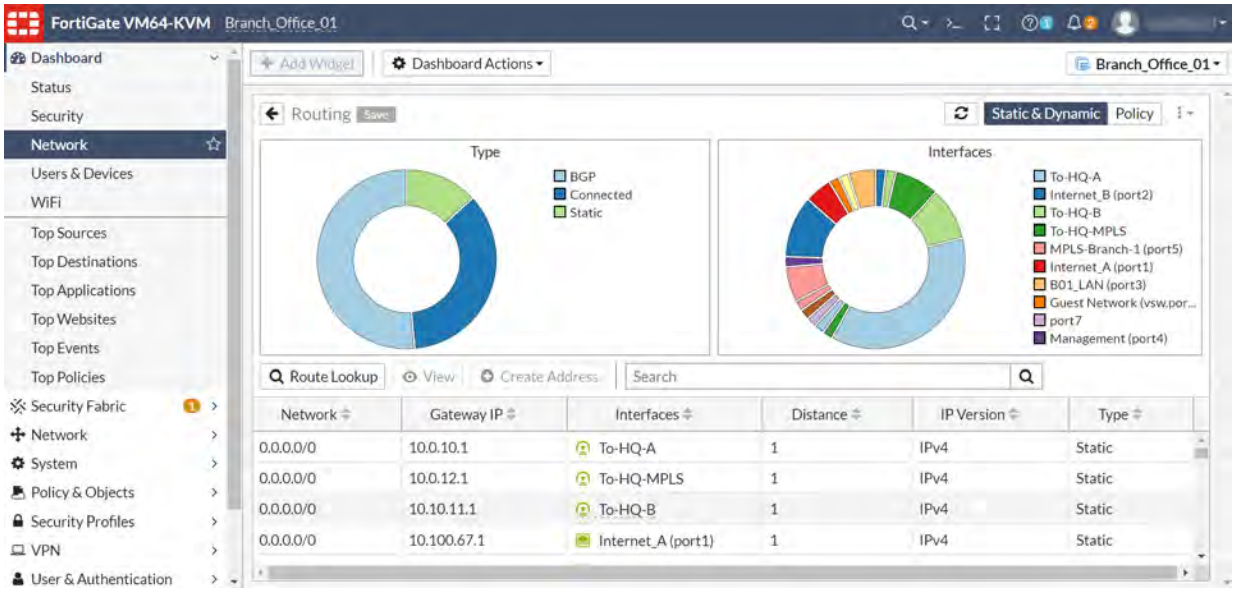
Each built-in dashboard contains multiple widgets which can be expanded for detail view. To save a view as its own dashboard, click *Save* > *Save as Standalone*. For more information, see Using widgets on page 60

To view the widgets available in each dashboard category, click *Add Widget*.

Category	Use these widgets to:
Monitor	• View users and devices connected to the network • Identify threats from individual users and devices, and quarantine them. • View FortiGuard and FortiClient data • Monitor traffic bandwidth over time
Network	• Monitor DHCP clients • Monitor IPsec VPN connections • Monitor current routing table • Monitor SD-WAN status • Monitor SSL-VPN connections
WiFi	• View FortiAP status, channel utilization, and clients • View login failures and signal strength • View the number of WiFi clients

Static & Dynamic Routing Monitor

The *Static & Dynamic Routing Monitor* displays the routing table on the FortiGate including all static and dynamic routing protocols in IPv4 and IPv6. You can also use this monitor to view the firewall policy route.



To view the routing monitor in the GUI:

1. Go to *Dashboard* > *Network*.
2. Hover over the *Static & Dynamic Routing* widget, and click *Expand to Full Screen*. The *Routing* monitor opens.
3. To view the Policy monitor, click the *Policy* tab.
4. To filter the *Interfaces* and *Type* columns:
   a. Hover over the column heading, and click the *Filter/Configure Column* icon.
   b. Click *Group By This Column*, then click *Apply*.
   5. (Optional) Click *Save* > *Save View Settings*, to save your settings.

To look up a route a route in the GUI:

1. Click *Route Lookup*.
2. Enter an IP address in the *Destination* field, then click *Search*. The matching route is highlighted on the *Routing* monitor.

To view the routing table in the CLI:

```
# get route info routing-table all
```

Sample output:

```
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
* - candidate default
```


Netskope Exhibit 1017

```
Routing table for VRF=0
S* 0.0.0.0/0 [1/0] via 10.0.10.1, To-HQ-A
[1/0] via 10.0.12.1, To-HQ-MPLS
[1/0] via 10.10.11.1, To-HQ-B
[1/0] via 10.100.67.1, port1
[1/0] via 10.100.67.9, port2
C 10.0.10.0/24 is directly connected, To-HQ-A
C 10.0.10.2/32 is directly connected, To-HQ-A
C 10.0.11.0/24 is directly connected, To-HQ-B
C 10.0.11.2/32 is directly connected, To-HQ-B
C 10.0.12.0/24 is directly connected, To-HQ-MPLS
C 10.0.12.2/32 is directly connected, To-HQ-MPLS
C 10.1.0.0/24 is directly connected, port3
C 10.1.0.2/32 is directly connected, port3
C 10.1.0.3/32 is directly connected, port3
C 10.1.100.0/24 is directly connected, vsw.port6
```

To look up a firewall route in the CLI:

```
# diag firewall proute list
```

Sample output:

```
list route policy info(vf=root):


id=0x7f450002 vwl_service=2(BusinessCritialCloudApp) vwl_mbr_seq=4 5 3 dscp_tag=0xff 0xff
     flags=0x0 tos=0x00 tos_mask=0x00 protocol=0 sport=0:65535 iif=0 dport=1-65535 oif=3
     (port1) oif=4(port2) oif=18(To-HQ-MPLS)
source(1): 0.0.0.0-255.255.255.255
destination wildcard(1): 0.0.0.0/0.0.0.0
internet service(4): Microsoft.Office.365(4294837472,0,0,0, 33182) Microsoft.Office.Online
     (4294837475,0,0,0, 16177) Salesforce(4294837976,0,0,0, 16920) GoToMeeting
     (4294836966,0,0,0, 16354)
hit_count=0 last_used=2020-03-30 10:50:18

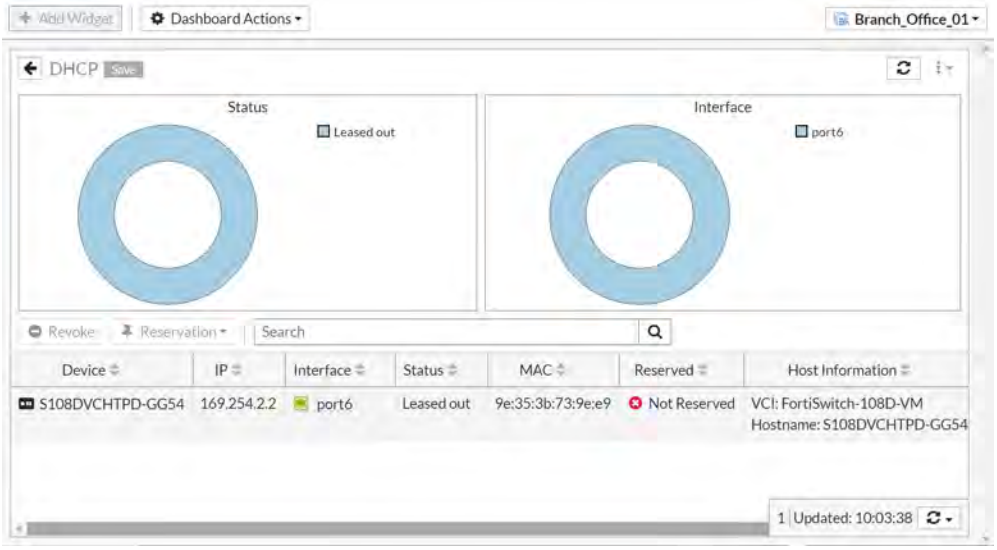
id=0x7f450003 vwl_service=3(NonBusinessCriticalCloudApp) vwl_mbr_seq=4 5 dscp_tag=0xff 0xff
     flags=0x0 tos=0x00 tos_mask=0x00 protocol=0 sport=0:65535 iif=0 dport=1-65535 oif=3
     (port1) oif=4(port2)
source(1): 0.0.0.0-255.255.255.255
destination wildcard(1): 0.0.0.0/0.0.0.0
internet service(2): Facebook(4294836806,0,0,0, 15832) Twitter(4294838278,0,0,0, 16001)
hit_count=0 last_used=2020-03-30 10:50:18

id=0x7f450004 vwl_service=4(Ping-Policy) vwl_mbr_seq=1 2 dscp_tag=0xff 0xff flags=0x0
     tos=0x00 tos_mask=0x00 protocol=1 sport=0:65535 iif=0 dport=1-65535 oif=16(To-HQ-A)
     oif=17(To-HQ-B)
```

DHCP monitor

The DHCP monitor displays all the addresses leased out by FortiGate's DHCP servers. You can use the monitor to revoke an address for a device, or create, edit, and delete address reservations.

Netskope Exhibit 1017



To view the DHCP monitor in the GUI:

1. Go to *Dashboard* > *Network*.
2. Hover over the *DHCP* widget, and click *Expand to Full Screen*.

To filter or configure a column in the table, hover over the column heading, and click *Filter/Configure Column*.

To revoke a lease:

1. Select a device in the table.
2. In the toolbar, click *Revoke*, or right-click the device, and click *Revoke*. The *Confirm* window opens.
3. Click *OK*.

A confirmation window opens only if there is an associated address reservation. If there is no address, the lease will be removed immediately upon clicking *Revoke*.

To create a DHCP reservation:

1. Select a server in the table.
2. In the toolbar, click *Reservation*, or right-click the device, and click *Create DHCP Reservation*. The *Create New DHCP Reservation* window opens.
3. Configure the DHCP settings.
4. Click *OK*.

To view top sources by bytes:

1. Right-click a device in the table, and click *Show in FortiView*. The *Top Sources by Bytes* widget opens.

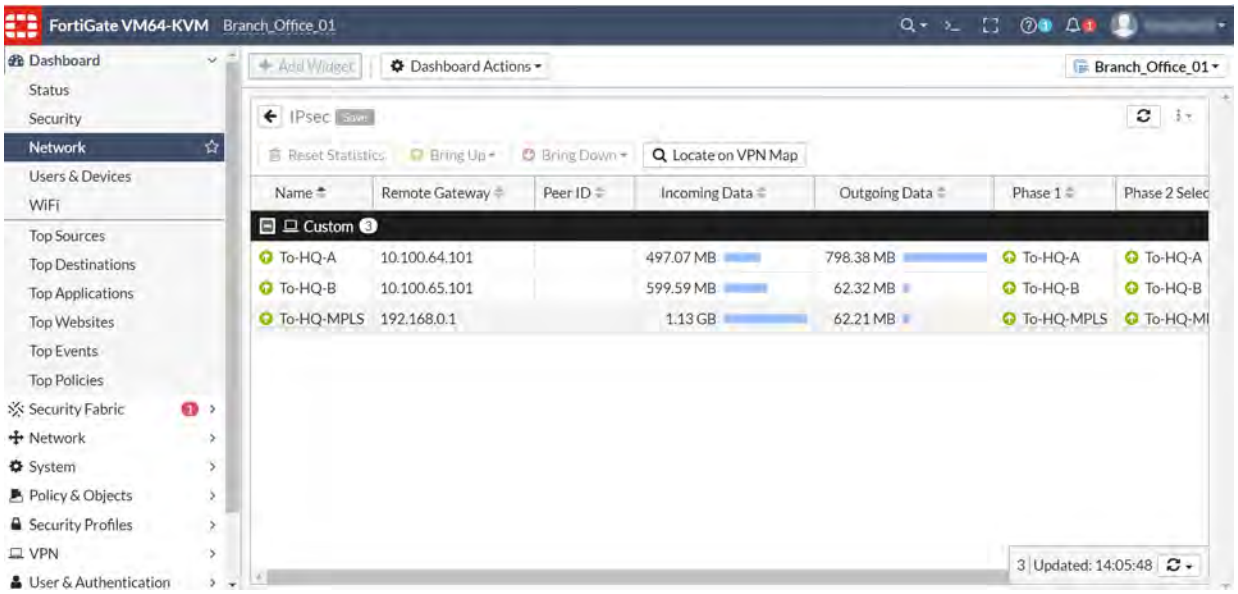
Netskope Exhibit 1017

To view the DHCP lease list in the CLI:

```
# exec dhcp lease-list
```

IPSEC monitor

The IPSEC monitor displays all connected Site to Site VPN and Dial-up VPNs. You can use the monitor to bring a phase 2 tunnel up or down, or disconnect dial-up users.



To view the IPSEC monitor in the GUI:

1. Go to *Dashboard* > *Network*.
2. Hover over the *IPSEC* widget, and click *Expand to Full Screen*.



To filter or configure a column in the table, hover over the column heading, and click *Filter/Configure Column*.

To reset statistics:

1. Select a tunnel in the table.
2. In the toolbar, click *Reset Statistics,* or right-click the tunnel, and click *Reset Statistics*. The *Confirm* window opens.
3. Click *OK*.

To bring a tunnel up:

1. Select a tunnel in the table.
2. Click *Bring Up*, or right-click the tunnel, and click *Bring Up*. The *Confirm* window opens.
3. Click *OK*.

To bring a tunnel down:

1. Select a tunnel in the table.
2. Click *Bring Down*, or right-click the tunnel, and click *Bring Down*. The *Confirm* window opens.
3. Click *OK*.

To locate a tunnel on the VPN Map:

1. Select a tunnel in the table.
2. Click *Locate on VPN Map*, or right-click the tunnel, and click *Locate on VPN Map*. You are taken to *VPN* > *VPN Location Map*.

To view the IPSEC monitor in the CLI:

```
# diag vpn tunnel list
```

Sample output:

```
list all ipsec tunnel in vd 0
------------------------------------------------------
name=fct-dialup ver=1 serial=4 10.100.67.5:0->0.0.0.0:0 dst_mtu=0
bound_if=3 lgwy=static/1 tun=intf/0 mode=dialup/2 encap=none/512 options[0200]=frag-rfc
     accept_traffic=1 overlay_id=0

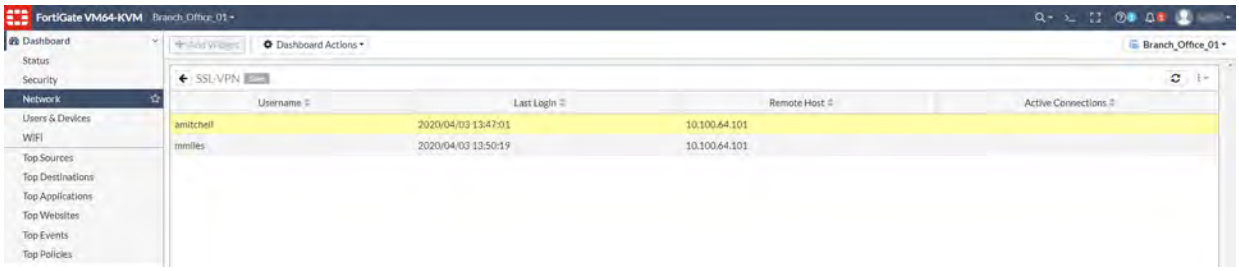
proxyid_num=0 child_num=0 refcnt=12 ilast=5545 olast=5545 ad=/0
stat: rxp=0 txp=0 rxb=0 txb=0
dpd: mode=on-demand on=0 idle=20000ms retry=3 count=0 seqno=0
natt: mode=none draft=0 interval=0 remote_port=0
run_tally=0
------------------------------------------------------
name=To-HQ-MPLS ver=2 serial=3 192.168.0.14:0->192.168.0.1:0 dst_mtu=1500
bound_if=7 lgwy=static/1 tun=intf/0 mode=auto/1 encap=none/528 options[0210]=create_dev
     frag-rfc accept_traffic=1 overlay_id=0

proxyid_num=1 child_num=0 refcnt=22 ilast=0 olast=0 ad=/0
stat: rxp=66693 txp=29183 rxb=33487128 txb=1908427
dpd: mode=on-demand on=1 idle=20000ms retry=3 count=0 seqno=0
natt: mode=none draft=0 interval=0 remote_port=0
proxyid=To-HQ-MPLS proto=0 sa=1 ref=6 serial=1 adr
src: 0:0.0.0.0/0.0.0.0:0
dst: 0:0.0.0.0/0.0.0.0:0
SA: ref=3 options=32203 type=00 soft=0 mtu=1438 expire=266/0B replaywin=2048
seqno=2c5e esn=0 replaywin_lastseq=00002ea3 itn=0 qat=0 hash_search_len=1
life: type=01 bytes=0/0 timeout=1773/1800
dec: spi=700c9198 esp=aes key=16 ebd04605de6148c8a92ced48b30930fa
ah=sha1 key=20 5f0201f67d7c714a046025a1df41d40376437f6a
enc: spi=5aaccc20 esp=aes key=16 13d5d4b46e5e9c42eef509f2d9879188
ah=sha1 key=20 2dde67ef7a2a78b622d9a7ec6d75ad3c55d241e1
dec:pkts/bytes=11938/5226964, enc:pkts/bytes=11357/1312184
```

SSL VPN monitor

The SSL VPN monitor displays user logins and active connections. You can use the monitor to disconnect a specific connection.

Netskope Exhibit 1017



To view the SSL-VPN monitor in the GUI:

1. Go *Dashboard* > *Network*.
2. Hover over the *SSL-VPN* widget, and click *Expand to Full Screen*.

To filter or configure a column in the table, hover over the column heading, and click *Filter/Configure Column*.

To disconnect a user:

1. Select a user in the table.
2. In the table, right-click the user, and click *End Session*. The Confirm window opens.
3. Click *OK*.

To monitor SSL-VPN users in the CLI:

```
# get vpn ssl monitor
```

Sample output

```
SSL VPN Login Users:
Index User Group Auth Type Timeout From HTTP in/out HTTPS in/out
0 amitchell TAC 1(1) 296 10.100.64.101 3838502/11077721 0/0
1 mmiles Dev 1(1) 292 10.100.64.101 4302506/11167442 0/0

SSL VPN sessions:
Index User Group Source IP Duration I/O Bytes Tunnel/Dest IP
```

Firewall Users Monitor

The Firewall Users monitor displays all firewall users currently logged in. You can use the monitor to diagnose user-related logons, or to highlight and deauthenticate a user.

To view the firewall monitor in the CLI:

1. Go to *Dashboard* > *Users & Devices*.
2. Hover over the *Firewall Users* widget, and click *Expand to Full Screen*.
3. To show FSSO Logons, click *Show all FSSO Logons*.



To filter or configure a column in the table, hover over the column heading, and click *Filter/Configure Column*.

To deauthenticate a user:

1. (Optional) Use the *Search* field to search for a specific user.
2. In the toolbar, click *Deauthenticate*, or right-click the user, and click *Deauthnticate*. The *Confirm* window opens.
3. Click *OK*.

To view firewall users in the CLI:

```
# diag firewall auth list
```

Device inventory

You can enable device detection to allow FortiOS to monitor your networks and gather information about devices operating on those networks, including:

- MAC address
- IP address
- Operating system
- Hostname
- Username
- When FortiOS detected the device and on which interface

You can enable device detection separately on each interface in *Network* > *Interfaces*.

Device detection is intended for devices directly connected to your LAN ports. If enabled on a WAN port, device detection may be unable to determine the OS on some devices. You can enable active scanning on the interface to find hosts whose device types FortiOS cannot determine passively.

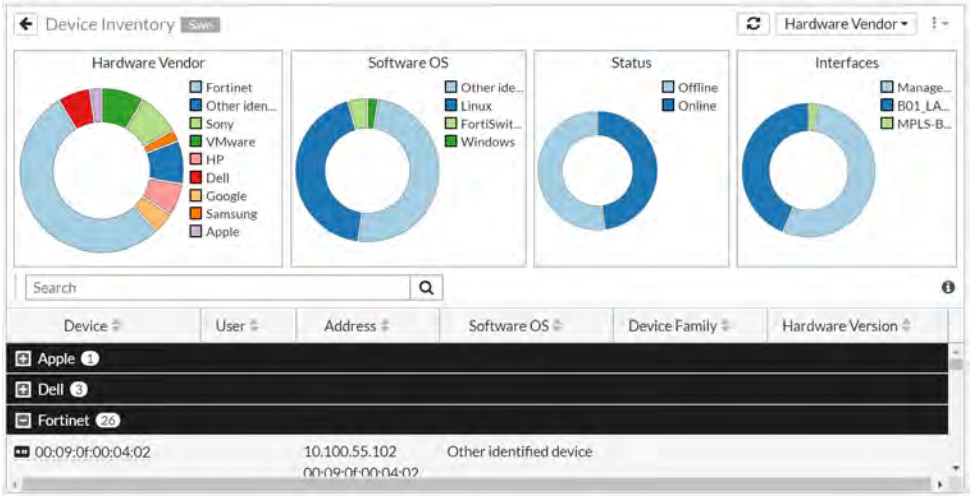
You can also manually add devices to Device Inventory to ensure that a device with multiple interfaces displays as a single device.

To view the device inventory monitor in the GUI:

1. Go to *Dashboard* > *Users & Devices*.
2. Hover over the *Device Inventory* widget, and click *Expand to Full Screen*. The *Device Inventory* pane appears.



To filter or configure a column in the table, hover over the column heading, and click *Filter/Configure Column*. See Device inventory and filtering on page 70.



Device inventory and filtering

The *Device Inventory* widget contains a series of summary charts that provide an overview of the hardware, operating system, status, and interfaces. You can use these clickable charts to simplify filtering among your devices.

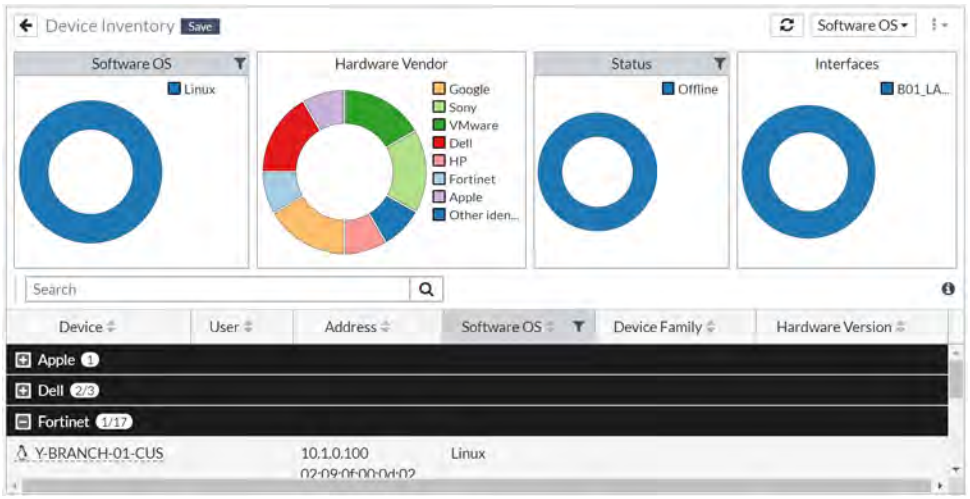
To view the device inventory and apply a filter:

1. Go to *Dashboard* > *Users & Devices*.
2. Hover over the *Device Inventory* widget, and click *Expand to Full Screen*. The *Device Inventory* pane appears.
3. To filter the order of the charts by operating system, click the dropdown at the right side of the page, and select *Software OS*.
4. To filter a chart, click an item in the legend or chart area. The table displays the filter results.
5. To combine filters, hover over a column heading, and click *Filter/Configure Column*.
6. Click the filter icon in the top-right corner of the chart to remove the filter.
7. (Optional) In the banner next to the widget name, click *Save* > *Save View Settings*.

Filter examples

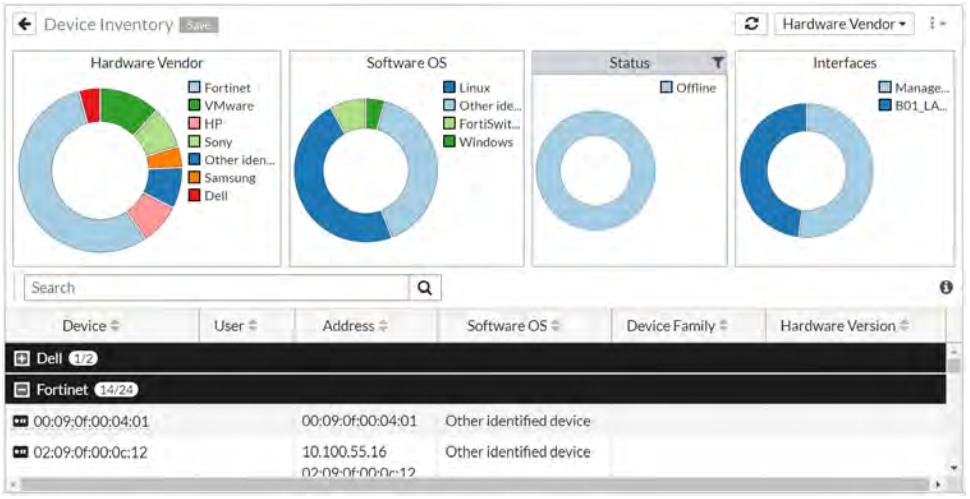
To filter all Linux devices:

1. In the *Software OS* chart, click *Linux* .



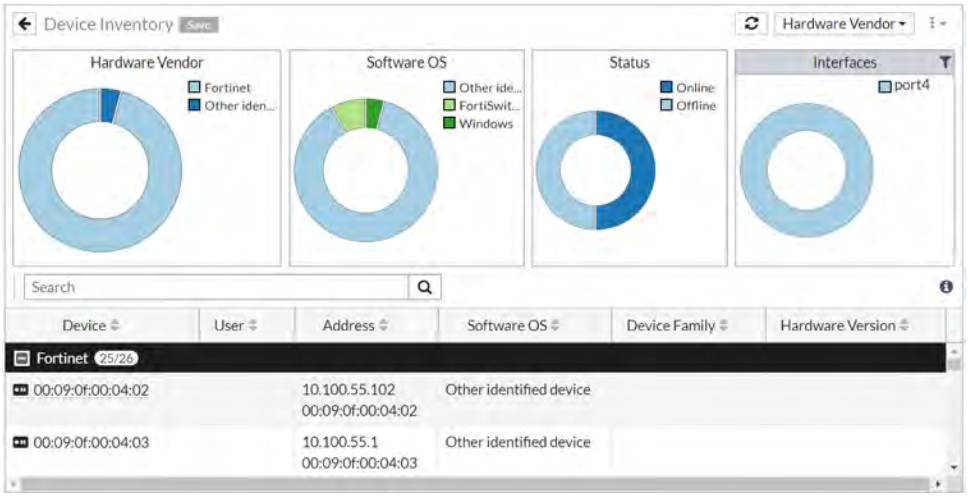
To filter all offline devices:

1. In the *Status* chart, click *Offline*.



To filter all devices discovered on port4:

1. In the *Interfaces* chart, click *port4*.

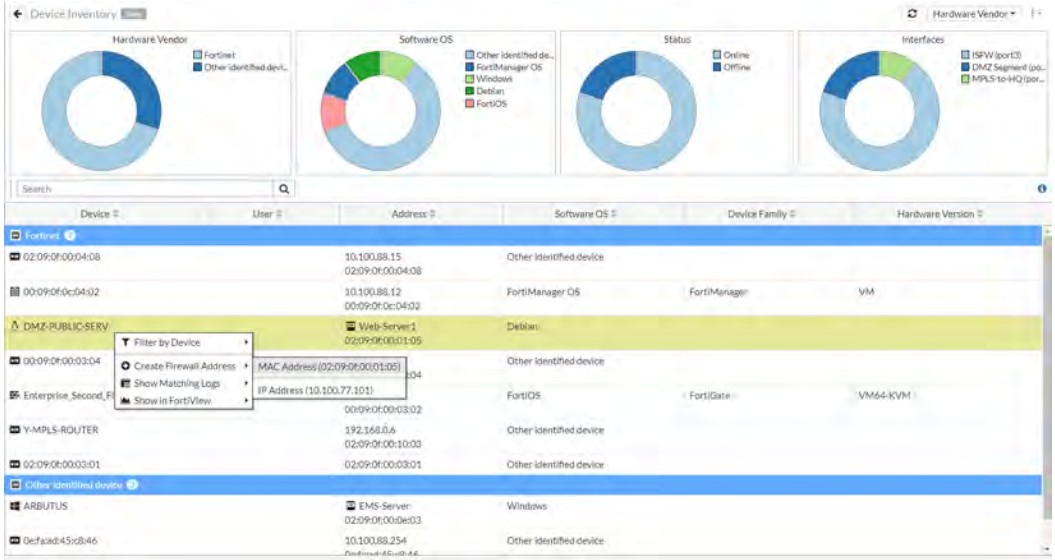


Adding MAC-based addresses to devices

Assets detected by device detection appear in the *Device Inventory* widget. You can manage policies around devices by adding a new device object (MAC-based address) to a device. Once you add the MAC-based address, the device can be used in address groups or directly in policies.

To add a MAC-based address to a device:

1. Go to *Dashboard* > *Users & Devices*.
2. Hover over the *Device Inventory* widget, and click *Expand to Full Screen*. The *Device Inventory* pane appears.
3. Right-click a device, and select *Create Firewall Address* > *MAC Address*.

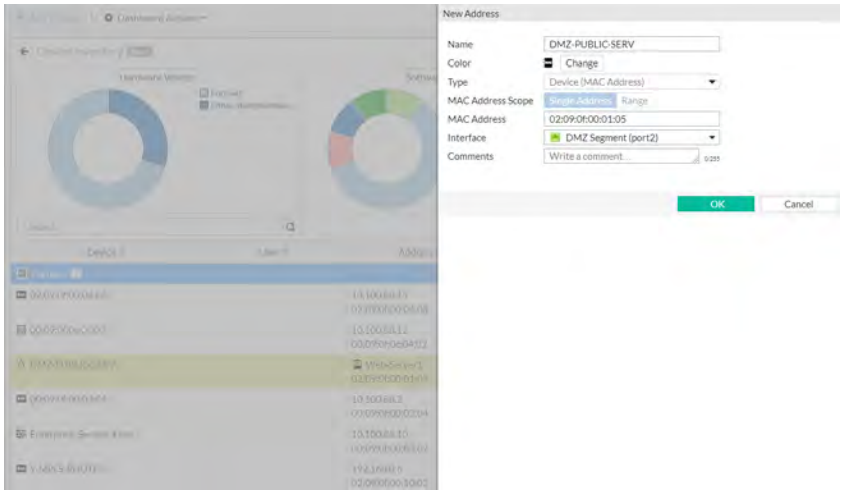


4. (Optional) In the *Name* field, give the device a descriptive name.

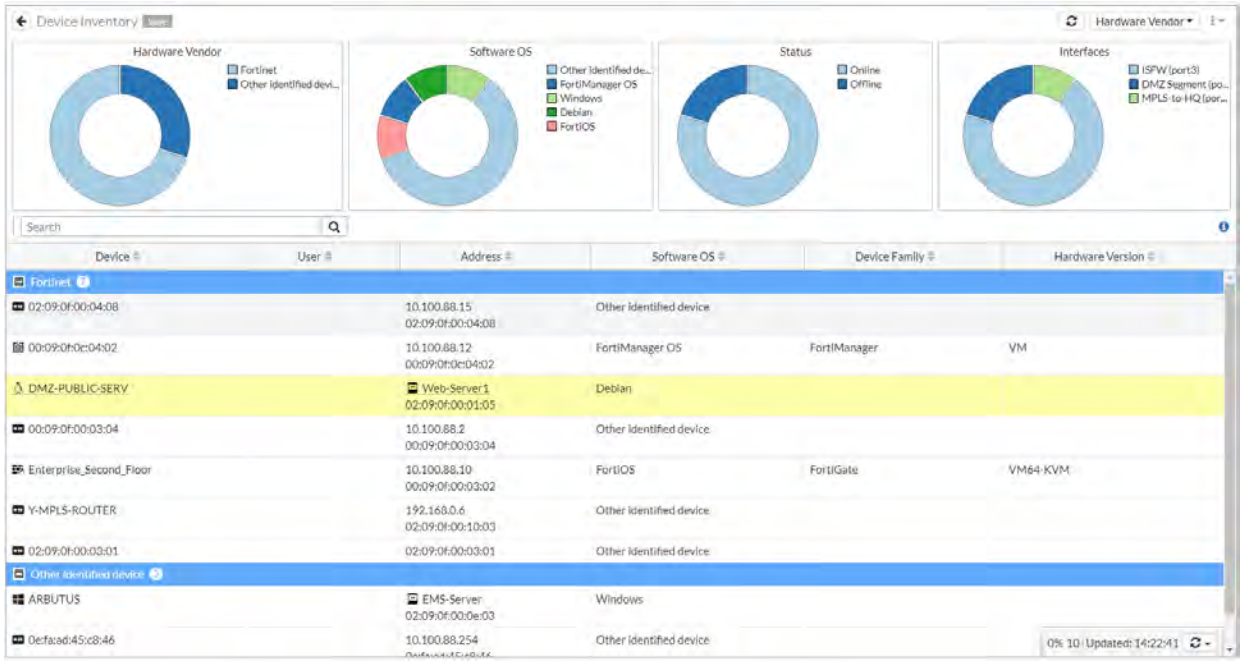


Use the *Name* field to assign a descriptive name to a device so it is easier to find in the *Device* column. After you finish configuring the device, refresh the page to see the new name in the dashboard.

5. In the *New Address* pane, enter an address name.



6. Click *OK*. The MAC address icon is now displayed in the *Address* column for the device.



FortiView

FortiView is the FortiOS log view tool and comprehensive monitoring system for your network. FortiView integrates real-time and historical data into a single view on your FortiGate. It can log and monitor network threats, keep track of

Netskope Exhibit 1017

administration activities, and more.

Use FortiView dashboards and widgets to investigate traffic activity such as user uploads and downloads, or videos watched on YouTube. You can view the traffic on the whole network, by user group, or by individual. FortiView displays the information in both text and visual format, giving you an overall picture of your network traffic activity so that you can quickly decide on actionable items.

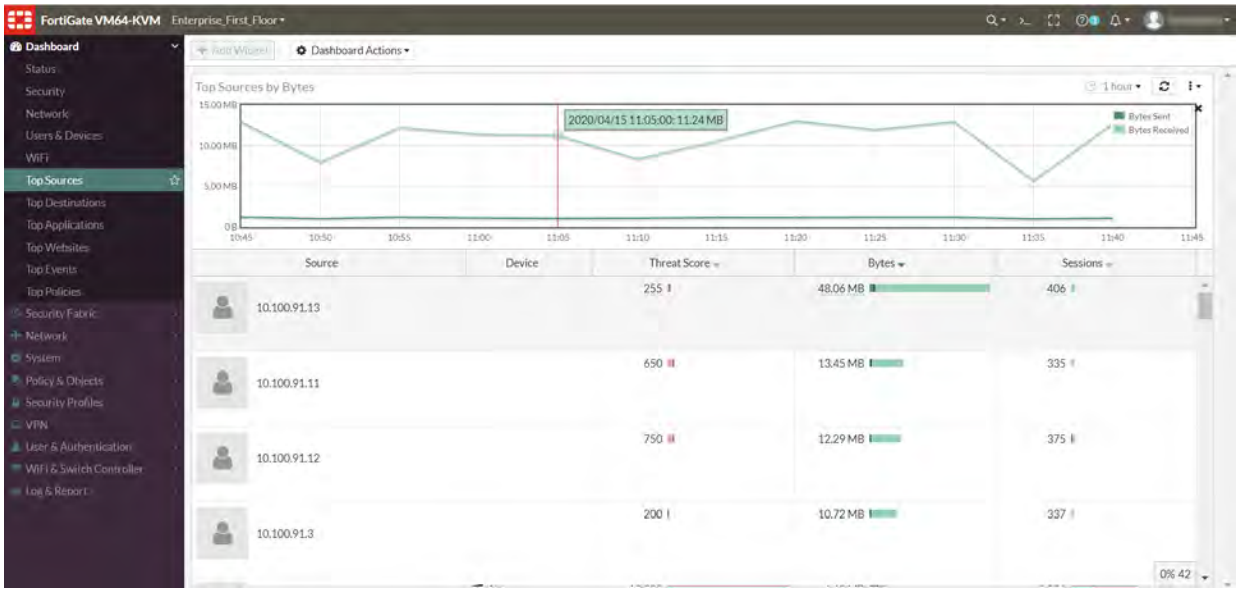
FortiView is integrated with many UTM functions and each release adds more features. For example, you can quarantine an IP address directly in FortiView or create custom devices and addresses from a FortiView entry.



The logging range and depth will depend on the FortiGate model.

FortiView dashboards and widgets

FortiView dashboards and widgets are available in the tree menu under the *Dashboards* module. The module contains several core dashboards for the top categories. Non-core FortiView pages are available as widgets that can be added to the dashboards. You can also use non-core pages to create standalone dashboards.



Core FortiView dashboards

The following core dashboards are available in the tree menu under the *Dashboard* console:

Dashboard	Usage
Top Sources	Displays Top Sources by traffic volume and drilldown by Source.
Top Destinations	Displays Top Destinations by traffic volume and drilldown by Destination.

Dashboard	Usage
Top Applications	Displays Top Applications by traffic volume and drilldown by Application.
Top Websites	Displays Top Websites by session count and drilldown by Domain.
Top Events	Displays Top Events by Level and drilldown by Event Name.
Top Policies	Displays Top Policies by traffic volume and drilldown by Policy number

Usage is based on default settings. The pages may be customized further and sorted by other fields.



With the exception of *Top Events*, you can quarantine a host and ban an IP from all of the core FortiView dashboards.

FortiView widgets

FortiView widgets allow you to create custom dashboards to monitor vulnerabilities, scan summaries, and top items from selected FortiView categories. You can also customize widgets to show information that is most important to you, such as the time range, source logging device, and other information. For information, see Using widgets on page 60

The following FortiView widgets are available in the *Add Dashboard Widget* menu:

Widget	Usage
FortiClient Detected Vulnerabilities	Displays a summary of vulnerabilities detected by FortiClient. FortiClient must be enabled.
FortiView Top N	Top items from the selected FortiView category. The categories include, *LAN/DMZ*, *WAN*, and *All Segments*. For information, seeAdding top FortiView widgets by category on page 75
Host Scan Summary	Displays a summary of hosts scanned. You can use this widget to view the FortiClient monitor and the device inventory.

Adding top FortiView widgets by category

Non-core FortiView pages are available in the *Add Dashboard* window. The *FortiView Top N* widget contains the top pages by category. You can add a top category widget to a dashboard, or use the widget to create a standalone dashboard.



You cannot add widgets to core FortiView dashboards.

To add a top category widget in the GUI:

1. In the tree menu, select a dashboard, and click *Add Widget*.
2. Under the *FortiView* section, select *FortiView Top N*. The *Add Dashboard Widget* pane opens.
3. In the *Title* field, enter a name for the widget.

4. From the *FortiView* dropdown, select the widget. There are three categories to choose from: *LAN/DMZ*, *WAN*, *All Segments*.
5. From the *Time Period* dropdown, select the time period. This option is not available in all widgets.
6. In the *Visualization* area, select *Table View* or *Bubble Chart*.
7. From the *Sort By* dropdown, select the sorting method.
8. Click *Add Widget*.

Widgets by category

Usage is based on the default settings. The widgets may be customized further and sorted by other fields.

LAN/DMZ

Widget	Sort by	Usage
Applications	Bytes/Sessions/Bandwidth/Packets	Displays top applications and drilldown by application.
Cloud Applications	Bytes/Sessions/Files(Up/Down)	Displays top cloud applications and drilldown by application .
Cloud Users	Bytes/Sessions/Files(Up/Down)	Displays top cloud users and drilldown by cloud user.
Compromised Hosts	Verdict	Displays compromised hosts and drilldown by source.
Countries/Regions	Bytes/Sessions/Bandwidth/Packets	Displays top countries/regions and drilldown by countries/regions.
Destination Firewall Objects	Bytes/Sessions/Bandwidth/Packets	Displays top destination firewall objects and drilldown by destination objects.
Destination Owners	Bytes/Sessions/Bandwidth/Packets	Displays top destination owners and drilldown by destination.
Destinations	Bytes/Sessions/Bandwidth/Packets	Displays top destinations and drilldown by destination.
Search Phrases	Count	Displays top search phrases and drilldown by search phrase.
Source Firewall Objects	Bytes/Sessions/Bandwidth/Packets	Displays top search phrases and drilldown by source object.
Sources	Bytes/Sessions/Bandwidth/Packets	Displays top sources and drilldown by source.
Threats	Threat level/Threat Score/Sessions	Displays top threats and drilldown by threat.
Traffic Shaping	Dropped Bytes/Bytes/Sessions/Bandwidth/Packets	Displays top traffic shaping and drilldown by shaper.
Web Categories	Bytes/Sessions/Bandwidth/Packets	Displays top web categories and drilldown by category.

Netskope Exhibit 1017

Widget	Sort by	Usage
Web Sites	Bytes/Sessions/Bandwidth/Packets	Displays top web sites and drilldown by domain.
WiFi Clients	Bytes/Sessions	Displays top WiFi clients and drilldown by source.

WAN

Widget	Sort by	Usage
Servers	Bytes/Sessions/Bandwidth/Packets	Displays top servers and drilldown by server address.
Sources	Bytes/Sessions/Bandwidth/Packets	Displays top sources and drilldown by device.
Threats	Threat Level/Threat Score/Sessions	Displays top threats and drilldown by threat.

All Segments

Widget	Sort by	Usage
Admin Logins	Configuraton Changes/Logins/Failed Logins	Displays top admin logins by username.
Destination Interfaces	Bytes/Sessions/Bandwidth/Packets	Displays top destination interfaces by destination interface.
Endpoint Vulnerabilities	Severity	Displays top endpoint vulnerabilities by vulnerability name.
Failed Authentication	Failed Attempts	Displays top failed authentications by failed authentication source.
FortiSandbox Files	Submitted	Displays top FortiSandbox files by file name.
Interface Pairs	Bytes/Sessions/Bandwidth/Packets	Displays top interface pairs by source interface.
Policies	Bytes/Sessions/Bandwidth/Packets	Displays top policies by policy.
Source Interfaces	Bytes/Sessions/Bandwidth/Packets	Displays top source interfaces by source interface.
System Events	Level/Events	Displays top system events by event name.
VPN	Connections/Bytes	Displays top VPN connections by user.
Vulnerable Endpoint Devices	Detected Vulnerabilities	Displays top vulnerable endpoint devices by device.



A maximum of 25 interfaces can be monitored at one time on a device.

Netskope Exhibit 1017

VDOMs and dashboards

Dashboards are created per VDOM when VDOM mode is enabled. Some features and widgets are not available depending on Multi or Split-task VDOM mode.

Multi-VDOM mode

The following widgets and dashboard setting are not available Multi VDOM mode because it does not support Security Fabric:

- Security Fabric related widgets
- *FortiGate Selection* option

Split-task mode

Split-task VDOM mode is limited to two VDOMs, the root VDOM and the FortiGate traffic VDOM. The root VDOM is for all management related settings and the FortiGate traffic VDOM is for all traffic related settings.

The *FortiGate Selection* option is available when you create a dashboard in Split-Task VDOM mode.

For information about VDOM modes, see Virtual Domains on page 610.

Examples

When VDOM mode is disabled, notice the *FortiGate Selection* option is available in the *Add Dashboard* window:

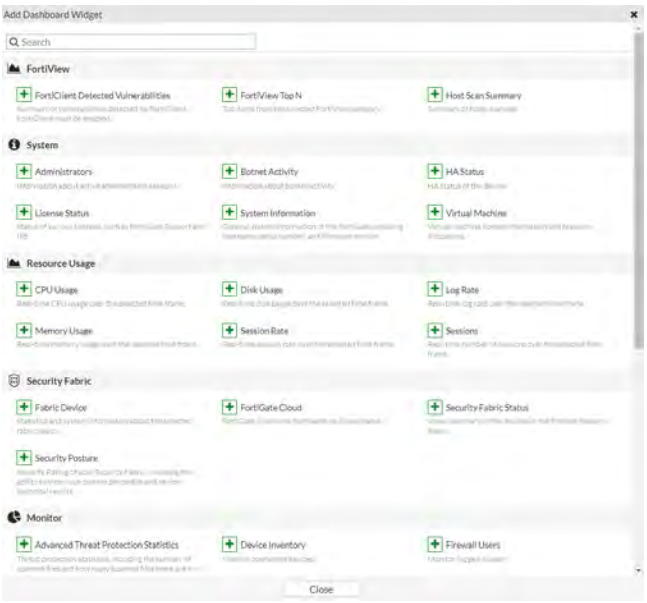


When Multi VDOM mode is enabled, notice the *FortiGate Selection* is not available in the *Add Dashboard* window:

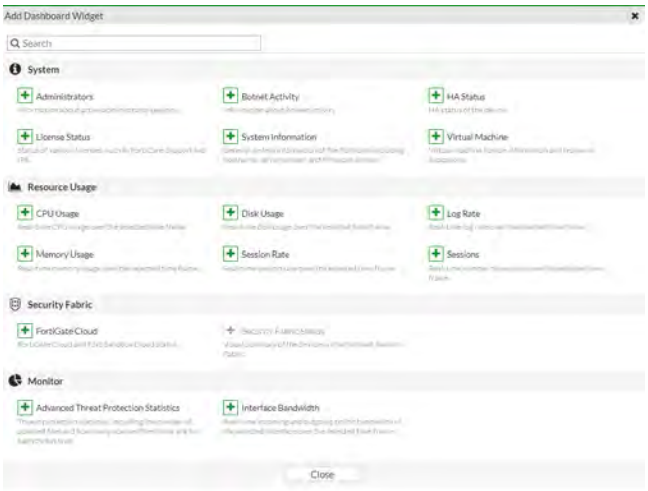


When Multi VDOM mode is disabled, notice all the widgets in the *Add Dashboard Widget* menu are enabled:

Netskope Exhibit 1017



When Multi VDOM mode is enabled, notice the *Security Fabric Status* widget is disabled:



FortiView interface

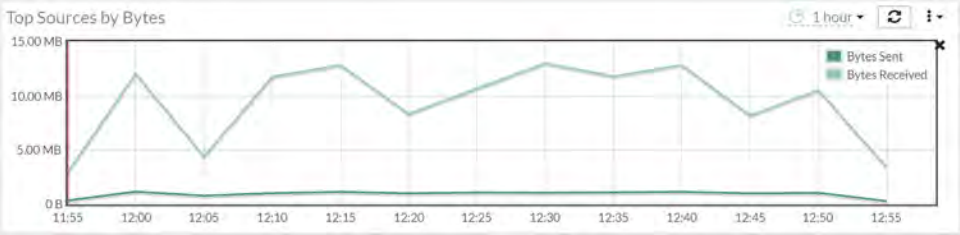
Use the FortiView interface to customize the view and visualizations within a dashboard to find the information you are looking for. The tools in the top menu bar allow you to change the time display, refresh the data, and customize the data source. You can also right-click a table in the dashboard to view drilldown information for an entry.

FortiOS 6.4.0 does not support search and filtering of results at this time.

Top menu bar

The top menu bar contains the following settings:

- A *time display* dropdown to switch between current and historical data.
- A *Refresh* button to update the displayed data.
- A *Settings* dropdown to change the information shown on the dashboard.

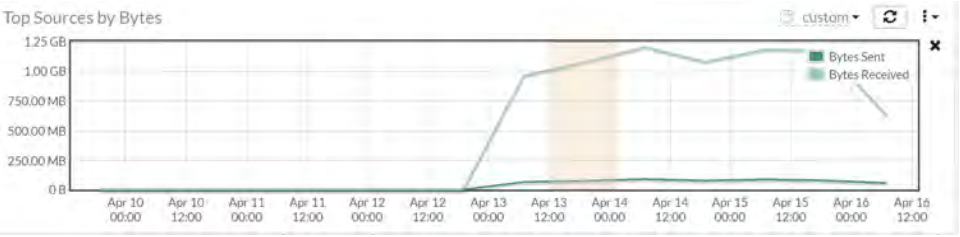


Time period display

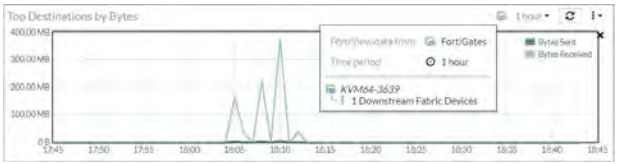
Use the *time display* dropdown to select the time period to display on the current dashboard. Time display options vary depending on the dashboard and can include current information (*now*) and historical information (*1 hour*, *24 hours*, and *7 days*).

Disk logging or remote logging must be enabled to view historical information.

You can use a chart to create a custom time display by selecting the time range with your cursor.



The icon next to the time period identifies the data source (FortiGate Disk, FortiAnalyzer, or FortiGate Cloud). You can hover over the icon to see a description of the device.

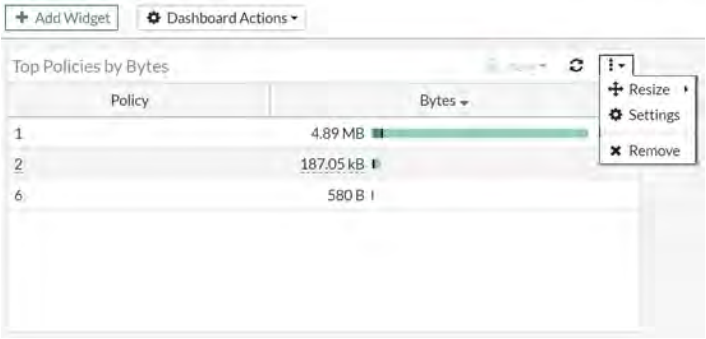


View settings

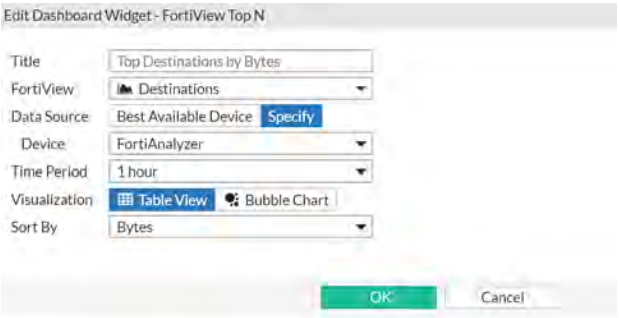
Use the *Settings* menu to change the data source, sort by information, and visualization.

To change the widget settings:

1. Click the dropdown menu at the right side of the top menu bar, and select *Settings*.



2. Configure the widget settings, and click *OK*.



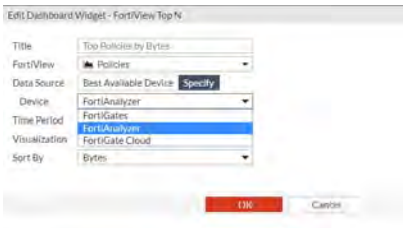
The *Data Source* dropdown only appears when FortiGate is connected to another data source.

For information about widget settings, see Adding top FortiView widgets by category on page 75

For dashboards with multiple widgets, you cannot access the settings dropdown when the widget is expanded to full screen. To change the settings, click the back button to return to the dashboard, and click the dropdown.

Data source

FortiView gathers information from a variety of data sources. If there are no log disk or remote logging configured, the data will be drawn from the FortiGate's session table, and the Time Period is set to *Now*.



Other data sources that can be configured are:

- FortiGates (disk)
- FortiAnalyzer
- FortiGate Cloud

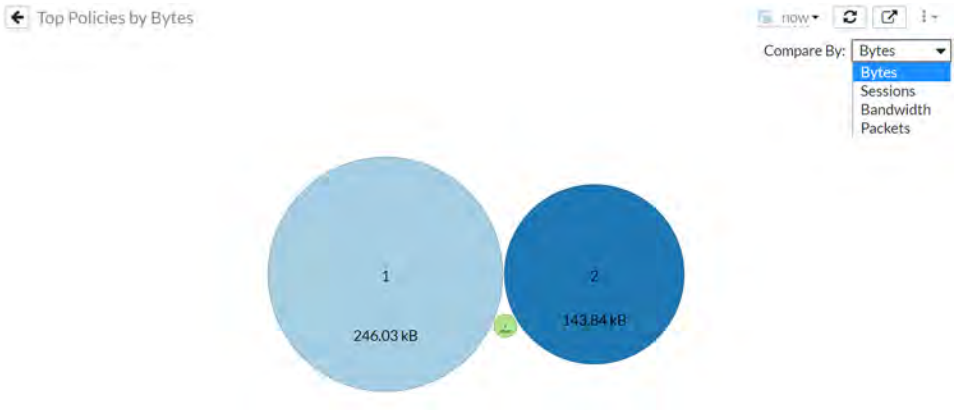
When *Data Source* is set to *Best Available Device*, FortiAnalyzer is selected when available, then FortiGate Cloud, and then FortiGate Disk.

Display types

Bubble charts

Display types include table view, bubble charts, and country maps. Not all display types are supported by all dashboards.

Bubble charts allow you to sort information using the *Compare By* dropdown menu. The size of each bubble represents the related amount of data. You can place your cursor over a bubble to display a tool-tip with detailed information on that item, and click on a bubble to drilldown into greater detail.



Country maps

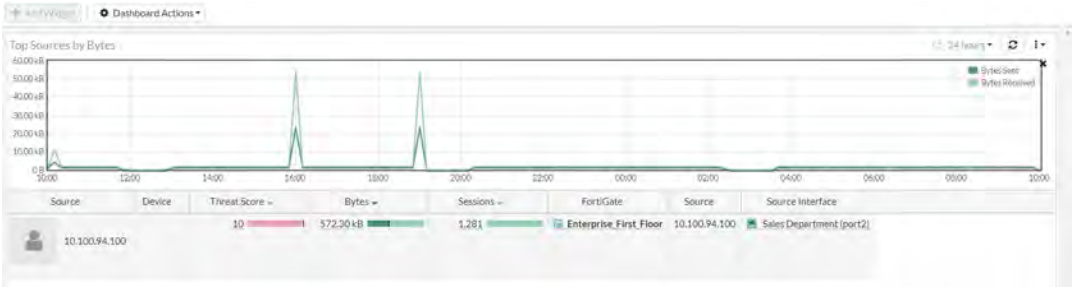
Country maps display traffic activity as regions on a map. Hover over the highlighted region to view information about the entry. You can also compare data by *Bytes*, *Sessions*, *Bandwidth*, and *Packets*. Country maps are not available in all dashboards and widgets.



Table view

Table view displays traffic activity as a graph and a table. To remove the table, click close, at the top right corner of the graph. To view the graph, click *Show Graph*.

Source view

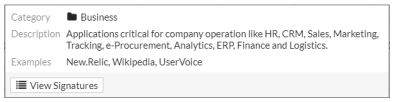


Time	• *Now* entries are determined by the FortiGate's system session list. • Historical or *1 hour* or later entries are determined by traffic logs, with additional information coming from UTM logs. The dropdown only shows *now* if there is no disk.
Graph	• The graph shows the bytes sent/received in the time frame. • Users can customize the time frame by selecting a time period within the graph.
Columns	• *Source* shows the IP address (and user as well as user avatar if configured) of the source device. • *Device* shows the device information as listed in the *Device Inventory* widget. Device detection should be enabled on the applicable interfaces for best function. For

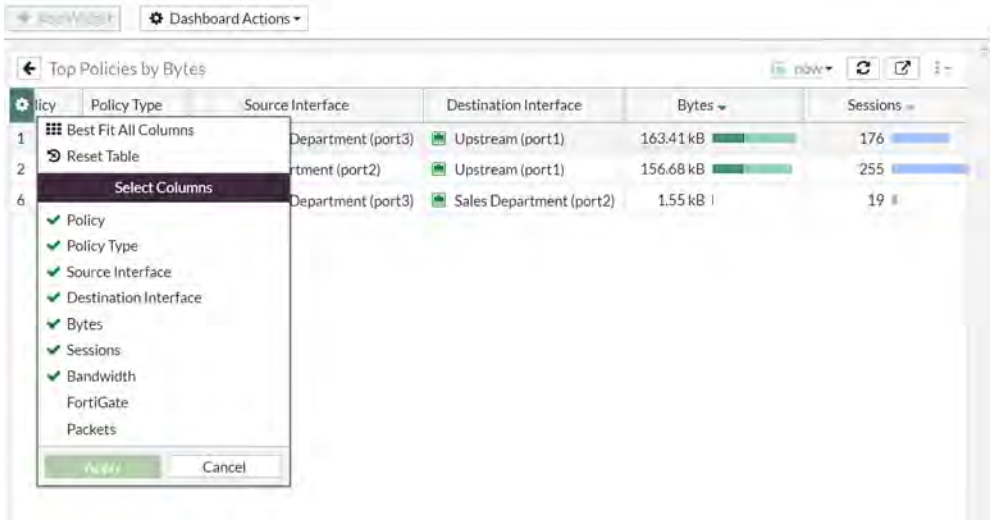
information about adding widgets, see Using widgets on page 60.

- *Threat Score* is the threat score of the source based on UTM features such as Web Filter and antivirus. It shows threat scores allowed and threat scores blocked.
- *Bytes* is the accumulated bytes sent/received. In realtime, this is calculated from the session list, and in historical it is from logs.
- *Sessions* is the total sessions blocked/allowed. In realtime, this is calculated from the session list, and in historical it is from logs.
- *Source* is a simplified version of the first column, including only the IP address without extra information.
- *Source Interface* is the interface from which the traffic originates. In realtime, this is calculated from the session list, and in historical it is from the logs.
- *FortiGate* is the name of the fabric device.
- More information can be shown in a tooltip while hovering over these entries.
- For realtime, two more columns are available, *Bandwidth* and *Packets*, both of which come from the session list.

Hover over linked items in an entry to view additional information. Some information windows provide links to other areas of FortiOS such as the application signatures page.



To select the columns displayed in a table, hover over the header in the first column, and click the *configure table* icon.

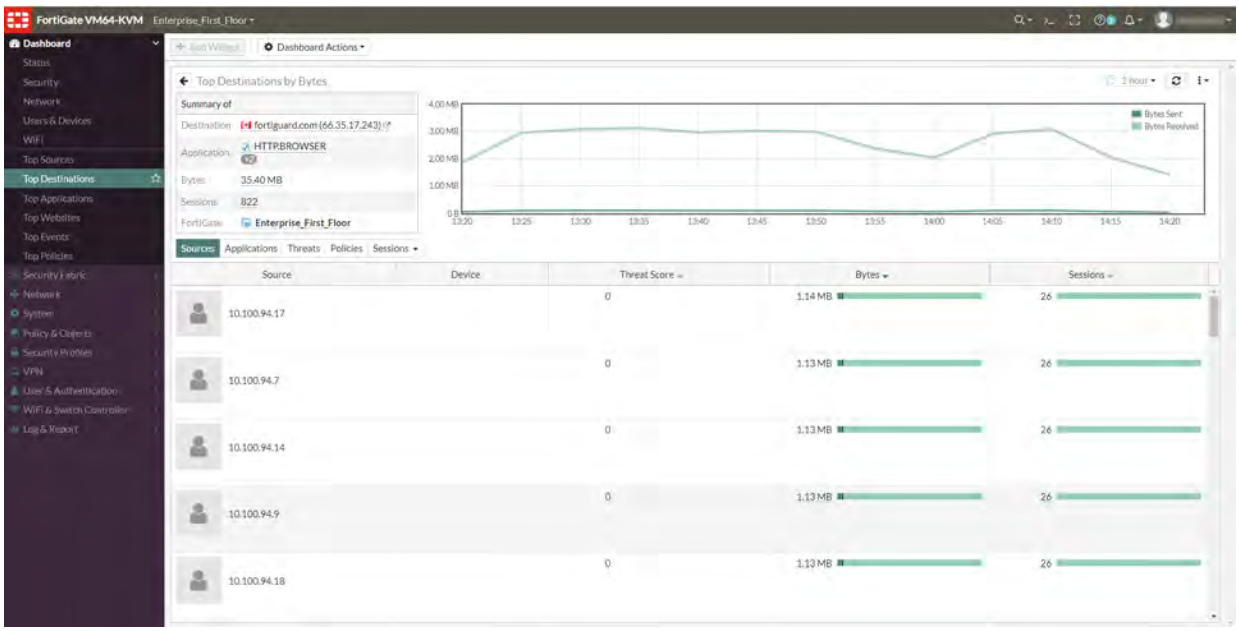


Drilldown information

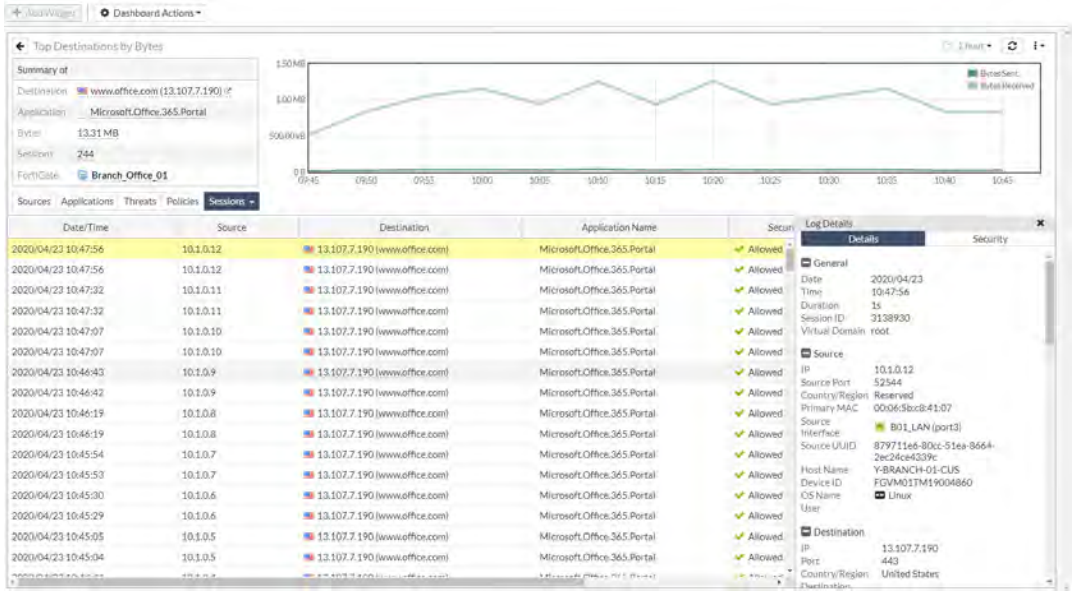
Double-click or right-click an entry in a FortiView dashboard and select *Drill Down to Details* to view additional details about the selected traffic activity. Click the *Back* icon in the toolbar to return to the previous view.

You can group drilldown information into different drilldown views. For example, you can group the drilldown information in the Top Destionations dashboards by *Sources*, *Applications*, *Threats*, and *Policies*.

Netskope Exhibit 1017



Double-click an entry to view the logs in *Sessions* view. Double-click a session to view the logs.



Graph	• The graph shows the bytes sent/received in the time frame. Realtime does not include a chart. • Users can customize the time frame by selecting a time period within the graph.
Summary Information	• Shows information such as the user/avatar, avatar/source IP, bytes, and sessions total for the time period. • Can quarantine host (access layer quarantine) if they are behind a FortiSwitch or FortiAP. • Can ban IP addresses, adds the source IP address into the quarantine list.
Tabs	• Drilling down entries in any of these tabs (except sessions tab) will take you to the

Netskope Exhibit 1017

underlying traffic log in the sessions tab.

- *Applications* shows a list of the applications attributed to the source IP. This can include scanned applications (using Application Control in a firewall policy or unscanned applications.

```
config log gui-display
   set fortiview-unscanned-apps enable
end
```

- *Destinations* shows destinations grouped by IP address/FQDN.
- *Threats* lists the threats caught by UTM profiles. This can be from antivirus, IPS, Web Filter, Application Control, etc.
- *Web Sites* contains the websites which were detected either with webfilter, or through FQDN in traffic logs.
- *Web Categories* groups entries into their categories as dictated by the Web Filter Database.
- *Policies* groups the entries into which polices they passed through or were blocked by.
- *Sessions* shows the underlying logs (historical) or sessions (realtime). Drilldowns from other tabs end up showing the underlying log located in this tab.
- *Search Phrases* shows entries of search phrases on search engines captured by a Web Filter UTM profile, with deep inspection enabled in firewall policy.
- More information can be shown in a tooltip while hovering over these entries.

To view matching logs or download a log, click the *Security* tab in the *Log Details* .



FortiView from disk

FortiView from disk is available on all FortiGates with an SSD disk.

Restrictions

Model	Supported view
Desktop models (100 series) with SSD	Five minutes and one hour
Medium models with SSD	Up to 24 hours
Large models (1500D and above) with SSD	Up to seven days

Model	Supported view

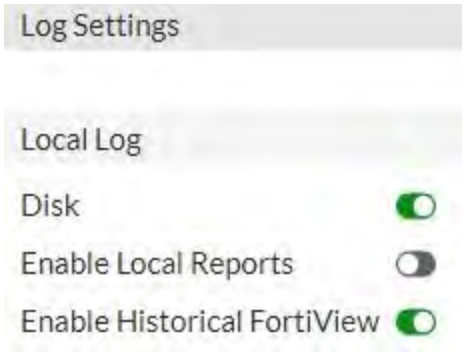
	To enable seven days view:
	`config log setting`
	`set fortiview-weekly-data enable`
	`end`

Configuration

A firewall policy needs to be in place with traffic logging enabled. For optimal operation with FortiView, internal interface roles should be clearly defined as LAN. DMZ and internet facing or external interface roles should be defined as WAN.

To configure logging to disk in the GUI:

1. Enable disk logging from the FortiGate GUI.
   a. Go to *Log & Report* > *Log Settings* > *Local Traffic Log*.
   b. Select the checkbox next to *Disk*.
2. Enable historical FortiView from the FortiGate GUI.
   a. Go to *Log & Report* > *Log Settings* > *Local Traffic Log*.
   b. Select the checkbox next to *Enable Historical FortiView*.



3. Click *Apply*.

To include sniffer traffic and local-deny traffic when FortiView from Disk:

This feature is only supported through the CLI.

```
config report setting
   set report-source forward-traffic sniffer-traffic local-deny-traffic
end
```

To configure the data source in the GUI:

1. Go to *Dashboard* > *Top Sources*.
2. Select a time range other than *now* from the drop-down list to view historical data.
3. In the top menu, click the dropdown, and select *Settings*. The *Edit Dashboard Widget* window opens.
   a. In the *Data Source* area, click *Specify*.
   b. From the dropdown, select *Disk*, and click *OK*.

   For information see, FortiView interface on page 79.

Troubleshooting

Use `execute report flush-cache` and `execute report recreate-db` to clear up any irregularities that may be caused by upgrading or cache issues.

FortiView from FortiAnalyzer

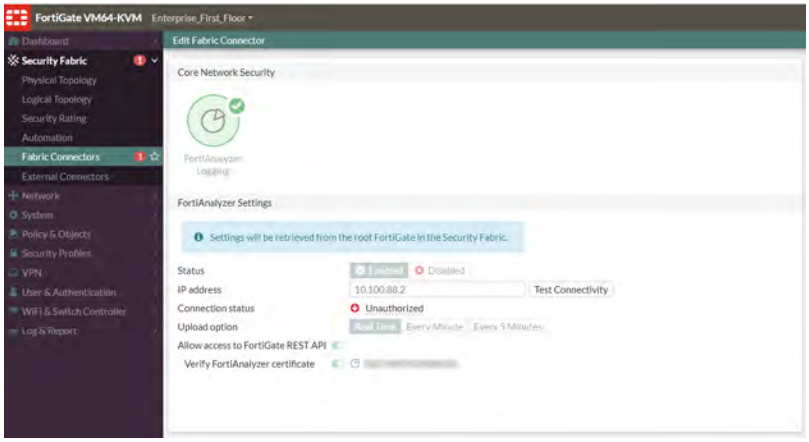
Attach a FortiAnlyzer to FortiGate to increase the functionality of FortiView. Adding a FortiAnalyzer is useful when adding widgets such as the *Compromised Hosts* widget. It also allows historical view for up to seven days.

Requirements

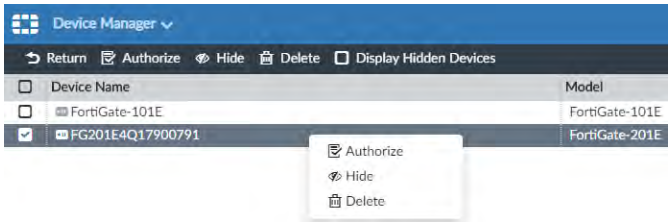
- A FortiGate or FortiOS
- A compatible FortiAnalyzer (see Compatibility with FortiOS)

To configure logging to the FortiAnalyzer:

1. On the FortiGate, go to *Security Fabric* > *Fabric Connectors*, and double-click the *FortiAnalyzer Logging* card.
2. Enter the IP address of the FortiAnalyzer device.
3. Click *Test Connectivity*. A message will be shown stating that the FortiGate is not authorized on the FortiAnalyzer.



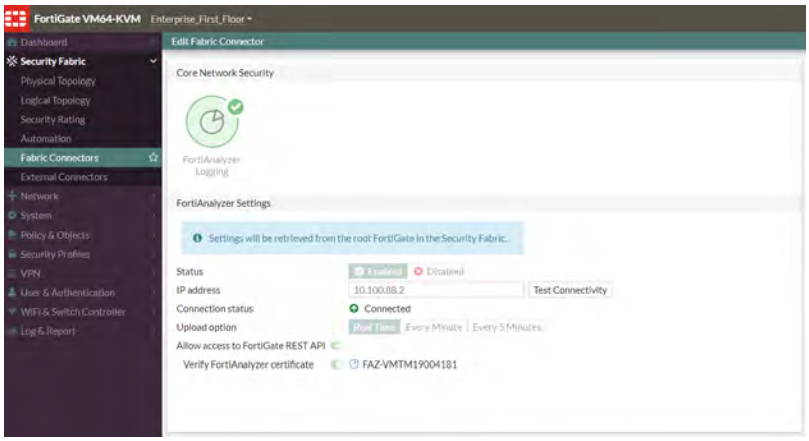
4. On the FortiAnalyzer, go to *Device Manager*, and click the *Unauthorized* tab.
5. In the device list, right-click the FortiGate, then click *Authorize*.



6. On FortiGate, go to *Security Fabric* > *Fabric Connectors*, and double-click the *FortiAnalyzer Logging* card.

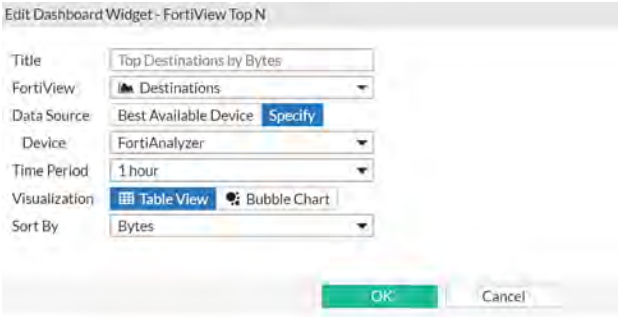
Netskope Exhibit 1017

7. Click *Test Connectivity* to confirm that the device is now authorized.



To enable FortiView from FortiAnalyzer:

1. Go to *Dashboard* > *Top Sources*.
2. Select a time range other than *now* from the drop-down list to view historical data.
3. In top menu, click the dropdown, and select *Settings*. The *Edit Dashboard Widget* window opens.
   a. In the *Data Source* area, click *Specify*.
   b. From the dropdown, select *FortiAnalyzer*, and click *OK*.



All the historical information now comes from the FortiAnalyzer.

When *Data Source* is set to *Best Available Device*, FortiAnalyzer is selected when available, then FortiGate Cloud, and then FortiGate Disk.

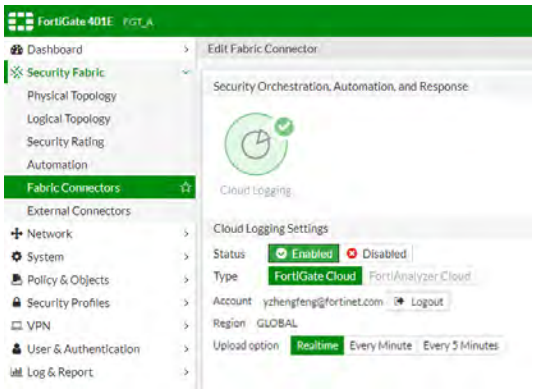
FortiView from FortiGate Cloud

This function requires a FortiGate that is registered and logged into a compatible FortiGate Cloud. When using FortiGate Cloud, the Time Period can be set to up to 24 hours.

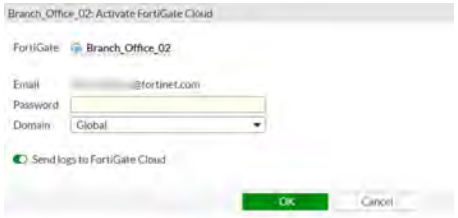
To configure logging to FortiGate Cloud:

1. Go to *Security Fabric* > *Fabric Connectors*, and double-click the *Cloud Logging* card.
2. For *Status*, click *Enabled*.

3. For *Type*, click *FortiGate Cloud*.
   If the FortiGate is registered and logged into FortiGate Cloud, the *Account* and *Region* is displayed.



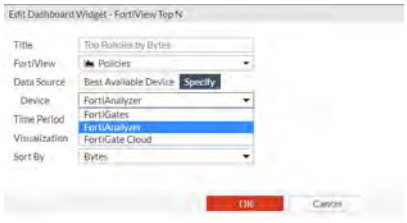
   If the FortiGate is logged out from FortiGate Cloud, click *Activate* and log in, and ensure *Send logs to FortiGate Cloud* is selected.



4. Go to *Log & Report* > *Log Settings* and set the following:
   a. Set *Event Logging* to *All*.
   b. Set *Local Traffic Log* to *All*.

To enable FortiView with log source as FortiGate Cloud:

1. Go to *Dashboard* > *Top Sources*.
2. In the top menu, click the dropdown, and select *Settings*. The *Edit Dashboard Widget* window opens.
   a. In the *Data Source* area, click *Specify*.
   b. From the dropdown, select *FortiGate Cloud*, and click *OK*.





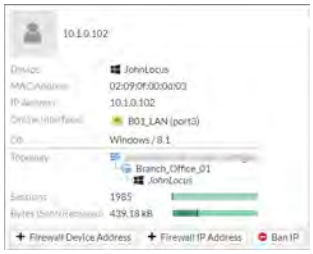
You can select FortiGate Cloud as the data source for all available FortiView pages and widgets.

FortiView top sources

The *Top Sources* dashboard displays top sources sorted by Bytes, Sessions or Threat Score. The information can be displayed in real time or historical views. You can use the dashboard to create or edit firewall device address or IP address definitions, and temporarily or permanently ban IPs.

To add a firewall device address:

1. In the *Device* column, hover over the device MAC address. An information window opens.



2. Click *Firewall Device Address*. The *New Address* window opens.
3. Configure the address settings, and click *Return*.

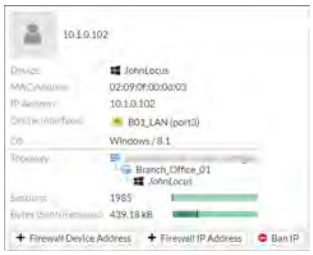


Use the *Name* field to assign a descriptive name to a device so it is easier to find in the *Device* column. After you finish configuring the device, refresh the page to see the new name in the dashboard.



To add a firewall IP address:

1. In the *Device* column, hover over the device MAC address. An information window opens.



2. Click *Firewall IP Address*. The *New Address* window opens.

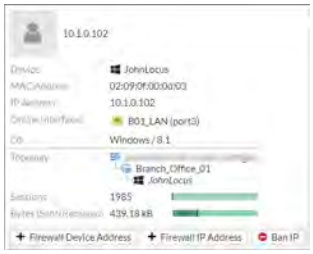
3. Configure the address settings, and click *Return*.

Use the *Name* field to assign a descriptive name to a device so it is easier to find in the *Device* column. After you finish configuring the device, refresh the page to see the new name in the dashboard.

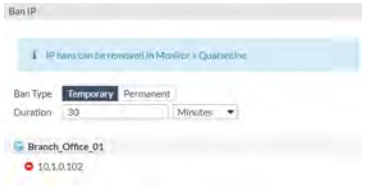


To ban an IP address:

1. In the *Device* column, hover over the device MAC address. An information window opens.



2. Click *Ban IP* . The *Ban IP* window opens.
3. Configure the ban IP settings, and click *OK*.

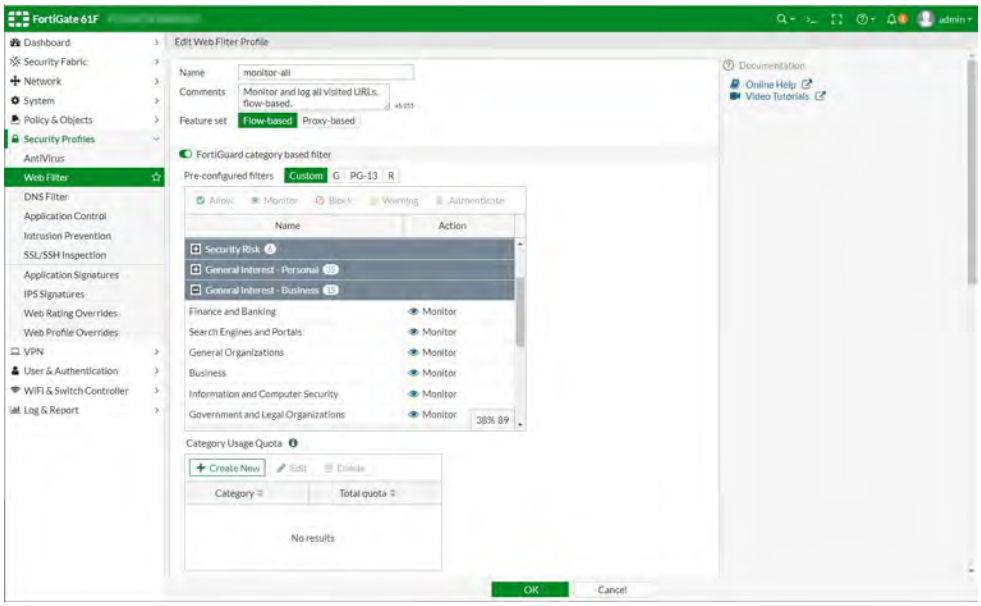


Viewing top websites and sources by category

You can use FortiGuard web categories to populate the category fields in various FortiViews such as *Top Web Categories*, *Top Websites* or *Top Sources*. To view the categories in a dashboard, the web filter profile must be configured to at least monitor for FortiGuard category based filter, and applied to a firewall policy for outbound traffic.

To verify the web filter profile is monitor-only:

1. Go to *Security Profiles* > *Web Filter*.
2. Double-click a web filter that is applied to an outbound traffic firewall policy. The *Edit Web Filter Profile* window opens.
3. Ensure *FortiGuard category based filter* is enabled.
   In the image below, the *General Interest - Business* categories are monitor-only.

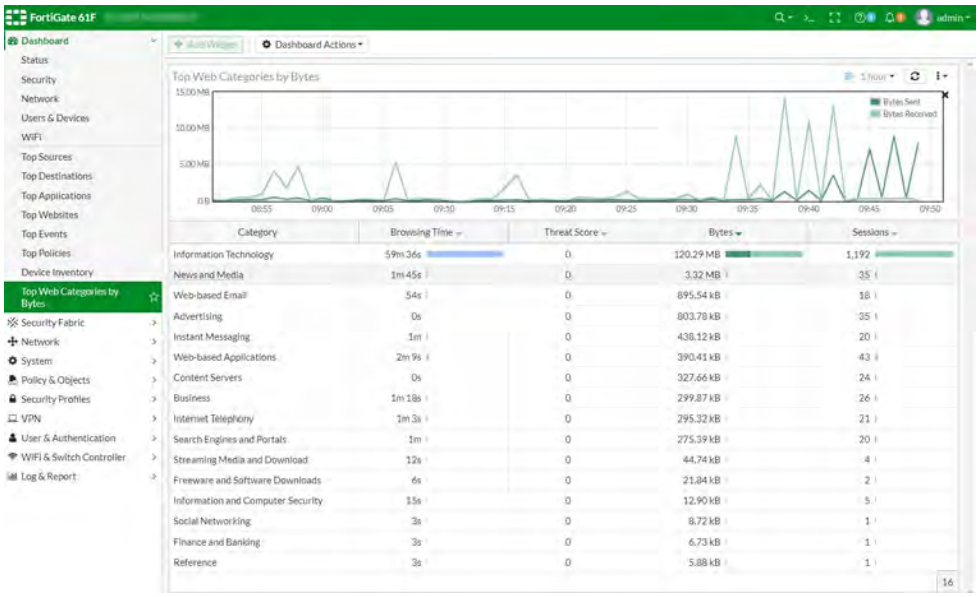


To create a top categories dashboard:

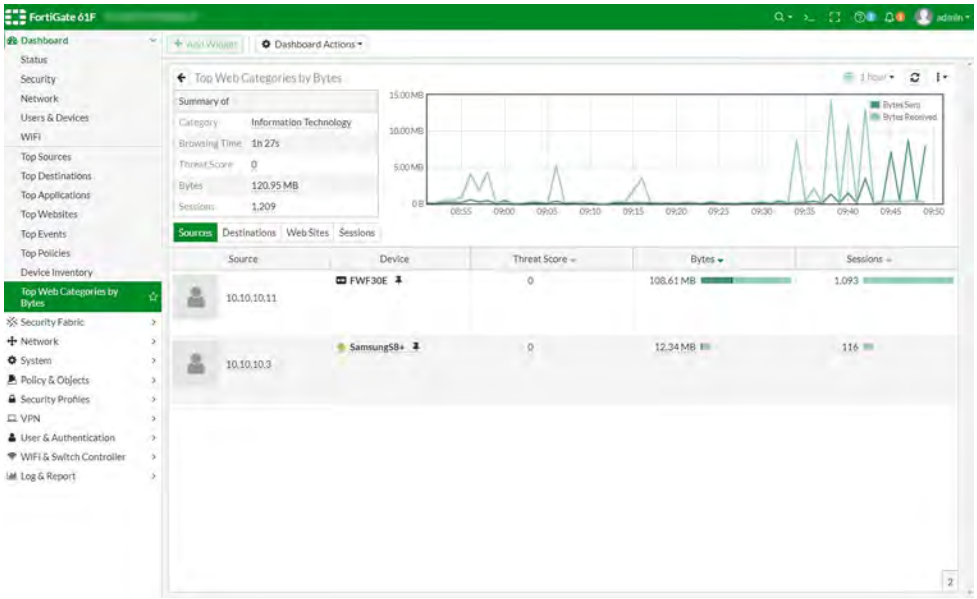
1. In the tree menu, open a dashboard, and click *Dashboard Actions* > *Create New*. The *Add Dashboard* window opens.
2. Configure the dashboard settings.
   a. In the *Name* field , enter a name such as *Top Categories by Bytes*.
   b. From the *Layout Type* select *Standalone*.
   c. Click *OK*. The dashboard is added to the tree menu.
3. Open the dashboard, and click *Add Widget*. The *Add Dashboard Widget* window opens.
4. Click *FortiView Top N*. The *Add Dashboard Widget - FortiView Top N* window opens.
5. From the *FortiView* dropdown, select *Web Categories*.
6. From the *Time Period* dropdown, select a time period greater than *Now*.
7. Click *Add Widget*. The widget is added to the dashboard.

Viewing the web filter category

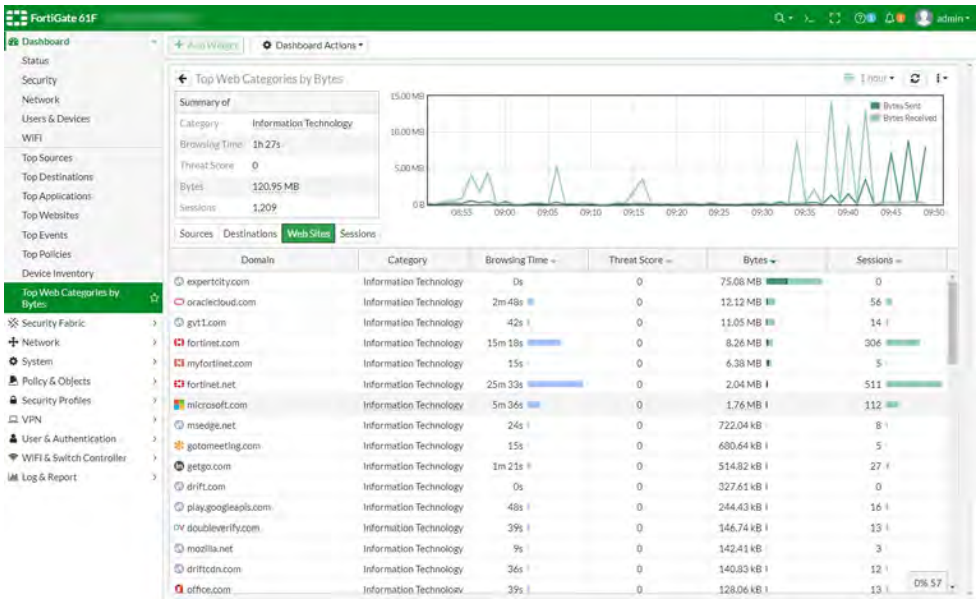
The web filter category name appears in the *Category* column of the dashboard.



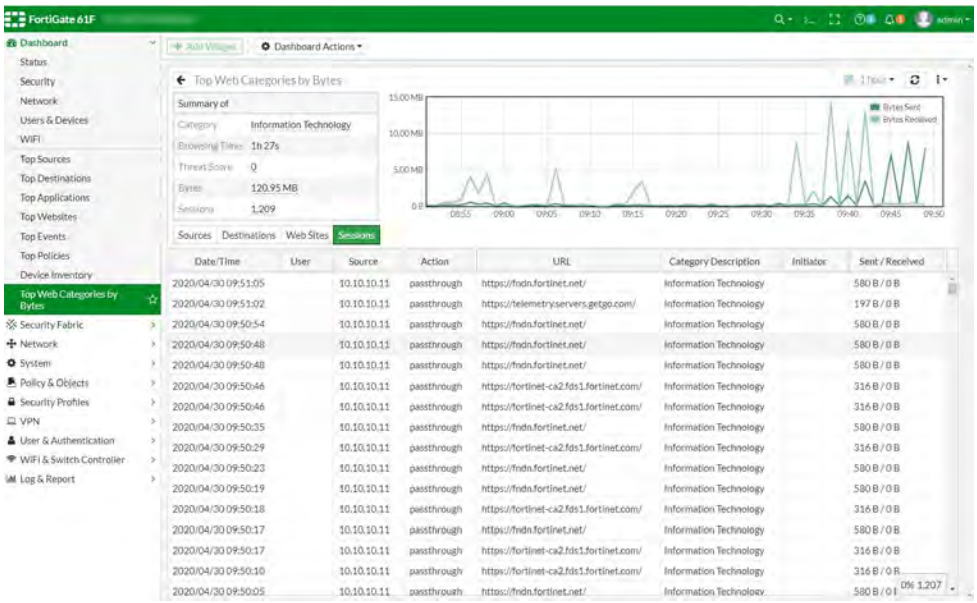
Click an entry in the table. The category name appears at the top of *Summary of* box.



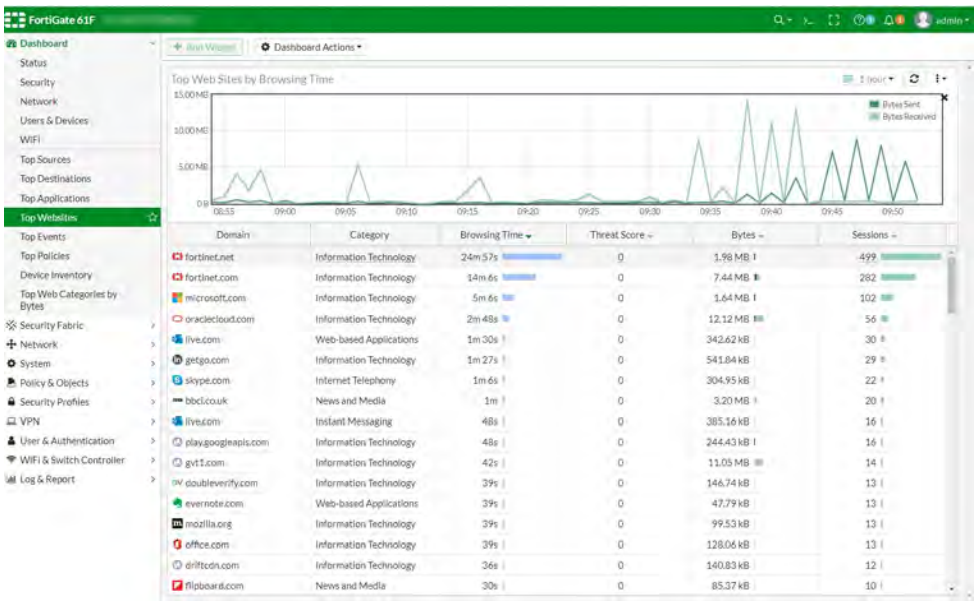
Click the *Web Sites* tab. The category name appears in the *Category* column.



Click the *Sessions* tab. The category name appears in the *Category Description* column.



The category name also appears in the *Category* column in the *Top Websites* and *Top Sources* dashboards.



Cloud application view

To see different cloud application views, set up the following:

- A FortiGate having a relative firewall policy with the Application Control security profile.
- A FortiGate with log data from the local disk or FortiAnalyzer.
- Optional but highly recommended: *SSL Inspection* set to *deep-inspection* on relative firewall policies.

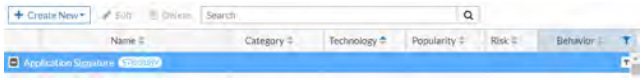
Viewing cloud applications

Cloud applications

All cloud applications require *SSL Inspection* set to *deep-inspection* on the firewall policy. For example, `Facebook_File.Download` can monitor Facebook download behavior which requires *SSL deep-inspection* to parse the deep information in the network packets.

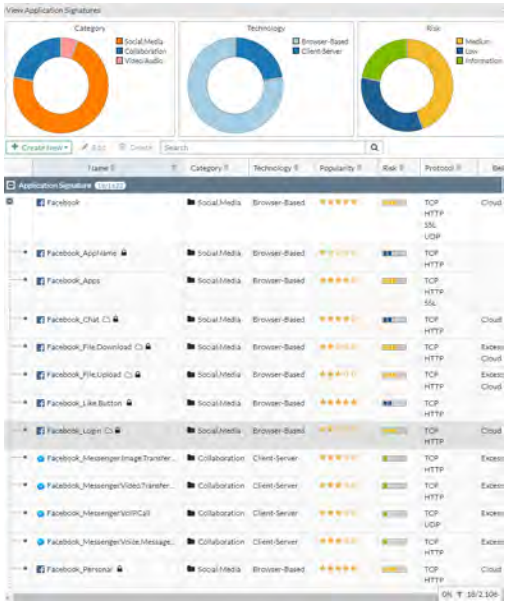
To view cloud applications:

1. Go to *Security Profiles* > *Application Control*.
2. Select a relative Application Control profile used by the firewall policy and click *Edit*.
3. On the *Edit Application Sensor* page, click *View Application Signatures*.
4. Hover over a column heading or the *Application Signature* bar. In the right gutter area, click the filter icon to filter the applications.

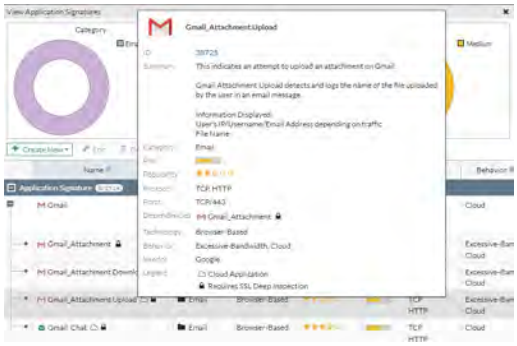


Cloud applications have a cloud icon beside them.

The lock icon indicates that the application requires SSL deep inspection.



5. Hover over an item to see its details.
   This example shows *Gmail_Attachment.Download*, a cloud application signature based sensor which requires SSL deep inspection. If any local network user behind the firewall logs into Gmail and downloads a Gmail attachment, that activity is logged.



Applications with cloud behavior

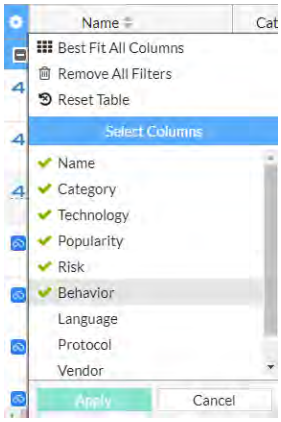
Applications with cloud behavior is a superset of cloud applications.

Some applications do not require SSL deep inspection, such as Facebook, Gmail, and YouTube. This means that if any traffic trigger application sensors for these applications, there is a FortiView cloud application view for that traffic.

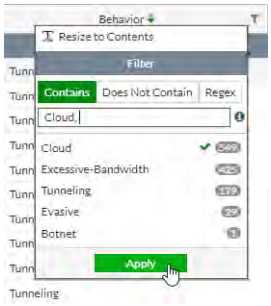
Other applications require SSL deep inspection, such as Gmail attachment, Facebook_Workplace, and so on.

To view applications with cloud behavior:

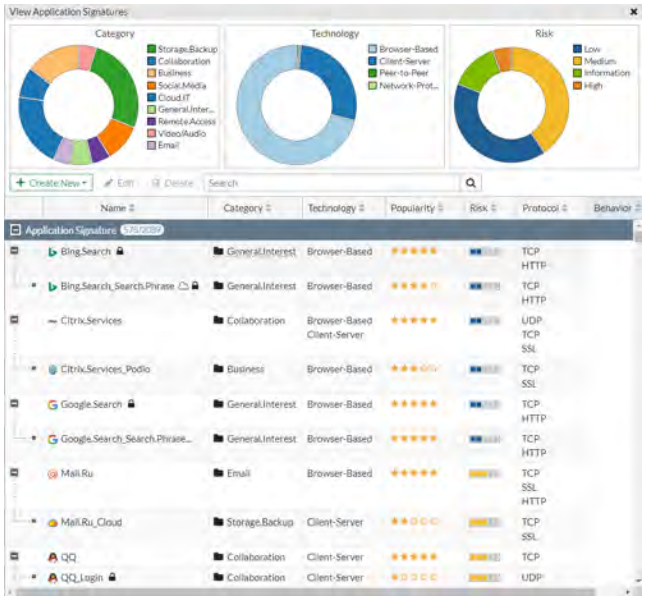
1. In the *Application Signature* page, ensure the *Behavior* column is displayed. If necessary, add the *Behavior* column.
   a. Hover over the left side of the table column headings to display the *Configure Table* icon.
   b. Click *Configure Table* and select *Behavior*.
   c. Click *Apply*.



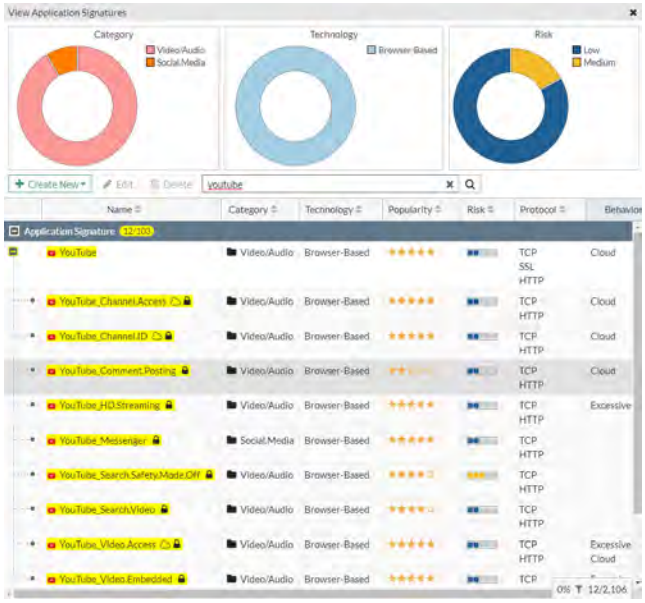
2. Click the filter icon in the *Behavior* column and select *Cloud* to filter by Cloud. Then click *Apply*.



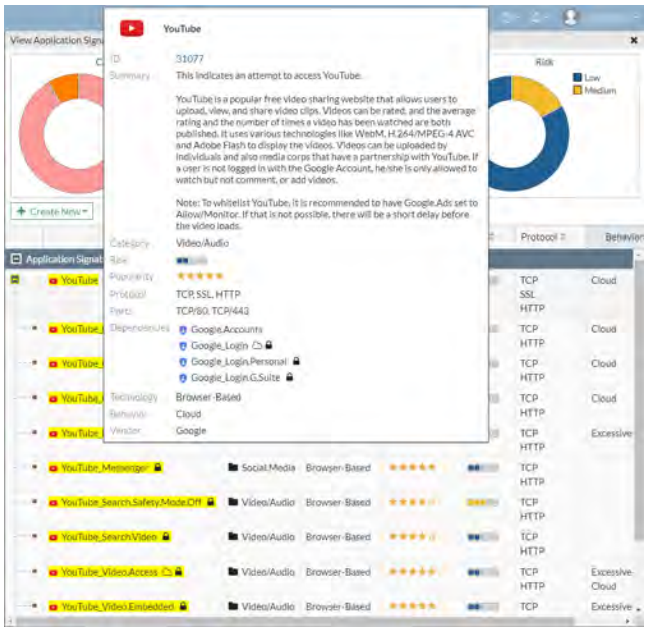
3. The *Application Signature* page displays all applications with cloud behavior.



4. Use the *Search* box to search for applications. For example, you can search for *youtube*.



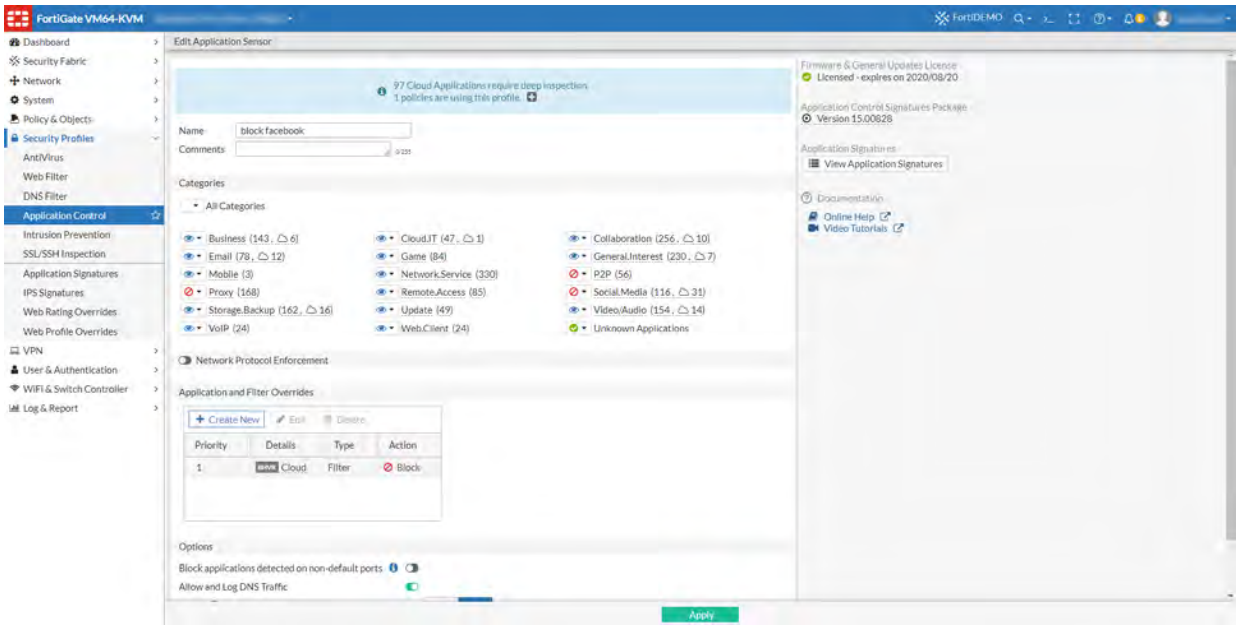
5. Hover over an item to see its details.
This example shows an application sensor with no lock icon which means that this application sensor does not require SSL deep inspection. If any local network user behind the firewall tries to navigate to the YouTube website, that activity is logged.



Configuring the Cloud Applications widget

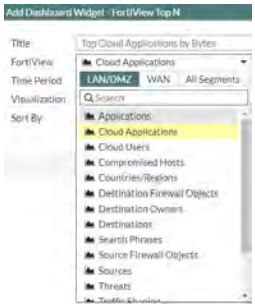
On the *Edit Application Sensor* page in the *Categories* section, the eye icon next to a category means that category is monitored and logged.

Netskope Exhibit 1017

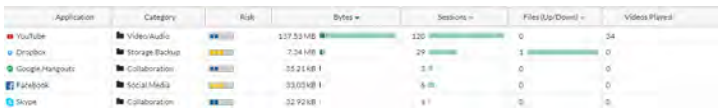


To add the Cloud Applications widget in the GUI:

1. Go to *Dashboard* , and select a dashboard in the tree menu.
2. In the dashboard banner, click *Add Widget*. The *Add Dashboard Widget* window opens.
3. Select the *FortiView Top N* widget.
4. In the *FortiView* dropdown list, select the *LAN/DMZ* tab and then select *Cloud Applications*.



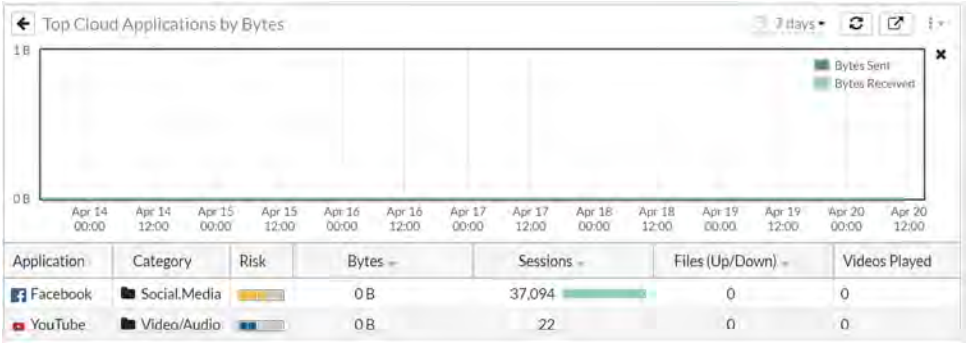
5. Click *Add Widget*. The widget is added to the dashboard.
6. Hover the widget and click *Expand to Full Screen*.
7. If SSL deep inspection is enabled on the relative firewall, then the widget shows the additional details that are logged, such as *Files (Up/Down)* and *Videos Played*.
   For YouTube, the *Videos Played* column is triggered by the *YouTube_Video.Play* cloud application sensor. This shows the number of local network users who logged into YouTube and played YouTube videos.
   For Dropbox, the *Files (Up/Down)* column is triggered by *Dropbox_File.Download* and *Dropbox_File.Upload* cloud application sensors. This shows the number of local network users who logged into Dropbox and uploaded or downloaded files.



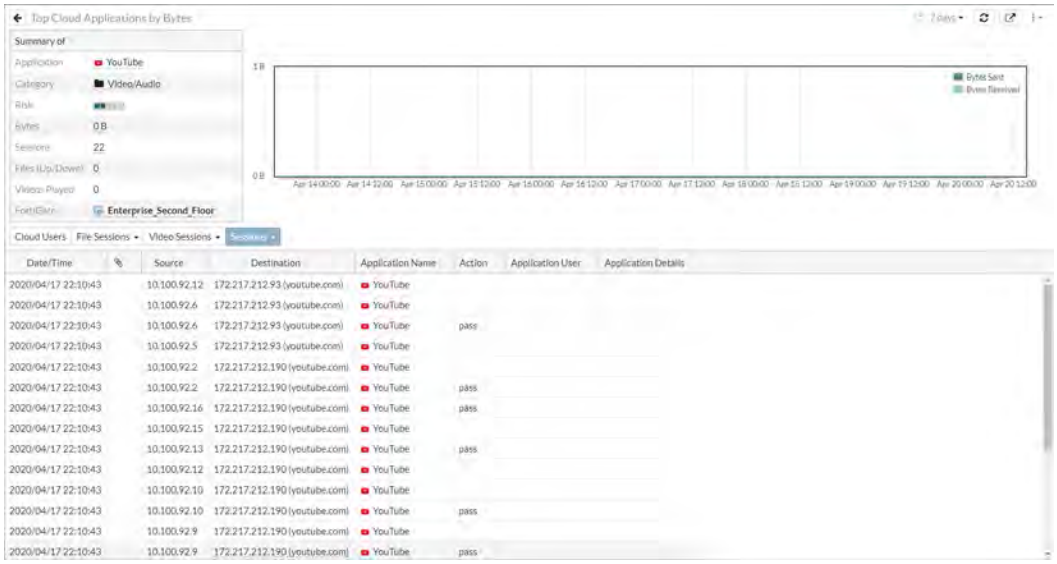
Using the Cloud Applications widget

To see additional information in the Cloud Applications widget:

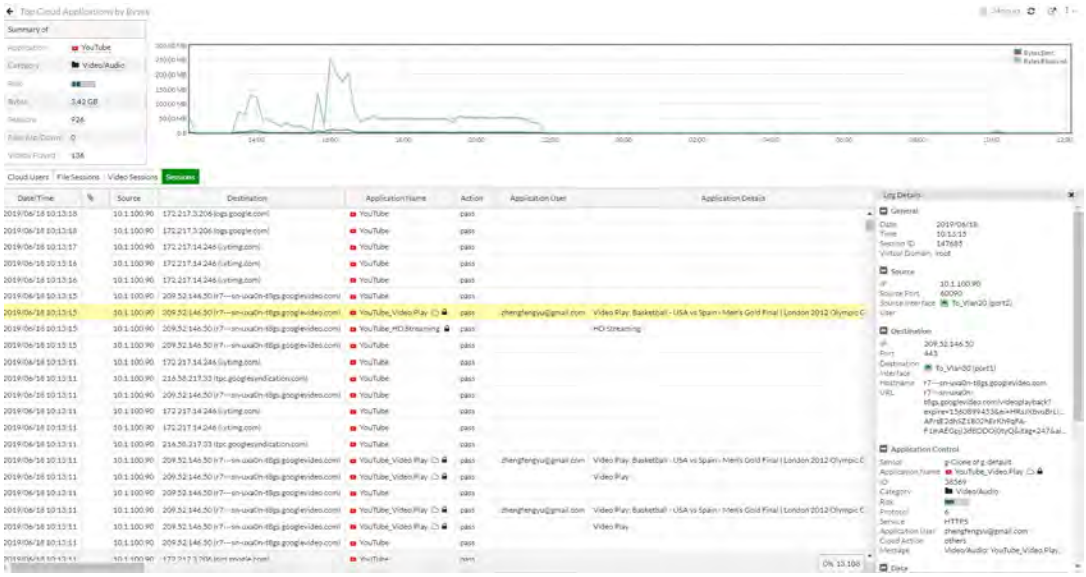
1. Hover over the widget in the dashboard, and click *Expand to full screen*.



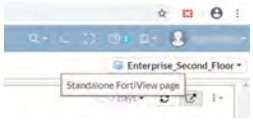
2. For details about a specific entry, double-click the entry or right-click the entry and select *Drill Down to Details*.
3. To see all the sessions for an application, click *Sessions*.
   In this example, the *Application Name* column shows all applications related to YouTube.



4. To view log details, double-click a session to display the *Log Details* pane.
   Sessions monitored by SSL deep inspection (in this example, Youtube_Video.Play) captured deep information such as *Application User*, *Application Details*, and so on. The *Log Details* pane also shows additional deep information such as application *ID*, *Message*, and so on.
   Sessions not monitored by SSL deep inspection (YouTube) did not capture the deep information.



5. In the top menu, click the *Standalone FortiView page* icon to see the page in standalone view.



6. To display a specific time period, select and drag in the timeline graph to display only the data for that time period.



Top application: YouTube example

Monitoring network traffic with SSL deep inspection

This example of monitors network traffic for YouTube using FortiView *Top Application* view with SSL deep inspection.

To monitor network traffic with SSL deep inspection:

1. Use a firewall policy with the following settings. If necessary, create a policy with these settings.
   - *Application Control* is enabled.
   - *SSL Inspection* is set to *deep-inspection*.

- *Log Allowed Traffic* is set to *All Sessions*.



2. Go to *Security Profiles* > *Application Control*.
3. Select a relative Application Control profile used by the firewall policy and click *Edit*.
4. Because YouTube cloud applications are categorized into *Video/Audio*, ensure the *Video/Audio* category is monitored.
   Monitored categories are indicate by an eye icon.
5. Click *View Application Signatures* and hover over YouTube cloud applications to view detailed information about YouTube application sensors.
6. Expand *YouTube* to view the Application Signatures associated with the application.

Application Signature	Description	Application ID
YouTube_Video.Access	An attempt to access a video on YouTube.	16420
YouTube_Channel.ID	An attempt to access a video on a specific channel on YouTube.	44956
YouTube_Comment.Posting	An attempt to post comments on YouTube.	31076
YouTube_HD.Streaming	An attempt to watch HD videos on YouTube.	33104
YouTube_Messenger	An attempt to access messenger on YouTube.	47858
YouTube_Video.Play	An attempt to download and play a video from YouTube.	38569
YouTube_Video.Upload	An attempt to upload a video to YouTube.	22564
YouTube	An attempt to access YouTube. This application sensor does not depend on SSL deep inspection so it does not have a cloud or lock icon.	31077
YouTube_Channel.Access	An attempt to access a video on a specific channel on YouTube.	41598

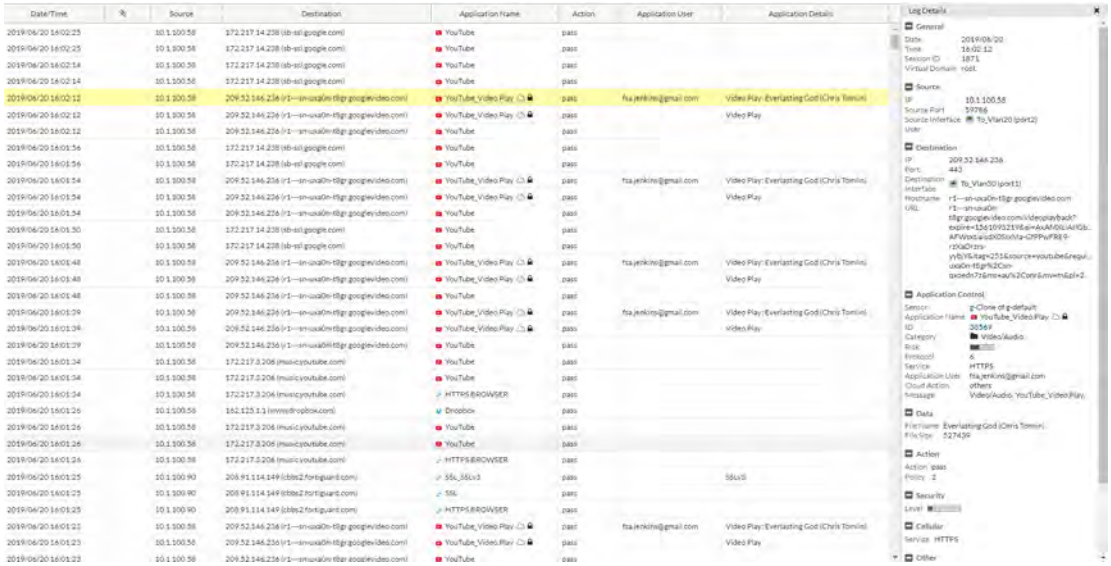


To view the application signature description, click the ID link in the information window.

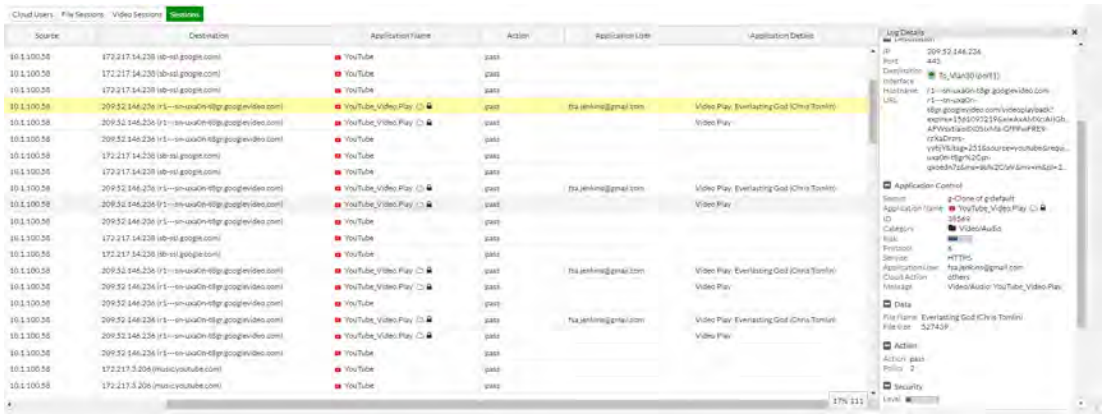
7. On the test PC, log into YouTube and play some videos.
8. On the FortiGate, go to *Log & Report* > *Application Control* and look for log entries for browsing and playing YouTube videos.

Netskope Exhibit 1017

In this example, note the *Application User* and *Application Details*. Also note that the *Application Control ID* is *38569* showing that this entry was triggered by the application sensor *YouTube_Video.Play*.



9. Go to *Dashboard* > *Top Applications*.
10. In the *Top Applications* dashboard, double-click *YouTube* to view the drilldown information.
11. Select the *Sessions* tab to see all the entries for the videos played. Check the sessions for *YouTube_Video.Play* with the ID *38569*.



Monitoring network traffic without SSL deep inspection

This example of monitors network traffic for YouTube using FortiView cloud application view without SSL deep inspection.

To monitor network traffic without SSL deep inspection:

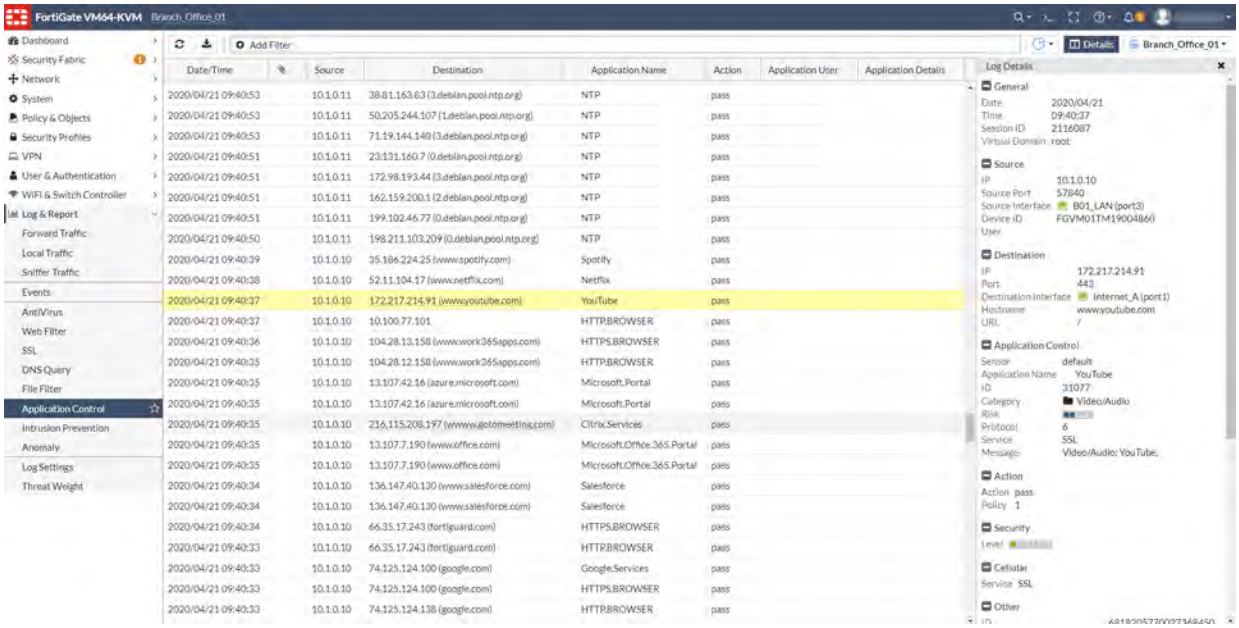
1. Use a firewall policy with the following settings. If necessary, create a policy with these settings.
   - *Application Control* is enabled.
   - *SSL Inspection* is set to *certificate-inspection*.

Netskope Exhibit 1017

- *Log Allowed Traffic* is set to *All Sessions*.

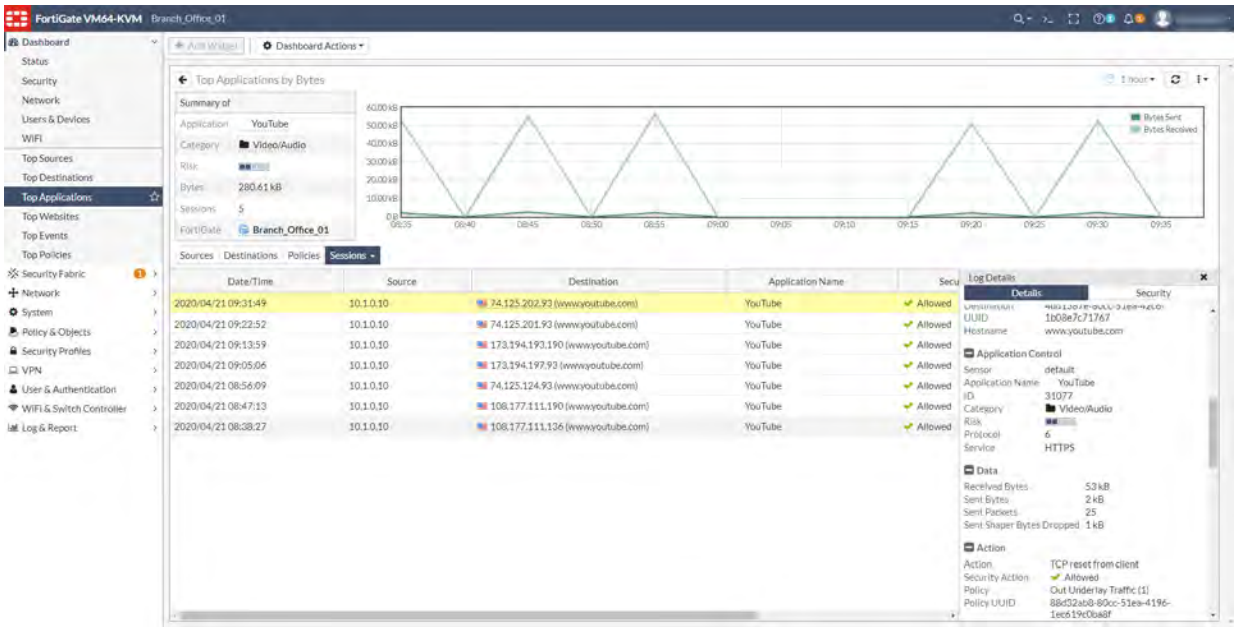


2. On the test PC, log into YouTube and play some videos.
3. On the FortiGate, go to *Log & Report* > *Application Control* and look for log entries for browsing and playing YouTube videos.
   In this example, the log shows only applications with the name YouTube. The log cannot show YouTube application sensors which rely on SSL deep inspection.



4. Go to *Dashboard* > *Top Applications*.
   The *Top Cloud Application by Bytes* dashbaord shows the YouTube cloud application without the video played information that requires SSL deep inspection.

5. Double-click *YouTube* and click the *Sessions* tab.
   These sessions were triggered by the application sensor *YouTube* with the ID *31077*. This is the application sensor with cloud behavior which does not rely on SSL deep inspection.



FortiView Top Source and Top Destination Firewall Objects widgets

The FortiView *Top Source Firewall Objects* and *Destination Firewall Objects* widgets leverage UUID to resolve firewall object address names for improved usability.

Requirements

To have a historical *Firewall Objects*-based view, address objects' UUIDs need to be logged.

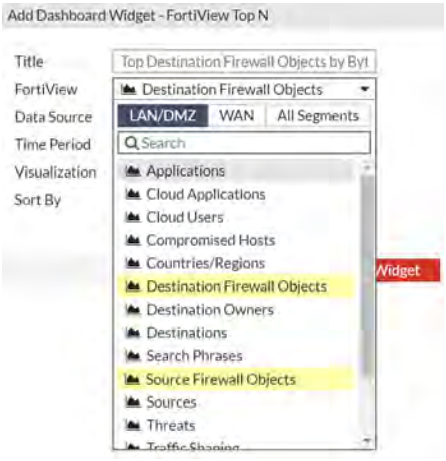
To enable address object UUID logging in the CLI:

```
config system global
     set log-uuid-address enable
end
```

To add a firewall object widget in the GUI:

1. Open a dashboard, and click *Add Widget*. The *Add Dashboard Widget* window opens.
2. In the *FortiView* area, click *Top FortiView N*.

3. In the *LAN/DMZ* tab, select *Destination Firewall Objects* or *Source Firewall Objects*.



4. In the *Data Source* area, select *Best Available* or *Specify*. For information about data sources, see FortiView interface on page 79.
5. From the *Time Period* dropdown, select the time period.
6. In the *Visualization* area, select *Table View* or *Bubble Chart*.
7. From the Sort By dropdown, select the data type.

Example

In this example, firewall addresses have been configured and associated with a unique UUID.

In the *Top Source Firewall Objects* and *Top Destination Firewall Objects* widgets, firewall objects can be displayed in real-time or in a historical chart. Objects can also be drilled down for more details.

To view Firewall Object-based charts in real-time:

1. Open a dashboard, and click *Add Widget*. The *Add Dashboard Widget* window opens.
2. In the *FortiView* area, click *Top FortiView N*.
   a. From the *FortiView* dropdown, select *Destination Firewall Objects*.
   b. From the *Time Period* dropdown, select *now*.
   c. Click *OK*.

To view Firewall Object-based charts over a historical period:

1. Open a dashboard, and click *Add Widget*. The *Add Dashboard Widget* window opens.
2. In the *FortiView* area, click *Top FortiView N*.
   a. From the *FortiView* dropdown, select *Destination Firewall Objects*.
   b. From the *Time Period* dropdown, select a time period other than *now*.
   c. Click *OK*.

To drill down Firewall Objects:

1. Right-click on any *Source* or *Destination Object* in the view results.
2. Select *Drill Down to Details*. More information displays about the object; there are additional criteria to filter data.

Configuration backups

Once you successfully configure the FortiGate, it is extremely important that you backup the configuration. In some cases, you may need to reset the FortiGate to factory defaults or perform a TFTP upload of the firmware, which will erase the existing configuration. In these instances, the configuration on the device will have to be recreated, unless a backup can be used to restore it. You should also backup the local certificates, as the unique SSL inspection CA and server certificates that are generated by your FortiGate by default are not saved in a system backup.

We also recommend that you backup the configuration after *any* changes are made, to ensure you have the most current configuration available. Also, backup the configuration before any upgrades of the FortiGate's firmware. Should anything happen to the configuration during the upgrade, you can easily restore the saved configuration.

Always backup the configuration and store it on the management computer or off-site. You have the option to save the configuration file to various locations including the local PC, USB key, FTP, and TFTP server. The last two are configurable through the CLI only.

If you have VDOMs, you can back up the configuration of the entire FortiGate or only a specific VDOM. Note that if you are using FortiManager or FortiGate Cloud, full backups are performed and the option to backup individual VDOMs will not appear.



You can also backup and restore your configuration using Secure File Copy (SCP). See How to download/upload a FortiGate configuration file using secure file copy (SCP).

You enable SCP support using the following command:

```
config system global
   set admin-scp enable
end
```

For more information about this command and about SCP support, see config system global.

Backing up the configuration

To backup the configuration using the GUI:

1. Click on the user name in the upper right-hand corner of the screen and select *Configuration* > *Backup*.
2. Direct the backup to your *Local PC* or to a *USB Disk*.
   The *USB Disk* option will not be available if no USB drive is inserted in the USB port. You can also backup to the FortiManager using the CLI.
3. If VDOMs are enabled, indicate whether the scope of the backup is the entire FortiGate configuration (*Global*) or only a specific VDOM configuration (*VDOM*).
   If backing up a VDOM configuration, select the VDOM name from the list.
4. Enable *Encryption*. Encryption must be enabled on the backup file to back up VPN certificates.
5. Enter a password, and enter it again to confirm it. This password will be required to restore the configuration.
6. Click *OK*.
7. When prompted, select a location on the PC or USB disk to save the configuration file. The configuration file will have a .conf extension.

To backup the configuration using the CLI:

Use one of the following commands:

```
execute backup config management-station <comment>
```


Netskope Exhibit 1017

or:

```
execute backup config usb <backup_filename> [<backup_password>]
```

or for FTP, note that port number, username are optional depending on the FTP site:

```
execute backup config ftp <backup_filename> <ftp_server> [<port>] [<user_name>]
     [<password>]
```

or for TFTP:

```
execute backup config tftp <backup_filename> <tftp_servers> <password>
```

Use the same commands to backup a VDOM configuration by first entering the commands:

```
config vdom
  edit <vdom_name>
```

Restoring a configuration

To restore the FortiGate configuration using the GUI:

1. Click on the user name in the upper right-hand corner of the screen and select *Configuration* > *Restore*.
2. Identify the source of the configuration file to be restored: your *Local PC* or a *USB Disk*.
   The *USB Disk* option will not be available if no USB drive is inserted in the USB port. You can restore from the FortiManager using the CLI.
3. Click Upload, locate the configuration file, and click *Open*.
4. Enter the password if required.
5. Click *OK*.

To restore the FortiGate configuration using the CLI:

```
execute restore config management-station normal 0
```

or:

```
execute restore config usb <filename> [<password>]
```

or for FTP, note that port number, username are optional depending on the FTP site:

```
execute restore config ftp <backup_filename> <ftp_server> [<port>] [<user_name>]
     [<password>]
```

or for TFTP:

```
execute restore config tftp <backup_filename> <tftp_server> <password>
```

The FortiGate will load the configuration file and restart. Once the restart has completed, verify that the configuration has been restored.

Troubleshooting

When restoring a configuration, errors may occur, but the solutions are usually straightforward.

Netskope Exhibit 1017

Error message	Reason and Solution
Configuration file error	This error occurs when attempting to upload a configuration file that is incompatible with the device. This may be due to the configuration file being for a different model or being saved from a different version of firmware. **Solution**: Upload a configuration file that is for the correct model of FortiGate device and the correct version of the firmware.
Invalid password	When the configuration file is saved, it can be protected by a password. The password entered during the upload process is not matching the one associated with the configuration file. **Solution**: Use the correct password if the file is password protected.

Configuration revision

You can manage multiple versions of configuration files on models that have a 512MB flash memory and higher. Revision control requires either a configured central management server or the local hard drive, if your FortiGate has this feature. Typically, configuration backup to local drive is not available on lower-end models.

The central management server can either be a FortiManager unit or FortiGate Cloud.

If central management is not configured on your FortiGate unit, a message appears instructing you to either

- Enable central management, or
- Obtain a valid license.

When revision control is enabled on your FortiGate unit, and configuration backups have been made, a list of saved revisions of those backed-up configurations appears.

Configuration revisions are viewed by clicking on the user name in the upper right-hand corner of the screen and selecting *Configuration* > *Revisions*.

Backup and restore the local certificates

This procedure exports a server (local) certificate and private key together as a password protected PKCS12 file. The export file is created through a customer-supplied TFTP server. Ensure that your TFTP server is running and accessible to the FortiGate before you enter the command.

To back up the local certificates:

Connect to the CLI and use the following command:

```
execute vpn certificate local export tftp <cert_name> <filename> <tftp_ip>
```

where:

- `<cert_name>` is the name of the server certificate.
- `<filename>` is a name for the output file.
- `<tftp_ip>` is the IP address assigned to the TFTP server host interface.

Netskope Exhibit 1017

To restore the local certificates using the GUI:

1. Move the output file from the TFTP server location to the management computer.
2. Go to *System* > *Certificates* and click *Import* > *Local*.
3. Select the certificate type, then click *Upload* in the *Certificate file* field.
4. On the management computer, browse to the file location, select it, and click *Open*.
5. If the *Type* is *Certificate*, upload the *Key file* as well.
6. If required, enter the *Password* that is required to upload the file or files.
7. Click *OK*.

To restore the local certificates using the CLI:

Connect to the CLI and use the following command:

```
execute vpn certificate local import tftp <filename> <tftp_ip>
```

Restore factory defaults

There may be a need to reset the FortiGate to its original defaults; for example, to begin with a fresh configuration. There are two options when restoring factory defaults. The first resets the entire device to the original out-of-the-box configuration.

You can reset the device with the following CLI command:

```
execute factoryreset
```

When prompted, type `y` to confirm the reset.

Alternatively, in the CLI you can reset the factory defaults but retain the interface and VDOM configuration with the following command:

```
execute factoryreset2
```


Netskope Exhibit 1017

Fortinet Security Fabric

The Fortinet Security Fabric provides an intelligent architecture that interconnects discrete security solutions into an integrated whole to detect, monitor, block, and remediate attacks across the entire attack surface. It delivers broad protection and visibility into every network segment and device, be they hardware, virtual, or cloud based.

- The physical topology view shows all connected devices, including access layer devices. The logical topology view shows information about the interfaces that each device is connected to.
- Security rating checks analyze the Security Fabric deployment to identify potential vulnerabilities and highlight best practices to improve the network configuration, deploy new hardware and software, and increase visibility and control of the network.
- Fabric connectors provide integration with multiple SDN, cloud, and partner technology platforms to automate the process of managing dynamic security updates without manual intervention.
- Automation pairs an event trigger with one or more actions to monitor the network and take the designated actions automatically when the Security Fabric detects a threat.

Security Fabric settings and usage

This section contains information about how to configure the following devices as part of the Fortinet Security Fabric:

- Components on page 113
- Configuring the root FortiGate and downstream FortiGates
- Configuring FortiAnalyzer
- Configuring other Security Fabric devices on page 120
- Using the Security Fabric
- Deploying the Security Fabric on page 158
- Synchronizing objects across the Security Fabric on page 166
- Security Fabric over IPsec VPN on page 170
- Leveraging LLDP to simplify security fabric negotiation on page 176

System requirements

To set up the Security Fabric, the devices that you want to include must meet the Product Integration and Support requirements in the FortiOS Release Notes.

Some features of the Security Fabric are only available in certain firmware versions and models. Not all FortiGate models can run the FortiGuard Security Rating Service if they are the root FortiGate in a Security Fabric. For more information, see the Special Notices in the FortiOS Release Notes.

Prerequisites

- If devices are not already installed in your network, complete basic installation and configuration tasks by following the instructions in the device documentation.

Netskope Exhibit 1017

- FortiGate devices must either have VDOMs disabled or be running in split-task VDOM mode in order to be added to the Security Fabric. See Virtual Domains on page 610.
- FortiGate devices must be operating in NAT mode.

Components

The Fortinet Security Fabric consists of different components that work together to secure you network.

The following devices are required to create a Security Fabric:

Device	Description
FortiGate	FortiGate devices are the core of the Security Fabric and can have one of the following roles: • **Root**: The root FortiGate is the main component in the Security Fabric. It is typically located on the edge of the network and connects the internal devices and networks to the Internet through your ISP. From the root FortiGate, you can see information about the entire Security Fabric on the Physical and Logical Topology pages in the GUI. • **Downstream**: After a root FortiGate is installed, all other FortiGate devices in the Security Fabric act as Internal Segmentation Firewalls (ISFWs), located at strategic points in your internal network, rather than on the network edge. This allows extra security measures to be taken around key network components, such as servers that contain valuable intellectual property. ISFW FortiGate devices create network visibility by sending traffic and information about the devices that are connected to them to the root FortiGate. See Configuring the root FortiGate and downstream FortiGates on page 115 for more information about adding FortiGate devices in the Security Fabric. FortiGate documentation: https://docs.fortinet.com/product/fortigate
FortiAnalyzer	FortiAnalyzer gives you increased visibility into your network, centralized monitoring, and awareness of threats, events, and network activity by collecting and correlating logs from all Security Fabric devices. This gives you a deeper and more comprehensive view across the entire Security Fabric. See Configuring FortiAnalyzer on page 120 for more information about adding FortiAnalyzer devices in the Security Fabric. FortiAnalyzer documentation: https://docs.fortinet.com/product/fortianalyzer

The following devices are recommended:

Device	Description
FortiADC	FortiADC devices optimize the availability, user experience, and scalability of enterprise application delivery. They enable fast, secure, and intelligent acceleration and distribution of even the most demanding enterprise applications. See Additional devices on page 143 for more information about adding FortiADC devices in the Security Fabric. FortiADC documentation: https://docs.fortinet.com/product/fortiadc

Netskope Exhibit 1017

Device	Description
FortiAP	Add FortiAP devices to extend the Security Fabric to your wireless devices. Devices connected to a FortiAP appear in the Physical and Logical Topology pages in the Security Fabric menu. See FortiAP and FortiSwitch on page 142 for more information about adding FortiAP devices in the Security Fabric. FortiAP documentation: https://docs.fortinet.com/product/fortiap
FortiClient	FortiClient adds endpoint control to devices that are located in the Security Fabric, allowing only traffic from compliant devices to flow through the FortiGate. FortiClient compliance profiles are applied by the first FortiGate that a device's traffic flows through. Device registration and on-net status information for a device that is running FortiClient appears only on the FortiGate that applies the FortiClient profile to that device. FortiClient documentation: https://docs.fortinet.com/product/forticlient
FortiClient EMS	FortiClient EMS is used in the Security Fabric to provide visibility across your network, securely share information, and assign security profiles to endpoints. See FortiClient EMS on page 132 for more information about adding FortiClient EMS devices in the Security Fabric. FortiClient EMS documentation: https://docs.fortinet.com/product/forticlient
FortiDDoS	FortiDDoS is a Network Behavior Anomaly (NBA) prevention system that detects and blocks attacks that intend to disrupt network service by overutilizing server resources. See Additional devices on page 143 for more information about adding FortiDDoS devices in the Security Fabric. FortiDDoS documentation: https://docs.fortinet.com/product/fortiddos
FortiMail	FortiMail antispam processing helps offload from other devices in the Security Fabric that would typically carry out this process. See Additional devices on page 143 for more information about adding FortiMail devices in the Security Fabric. FortiMail documentation: https://docs.fortinet.com/product/fortimail
FortiManager	Add FortiManager to simplify the network management of devices in the Security Fabric by centralizing management access in a single device. This allows you to easily control the deployment of security policies, FortiGuard content security updates, firmware revisions, and individual configurations for devices in the Security Fabric. See FortiManager on page 127 for more information about adding FortiManager devices in the Security Fabric. FortiManager documentation: https://docs.fortinet.com/product/fortimanager
FortiSandbox	Add FortiSandbox to your Security Fabric to improve security with sandbox inspection. Sandbox integration allows FortiGate devices in the Security Fabric to automatically receive signature updates from FortiSandbox and add the originating URL of any malicious file to a blocked URL list. See FortiSandbox on page 130 for more information about adding FortiSandbox devices in the Security Fabric. FortiSandbox documentation: https://docs.fortinet.com/product/fortisandbox

Netskope Exhibit 1017

Device	Description
FortiSwitch	A FortiSwitch can be added to the Security Fabric when it is managed by a FortiGate that is in the Security Fabric with the FortiLink protocol, and connected to an interface with *Security Fabric Connection* enabled. FortiSwitch ports to become logical extensions of the FortiGate. Devices connected to the FortiSwitch appear in the Physical and Logical Topology pages in the Security Fabric menu, and security features, such as FortiClient compliance profiles, are applied to them. See FortiAP and FortiSwitch on page 142 for more information about adding FortiSwitch devices in the Security Fabric. FortiSwitch documentation: https://docs.fortinet.com/product/fortiswitch
FortiWeb	Add FortiWeb to defend the application attack surface from attacks that target application exploits. You can also configure FortiWeb to apply web application firewall features, virus scanning, and web filtering to HTTP traffic to help offload from other devices in the Security Fabric that would typically carry out these processes. See Additional devices on page 143 for more information about adding FortiWeb devices in the Security Fabric. FortiWeb documentation: https://docs.fortinet.com/product/fortiweb
FortiWLC	FortiWLC delivers seamless mobility and superior reliability with optimized client distribution and channel utilization. Both single and multi channel deployment options are supported, maximizing efficiency to make the most of available wireless spectrum. See Additional devices on page 143 for more information about adding FortiWLC devices in the Security Fabric. FortiWLC documentation: https://docs.fortinet.com/product/wireless-controller

The following devices are optional:

Device	Description
Other Fortinet products	Many other Fortinet products can be added to the Security Fabric, including FortiAuthenticator, FortiToken, FortiCache, and FortiSIEM. Documentation: https://docs.fortinet.com/
Third-party products	Third-party products that belong to the Fortinet Fabric-Ready Partner Program can be added to the Security Fabric.

Configuring the root FortiGate and downstream FortiGates

The following procedures include configuration steps for a typical Security Fabric implementation, where the edge FortiGate is the root FortiGate, and the downstream FortiGate devices are all devices that are downstream from the root FortiGate.

Prerequisites

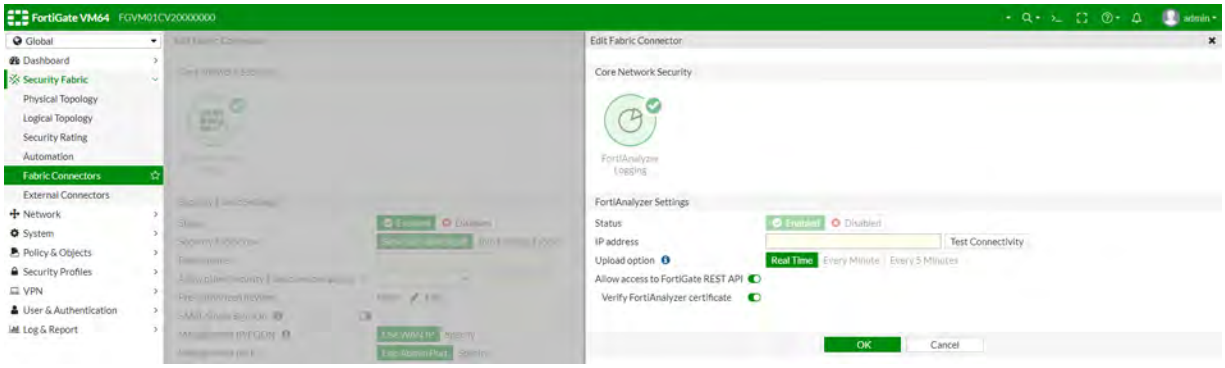
- FortiGate devices must either have VDOMs disabled or be running in split-task VDOM mode in order to be added to the Security Fabric. See Virtual Domains on page 610.
- FortiGate devices must be operating in NAT mode.

Configure the root FortiGate

The edge FortiGate is typically configured as the root FortiGate, as this allows you to view the full topology of the Security Fabric from the top down.

To configure the root FortiGate:

1. On the root FortiGate, go to *Security Fabric* > *Fabric Connectors* and double-click the *Security Fabric Setup* card.
2. For *Status*, click *Enable*.
3. Set the *Security Fabric role* to *Serve as Fabric Root*. FortiAnalyzer logging is automatically enabled and the settings can be configured.



4. Enter the FortiAnalyzer IP and select and *Upload option*.
5. In the *FortiAnalyzer Logging* section, in the *IP address* field, enter the IP address of the FortiAnalyzer.
6. If required, enable *Allow access to FortiGate REST API* and, optionally, *Verify FortiAnalyzer certificate*.
   The REST API accesses the FortiGate topology and shares data and results. The FortiGate will verify the FortiAnalyzer by retrieving its serial number and checking it against the FortiAnalyzer certificate. When verified, the FortiAnalyzer serial number is stored in the FortiGate configuration. When authorizing the FortiGate on the FortiAnalyzer, the FortiGate admin credentials do not need to be entered.
7. Click *Test Connectivity*.
   If you select *Test Connectivity* and this is the first time that you are connecting the FortiGate to the FortiAnalyzer, you will receive a warning message because the FortiGate has not yet been authorized on the FortiAnalyzer. You can configure this authorization when you configure the FortiAnalyzer. See Configuring FortiAnalyzer on page 120.
8. Click *OK*. The FortiAnalyzer serial number is verified.
9. Enter a *Fabric name*.
10. Ensure *Allow other Security Fabric devices to join* is enabled and add the interfaces.
11. Click *OK*.

Using the root FortiGate with disk to store historic user and device information

This backend implementation allows the root FortiGate in a Security Fabric to store historic user and device information in a database on its disk. This will allow administrators to visualize users and devices over a period of time.

A new daemon, user_info_history, stores this data on the disk. The information source for the historical data will be the user_info daemon, which would be recorded on the disk when user_info notifies user_info_history that a user has logged out or the device is no longer connected.

Add downstream devices

Downstream FortiGate devices can be securely added to the Security Fabric without sharing the password of the root FortiGate. Downstream device serial numbers can be authorized from the root FortiGate, or allowed to join by request. New authorization requests include the device serial number, IP address, and HA members. HA members can include up to four serial numbers and is used to ensure that, in the event of a fail over, the secondary FortiGate is still authorized.

Pre-authorizing the downstream FortiGate

When a downstream Fortinet device's serial number is added to the trusted list on the root FortiGate, the device can join the Security Fabric as soon as it connects. After the new device is authorized, connected FortiAP and FortiSwitch devices are automatically included in the topology, where they can be authorized with one click.

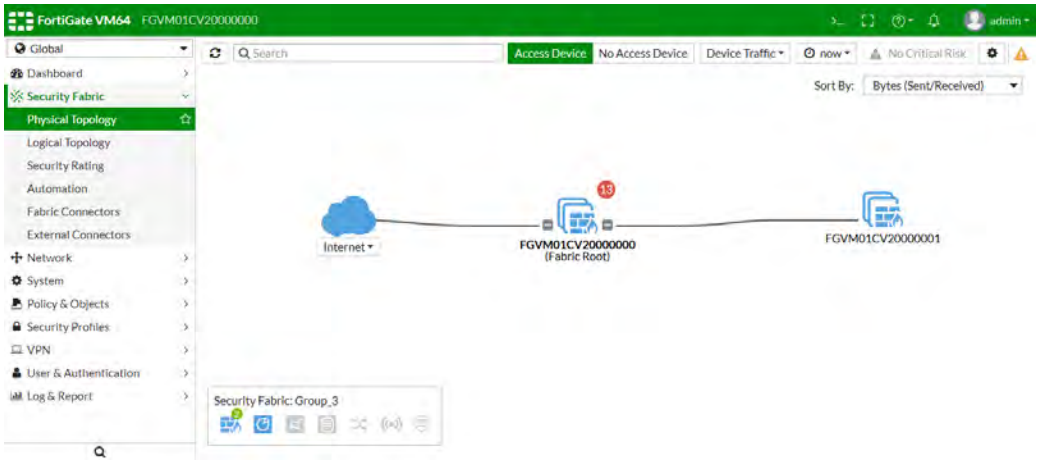
The interface that connects to the downstream FortiGate must have *Security Fabric Connection* enabled.

To pre-authorize a FortiGate:

1. On the root FortiGate, go to *Security Fabric* > *Fabric Connectors* and double-click the *Security Fabric Setup* card.
2. In the *Pre-authorized devices*, click *Edit*. The *Pre-Authorized Devices* window opens.
3. Add a new FortiGate to the list using the downstream device's serial number.



4. On the downstream FortiGate, go to *Security Fabric* > *Fabric Connectors* and double-click the *Security Fabric Setup* card.
5. For *Status*, click *Enable*.
6. Set *Security Fabric role* to *Join Existing Fabric*.
7. Enter the IP address of the root FortiGate in the *Upstream FortiGate IP* field.
8. Click *OK*.
9. On the root FortiGate, go to *Security Fabric* > *Physical Topology* and verify that the downstream FortiGate that you added appears in the Security Fabric topology.



Using LLDP

You can automatically prompt downstream FortiGate devices to join the Security Fabric using Link Layer Discovery Protocol (LLDP) and interface role assignments.

1. On the root FortiGate, assign the LAN role to all interfaces that may connect to downstream FortiGate devices. When the LAN role is assigned to an interface, LLDP transmission is enabled by default.
2. When a downstream FortiGate is installed, assign the WAN role to the interface that connects to the upstream FortiGate.
   When the WAN role is assigned, LLDP reception is enabled by default. The newly installed FortiGate uses LLDP to discover the upstream FortiGate, and the administrator is prompted to configure the FortiGate to join the Security Fabric.
3. On the root FortiGate, the new FortiGate must be authorized before it can join the Security Fabric.



If the network contains switches or routers, LLDP may not function as expected because some devices do not pass LLDP packets.

Device request

A device can request to join the Security Fabric from another FortiGate, but it must have the IP address of the root FortiGate. The administrator of the root FortiGate must also authorize the device before it can join the Security Fabric.

The root FortiGate must have *Security Fabric Connection* enabled on the interface that the device connects to.

To enable FortiTelemetry on an interface:

1. Go to *Network* > *Interfaces*.
2. Edit the interface that the device that you authorizing to join the Security Fabric is connected to.
3. Under *Administrative Access*, enable *Security Fabric Connection*.
4. Under *Network*, turn on *Device Detection*.

To join the Security Fabric by device request:

1. Connect to the unauthorized FortiGate or FortiWiFi device, and go to *Security Fabric* > *Fabric Connectors* and double-click the *Security Fabric Setup* card.
2. For *Status*, click *Enable*.
3. Set *Security Fabric role* to *Join Existing Fabric*.
4. Set *Upstream FortiGate IP* to the IP address of the upstream FortiGate.
5. Connect to the root FortiGate and go to *Security Fabric* > *Fabric Connectors*. The new FortiGate appears in the topology tree as unauthorized.
6. Click the unauthorized device and select *Authorize* to authorize the device.

CLI commands

Use the following commands to view, accept, and deny authorization requests, to view upstream and downstream devices, and to list or test fabric devices:

Netskope Exhibit 1017

Command	Description
diagnose sys csf authorization pending-list	View pending authorization requests on the root FortiGate.
diagnose sys csf authorization accept <serial-number-value>	Authorize a device to join the Security Fabric.
diagnose sys csf authorization deny <serial-number-value>	Deny a device from joining the Security Fabric.
diagnose sys csf downstream	Show connected downstream devices.
diagnose sys csf upstream	Show connected upstream devices.
diagnose sys csf fabric-device list	List all known fabric devices.
diagnose sys csf fabric-device test	Test connections to locally configured fabric devices.

Desynchronizing settings

By default, the settings for FortiAnalyzer logging, central management, sandbox inspection, and FortiClient EMS are synchronized between all FortiGate devices in the Security Fabric. To disable the automatic synchronization of these settings, use the following CLI command:

```
config system csf
   set configuration-sync local
end
```

Deauthorizing a device

A device can be deauthorized to remove it from the Security Fabric.

To deauthorize a device:

1. On the root FortiGate, go to *Security Fabric* > *Fabric Connectors*
2. In the topology tree, click the device and select *Deauthorize*.

After devices are deauthorized, the devices' serial numbers are saved in a trusted list that can be viewed in the CLI using the `show system csf` command. For example, this result shows a deauthorized FortiSwitch:

```
show system csf
   config system csf
      set status enable
      set group-name "Office-Security-Fabric"
      set group-password ENC 1Z2X345V678
      config trusted-list
         edit "FGT6HD391806070"
         next
         edit "S248DF3X17000482"
            set action deny
         next
      end
   end
```


Netskope Exhibit 1017

```
end
```

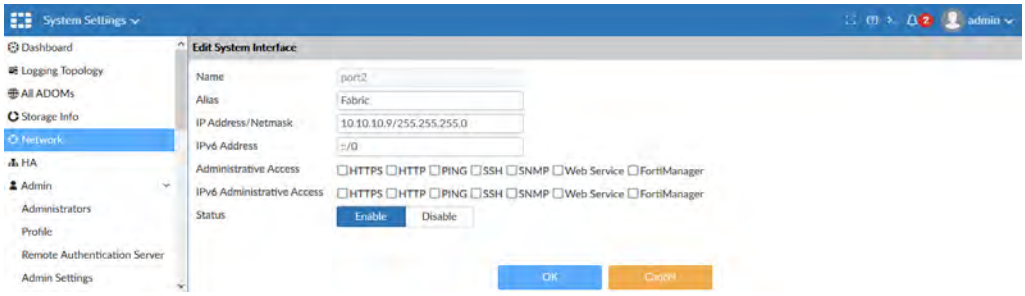
Configuring FortiAnalyzer

FortiAnalyzer is a required component for the Security Fabric. It allows the Security Fabric to show historical data for the Security Fabric topology and logs for the entire Security Fabric.

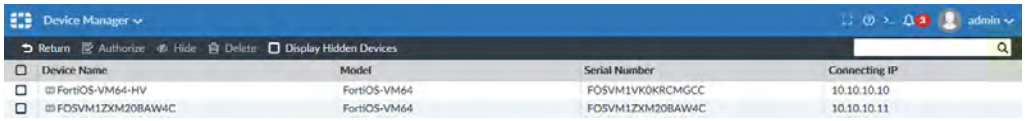
For more information about using FortiAnalyzer, see the FortiAnalyzer Administration Guide.

To connect a FortiAnalyzer to the Security Fabric:

1. Enable *FortiAnalyzer Logging* on the root FortiGate. See Configure the root FortiGate on page 116.
2. On the FortiAnalyzer, go to *System Settings* > *Network* and click *All Interfaces*.
3. Edit the port that connects to the root FortiGate.
4. Set the *IP Address/Netmask* to the IP address that is used for the Security Fabric on the root FortiGate.



5. Click *OK*.
   If the FortiGates have already been configured, it will now be listed as an unauthorized device.
6. Go to *Device Manager* > *Devices Unauthorized*. The unauthorized FortiGate devices are listed.



7. Select the root FortiGate and downstream FortiGate devices in the list, then click *Authorize*. The *Authorize Device* page opens.
8. Click *OK* to authorize the selected devices.



9. On the FortiGate devices, go to *Security Fabric* > *Fabric Connectors* and double-click the *FortiAnalyzer Logging* card. The page will now show the ADOM on the FortiAnalyzer that the FortiGate is in, and the storage, analytics, and archive usage.

Configuring other Security Fabric devices

This section contains information about configuring the following devices as part of the Fortinet Security Fabric:

- FortiGate Cloud on page 121
- FortiAnalyzer Cloud service on page 123
- FortiManager on page 127
- FortiManager Cloud service on page 128
- FortiSandbox on page 130
- FortiClient EMS on page 132
- FortiNAC on page 140
- FortiAP and FortiSwitch on page 142
- Additional devices on page 143

Prerequisites

- FortiGate devices must either have VDOMs disabled or be running in split-task VDOM mode in order to be added to the Security Fabric. See Virtual Domains on page 610.
- FortiGate devices must be operating in NAT mode.

FortiGate Cloud

FortiGate Cloud is a hosted security management and log retention service for FortiGate devices. It provides centralized reporting, traffic analysis, configuration management, and log retention without the need for additional hardware or software.

FortiGate Cloud offers a wide range of features:

- **Simplified central management**
  FortiGate Cloud provides a central GUI to manage individual or aggregated FortiGate and FortiWiFi devices. Adding a device to the FortiGate Cloud management subscription is straightforward. FortiGate Cloud has detailed traffic and application visibility across the whole network.
- **Hosted log retention with large default storage allocated**
  Log retention is an integral part of any security and compliance program, but administering a separate storage system is onerous. FortiGate Cloud takes care of this automatically and stores the valuable log information in the cloud. Different types of logs can be stored, including Traffic, System Events, Web, Applications, and Security Events.
- **Monitoring and alerting in real time**
  Network availability is critical to a good end-user experience. FortiGate Cloud enables you to monitor your FortiGate network in real time with different alerting mechanisms to pinpoint potential issues. Alerting mechanisms can be delivered via email.
- **Customized or pre-configured reporting and analysis tools**
  Reporting and analysis are your eyes and ears into your network's health and security. Pre-configured reports are available, as well as custom reports that can be tailored to your specific reporting and compliance requirements. The reports can be emailed as PDFs, and can cover different time periods.
- **Maintain important configuration information uniformly**
  The correct configuration of the devices within your network is essential for maintaining optimum performance and security posture. In addition, maintaining the correct firmware (operating system) level allows you to take advantage of the latest features.
- **Service security**
  All communication (including log information) between the devices and the cloud is encrypted. Redundant data centers are always used to give the service high availability. Operational security measures have been put in place to make sure your data is secure — only you can view or retrieve it.

Netskope Exhibit 1017

Registration and activation



Before you can activate a FortiGate Cloud account, you must first register your device.

FortiGate Cloud accounts can be registered manually through the FortiGate Cloud website, https://www.forticloud.com, or you can easily register and activate your account directly from your FortiGate.

To activate your FortiGate Cloud account:

1. On your device, go to *Dashboard* > *Status*.
2. In the *FortiGate Cloud* widget, click the *Not Activated* > *Activate* button in the *Status* field.
3. A pane will open asking you to register your FortiGate Cloud account. Click *Create Account*, enter your information, view and accept the terms and conditions, and then click *OK*.
4. A second dialogue window open , asking you to enter your information to confirm your account. This sends a confirmation email to your registered email. The dashboard widget then updates to show that confirmation is required.
5. Open your email, and follow the confirmation link it contains.
   A FortiGate Cloud page will open, stating that your account has been confirmed. The *Activation Pending* message on the dashboard will change to state the type of account you have, and will provide a link to the FortiGate Cloud portal.

Enabling logging to FortiGate Cloud

To enable logging to FortiGate Cloud:

1. Go to *Security Fabric* > *Fabric Connectors* > *Cloud Logging* or *Log & Report* > *Log Settings*.
2. Enable *Cloud Logging*.
3. Select an upload option: *Realtime*, *Every Minute*, or *Every 5 Minutes* (default).
4. Click *Apply*.

Logging into the FortiGate Cloud portal

Once logging has been configured and you have registered your account, you can log into the FortiGate Cloud portal and begin viewing your logging results. There are two methods to reach the FortiGate Cloud portal:

- If you have direct network access to the FortiGate:
  a. Go to *Dashboard* > *Status*.
  b. In the *FortiGate Cloud* widget, in the *Status* field, click *Activated* > *Launch Portal*, or, in the *Licenses* widget, click *FortiCare Support* > *Launch Portal*.
- If you do not have access to the FortiGate's interface, visit the FortiGate Cloud website (https://www.forticloud.com) and log in remotely, using your email and password. It will ask you to confirm the FortiGate Cloud account you are connecting to and then you will be granted access.

Cloud sandboxing

FortiGate Cloud can be used for automated sample tracking, or sandboxing, for files from a FortiGate. This allows suspicious files to be sent to be inspected without risking network security. If the file exhibits risky behavior, or is found to

Netskope Exhibit 1017

contain a virus, a new virus signature is created and added to the FortiGuard antivirus signature database.

To configure cloud sandboxing:

1. Go to *Security Fabric* > *Fabric Connectors* and double-click the *FortiSandbox* card.
2. For status, click *Enable*.
3. Set the *Type* to *FortiSandbox Cloud*.



By default, the *FortiSandbox Cloud* option is not visible. See Feature visibility on page 694 for instructions on making it visible.

4. Select the *FortiSandbox cloud region*.
5. Click *OK*.
   Sandboxing results are shown on the *Sandbox* tab in the FortiGate Cloud portal.

For more information about FortiGate Cloud, see the FortiGate Cloud documentation.

FortiAnalyzer Cloud service

FortiGate supports the FortiAnalyzer Cloud service for event logging.



Traffic and security logs are not supported in the initial version of FortiAnalyzer Cloud.

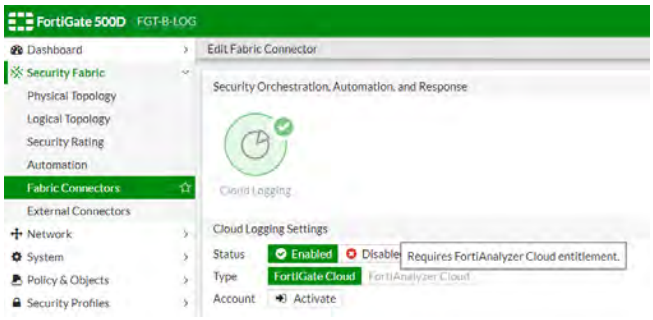
When FortiAnalyzer Cloud is licensed and enabled (see Deploying FortiAnalyzer Cloud for more information), all event logs are sent to FortiAnalyzer Cloud by default. All traffic logs, security logs, and archive files are not sent to FortiAnalyzer Cloud.

FortiAnalyzer Cloud differs from FortiAnalyzer in the following ways:

- You cannot enable FortiAnalyzer Cloud in `vdom override-setting` when global FortiAnalyzer Cloud is disabled.
- You must use the CLI to retrieve and display logs sent to FortiAnalyzer Cloud. The FortiOS GUI is not supported.
- You cannot enable FortiAnalyzer Cloud and FortiGate Cloud at the same time.

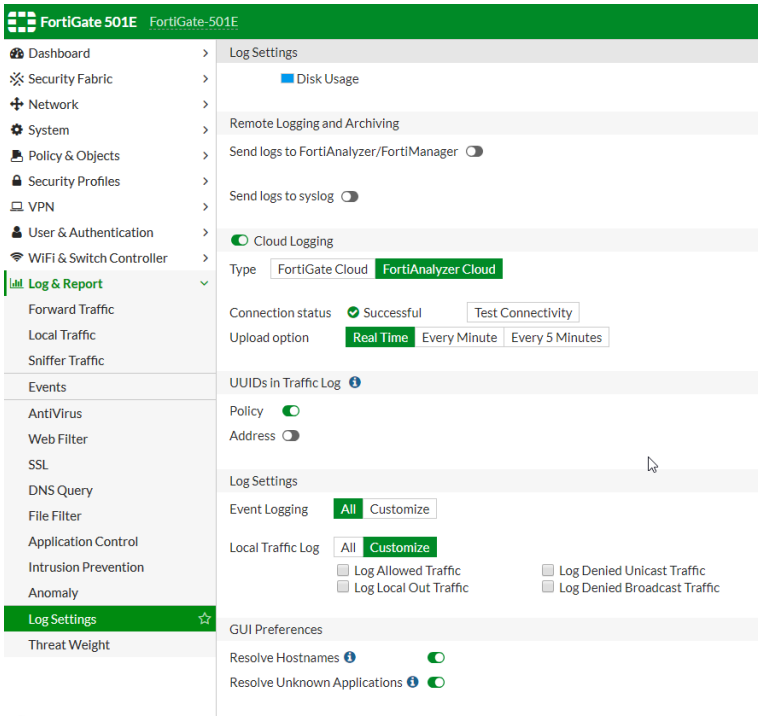
Sample settings panes

In the FortiOS *Security Fabric* > *Fabric Connectors* > *Cloud Logging* card settings page, *FortiAnalyzer Cloud* is grayed out when you do not have a FortiAnalyzer Cloud entitlement.



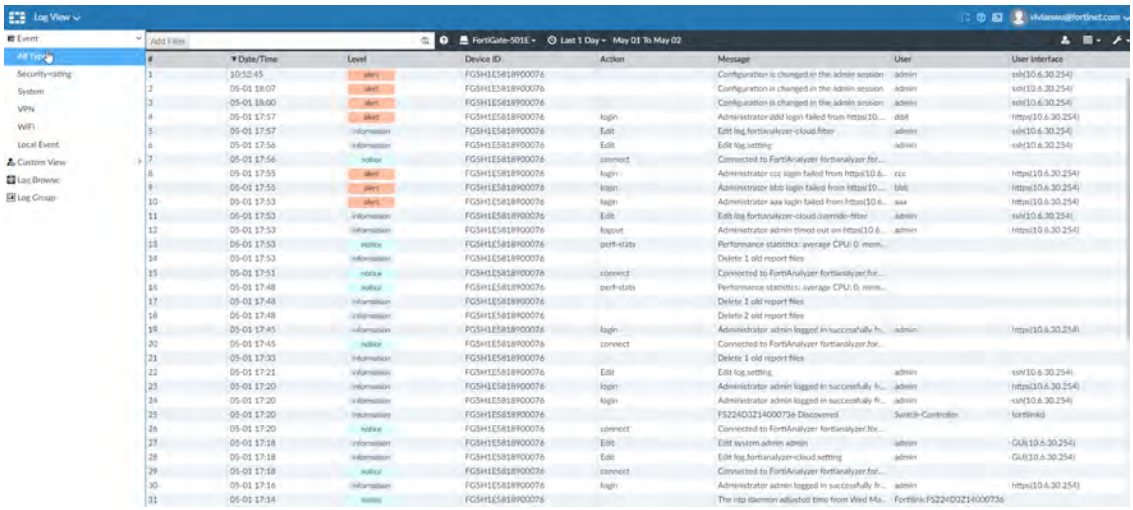
When you have a FortiAnalyzer Cloud entitlement, *FortiAnalyzer Cloud* is available.

You can also view the FortiAnalyzer Cloud settings in the *Log & Report* > *Log Settings* pane.



In FortiAnalyzer Cloud, you can view logs from *FortiOS* in the *Event* > *All Types* pane.

Netskope Exhibit 1017



To enable fortianalyzer-cloud using the CLI:

```
config log fortianalyzer-cloud setting
   set status enable
   set ips-archive disable
   set access-config enable
   set enc-algorithm high
   set ssl-min-proto-version default
   set conn-timeout 10
   set monitor-keepalive-period 5
   set monitor-failure-retry-period 5
   set certificate ''
   set source-ip ''
   set upload-option realtime
end
config log fortianalyzer-cloud filter
   set severity information
   set forward-traffic disable
   set local-traffic disable
   set multicast-traffic disable
   set sniffer-traffic disable
   set anomaly disable
   set voip disable
   set dlp-archive disable
   set dns disable
   set ssh disable
   set ssl disable
   set cifs disable
   set filter ''
   set filter-type include
end
```

To disable fortianalyzer-cloud for a specific VDOM using the CLI:

```
config log setting
   set faz-override enable
end
config log fortianalyzer-cloud override-setting
   set status disable
```


Netskope Exhibit 1017

```
end
```

To set fortianalyzer-cloud filter for a specific vdom using the CLI:

```
config log setting
   set faz-override enable
end
config log fortianalyzer-cloud override-setting
   set status enable
end
config log fortianalyzer-cloud override-filter
   set severity information
   set forward-traffic disable
   set local-traffic disable
   set multicast-traffic disable
   set sniffer-traffic disable
   set anomaly disable
   set voip disable
   set dlp-archive disable
   set dns disable
   set ssh disable
   set ssl disable
   set cifs disable
   set filter ''
   set filter-type include
end
```

To display fortianalyzer-cloud log using the CLI:

```
execute log filter device fortianalyzer-cloud
execute log filter category event
execute log display
```

Sample log

```
date=2019-05-01 time=17:57:45 idseq=60796052214644736 bid=100926 dvid=1027 itime="2019-05-01
     17:57:48" euid=3 epid=3 dsteuid=0 dstepid=3 logver=602000890 logid=0100032002
     type="event" subtype="system" level="alert" srcip=10.6.30.254 dstip=10.6.30.9
     action="login" msg="Administrator ddd login failed from https(10.6.30.254) because of
     invalid user name" logdesc="Admin login failed" sn="0" user="ddd" ui="https
     (10.6.30.254)" status="failed" reason="name_invalid" method="https"
     eventtime=1556758666274548325 devid="FG5H1E5818900076" vd="root" dtime="2019-05-01
     17:57:45" itime_t=1556758668 devname="FortiGate-501E"
date=2019-05-01 time=17:57:21 idseq=60796052214644736 bid=100926 dvid=1027 itime="2019-05-01
     17:57:23" euid=3 epid=3 dsteuid=0 dstepid=3 logver=602000890 logid=0100044546
     type="event" subtype="system" level="information" action="Edit" msg="Edit
     log.fortianalyzer-cloud.filter " logdesc="Attribute configured" user="admin" ui="ssh
     (10.6.30.254)" cfgtid=164757536 cfgpath="log.fortianalyzer-cloud.filter"
     cfgattr="severity[information->critical]" eventtime=1556758642413367644
     devid="FG5H1E5818900076" vd="root" dtime="2019-05-01 17:57:21" itime_t=1556758643
     devname="FortiGate-501E"
```


Netskope Exhibit 1017

FortiManager

When a FortiManager device is added to the Security Fabric, it automatically synchronizes with any connected downstream devices.

To add a FortiManager to the Security Fabric, configure it on the root FortiGate. The root FortiGate then pushes this configuration to downstream FortiGate devices. The FortiManager provides remote management of FortiGate devices over TCP port 541. The FortiManager must have internet access for it to join the Security Fabric.

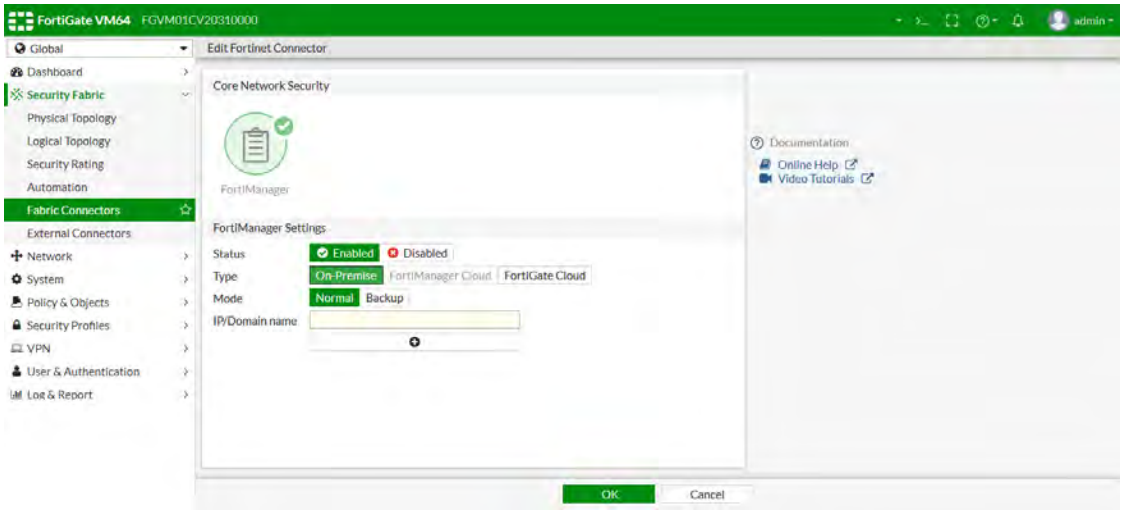
Once configured, the FortiGate can receive antivirus and IPS updates, and allows remote management through FortiManager or the FortiGate Cloud service. The FortiGate management option must be enabled so that the FortiGate can accept management updates to its firmware and FortiGuard services.

To add a FortiManager to the Security Fabric using the CLI:

```
config system central-management
   set type fortimanager
   set fmg {<IP_address> | <FQDN_address>}
end
```

To add a FortiManager to the Security Fabric using the GUI:

1. On the root FortiGate, go to *Security Fabric* > *Fabric Connectors* and double-click the *FortiManager* card.
2. For *Status*, click *Enable*.
3. For *Type*, click *On-Premise*.



4. Enter the *IP/Domain Name* of the FortiManager.
5. Click *OK*.
6. On the FortiManager, go to *Device Manager* and find the FortiGate in the *Unauthorized Devices* list.
7. Select the FortiGate device or devices, and click *Authorize* in the toolbar.
8. In the *Authorize Device* pop-up, adjust the device names as needed, then click *OK*.
   For more information about using FortiManager, see the FortiManager Administration Guide.

FortiManager Cloud service

This cloud-based SaaS management service is available through FortiManager. This service is included in FortiCloud accounts with a FortiManager Cloud account level subscription (ALCI).

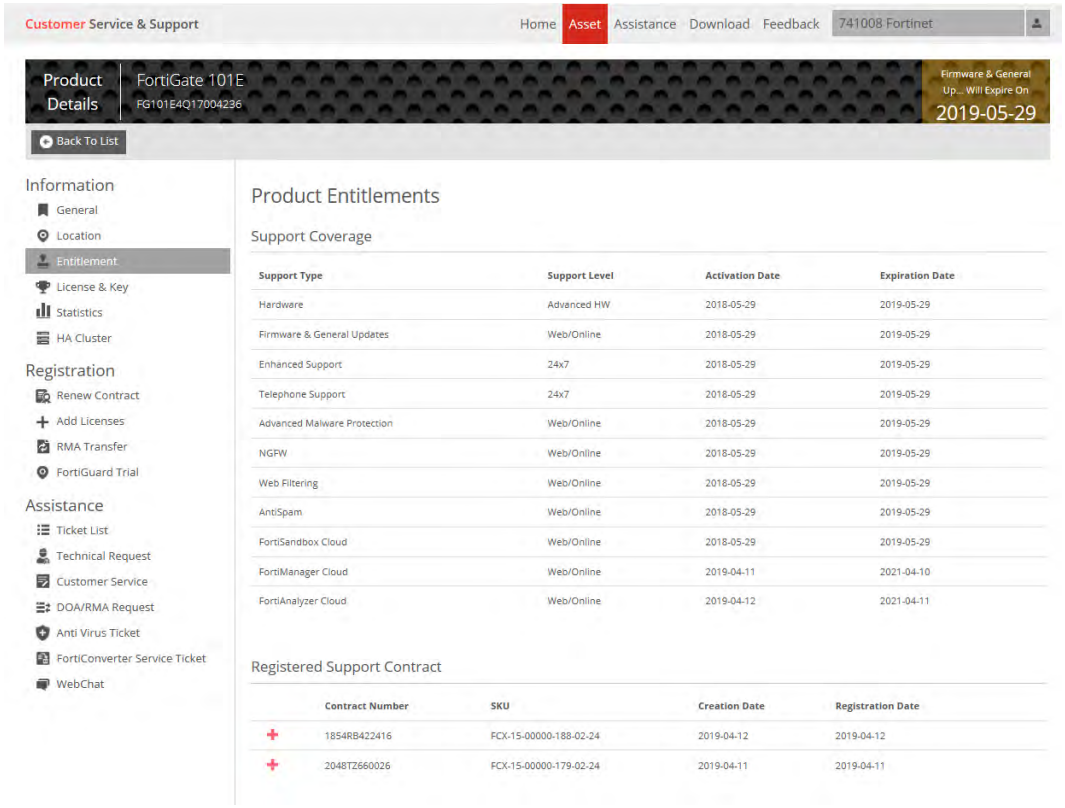
Once the FortiGate has acquired a contract named *FortiManager Cloud*, FortiCloud creates a cloud-based FortiManager instance under the user account. You can launch the portal for the cloud-based FortiManager from FortiCloud, and its URL starts with the User ID.

You can use a FortiGate with a contract for *FortiManager Cloud* to configure central management by using the FQDN of *fortimanager.forticloud.com*. A FortiGate-FortiManager tunnel is established between FortiGate and the FortiManager instance.

After the tunnel is established, you can execute FortiManager functions from the cloud-based FortiManager portal.

To manage FortiGate with FortiManager Cloud:

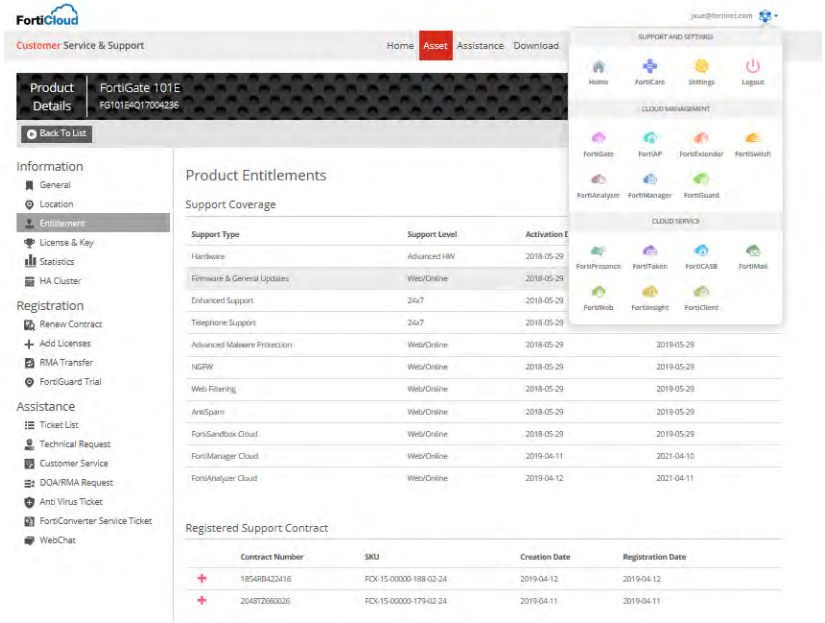
1. In the FortiCare portal, ensure you have a product entitlement for *FortiManager Cloud*.
   a. Go to *Asset* > *Information* > *Entitlement*.
   b. Take note of your user ID number in the top-right corner.



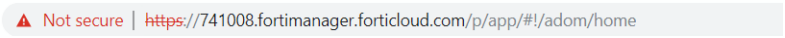
2. Click the *FortiCloud* icon in the upper-right corner.

Netskope Exhibit 1017

3. Under *CLOUD MANAGEMENT*, select *FortiManager*.



The FortiManager portal opens. The URL incorporates the user ID for the dedicated instance.



4. In FortiOS, enable FortiManager Cloud.
   a. Go to *Security Fabric* > *Fabric Connectors* and double-click the *FortiManager* card.
   b. For *Status*, click *Enable*.
   c. For *Type*, click *FortiManager Cloud*.
   d. Click *OK*.

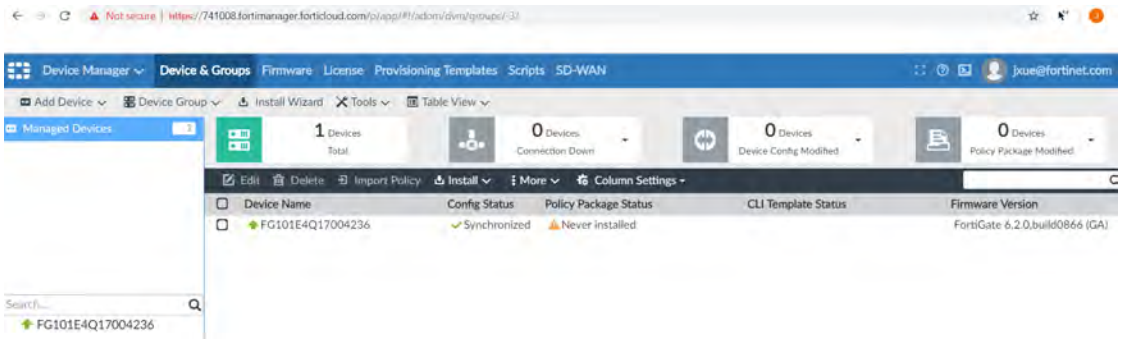


The *FortiManager Cloud* button can only be selected if you have a FortiManager Cloud product entitlement.

5. In the FortiManager Cloud instance, go to *Device Manager* and authorize the FortiGate. See Authorizing devices for more information.
   When using FortiGate to enable FortiManager Cloud, the FortiGate appears as an unauthorized device.



After authorizing the FortiGate, it becomes a managed device.



In FortiOS, the *Security Fabric* > *Fabric Connectors* page now displays green arrow in the *FortiManager* card because FortiManager Cloud is registered.

FortiSandbox

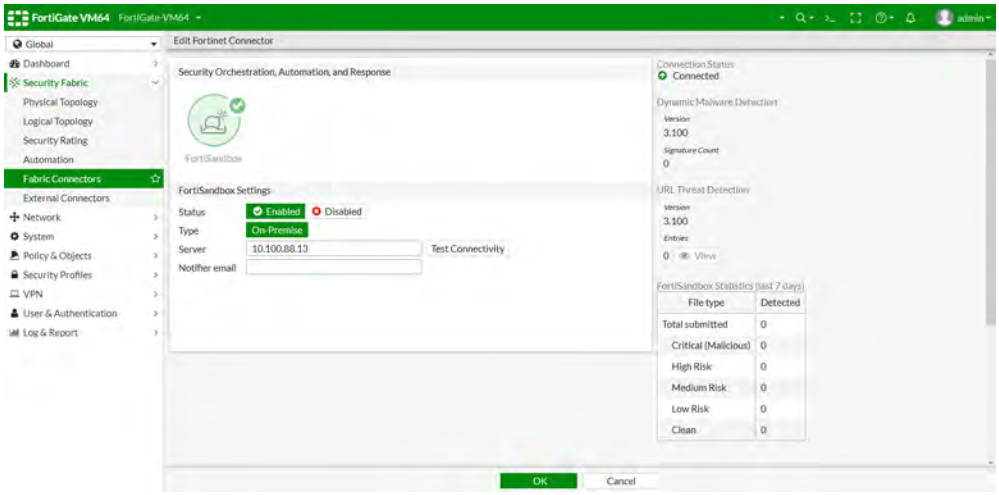
The Security Fabric supports FortiSandbox appliances and FortiSandbox Cloud. A FortiGate Cloud account is not required.

To use FortiSandbox in a Security Fabric, connect the FortiSandbox to the Security Fabric, then configure an antivirus profile to send files to the FortiSandbox. Sandbox inspection can also be used in web filter profiles.

FortiSandbox settings are configured on the root FortiGate of the Security Fabric. After configuration, the root FortiGate pushes the settings to other FortiGate devices in the Security Fabric.

To add a FortiSandbox appliance to the Security Fabric:

1. On the root FortiGate, go to *Security Fabric* > *Fabric Connectors* and double-click the *FortiSandbox* card.
2. For status, click *Enable*.
3. For *Type*, click *On-Premise*.
4. In the *Server* field, enter the FortiSandbox device's IP address.



5. Optionally, enter a *Notifier email*.
6. Click *OK*.
7. On the FortiSandbox device, go to *Scan Input* > *Device*.
8. Edit the root FortiGate.
9. Under *Permissions*, check the *Authorized* box.

Netskope Exhibit 1017

10. Click *OK*.
11. Authorize the rest of the FortiGate devices that are in the Security Fabric.

To add a FortiSandbox cloud instance to the Security Fabric:

1. On the root FortiGate, go to *Security Fabric* > *Fabric Connectors* and double-click the *FortiSandbox* card.
2. For status, click *Enable*.
3. For *Type*, click *FortiSandbox Cloud*.
4. Select the *FortiSandbox cloud region* from the dropdown list. Data from your network will only be sent to servers in the selected region.
5. Click *OK*.

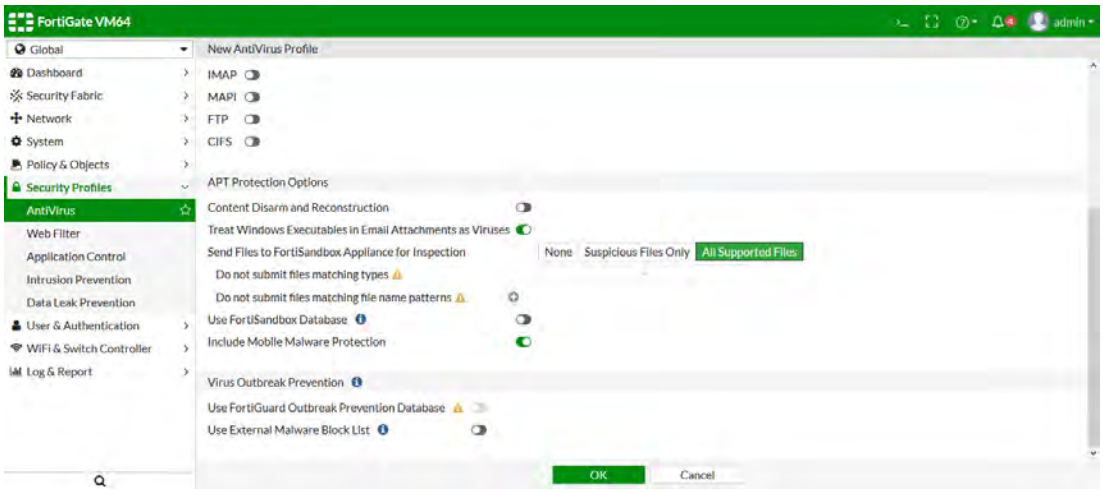
If *FortiSandbox Cloud* is not visible in the GUI, run the `execute forticloud-sandbox region` and `execute forticloud-sandbox update` commands.

Antivirus profiles

An antivirus profile must be configured to send files to the FortiSandbox.

To configure an antivirus profile:

1. On the FortiGate, go to *Security Profile* > *AntiVirus*.
2. Create, edit, or clone an antivirus profile.



3. Under *APT Protection Options*, set *Send Files to FortiSandbox Appliance for Inspection* to *All Supported Files*.
4. Optionally, configure file exceptions.
5. Enable *Use FortiSandbox Database*.
6. Click *OK*.

Web Filter profiles

Sandbox inspection can be used in Web Filter profiles.

To configure a web filter profile:

1. On the FortiGate, go to *Security Profiles* > *Web Filter*.
2. Create, edit, or clone a profile.
3. Under *Static URL Filter*, enable *Block malicious URLs discovered by FortiSandbox*.
4. Click *OK*.

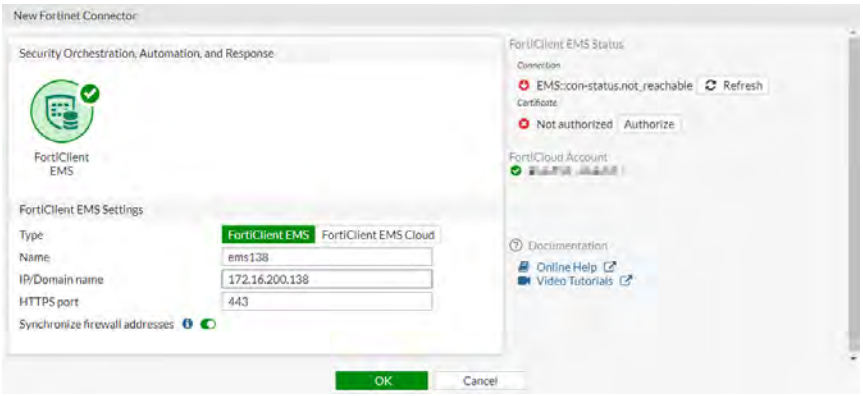
FortiClient EMS

The FortiGate Security Fabric root device can link to FortiClient Endpoint Management System (EMS) and FortiClient EMS Cloud (a cloud-based EMS solution) for endpoint connectors and automation. Up to three EMS servers can be added to the Security Fabric, including a FortiClient EMS Cloud server. EMS settings are synchronized between all fabric members.

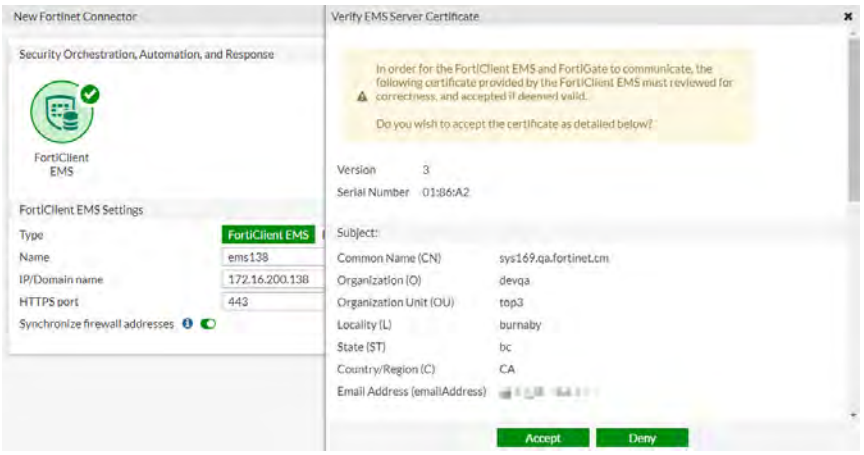
To enable cloud-based EMS services, the FortiGate must be registered to FortiCloud with an appropriate user account. The following examples presume the EMS certificate has already been configured.

To add an on-premise FortiClient EMS server to the Security Fabric in the GUI:

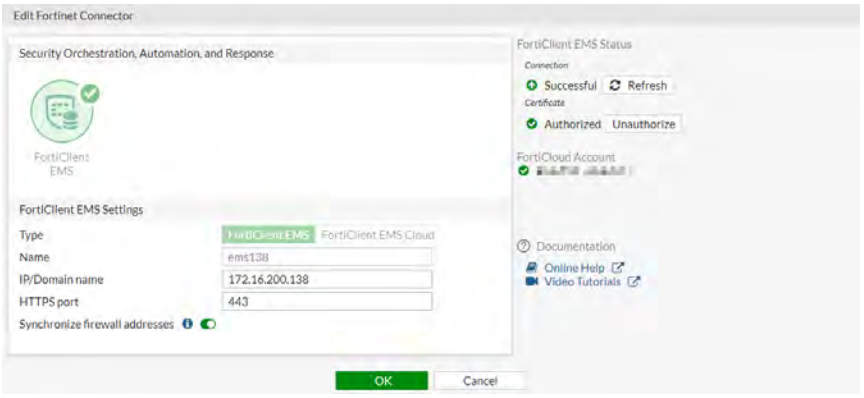
1. On the root FortiGate, go to *System* > *Feature Visibility* and enable *Endpoint Control*.
2. Go to *Security Fabric* > *Fabric Connectors*.
3. Click *Create New* and click *FortiClient EMS*.
4. For *Type*, click *FortiClient EMS*.
5. Enter a name and IP address.
6. Click *OK*.



A window appears to verify the EMS server certificate:

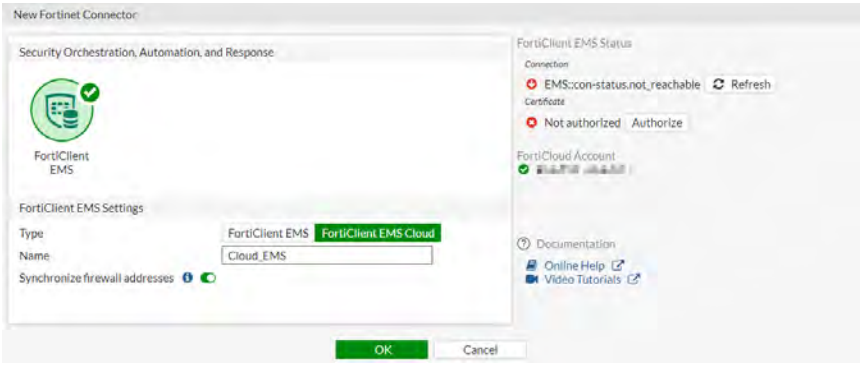


7. Click *Accept*.
   The *FortiClient EMS Status* section displays a *Successful* connection and an *Authorized* certificate:



To add a FortiClient EMS Cloud server to the Security Fabric in the GUI:

1. Go to *Security Fabric* > *Fabric Connectors*.
2. Click *Create New* and click *FortiClient EMS*.
3. For *Type*, click *FortiClient EMS Cloud*.
4. Enter a name.
5. Click *OK*.



   A window appears to verify the EMS server certificate.
6. Click *Accept*.
   The *FortiClient EMS Status* section displays a *Successful* connection and an *Authorized* certificate.

To test connectivity with the EMS server:

1. Go to *Security Fabric* > *Fabric Connectors* and double-click the *FortiClient EMS* or *FortiClient EMS Cloud* card.
2. In the *FortiClient EMS Status* section under *Connection*, click *Refresh*.

To add an on-premise FortiClient EMS server to the Security Fabric in the CLI:

```
config endpoint-control fctems
    edit <ems_name>
        set server <ip_address>
        set certificate <string>
        set https-port <integer>
        set source-ip <ip_address>
    next
end
```

The `https-port` is the EMS HTTPS access port number, and the `source-ip` is the REST API call source IP address.

To add a FortiClient EMS Cloud server to the Security Fabric in the CLI:

```
config endpoint-control fctems
    edit <name>
        set fortinetone-cloud-authentication enable
        set certificate <string>
    next
end
```

To verify an EMS certificate in the CLI:

```
# execute fctems verify ems137

        Subject:     C = CA, ST = bc, L = burnaby, O = devqa, OU = top3, CN =
sys169.qa.fortinet.cm, emailAddress = xxxx@xxxxxxxx.xxx
        Issuer:      CN = 155-sub1.fortinet.com
        Valid from:  2017-12-05 00:37:57  GMT
        Valid to:    2027-12-02 18:08:13  GMT
        Fingerprint: D3:7A:1B:84:CC:B7:5C:F0:A5:73:3D:BB:ED:21:F2:E0
        Root CA:     No
        Version:     3
        Serial Num:
                01:86:a2
        Extensions:
                Name:     X509v3 Basic Constraints
                Critical: yes
                Content:
                CA:FALSE

                Name:     X509v3 Subject Key Identifier
                Critical: no
                Content:
                35:B0:E2:62:AF:9A:7A:E6:A6:8E:AD:CB:A4:CF:4D:7A:DE:27:39:A4

                Name:     X509v3 Authority Key Identifier
                Critical: no
                Content:
```


Netskope Exhibit 1017

```
                keyid:66:54:0F:78:78:91:F2:E4:08:BB:80:2C:F6:BC:01:8E:3F:47:43:B1

DirName:/C=CA/ST=bc/L=burnaby/O=devqa/OU=top3/CN=fac155.fortinet.com/emailAddress=xyguo@fort
inet.com
serial:01:86:A4

                Name:     X509v3 Subject Alternative Name
                Critical: no
                Content:
                DNS:sys169.qa.fortinet.cm

                Name:     X509v3 Key Usage
                Critical: no
                Content:
                Digital Signature, Non Repudiation, Key Encipherment, Data Encipherment, Key
Agreement, Certificate Sign, CRL Sign, Encipher Only, Decipher Only

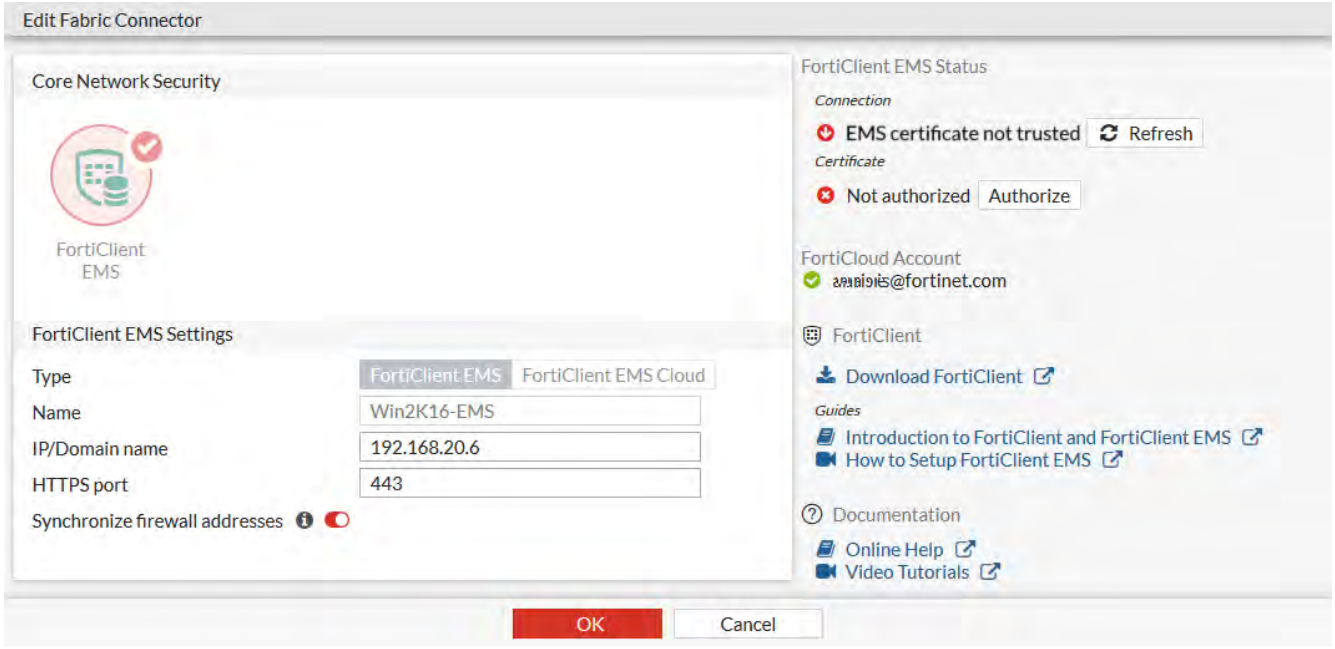
                Name:     X509v3 Extended Key Usage
                Critical: no
                Content:
                TLS Web Server Authentication, TLS Web Client Authentication

EMS configuration needs user to confirm server certificate.
Do you wish to add the above certificate to trusted remote certificates? (y/n)y
```

Troubleshooting

Certificate not trusted

When configuring a new connection to an EMS server, the certificate might not be trusted.



When you click *Authorize*, a warning is shown: *The server certificate cannot be authenticated with installed CA certificates. Please install its CA certificates on this FortiGate*.

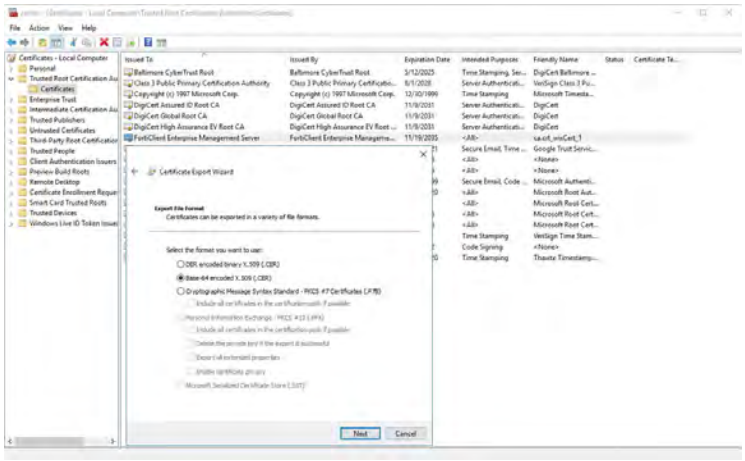
In the CLI, an error message is shown when you try to verify the certificate:

```
# execute fctems verify Win2K16-EMS
certificate not configured/verified: 2
Could not verify server certificate based on current certificate authorities.
Error 1--92-60-0 in get SN call: EMS Certificate is not signed by a known CA.
```

The default FortiClient EMS certificate that is used for the SDN connection is signed by the CA certificate that is saved on the Windows server when FortiClient EMS is first installed. It can be manually exported and installed on the FortiGate.

To manually export and install the certificate on to the FortiGate:

1. Export the EMS certificate on the server that EMS is installed on:
   a. On the Windows server that EMS is installed on, go to *Settings* > *Manage computer certificates*.
   b. In the certificate management module, go to *Trusted Root Certification Authorities* > *Certificates*.
   c. Right click on the certificate issued by FortiClient Enterprise Management Server and select *All Tasks* > *Export*.
   d. The *Certificate Export Wizard* opens. Click *Next*.
   e. Select *Base-64 encoded X.509*, then click *Next*.



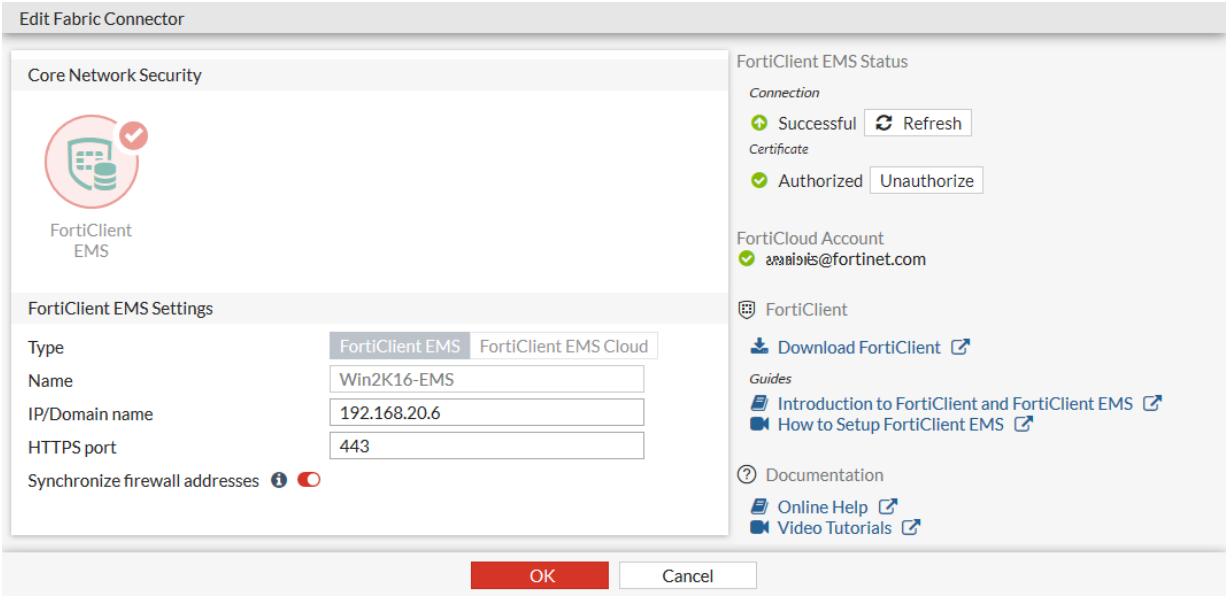
   f. Enter a file name for the certificate and click *Browse* to select the folder where it will be located, then click *Next*.
   g. Review the settings, then click *Finish*. The certificate is downloaded to the specified folder.
2. On the FortiGate, import the certificate:
   a. Go to *System* > *Certificate*. By default, the *Certificate* option is not visible, see Feature visibility on page 694 for information.
   b. Click *Import* > *CA Certificate*.
   c. Set *Type* to *File*, and click *Upload* to import the certificate from the management computer.
   d. Click *OK*. The imported certificate is shown in the *Remote CA Certificate* section of the certificate table.
3. Try to authorize the certificate on the FortiGate:
   a. Go to *Security Fabric* > *Fabric Connectors* and edit the FortiClient EMS connector. The connection status should now say that the certificate is not authorized.
   b. Click *Authorize*. The following warning is shown:



The warning can also be seen in the CLI:

```
# execute fctems verify Win2K16-EMS
failure in certificate configuration/verification: -4
Could not verify EMS. Error 1--94-0-401 in get SN call: Authentication denied.
```

4. Authorize the FortiGate on EMS:
   a. Log in to the EMS server console and go to *Administration* > *Fabric Devices*.
   b. Select the serial number of the FortiGate device, then click *Authorize*.
5. Try to authorize the certificate on the FortiGate again:
   a. On the FortiGate, go to *Security Fabric* > *Fabric Connectors* and edit the FortiClient EMS connector.
   b. Click *Authorize*.
   c. When presented with the EMS server certificate, click *Accept* to accept the certificate.
   Your connection should now be successful and authorized.



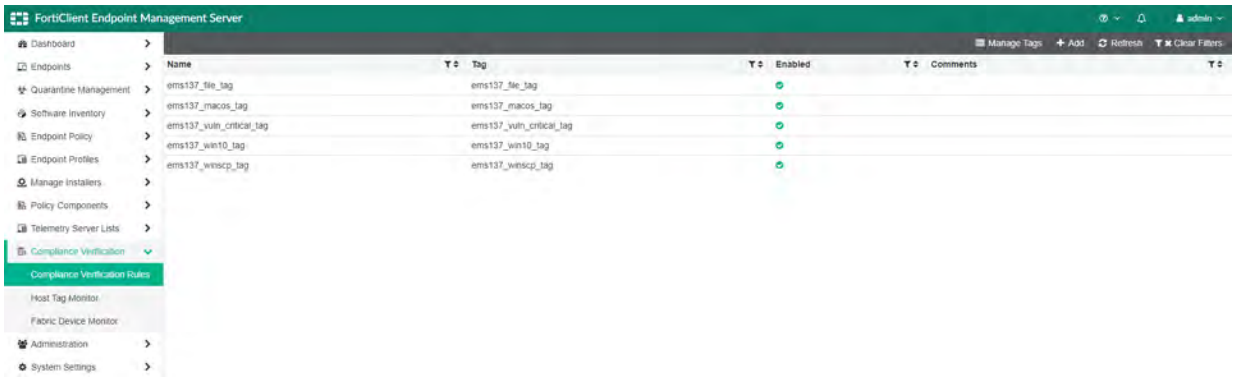
   d. Click *OK*.

Synchronizing FortiClient EMS tags and configurations

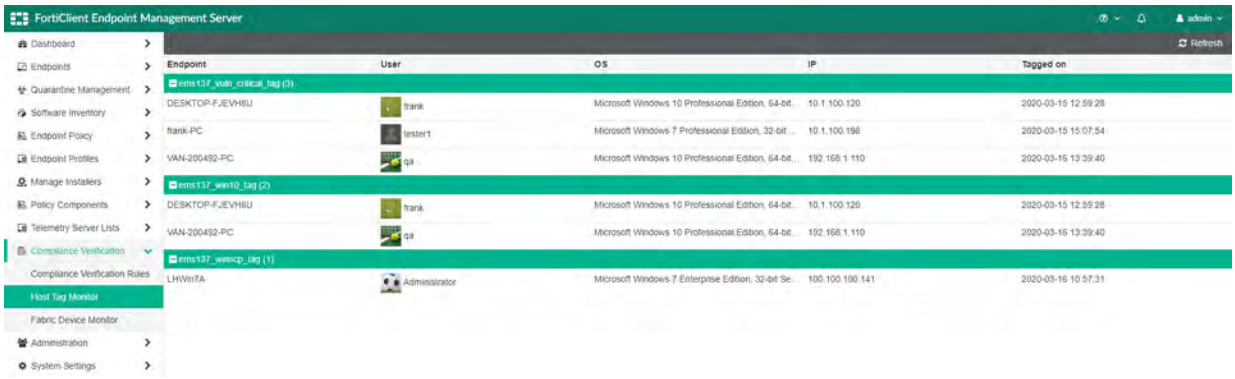
A new option under the FortiClient EMS settings consolidates the setup of EMS connectors to support EMS tags. EMS tags are pulled and automatically synced with the EMS server. They are converted into read-only dynamic firewall addresses that can be used in firewall policies, routing, and so on.

These examples presume the following have been configured in FortiClient EMS:

- Tags have been created on the *Compliance Verification* > *Compliance Verification Rules* page.

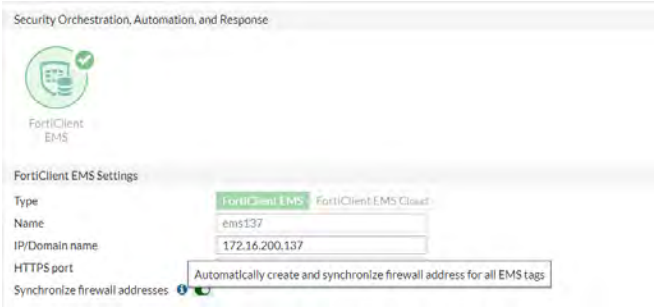


- There are registered users who match the defined tags that are visible on the *Compliance Verification* > *Host Tag Monitor* page.



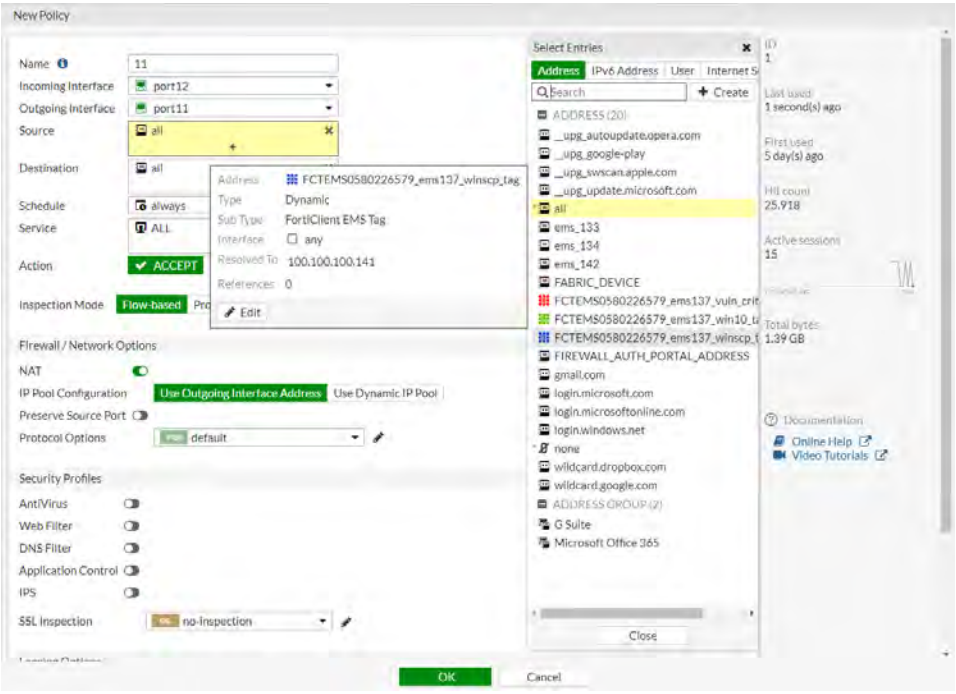
To configure FortiClient EMS with tag synchronization in the GUI:

1. Configure the EMS Fabric Connector:
   a. On the root FortiGate, go to *Security Fabric* > *Fabric Connectors*.
   b. Click *Create New* and click *FortiClient EMS*.
   c. Enable *Synchronize firewall addresses*.



   d. Configure the other settings as needed and validate the certificate.
   e. Click *OK*.
2. Go to *Policy & Objects* > *Addresses* and hover over the EMS tag to view which IPs it resolves to.

3. Configure a firewall policy:
   a. Go to *Policy & Objects* > *Firewall Policy* and create a new policy.
   b. For the *Source Address*, add the EMS tag dynamic address.



   c. Configure the other settings as needed.
   d. Click *OK*.

To configure FortiClient EMS with tag synchronization in the CLI:

1. Configure the EMS Fabric Connector:

```
config endpoint-control fctems
    edit "ems137"
        set fortinetone-cloud-authentication disable
        set server "172.16.200.137"
        set https-port 443
        set source-ip 0.0.0.0
        set pull-sysinfo enable
        set pull-vulnerabilities enable
        set pull-avatars enable
        set pull-tags enable
        set call-timeout 5000
        set certificate "REMOTE_Cert_1"
    next
end
```

2. Verify which IPs the dynamic firewall address resolves to:

```
# diagnose firewall dynamic list
List all dynamic addresses:
FCTEMS0580226579_ems137_vuln_critical_tag: ID(118)
        ADDR(10.1.100.120)
        ADDR(10.1.100.198)
```

```
FCTEMS0580226579_ems137_winscp_tag: ID(155)
        ADDR(100.100.100.141)

FCTEMS0580226579_ems137_win10_tag: ID(182)
        ADDR(10.1.100.120)

# diagnose firewall dynamic address FCTEMS0580226579_ems137_vuln_critical_tag
FCTEMS0580226579_ems137_vuln_critical_tag: ID(118)
        ADDR(10.1.100.120)
        ADDR(10.1.100.198)

Total dynamic list entries: 1.
Total dynamic addresses: 2
Total dynamic ranges: 0
```

3. Configure a firewall policy that uses the EMS tag dynamic firewall address as a source.

FortiNAC

A FortiNAC device can be added to the Security Fabric on the root FortiGate. After the device has been added and authorized, you can log in to the FortiNAC from the FortiGate topology views.



Adding a FortiNAC to the Security Fabric requires a FortiNAC with a license issued in the year 2020 or later that includes an additional certificate. The device cannot be added if it has an older license. Use the `licensetool` in the FortiNAC CLI to determine if your license includes the additional certificate.

To add a FortiNAC to the Security Fabric:

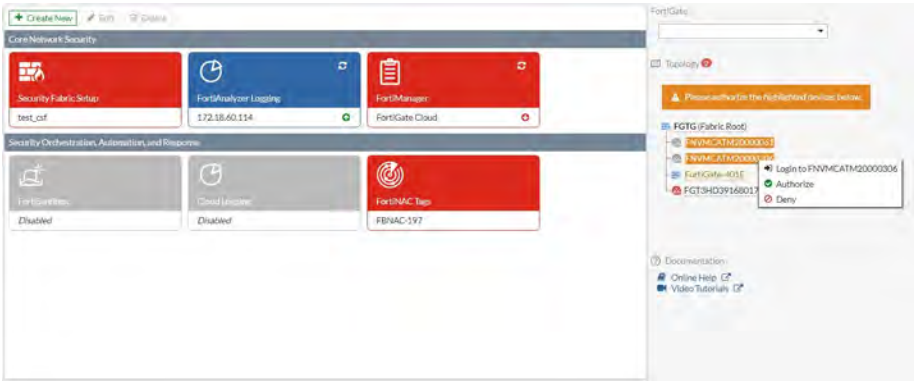
1. On the FortNAC, configure telemetry and input the IP address of the root FortiGate. See Security Fabric Connection in the *FortiNAC Administration Guide* for more information.
2. On the root FortiGate, authorize the FortiNAC.
3. Verify the connection status in the topology views.

To authorize the FortiNAC on the root FortiGate in the GUI:

1. Go to *Security Fabric* > *Fabric Connectors*.
2. The FortiNAC device will be highlighted in the topology list in the right panel with the status *Waiting for Authorization*.

Netskope Exhibit 1017

3. Click on the highlighted FortiNAC and select *Authorize*.



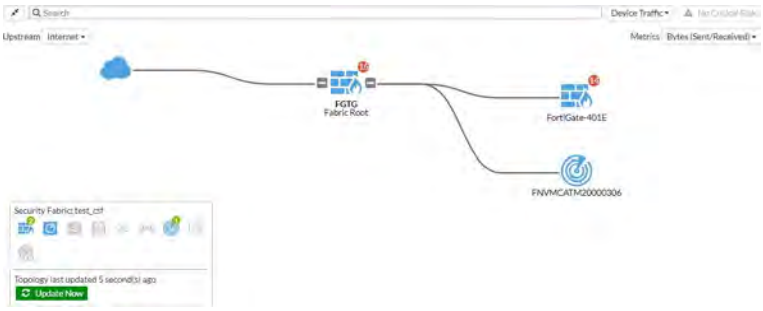
Optionally, you can also deny authorization to the FortiNAC to remove it from the list.

To authorize the FortiNAC on the root FortiGate in the CLI:

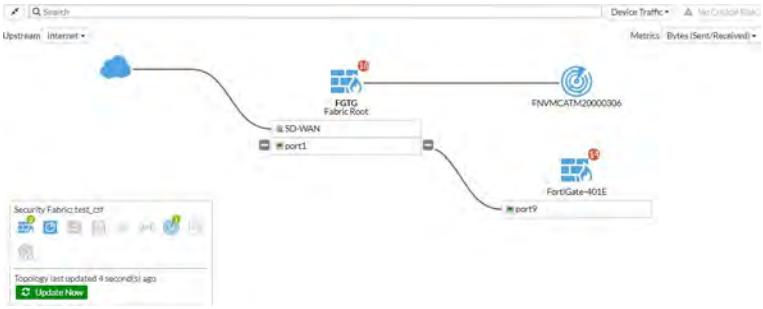
```
config system csf
    config trusted-list
        edit "FNVMCATM20000306"
            set action accept
        next
    end
end
```

To verify the connection status:

1. After the FortiNAC is authorized, go to *Security Fabric* > *Physical Topology* and confirm that it is included in the topology.



2. Go to *Security Fabric* > *Logical Topology* and confirm the FortiNAC is also displayed there.



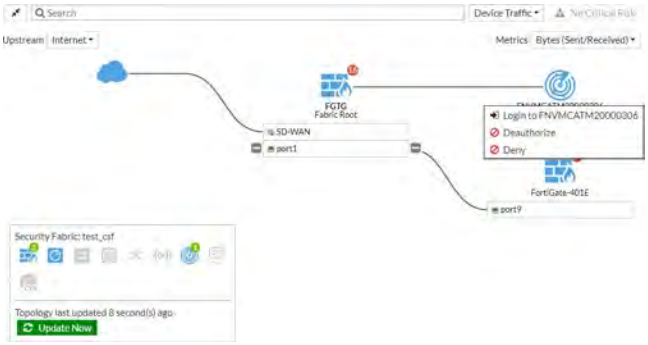
Netskope Exhibit 1017

3. Run the following command in the CLI to view information about the FortiNAC device's status:

```
# diagnose sys csf downstream-devices fortinac
{
  "path":"FG5H1E5818900126:FNVMCATM20000306",
  "mgmt_ip_str":"10.1.100.197",
  "mgmt_port":0,
  "admin_port":8443,
  "serial":"FNVMCATM20000306",
  "host_name":"adnac",
  "device_type":"fortinac",
  "upstream_intf":"port2",
  "upstream_serial":"FG5H1E5818900126",
  "is_discovered":true,
  "ip_str":"10.1.100.197",
  "downstream_intf":"eth0",
  "authorizer":"FG5H1E5818900126",
  "idx":1
}
```

To log in to the FortiNAC from the FortiGate:

1. On the FortiGate, go to *Security Fabric* > *Physical Topology* or *Security Fabric* > *Logical Topology*.
2. Click on the FortiNAC and select *Login to <serial_number>*.



A new tab will open to the FortiNAC log in page.

3. Enter the username and password to log in to the FortiNAC.

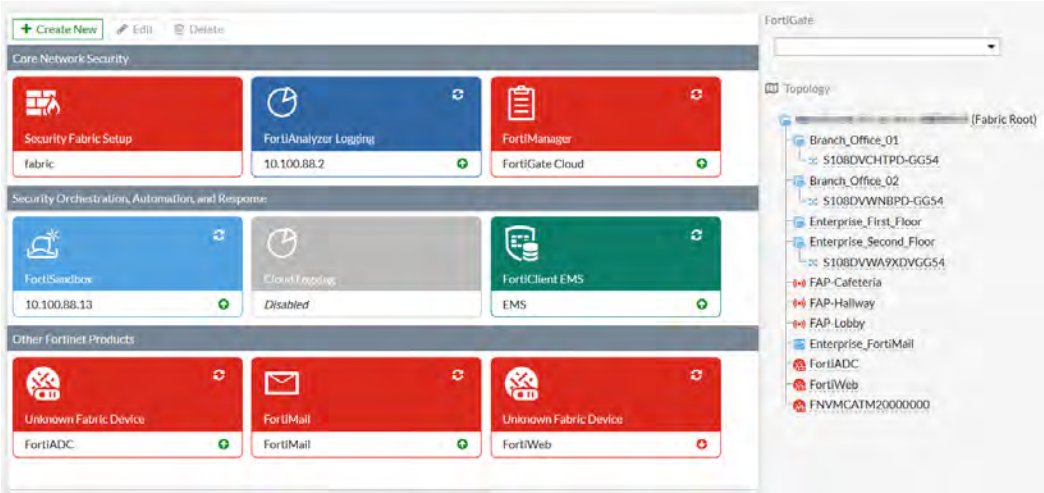
FortiAP and FortiSwitch

FortiAP and FortiSwitch devices can be authorized in the Security Fabric with one click. After connecting a FortiAP or FortiSwitch device to an authorized FortiGate, it will automatically be listed in the topology tree.

To authorize FortiAP and FortiSwitch devices:

1. Connect the FortiAP or FortiSwitch device to a FortiGate.
2. On the root FortiGate, go to *Security Fabric* > *Fabric Connectors*. The new device is shown in the *Topology* tree.

3. Click the device and select *Authorize*.



Additional devices

The following Fortinet devices are supported by the Security Fabric:

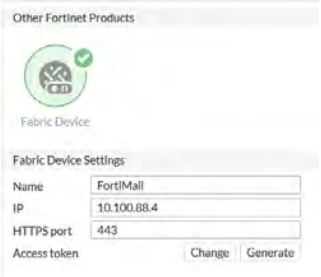
- FortiADC
- FortiDDoS
- FortiSandbox
- FortiMail
- FortiWeb
- FortiWLC

Security Fabric supports standalone FortiSandbox devices, FortiSandbox HA-Cluster primary, and FortiSandbox cluster IP.

In FortiOS, the device details can be shown in the *Security Fabric* and *Fabric Device* dashboard widgets, as well as the Fabric Connectors page, and physical and logical topologies.

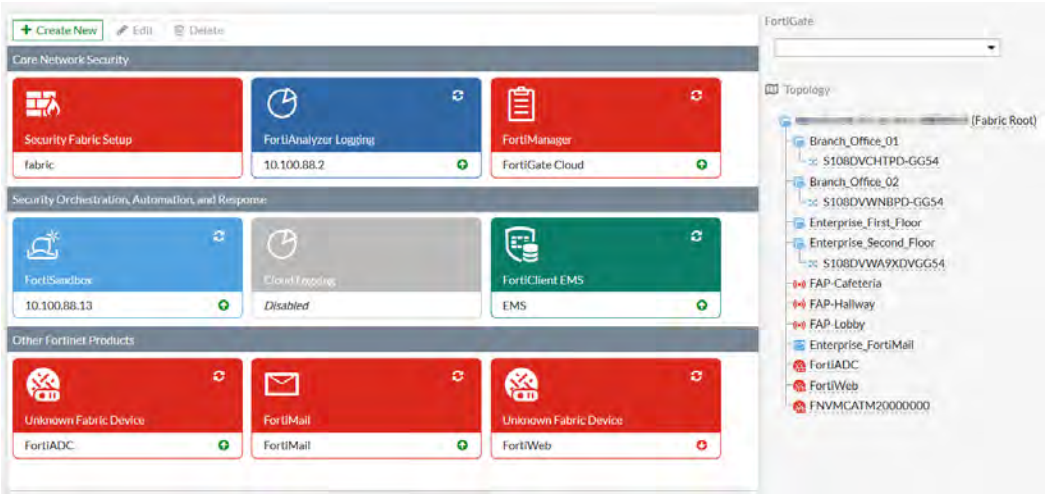
To add one or more of the devices to the Security Fabric in the GUI:

1. On the root FortiGate, go to *Security Fabric* > *Fabric Connectors*.
2. Click *Create New* and select *Fabric Device*.
3. Enter the *Name*, *IP*, *HTTPS port* for the device.



FortiSandbox only supports HTTPS port 443.

4. Click Generate to generate an access token. The *Generate Access Token* pane opens.
   a. Enter the device's username and password.
   b. Click *OK*.
5. Click *OK*.
6. Add more devices as required.
   The additional devices are shown on the *Fabric Connectors* page under *Other Fortinet Products* and in the *Topology* tree.



To add one or more of the devices to the Security Fabric in the CLI:

```
config system csf
    ...
    config fabric-device
        edit "FortiMail"
            set device-ip 10.100.88.4
            set access-token xxxxxx
        next
    end
end
```

Using the Security Fabric

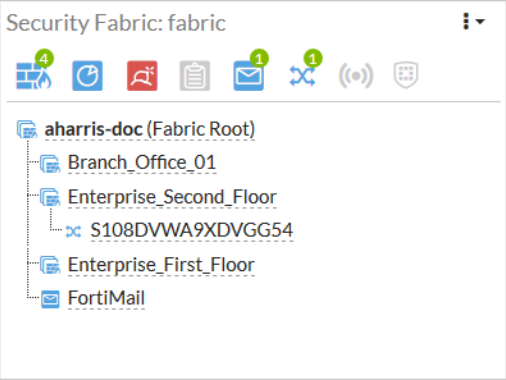
Dashboard widgets

Security Fabric widgets can be added to FortiGate dashboards, including:

- Security Fabric status on page 145
- Security Posture on page 145
- Fabric Device on page 146
- FortiGate Cloud on page 147

Security Fabric status

The Security Fabric status widget shows a summary of the devices in the Security Fabric.



Hover the cursor over the top icons to view pop-ups showing the statuses of the devices in the fabric.

The device tree shows devices that are connected, or could be connected, to you Security Fabric, according to the following color scheme:

- Blue: connected to the network
- Gray: not configured or not detected
- Red: no longer connected or not authorized

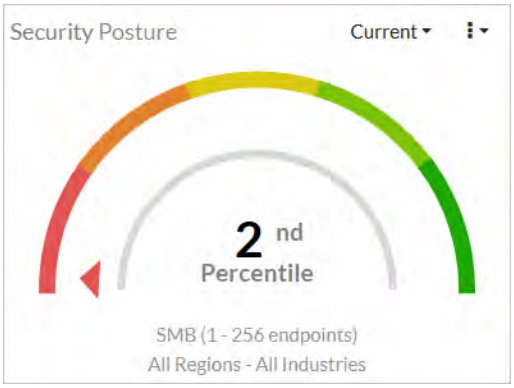
Hover over a device in the tree to view details about the device, such as it's serial number, operation mode, IP address, CPU and memory usage, and others, depending on the device type.

Unauthorized FortiAP and FortiSwitch devices are highlighted in the list, and can be authorized by clicking on the device name.

Security Posture

The Security Posture widget shows the security rating for your Security Fabric. It can show the current rating percentile, or historical security rating score or percentile charts.

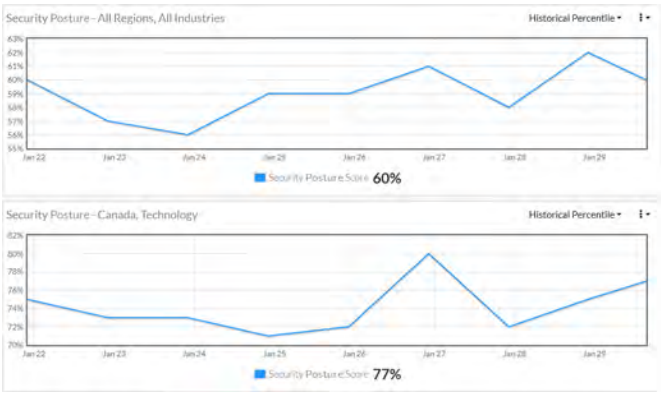
The widget can be configured to show how your organization's security rating compares to the ratings of either all organizations, or only organizations that are in the same industry and/or geographic region as you (determined from your FortiCare account settings).



To switch to the historical security rating score view, select *Historical Score* from the view dropdown menu. Making the widget wider makes the graph easier to read.



To switch to the historical percentile view, select *Historical Percentile* from the view dropdown menu. Making the widget wider makes the graph easier to read.



To receive a security rating score, all FortiGate devices that are in the Security Fabric must have a valid Security Rating License.

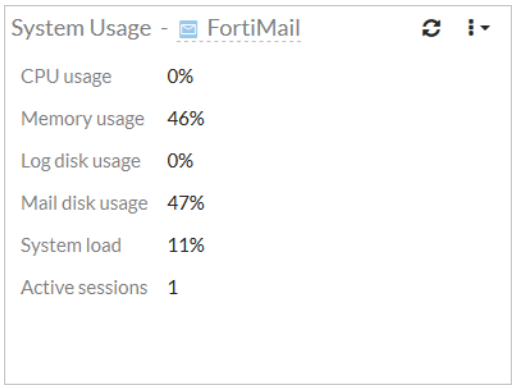
Fabric Device

The Fabric Device widget show statistics and system information about the selected fabric device.

For a FortiMail device, the widget can show:

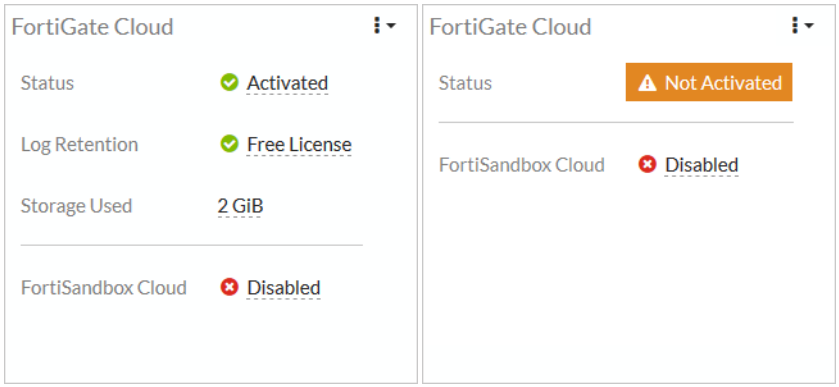
- Mail Statistics: a chart of the total messages and total spam messages over time.
- Statistics Summary: a pie chart summarizes mail statistics.
- System Information: The FortiMail System Information widget
- System Usage: System usage information, such as CPU, memory, and disk usage, as well as the number of active

sessions.



FortiGate Cloud

The FortiGate Cloud widget shows the FortiGate Cloud status and information. If your account is not activated, you can activate it from the widget.



To activate your FortiGate Cloud account:

1. Click on the *Not Activated* button and select *Activate*. The *Activate FortiGate Cloud* pane opens.
2. If you already have an account:
   a. Fill in your email address, password, country or region, and reseller.
   b. Click *OK*.
3. If you are creating an account:
   a. In the *FortiCloud* field select *Create Account*.
   b. Fill in all of the required information.
   c. Click *OK*.

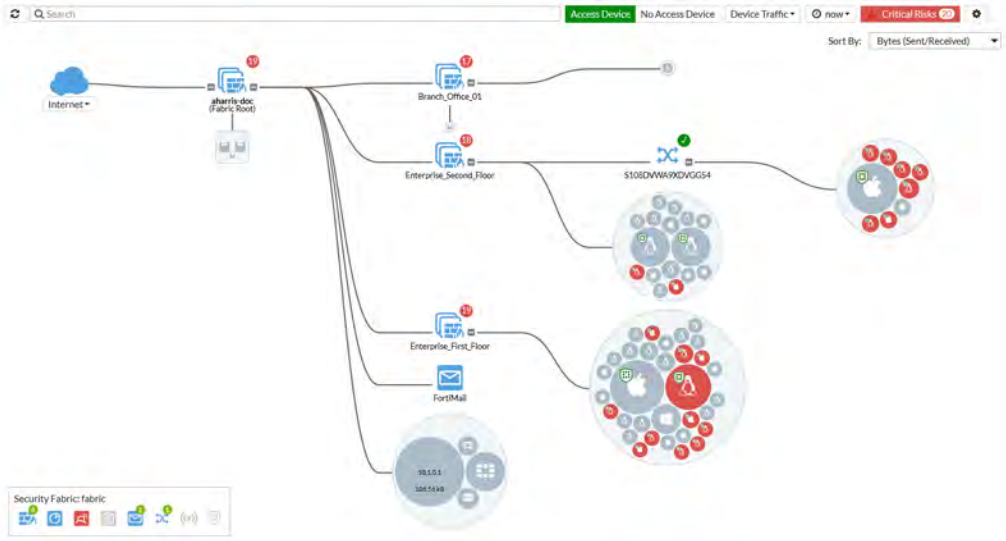
Topology

The full Security Fabric topology can be viewed on the root FortiGate. Downstream FortiGate devices' topology views do not include upstream devices.

The Physical Topology shows the physical structure of your network, including all connected devices and the connections between them. The Logical Topology shows information about the interfaces that connect devices to the Security Fabric. Only Fortinet devices are shown in the topologies.

In both topology pages, you can use filtering and sorting options to control the information that is shown. Hover the cursor over a device icon, port number, or endpoint to open a tooltip that shows information about that specific device, port, or endpoint. Right-click on a device to log in to it or to deauthorize it. Right-click on an endpoint to perform various tasks, including drilling down for more details on sources or compromised hosts, quarantining the host, and banning the IP address.

The small number that might be shown on the top right corner of a device icon is the number of security ratings recommendations or warnings for that device. The color of the circle shows the severity of the highest security rating check that failed. Clicking it opens the Security Rating page. See Security rating on page 195 for more information.



Servers and server clusters are represented by squares with rounded corners. They are grouped separately from circular endpoints. Devices are grouped by type and are colored based on their risk level. Endpoint groups are represented by donut charts or bubble packs depending on the current view settings (see Endpoint groups for more information). The size of the bubbles in the topology vary based on traffic volume.

AWS assets are grouped by AWS security groups or subnets, and information about detected Common Vulnerabilities and Exposures (CVEs), as well as the instance details and ID, are shown.

Views

The topology views can be focused using filters and by sorting in different ways to help you locate the information that you need.

Select one of *Access Device* or *No Access Device* to only show access or no access devices in the physical topology.

From the *Endpoint Option* dropdown list, select one of the following views:

- *Device Traffic*: Organize devices by traffic.
- *Device Count*: Organize devices by the number of devices connected to it.
- *Device Operating System*: Organize devices by operating system.
- *Device Hardware Vendor*: Organize devices by hardware vendor.

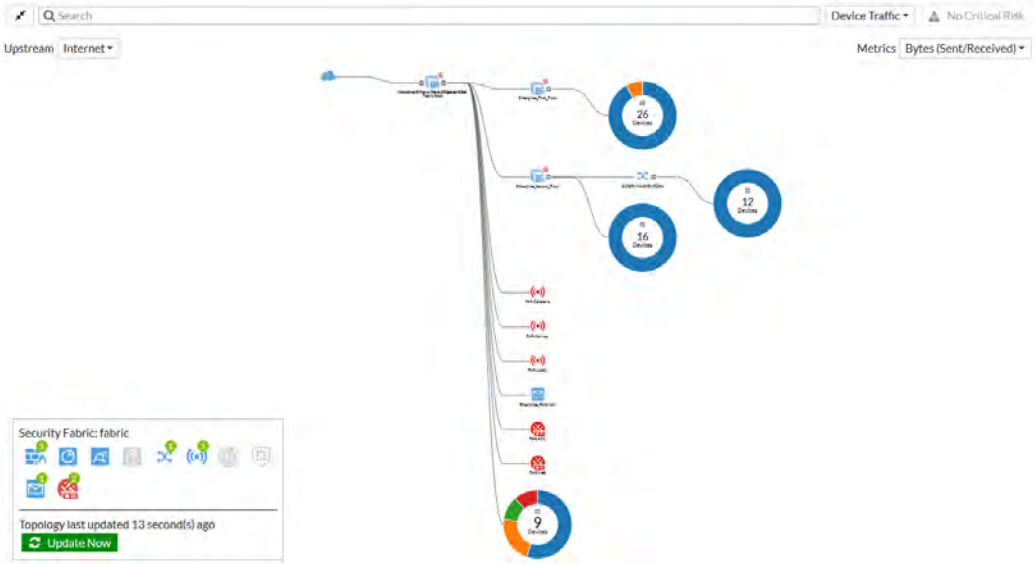
Netskope Exhibit 1017

- *Risk*: Only include devices that have endpoints with medium, high, or critical risk values of the specified type: *All*, *Compromised Host*, *Vulnerability*, or *Threat Score*.
- *No Devices*: Do not show endpoints.

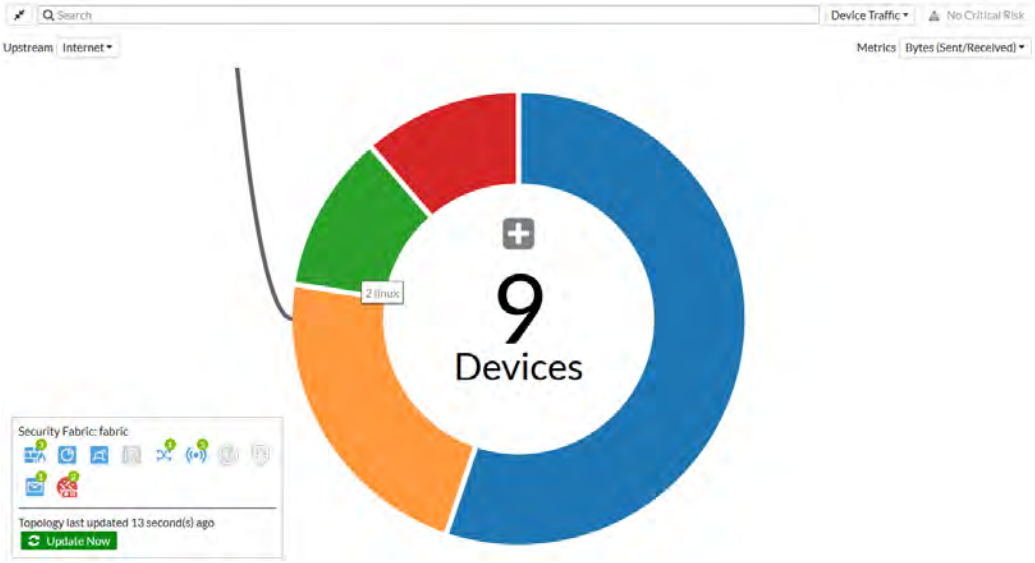
The time period dropdown list filters the view by time. Options include: *now* (real time), *5 minutes*, *1 hour*, *24 hours*, or *7 days*.

Endpoint groups

The Device Traffic and Device Count views display endpoint groups as donut charts, with the total number of endpoints in the group in the center of the chart. Each sector of the donut chart represents a different endpoint operating system.



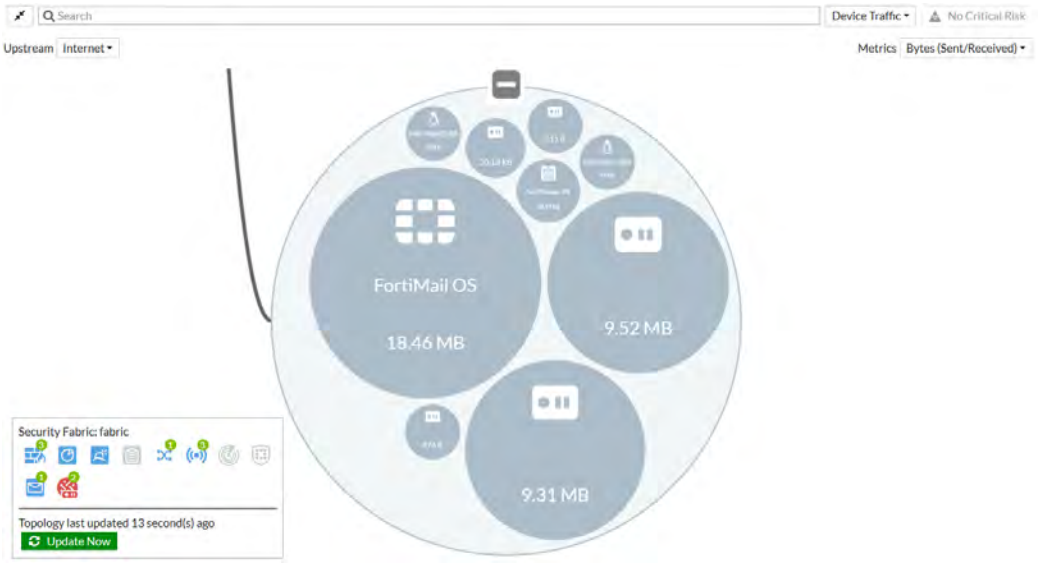
To zoom in on a donut chart, click any chart sector. Each sector represents a different endpoint OS. Hovering over each sector allows you to see the OS that the sector represents and the number of endpoints that have that OS installed.



In this example, the endpoint group contains a total of nine endpoints, with the following OSes installed:

Donut sector color	OS	Number of endpoints
Orange	Linux	2
Green	FortiMail	1
Red	FortiManager	1
Blue	Other	5

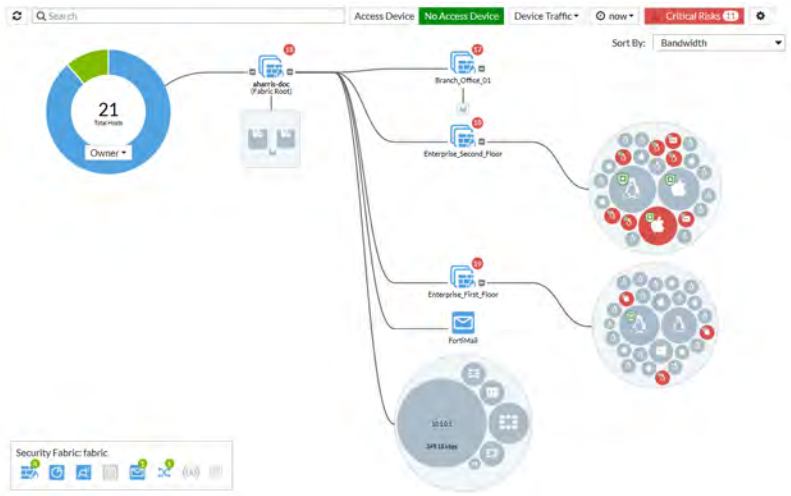
To view the endpoint group in a bubble pack display, click the + button in the center of the donut chart. You can view each individual endpoint in the bubble pack view.



WAN cloud

The WAN cloud icon includes a dropdown menu for selecting where the destination data comes from. The available options are: *Internet*, *Owner*, *IP Address*, and *Country/Region*. These options are only available when the filtering is based on *Device Traffic*.

When *Owner* is selected, the destination hosts are shown as donut charts that show the percentage of internal (with private IP addresses) and Internet hosts. Hover over either color in the chart to see additional information.



FortiAP and FortiSwitch devices

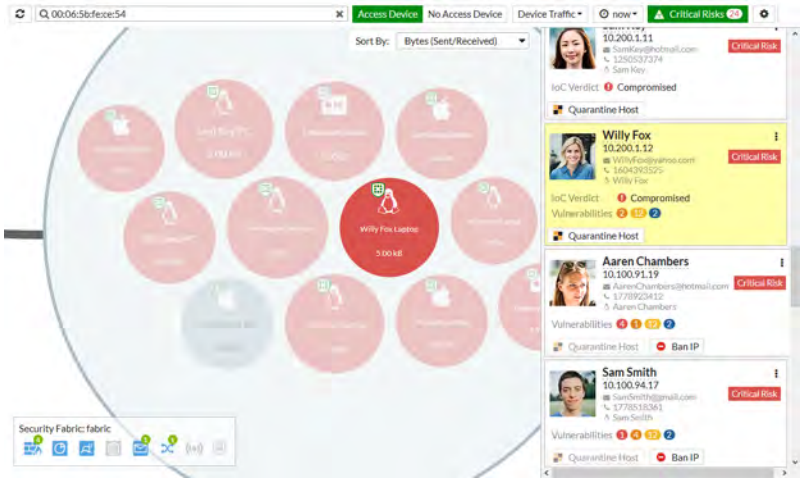
Newly discovered FortiAP and FortiSwitch devices are initial shown in the topologies with gray icons to indicate that they have not been authorized. To authorize a device, click on the device icon or name and select *Authorize*. Once authorized, the device icon will turn blue.

Right-click on an authorized FortiAP device to *Deauthorize* or *Restart* the device. Right-click on a FortiSwitch device to *Deauthorize*, *Restart*, or *Upgrade* the device, or to *Connect to the CLI*.

FortiAP and FortiSwitch links are enhanced to show link aggregation groups for the inter-switch link (ISL-LAG). To differentiate them from physical links, ISL-LAG links are shown with a thicker line. The endpoint circles can also be used as a reference to identify ISL-LAG groups that have more than two links.

Critical risks

Click the *Critical Risks* button to see a list of endpoints that are deemed critical risks, organized by threat severity. These are the red endpoints in the current topology view.



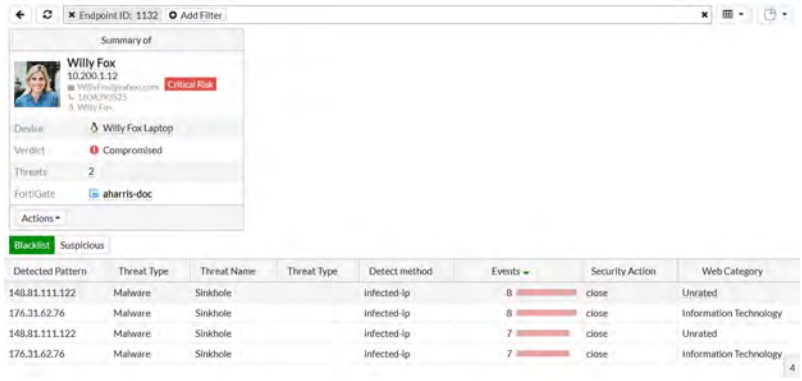
For each endpoint, the user's photo, name, IP address, email address, and phone number are shown. The number of vulnerabilities of each severity is shown, and if the IoC verdict is that the endpoint is compromised.

If applicable, the endpoint's host can be quarantined or their IP address banned, by clicking the *Quarantine Host* on *Ban IP* button.

The dropdown menu also provides options to drill down to more information on compromised hosts or endpoint vulnerabilities.



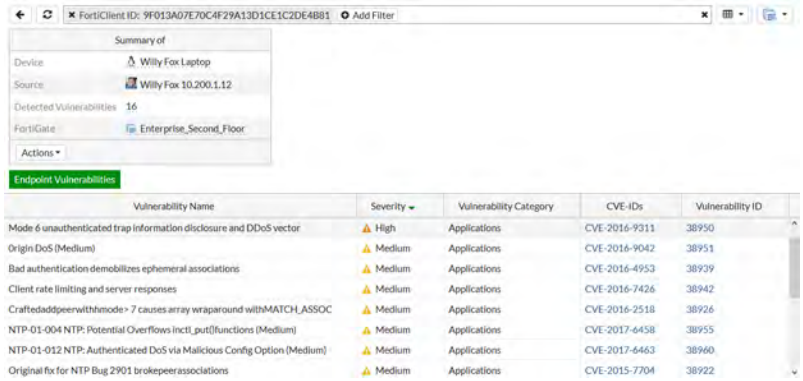
Click *Drill Down to Compromised Hosts* to open the *Top Compromised Hosts* page that shows a summary for the selected endpoint.



Compromised host information can also be viewed on the FortiAnalyzer in *SOC* > *FortiView* > *Threats* > *Compromised Hosts*.

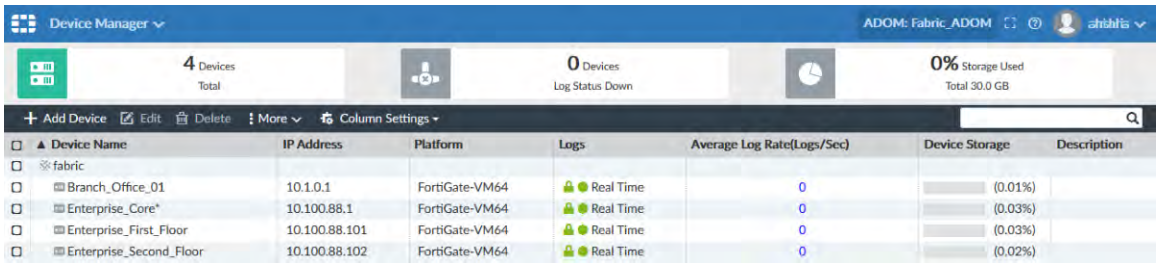
The FortiAnalyzer must have a FortiGuard Indicators of Compromise service license in order to see compromised hosts.

Click *Drill Down to Endpoint Vulnerability* to open the vulnerabilities page that shows a summary of the vulnerabilities on the selected endpoint.



FortiAnalyzer

The Security Fabric topology can also be seen on the FortiAnalyzer device. In the *Device Manager*, FortiGate devices are shown as part of a Security Fabric group with an asterisk next to the name of the root FortiGate.



To view the Security Fabric topology, right-click on the fabric group and select *Fabric Topology*. Only Fortinet devices are shown in the Security Fabric topology views.



Topology view — consolidated risk

The topology view shows endpoints based on their highest severity event.

In the default topology view, you can view hosts with critical vulnerabilities and compromised hosts identified as critical risks.

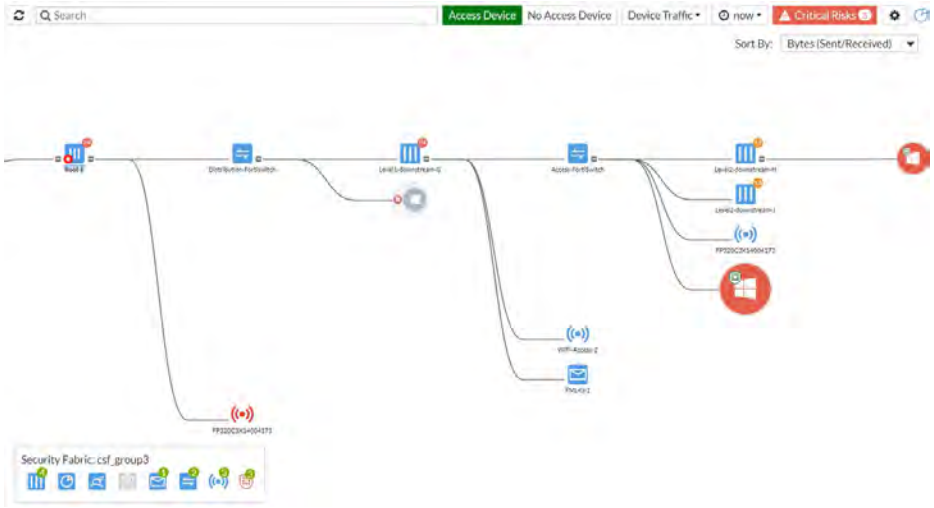
The consolidated *Risk* view mode displays different risks within the Security Fabric topology. You can use the *Risk* view mode to filter threats by *Compromised Hosts*, *Vulnerability*, and *Threat Score*.



To access the default topology view:

1. Go to *Security Fabric* > *Physical Topology*.
   The default topology view highlights hosts with critical vulnerabilities and compromised hosts as critical risks (three critical risks in the example).

a. Hover over the tooltips for more details.



2. To view the critical risk summary, click *Critical Risks*.
The *Critical Risks* pane displays on the right-side of the screen.



To access the consolidated Risk view mode:

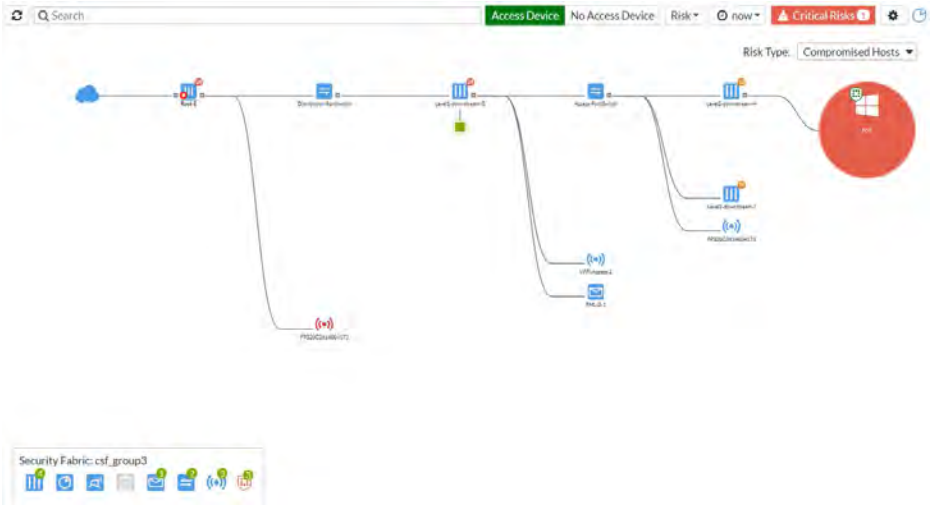
1. In the view option dropdown button, select *Risk*.



2. Select one of the following options from the *Risk Type* dropdown menu:
   - *All*



   - *Compromised Hosts*



   - *Vulnerability*



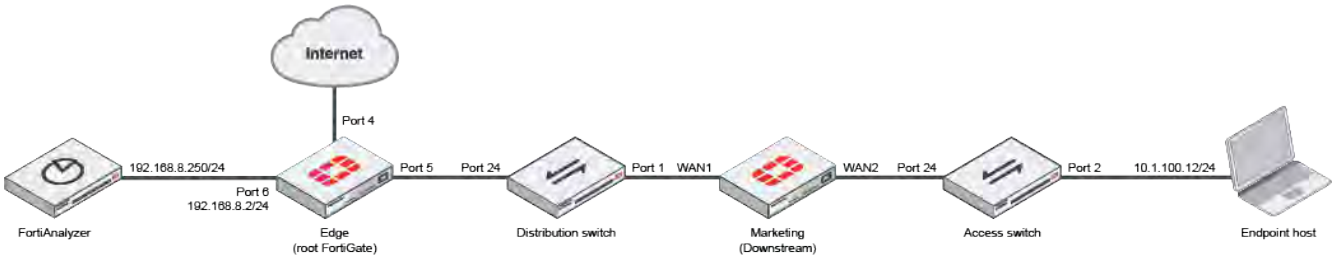
- *Threat Score*



Viewing and controlling network risks via topology view

This topic shows how to view and control compromised hosts via the *Security Fabric* > *Physical Topology* or *Security Fabric* > *Logical Topology* view.

In the following topology, the downstream FortiGate (Marketing) is connected to the root FortiGate (Edge) through a FortiSwitch (Distribution). The Endpoint Host is connected to the downstream FortiGate (Marketing) through another FortiSwitch (Access).



This topic consists of the following steps:

1. View the compromised endpoint host.
2. Quarantine the compromised endpoint host.
3. Run diagnose commands.

To view the compromised endpoint host:

1. Test that FortiGate detects a compromised endpoint host by opening a browser on the endpoint host and entering a malicious website URL. The browser displays a *Web Page Blocked!* warning and does not allow access to the website.
2. In FortiOS on the root FortiGate, go to *Security Fabric* > *Physical Topology*. The endpoint host, connected to the Access FortiSwitch, is highlighted in red. Mouse over the endpoint host to view a tooltip that shows the IoC verdict. The endpoint host is compromised.



3. Go to *Security Fabric* > *Logical Topology*. The endpoint host, connected to the downstream FortiGate, is highlighted in red. Mouse over the endpoint host to view a tooltip that shows the IoC verdict. The endpoint host is compromised.



To quarantine the compromised endpoint host:

1. In FortiOS on the root FortiGate, go to *Security Fabric* > *Physical Topology*.
2. Right-click the endpoint host and select *Quarantine Host*. Click *OK* to confirm the confirmation dialog.
3. Go to *Monitor* > *Quarantine Monitor*. From the dropdown list at the top right corner, select *All FortiGates*. The quarantined endpoint host displays in the content pane.
4. On the endpoint host, open a browser and visit a website such as https://www.fortinet.com/. If the website cannot be accessed, this confirms that the endpoint host is quarantined.

To run diagnose commands:

1. To show the downstream FortiGate after it joins the Security Fabric, run the `diagnose sys csf downstream` command in the root FortiGate (Edge) CLI. The output should resemble the following:

```
Edge # diagnose sys csf downstream
1: FG101ETK18002187 (192.168.7.3) Management-IP: 0.0.0.0 Management-port:0 parent:
     FG201ETK18902514
path:FG201ETK18902514:FG101ETK18002187
data received: Y downstream intf:wan1 upstream intf:vlan70 admin-port:443
authorizer:FG201ETK18902514
```

2. To show the upstream FortiGate after the downstream FortiGate joins the Security Fabric, run the `diagnose sys csf upstream` command in the downstream FortiGate (Marketing) CLI. The output should resemble the following:

```
Marketing # diagnose sys csf upstream
Upstream Information:
Serial Number:FG201ETK18902514
IP:192.168.7.2
Connecting interface:wan1
Connection status:Authorized
```

3. To show the quarantined endpoint host in the connected FortiGate, run the following commands in the downstream FortiGate (Marketing) CLI:

```
Marketing # show user quarantine
config user quarantine
   config targets
      edit "PC2"
         set description "Manually quarantined"
         config macs
            edit 00:0c:29:3d:89:39
```

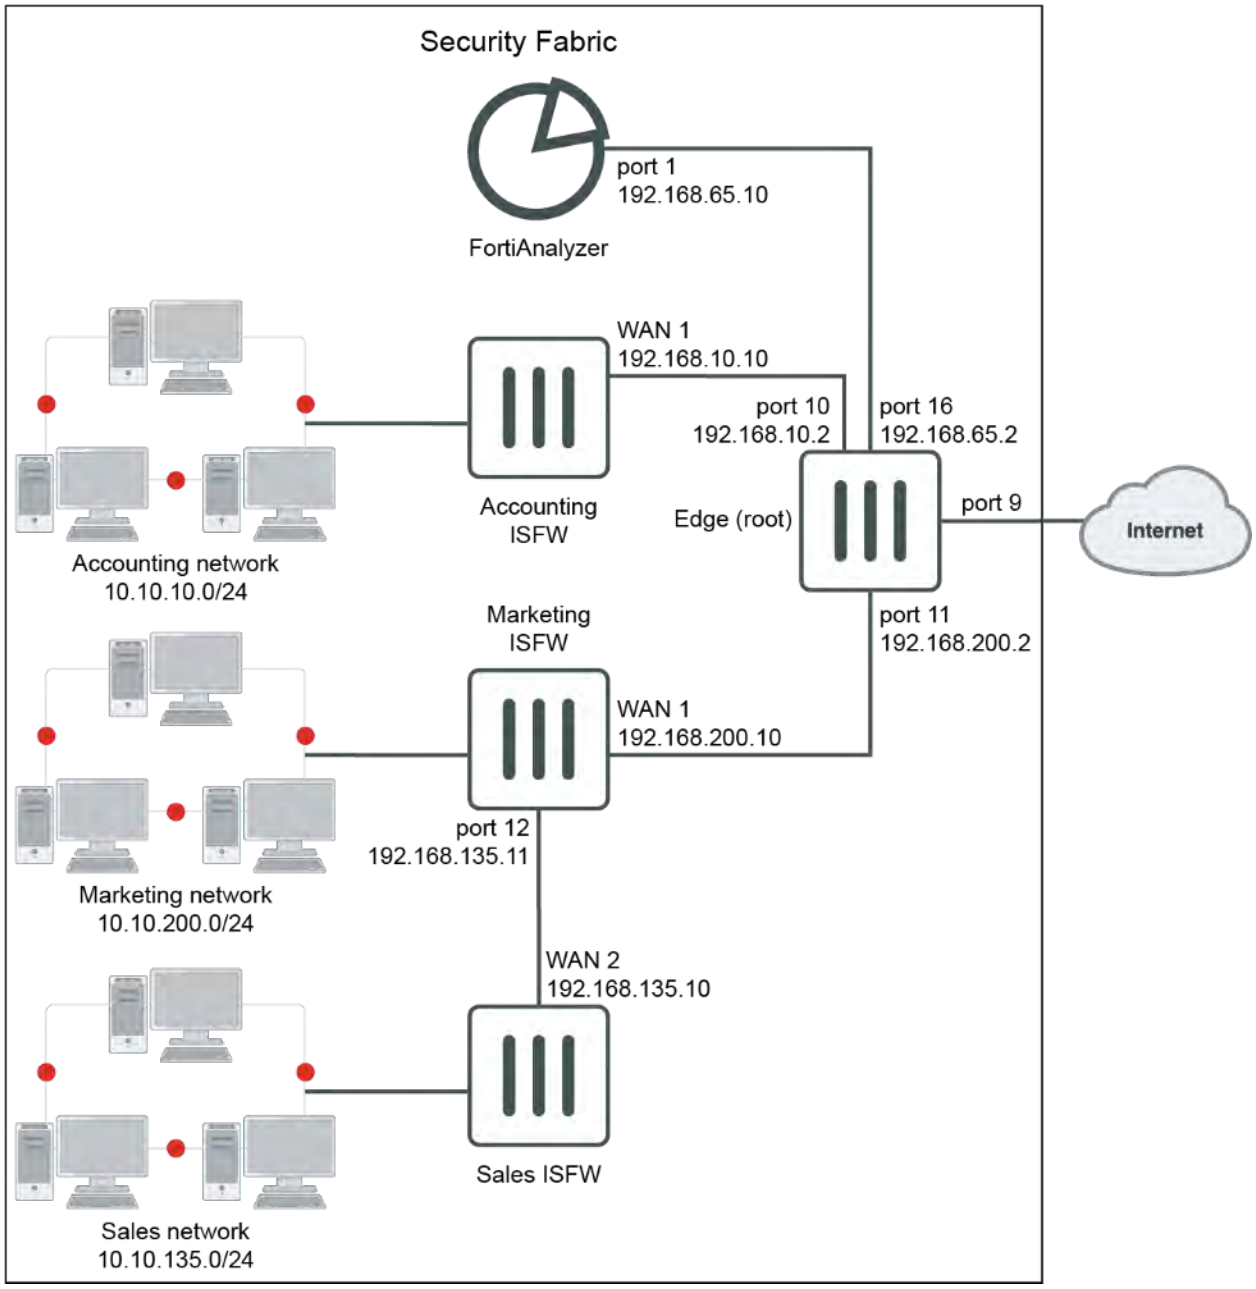

Netskope Exhibit 1017

```
                set description "manual-qtn Hostname: PC2"
            next
        end
    next
  end
end
```

Deploying the Security Fabric

This topic provides an example of deploying Security Fabric with three downstream FortiGates connecting to one root FortiGate. To deploy Security Fabric, you need a FortiAnalyzer running firmware version 6.2 or later.

The following shows a sample network topology with three downstream FortiGates (Accounting, Marketing, and Sales) connected to the root FortiGate (Edge).



To configure the root FortiGate (Edge):

1. Configure interfaces:
   a. In the root FortiGate (Edge), go to *Network* > *Interfaces*.
   b. Edit *port16*:
      - Set *Role* to *DMZ*.
      - For the interface connected to FortiAnalyzer, set the *IP/Network Mask* to *192.168.65.2/255.255.255.0*

c. Edit *port10*:
- Set *Role* to *LAN*.
- For the interface connected to the downstream FortiGate (Accounting), set the *IP/Network Mask* to *192.168.10.2/255.255.255.0*

d. Edit *port11*:
- Set *Role* to *LAN*.
- For the interface connected to the downstream FortiGate (Marketing), set the *IP/Network Mask* to *192.168.200.2/255.255.255.0*

2. Configure Security Fabric:
a. In the root FortiGate (Edge), go to *Security Fabric* > *Fabric Connectors* and double-click the *Security Fabric Setup* card.
b. For *Status*, click *Enable*.
c. Set the *Security Fabric role* to *Serve as Fabric Root*. The FortiAnalyzer settings can be configured.
d. Enter the FortiAnalyzer IP (*192.168.65.10*) and select and *Upload option* (the default is *Real Time*).
e. Click *Test Connectivity*.
A warning message indicates that the FortiGate is not authorized on the FortiAnalyzer. The authorization is configured in a later step on the FortiAnalyzer.
f. Click *OK*. The FortiAnalyzer serial number is verified.
g. Enter a *Fabric name*, such as *Office-Security-Fabric*.
h. Ensure *Allow other Security Fabric devices to join* is enabled and add *port10* and *port11*.
i. Click *OK*.
3. Create a policy to allow the downstream FortiGate (Accounting) to access the FortiAnalyzer:
a. In the root FortiGate (Edge), go to *Policy & Objects* > *Addresses*.
b. Click *Create New*.
- Set *Name* to *FAZ-addr*.
- Set *Type* to *Subnet*.
- Set *Subnet/IP Range* to *192.168.65.10/32*.
- Set *Interface* to *any*.

c. Click *OK*.
d. Click *Create New*.
- Set *Name* to *Accounting*.
- Set *Type* to *Subnet*.
- Set *Subnet/IP Range* to *192.168.10.10/32*.
- Set *Interface* to *any*.

e. Click *OK*.
f. In the root FortiGate (Edge), go to *Policy & Objects* > *Firewall Policy* and click *Create New*.
- Set *Name* to *Accounting-to-FAZ*.
- Set *srcintf* to *port10*.
- Set *dstintf* to *port16*.
- Set *srcaddr* to *Accounting-addr*.
- Set *dstaddr* to *FAZ-addr*.
- Set *Action* to *Accept*.
- Set *Schedule* to *Always*.
- Set *Service* to *All*.
- Enable *NAT*.
- Set *IP Pool Configuration* to *Use Outgoing Interface Address*.

Netskope Exhibit 1017

g. Click *OK*.

4. Create a policy to allow the two downstream FortiGates (Marketing and Sales) to access the FortiAnalyzer:
   a. In the root FortiGate (Edge), go to *Policy & Objects* > *Addresses* and click *Create New*.
      - Set *Name* to *Marketing-addr*.
      - Set *Type* to *Subnet*.
      - Set *Subnet/IP Range* to *192.168.200.10/32*.
      - Set *Interface* to *any*.
   b. Click *OK*.
   c. In the root FortiGate (Edge), go to *Policy & Objects* > *Firewall Policy* and click *Create New*.
      - Set *Name* to *Marketing-to-FAZ*.
      - Set *srcintf* to *port11*.
      - Set *dstintf* to *port16*.
      - Set *srcaddr* to *Marketing-addr*.
      - Set *dstaddr* to *FAZ-addr*.
      - Set *Action* to *Accept*.
      - Set *Schedule* to *Always*.
      - Set *Service* to *All*.
      - Enable *NAT*.
      - Set *IP Pool Configuration* to *Use Outgoing Interface Address*.
   d. Click *OK*.

To configure the downstream FortiGate (Accounting):

1. Configure interface:
   a. In the downstream FortiGate (Accounting), go to *Network* > *Interfaces*.
   b. Edit interface *wan1*:
      - Set *Role* to *WAN*.
      - For the interface connected to root, set the *IP/Network Mask* to *192.168.10.10/255.255.255.0*
2. Configure the default static route to connect to the root FortiGate (Edge):
   a. In the downstream FortiGate (Accounting), go to *Network* > *Static Routes* and click *Create New* or *Create New* > *IPv4 Static Route*.
      - Set *Destination* to *0.0.0.0/0.0.0.0*.
      - Set *Interface* to *wan1*.
      - Set *Gateway Address* to *192.168.10.2*.
   b. Click *OK*.
3. Configure Security Fabric:
   a. In the downstream FortiGate (Accounting), go to *Security Fabric* > *Fabric Connectors* and double-click the *Security Fabric Setup* card.
   b. For *Status*, click *Enable*.
      FortiAnalyzer automatically enables logging. Settings for the FortiAnalyzer are retrieved from the root FortiGate (Edge) when FortiGate (Accounting) connects to the root FortiGate (Edge).
   c. Set the *Security Fabric role* to *Join Existing Fabric*.
   d. *Upstream FortiGate IP* is filled in automatically with the default static route *Gateway Address* of *192.168.10.2* set in the previous step.
   e. Disable *Allow other FortiGates to join*, because there is no downstream FortiGate connecting to it.
   f. Click *OK*.

Netskope Exhibit 1017

To configure the downstream FortiGate (Marketing):

1. Configure interface:
   a. In the downstream FortiGate (Marketing), go to *Network* > *Interfaces*.
   b. Edit *port12*:
      - Set *Role* to *LAN*.
      - For the interface connected to the downstream FortiGate (Sales), set the *IP/Network Mask* to *192.168.135.11/255.255.255.0*.
   c. Edit *wan1*:
      - Set *Role* to *WAN*.
      - For the interface connected to the root FortiGate (Edge), set the *IP/Network Mask* to *192.168.200.10/255.255.255.0*.
2. Configure the default static route to connect to the root FortiGate (Edge):
   a. In the downstream FortiGate (Marketing), go to *Network* > *Static Routes* and click *Create New* or *Create New* > *IPv4 Static Route*.
      - Set *Destination* to *0.0.0.0/0.0.0.0*.
      - Set *Interface* to *wan1*.
      - Set *Gateway Address* to *192.168.200.2*.
   b. Click *OK*.
3. Configure Security Fabric:
   a. In the downstream FortiGate (Marketing), go to *Security Fabric* > *Fabric Connectors* and double-click the *Security Fabric Setup* card.
   b. For *Status*, click *Enable*.
   FortiAnalyzer automatically enables logging. Settings for the FortiAnalyzer are retrieved from the root FortiGate (Edge) when FortiGate (Marketing) connects to the root FortiGate (Edge).
   c. Set the *Security Fabric role* to *Join Existing Fabric*.
   d. *Upstream FortiGate IP* is filled in automatically with the default static route *Gateway Address* of *192.168.200.2* set in the previous step.
   e. Enable *Allow other FortiGates to join* and add *port12*.
   f. Click *OK*.
4. Create a policy to allow another downstream FortiGate (Sales) going through FortiGate (Marketing) to access the FortiAnalyzer:
   a. In the downstream FortiGate (Marketing), go to *Policy & Objects* > *Addresses* and click *Create New*.
      - Set *Name* to *FAZ-addr*.
      - Set *Type* to *Subnet*.
      - Set *Subnet/IP Range* to *192.168.65.10/32*.
      - Set *Interface* to *any*.
   b. Click *OK*.
   c. Click *Create New*.
      - Set *Name* to *Sales-addr*.
      - Set *Type* to *Subnet*.
      - Set *Subnet/IP Range* to *192.168.135.10/32*.
      - Set *Interface* to *any*.
   d. Click *OK*.

Netskope Exhibit 1017

e. In the downstream FortiGate (Marketing), go to *Policy & Objects* > *Firewall Policy* and click *Create New*.
- Set *Name* to *Sales-to-FAZ*.
- Set *srcintf* to *port12*.
- Set *dstintf* to *wan1*.
- Set *srcaddr* to *Sales-addr*.
- Set *dstaddr* to *FAZ-addr*.
- Set *Action* to *Accept*.
- Set *Schedule* to *Always*.
- Set *Service* to *All*.
- Enable *NAT*.
- Set *IP Pool Configuration* to *Use Outgoing Interface Address*.

f. Click *OK*.

To configure the downstream FortiGate (Accounting):

1. Configure interface:
   a. In the downstream FortiGate (Accounting), go to *Network* > *Interfaces*.
   b. Edit interface *wan1*:
   - Set *Role* to *WAN*.
   - For the interface connected to root, set the *IP/Network Mask* to *192.168.10.10/255.255.255.0*
2. Configure the default static route to connect to the root FortiGate (Edge):
   a. In the downstream FortiGate (Accounting), go to *Network* > *Static Routes* and click *Create New* or *Create New* > *IPv4 Static Route*.
   - Set *Destination* to *0.0.0.0/0.0.0.0*.
   - Set *Interface* to *wan1*.
   - Set *Gateway Address* to *192.168.10.2*.

   b. Click *OK*.
3. Configure Security Fabric:
   a. In the downstream FortiGate (Accounting), go to *Security Fabric* > *Fabric Connectors* and double-click the *Security Fabric Setup* card.
   b. For *Status*, click *Enable*.
   FortiAnalyzer automatically enables logging. Settings for the FortiAnalyzer are retrieved from the root FortiGate (Edge) when FortiGate (Accounting) connects to the root FortiGate (Edge).
   c. Set the *Security Fabric role* to *Join Existing Fabric*.
   d. *Upstream FortiGate IP* is filled in automatically with the default static route *Gateway Address* of *192.168.10.2* set in the previous step.
   e. Disable *Allow other FortiGates to join*, because there is no downstream FortiGate connecting to it.
   f. Click *OK*.

To configure the downstream FortiGate (Sales):

1. Configure interface:
   a. In the downstream FortiGate (Sales), go to *Network* > *Interfaces*.
   b. Edit *wan2*:
   - Set *Role* to *WAN*.
   - For the interface connected to the upstream FortiGate (Marketing), set the *IP/Network Mask* to *192.168.135.10/255.255.255.0*.

Netskope Exhibit 1017

2. Configure the default static route to connect to the upstream FortiGate (Marketing):
   a. In the downstream FortiGate (Sales), go to *Network* > *Static Routes* and click *Create New* or *Create New* > *IPv4 Static Route*.
      - Set *Destination* to *0.0.0.0/0.0.0.0*.
      - Set *Interface* to *wan2*.
      - Set *Gateway Address* to *192.168.135.11*.
   b. Click *OK*.
3. Configure Security Fabric:
   a. In the downstream FortiGate (Sales), go to *Security Fabric* > *Fabric Connectors* and double-click the *Security Fabric Setup* card.
   b. For *Status*, click *Enable*.
   FortiAnalyzer automatically enables logging. Settings for the FortiAnalyzer are retrieved from the root FortiGate (Edge) when FortiGate (Sales) connects to the root FortiGate (Edge).
   c. Set the *Security Fabric role* to *Join Existing Fabric*.
   d. *Upstream FortiGate IP* is filled in automatically with the default static route *Gateway Address* of *192.168.135.11* set in the previous step.
   e. Disable *Allow other FortiGates to join*, because there is no downstream FortiGate connecting to it.
   f. Click *OK*.

To authorize downstream FortiGates (Accounting, Marketing, and Sales) on the root FortiGate (Edge):

1. In the root FortiGate (Edge), go to *Security Fabric* > *Fabric Connectors* and double-click the *Security Fabric Setup* card.
   The *Topology* tree highlights two connected FortiGates with their serial numbers and asks you to authorize the highlighted devices.
2. Select the highlighted FortiGates and select *Authorize*.
   After they are authorized, the two downstream FortiGates (Accounting and Marketing) appear in the *Topology* tree in the *Security Fabric* > *Fabric Connectors* > *Security Fabric Setup* page. This means that the two downstream FortiGates (Accounting and Marketing) have successfully joined the Security Fabric.
3. The *Topology* tree now highlights the FortiGate with the serial number that is connected to the downstream FortiGate (Marketing) and asks you to authorize the highlighted device.
4. Select the highlighted FortiGates and select *Authorize*.
   After it is authorized, the downstream FortiGate (Sales) appears in the *Topology* tree in the *Security Fabric* > *Fabric Connectors* > *Security Fabric Setup* page. This means that the downstream FortiGates (Sales) has successfully joined the Security Fabric.

To use FortiAnalyzer to authorize all the Security Fabric FortiGates:

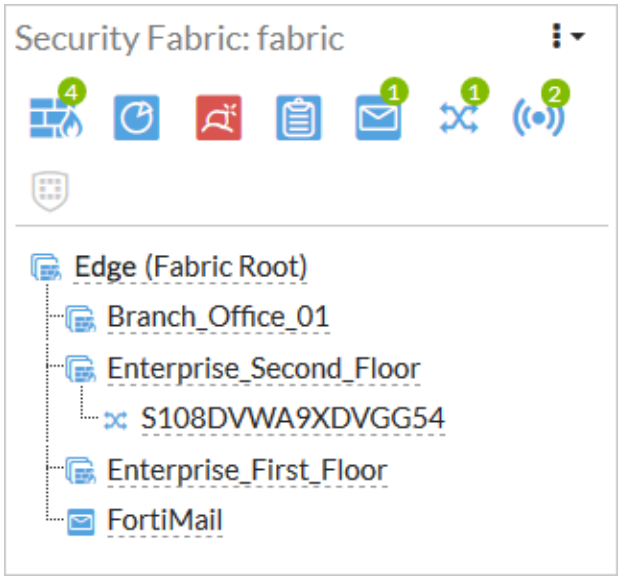
1. Authorize all the Security Fabric FortiGates on the FortiAnalyzer side:
   a. On the FortiAnalyzer, go to *System Settings* > *Network* > *All Interfaces*.
   b. Edit *port1* and set *IP Address/Netmask* to *192.168.65.10/255.255.255.0*.
   c. Go to *Device Manager* > *Unauthorized*. All of the FortiGates are listed as unauthorized.
      i. Select all the FortiGates and select *Authorize*. The FortiGates are now listed as authorized.
      After a moment, a warning icon appears beside the root FortiGate (Edge) because the FortiAnalyzer needs administrative access to the root FortiGate (Edge) in the Security Fabric.
      ii. Click the warning icon and enter the admin username and password of the root FortiGate (Edge).

Netskope Exhibit 1017

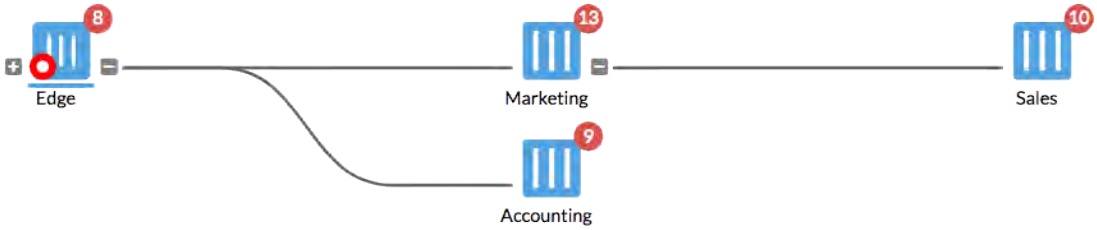
2. Check FortiAnalyzer status on all the Security Fabric FortiGates:
   a. On each FortiGate, go to *Security Fabric* > *Fabric Connectors* and double-click the *FortiAnalyzer Logging* card.
   b. Check that *Storage usage* information is shown.

To check Security Fabric deployment result:

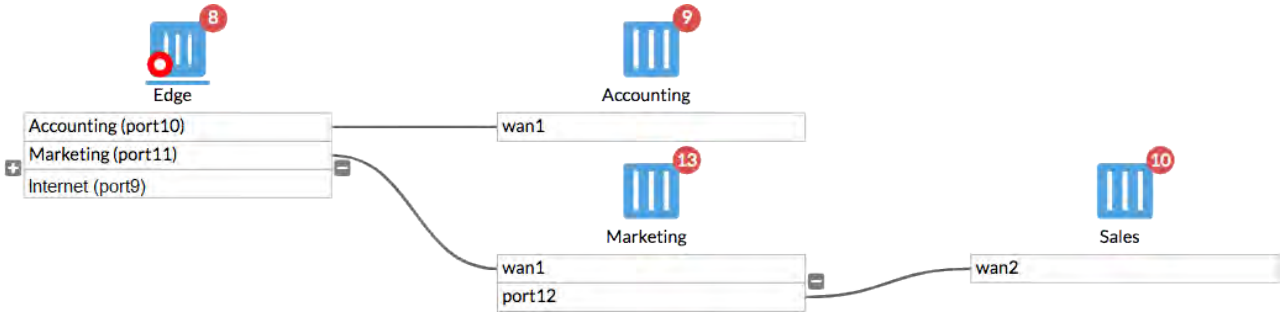
1. On FortiGate (Edge), go to *Dashboard* > *Status*.
   The *Security Fabric* widget displays all the FortiGates in the Security Fabric.



2. On FortiGate (Edge), go to *Security Fabric* > *Physical Topology*.
   This page shows a visualization of access layer devices in the Security Fabric.



3. On FortiGate (Edge), go to *Security Fabric* > *Physical Topology*.
   This dashboard shows information about the interfaces of each device in the Security Fabric.



To run diagnose commands:

1. Run the `diagnose sys csf authorization pending-list` command in the root FortiGate to show the downstream FortiGate pending for root FortiGate authorization:

```
Edge # diagnose sys csf authorization pending-list
Serial                  IP Address      HA-Members    Path
-------------------------------------------------------------------------------
FG201ETK18902514        0.0.0.0                       FG3H1E5818900718:FG201ETK18902514
```

2. Run the `diagnose sys csf downstream` command in the root or middle FortiGate to show the downstream FortiGates after they join Security Fabric:

```
Edge # diagnose sys csf downstream
1:     FG201ETK18902514 (192.168.200.10) Management-IP: 0.0.0.0 Management-port:0
parent: FG3H1E5818900718
        path:FG3H1E5818900718:FG201ETK18902514
        data received: Y downstream intf:wan1 upstream intf:port11 admin-port:443
        authorizer:FG3H1E5818900718
2:     FGT81ETK18002246 (192.168.10.10) Management-IP: 0.0.0.0 Management-port:0 parent:
FG3H1E5818900718
        path:FG3H1E5818900718:FGT81ETK18002246
        data received: Y downstream intf:wan1 upstream intf:port10 admin-port:443
        authorizer:FG3H1E5818900718
3:     FG101ETK18002187 (192.168.135.10) Management-IP: 0.0.0.0 Management-port:0
parent: FG201ETK18902514
        path:FG3H1E5818900718:FG201ETK18902514:FG101ETK18002187
        data received: Y downstream intf:wan2 upstream intf:port12 admin-port:443
        authorizer:FG3H1E5818900718
```

3. Run the `diagnose sys csf upstream` command in any downstream FortiGate to show the upstream FortiGate after downstream FortiGate joins Security Fabric:

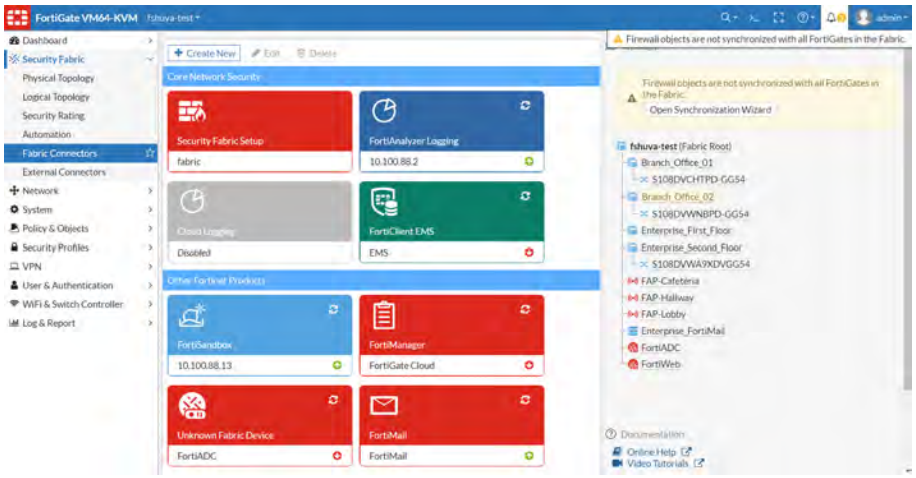
```
Marketing # diagnose sys csf upstream
Upstream Information:
Serial Number:FG3H1E5818900718
IP:192.168.200.2
Connecting interface:wan1
Connection status:Authorized
```

Synchronizing objects across the Security Fabric

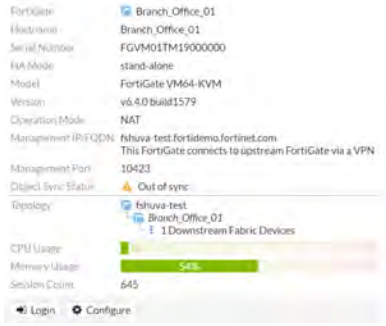
When the Security Fabric is enabled, various objects such as addresses, services, and schedules are synced from the upstream FortiGate to all downstream devices by default. The firewall object synchronization wizard helps identify objects that are out of sync and resolves any conflicts. Objects that are out of sync are highlighted in yellow in the GUI.

In this example, the notifications icon displays a message that *Firewall objects are not synchronized with all the FortiGates in the Fabric*. In the topology tree, *Branch_Office_02* is highlighted in yellow because it is out of sync.

Netskope Exhibit 1017

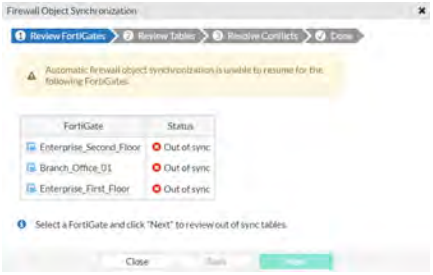


In this example, the tooltip displays a caution icon that the device is out of sync.



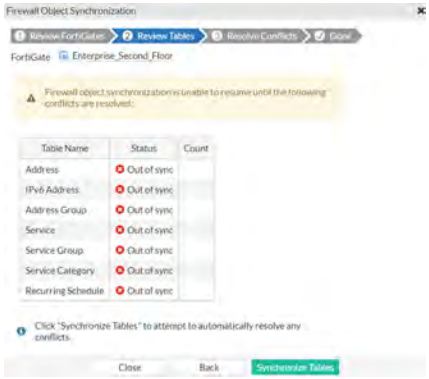
To use the firewall object synchronization wizard in the GUI:

1. Go to *Security Fabric >Fabric Connectors* and click *Open Synchronization Wizard*.
   A list of FortiGates and their synchronization status displays.

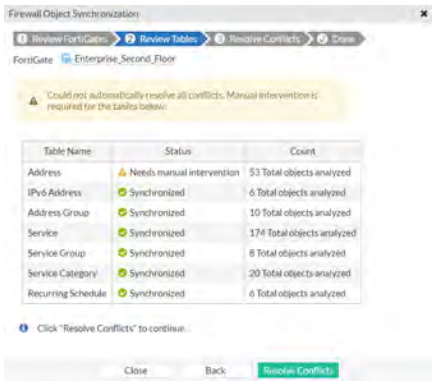


2. Select a FortiGate that is *Out of sync* and click *Next*.
   A list of tables and their synchronization status displays.

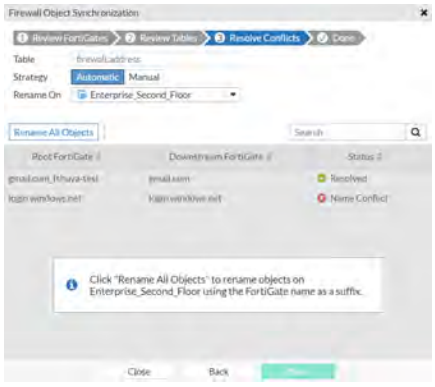
3. Click *Synchronize Tables*.



The FortiGate attempts to automatically resolve the conflicts. In this example, the address table requires manual intervention.



4. Click *Resolve Conflicts*.
5. For *Strategy*, choose one of the following.
   a. Automatic resolve (automatically resolves all the name conflicts and renames them on the selected FortiGate using the FortiGate name as a suffix):
      i. Click *Automatic*.
      ii. Click *Rename All Objects*.

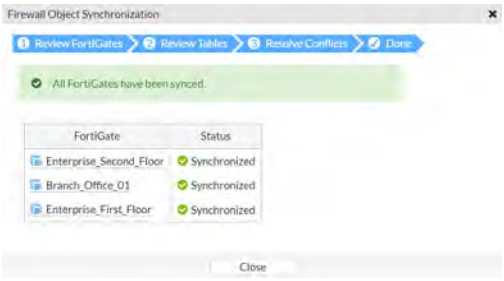


   b. Manual resolve:
      i. Click *Manual*.
      ii. Double-click an object and re-name it.

iii. Click *OK*.



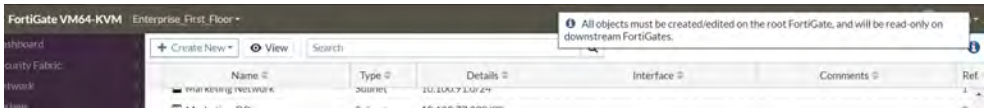
6. Click *Next*.
   An updated list of FortiGates and their synchronization status displays.



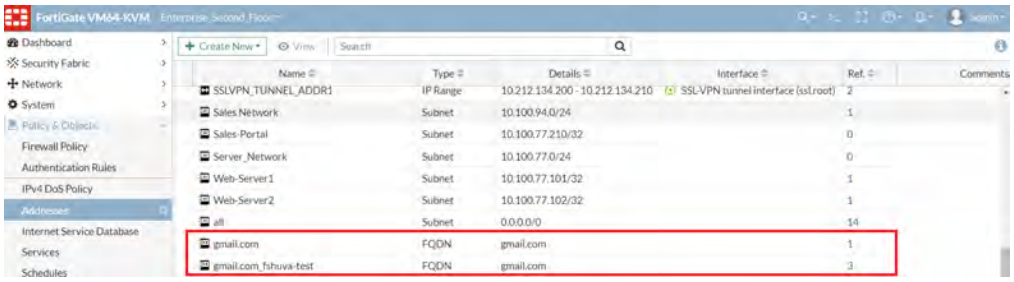
7. Click *Close*.

To verify object synchronization on downstream devices:

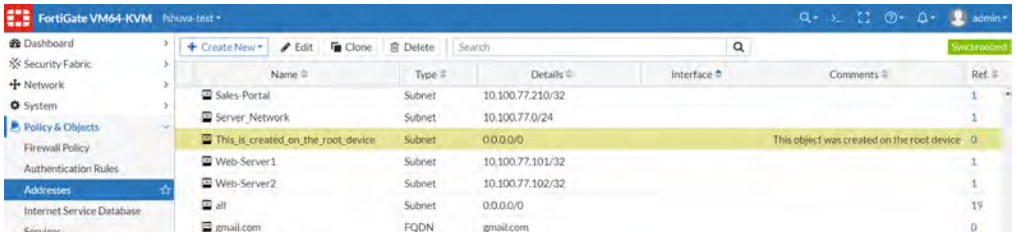
1. Log in to a downstream device.
2. Go to *Policy & Objects* > *Addresses*.
   - An information bubble displays the following: *All objects must be created/edited on the root FortiGate, and will be read-only on downstream FortiGates*.



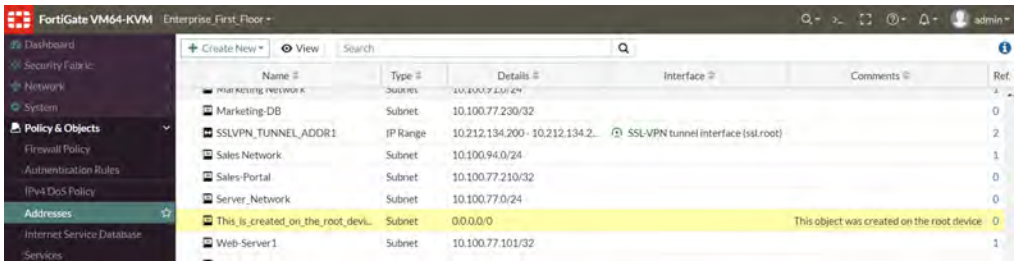
   - The following example shows an object that exists on both upstream (Enterprise_Second_Floor) and downstream (fshuva-test) FortiGates. On the downstream device, there is an existing *gmail.com*, and another object, *gmail.com_fshuva-test*, that was resolved by adding the suffix of the upstream FortiGate name to the end.



- In this example, an object created on the upstream FortiGate is synchronized to a downstream FortiGate.



The same object appears automatically on the downstream device.



CLI commands

Object synchronization can be configured with the following commands:

```
config system csf
    set fabric-object-unification [default | local]
    set configuration-sync [default | local]
    ...
    next
end
```

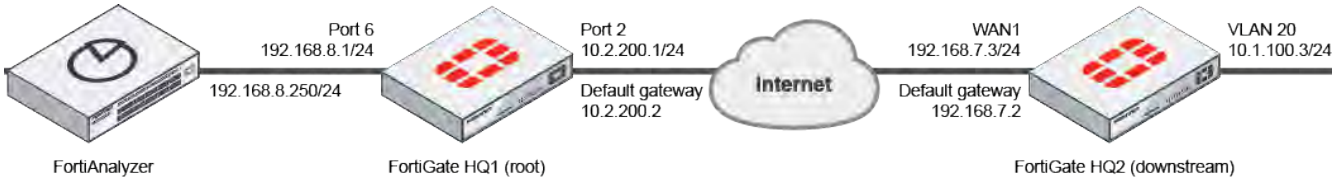
Parameter	Description
fabric-object-unification	*default*: Global CMDB objects will be synchronized in Security Fabric. *local*: Global CMDB objects will not be synchronized to and from this device.
configuration-sync	*default*: Synchronize configuration for FortiAnalyzer, FortiSandbox, and Central Management to root node. *local*: Do not synchronize configuration with root node.

Security Fabric over IPsec VPN

This is an example of configuring Security Fabric over IPsec VPN.

Sample topology

This sample topology shows a downstream FortiGate (HQ2) connected to the root FortiGate (HQ1) over IPsec VPN to join Security Fabric.



Sample configuration

To configure the root FortiGate (HQ1):

1. Configure interface:
   a. In the root FortiGate (HQ1), go to *Network* > *Interfaces*.
   b. Edit *port2*:
      - Set *Role* to *WAN*.
      - For the interface connected to the Internet, set the *IP/Network Mask* to *10.2.200.1/255.255.255.0*
   c. Edit *port6*:
      - Set *Role* to *DMZ*.
      - For the interface connected to FortiAnalyzer, set the *IP/Network Mask* to *192.168.8.250/255.255.255.0*
2. Configure the static route to connect to the Internet:
   a. Go to *Network* > *Static Routes* and click *Create New* or *Create New* > *IPv4 Static Route*.
      - Set *Destination* to *0.0.0.0/0.0.0.0*.
      - Set *Interface* to *port2*.
      - Set *Gateway Address* to *10.2.200.2*.
   b. Click *OK*.
3. Configure IPsec VPN:
   a. Go to *VPN* > *IPsec Wizard*.
      - Set *Name* to *To-HQ2*.
      - Set *Template Type* to *Custom*.
      - Click *Next*.
      - Set *Authentication* to *Method*.
      - Set *Pre-shared Key* to *123456*.
   b. Leave all other fields in their default values and click *OK*.
4. Configure the IPsec VPN interface IP address which will be used to form Security Fabric:
   a. Go to *Network* > *Interfaces*.
   b. Edit *To-HQ2*:
      - Set *Role* to *LAN*.
      - Set the *IP/Network Mask* to *10.10.10.1/255.255.255.255*.
      - Set *Remote IP/Network Mask* to *10.10.10.3/255.255.255.0*.

5. Configure IPsec VPN local and remote subnet:
   a. Go to *Policy & Objects* > *Addresses*.
   b. Click *Create New*
      - Set *Name* to *To-HQ2_remote_subnet_2*.
      - Set *Type* to *Subnet*.
      - Set *IP/Network Mask* to *10.10.10.3/32*.
   c. Click *OK*.
   d. Click *Create New*
      - Set *Name* to *To-HQ2_local_subnet_1*.
      - Set *Type* to *Subnet*.
      - Set *IP/Network Mask* to *192.168.8.0/24*.
   e. Click *OK*.
   f. Click *Create New*
      - Set *Name* to *To-HQ2_remote_subnet_1*.
      - Set *Type* to *Subnet*.
      - Set *IP/Network Mask* to *10.1.100.0/24*.
   g. Click *OK*.
6. Configure IPsec VPN static routes:
   a. Go to *Network* > *Static Routes*
   b. Click *Create New* or *Create New* > *IPv4 Static Route*.
      - For *Named Address*, select *Type* and select *To-HQ2_remote_subnet_1*.
      - Set *Interface* to *To-HQ2*.

      Click *OK*.
   c. Click *Create New* or *Create New* > *IPv4 Static Route*.
      - For *Named Address*, select *Type* and select *To-HQ2_remote_subnet_1*.
      - Set *Interface* to *Blackhole*.
      - Set *Administrative Distance* to *254*.
   d. Click *OK*.
7. Configure IPsec VPN policies:
   a. Go to *Policy & Objects* > *Firewall Policy*
   b. Click *Create New*.
      - Set *Name* to *vpn_To-HQ2_local*.
      - Set *Incoming Interface* to *port6*.
      - Set *Outgoing Interface* to *To-HQ2*.
      - Set *Source* to *To-HQ2_local_subnet_1*.
      - Set *Destination* to *To-HQ2_remote_subnet_1*.
      - Set *Schedule* to *Always*.
      - Set *Service* to *All*.
      - Disable *NAT*.
   c. Click *OK*.
   d. Click *Create New*.
      - Set *Name* to *vpn_To-HQ2_remote*.
      - Set *Incoming Interface* to *To-HQ2*.
      - Set *Outgoing Interface* to *port6*.
      - Set *Source* to *To-HQ2_remote_subnet_1, To-HQ2_remote_subnet_2*.

Netskope Exhibit 1017

- Set *Destination* to *To-HQ2_local_subnet_1*.
- Set *Schedule* to *Always*.
- Set *Service* to *All*.
- Enable *NAT*.
- Set *IP Pool Configuration* to *Use Outgoing Interface Address*.

e. Click *OK*.

8. Configure Security Fabric:
   a. Go to *Security Fabric* > *Fabric Connectors* and double-click the *Security Fabric Setup* card.
   b. For *Status*, click *Enable*.
   After *FortiGate Telemetry* is enabled, FortiAnalyzer automatically enables *Logging* and *Upload* is set to *Real Time*.
   c. Set the *Security Fabric role* to *Serve as Fabric Root*. The FortiAnalyzer settings can be configured.
   d. Enter the FortiAnalyzer IP (*192.168.8.250*).
   e. Click *OK*. The FortiAnalyzer serial number is verified.
   f. Enter a *Fabric name*, such as *Office-Security-Fabric*.
   g. Ensure *Allow other Security Fabric devices to join* is enabled and add VPN interface *To-HQ2*.
   h. Click *OK*.

To configure the downstream FortiGate (HQ2):

1. Configure interface:
   a. Go to *Network* > *Interfaces*.
   b. Edit interface *wan1*:
   - Set *Role* to *WAN*.
   - For the interface connected to the Internet, set the *IP/Network Mask* to *192.168.7.3/255.255.255.0*.

   c. Edit interface *vlan20*:
   - Set *Role* to *LAN*.
   - For the interface connected to local endpoint clients, set the *IP/Network Mask* to *10.1.100.3/255.255.255.0*.

2. Configure the static route to connect to the Internet:
   a. Go to *Network* > *Static Routes* and click *Create New* or *Create New* > *IPv4 Static Route*.
   - Set *Destination* to *0.0.0.0/0.0.0.0*.
   - Set *Interface* to *wan1*.
   - Set *Gateway Address* to *192.168.7.2*.

   b. Click *OK*.
3. Configure IPsec VPN:
   a. Go to *VPN* > *IPsec Wizard*.
   - Set *VPN Name* to *To-HQ1*.
   - Set *Template Type* to *Custom*.
   - Click *Next*.
   - In the *Network IP Address*, enter *10.2.200.1*.
   - Set *Interface* to *wan1*.
   - Set *Authentication* to *Method*.
   - Set *Pre-shared Key* to *123456*.

   b. Leave all other fields in their default values and click *OK*.

Netskope Exhibit 1017

4. Configure the IPsec VPN interface IP address which will be used to form Security Fabric:
   a. Go to *Network* > *Interfaces*.
   b. Edit *To-HQ1*:
      - Set *Role* to *WAN*.
      - Set the *IP/Network Mask* to *10.10.10.3/255.255.255.255*.
      - Set *Remote IP/Network Mask* to *10.10.10.1/255.255.255.0.0*.
5. Configure IPsec VPN local and remote subnet:
   a. Go to *Policy & Objects* > *Addresses*.
   b. Click *Create New*
      - Set *Name* to *To-HQ1_local_subnet_1*.
      - Set *Type* to *Subnet*.
      - Set *IP/Network Mask* to *10.1.100.0/24*.
   c. Click *OK*.
   d. Click *Create New*
      - Set *Name* to *To-HQ1_remote_subnet_1*.
      - Set *Type* to *Subnet*.
      - Set *IP/Network Mask* to *192.168.8.0/24*.
   e. Click *OK*.
6. Configure IPsec VPN static routes:
   a. Go to *Network* > *Static Routes* and click *Create New* or *Create New* > *IPv4 Static Route*.
      - For *Named Address*, select *Type* and select *To-HQ1_remote_subnet_1*.
      - Set *Interface* to *To-HQ1*.
   b. Click *OK*.
   c. Click *Create New* or *Create New* > *IPv4 Static Route*.
      - For *Named Address*, select *Type* and select *To-HQ1_remote_subnet_1*.
      - Set *Interface* to *Blackhole*.
      - Set *Administrative Distance* to *254*.
   d. Click *OK*.
7. Configure IPsec VPN policies:
   a. Go to *Policy & Objects* > *Firewall Policy* and click *Create New*.
      - Set *Name* to *vpn_To-HQ1_local*.
      - Set *Incoming Interface* to *vlan20*.
      - Set *Outgoing Interface* to *To-HQ1*.
      - Set *Source* to *To-HQ1_local_subnet_1*.
      - Set *Destination* to *To-HQ1_remote_subnet_1*.
      - Set *Schedule* to *Always*.
      - Set *Service* to *All*.
      - Disable *NAT*.
   b. Click *OK*.
   c. Click *Create New*.
      - Set *Name* to *vpn_To-HQ1_remote*.
      - Set *Incoming Interface* to *To-HQ1*.
      - Set *Outgoing Interface* to *vlan20*.
      - Set *Source* to *To-HQ1_remote_subnet_1*.
      - Set *Destination* to *-HQ1_local_subnet_1*.

Netskope Exhibit 1017

- Set *Schedule* to *Always*.
- Set *Service* to *All*.
- Disable *NAT*.

d. Click *OK*.

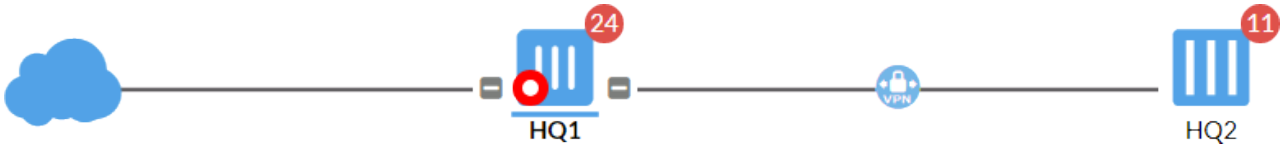
8. Configure Security Fabric:
   a. Go to *Security Fabric* > *Fabric Connectors* and double-click the *Security Fabric Setup* card.
   b. For *Status*, click *Enable*.
   FortiAnalyzer automatically enables logging. FortiAnalyzer settings will be retrieved when the downstream FortiGate connects to the root FortiGate.
   c. Set the *Security Fabric role* to *Join Existing Fabric*.
   d. Set the *Upstream FortiGate IP* to *10.10.10.1*.
   e. Click *OK*.

To authorize the downstream FortiGate (HQ2) on the root FortiGate (HQ1):

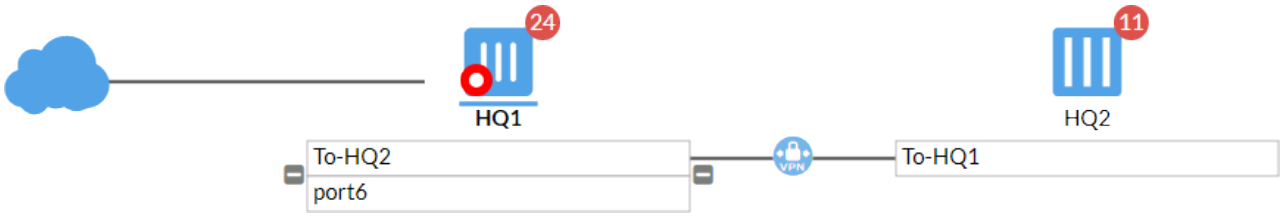
1. In the root FortiGate (HQ1), go to *Security Fabric* > *Fabric Connectors* and double-click the *Security Fabric Setup* card.
   The *Topology* tree highlights the connected FortiGate (HQ2) with the serial number and asks you to authorize the highlighted device.
2. Select the highlighted FortiGates and select *Authorize*.
   After authorization, the downstream FortiGate (HQ2) appears in the *Topology* tree in the *Security Fabric* > *Fabric Connectors* > *Security Fabric Setup* page. This means the downstream FortiGate (HQ2) has successfully joined the Security Fabric.

To check Security Fabric over IPsec VPN:

1. On the root FortiGate (HQ1), go to *Security Fabric* > *Physical Topology*.
   The root FortiGate (HQ1) is connected by the downstream FortiGate (HQ2) with VPN icon in the middle.



2. On the root FortiGate (HQ1), go to *Security Fabric* > *Logical Topology*.
   The root FortiGate (HQ1) VPN interface *To-HQ2* is connected by downstream FortiGate (HQ2) VPN interface *To-HQ1* with VPN icon in the middle.



To run diagnose commands:

1. Run the `diagnose sys csf authorization pending-list` command in the root FortiGate (HQ1) to show the downstream FortiGate pending for root FortiGate authorization:

```
HQ1 # diagnose sys csf authorization pending-list
Serial                  IP Address      HA-Members
Path
--------------------------------------------------------------------------------
FG101ETK18002187        0.0.0.0
FG3H1E5818900718:FG101ETK18002187
```

2. Run the `diagnose sys csf downstream` command in the root FortiGate (HQ1) to show the downstream FortiGate (HQ2) after it joins Security Fabric:

```
HQ1 # diagnose sys csf downstream
 1:     FG101ETK18002187 (10.10.10.3) Management-IP: 0.0.0.0 Management-port:0 parent:
FG3H1E5818900718
        path:FG3H1E5818900718:FG101ETK18002187
        data received: Y downstream intf:To-HQ1 upstream intf:To-HQ2 admin-port:443
        authorizer:FG3H1E5818900718
```

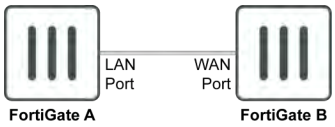
3. Run the `diagnose sys csf upstream` command in the downstream FortiGate (HQ2) to show the root FortiGate (HQ1) after the downstream FortiGate joins Security Fabric:

```
HQ2 # diagnose sys csf upstream
Upstream Information:
Serial Number:FG3H1E5818900718
IP:10.10.10.1
Connecting interface:To-HQ1
Connection status:Authorized
```

Leveraging LLDP to simplify security fabric negotiation

This feature enables LLDP reception on WAN interfaces, and prompts FortiGates that are joining the Security Fabric if the upstream FortiGate asks.

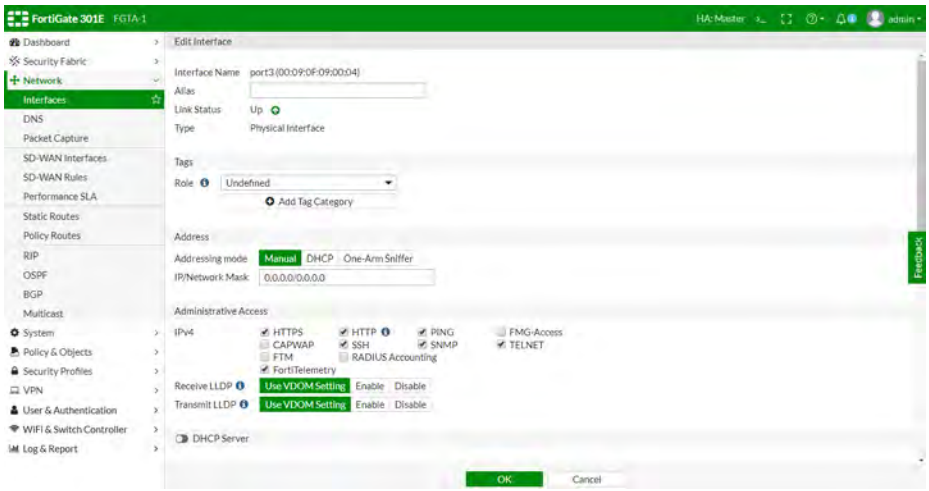
- If an interface's role is undefined, LLDP reception and transmission inherit settings from the VDOM.
- If an interface's role is WAN, LLDP reception is enabled.
- If an interface's role is LAN, LLDP transmission is enabled.



When a FortiGate B's WAN interface detects that FortiGate A's LAN interface is immediately upstream (through the default gateway), and FortiGate A has Security Fabric enabled, FortiGate B will show a notification on the GUI asking to join the Security Fabric.

To configure LLDP reception and join a Security Fabric:

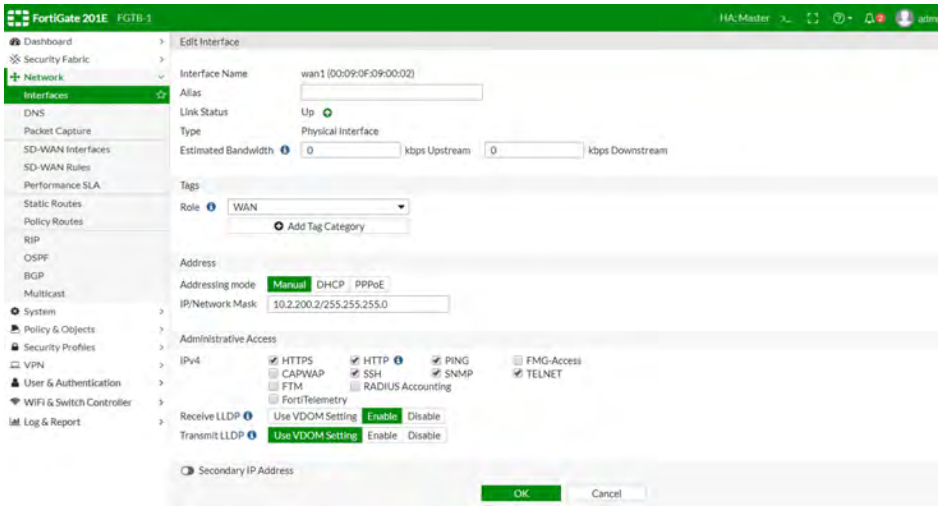
1. Go To *Network* > *Interfaces*.
2. Configure an interface:
   - If the interface's role is undefined, under *Administrative Access*, set *Receive LLDP* and *Transmit LLDP* to *Use VDOM Setting*.



Using the CLI:

```
config system interface
    edit "port3"
        set lldp-reception vdom
        set lldp-transmission vdom
        set role undefined
        ...
    next
end
```

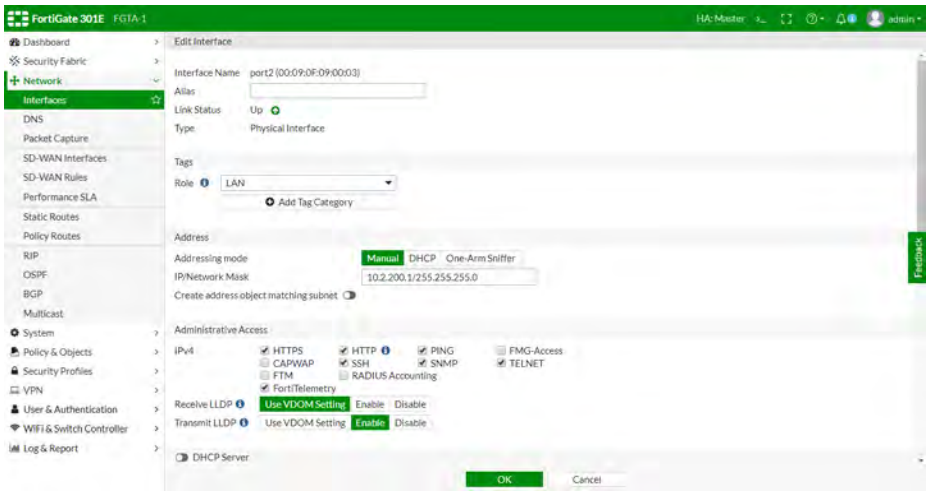
- If the interface's role is WAN, under *Administrative Access*, set *Receive LLDP* to *Enable* and *Transmit LLDP* to *Use VDOM Setting*.



Using the CLI:

```
config system interface
    edit "wan1"
        set lldp-reception enable
        set lldp-transmission vdom
        set role wan
        ...
    next
end
```

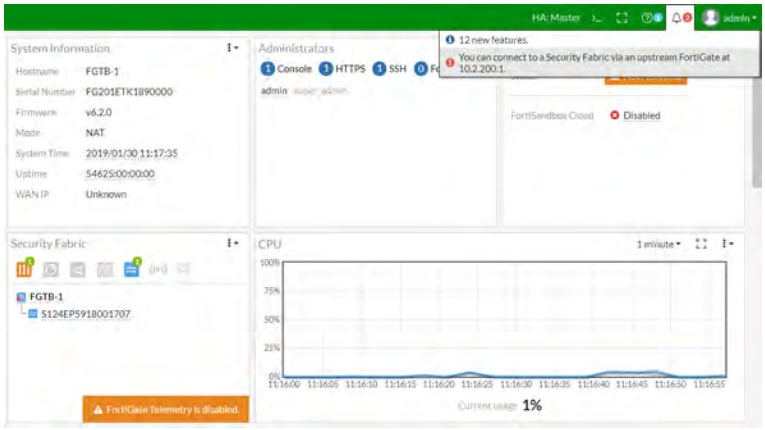
- If the interface's role is LAN, under *Administrative Access*, set *Receive LLDP* to *Use VDOM Setting* and *Transmit LLDP* to *Enable*.



Using the CLI:

```
config system interface
    edit "port2"
        set lldp-reception vdom
        set lldp-transmission enable
        set role lan
        ...
    next
end
```

A notification will be shown on FortiGate B.



3. Click the notification. The *Core Network Security* page with the Security Fabric settings opens. All the required settings automatically configured.
4. Click *OK* to apply the settings, or use the following CLI commands:

```
config system csf
    set status enable
    set upstream-ip 10.2.200.1
end
```

Configuring the Security Fabric with SAML

Security Assertion Markup Language (SAML) is an open standard for exchanging authentication and authorization data between one Identity Provider (IdP) and one or more Service Providers (SP). Both parties exchange messages using the XML protocol as transport. FortiGate firewall devices can be configured as IdPs or SPs.

When the Security Fabric is enabled, you can configure the root FortiGate as the IdP. You can also configure downstream FortiGates to be automatically configured as SPs, with all links required for SAML communication, when added to the Security Fabric. Administrators must still be authorized on each device. Credentials are verified by the root FortiGate, and login credentials are shared between devices. Once authorized, an administrator can move between fabric devices without logging in again.

Optionally, the downstream FortiGate can also be manually configured as an SP, and then linked to the root FortiGate.

The authentication service is provided by the root FortiGate using local system admin accounts for authentication. Any of the administrator account types can be used for SAML log in. After successful authentication, the administrator logs in to the first downstream FortiGate SP, and can then connect to other downstream FortiGates that have the SSO account properly configured, without needing to provide credentials again, as long as admins use the same browser session. In summary, the root FortiGate IdP performs SAML SSO authentication, and individual device administrators define authorization on FortiGate SPs by using security profiles.

Configuring single-sign-on in the Security Fabric

SAML SSO enables a single FortiGate device to act as the identify provider (IdP), while other FortiGate devices act as service providers (SP) and redirect logins to the IdP.



Only the root FortiGate can be the identity provider (IdP). The downstream FortiGates can be configured as service providers (SP).

The process is as follows:

1. Configuring the root FortiGate as the IdP on page 179
2. Configuring a downstream FortiGate as an SP on page 180
3. Configuring certificates for SAML SSO on page 182
4. Verifying the single-sign-on configuration on page 183

You can also use the CLI. See CLI commands for SAML SSO on page 184.

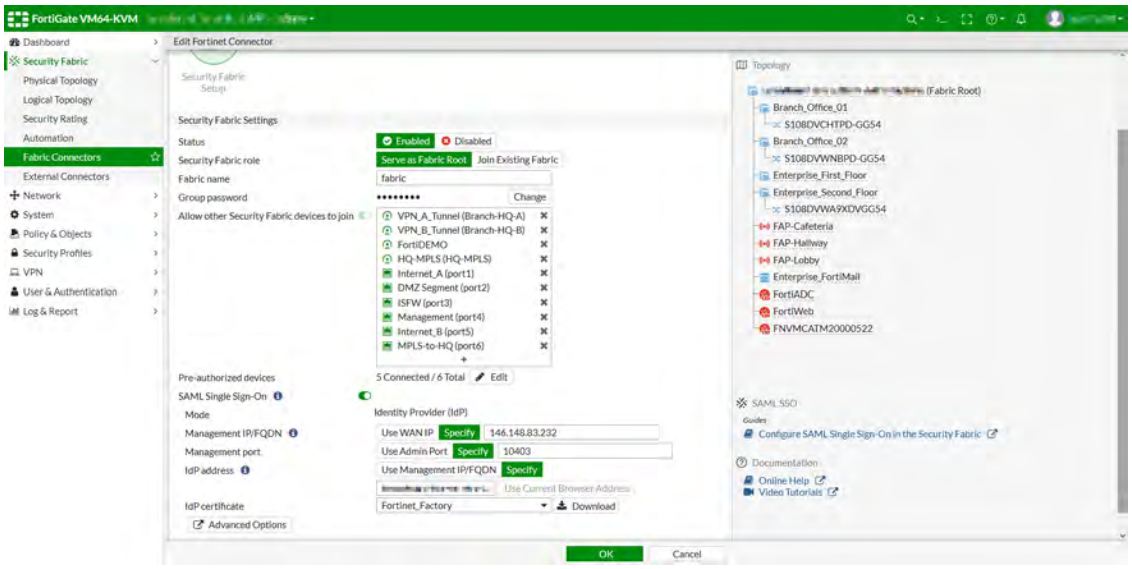
Configuring the root FortiGate as the IdP

To configure the root FortiGate as the IdP:

1. Log in to the root FortiGate.
2. Go to *Security Fabric* > *Fabric Connectors* and double-click the *Security Fabric Setup* card.
3. Enable *SAML Single Sign-On*. The *Mode* field is automatically populated as *Identity Provider (IdP)*.
4. Enter an IP address in the *Management IP/FQDN* box.

Netskope Exhibit 1017

5. Enter a management port in the *Management port* box.
   The *Management IP/FQDN* will be used by the SPs to redirect the login request. The *Management IP/FQDN* and *Management port* must be reachable from the user's device.
6. Select the *IdP certificate*.
7. Click *OK*.



Configuring a downstream FortiGate as an SP

There are two ways to configure the downstream FortiGate:

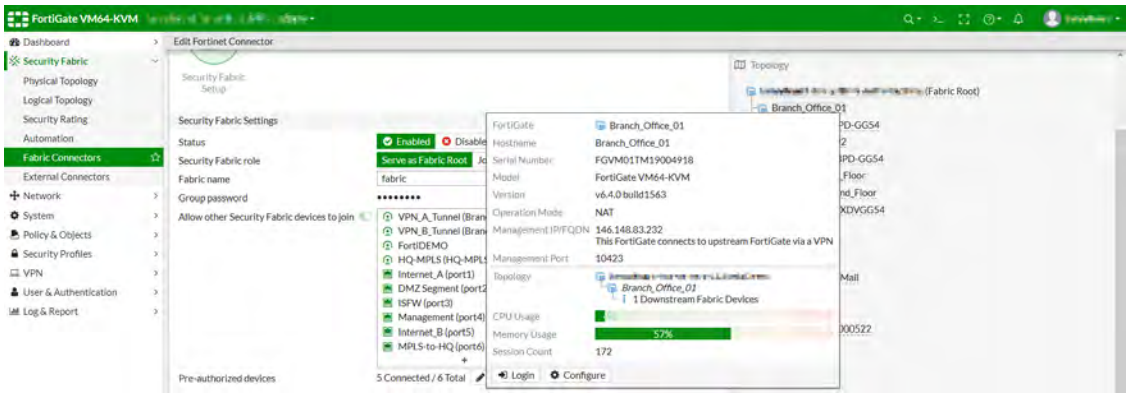
- From the root FortiGate
- From within the downstream device

An SP must be a member of the Security Fabric before you configure it.

To configure the downstream FortiGate from the root FortiGate:

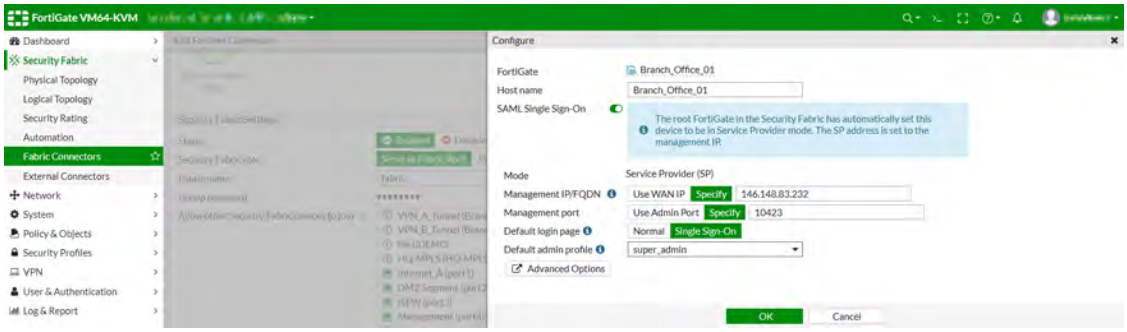
1. Log in to the root FortiGate.
2. Go to *Security Fabric* > *Fabric Connectors* and double-click the *Security Fabric Setup* card.

3. In the *Topology* tree, hover over a FortiGate and click *Configure*.



The *Configure* pane opens.

4. Enable *SAML Single Sign-On*. The *Mode* field is automatically populated as *Service Provider (SP)*.
5. Enter an IP address in the *Management IP/FQDN* box.
6. Enter a management port in the *Management port* box.
   The *Management IP/FQDN* will be used by the IdP and so other SPs can redirect to each other. The *Management port* must be reachable from the user's device.
7. Select a *Default login page* option.
8. Select one of the following *Default admin profile* types: *prof_admin*, *super_admin*, or *super_admin_readonly*. The *no_access_admin profile* is set as the default.
9. Click *OK*.



To configure the downstream FortiGate within the device:

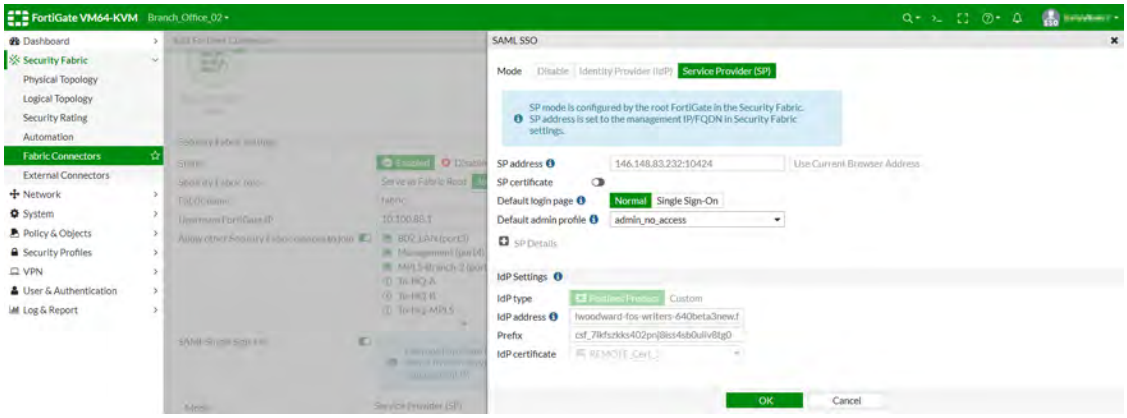
1. Log in to the downstream FortiGate.
2. Go to *Security Fabric* > *Fabric Connectors* and double-click the *Security Fabric Setup* card.
3. Enable *SAML Single Sign-On*. The *Mode* field is automatically populated as *Service Provider (SP)*.
4. Enter an IP address in the *Management IP/FQDN* box.
5. Enter a management port in the *Management port* box.
   The *Management IP/FQDN* will be used by the IdP and so other SPs can redirect to each other. The *Management port* must be reachable from the user's device.
6. Select a *Default login page* option.
7. Select one of the following *Default admin profile* types: *prof_admin*, *super_admin*, or *super_admin_readonly*. The *no_access_admin profile* is set as the default.
8. Click *OK*.

Netskope Exhibit 1017

Configuring certificates for SAML SSO

Because communication between the root FortiGate IdP and FortiGate SPs is secured, you must select a local server certificate in the *IdP certificate* option on the root FortiGate. When downstream SPs join the IdP (root FortiGate), the SP automatically obtains the certificate.

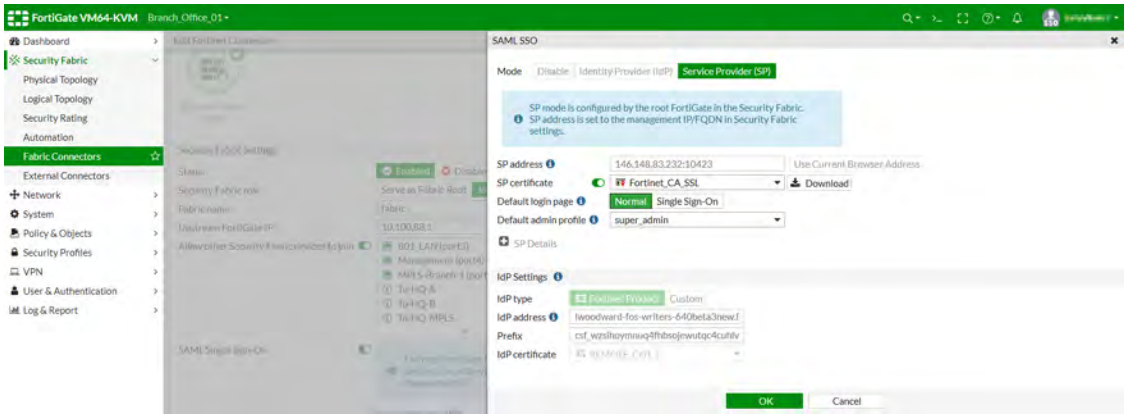
In the following SP example, the *IdP certificate* displays *REMOTE_Cert_2*, which is the root server certificate for the IdP:



It is possible to manually import a certificate from an SP to the IdP so it can be used for authentication.

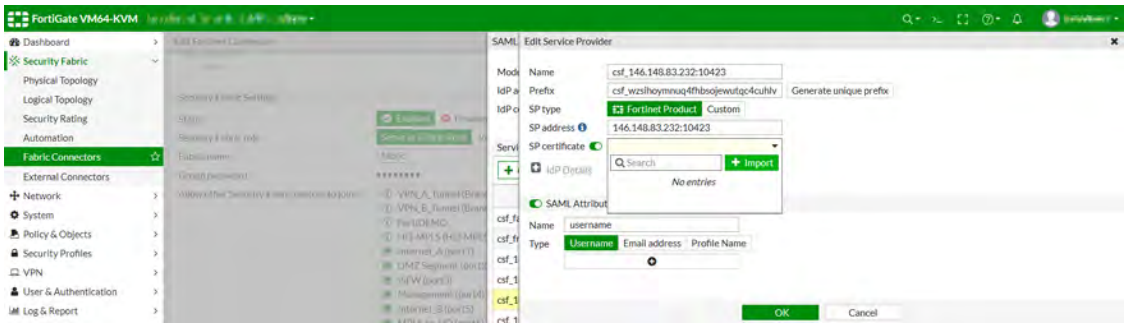
To manually import an SP certificate to an IdP:

1. Add the certificate:
   a. On the SP, go to *Security Fabric* > *Fabric Connectors* and double-click the *Security Fabric Setup* card.
   b. Click *Advanced Options*. The *SAML SSO* pane opens.
   c. Enable *SP certificate* and select a certificate from the dropdown box.
   d. Click *Download*. The certificate is downloaded on the local file system.
   e. Click *OK*.



2. Import the certificate:
   a. On the IdP, go to *Security Fabric* > *Fabric Connectors* and double-click the *Security Fabric Setup* card.
   b. Click *Advanced Options*. The *SAML SSO* pane opens.
   c. In the *Service Providers* table, select the SP from step 1 and click *Edit*.

d. Enable *SP certificate* and in the dropdown box, click *Import*.



The *Upload Remote Certificate* window opens.

e. Click *Upload* and select the certificate downloaded in step 1.

f. Click *OK*. The certificate is imported.

g. Click *OK*.

h. In the *IdP certificate* list, select the certificate that you imported.

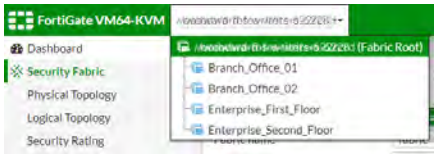
i. Click *OK*.

Verifying the single-sign-on configuration

After you have logged in to a Security Fabric member using SSO, you can navigate between any Security Fabric member with SSO configured.

To navigate between Security Fabric members:

1. Log in to a Security Fabric member that is using SSO.
2. In the top banner, click the name of the device you are logged in to. A list of Security Fabric members displays.



3. Click a Security Fabric member. The login page appears.
4. Select the option to log in *via Single-Sign-On*.



You are now logged in to the Security Fabric member with SSO. The letters "SSO" also display beside the user name in the top banner.

5. Go to *System* > *Administrators* > *Single-Sign-On Administrator* to view the list of SSO admins created.



CLI commands for SAML SSO

To enter a question mark (?) or a tab, Ctrl + V must be entered first. Question marks and tabs cannot be typed or copied into the CLI Console or some SSH clients.

To configure the IdP:

```
config system saml
    set status enable
    set role identity-provider
    set cert "Fortinet_Factory"
    set server-address "172.16.106.74"
    config service-providers
        edit "csf_172.16.106.74:12443"
            set prefix "csf_ngczjwqxujfsbhgr9ivhehwu37fml20"
            set sp-entity-id "http://172.16.106.74/metadata/"
            set sp-single-sign-on-url "https://172.16.106.74/saml/?acs"
            set sp-single-logout-url "https://172.16.106.74/saml/?sls"
            set sp-portal-url "https://172.16.106.74/saml/login/"
            config assertion-attributes
                edit "username"
                next
                edit "tdoc@fortinet.com"
                    set type email
                next
            end
        next
    end
end
```

To configure an SP:

```
config system saml
    set status enable
    set cert "Fortinet_Factory"
    set idp-entity-id "http://172.16.106.74/saml-idp/csf_
ngczjwqxujfsbhgr9ivhehwu37fml20/metadata/"
    set idp-single-sign-on-url "https://172.16.106.74/csf_
ngczjwqxujfsbhgr9ivhehwu37fml20/login/"
    set idp-single-logout-url "https://172.16.106.74/saml-idp/csf_
ngczjwqxujfsbhgr9ivhehwu37fml20/logout/"
    set idp-cert "REMOTE_Cert_1"
    set server-address "172.16.106.74:12443"
end
```

To configure an SSO administrator:

```
config system sso-admin
    edit "SSO-admin-name"
        set accprofile <SSO admin user access profile>
        set vdom <Virtual domain(s) that the administrator can access>
    next
end
```


Netskope Exhibit 1017

SAML SSO with pre-authorized FortiGates

You can set up SAML SSO authentication in a Security Fabric environment by starting with a root FortiGate that has one or more pre-authorized FortiGates.

After the initial configuration, you can add more downstream FortiGates to the Security Fabric, and they are automatically configured with default values for a service provider.

To set up basic SAML SSO for the Security Fabric:

1. Log in to the root FortiGate of the Security Fabric.
2. Go to *Security Fabric* > *Fabric Connectors* and double-click the *Security Fabric Setup* card.
3. Join two pre-authorized FortiGates to the root FortiGate.



4. Configure the IdP (see Configuring the root FortiGate as the IdP on page 179).
5. Configure the SPs (see Configuring a downstream FortiGate as an SP on page 180).

Navigating between Security Fabric members with SSO

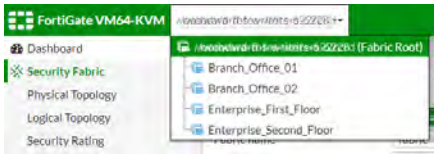
After you have logged in to a Security Fabric member by using SSO, you can navigate between any Security Fabric member with SSO configured. This can be done using the Security Fabric members dropdown menu or by logging in to a FortiGate SP from the root FortiGate IdP.

Security Fabric members dropdown

The Security Fabric members dropdown menu allows you to easily switch between all FortiGate devices that are connected to the Security Fabric. You can also use this menu to customize a FortiGate in the Security Fabric.

To navigate between Security Fabric members:

1. Log in to a Security Fabric member by using SSO.
2. In the top banner, click the name of the device you are logged into with SSO.
   A list of Security Fabric members is displayed.

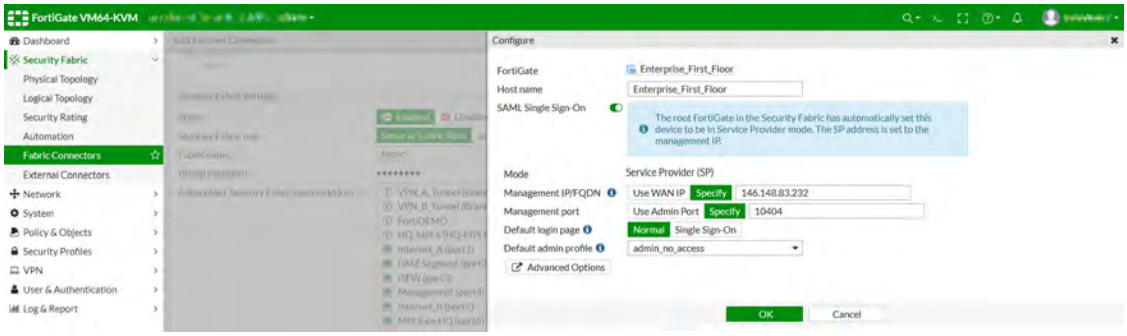


3. Click the Security Fabric member.
   You are logged in to the Security Fabric member without further authentication.

Netskope Exhibit 1017

To customize a FortiGate in the Security Fabric:

1. In the Security Fabric members dropdown menu, hover the cursor over a FortiGate so the tooltip is shown.
2. Click *Configure*. The *Configure* pane opens.



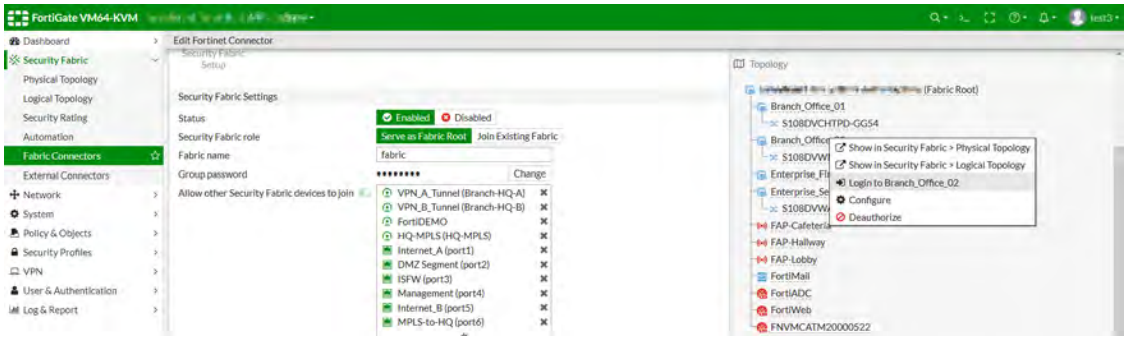
3. Edit the settings as required.
4. Click *OK*.

Logging in to an SP from the root IdP

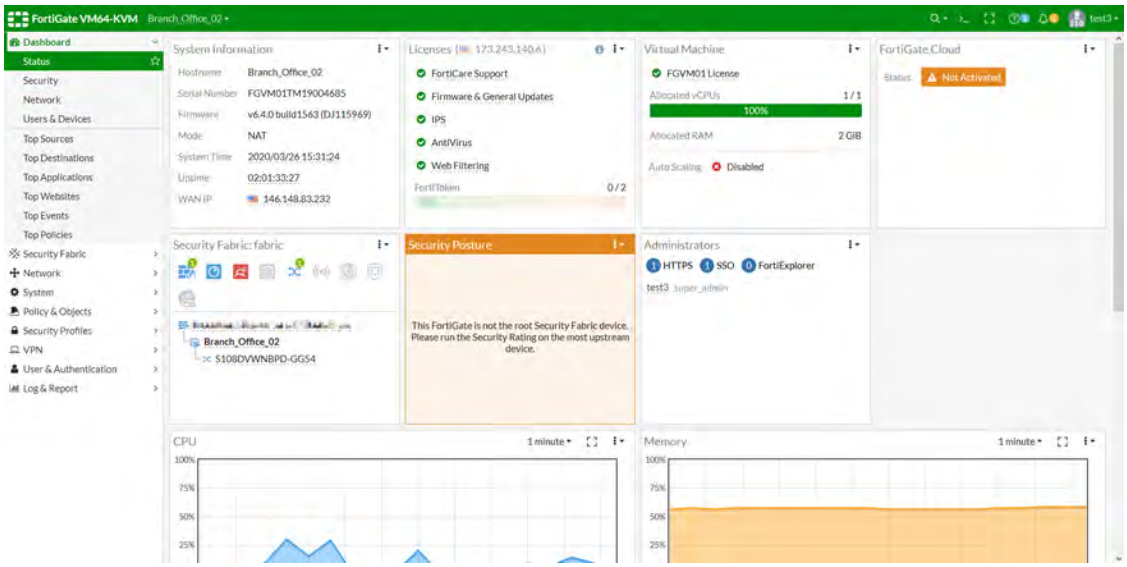
The following example describes how to log in to a root FortiGate IdP, and navigate to other FortiGate SPs in the Security Fabric without further authentication. The local administrator account is named *test3*. The local administrator account must also be available as an SSO administrator account on all downstream FortiGate SPs. Different tabs of the same browser are used to log in to the various FortiGates.

To log in to a FortiGate SP from a root FortiGate IdP:

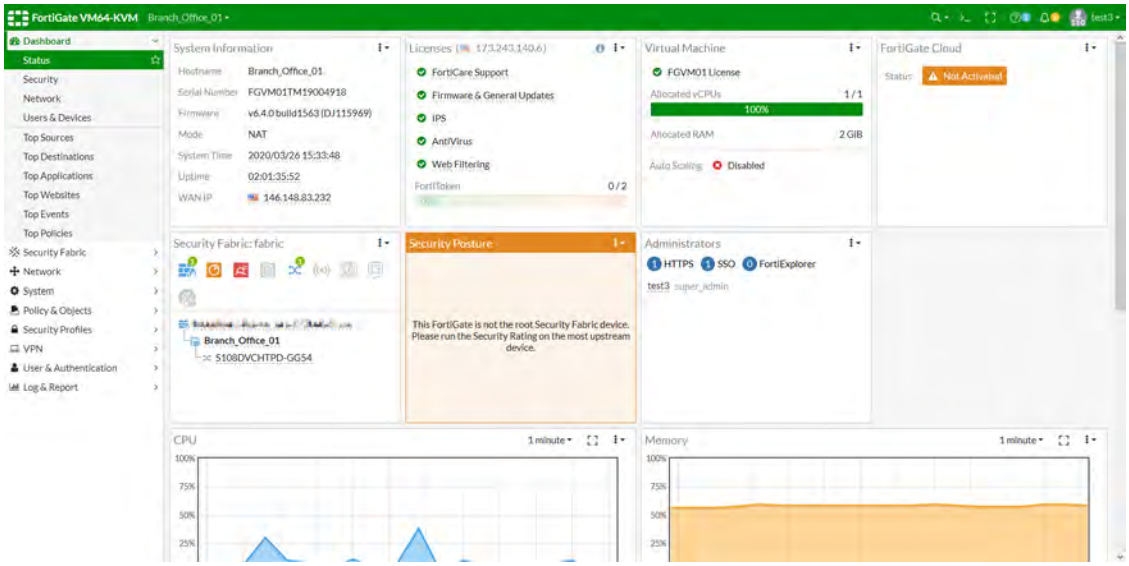
1. Log in to the root FortiGate IdP by using the local administrator account.
   In this example, the local administrator account is named *test3*.
2. Go to *Security Fabric* > *Fabric Connectors* and double-click the *Security Fabric Setup* card.
3. In the *Topology* tree, click one of the downstream FortiGate SPs, and select *Login to <name of FortiGate>*.



   The login screen is displayed.
4. In the login screen, select *Single Sign-On*.
   By using cookies in your local browser for the already-authenticated SSO administrator, FortiGate logs you in to the downstream FortiGate SP as the SSO administrator. In this example, the SSO administrator name is *test3*.



5. While still logged into the root FortiGate IdP in your browser, go to the browser tab for the root FortiGate IdP, and log in to another FortiGate SP that is displayed on the *Security Fabric* pane in the GUI.



SAML SSO login uses *SAML_IDP* session cookies of already authenticated admin users in your local browser cache to send to the root FortiGate IdP for authentication. If your browser cache is manually cleared, or you close your browser, you must authenticate again.

It is possible to log in to one downstream FortiGate SP in a Security Fabric, and then open another tab in your browser to connect to another FortiGate SP that is not a member of the Security Fabric.

This is useful in cases where the SSO administrator and the local system administrator on the FortiGate SP both have the same login name, but are two different entities.

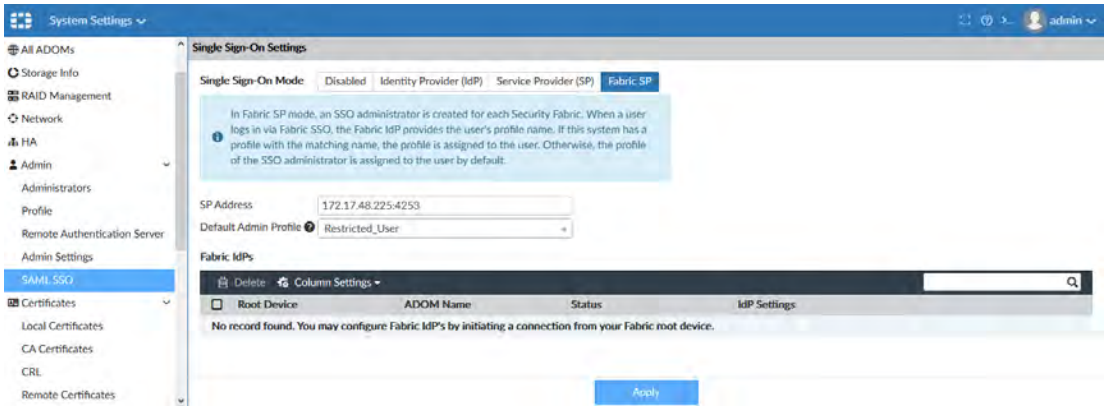
Integrating FortiAnalyzer management using SAML SSO

When a FortiGate is configured as the SAML SSO IdP, FortiAnalyzer can register itself as the SP (FortiAnalyzer must be running version 6.4.0). Once registered, FortiAnalyzer will be added automatically to the Security Fabric navigation in FortiOS. A similar dropdown navigation is displayed in FortiAnalyzer where users can navigate to the FortiGate using SAML SSO.

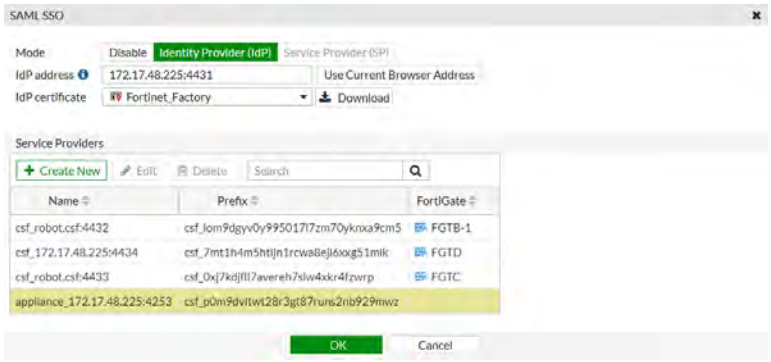
The following example assumes the root FortiGate (FGTA-1, server address 172.17.48.225:4431) has been configured as the SAML SSO IdP, and FortiAnalyzer logging has been enabled in the Security Fabric settings.

To enable FortiAnalyzer as a Fabric SP in the GUI:

1. In FortiAnalyzer, go to *System Settings* > *Admin* > *SAML SSO*.
2. For *Single Sign-On Mode*, click *Fabric SP* and enter the *SP Address*.



3. Click *Apply*.
   FortiAnalyzer will automatically register itself on the FortiGate as an appliance visible in the list of SPs. Go to *Security Fabric* > *Fabric Connectors*, edit the *Security Fabric Setup* connector, then click *Advanced Options* to view the list of SPs.



To enable FortiAnalyzer as a Fabric SP in the CLI:

1. In FortiAnalyzer, enable the device as a Fabric SP:

```
config system saml
    set status enable
    set role FAB-SP
```

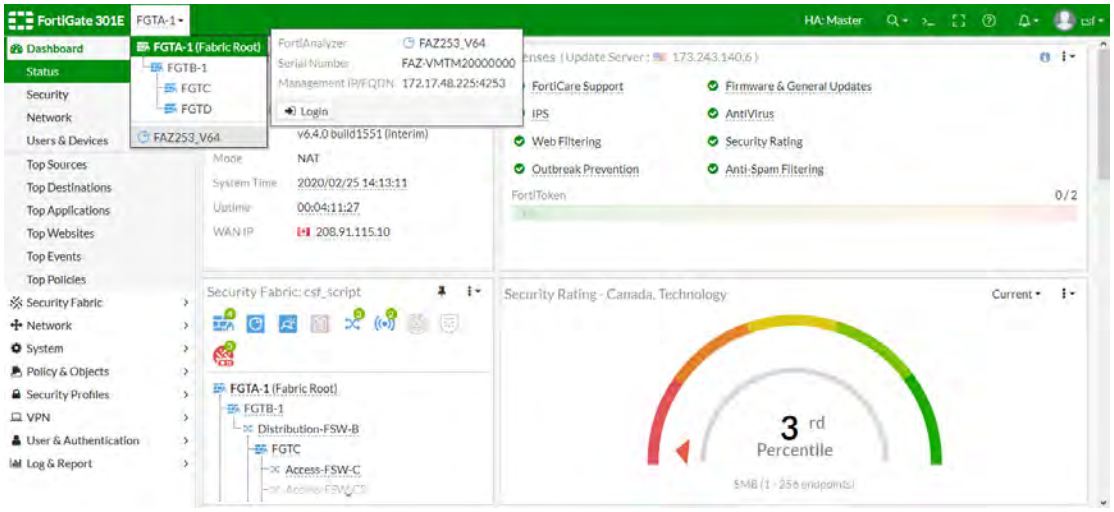
```
    set server-address "172.17.48.225:4253"
end
```

FortiAnalyzer will register itself on the FortiGate as an appliance. To view the configuration in FortiOS:

```
show system saml
    config service-providers
        edit "appliance_172.17.48.225:4253"
            set prefix "csf_p0m9dvltwt28r3gt87runs2nb929mwz"
            set sp-entity-id "http://172.17.48.225:4253/metadata/"
            set sp-single-sign-on-url "https://172.17.48.225:4253/saml/?acs"
            set sp-single-logout-url "https://172.17.48.225:4253/saml/?sls"
            set sp-portal-url "https://172.17.48.225:4253/saml/login/"
            config assertion-attributes
                edit "username"
                next
                edit "profilename"
                    set type profile-name
                next
            end
        next
    end
```

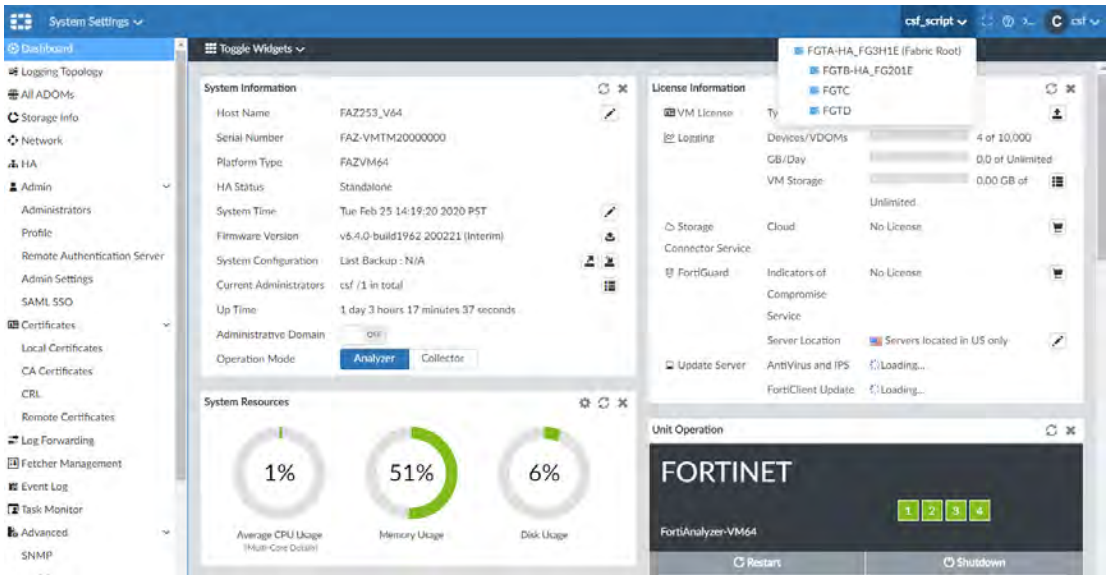
To navigate between devices using SAML SSO:

1. Log in to the root FortiGate.
2. In the toolbar, click the device name to display the Security Fabric members dropdown.
3. Hover over the FortiAnalyzer and click *Login*.



4. Log in to the FortiAnalyzer using SAML SSO.

5. In the toolbar, click the Security Fabric members dropdown to navigate between other FortiGates.

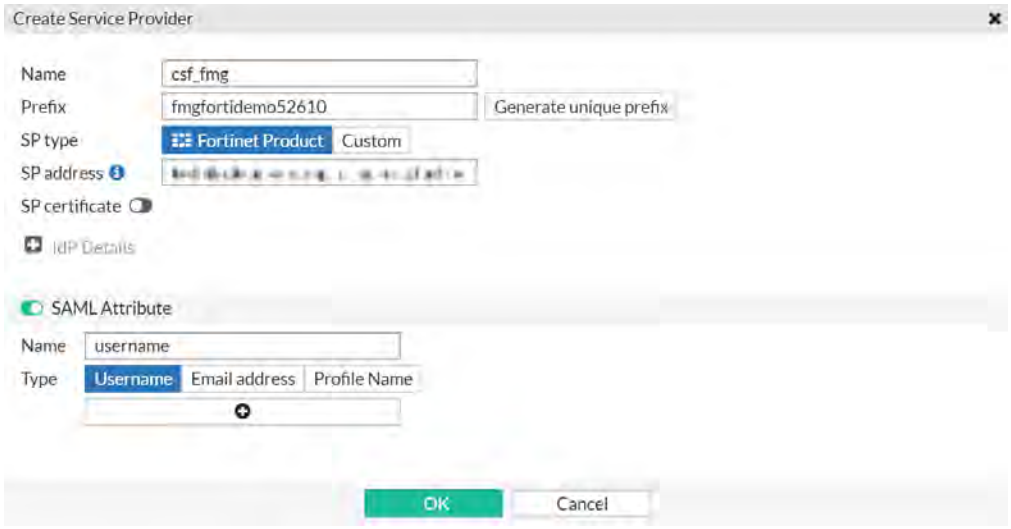


Integrating FortiManager management using SAML SSO

When a FortiGate is configured as the SAML SSO IdP, FortiManager can be added as an SP.

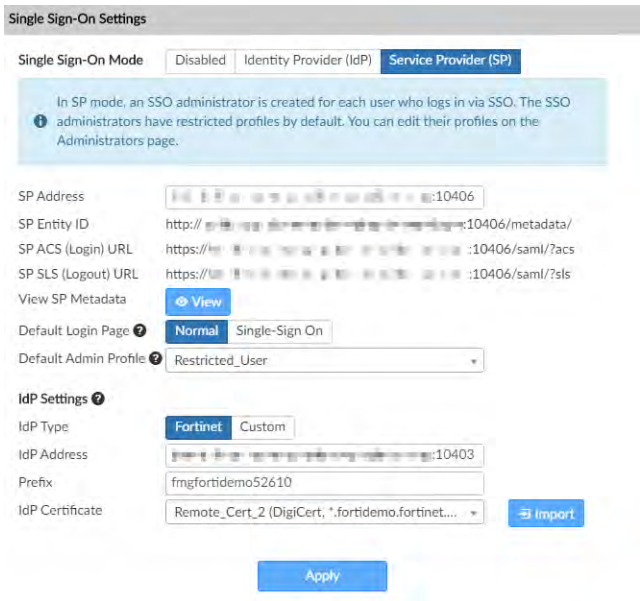
To configure FortiManager as a Fabric SP:

1. On the root FortiGate, go to *Security Fabric* > *Fabric Connectors*, and edit the *Security Fabric Setup* connector.
2. In the *Security Fabric Settings* section, click *Advanced Options*.
3. In the *Service Providers* section, click *Create New*.
4. Enter a name and a prefix for the SP. FortiOS generates a unique prefix, but you can enter your own.
5. In *SP address*, enter the FortiManager address including the port number.

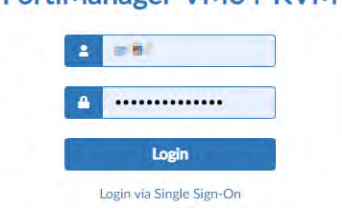


6. Click *OK*.
7. In FortiManager, go to *System Settings* > *Admin* > *SAML SSO* and in the *Single Sign-On Mode* section, click *Service Provider (SP)*.

8. Configure the *IdP Settings*:
   a. For *IdP Type*, click *Fortinet*.
   b. For *IdP Address*, enter the root FortiGate address including the port number.
   c. Enter the *Prefix* of the SP.
   d. For *IdP Certificate*, import the same certificate used on the root FortiGate.
   e. Click *Apply*.



9. To verify that the configuration works, log out of FortiManager and log in using the *Login via Single-Sign-On* link.

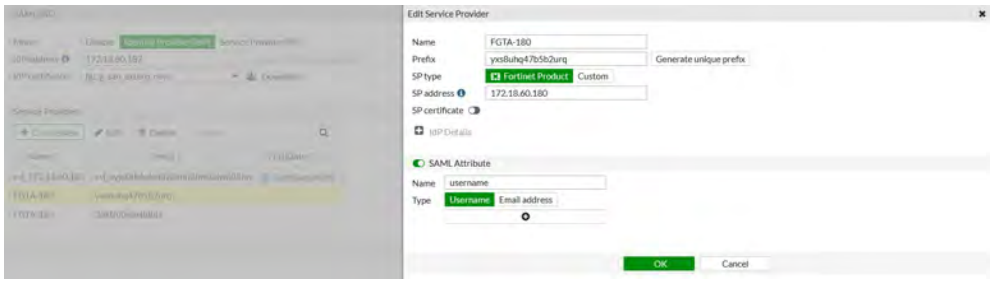


Advanced option - FortiGate SP changes

From a root FortiGate IdP, you can edit each of the FortiGate SPs. For example, you can edit a FortiGate SP to generate a new prefix, or you can add or modify SAML attributes. When you generate a new prefix value, it is propagated to the respective downstream FortiGates.

To edit an SP from the root FortiGate (IdP):

1. Go to *Security Fabric* > *Fabric Connectors* and double-click the *Security Fabric Setup* card.
2. Click *Advanced Options*. The *SAML SSO* pane opens.
3. In the *Service Providers* table, select a device, and click *Edit*. The *Edit Service Provider* pane opens.
4. Edit the settings as needed.
5. Click *OK*.



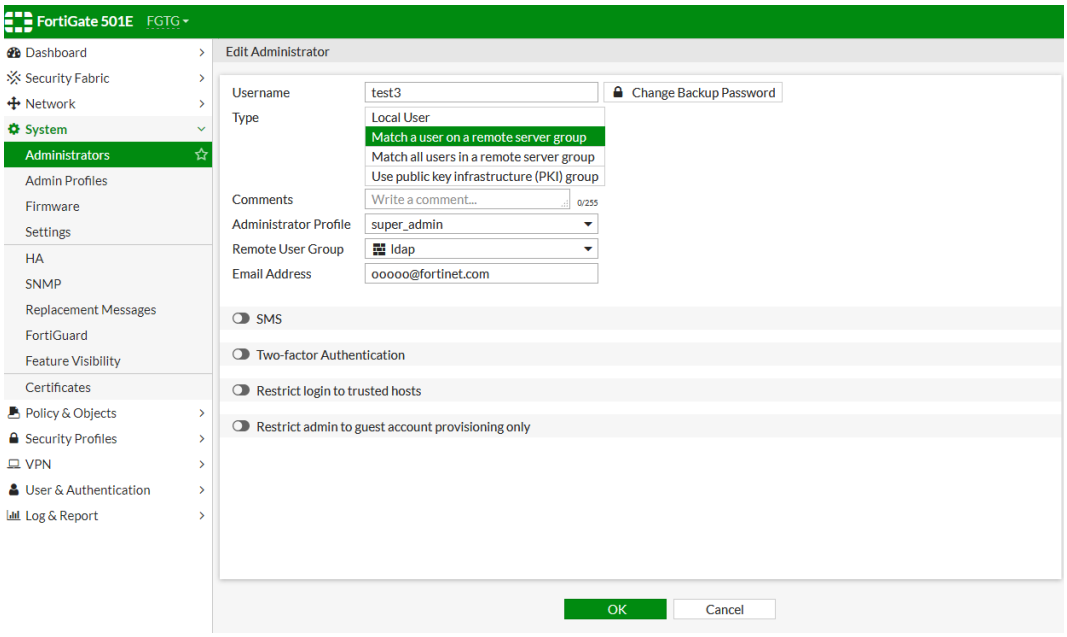
Advanced option - unique SAML attribute types

The default SAML attribute type is *username*. When the attribute type is set to *username*, SSO administrator accounts created on FortiGate SPs use the login username that is provided by the user for authentication on the root FortiGate IdP.

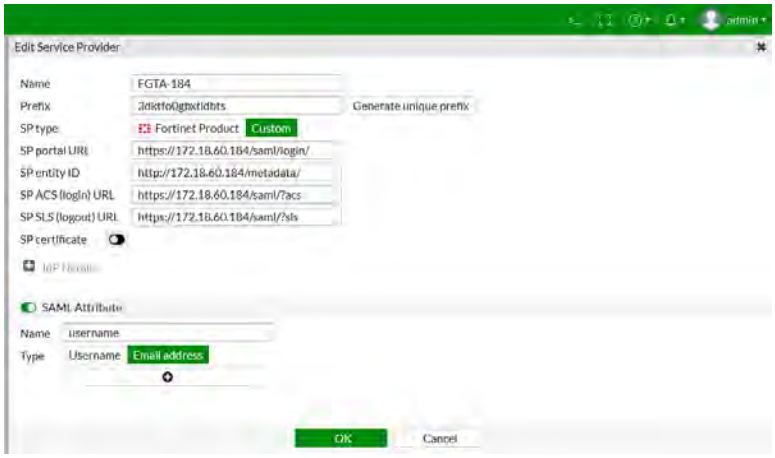
Because user names might not be unique, cases can occur where the user name is the same for the SSO administrator and the local administrator on the FortiGate SP. As a result, you might be unable to distinguish between actions taken by the local administrator and the SSO administrator on the FortiGate SP when looking at the system log. By using a unique SAML attribute type, such as an email address, you can create unique user names to better track what actions were taken by each administrator.

To configure a unique SAML attribute using the GUI:

1. On the root FortiGate (IdP), assign a unique email address to local administrator.
   In this example, the local administrator name is *test3*.
   a. Go to *System* > *Administrators*, and expand the list of local users.
   b. Select the local user, and click *Edit*.
   c. In the *Type* field, select *Match a user on a remote server group*.
   d. In the *Remote User Group* field, select a group.
   e. In the *Email Address* field, enter the email address.
   f. Click *OK*.



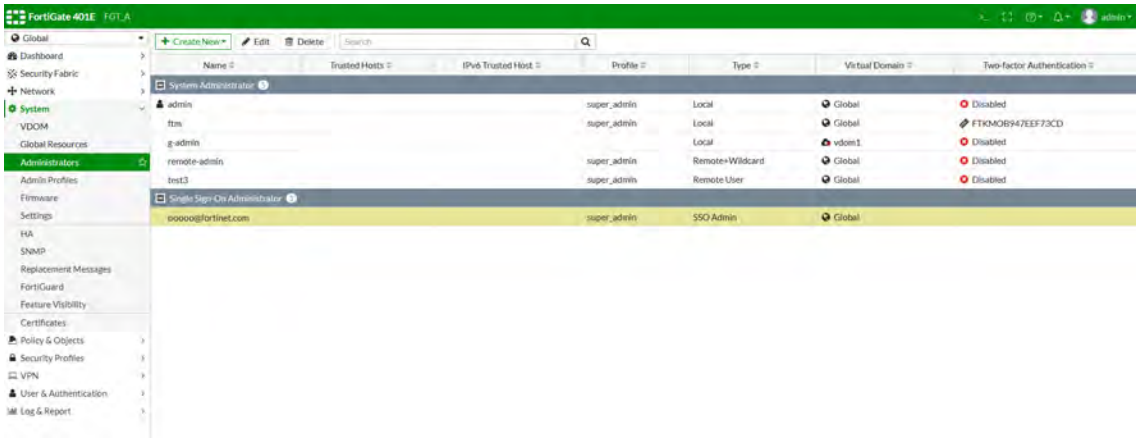
2. On the root FortiGate (IdP), update the SAML configuration:
   a. Go to *Security Fabric* > *Fabric Connectors* and double-click the *Security Fabric Setup* card.
   b. Click *Advanced Options*. The *SAML SSO* pane opens.
   c. In the *Service Providers* table, select the FortiGate, and click *Edit*. The *Edit Service Provider* pane opens.
   d. For *SP type*, select *Custom*.
   e. In the *SAML Attribute* section for *Type*, select *Email address*.
   f. Beside *Type*, select *Email address*.
   g. Click *OK*.



After the administrator (test3) logs in to the FortiGate SP for the first time, SAML authentication occurs on FortiGate SP. A new SSO administrator account is created, and the account name is now the email address instead of the login name (test3).

To view the new SSO administrator account:

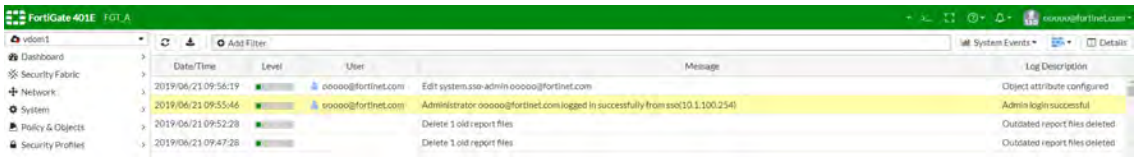
1. In the SP, go to *System* > *Administrators*, and expand the list of SSO administrators.
   The email address (*ooooo@fortinet.com*) is listed as the account name:



   If the SAML attribute had been set to the default setting of *username*, the user name for the SSO administrator account would have been (*test3*).

To view the SSO administrator activity in the log files:

1. In the SP, go to *Log & Report* > *Events*.
   Because the SAML attribute is set to *Custom*, the SSO administrator account *ooooo@fortinet.com* is used as the user name on the FortiGate SP, and it appears in the log files:



To configure a unique SAML attribute using the CLI:

```
config system saml
   set status enable
   set role identity-provider
   set cert "fgt_g_san_extern_new"
   set server-address "172.18.60.187"
   config service-providers
      edit "csf_172.18.60.185"
         set prefix "csf_avju0tk4oiodifz3kbh2fms8dw688hn"
         set sp-entity-id "http://172.18.60.185/metadata/"
         set sp-single-sign-on-url "https://172.18.60.185/saml/?acs"
         set sp-single-logout-url "https://172.18.60.185/saml/?sls"
         set sp-portal-url "https://172.18.60.185/saml/login/"
         config assertion-attributes
            edit "username"
            next
         end
      next
      edit "FGTA-180"
         set prefix "yxs8uhq47b5b2urq"
         set sp-entity-id "http://172.18.60.180/metadata/"
```


Netskope Exhibit 1017

```
            set sp-single-sign-on-url "https://172.18.60.180/saml/?acs"
            set sp-single-logout-url "https://172.18.60.180/saml/?sls"
            set sp-portal-url "https://172.18.60.180/saml/login/"
            config assertion-attributes
               edit "username"
               next
            end
         next
         edit "FGTA-184"
            set prefix "3dktfo0gbxtldbts"
            set sp-entity-id "http://172.18.60.184/metadata/"
            set sp-single-sign-on-url "https://172.18.60.184/saml/?acs"
            set sp-single-logout-url "https://172.18.60.184/saml/?sls"
            set sp-portal-url "https://172.18.60.184/saml/login/"
            config assertion-attributes
               edit "username"
                  set type email
               next
            end
         next
      end
   end
```

The `csf_172.18.60.185` service provider was automatically added when the FortiGate SP 172.18.60.185 joined the root FortiGate IdP in the Security Fabric.

All `sp-*` options, such as `sp-portal-url`, are set with default values when a service provider is created, but can be modified using the CLI or GUI.

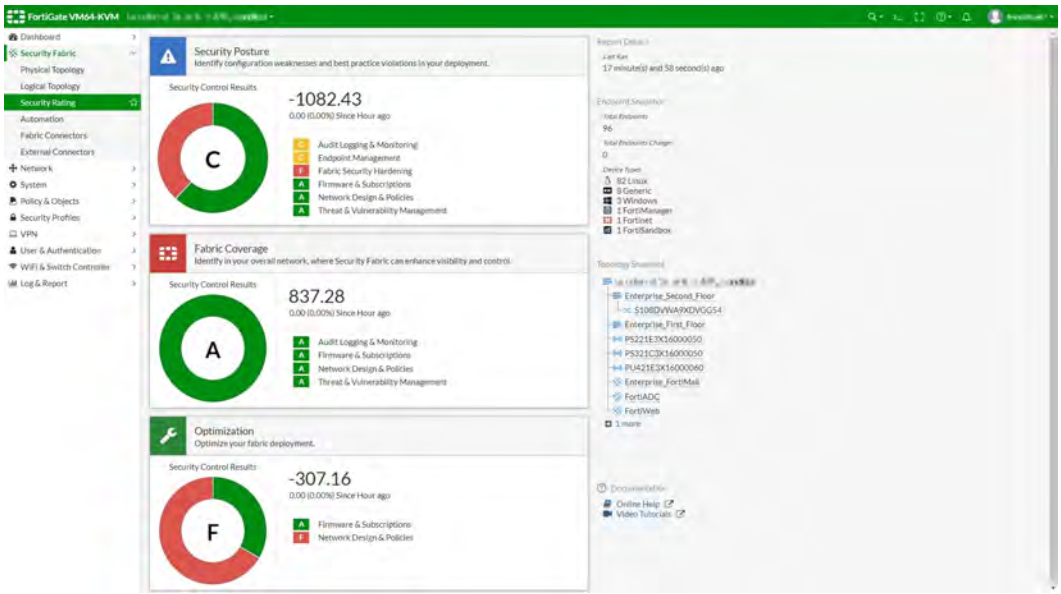
Security rating

The security rating uses real-time monitoring to analyze your Security Fabric deployment, identify potential vulnerabilities, highlight best practices that can be used to improve the security and performance of your network, and calculate Security Fabric scores.

To view the security rating, go to *Security Fabric* > *Security Rating* on the root FortiGate.

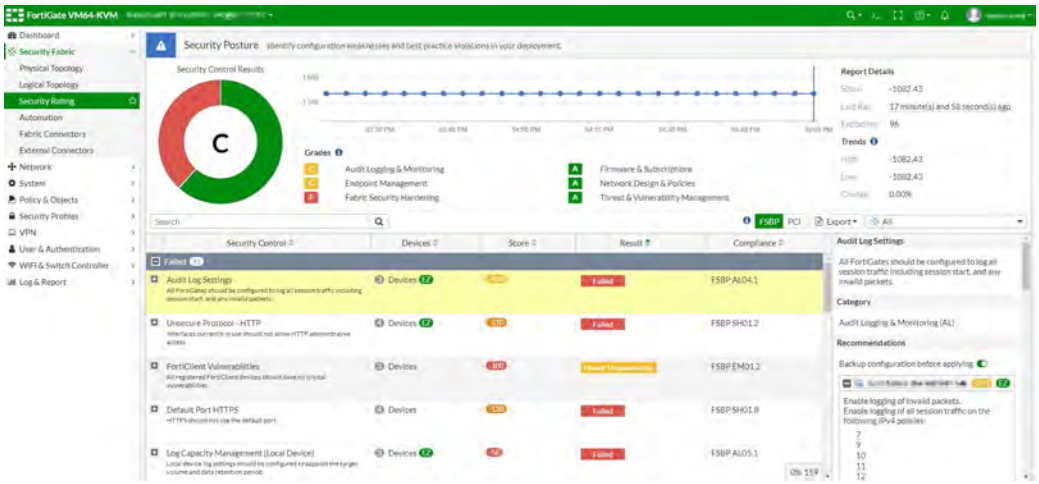
The *Security Rating* page is separated into three major scorecards: *Security Posture*, *Fabric Coverage*, and *Optimization*, which provide an executive summary of the three largest areas of security focus in the Security Fabric.

Netskope Exhibit 1017



The scorecards show an overall letter grade and breakdown of the performance in sub-categories. Clicking a scorecard drills down to a detailed report of itemized results and compliance recommendations. The point score represents the net score for all passed and failed items in that area. In the drill down report, hover the cursor over a score to view the calculation breakdown.

The report includes the security controls that were tested against, linking to specific FSBP or PCI compliance policies. Click the *FSBP* and *PCI* buttons to reference the corresponding standard. Users can search or filter the report results.



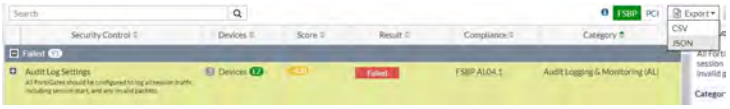
Certain remediations marked with an *EZ* symbol represent configuration recommendations that support *Easy Apply*. In the panel on the right, in the *Recommendations* section, click *Apply* to apply the changes to resolve the failed security control.



The report table can be customized by adding more columns, such as *Category*, to view, filter, or sort the results based on scorecard categories. Click the gear icon to customize the table.



Users can also export the reports as CSV or JSON files by clicking the *Export* dropdown.



To exit the current view, click the icon beside the scorecard title to return to the summary view.

For more information about security ratings, and details about each of the checks that are performed, go to Security Best Practices & Security Rating Feature.

Security rating licenses are required to run security rating checks across all the devices in the Security Fabric. It also allows ratings scores to be submitted to and received from FortiGuard for ranking networks by percentile.

See https://www.fortinet.com/support/support-services/fortiguard-security-subscriptions/security-rating.html for information.

Security rating check scheduling

Security rating checks by default are scheduled to run automatically every four hours.

To disable automatic security checks using the CLI:

```
config system global
    security-rating-run-on-schedule disable
end
```

To manually run a report using the CLI:

```
# diagnose report-runner trigger
```

Opt out of ranking

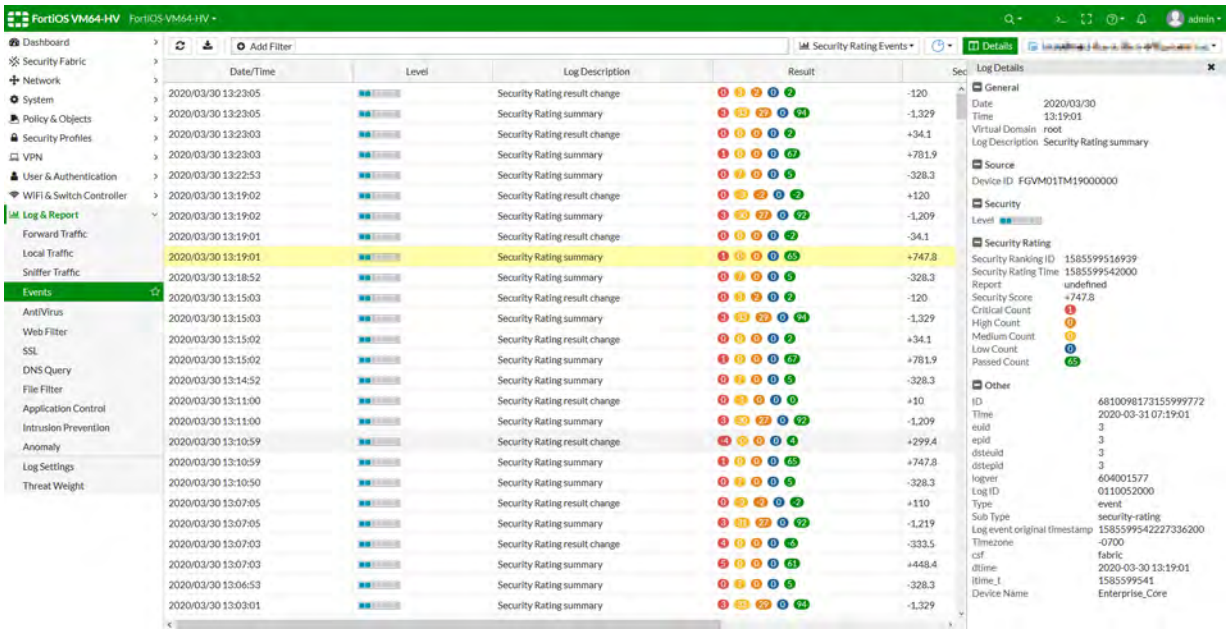
Security rating scores can be submitted to FortiGuard for comparison with other organizations' scores, allowing a percentile score to be calculated. If you opt out of submitting your score, only an absolute score will be available.

To opt out of submitting the score using the CLI:

```
config system global
    set security-rating-result-submission {enable | disable}
end
```

Logging the security rating

The results of past security checks is available in *Log & Report* > *Events* by selecting *Security Rating Events* from the event type dropdown list.



An event filter subtype can be created for the Security Fabric rating so that event logs are created on the root FortiGate that summarize the results of a check, and show detailed information for the individual tests.

To configure security rating logging using the CLI:

```
config log eventfilter
    set security-rating enable
end
```

Security Fabric score

The Security Fabric score is calculated when a security rating check is run, based on the severity level of the checks that are passed or failed. A higher scores represents a more secure network. Points are added for passed checks and removed for failed checks.

Severity level	Weight (points)
Critical	50

Severity level	Weight (points)
High	25
Medium	10
Low	5

To calculate the number of points awarded to a device for a passed check, the following equation is used:

$$\text{score} = \frac{\text{<severity level weight>}}{\text{<\# of FortiGates>}} \times \text{<secure FortiGate multiplier>}$$

The secure FortiGate multiplier is determined using logarithms and the number of FortiGate devices in the Security Fabric.

For example, if there are four FortiGate devices in the Security Fabric that all pass the compatible firmware check, the score for each FortiGate device is calculated with the following equation:

$$\frac{50}{4} \times 1.292 = 16.15 \text{ points}$$

All of the FortiGate devices in the Security Fabric must pass the check in order to receive the points. If any one of the FortiGate devices fails a check, the devices that passed are not awarded any points. For the device that failed the check, the following equation is used to calculated the number of points that are lost:

$$\text{score} = \text{<severity level weight>} \times \text{<secure FortiGate multiplier>}$$

For example, if the check finds two critical FortiClient vulnerabilities, the score is calculated with the following equation:

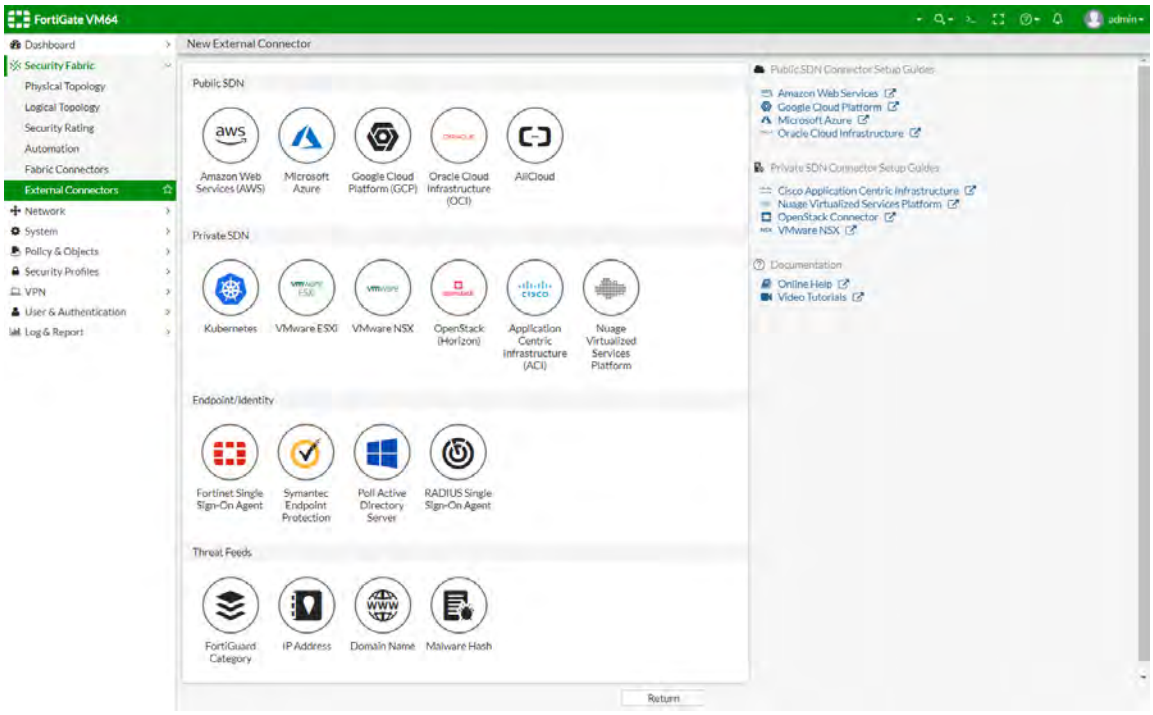
$$-50 \times 2 = -100 \text{ points}$$

Scores are not affected by checks that do not apply to your network. For example, if there are no FortiAP devices in the Security Fabric, no points will be added or subtracted for the FortiAP firmware version check.

External connectors

The following external connector categories are available in the Security Fabric: Public SDN, Private SDN, Endpoint/Identity, and Threat Feeds.

Netskope Exhibit 1017



If VDOMs are enabled, SDN and Threat Feeds connectors are in the global settings, and Endpoint/Identity connectors are per VDOM.

SDN connectors

Fabric connectors to SDNs provide integration and orchestration of Fortinet products with SDN solutions. Fabric connectors ensure that any changes in the SDN environment are automatically updated in your network.

There are four steps to creating and using an SDN connector:

1. Gather the required information
2. Create the fabric connector on page 200
3. Create a fabric connector address on page 201
4. Add the address to a firewall policy on page 202

An example of creating a Microsoft Azure SDN connector is available at https://docs.fortinet.com/vm/azure/fortigate/6.4/azure-cookbook/6.4.0/502895/configuring-a-fabric-connector-in-azure.

Create the fabric connector

To create an SDN Fabric connector in the GUI:

1. Go to *Security Fabric* > *External Connectors*.
2. Click *Create New*.
3. Click the service that you are using.

4. Enter the *Name*, *Status*, and *Update Interval* for the connector.
5. Enter the previously collected information for the specific connector that you are creating.
6. Click *OK*.

To create an SDN Fabric connector in the CLI:

```
config system sdn-connector
    edit <name>
        set status {enable | disable}
        set type {connector type}
        ...
        set update-interval <integer>
    next
end
```



The available CLI commands will vary depending on the selected SDN connector type.

Create a fabric connector address

A fabric connector address can be used in the following ways:

- As the source or destination address for firewall policies.
- To automatically update changes to addresses in the environment of the service that you are using, based on specified filtering conditions.
- To automatically apply changes to firewall policies that use the address, based on specified filtering conditions.

To create a fabric connector address in the GUI:

1. Go to *Policy & Objects* > *Addresses*.
2. Click *Create New* > *Address*.
3. Enter a name for the address.
4. Set the *Type* to *Fabric Connector Address*.
5. Select an *SDN Connector* from the dropdown list, or click *Create New* to make a new one.
6. Set the *SDN address type*. Only addresses of the selected type will be collected.
7. Configure the connector specific settings.
8. Select an *Interface* for the address, or leave it as *any*, enable or disable *Show in Address List*, and optionally add *Comments*.
9. Add tags.
10. Click *OK*.

To create a fabric connector address in the CLI:

```
config firewall address
    edit <name>
        set type dynamic
        set sdn <sdn_connector>
        set visibility enable
```


Netskope Exhibit 1017

```
        set associated-interface <interface_name>
        set color <integer>
        ...
        set comment <comment>
        config tagging
            edit <name>
                set category <string>
                set tags <strings>
            next
        end
    next
end
```



The available CLI commands will vary depending on the selected SDN connector type.

Add the address to a firewall policy

A fabric connector address can be used as either the source or destination address.

To add the address to a firewall policy in the GUI:

1. Go to *Policy & Objects* > *Firewall Policy*.
2. Click *Create New*.
3. Enter a name for the policy.
4. Set the incoming and outgoing interfaces.
5. Use the fabric connector address as the source or destination address.
6. Configure the remaining settings as needed.
7. Click *OK*.

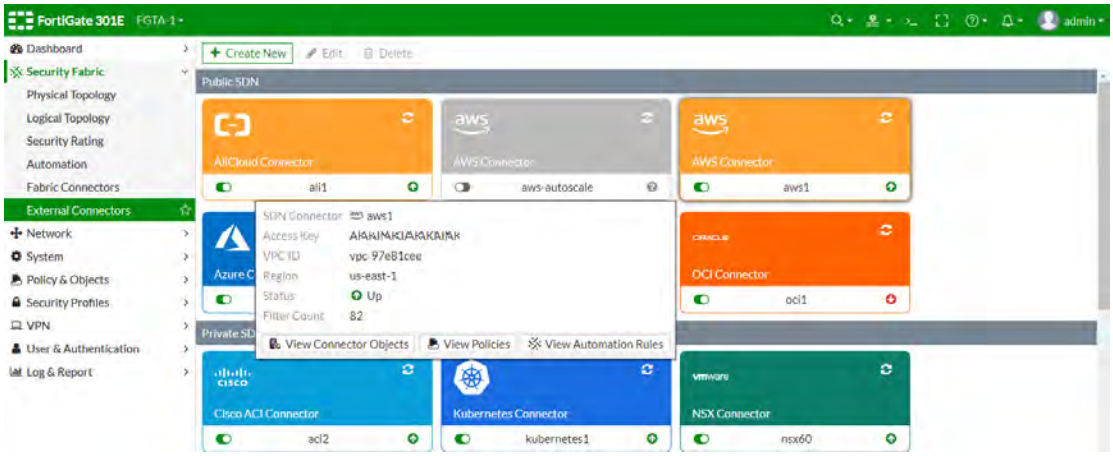
To add the address to a firewall policy in the CLI:

```
config firewall policy
    edit 0
        set name <name>
        set srcintf <port_name>
        set dstintf <port_name>
        set srcaddr <firewall_address>
        set dstaddr <firewall_address>
        set action accept
        set schedule <schedule>
        set service <service>
    next
            end
```


Netskope Exhibit 1017

Connector tooltips

In the *Security Fabric* > *External Connectors* page, hover over an SDN connector to view a tooltip that shows basic configuration information.



There are three buttons that provide additional information:

- *View Connector Objects* shows the connector's dynamic objects, such as filters and instances.
- *View Policies* shows a list of policies that use the dynamic addresses from the connector.
- *View Automation Rules* shows a list of automation actions that use the connector.

AliCloud SDN connector

FortiOS automatically updates dynamic addresses for AliCloud using an AliCloud SDN connector, including mapping the following attributes from AliCloud instances to dynamic address groups in FortiOS:

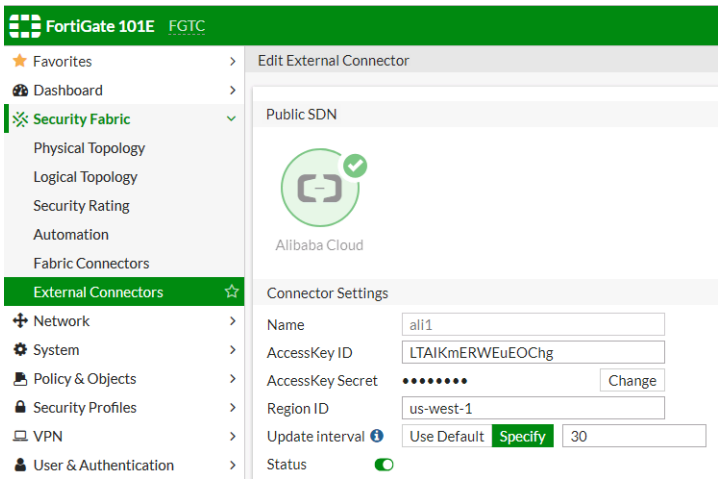
- ImageId
- InstanceId
- SecurityGroupId
- VpcId
- VSwitchId
- TagKey
- TagValue

To configure AliCloud SDN connector using the GUI:

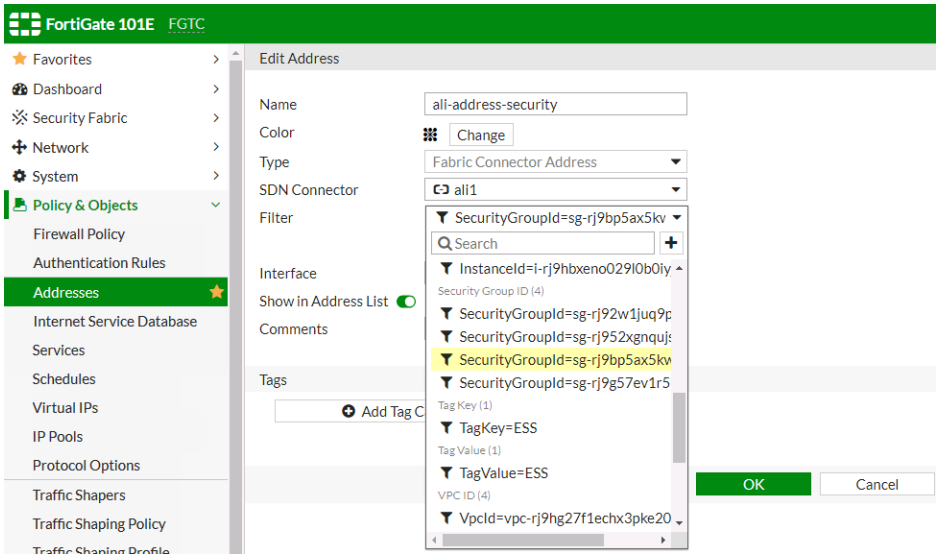
1. Configure the AliCloud SDN connector:
   a. Go to *Security Fabric* > *External Connectors*.
   b. Click *Create New*, and select *AliCloud*.
   c. Configure as shown, substituting the access key, secret, and region ID for your deployment. The update

Netskope Exhibit 1017

interval is in seconds.

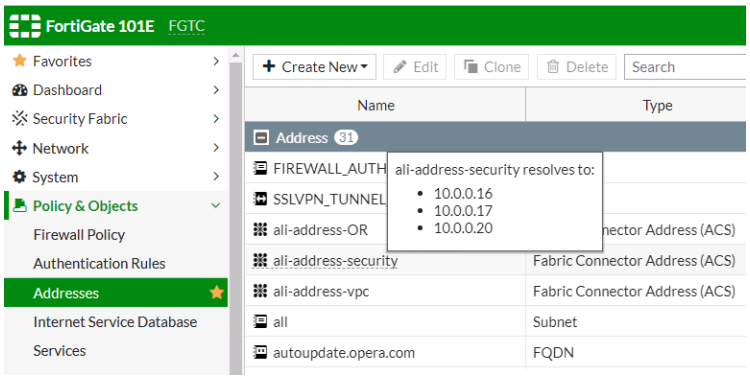


2. Create a dynamic firewall address for the configured AliCloud SDN connector:
   a. Go to *Policy & Objects* > *Addresses*.
   b. Click *Create New*, then select *Address*.
   c. Configure the address as shown, selecting the desired filter in the *Filter* dropdown list. In this example, the AliCloud SDN connector will automatically populate and update IP addresses only for instances that belong to the specified security group:



3. Ensure that the AliCloud SDN connector resolves dynamic firewall IP addresses:
   a. Go to *Policy & Objects* > *Addresses*.
   b. Hover over the address created in step 2 to see a list of IP addresses for instances that belong to the security

group configured in step 2:



To configure AliCloud SDN connector using CLI commands:

1. Configure the AliCloud SDN connector:

```
config system sdn-connector
   edit "ali1"
      set type acs
      set access-key "LTAIKmERWEuEOChg"
      set secret-key xxxxx
      set region "us-west-1"
      set update-interval 30
   next
end
```

2. Create a dynamic firewall address for the configured AliCloud SDN connector with the supported AliCloud filter. In this example, the AliCloud SDN Connector will automatically populate and update IP addresses only for instances that belong to the specified security group:

```
config firewall address
   edit "ali-address-security"
      set type dynamic
      set sdn "ali1"
      set filter "SecurityGroupId=sg-rj9bp5ax5kwy3gqdizqb"
   next
end
```

3. Confirm that the AliCloud SDN connector resolves dynamic firewall IP addresses using the configured filter:

```
config firewall address
   edit "ali-address-security"
      set type dynamic
      set sdn "ali1"
      set filter "SecurityGroupId=sg-rj9bp5ax5kwy3gqdizqb"
      config list
         edit "10.0.0.16"
         next
         edit "10.0.0.17"
         next
         edit "10.0.0.20"
         next
      end
   next
end
```

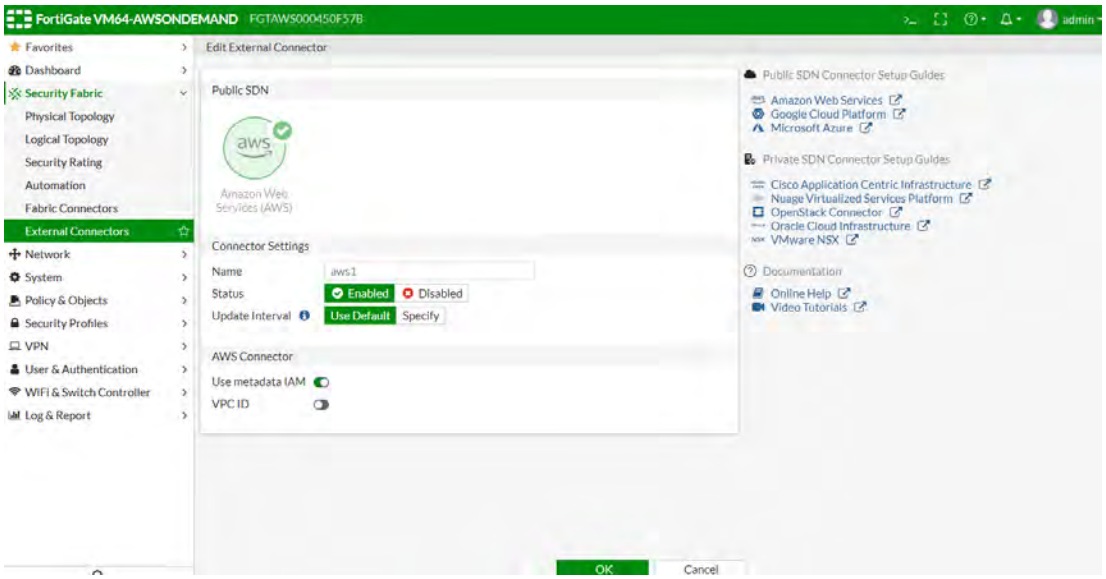
AWS SDN connector with IAM credentials

For instances running in AWS (on demand or BYOL), you can set up the AWS SDN connector using AWS Identify and Access Management (IAM) credentials.

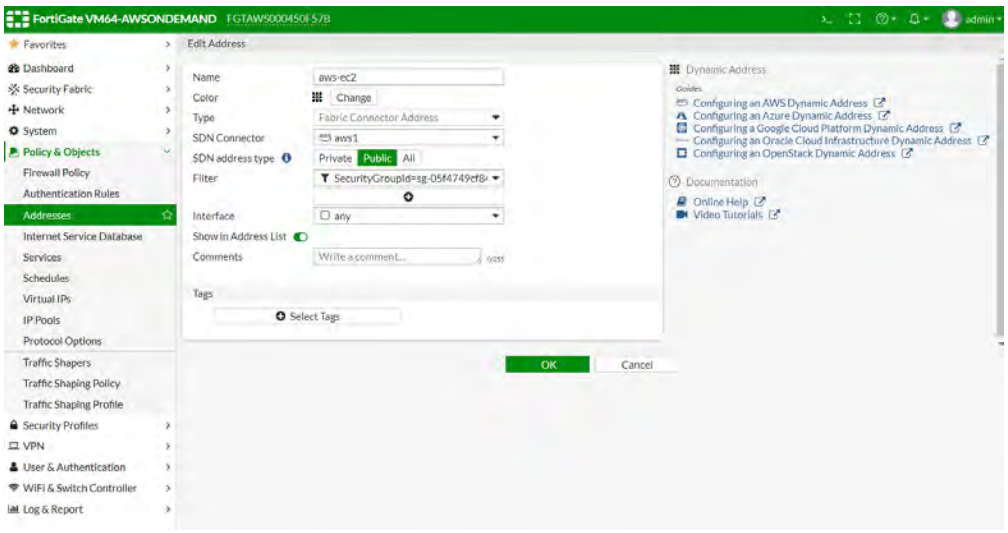
IAM authentication is available only for FGT-AWS and FGT-AWSONDEMAND platforms.

To configure AWS SDN connector using the GUI:

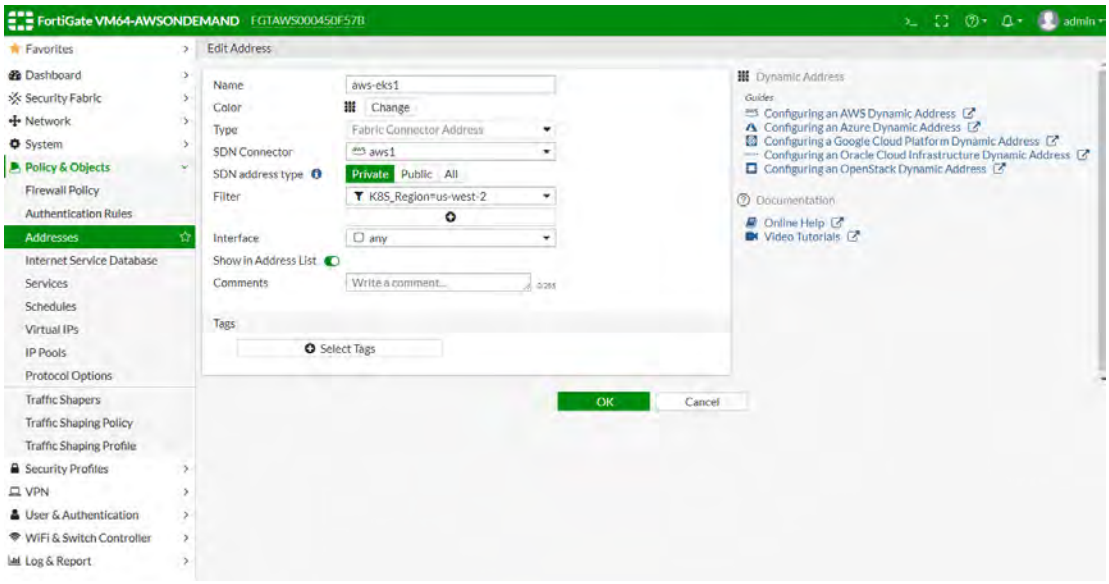
1. Configure the AWS SDN connector:
   a. Go to *Security Fabric* > *External Connectors*.
   b. Click *Create New*, and select *Amazon Web Services (AWS)*.
   c. Configure as shown:



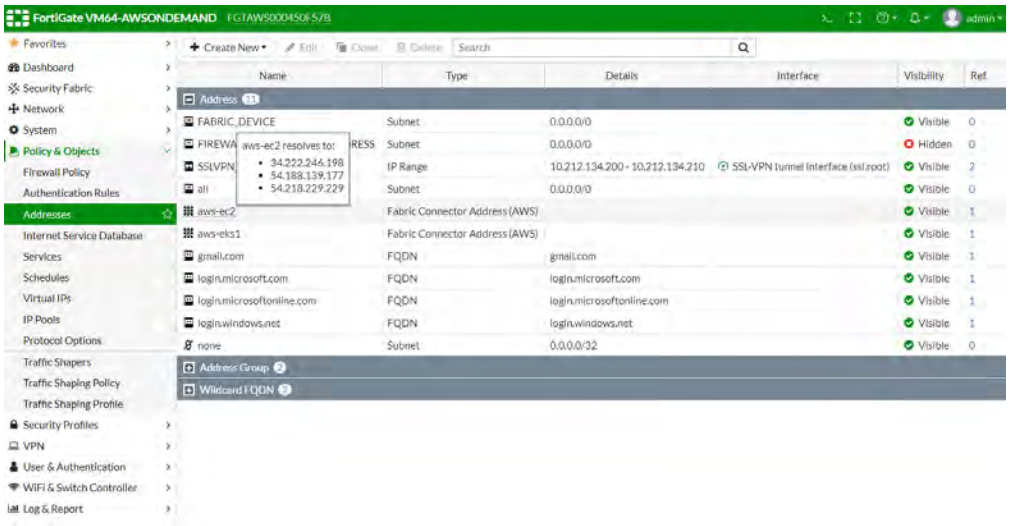
2. Create a dynamic firewall address for the configured AWS SDN connector:
   a. Go to *Policy & Objects* > *Addresses*.
   b. Click *Create New*, then select *Address*.
   c. Configure the address as shown, selecting the desired filter in the *Filter* dropdown list.
   The following is an example for a public SDN address type:



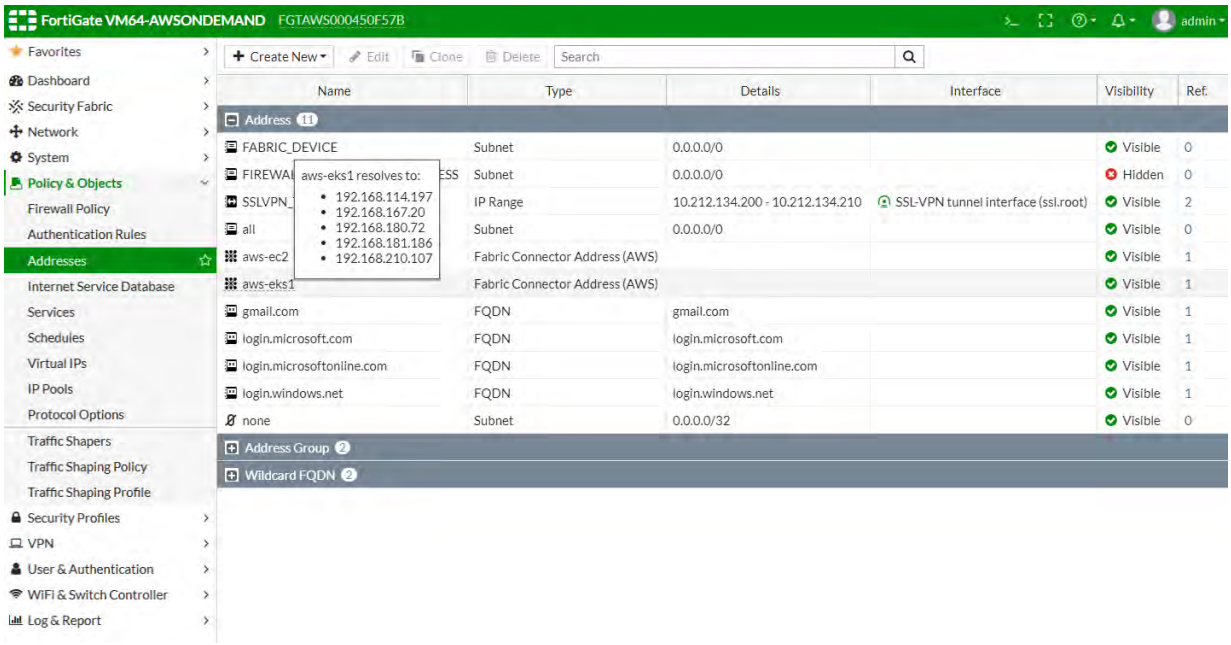
   The following is an example for a private SDN address type:



3. Ensure that the AWS SDN connector resolves dynamic firewall IP addresses:
   a. Go to *Policy & Objects* > *Addresses*.
   b. Hover over the address created in step 2 to see a list of IP addresses for instances that belong to the security group configured in step 2.
   The following is an example for a public SDN address type:



   The following is an example for a private SDN address type:



To configure AWS SDN connector using CLI commands:

1. Configure the AWS connector:

```
config system sdn-connector
   edit "aws1"
      set status enable
      set type aws
      set use-metadata-iam enable
      set update-interval 60
   next
end
```

2. Create a dynamic firewall address for the configured AWS SDN connector with the supported filter: Dynamic firewall address IPs are resolved by the SDN connector.

```
config firewall address
   edit "aws-ec2"
      set type dynamic
      set sdn "aws1"
      set filter "SecurityGroupId=sg-05f4749cf84267548"
      set sdn-addr-type public
   next
   edit "aws-eks1"
      set type dynamic
      set sdn "aws1"
      set filter "K8S_Region=us-west-2"
   next
end
```

3. Confirm that the AWS SDN connector resolves dynamic firewall IP addresses using the configured filter:

```
config firewall address
   edit "aws-ec2"
      set type dynamic
      set sdn "aws1"
      set filter "SecurityGroupId=sg-05f4749cf84267548"
      set sdn-addr-type public
```


Netskope Exhibit 1017

```
        config list
          edit "34.222.246.198"
          next
          edit "54.188.139.177"
          next
          edit "54.218.229.229"
          next
        end
      next
      edit "aws-eks1"
        set type dynamic
        set sdn "aws1"
        set filter "K8S_Region=us-west-2"
        config list
          edit "192.168.114.197"
          next
          edit "192.168.167.20"
          next
          edit "192.168.180.72"
          next
          edit "192.168.181.186"
          next
          edit "192.168.210.107"
          next
        end
      next
  end
```

Azure Stack SDN connector

FortiOS automatically updates dynamic addresses for Azure Stack on-premise environments using an Azure Stack SDN connector, including mapping the following attributes from Azure Stack instances to dynamic address groups in FortiOS:

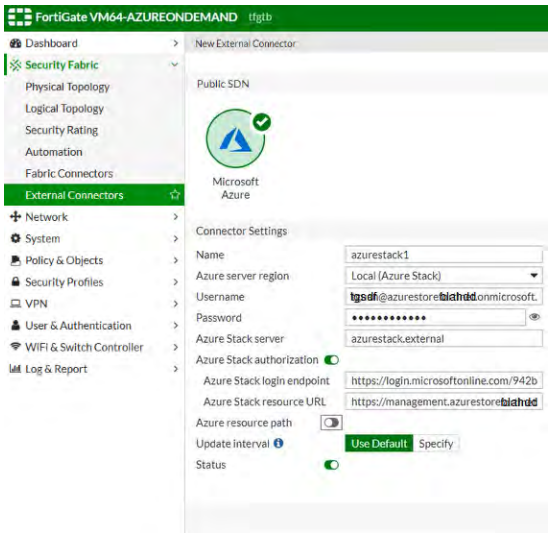
- vm
- tag
- size
- securitygroup
- vnet
- subnet
- resourcegroup
- vmss

To configure Azure Stack SDN connector using the GUI:

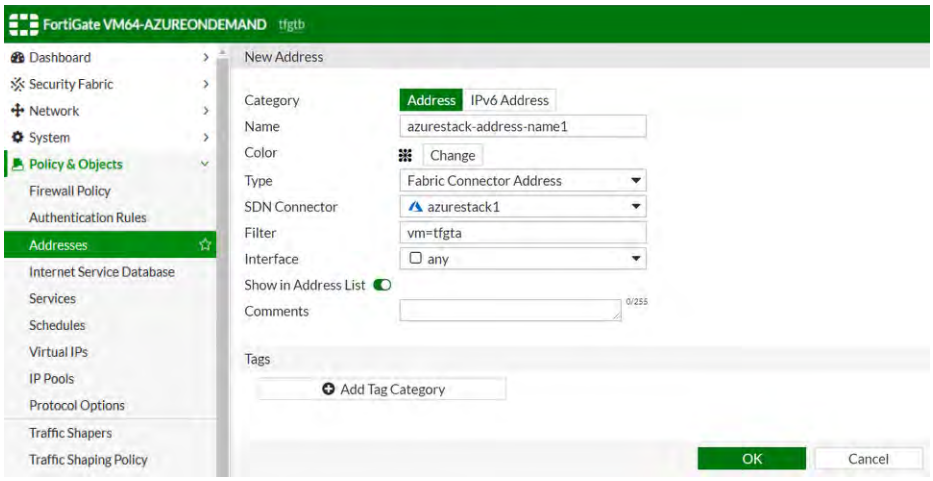
1. Configure the Azure Stack SDN connector:
   a. Go to *Security Fabric* > *External Connectors*.
   b. Click *Create New*, and select *Microsoft Azure*.
   c. Configure as shown, substituting the Azure Stack settings for your deployment. The update interval is in

Netskope Exhibit 1017

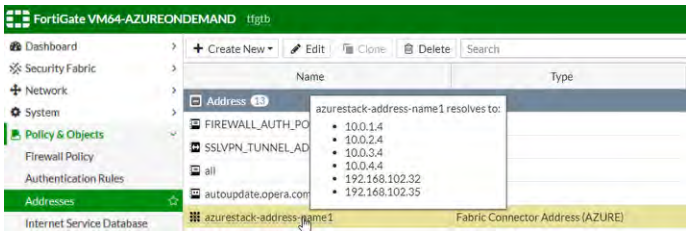
seconds.



2. Create a dynamic firewall address for the configured Azure Stack SDN connector:
   a. Go to *Policy & Objects* > *Addresses*.
   b. Click *Create New*, then select *Address*.
   c. Configure the address as shown, selecting the desired filter in the *Filter* dropdown list. In this example, the Azure Stack SDN connector will automatically populate and update IP addresses only for instances that are named tfgta:



3. Ensure that the Azure Stack SDN connector resolves dynamic firewall IP addresses:
   a. Go to *Policy & Objects* > *Addresses*.
   b. Hover over the address created in step 2 to see a list of IP addresses for instances that are named tftgta as configured in step 2:



To configure Azure Stack SDN connector using CLI commands:

1. Configure the Azure Stack SDN connector:

```
config system sdn-connector
   edit "azurestack1"
      set type azure
      set azure-region local
      set server "azurestack.external"
      set username "username@azurestoreexamplecompany.onmicrosoft.com"
      set password xxxxx
      set log-in endpoint "https://login.microsoftonline.com/942b80cd-1b14-42a1-8dcf-
           4b21dece61ba"
      set resource-url
           "https://management.azurestoreexamplecompany.onmicrosoft.com/12b6fedd-9364-
           4cf0-822b-080d70298323"
      set update-interval 30
   next
end
```

2. Create a dynamic firewall address for the configured Azure Stack SDN connector with the supported Azure Stack filter. In this example, the Azure Stack SDN Connector will automatically populate and update IP addresses only for instances that are named tfgta:

```
config firewall address
   edit "azurestack-address-name1"
      set type dynamic
      set sdn "azurestack1"
      set filter "vm=tfgta"
   next
end
```

3. Confirm that the Azure Stack fabric connector resolves dynamic firewall IP addresses using the configured filter:

```
config firewall address
   edit "azurestack-address-name1"
      set type dynamic
      set sdn "azurestack1"
      set filter "vm=tfgta"
      config list
         edit "10.0.1.4"
         next
         edit "10.0.2.4"
         next
         edit "10.0.3.4"
         next
         edit "10.0.4.4"
         next
         edit "192.168.102.32"
         next
         edit "192.168.102.35"
         next
      end
   next
end
```

Azure SDN connector for non-VM resources

IP address resolving functionality is available for the following Azure resources:

Netskope Exhibit 1017

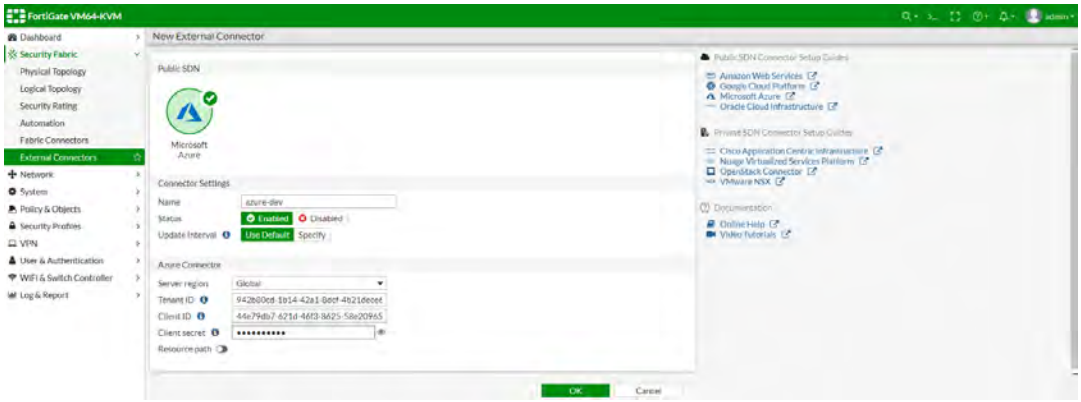
- VM network interfaces (including VMSS)
- Internet-facing load balancers
- Internal load balancers
- Application gateways

VPN gateways are currently not supported.

The following example demonstrates configuring an internet-facing load balancer.

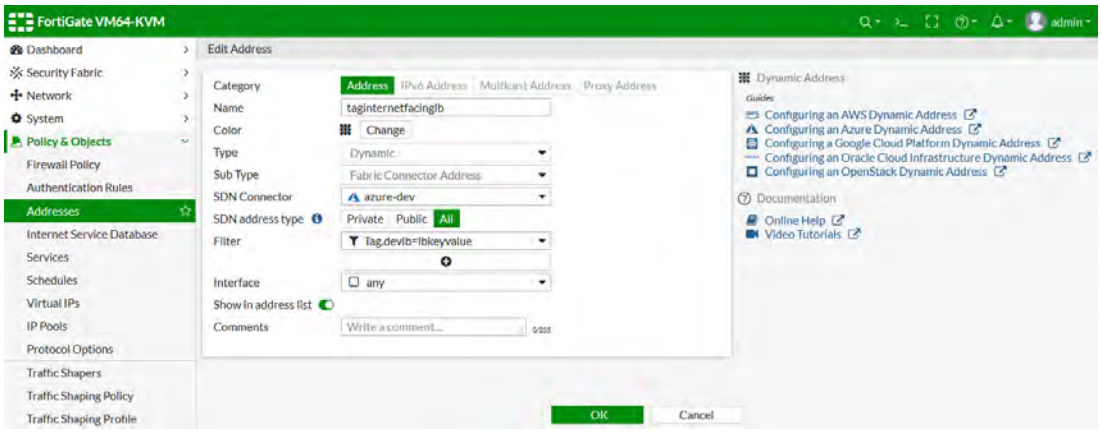
To configure an internet-facing load balancer address in the GUI:

1. Configure the Azure SDN connector:
   a. Go to *Security Fabric* > *External Connectors*.
   b. Click *Create New*, and select *Microsoft Azure*.
   c. Enter the settings based on your deployment, and click *OK*. The update interval is in seconds.



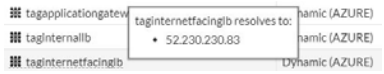
2. Create the dynamic firewall address:
   a. Go to *Policy & Objects* > *Addresses*.
   b. Click *Create New* > *Address* and enter a name.
   c. Configure the following settings:
      i. For *Type*, select *Dynamic*.
      ii. For *Sub Type*, select *Fabric Connector Address*.
      iii. For *SDN Connector*, select *azure-dev*.
      iv. For *SDN address type*, select *All*.
      v. For *Filter*, enter `Tag.devlb=lbkeyvalue`.

d. Click *OK*.



The corresponding IP addresses are dynamically updated and resolved after applying the tag filter.

3. Ensure that the connector resolves the dynamic firewall IP address:
   a. Go to *Policy & Objects* > *Addresses*.
   b. In the address table, hover over the address created in step 2 to view what IP it resolves to:



   c. In Azure, verify to confirm the IP address matches:



To configure an internet-facing load balancer in the CLI:

1. Configure the Azure SDN connector:

```
config system sdn-connector
    edit "azure-dev"
        set status enable
        set type azure
        set azure-region global
        set tenant-id "942b80cd-1b14-42a1-8dcf-4b21dece61ba"
        set client-id "44e79db7-621d-46f3-8625-58e209654e58"
        set client-secret xxxxxxxxxx
        set update-interval 60
    next
end
```

2. Create the dynamic firewall address:

```
config firewall address
    edit "taginternetfacinglb"
        set type dynamic
        set sdn "azure-dev"
        set filter "Tag.devlb=lbkeyvalue"
        set sdn-addr-type all
    next
end
```

The corresponding IP addresses are dynamically updated and resolved after applying the tag filter.

3. Confirm that the connector resolves the dynamic firewall IP address:

```
config firewall address
    edit "taginternetfacinglb"
        show
            config firewall address
                edit "taginternetfacinglb"
                    set uuid df391760-3bb6-51ea-f775-421df18f368d
                    set type dynamic
                    set sdn "azure-dev"
                    set filter "Tag.devlb=lbkeyvalue"
                    set sdn-addr-type all
                    config list
                        edit "52.230.230.83"
                        next
                    end
                next
            end
        next
end
```

VMware ESXi SDN connector

Dynamic addresses for VMware ESXi and vCenter servers can be automatically updated by using a VMware ESXi SDN connector, including mapping the following attributes from VMware ESXi and vCenter objects to dynamic address groups in FortiOS:

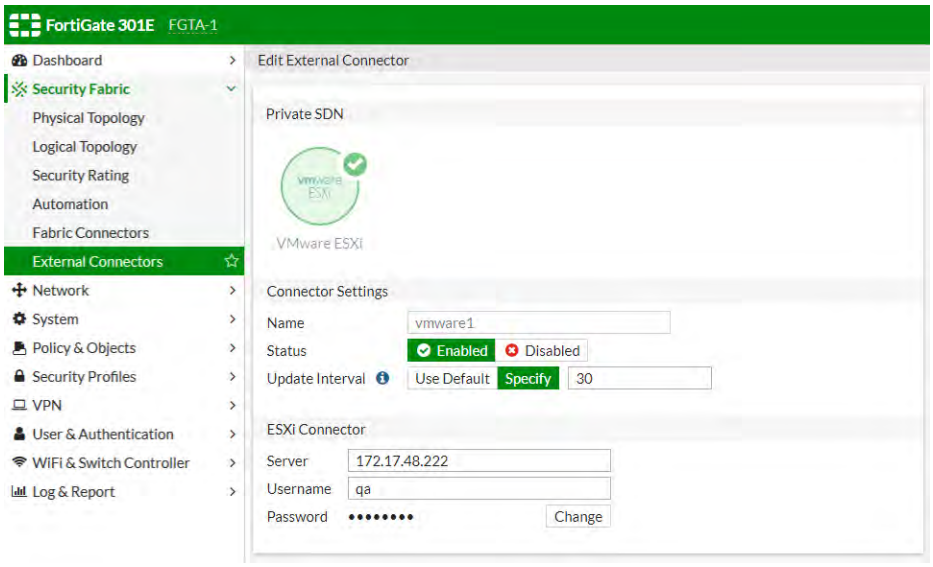
- vmid
- host
- name
- uuid
- vmuuid
- vmnetwork
- guestid
- guestname
- annotation

To configure VMware ESXi SDN connector using the GUI:

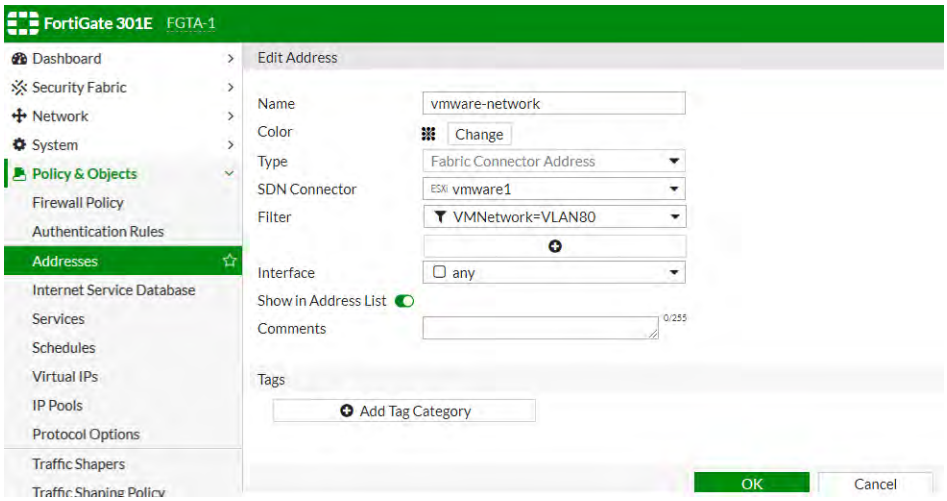
1. Configure the VMware ESXi SDN connector:
   a. Go to *Security Fabric* > *External Connectors*.
   b. Click *Create New*, and select *VMware ESXi*.
   c. Configure as shown, substituting the server IP address, username, and password for your deployment. The

Netskope Exhibit 1017

update interval is in seconds.



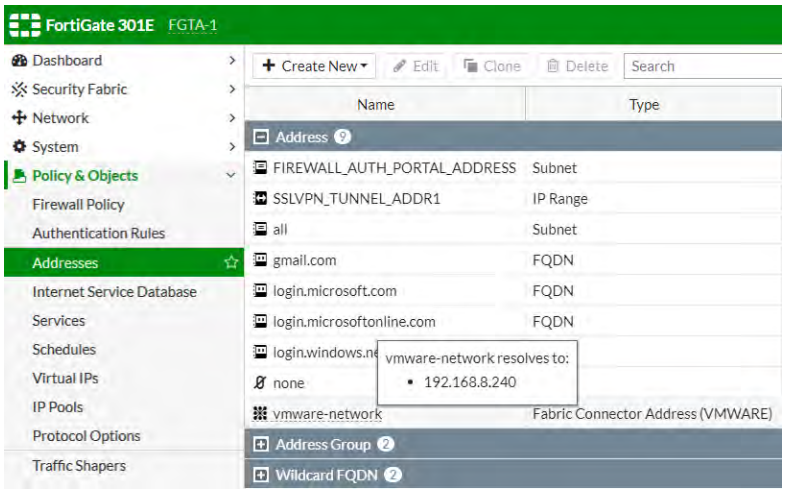
2. Create a dynamic firewall address for the configured VMware ESXi SDN connector:
   a. Go to *Policy & Objects* > *Addresses*.
   b. Click *Create New*, then select *Address*.
   c. Configure the address as shown, selecting the desired filter in the *Filter* dropdown list. In this example, the VMware ESXi fabric connector will automatically populate and update IP addresses only for instances that belong to VLAN80:



3. Ensure that the VMware ESXi SDN connector resolves dynamic firewall IP addresses:
   a. Go to *Policy & Objects* > *Addresses*.
   b. Hover over the address created in step 2 to see a list of IP addresses for instances that belong to VLAN80 as

Netskope Exhibit 1017

configured in step 2:



To configure VMware ESXi SDN connector using CLI commands:

1. Configure the VMware ESXi SDN connector:

```
config system sdn-connector
   edit "vmware1"
      set type vmware
      set server "172.17.48.222"
      set username "example_username"
      set password xxxxx
      set update-interval 30
   next
end
```

2. Create a dynamic firewall address for the configured VMware ESXi SDN connector with the supported VMware ESXi filter. In this example, the VMware ESXi SDN connector will automatically populate and update IP addresses only for instances that belong to the specified VLAN:

```
config firewall address
   edit "vmware-network"
      set type dynamic
      set sdn "vmware1"
      set filter "vmnetwork=VLAN80"
   next
end
```

3. Confirm that the VMware ESXi SDN connector resolves dynamic firewall IP addresses using the configured filter:

```
config firewall address
   edit "vmware-network"
      set type dynamic
      set sdn "vmware1"
      set filter "vmnetwork=VLAN80"
      config list
         edit "192.168.8.240"
         next
      end
   next
end
```

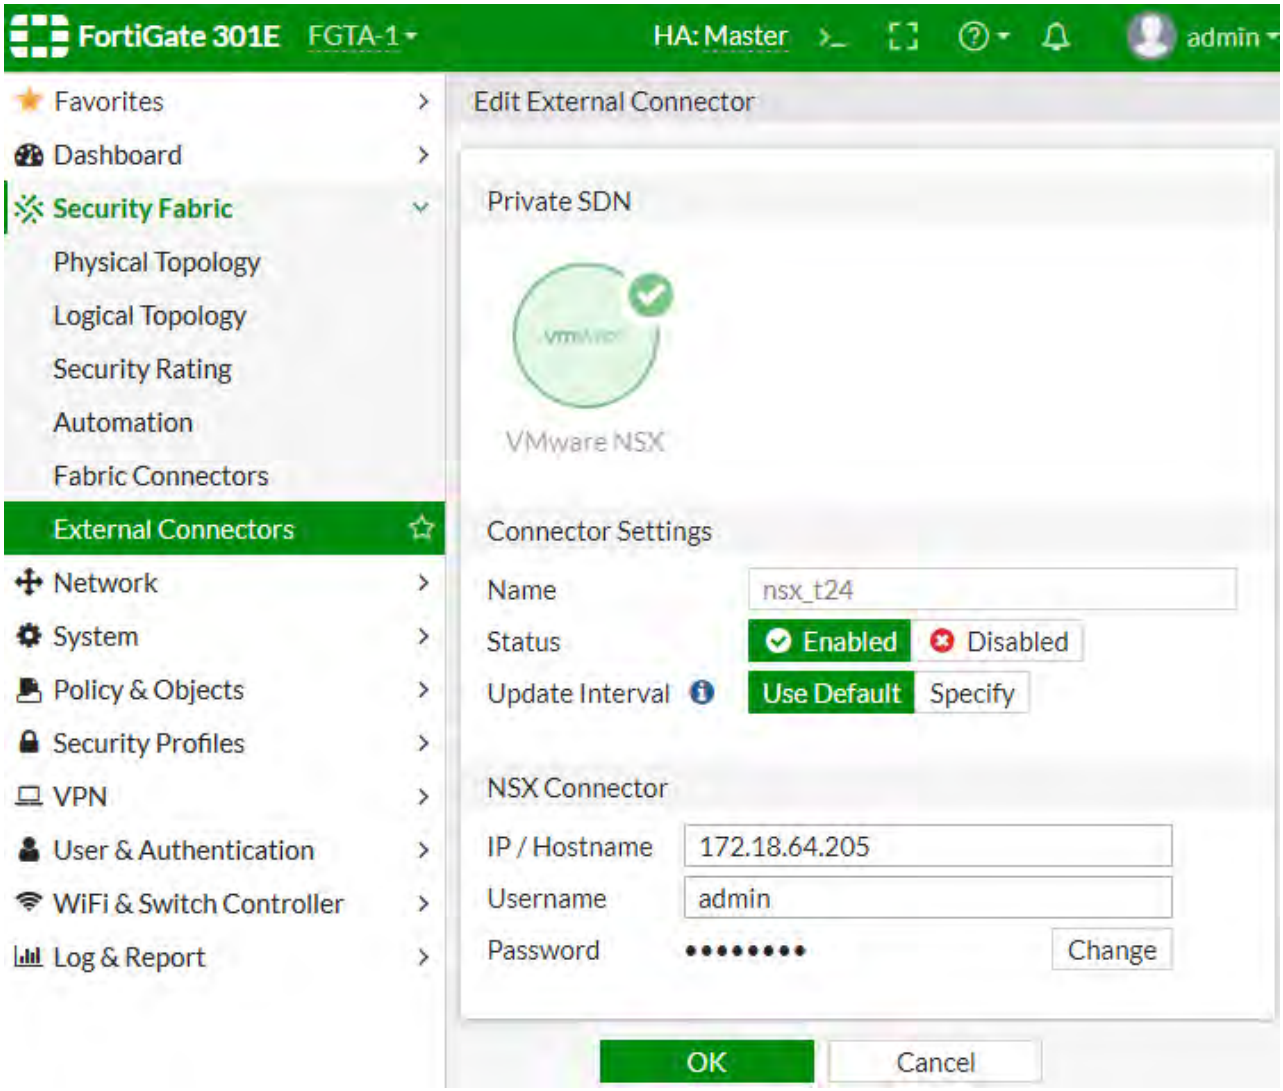

Netskope Exhibit 1017

VMware NSX-T manager SDN connector

This feature provides SDN connector configuration for VMware NSX-T manager. You can import specific groups, or all groups from the NSX-T manager.

To configure SDN connector for NSX-T manager in the GUI:

1. Go to *Security Fabric* > *External Connectors* and click *Create New*.
2. In the *Private SDN* section, click *VMware NSX*.
3. Enter the settings and click *OK*.



To configure SDN connector for NSX-T manager in the CLI:

```
config system sdn-connector
    edit "nsx_t24"
        set type nsx
        set server "172.18.64.205"
        set username "admin"
```

```
        set password xxxxxx
    next
end
```

To import a specific group from the NSX-T manager:

```
# execute nsx group import nsx_t24 root csf_ns_group
[1] 336914ba-0660-4840-b0f1-9320f5c5ca5e csf_ns_group:
  Name:csf_ns_group
  Address:1.1.1.0
  Address:1.1.1.1
  Address:172.16.10.104
  Address:172.16.20.104
  Address:172.16.30.104
  Address:2.2.2.0
  Address:2.2.2.2
  Address:4.4.4.0
  Address:5.5.5.0
  Address:6.6.6.6
  Address:7.7.7.7
```

To import all groups from NSX-T manager:

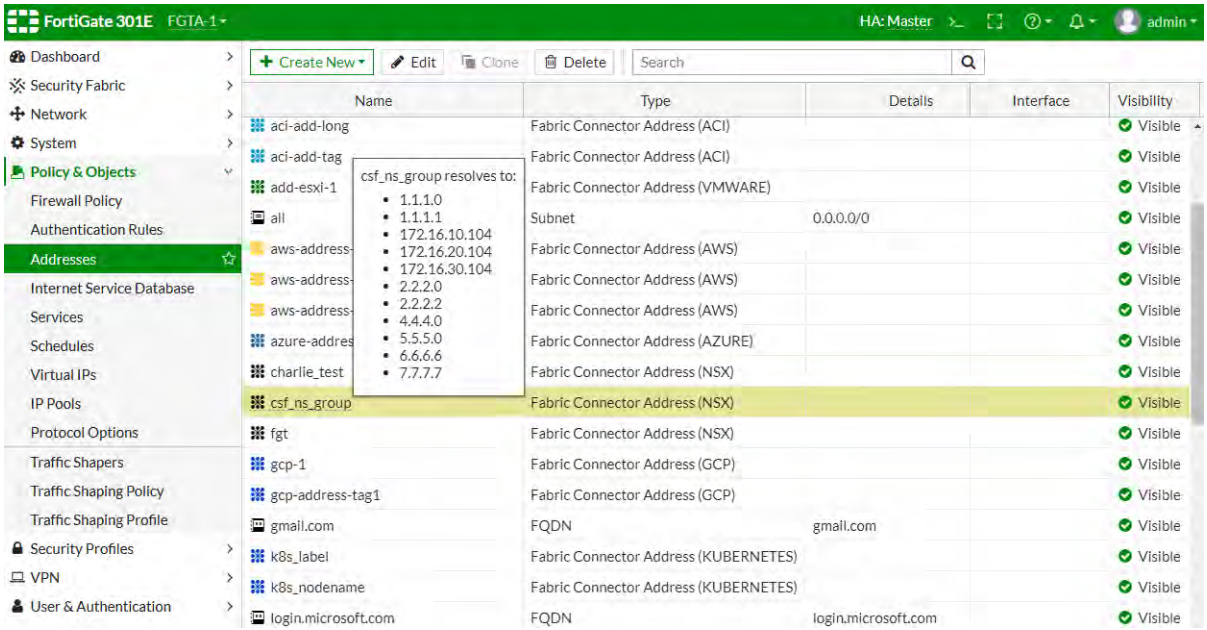
```
# execute nsx group import nsx_t24 root
[1] 663a7686-b9a3-4659-b06f-b45c908349a0 ServiceInsertion_NSGroup:
  Name:ServiceInsertion_NSGroup
  Address:10.0.0.2
[2] 336914ba-0660-4840-b0f1-9320f5c5ca5e csf_ns_group:
  Name:csf_ns_group
  Address:1.1.1.0
  Address:1.1.1.1
  Address:172.16.10.104
  Address:172.16.20.104
  Address:172.16.30.104
  Address:2.2.2.0
  Address:2.2.2.2
  Address:4.4.4.0
  Address:5.5.5.0
  Address:6.6.6.6
  Address:7.7.7.7
[3] c462ec4d-d526-4ceb-aeb5-3f168cecd89d charlie_test:
  Name:charlie_test
  Address:1.1.1.1
  Address:2.2.2.2
  Address:6.6.6.6
  Address:7.7.7.7
[4] ff4dcb08-53cf-46bd-bef4-f7aeda9c0ad9 fgt:
  Name:fgt
  Address:172.16.10.101
  Address:172.16.10.102
  Address:172.16.20.102
  Address:172.16.30.103
[5] 3dd7df0d-2baa-44e0-b88f-bd21a92eb2e5 yongyu_test:
  Name:yongyu_test
  Address:1.1.1.0
  Address:2.2.2.0
```


Netskope Exhibit 1017

```
Address:4.4.4.0
Address:5.5.5.0
```

To view the dynamic firewall IP addresses that are resolved by the SDN connector in the GUI:

1. Go to *Policy & Objects* > *Addresses* to view the IP addresses resolved by an SDN connector.



To view the dynamic firewall IP addresses that are resolved by the SDN connector in the CLI:

```
# show firewall address csf_ns_group
config firewall address
    edit "csf_ns_group"
        set uuid ee4a2696-bacd-51e9-f828-59457565b880
        set type dynamic
        set sdn "nsx_t24"
        set obj-id "336914ba-0660-4840-b0f1-9320f5c5ca5e"
        config list
            edit "1.1.1.0"
            next
            edit "1.1.1.1"
            next
            edit "172.16.10.104"
            next
            edit "172.16.20.104"
            next
            edit "172.16.30.104"
            next
            edit "2.2.2.0"
            next
            edit "2.2.2.2"
            next
            edit "4.4.4.0"
            next
            edit "5.5.5.0"
            next
```

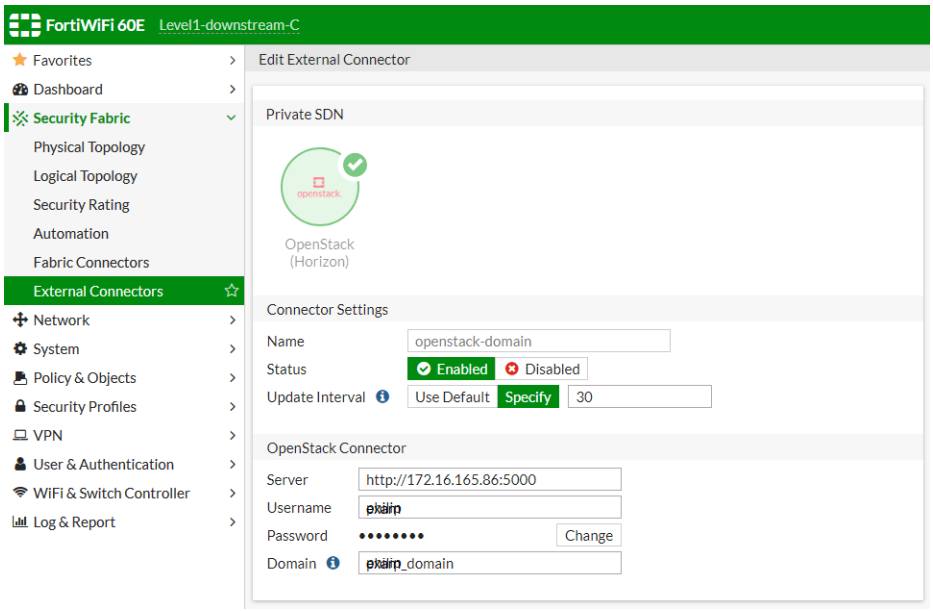
```
            edit "6.6.6.6"
            next
            edit "7.7.7.7"
            next
        end
    next
end
```

OpenStack (Horizon) SDN connector with domain filter

You can select a domain attribute when configuring an OpenStack SDN connector in FortiOS. When a domain is configured for the OpenStack SDN connector, FortiOS resolves OpenStack dynamic firewall addresses from the specified OpenStack domain. If a domain is not specified, FortiOS resolves the dynamic firewall addresses using the default OpenStack domain.

To configure OpenStack SDN connector with a domain filter using the GUI:

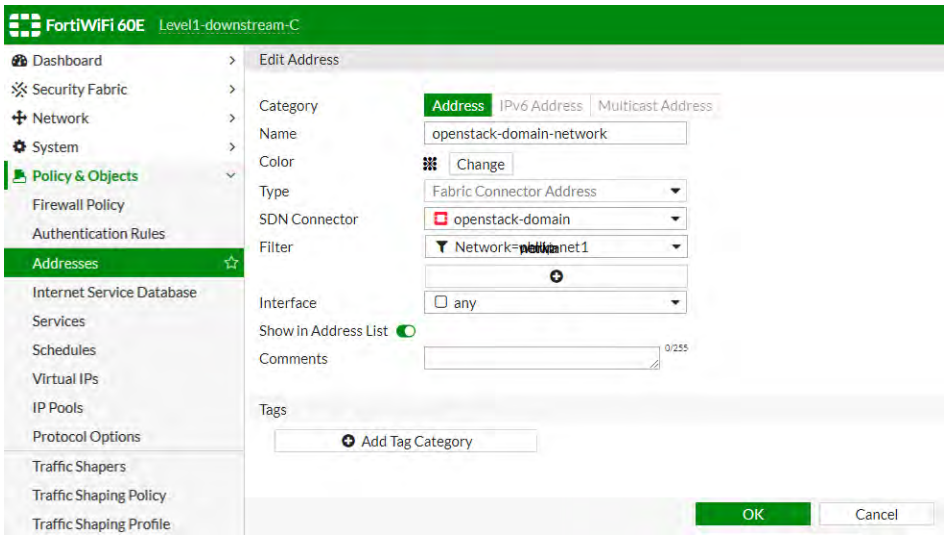
1. Configure the OpenStack SDN connector:
   a. Go to *Security Fabric* > *External Connectors*.
   b. Click *Create New*, and select *Openstack (Horizon)*.
   c. In the *Domain* field, enter the desired domain name from OpenStack. The fabric connector will only resolve IP addresses for instances that belong to the specified domain.
   d. Configure as shown, substituting the server IP address, username, and password for your deployment. The update interval is in seconds.



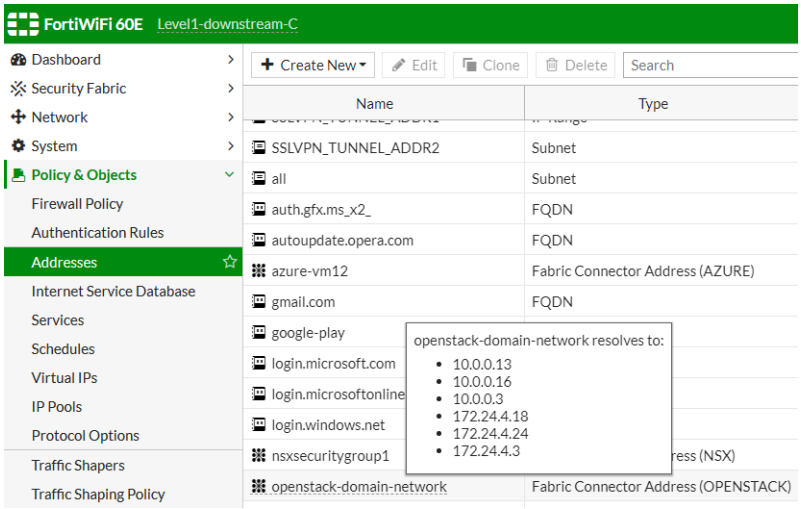
2. Create a dynamic firewall address for the configured OpenStack SDN connector:
   a. Go to *Policy & Objects* > *Addresses*.
   b. Click *Create New*, then select *Address*.
   c. Configure the address as shown, selecting the desired filter in the *Filter* dropdown list. The OpenStack SDN connector will automatically populate and update IP addresses only for instances that belong to the

Netskope Exhibit 1017

specified domain and network:



3. Ensure that the OpenStack SDN connector resolves dynamic firewall IP addresses:
   a. Go to *Policy & Objects* > *Addresses*.
   b. Hover over the address created in step 2 to see a list of IP addresses for instances that belong to the specified domain and specified network as configured in steps 1 and 2:



To configure OpenStack SDN connector with a domain filter using CLI commands:

1. Configure the OpenStack SDN connector. The SDN connector will only resolve IP addresses for instances that belong to the specified domain:

```
config system sdn-connector
   edit "openstack-domain"
      set type openstack
      set server "http://172.16.165.86:5000"
      set username "example_username"
      set password xxxxx
      set domain "example_domain"
      set update-interval 30
   next
end
```


Netskope Exhibit 1017

2. Create a dynamic firewall address for the configured OpenStack SDN connector with the supported OpenStack filter. The OpenStack SDN connector will automatically populate and update IP addresses only for instances that belong to the specified domain and the specified network:

```
config firewall address
   edit "openstack-domain-network"
      set type dynamic
      set sdn "openstack-domain"
      set filter "Network=example-net1"
   next
end
```

3. Confirm that the OpenStack SDN connector resolves dynamic firewall IP addresses using the configured domain and filter:

```
config firewall address
   edit "openstack-domain-network"
      set type dynamic
      set sdn "openstack-domain"
      set filter "Network=example-net1"
      config list
         edit "10.0.0.13"
         next
         edit "10.0.0.16"
         next
         edit "10.0.0.3"
         next
         edit "172.24.4.18"
         next
         edit "172.24.4.24"
         next
         edit "172.24.4.3"
         next
      end
   next
end
```

OCI SDN connector

You can configure Security Fabric connector integration with Oracle Cloud Infrastructure (OCI).

To configure an OCI SDN connector in the CLI:

1. Configure an SDN connector:

```
config system sdn-connector
    edit "oci1"
        set status enable
        set type oci
        set tenant-id
"ocid1.tenancy.oc1..aaaaaaaaaaa3aaaaaaaaaaaaaaaaa77xxxxxx54bbbbbb4xxxx35xx55xxxx"
        set user-id
"ocid1.user.oc1..aaaaaaaaa2laaaaa3aaaaaaaaaabbbbbbbbbbcccc3ccccccccccxxxxxxxx"
        set compartment-id
"ocid1.compartment.oc1..aaaaaaaaaaaaaaaaa7bbbbbbbbbbccccccccc6xxx53xxxx7xxxxxxxxx"
        set oci-region "us-ashburn-1"
        set oci-region-type commercial
```


Netskope Exhibit 1017

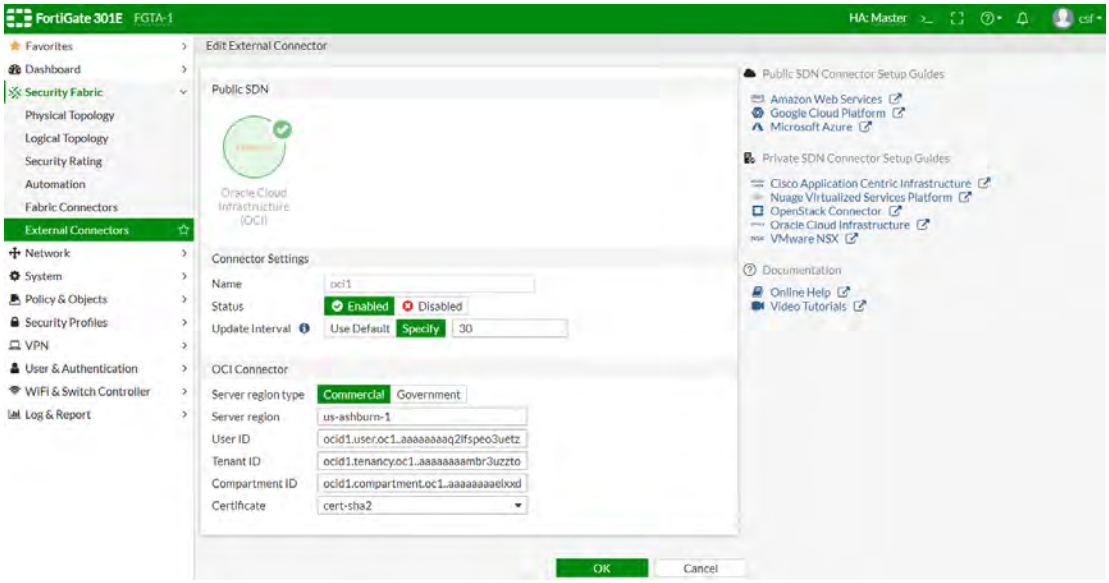
```
        set oci-cert "cert-sha2"
        set update-interval 30
    next
end
```

2. Create a dynamic firewall address for the SDN connector with a supported filter:

```
config firewall address
    edit "oci-address-1"
        set type dynamic
        set sdn "oci1"
        set filter "CompartmentName=DevelopmentEngineering"
    next
end
```

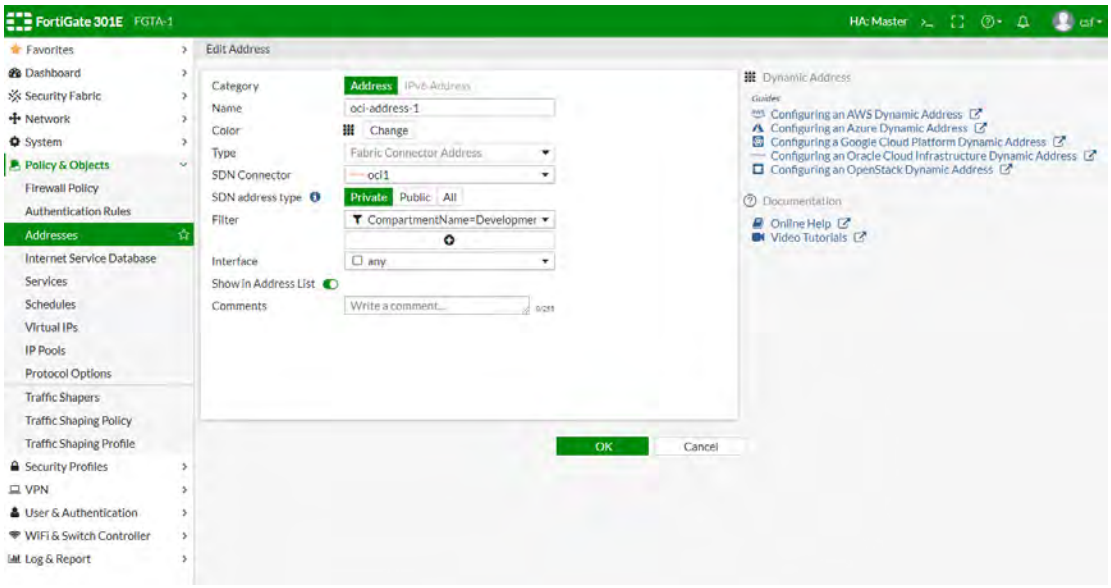
To configure an OCI SDN connector in the GUI:

1. Go to *Security Fabric* > *External Connectors* and click *Create New*.
2. In the *Public SDN* section, select *Oracle Cloud Infrastructure (OCI)*.
3. Configuration the connector as needed.



4. Click *OK*.
5. Go to *Policy & Objects* > *Addresses* and click *Create New* > *Address*.

6. Configure the address as needed, selecting the OCI connector in the *SDN Connector* field.



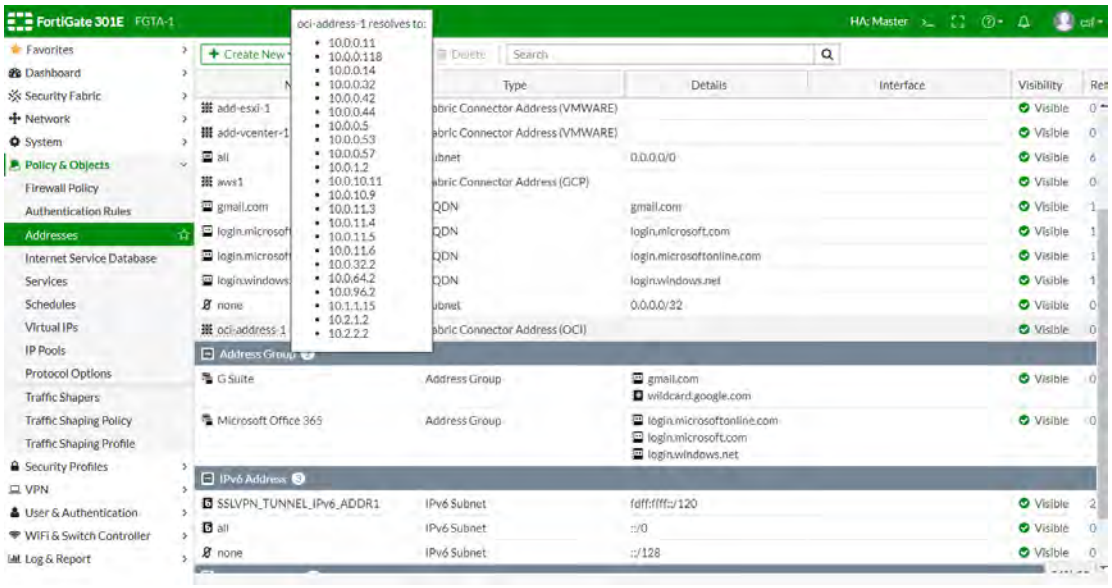
7. Click *OK*.

To confirm that dynamic firewall addresses are resolved by the SDN connector:

1. In the CLI, check that the addresses are listed:

```
config firewall address
    edit "oci-address-1"
        set type dynamic
        set sdn "oci1"
        set filter "CompartmentName=DevelopmentEngineering"
        config list
            edit "10.0.0.11"
            next
            edit "10.0.0.118"
            next
            ...
            next
        end
    next
end
```

2. In the GUI, go to *Policy & Objects* > *Addresses* and hover the cursor over the address name.



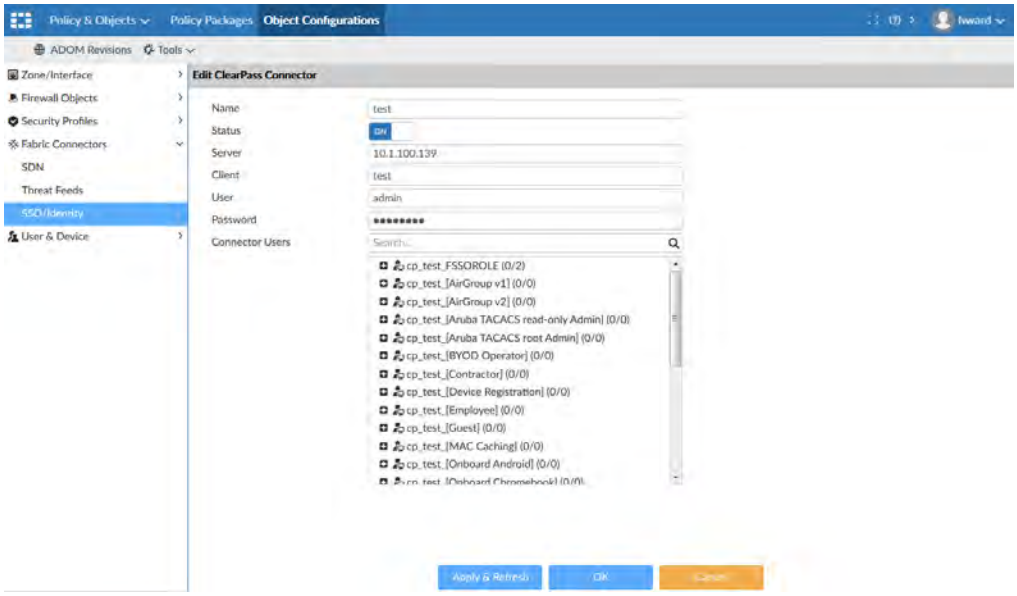
ClearPass endpoint connector via FortiManager

ClearPass Policy Manager (CPPM) is a network access system that can send information about authenticated users to third party systems, such as a FortiGate or FortiManager.

In this example, communications are established between CPPM and FortiManager, and then the FortiManager forwards information to a managed FortiGate. On the FortiGate, the user information can be used in firewall policies and added to FSSO dynamic addresses.

Configure the FortiManager

Establish communications between FortiManager and CPPM so that FortiManager can synchronize CPPM user groups. See Creating a ClearPass connector in the FortiManager Administration Guide.

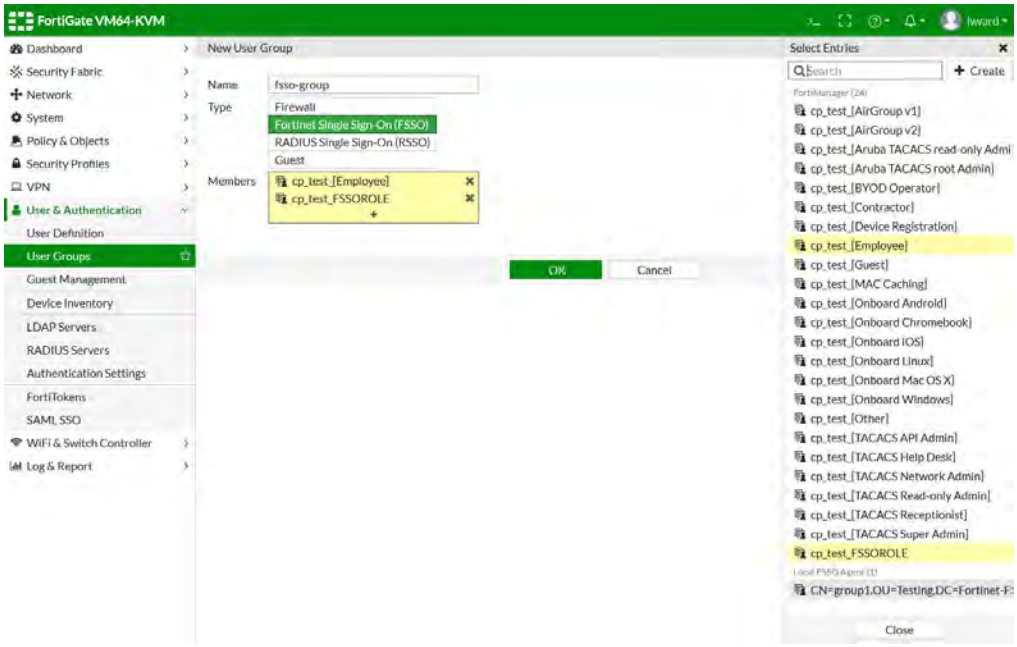


FortiManager forwards the group information to managed FortiGates.

Add CPPM FSSO user groups to a local user group

To add CPPM user groups to a local user group in the GUI:

1. On the FortiGate, go to *User & Authentication* > *User Groups*.
2. Click *Create New*.
3. Enter a name for the group and set *Type* to *Fortinet Single Sign-On (FSSO)*.
4. Click the *Members* field, and add one or more FSSO groups.
   FSSO groups can come from multiple sources; CPPM FSSO groups are prefixed with *cp_* and are listed under the *FortiManager* heading.



5. Click *OK*.

To add CPPM user groups to a local user group in the CLI:

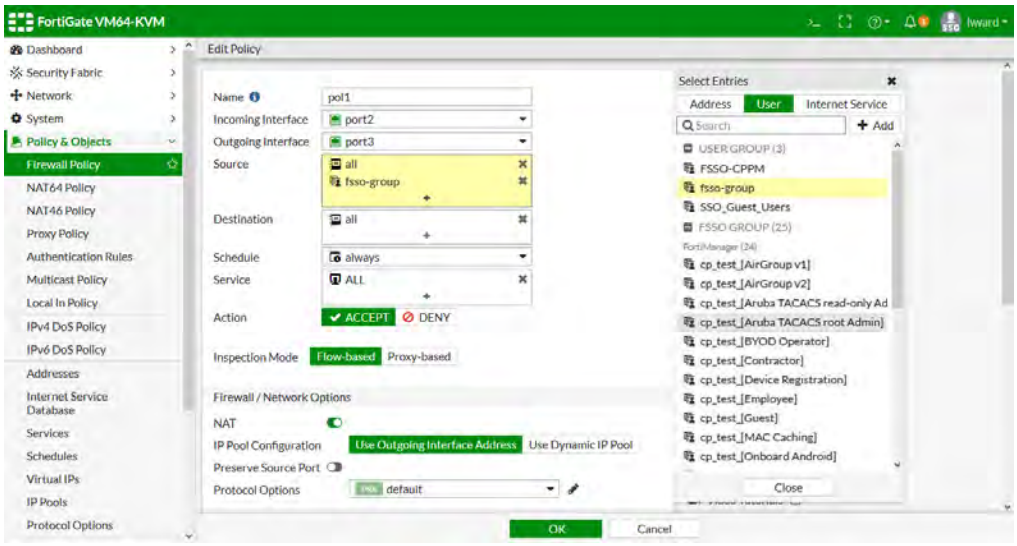
```
config user group
    edit fsso-group
        set group-type fsso-service
        set member "cp_test_[Employee]" "cp_test_FSSOROLE"
    next
end
```

Use the local FSSO user group in a firewall policy

To add the local FSSO user group to a firewall policy in the GUI:

1. Go to *Policy & Objects* > *Firewall Policy*.
2. Create a new policy, or edit an existing one.

3. Click in the *Source* field and add the *fsso-group* user group.



   CPPM user groups can also be added directly to the policy.

4. Click *OK*.

To add the local FSSO user group to a firewall policy in the CLI:

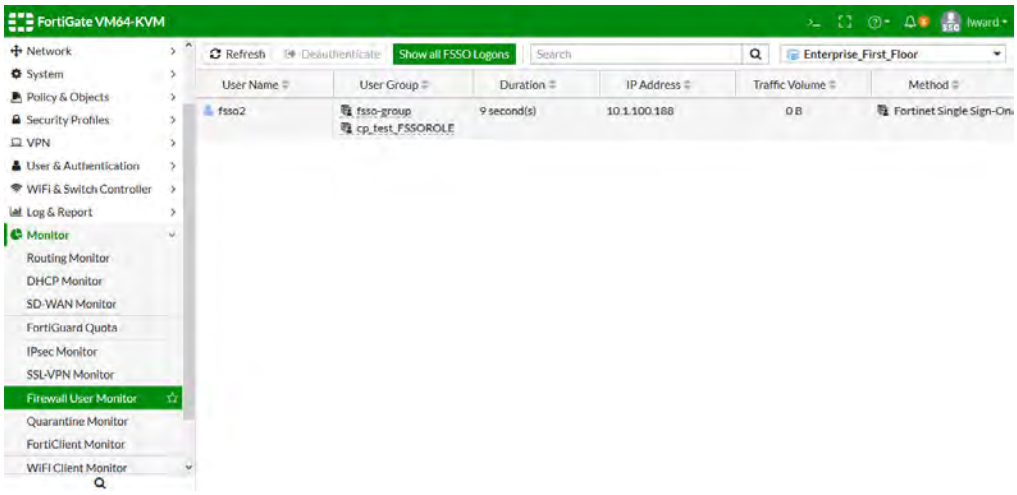
```
config firewall policy
    edit 1
        set name "pol1"
        set srcintf "port2"
        set dstintf "port3"
        set srcaddr "all"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ALL"
        set logtraffic all
        set groups "fsso-group"
        set nat enable
    next
end
```

Verification

To verify that a user was added to the FSSO list on the FortiGate:

1. Log on to the client and authenticate with CPPM.
   After successful authentication, the user is added to the FSSO list on the FortiGate.

2. On the FortiGate, go to *Monitor* > *Firewall User Monitor* to verify that the user was added.



The user group *cp_test_FSSOROLE* is listed separately because the user is a member of that group on the CPPM.

To verify that traffic can pass the firewall:

1. Log on to the client and browse to an external website.
2. On the FortiGate, go to *FortiView* > *Sources*.
3. Double-click on the user and select the *Destinations* tab to verify that traffic is being passed by the firewall.

To verify the user address groups:

```
show user adgrp
    config user adgrp
        edit "cp_test_FSSOROLE"
            set server-name "FortiManager"
        next
        edit "cp_test_[AirGroup v1]"
            set server-name "FortiManager"
        next
        edit "cp_test_[AirGroup v2]"
            set server-name "FortiManager"
        next
        edit "cp_test_[Aruba TACACS read-only Admin]"
            set server-name "FortiManager"
        next
        edit "cp_test_[Aruba TACACS root Admin]"
            set server-name "FortiManager"
        next
        edit "cp_test_[BYOD Operator]"
            set server-name "FortiManager"
        next
        edit "cp_test_[Contractor]"
            set server-name "FortiManager"
        next
        edit "cp_test_[Device Registration]"
            set server-name "FortiManager"
        next
        ...
```

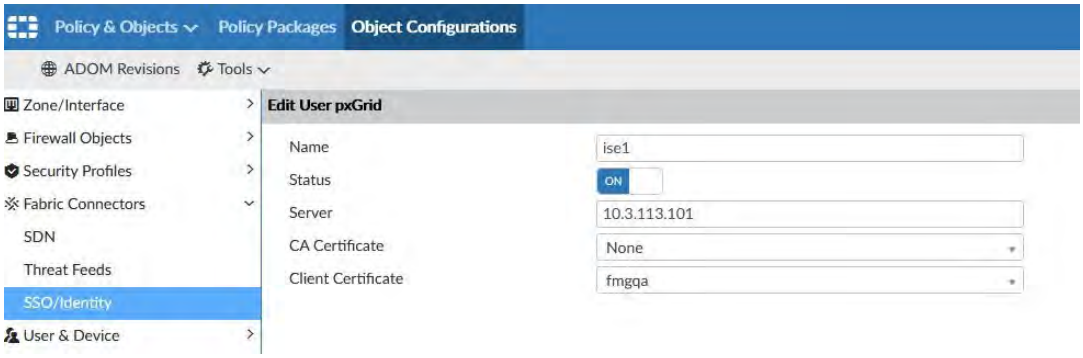
```
    edit "CN=group1,OU=Testing,DC=Fortinet-FSSO,DC=COM"
        set server-name "Local FSSO Agent"   <----- !!!
    next
end
```

Cisco pxGrid fabric connector

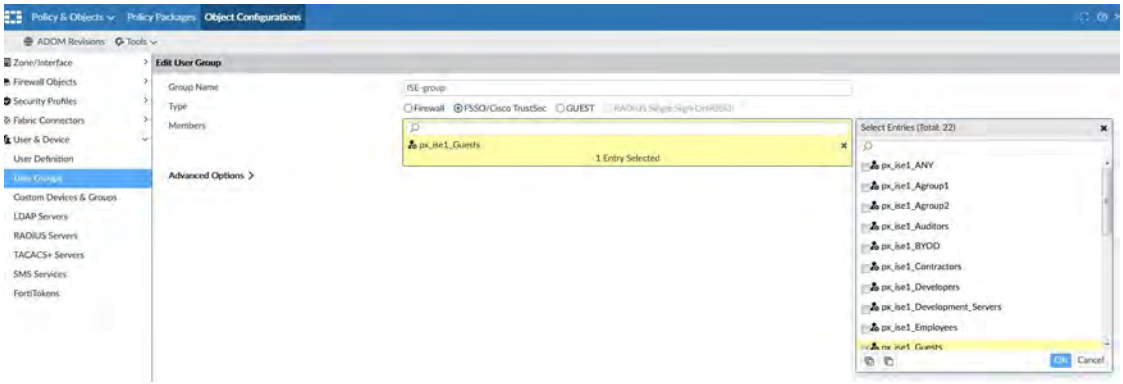
You can create an endpoint connector to Cisco pxGrid by using FortiManager. FortiManager dynamically collects updates from pxGrid and forwards them to FortiGate by using the Fortinet Single Sign On (FSSO) protocol.

To create a Cisco pxGrid fabric connector:

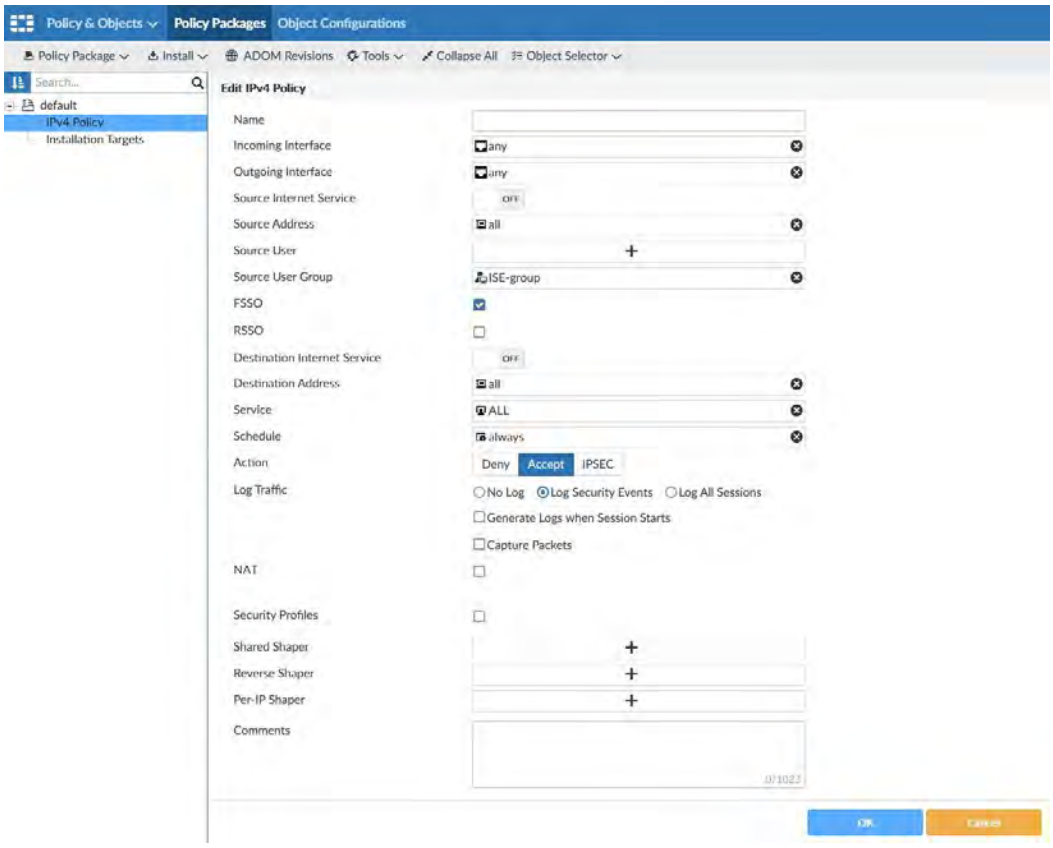
1. On FortiManager, create an SSO Connector to Cisco ISE.
   Communication between FortiManager and Cisco ISE is secured by using TLS. FortiManager requires a client certificate issued by Cisco ISE. FortiManager uses the certificate to authenticate to Cisco ISE.



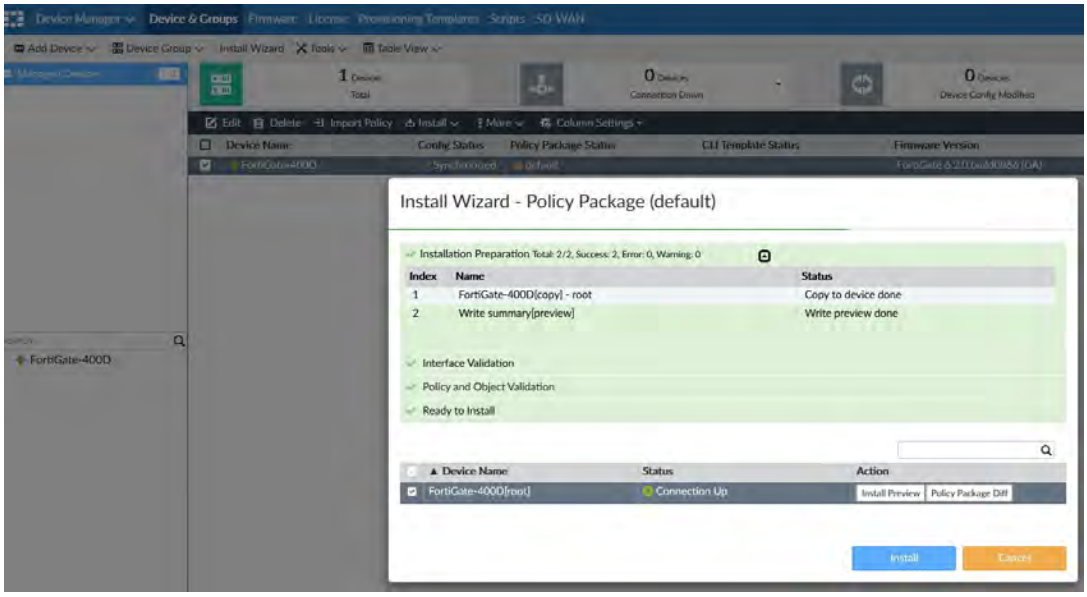
2. On FortiManager, map Cisco ISE groups to a Fortinet FSSO group.
   Once a secured communication channel is established, Cisco sends all user groups to FortiManager.
   The FortiManager administrator can select specific groups and map them to Fortinet FSSO groups.



3. On FortiManager, add Fortinet FSSO group to a firewall policy in a policy package.



4. On FortiManager, synchronize the policy package to the firewall for the managed FortiGate.

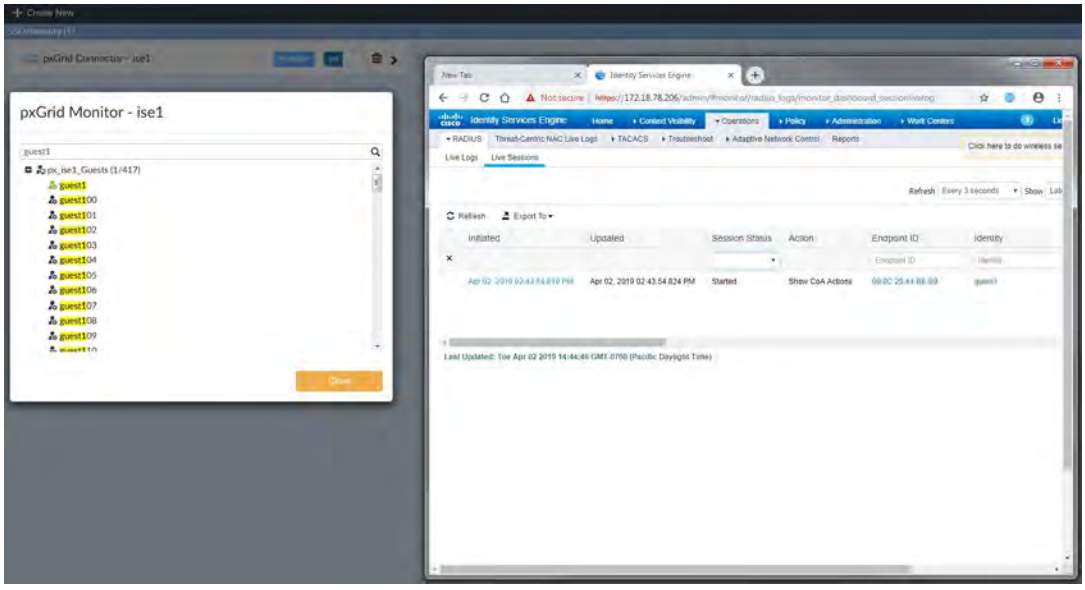


5. On FortiGate, verify that the synced firewall policy contains the correct FSSO group and that all FSSO-related information in `user adgrp` is correct.

```
config firewall policy
    edit 1
        set uuid b803052e-562a-51e9-0561-82525c8bcaa9
        set srcintf "any"
        set dstintf "any"
        set srcaddr "all"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ALL"
        set groups "ISE-group"
    next
end

FortiGate-400D # show user adgrp
config user adgrp
    edit "px_ise1_ANY"
        set server-name "FortiManager"
    next
    edit "px_ise1_Agroup1"
        set server-name "FortiManager"
    next
    edit "px_ise1_Agroup2"
        set server-name "FortiManager"
    next
    edit "px_ise1_Auditors"
        set server-name "FortiManager"
    next
    edit "px_ise1_BYOD"
        set server-name "FortiManager"
    next
    edit "px_ise1_Contractors"
        set server-name "FortiManager"
    next
    edit "px_ise1_Developers"
        set server-name "FortiManager"
    next
    edit "px_ise1_Development_Servers"
        set server-name "FortiManager"
    next
    edit "px_ise1_Employees"
        set server-name "FortiManager"
    next
    edit "px_ise1_Guests"
        set server-name "FortiManager"
    next
    edit "px_ise1_HR"
        set server-name "FortiManager"
    next
    edit "px_ise1_Network_Services"
        set server-name "FortiManager"
```

6. After successful user authentication on Cisco ISE, verify that information is forwarded to FortiManager. On FortiManager, the icon next to the authenticated user in *pxGrid Monitor* should be green.



FortiGate should have two entries: one in the firewall-authenticated user list and one in the FSSO logged-on user list.

In the FSSO logged-on user list, you can view both groups. You view the group that the user belongs to on Cisco ISE and the Fortinet FSSO group.

```
FortiGate-400D #
FortiGate-400D # dia deb authd fsso l
----FSSO logons----
IP: 10.1.100.188  User: guest1  Groups: px_isel_Guests  Workstation:  MemberOf: ISE-group
Total number of logons listed: 1, filtered: 0
----end of FSSO logons----

FortiGate-400D # dia firewall auth l

10.1.100.188, guest1
       type: fsso, id: 0, duration: 59694, idled: 59694
       server: FortiManager
       packets: in 0 out 0, bytes: in 0 out 0
       group_id: 2
       group_name: ISE-group

----- 1 listed, 0 filtered ------
```

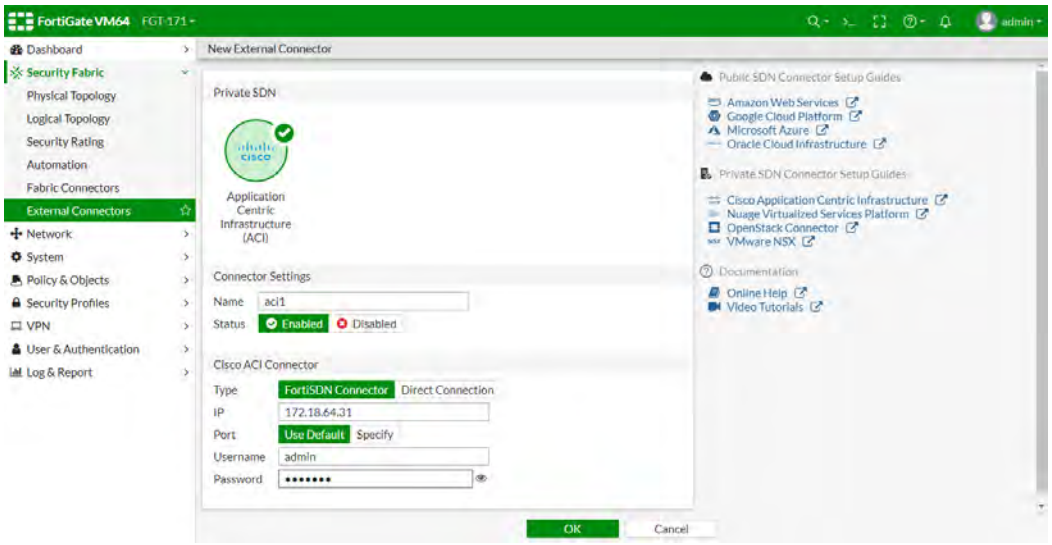
Cisco ACI SDN connector

Cisco ACI (Application Centric Infrastructure) SDN connectors can be used in dynamic firewall addresses.

The Fortinet SDN Connector for Cisco ACI and Nuage Networks is a standalone connector that connects to SDN controllers within Cisco ACI and Nuage Networks. You must configure a connection to the Fortinet SDN connector in FortiOS to query the dynamic addresses.

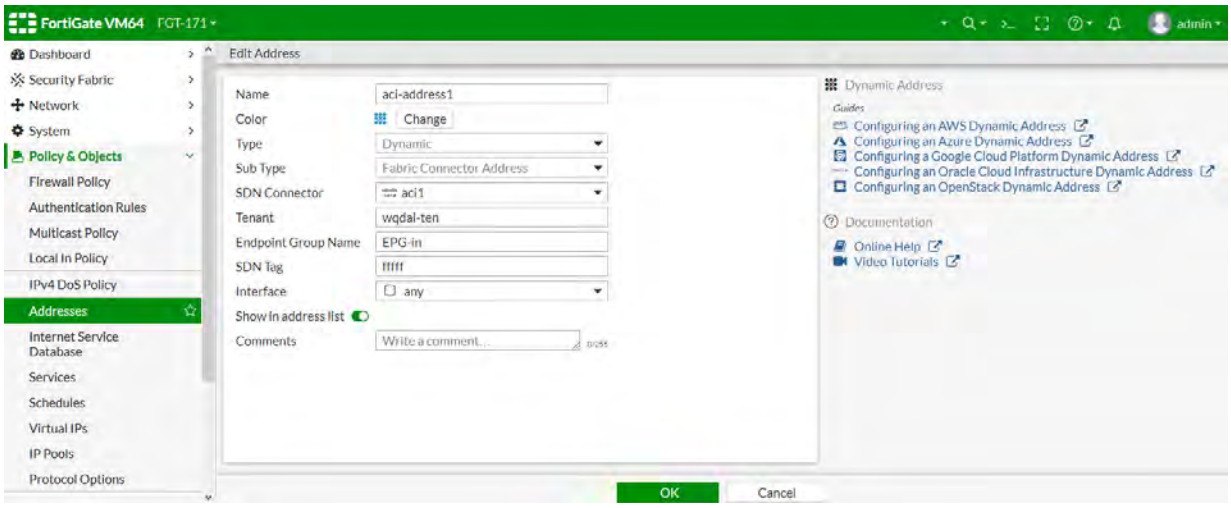
To configure a Cisco ACI connector in the GUI:

1. Create the Cisco ACI SDN connector:
   a. Go to *Security Fabric* > *External Connectors* and click *Create New*.
   b. In the *Private SDN* section, click *Application Centric Infrastructure (ACI)*.
   c. In the *Cisco ACI Connector* section, for *Type*, select *Fortinet SDN Connector* and configure the remaining settings as needed.
   d. Click *OK*.



2. Create the dynamic firewall address for the connector:
   a. Go to *Policy & Objects* > *Addresses* and click *Create New* > *Address*.
   b. Configure the following settings:
      i. For *Type*, select *Dynamic*.
      ii. For *Sub Type*, select *Fabric Connector Address*.
      iii. For *SDN Connector*, select the first ACI connector.
      iv. Configure the remaining settings as needed.

c. Click *OK*.



To verify the dynamic firewall IPs are resolved by the SDN connector in the GUI:

1. Go to *Policy & Objects* > *Addresses*.
2. In the address table, hover over the address to view which IPs it resolves to.

To configure a Cisco ACI connector in the CLI:

1. Create the SDN connector:

```
config system sdn-connector
    edit "aci1"
        set type aci
        set server "172.18.64.31"
        set username "admin"
        set password xxxxxxx
    next
end
```

2. Create the dynamic firewall address for the connector:

```
config firewall address
    edit "aci-address1"
        set type dynamic
        set sdn "aci1"
        set color 17
        set tenant "wqdai-ten"
        set epg-name "EPG-in"
        set sdn-tag "fffff"
    next
end
```

To verify the dynamic firewall IPs are resolved by the SDN connector in the CLI:

```
# diagnose firewall dynamic list

List all dynamic addresses:
```

```
aci1.aci.wqdai-ten.EPG-in.fffff: ID(171)
        ADDR(192.168.100.20)
```

Cisco ACI SDN connector with direct connection

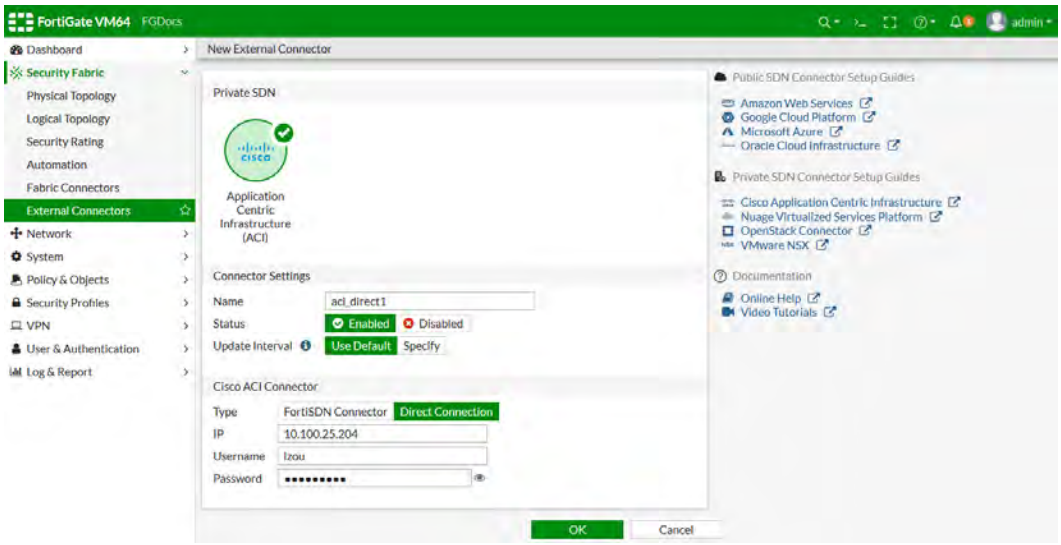
The Cisco ACI (Application Centric Infrastructure) connector can be used for northbound API integration with a direct connection.

The following address filters are supported:

- Tenant
- Application
- Endpoint group
- Tag

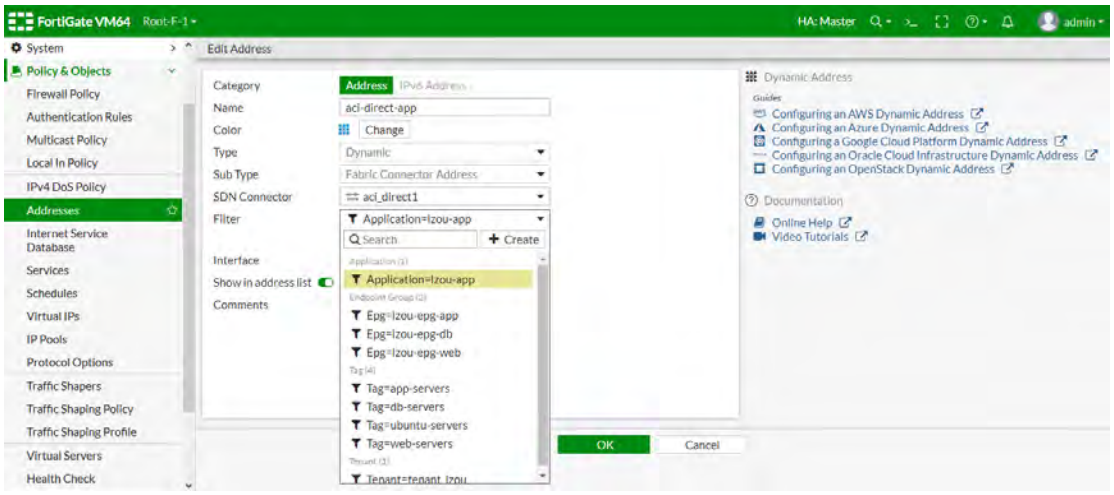
To configure a Cisco ACI connector in the GUI:

1. Create the Cisco ACI SDN connector:
   a. Go to *Security Fabric* > *External Connectors* and click *Create New*,.
   b. In the *Private SDN* section, click *Application Centric Infrastructure (ACI)*.
   c. Configure the *Connector Settings* as needed. The update interval is in seconds.
   d. In the *Cisco ACI Connector* section, for *Type*, select *Direct Connection* and configure the remaining settings as needed.
   e. Click *OK*.

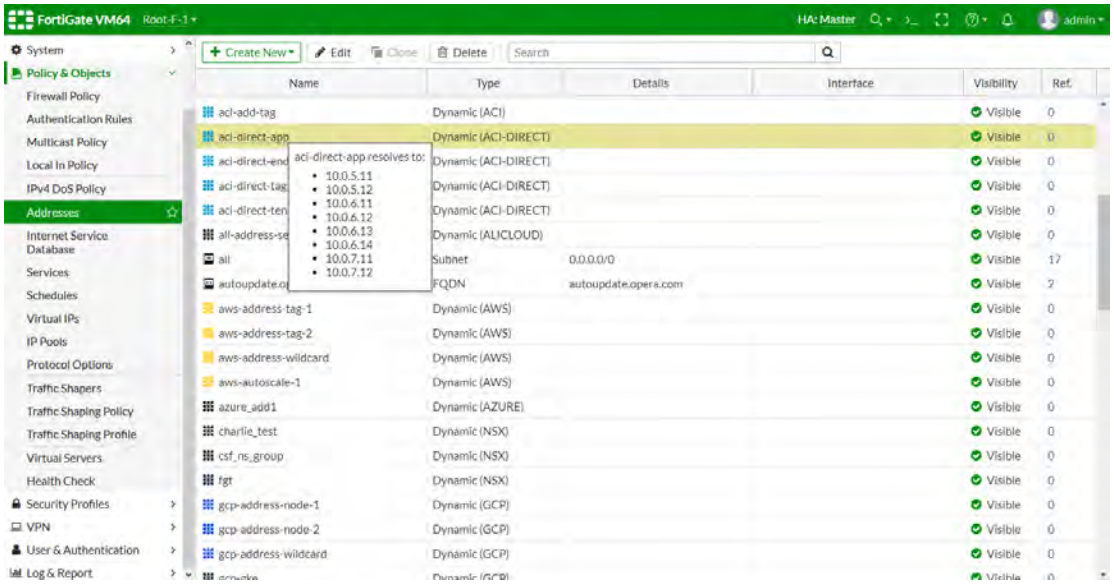


2. Create a dynamic firewall address for the connector:
   a. Go to *Policy & Objects* > *Addresses*.
   b. Click *Create New* > *Address* and enter a name.
   c. Configure the following settings:
      i. For *Type*, select *Dynamic*.
      ii. For *Sub Type*, select *Fabric Connector Address*.
      iii. For *SDN Connector*, select the connector created in step 1.
      iv. For *Filter*, select an entry from the dropdown list. In this example, *Application* is selected.

d. Click *OK*.



3. Confirm that the connector resolves the dynamic firewall IP addresses:
   a. Go to *Policy & Objects* > *Addresses*.
   b. In the address table, hover over the address created in step 2 to view which IPs it resolves to:



To configure a Cisco ACI connector in the CLI:

1. Create the Cisco ACI SDN connector:

```
config system sdn-connector
    edit "aci_direct1"
        set status enable
        set type aci-direct
        set server "10.100.25.204"
        set username "lzou"
        set password xxxxxxx
        set update-interval 60
    next
end
```

2. Create a dynamic firewall address for the connector:

```
config firewall address
    edit "aci-direct-app"
        set type dynamic
        set sdn "aci_direct1"
        set color 17
        set filter "Application=lzou-app"
    next
end
```

3. Confirm that the connector resolves the dynamic firewall IP addresses:

```
config firewall address
    edit "aci-direct-app"
        show
            config firewall address
                edit "aci-direct-app"
                    set uuid 794aaf20-3e33-51ea-57e1-10b5badf3fc7
                    set type dynamic
                    set sdn "aci_direct1"
                    set color 17
                    set filter "Application=lzou-app"
                    config list
                        edit "10.0.5.11"
                        next
                        edit "10.0.5.12"
                        next
                        edit "10.0.6.11"
                        next
                        edit "10.0.6.12"
                        next
                        edit "10.0.6.13"
                        next
                        edit "10.0.6.14"
                        next
                        edit "10.0.7.11"
                        next
                        edit "10.0.7.12"
                        next
                    end
                next
            end
        next
end
```

Nuage SDN connector

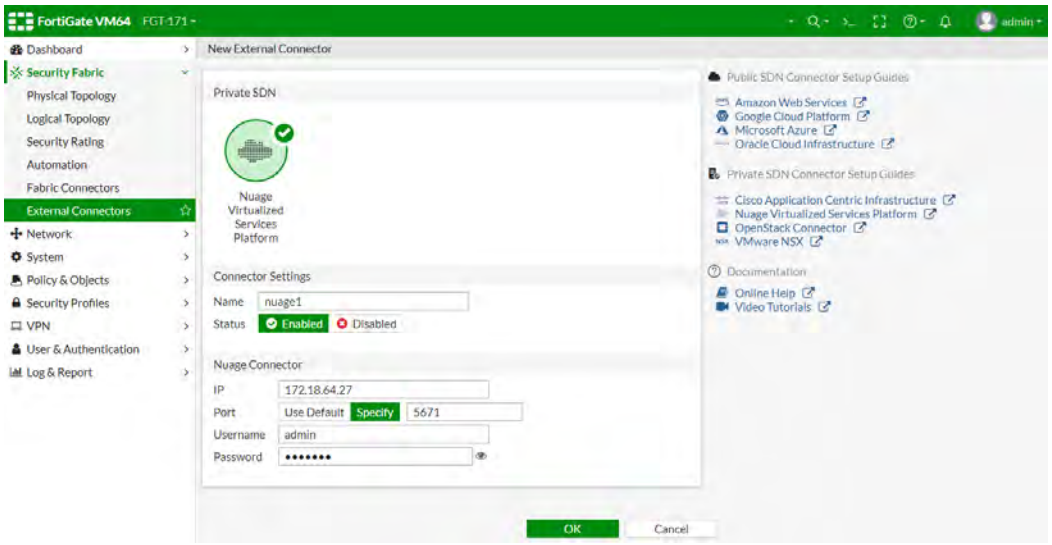
Nuage SDN connectors can be used in dynamic firewall addresses.

The Fortinet SDN Connector for Cisco ACI and Nuage Networks is a standalone connector that connects to SDN controllers within Cisco ACI and Nuage Networks. You must configure a connection to the Fortinet SDN connector in FortiOS to query the dynamic addresses.

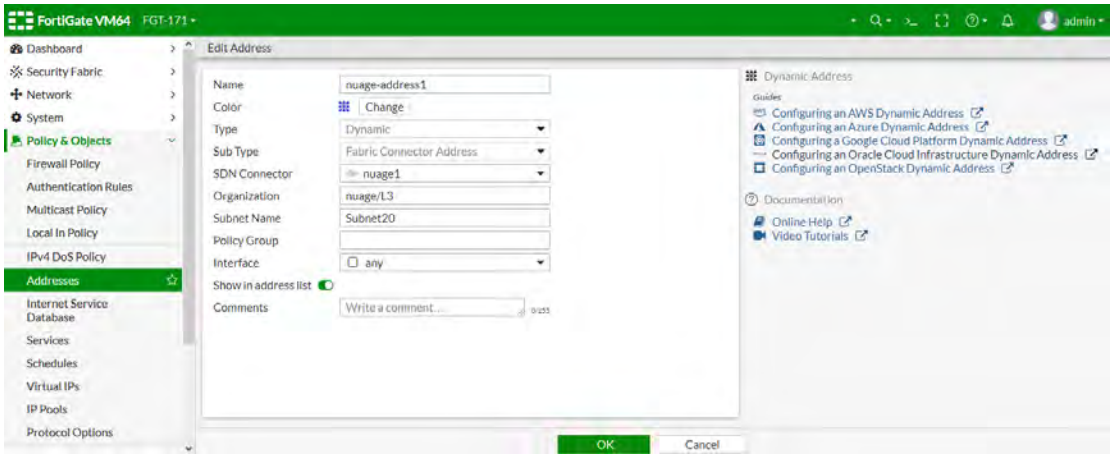
Netskope Exhibit 1017

To configure a Nuage connector in the GUI:

1. Create the Nuage SDN connector:
   a. Go to *Security Fabric* > *External Connectors* and click *Create New*.
   b. In the *Private SDN* section, click *Nuage Virtualized Services Platform*.
   c. Configure the settings as needed.
   d. Click *OK*.



2. Create the dynamic firewall address for the connector:
   a. Go to *Policy & Objects* > *Addresses* and click *Create New* > *Address*.
   b. Configure the following settings:
      i. For *Type*, select *Dynamic*.
      ii. For *Sub Type*, select *Fabric Connector Address*.
      iii. For *SDN Connector*, select the first the first Nuage connector.
      iv. Configure the remaining settings as needed.
   c. Click *OK*.



To verify the SDN connector resolves the dynamic firewall IP addresses in the GUI:

1. Go to *Policy & Objects* > *Addresses*.
2. In the address table, hover over an address to view which IPs it resolves to.

To configure a Nuage connector in the CLI:

1. Create the SDN connector:

```
config system sdn-connector
    edit "nuage1"
        set type nuage
        set server "172.18.64.27"
        set server-port 5671
        set username "admin"
        set password xxxxxxx
    next
end
```

2. Create the dynamic firewall address for the connector:

```
config firewall address
    edit "nuage-address1"
        set type dynamic
        set sdn "nuage1"
        set color 19
        set organization "nuage/L3"
        set subnet-name "Subnet20"
    next
end
```

To verify the SDN connector resolves the dynamic firewall IP addresses in the CLI:

```
# diagnose firewall dynamic list

List all dynamic addresses:
nuage1.nuage.nuage/L3.Subnet20.*: ID(196)
        ADDR(192.168.20.92)
        ADDR(192.168.20.240)
```

Multiple concurrent SDN connectors

You can configure multiple instances configured for every SDN connector. The specific connector instance must be specified when creating a dynamic firewall address.

This topic provides examples of how to create two Microsoft Azure SDN connectors and use them in new dynamic firewall addresses.

To create and use two new SDN connectors with the CLI:

1. Create two new SDN connectors:

```
config system sdn-connector
    edit "azure1"
        set type azure
```

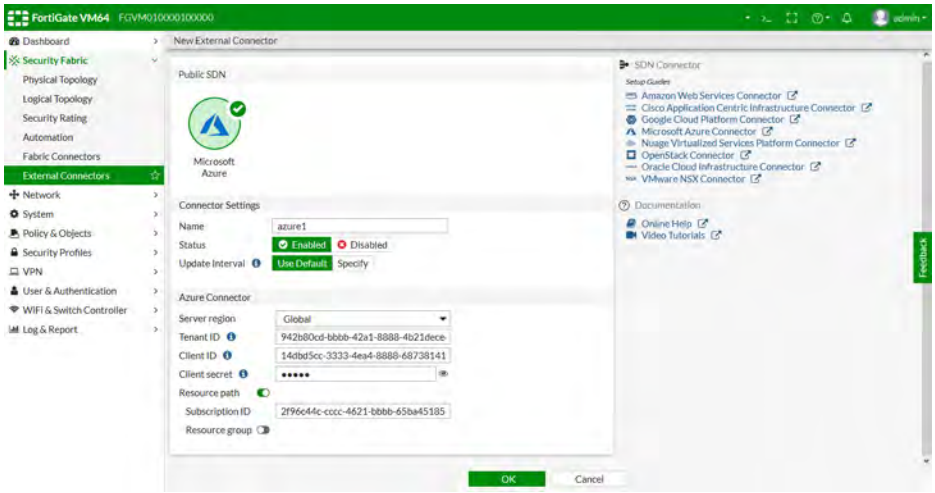
```
        set tenant-id "942b80cd-bbbb-42a1-8888-4b21dece61ba"
        set subscription-id "2f96c44c-cccc-4621-bbbb-65ba45185e0c"
        set client-id "14dbd5cc-3333-4ea4-8888-68738141feb1"
        set client-secret xxxxx
        set update-interval 30
    next
    edit "azure2"
        set type azure
        set tenant-id "942b80cd-bbbb-42a1-8888-4b21dece61ba"
        set client-id "3baa0acc-ffff-4444-b292-0777a2c36be6"
        set client-secret xxxxx
        set update-interval 30
    next
end
```

2. Create new dynamic firewall addresses that use the new connectors:

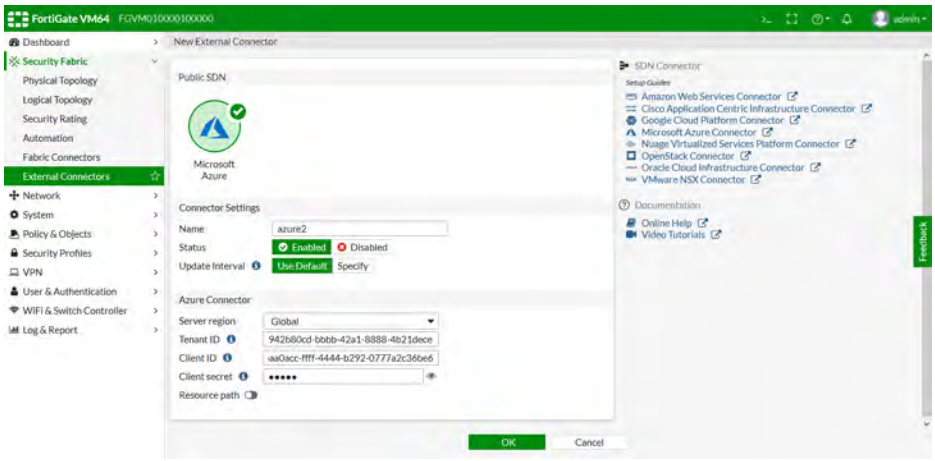
```
config firewall address
    edit "azure-address-location1"
        set type dynamic
        set color 2
        set sdn azure1
        set filter "location=WestUs"
    next
    edit "azure-address-location2"
        set type dynamic
        set color 2
        set sdn azure2
        set filter "location=NorthEurope"
    next
end
```

To create and use two new SDN connectors with the GUI:

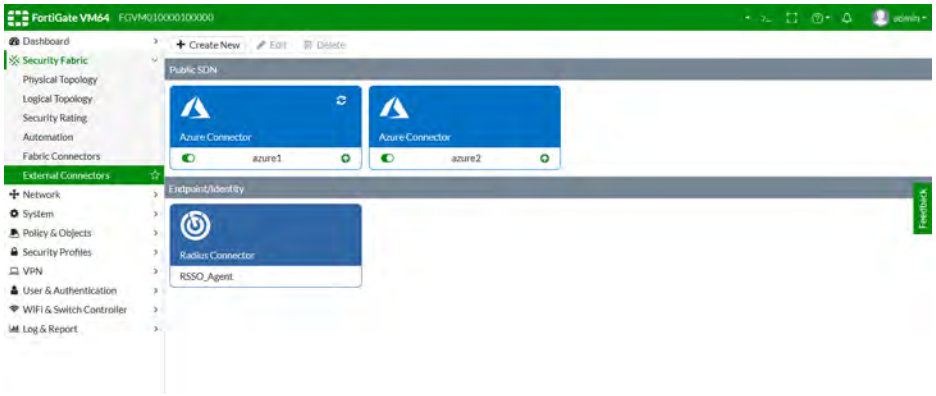
1. Create two new SDN connectors:
   a. Go to *Security Fabric* > *External Connectors*, and click *Create New* in the toolbar.
   b. Click on *Microsoft Azure*.
   c. Fill in the required information, then click *OK*.



d. Repeat the above steps for the second connector.



Two Microsoft Azure connectors will now be created.

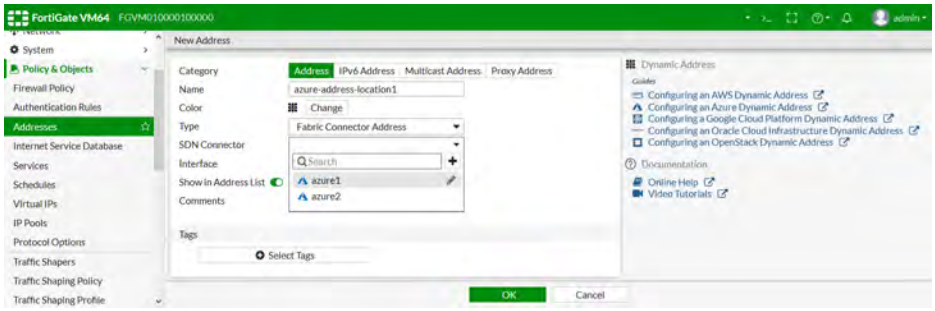


2. Create new dynamic firewall addresses that use the new connectors:

a. Go to *Policy and Objects* > *Addresses* and click *Create New* > *Address* in the toolbar.

b. Enter a name for the address, and select *Fabric Connector Address* for the *Type*.

c. Select one of the previously created SDN connectors from the *SDN Connector* drop down list.



d. Configure the rest of the required information, then click *OK* to create the address.

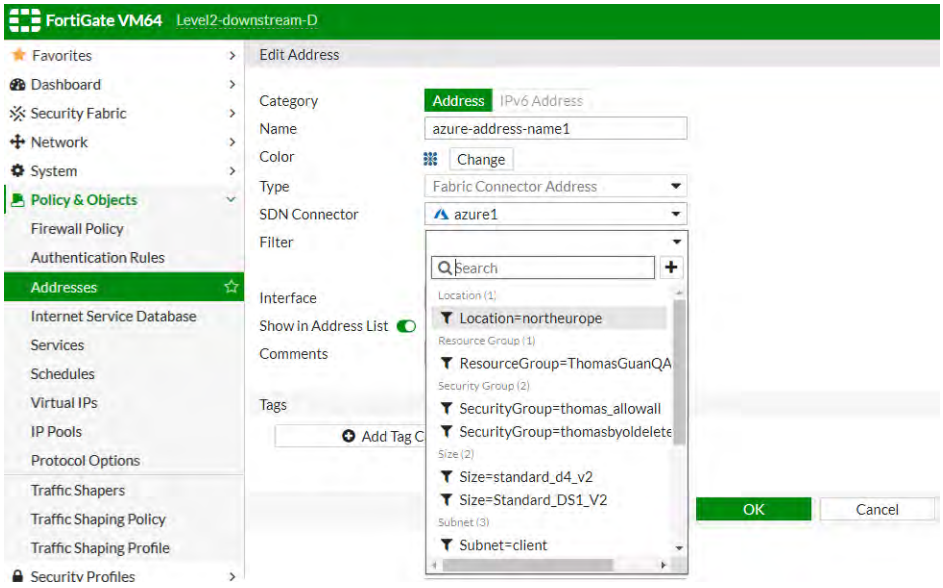
e. Repeat the above steps to create the second address, selecting the other Microsoft Azure SDN connector.

Filter lookup in SDN connectors

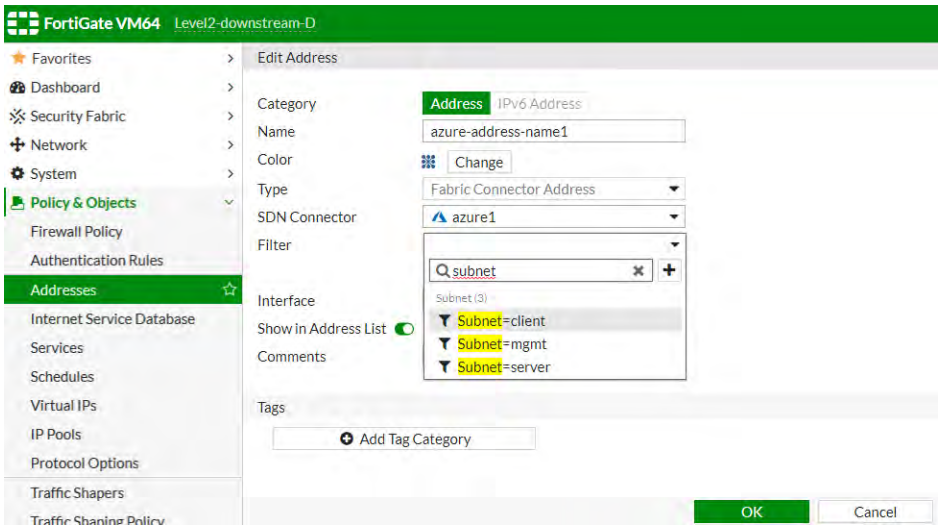
When configuring dynamic address mappings for filters in SDN connectors for Azure, GCP, OpenStack, Kubernetes, and AliCloud, FortiGate can query the filters automatically.

To use the filter lookup:

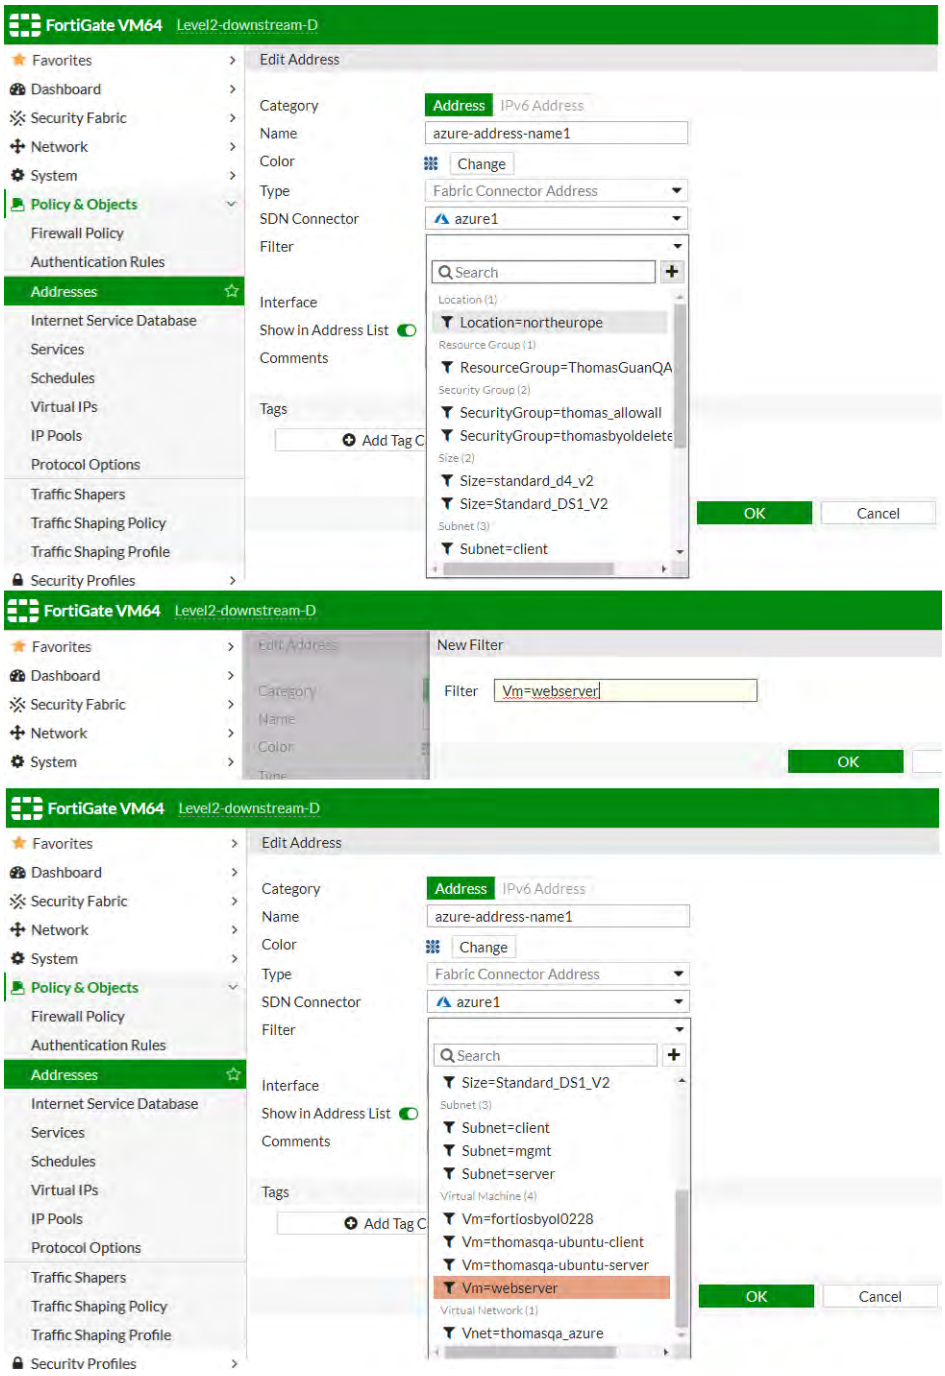
1. Navigate to *Policy & Objects* > *Addresses*.
2. Create or edit an SDN connector type dynamic IP address.
   Supported SDN connector types include: AWS, Azure, GCP, OpenStack, Kubernetes, and AliCloud. The example below is for an Azure SDN connector.
3. In the address *Filter* field, you can perform the following actions:
   - List all available filters.



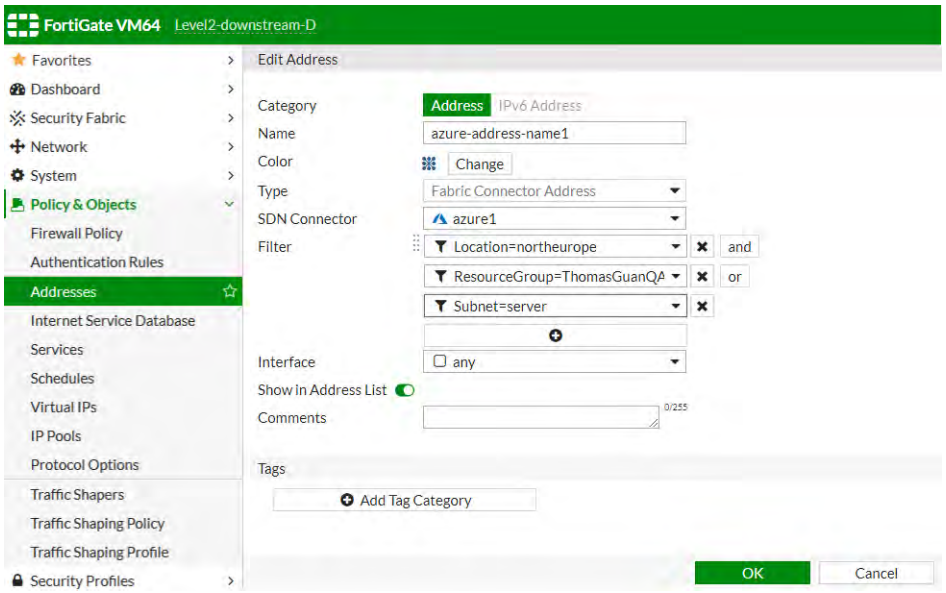
   - Search the available filters.



- Create custom filters.



- Set filter logic [and|or].



Support for wildcard SDN connectors in filter configurations

Wildcards are supported for SDN connectors when configuring dynamic address filters.

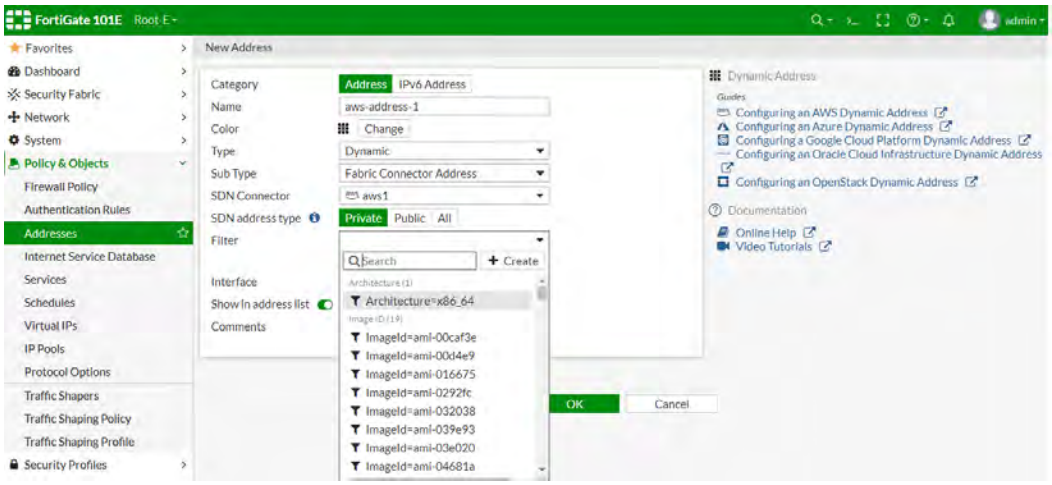
The following SDN connector types are currently supported:

- AWS
- Azure
- Google Cloud Platform
- Kubernetes
- OpenStack
- Oracle Cloud Infrastructure
- VMware ESXi

To configure a dynamic address filter for AWS in the GUI:

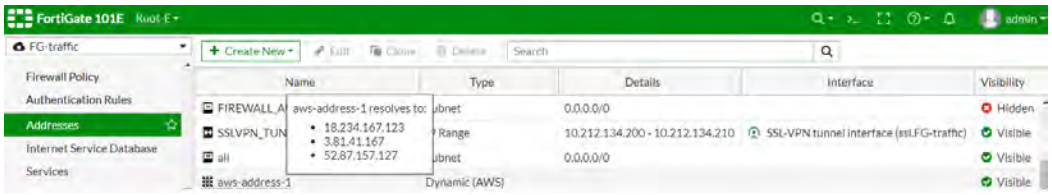
1. Create the SDN connector:
   a. Go to *Security Fabric* > *External Connectors*.
   b. Click *Create New*.
   c. In the *Public SDN* section, click *Amazon Web Services (AWS)*.
   d. Configure the settings as needed.
   e. Click *OK*.
2. Create the dynamic firewall address:
   a. Go to *Policy & Objects* > *Addresses*.
   b. Click *Create New* > *Address*
   c. Enter a name for the address, then configure the following settings:
      - Set *Type* to *Dynamic*.
      - Set *Sub Type* to *Fabric Connector Address*.
      - Set *SDN Connector* to *aws1*.

- Set *SDN address type* to *Private*.
- For *Filter*, click *Create*, enter `Tag.Name=aws*`, the click *OK*.

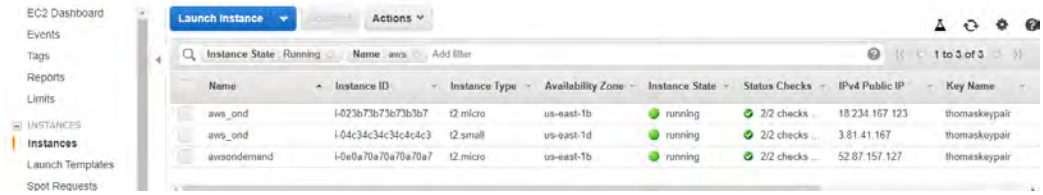


d. Click *OK*.

3. In the address table, hover over the address to view what IPs it resolves to.



4. In AWS, verify to confirm the IP addresses match.



To configure a dynamic address filter for AWS in the CLI:

1. Configure the SDN connector:

```
config firewall address
    edit "aws-address-1"
        set type dynamic
        set sdn "aws1"
        set filter "Tag.Name=aws*"
        set sdn-addr-type public
    next
end
```

2. Create the dynamic firewall address and verify where the IP addresses resolve to:

```
config firewall address
    edit "aws-address-1"
        set type dynamic
        set sdn "aws1"
```

```
            set filter "Tag.Name=aws*"
            set sdn-addr-type public
            config list
                edit "18.234.167.123"
                next
                edit "3.81.41.167"
                next
                edit "52.87.157.127"
                next
            end
        next
    end
```

3. In AWS, verify that the IP addresses match.

Kubernetes (K8s) SDN connectors

The following topics provide information about configuring Kubernetes SDN connectors:

- Private Cloud K8s SDN connector on page 245
- AWS Kubernetes (EKS) SDN connector on page 247
- Azure Kubernetes (AKS) SDN connector on page 253
- GCP Kubernetes (GKE) SDN connector on page 250
- Oracle Kubernetes (OKE) SDN connector on page 255

Private Cloud K8s SDN connector

FortiOS automatically updates dynamic and cluster IP addresses for Kubernetes (K8s) by using a K8s SDN connector, enabling FortiOS to manage K8s pods as global address objects, as with other connectors. This includes mapping the following attributes from K8s instances to dynamic address groups in FortiOS:

Filter	Description
Namespace	Filter service IP addresses in a given namespace.
ServiceName	Filter service IP addresses by the given service name.
NodeName	Filter node IP addresses by the given node name.
PodName	Filter IP addresses by the pod name.
Label.XXX	Filter service or node IP addresses with the given label XXX. For example: `K8S_Label.app=nginx`.

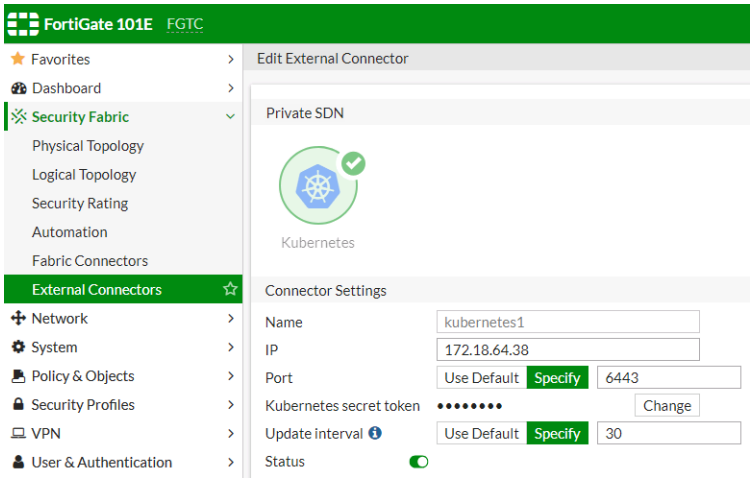
FortiOS 6.2.3 and later collects cluster IP addresses in addition to external IP addresses for exposed Kubernetes services.



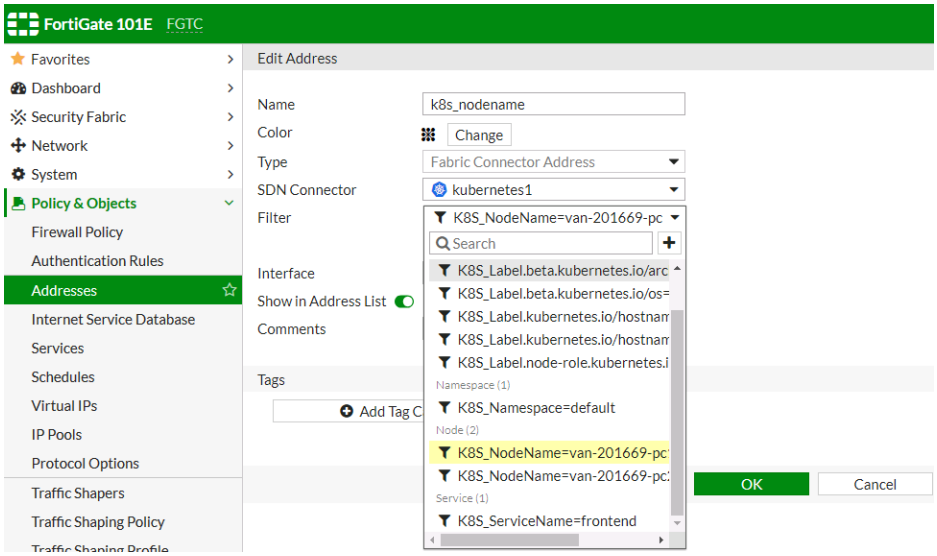
There is no maximum limit for the number of IP addresses populated with the filters.

To configure K8s SDN connector using the GUI:

1. Configure the K8s SDN connector:
   a. Go to *Security Fabric* > *External Connectors*.
   b. Click *Create New*, and select *Kubernetes*.
   c. Configure as shown substituting the IP address, port number, and secret token for your deployment. The update interval is in seconds.

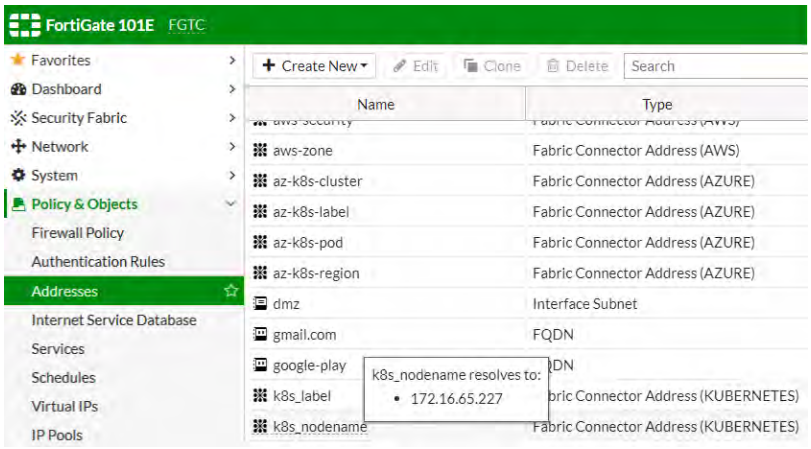


2. Create a dynamic firewall address for the configured K8S SDN connector:
   a. Go to *Policy & Objects* > *Addresses*.
   b. Click *Create New*, then select *Address*.
   c. Configure the address as shown, selecting the desired filter in the *Filter* dropdown list. In this example, the K8s SDN connector will automatically populate and update IP addresses only for node instances that match the specified node name:



3. Ensure that the K8s SDN connector resolves dynamic firewall IP addresses:
   a. Go to *Policy & Objects* > *Addresses*.
   b. Hover over the address created in step 2 to see a list of IP addresses for node instances that match the node

name configured in step 2:



To configure K8s SDN connector using CLI commands:

1. Configure the K8s SDN connector:

```
config system sdn-connector
   edit "kubernetes1"
      set type kubernetes
      set server "172.18.64.38"
      set server-port 6443
      set secret-token xxxxx
      set update-interval 30
   next
end
```

2. Create a dynamic firewall address for the configured K8s SDN connector with the supported K8s filter. In this example, the K8s SDN connector will automatically populate and update IP addresses only for node instances that match the specified node name:

```
config firewall address
   edit "k8s_nodename"
      set type dynamic
      set sdn "kubernetes1"
      set filter "K8S_NodeName=van-201669-pc1"
   next
end
```

3. Confirm that the K8s SDN connector resolves dynamic firewall IP addresses using the configured filter:

```
config firewall address
   edit "k8s_nodename"
      set type dynamic
      set sdn "kubernetes1"
      set filter "K8S_NodeName=van-201669-pc1"
      config list
         edit "172.16.65.227"
         next
      end
   next
end
```

AWS Kubernetes (EKS) SDN connector

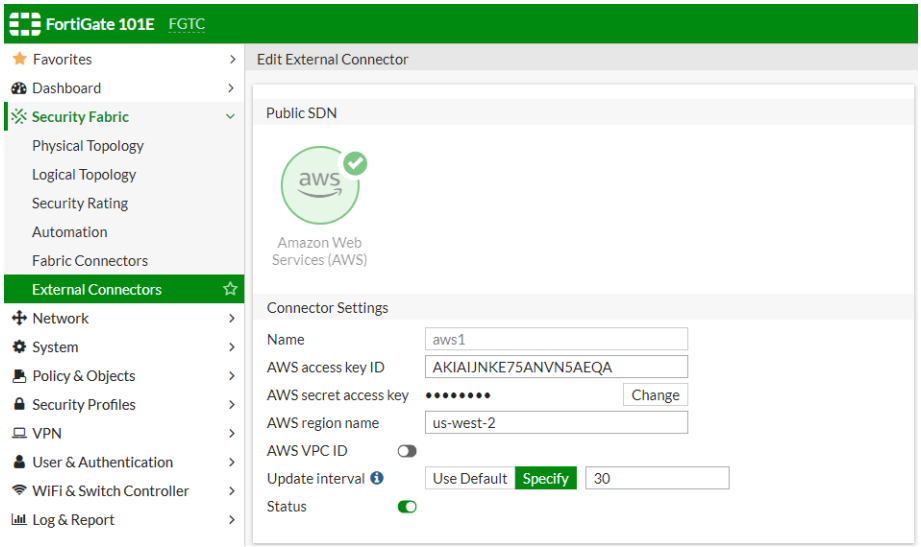
AWS SDN connectors support dynamic address groups based on AWS Kubernetes (EKS) filters.

To filter out the Kubernetes IP addresses, the following address filters have been introduced:

k8s_cluster	Name of Kubernetes cluster.
k8s_namespace	Namespace of a Kubernetes service or pod.
k8s_svcname	Name of a Kubernetes service.
k8s_nodename	Name of a Kubernetes node.
k8s_zone	Zone of a Kubernetes node.
k8s_region	Region of a Kubernetes node.
k8s_podname	Name of a Kubernetes pod.
k8s_label.xxx	Name of label of a Kubernetes resource (cluster/service/node/Pod).

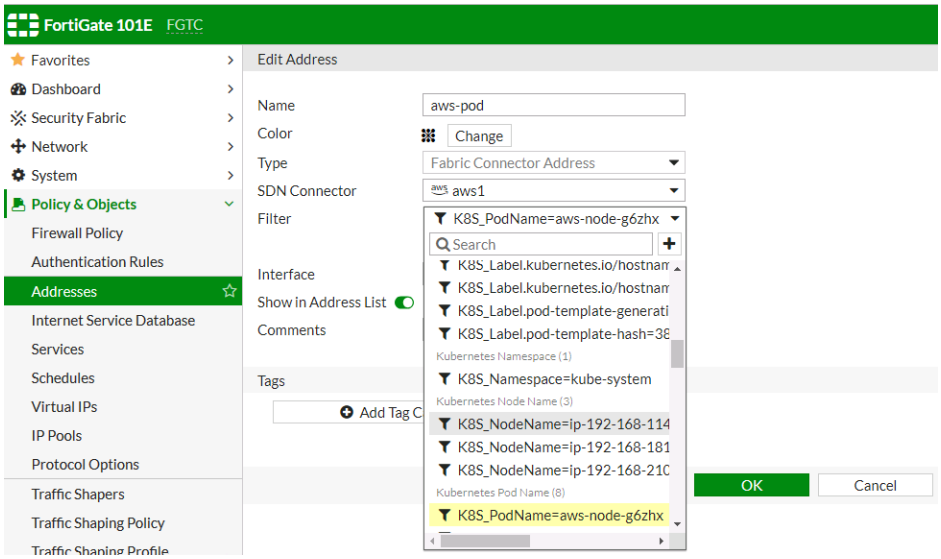
To enable an AWS SDN connector to fetch IP addresses from AWS Kubernetes:

1. Go to *Security Fabric* > *External Connectors*, and configure an SDN connector for AWS Kubernetes.

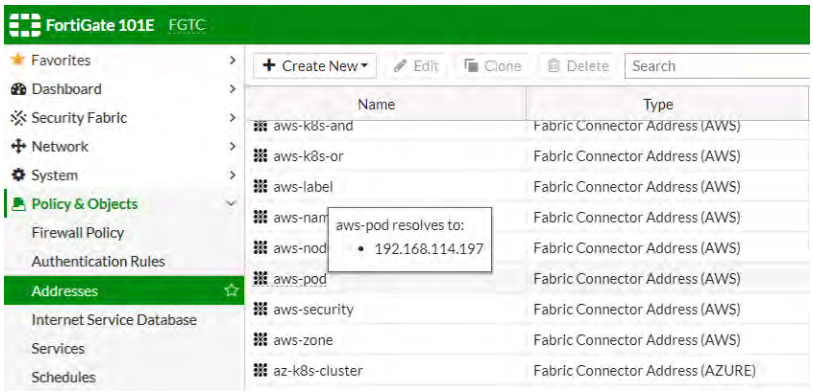


2. Go to *Policies & Objects* > *Addresses* and create a dynamic firewall address for the configured SDN connector using the supported Kubernetes filter.

3. To filter out the Kubernetes IP addresses, select the address filter or filters.



4. Configure the rest of the settings, then click *OK*.
The dynamic firewall address IP is resolved by the SDN connector.



To configure an AWS Kubernetes SDN connector through the CLI:

1. Configure an SDN connector for Kubernetes:

```
config system sdn-connector
   edit "aws1"
      set type aws
      set access-key "AKIAIJNKE75ANVN5AEQA"
      set secret-key xxxxx
      set region "us-west-2"
      set update-interval 30
   next
end
```

2. Create a dynamic firewall address for the SDN connector with a supported Kubernetes filter:

```
config firewall address
   edit "aws-pod"
      set type dynamic
      set sdn "aws1"
      set filter "K8S_PodName=aws-node-g6zhx"
   next
```


Netskope Exhibit 1017

```
end
```

The dynamic firewall address IP is resolved by the SDN connector:

```
config firewall address
   edit "aws-pod"
      set type dynamic
      set sdn "aws1"
      set filter "K8S_PodName=aws-node-g6zhx"
      config list
         edit "192.168.114.197"
         next
      end
   next
end
```

GCP Kubernetes (GKE) SDN connector

Google Cloud Platform (GCP) SDN connectors support dynamic address groups based on GCP Kubernetes Engine (GKE) filters.

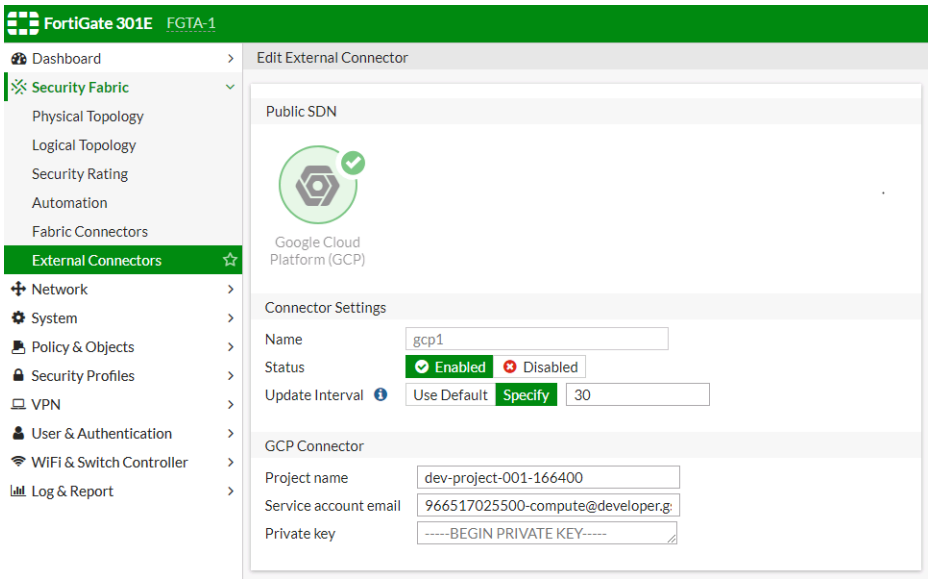
To filter out the Kubernetes IP addresses, the following address filters have been introduced:

k8s_cluster	Name of Kubernetes cluster.
k8s_nodepool	Name of node pool for a Kubernetes cluster.
k8s_namespace	Namespace of a Kubernetes service or pod.
k8s_servicename	Name of a Kubernetes service.
k8s_nodename	Name of a Kubernetes node.
k8s_zone	Zone of a Kubernetes node.
k8s_region	Region of a Kubernetes node.
k8s_podname	Name of a Kubernetes pod.
k8s_label.xxx	Name of label of a Kubernetes resource (cluster/service/node/Pod).

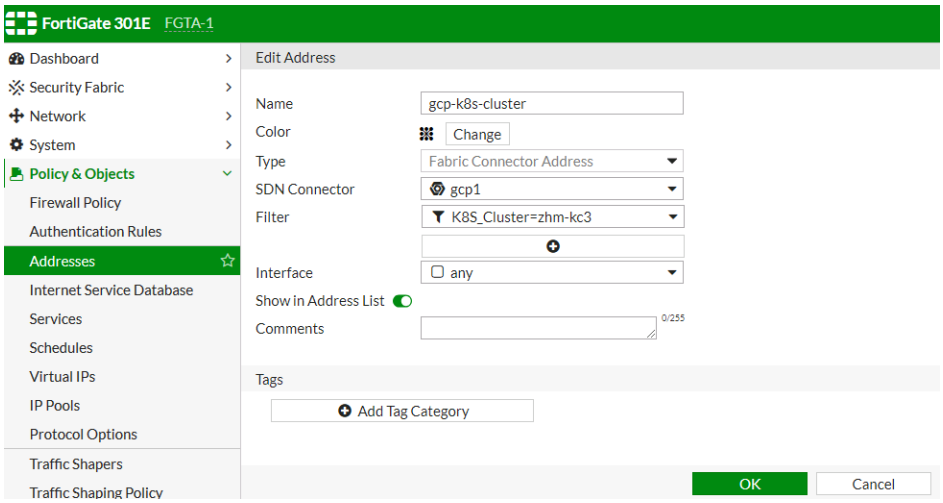
Netskope Exhibit 1017

To enable a GCP SDN connector to fetch IP addresses from GKE:

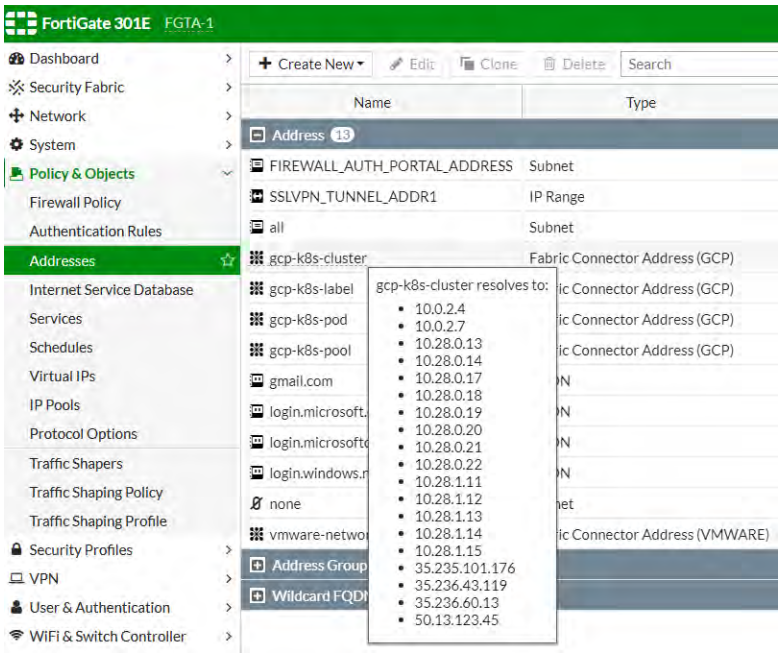
1. Go to *Security Fabric* > *External Connectors*, and configure an SDN connector for GCP.



2. Go to *Policies & Objects* > *Addresses* and create a dynamic firewall address for the configured SDN connector using the supported Kubernetes filter.
3. To filter out the Kubernetes IP addresses, select the address filter or filters. In this example, the GCP SDN connector will automatically populate and update IP addresses only for instances that belong to the zhm-kc3 cluster:



4. Configure the rest of the settings, then click *OK*.
   The dynamic firewall address IP is resolved by the SDN connector.



To configure a GCP Kubernetes SDN connector through the CLI:

1. Configure an SDN connector for Kubernetes:

```
config system sdn-connector
   edit "gcp1"
      set type gcp
      set gcp-project "dev-project-001-166400"
      set service-account "966517025500-compute@developer.gserviceaccount.com"
      set update-interval 30
   next
end
```

2. Create a dynamic firewall address for the SDN connector with a supported Kubernetes filter:

```
config firewall address
   edit "gcp-k8s-cluster"
      set type dynamic
      set sdn "gcp1"
      set filter "K8S_Cluster=zhm-kc3"
   next
end
```

The dynamic firewall address IP is resolved by the SDN connector:

```
config firewall address
   edit "gcp-k8s-cluster"
      set type dynamic
      set sdn "gcp1"
      set filter "K8S_Cluster=zhm-kc3"
      config list
         edit "10.0.2.4"
         next
         edit "10.0.2.7"
         next
         edit "10.28.0.13"
         next
      end
```

```
    next
end
```

Azure Kubernetes (AKS) SDN connector

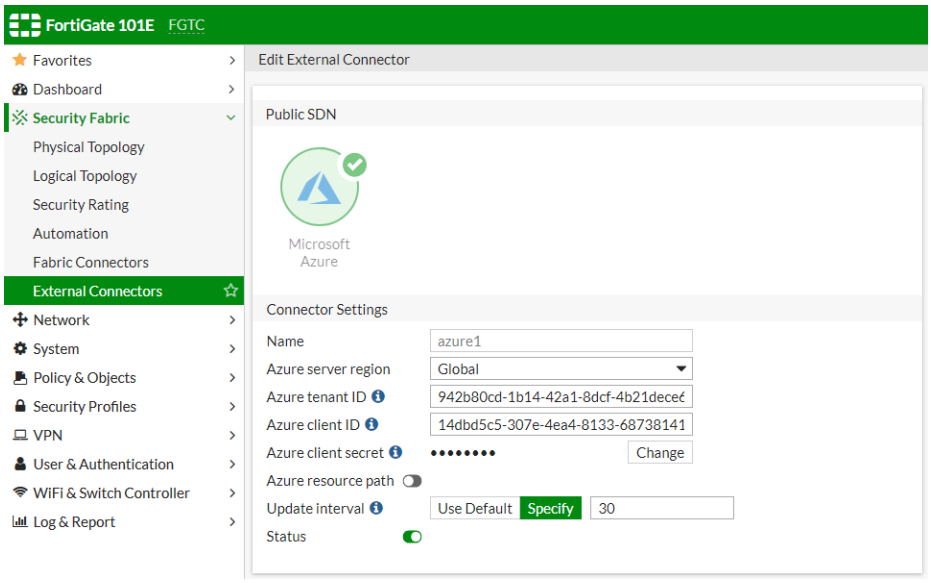
Azure SDN connectors support dynamic address groups based on Azure Kubernetes (AKS) filters.

To filter out the Kubernetes IP addresses, the following address filters have been introduced:

k8s_cluster	Kubernetes cluster name.
k8s_namespace	Namespace of a Kubernetes service or pod.
k8s_svcname	Kubernetes service name.
k8s_nodename	Kubernetes node name.
k8s_region	Kubernetes node region.
k8s_podname	Kubernetes pod name.
k8s_label.xxx	Name of label of a Kubernetes resource (cluster/service/node/Pod).

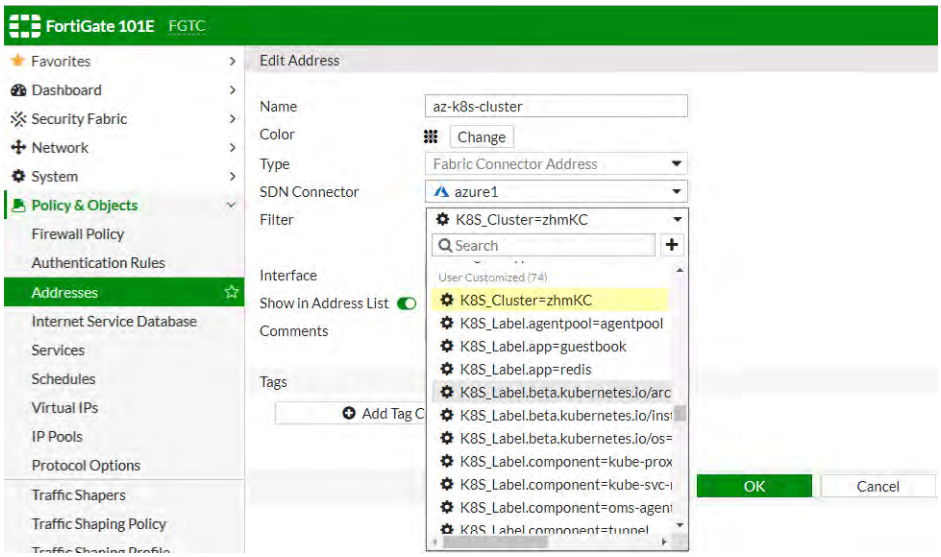
To enable an Azure SDN connector to fetch IP addresses from Azure Kubernetes:

1. Configure the Azure SDN connector:
   a. Go to *Security Fabric* > *External Connectors*.
   b. Click *Create New*, and select *Azure*.
   c. Configure as shown substituting the region, tenant and client IDs, and client secret for your deployment. The update interval is in seconds.

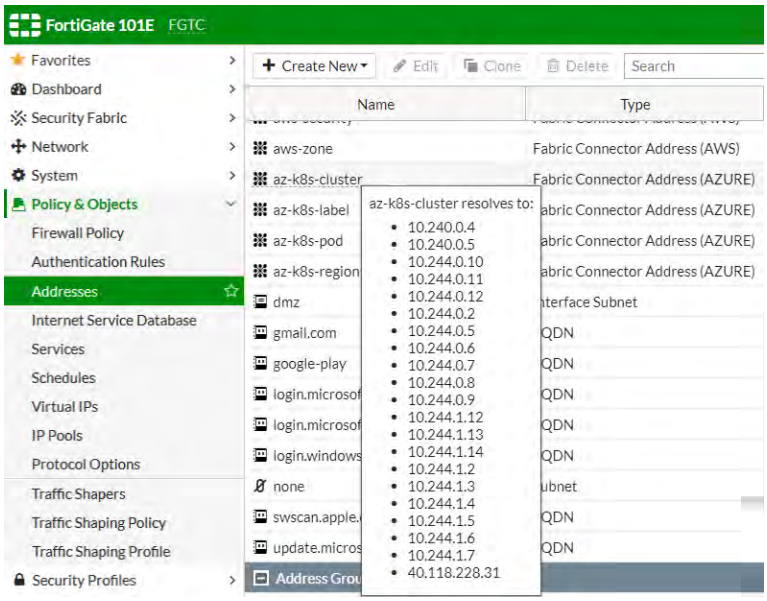


2. Create a dynamic firewall address for the configured K8S SDN connector:
   a. Go to *Policy & Objects* > *Addresses*.
   b. Click *Create New*, then select *Address*.
   c. Configure the address as shown, selecting the desired filter in the *Filter* dropdown list. In this example, the Azure SDN connector will automatically populate and update IP addresses only for instances that belong to the

zhmKC cluster:



3. Ensure that the K8S SDN connector resolves dynamic firewall IP addresses:
   a. Go to *Policy & Objects* > *Addresses*.
   b. Hover over the address created in step 2 to see a list of IP addresses for instances that belong to the zhmKC cluster as configured in step 2:



To configure an Azure Kubernetes SDN connector through the CLI:

1. Configure an SDN connector for Kubernetes:

```
config system sdn-connector
  edit "azure1"
    set type azure
    set tenant-id "942b80cd-1b14-42a1-8dcf-4b21dece61ba"
    set client-id "14dbd5c5-307e-4ea4-8133-68738141feb1"
    set client-secret xxxxx
    set update-interval 30
  next
```

```
end
```

2. Create a dynamic firewall address for the SDN connector with a supported Kubernetes filter. In this example, the Azure SDN connector will automatically populate and update IP addresses only for instances that belong to the zhmKC cluster:

```
config firewall address
   edit "az-k8s-cluster"
      set type dynamic
      set sdn "azure1"
      set filter "K8S_Cluster=zhmKC"
   next
end
```

3. Confirm that the Azure SDN connector resolves dynamic firewall IP addresses using the configured filter:

```
config firewall address
   edit "az-k8s-cluster"
      set type dynamic
      set sdn "azure1"
      set filter "K8S_Cluster=zhmKC"
      config list
         edit "10.240.0.4"
         next
         edit "10.240.0.5"
         next
         edit "10.244.0.10"
         next
      end
   next
end
```

Oracle Kubernetes (OKE) SDN connector

OCI SDN connectors support dynamic address groups based on Oracle Kubernetes (OKE) filters.

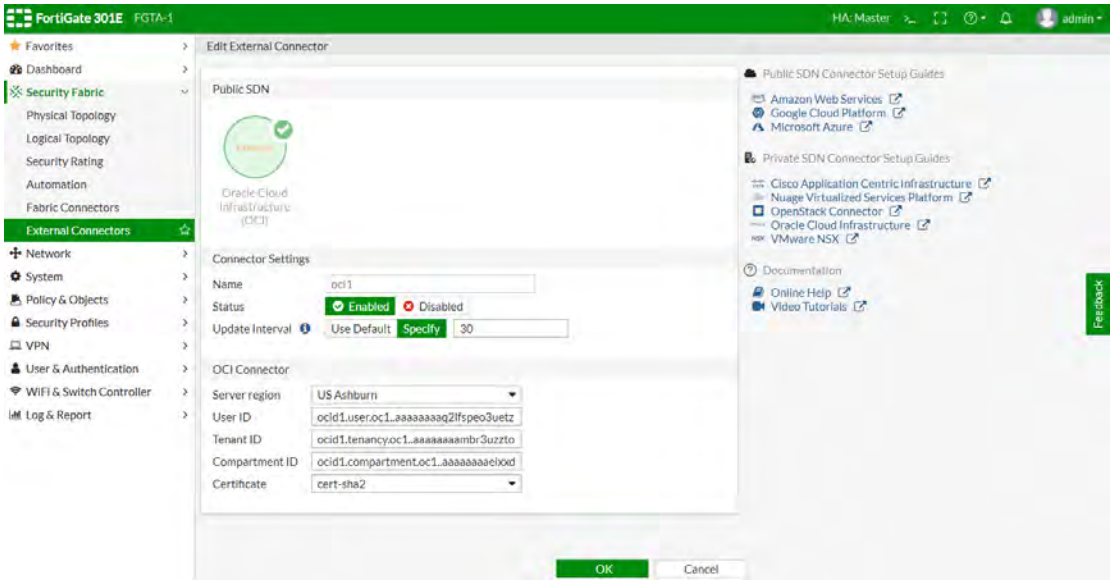
To filter out the Kubernetes IP addresses, the following address filters have been introduced:

k8s_compartment	Name of compartment that the Kubernetes cluster created in.
k8s_cluster	Name of Kubernetes cluster.
k8s_namespace	Namespace of a Kubernetes service or pod.
k8s_servicename	Name of a Kubernetes service.
k8s_nodename	Name of a Kubernetes node.
k8s_region	Region of a Kubernetes node.
k8s_zone	Zone of a Kubernetes node.
k8s_podname	name of a Kubernetes pod.
k8s_label.xxx	Name of label of a Kubernetes resource (cluster/service/node/Pod)

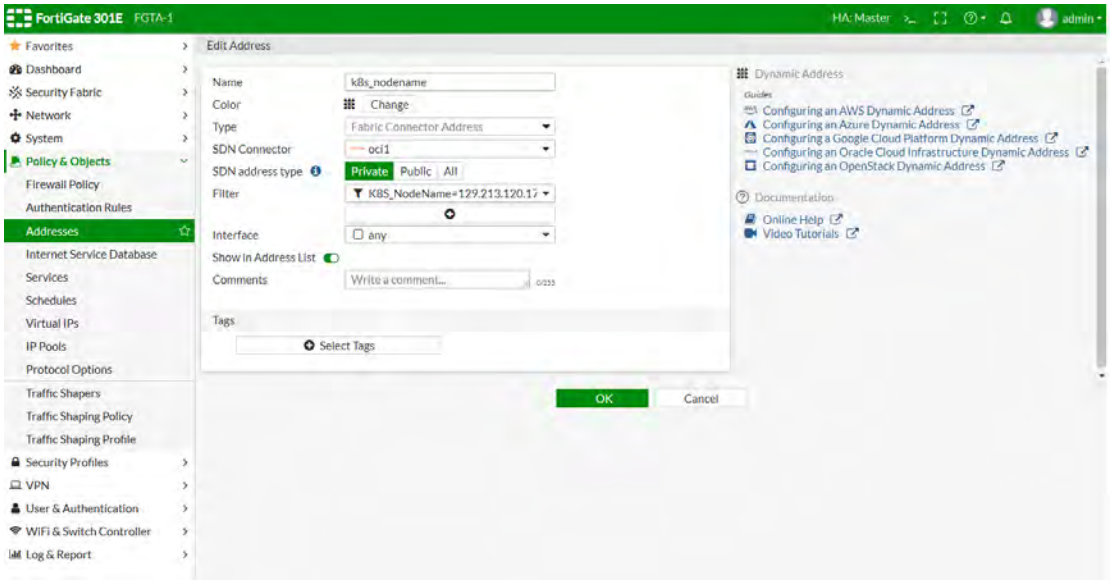
Netskope Exhibit 1017

To enable an OCI SDN connector to fetch IP addresses from Oracle Kubernetes:

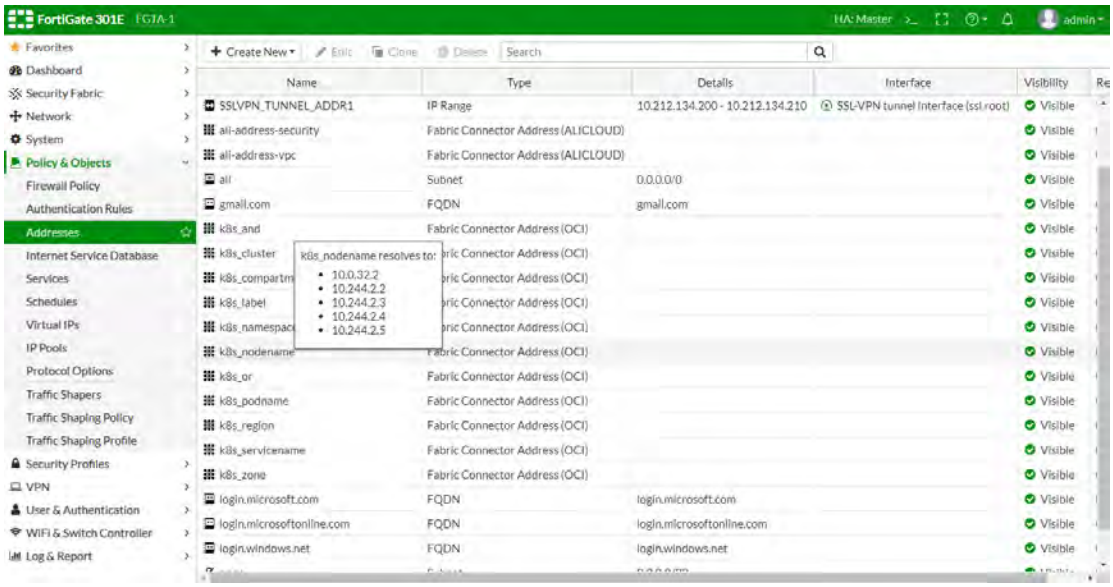
1. Configure the OCI SDN connector:
   a. Go to *Security Fabric* > *External Connectors*.
   b. Click *Create New*, and select *Oracle Cloud Infrastructure (OCI)*.
   c. Configure as shown substituting the region, tenant and client IDs, and client secret for your deployment. The update interval is in seconds.



2. Create dynamic firewall addresses for the configured SDN connector with supported Kubernetes filter:
   a. Go to *Policy & Objects* > *Addresses*.
   b. Click *Create New*, then select *Address*.
   c. Configure the addresses.



3. Confirm that the SDN connector resolves dynamic firewall IP addresses:
   a. Go to *Policy & Objects* > *Addresses*.
   b. Hover over the address created in step 2 to see a list of IP addresses for instances:



To configure an SDN connector through the CLI:

1. Configure the OCI SDN connector:

```
config system sdn-connector
   edit "oci1"
      set type oci
      set tenant-id
           "ocid1.tenancy.oc1..aaaaaaaambr3uzztoyhweohbzqqdo775h7d3t54zpmzkp4b2cf35vs5
           5cxxx"
      set user-id
           "ocid1.user.oc1..aaaaaaaaq2lfspeo3uetzbzpiv2pqvzzevozccnys347stwssvizqlatfx
           xx"
      set compartment-id
           "ocid1.compartment.oc1..aaaaaaaaelxxdjazqo7nzczpgypyiqcgkmytjry6nfq5345vw7e
           avpwnmxxx"
      set oci-region ashburn
      set oci-cert "cert-sha2"
      set update-interval 30
   next
end
```

2. Create dynamic firewall addresses for the configured SDN connector with supported Kubernetes filter:

```
config firewall address
   edit "k8s_nodename"
      set type dynamic
      set sdn "oci1"
      set filter "K8S_NodeName=129.213.120.172"
   next
end
```

3. Confirm that the SDN connector resolves dynamic firewall IP addresses:

```
config firewall address
   edit "k8s_nodename"
      set type dynamic
```

```
        set sdn "oci1"
        set filter "K8S_NodeName=129.213.120.172"
        config list
           edit "10.0.32.2"
           next
           edit "10.244.2.2"
           next
           edit "10.244.2.3"
           next
           edit "10.244.2.4"
           next
           edit "10.244.2.5"
           next
        end
     next
  end
```

Endpoint/Identity connectors

SSO fabric connectors integrate SSO authentication into the network. This allows users to enter their credentials only once, and have those credentials reused when accessing other network resources through the FortiGate.

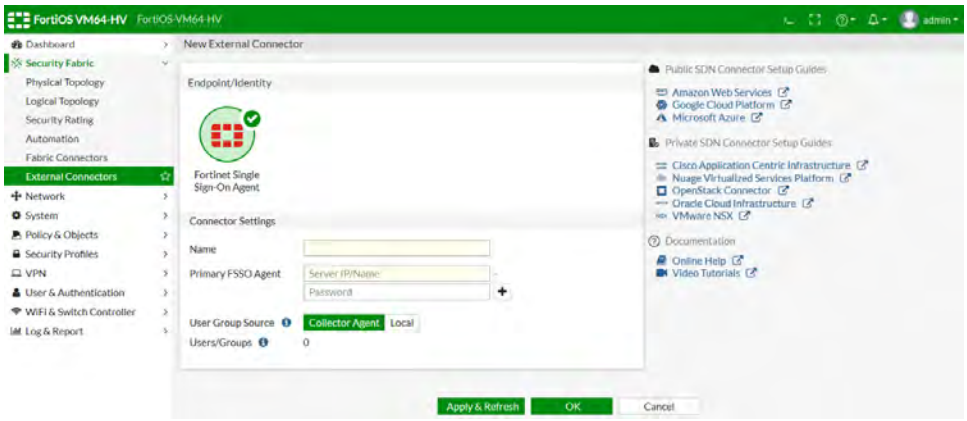
The following fabric connectors are available:

- Fortinet single sign-on agent on page 258
- Poll Active Directory server on page 259
- Symantec endpoint connector on page 259
- RADIUS single sign-on (RSSO) agent on page 266

Fortinet single sign-on agent

To create an FSSO agent connector in the GUI:

1. Go to *Security Fabric* > *External Connectors*.
2. Click *Create New*.
3. In the *Endpoint/Identity* section, click *Fortinet Single Sign-On Agent*.



4. Fill in the *Name*, and *Primary FSSO Agent* server IP address or name and *Password*.
5. Optionally, add more FSSO agents by clicking the plus icon.

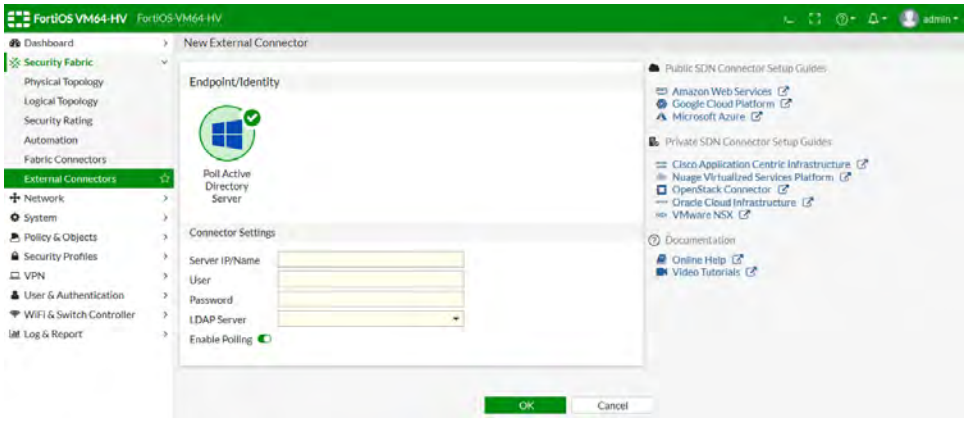
6. Select the *User Group Source*:
   - Collector Agent: User groups will be pushed to the FortiGate from the collector agent. Click *Apply & Refresh* to fetch group filters from the collector agent.
   - Local: User groups will be specified in the FortiGate unit's configuration. Select the LDAP server from the dropdown list, then select the *Users*, *Groups*, and *Organizational Units*. Enable or disable *Recursive* as required.
7. Click *OK*.

Poll Active Directory server

The FortiGate unit can authenticate users and allow them network access based on groups membership in Windows Active Directory (AD).

To create an AD server connector in the GUI:

1. Go to *Security Fabric* > *External Connectors*.
2. Click *Create New*.
3. In the *Endpoint/Identity* section, click *Poll Active Directory Server*.



4. Fill in the *Server IP/Name*, *User*, and *Password* for the AD server.
5. Select the LDAP server from the list.
6. If necessary, disable *Enable Polling*. This can be used to temporarily stop the FortiGate from polling security event logs on the Windows logon server, for troubleshooting purposes.
7. Click *OK*.

Symantec endpoint connector

With the Fabric connector for Symantec Endpoint Protection Manager (SEPM), you can use the client IP information from SEPM to assign to dynamic IP addresses on FortiOS.

When communication between FortiGate and SEPM is established, FortiGate polls every minute for updates via TLS over port 8446. You can use the CLI to change the default one minute polling interval.

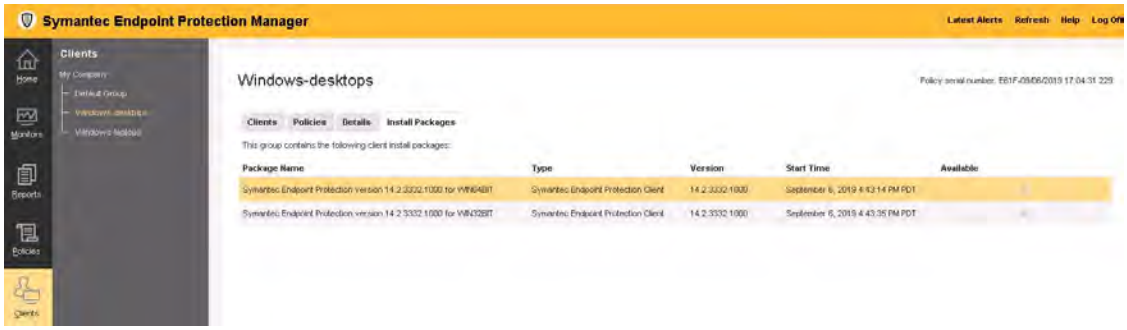
For example, you can create a dynamic Fabric Connector IP address subtype and use it in firewall policies as the source address. The dynamic IP address contains all IP addresses sent by SEPM.

This example shows a dynamic IP address with SEPM and one client PC managed by SEPM using FortiGate as the default gateway.

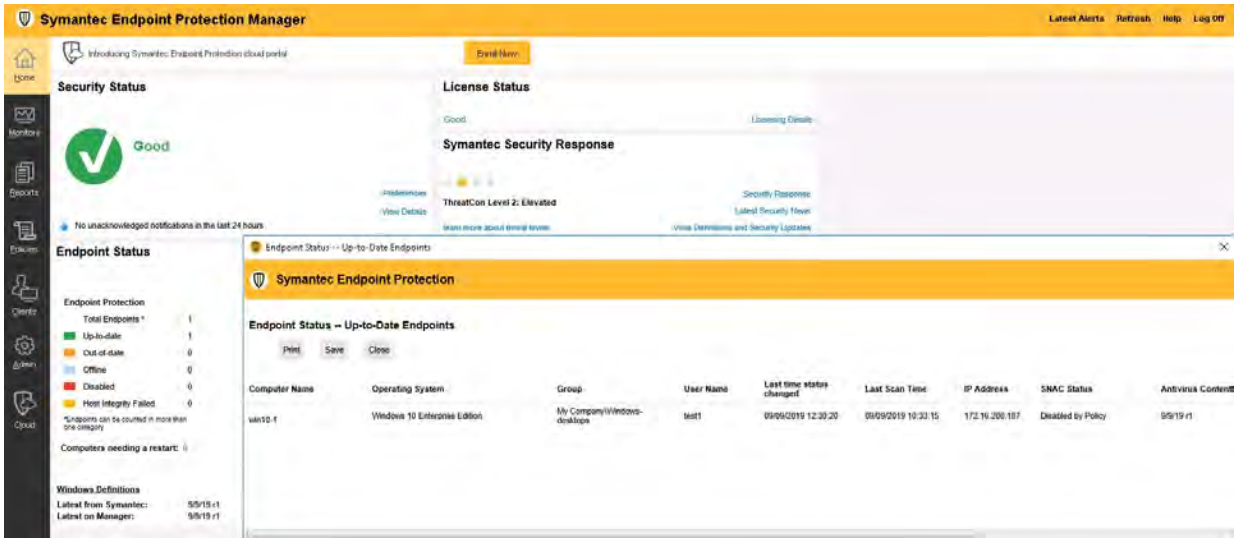
Netskope Exhibit 1017

To configure SEPM on a managed client PC:

1. In SEPM, create client packages for client hosts and group them into SEPM groups.
   You can install packages locally on clients or download them directly from SEPM.

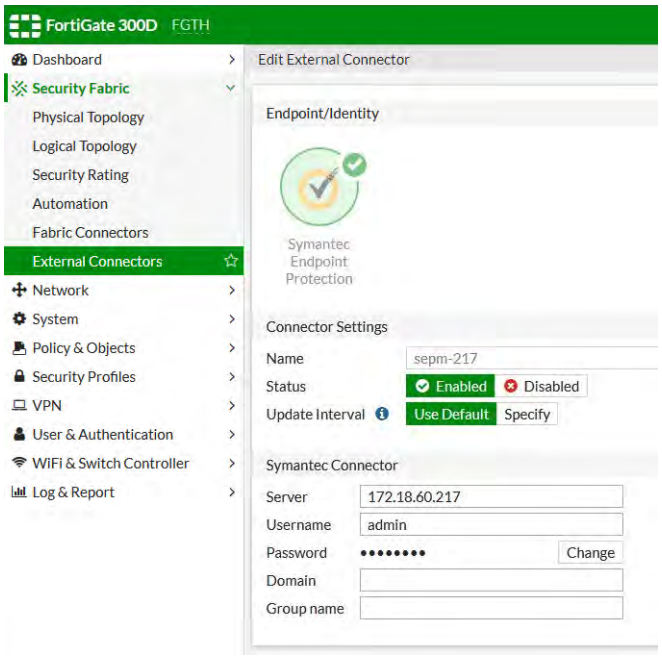


2. When a package is installed on the client host, the host is considered managed by SEPM.
   Even if the host has multiple interfaces, only one IP per host is displayed.

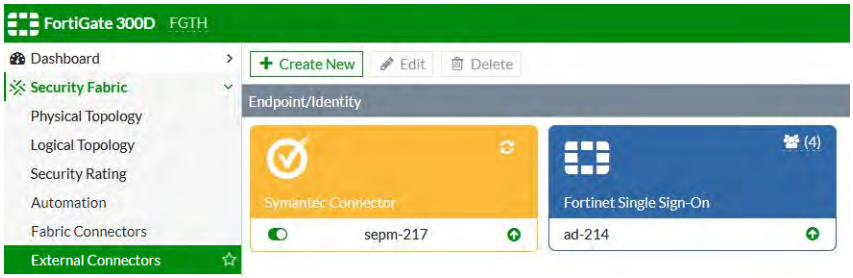


To configure Symantec endpoint connector on FortiGate in the GUI:

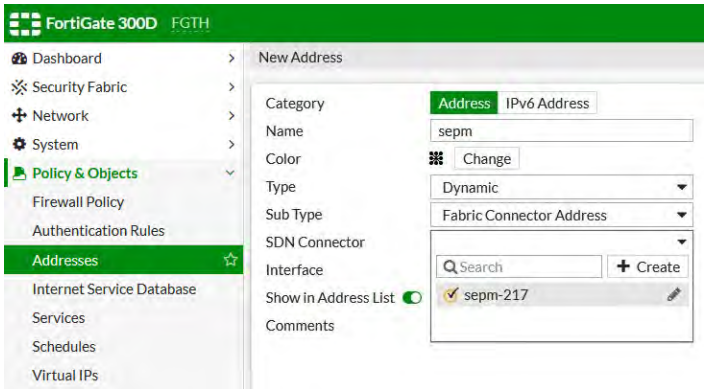
1. Go to *Security Fabric* > *External Connectors*.
2. Click *Create New*.
3. Click *Symantec Endpoint Protection*.
   - In the *Connector Settings* section, if options are left empty, then all SEPM domains and groups are monitored.
   - In the *Symantec Connector* section:
     - In the *Server* field, enter the SEPM IP address.
     - Enter the *Username* and *Password*.
     - If you want to limit the domains or groups that are monitored, enter the information in *Domain* and *Group name*.



4. Click *OK*.
   When the connection is established, you can see a green up arrow in the bottom right of the card. You might need to refresh your browser to see the established connection.

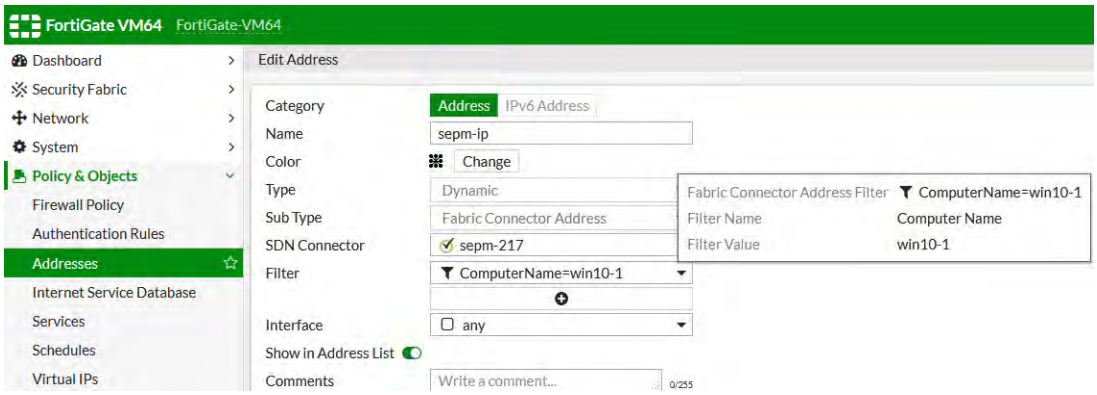


5. Go to *Policy & Objects* > *Addresses*.
6. Click *Create New* > *Address*.
   - Set *Type* to *Dynamic*.
   - Set *Sub Type* to *Fabric Connector Address*.
   - Set *SDN Connector* to the Fabric Connector that you just created.



7. Click *OK*.

8. Edit the address to see the configuration.
   - *Filter* shows the hostnames of the client PCs managed by SEPM. The GUI shows the ComputerName by default. You can change this using the CLI; see Specify filters for details.

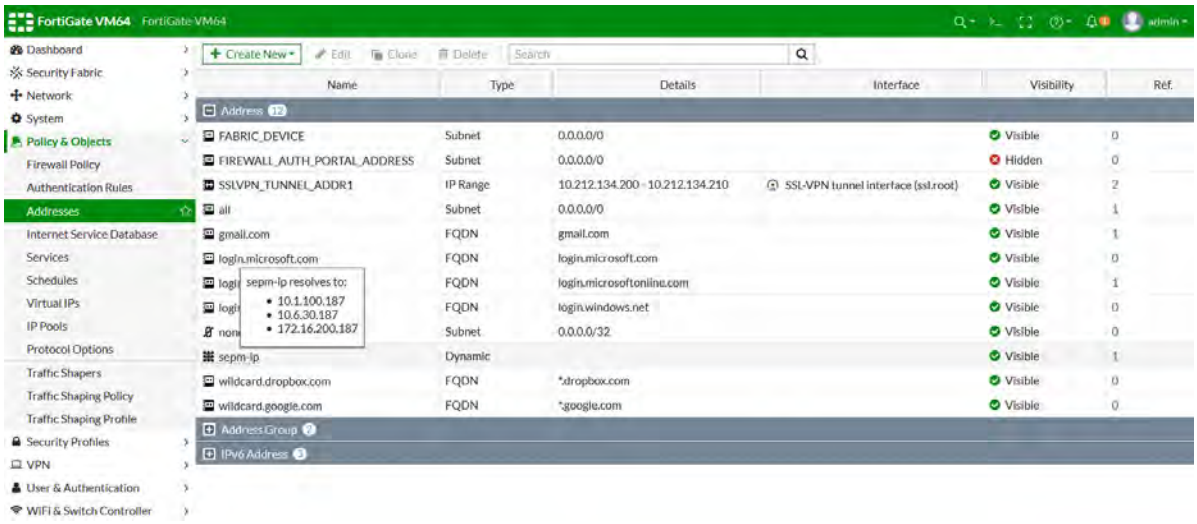




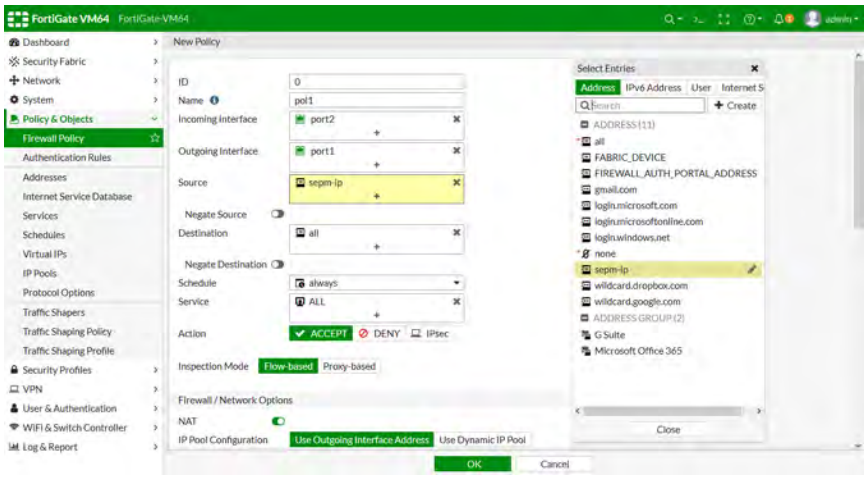
Filter options are only available for active computers that are configured and registered in SEPM. Free-form filters can be created manually by clicking *Create* and entering the filter, in the format: `filter_type=value`.

Possible manual filter types are: `GroupName`, `GroupID`, `ComputerName`, `ComputerUUID`, and `OSName`. For example: `GroupName=MyGroup`.

9. In *Policy & Objects* > *Addresses*, you can see all the IP addresses of the host.

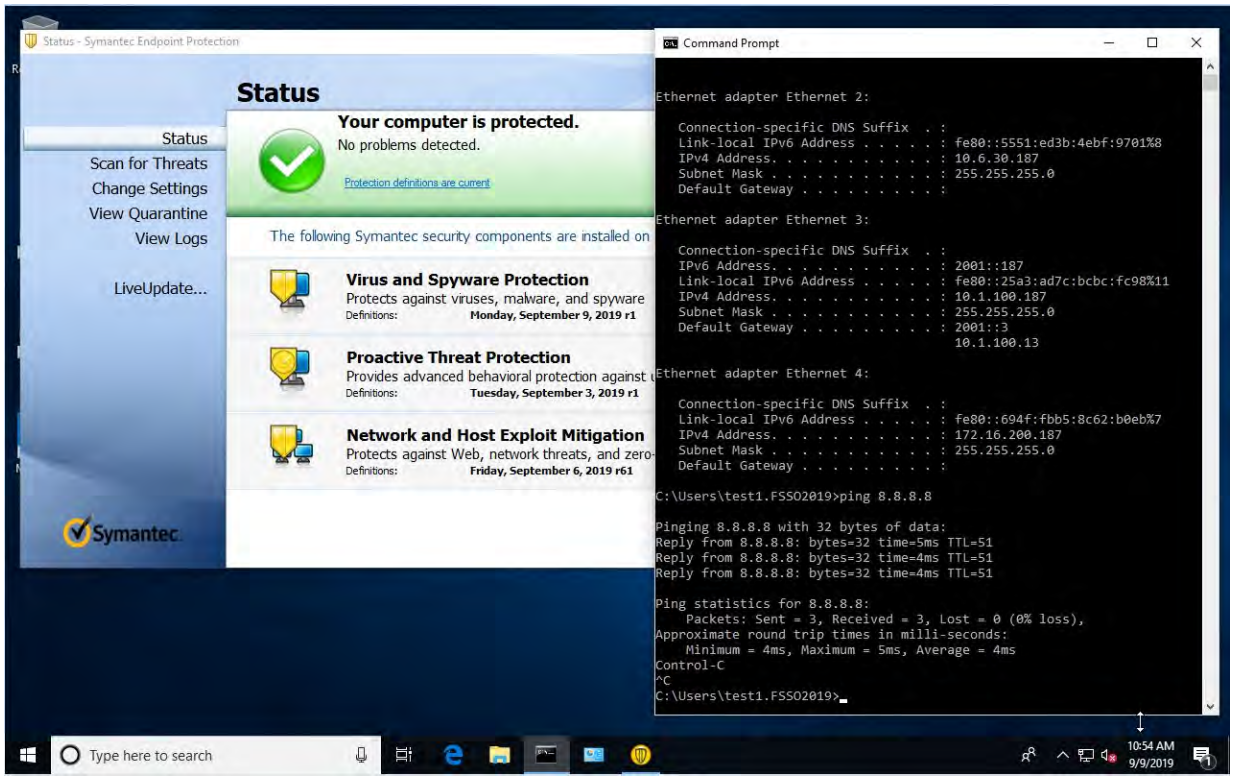


10. Go to *Policy & Objects* > *Firewall Policy*, click *Create New* and add the dynamic IP address to the firewall policy.



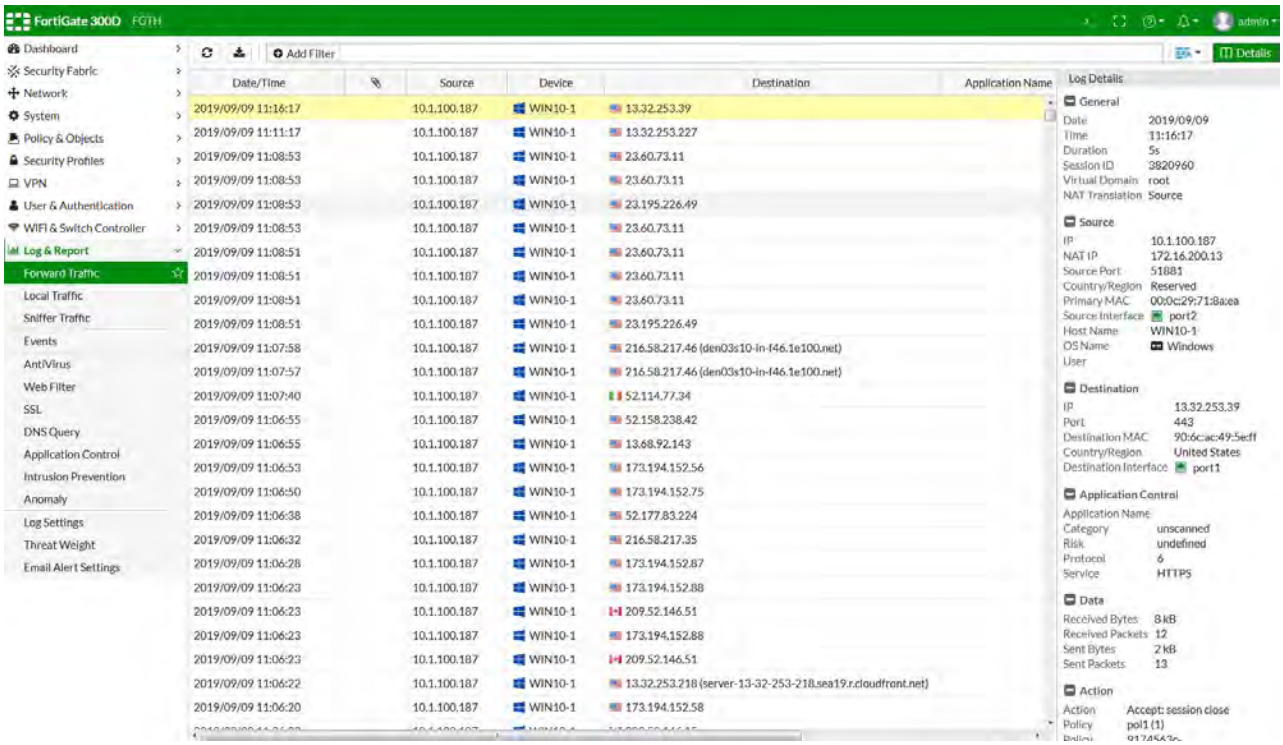
To verify the configuration:

1. On the client PC, check that it is managed by SEPM to access the Internet.



2. In FortiGate, you can check in *FortiView* > *Sources*.

3. In FortiGate, you can also check in *Log & Report* > *Forward Traffic*.





Since this traffic is not authenticated traffic but is based on source IP address only, this traffic is not shown in the GUI firewall monitor or in the CLI `diagnose firewall auth list` command.

To configure Symantec endpoint connector on FortiGate in the CLI:

1. Create the fabric connector:

```
config system sdn-connector
    edit "sepm-217"
        set type sepm
        set server "172.18.60.217"
        set username "admin"
        set password **********
        set status enable
    next
end
```

2. Create the dynamic IP address:

```
config firewall address
    edit "sepm-ip"
        set type dynamic
        set sdn "sepm-217"
        set filter "ComputerName=win10-1"
        config list
            edit "10.1.100.187"
            next
```

```
            edit "10.6.30.187"
            next
            edit "172.16.200.187"
            next
        end
    next
end
```

You can specify other filters and combine them with | and &, for example:

```
FGTH (sepm-ip) # set filter
<key1=value1>    [& <key2=value2>] [| <key3=value3>]
Available filter keys are:
        <ComputerName><ComputerUuid><GroupId><GroupName> <DomainId><DomainName><OsName>
```

3. Add the dynamic IP address to the firewall policy:

```
config firewall policy
    edit 1
        set name "pol1"
        set srcintf "port2"
        set dstintf "port1"
        set srcaddr "sepm-ip"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ALL"
        set utm-status enable
        set ssl-ssh-profile "certificate-inspection"
        set av-profile "default"
        set logtraffic all
        set fsso disable
        set nat enable
    next
end
```

To troubleshoot Symantec SD connector in the CLI:

```
# diagnose debug application sepmd -1

Output is sent every minute (default). All IPv4 learned from SEPM. IPv6 also sent but not
yet supported.

2019-09-09 12:01:09 sepmd sdn connector sepm-217 start updating IP addresses
2019-09-09 12:01:09 sepmd checking firewall address object sepm-ip, vd 0
2019-09-09 12:01:09 sepmd sdn connector sepm-217 finish updating IP addresses
2019-09-09 12:01:09 sepmd reap child pid: 18079
2019-09-09 12:02:09 sepmd sdn connector sepm-217 prepare to update
2019-09-09 12:02:09 sepmd sdn connector sepm-217 start updating
2019-09-09 12:02:09 sepm-217 sdn connector will retrieve token after 9526 secs
2019-09-09 12:02:09 sym_new_ip_addr ComputerName win10-1
    ComputerUuid AC894D56-BD86-A786-7DDB-7FD98B718AE0, OsName Windows 10
    IP 172.16.200.187
    GroupName My Company\Windows-desktops, GroupId E61FDEA2AC10C80E46D0B31BB58D7CB3
    DomainName Default, DomainId 6C507580AC10C80E5F3CAED5B1711A8E
2019-09-09 12:02:09 sym_new_ip_addr ComputerName win10-1
    ComputerUuid AC894D56-BD86-A786-7DDB-7FD98B718AE0, OsName Windows 10
    IP 10.6.30.187
```


Netskope Exhibit 1017

```
    GroupName My Company\Windows-desktops, GroupId E61FDEA2AC10C80E46D0B31BB58D7CB3
    DomainName Default, DomainId 6C507580AC10C80E5F3CAED5B1711A8E
2019-09-09 12:02:09 sym_new_ip_addr ComputerName win10-1
    ComputerUuid AC894D56-BD86-A786-7DDB-7FD98B718AE0, OsName Windows 10
    IP 10.1.100.187
    GroupName My Company\Windows-desktops, GroupId E61FDEA2AC10C80E46D0B31BB58D7CB3
    DomainName Default, DomainId 6C507580AC10C80E5F3CAED5B1711A8E
2019-09-09 12:02:09 2001:0000:0000:0000:0000:0000:0000:0187 is not in IPv4 presentation
format


2019-09-09 12:02:09 sepmd sdn connector sepm-217 start updating IP addresses
2019-09-09 12:02:09 sepmd checking firewall address object sepm-ip, vd 0
2019-09-09 12:02:09 sepmd sdn connector sepm-217 finish updating IP addresses
2019-09-09 12:02:09 sepmd reap child pid: 18089
2019-09-09 12:03:09 sepmd sdn connector sepm-217 prepare to update
2019-09-09 12:03:09 sepmd sdn connector sepm-217 start updating
2019-09-09 12:03:09 sepm-217 sdn connector will retrieve token after 9466 secs
2019-09-09 12:03:09 sym_new_ip_addr ComputerName win10-1
    ComputerUuid AC894D56-BD86-A786-7DDB-7FD98B718AE0, OsName Windows 10
    IP 172.16.200.187
    GroupName My Company\Windows-desktops, GroupId E61FDEA2AC10C80E46D0B31BB58D7CB3
    DomainName Default, DomainId 6C507580AC10C80E5F3CAED5B1711A8E
2019-09-09 12:03:09 sym_new_ip_addr ComputerName win10-1
    ComputerUuid AC894D56-BD86-A786-7DDB-7FD98B718AE0, OsName Windows 10
    IP 10.6.30.187
    GroupName My Company\Windows-desktops, GroupId E61FDEA2AC10C80E46D0B31BB58D7CB3
    DomainName Default, DomainId 6C507580AC10C80E5F3CAED5B1711A8E
2019-09-09 12:03:09 sym_new_ip_addr ComputerName win10-1
    ComputerUuid AC894D56-BD86-A786-7DDB-7FD98B718AE0, OsName Windows 10
    IP 10.1.100.187
    GroupName My Company\Windows-desktops, GroupId E61FDEA2AC10C80E46D0B31BB58D7CB3
    DomainName Default, DomainId 6C507580AC10C80E5F3CAED5B1711A8E
2019-09-09 12:03:09 2001:0000:0000:0000:0000:0000:0000:0187 is not in IPv4 presentation
format
```

To list the SEPM daemon SDN connectors:

```
diagnose test application sepmd 1
sepm SDN connector list:
  name: sepm-217, status: enabled, updater_interval: 60
```

To list the SEPM daemon SDN filters:

```
diagnose test application sepmd 2
sepm SDN connector sepm-217 filter list:
  name: sepm-ip, vd 0, filter 'ComputerName=win10-1'
```

RADIUS single sign-on (RSSO) agent

With RSSO, a FortiGate can authenticate users who have authenticated on a remote RADIUS server. Based on which user group the user belongs to, the security policy applies the appropriate UTM profiles.

The FortiGate does not interact with the remote RADIUS server; it only monitors RADIUS accounting records that the server forwards (originating from the RADIUS client). These records include the user IP address and user group. The remote RADIUS server sends the following accounting messages to the FortiGate:

Message	Action
Start	If the information in the start message matches the RSSO configuration on the FortiGate, the user is added to the local list of authenticated firewall users.
Stop	The user is removed from the local list of authenticated firewall users because the user session no longer exists on the RADIUS server.

You can configure an RSSO agent connector using the FortiOSGUI; however, in most cases, you will need to use the CLI. There are some default options you may need to modify, which can only be done in the CLI.

To configure an RSSO agent connector:

1. Create the new connector:
   a. Go to *Security Fabric* > *External Connectors*.
   b. Click *Create New*.
   c. In the *Endpoint/Identity* section, click *RADIUS Single Sign-On Agent*. The *New Fabric Connector* pane opens.
   d. Enter the connector name.
   e. Enable *Use RADIUS Shared Secret*.



The value entered in *Use RADIUS Shared Secret* must be identical to what the remote RADIUS server uses to authenticate when it sends RADIUS accounting messages to the FortiGate.

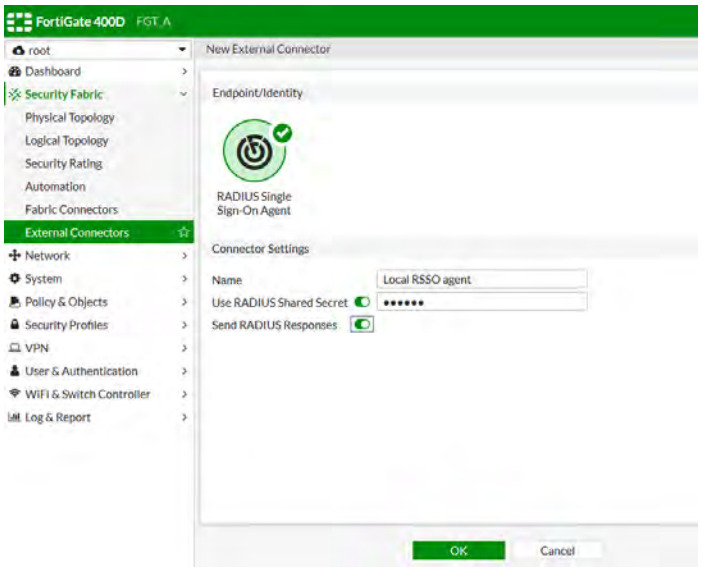
   f. Enable *Send RADIUS Responses*.



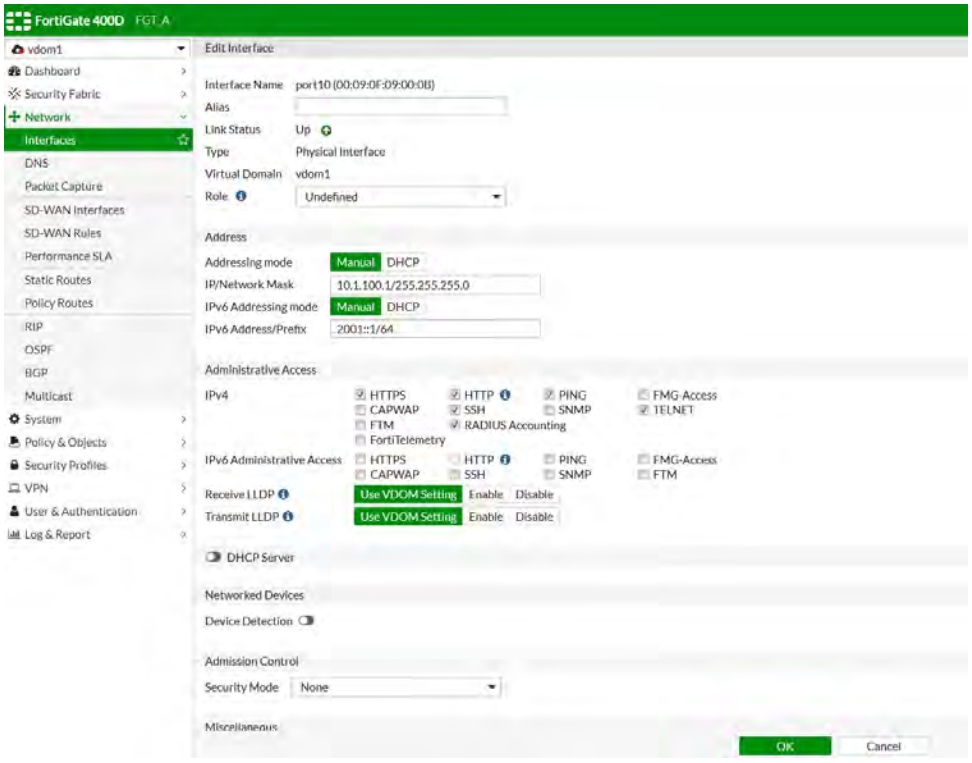
You should enable *Send RADIUS Responses* because some RADIUS servers continue to send the same RADIUS accounting message several times if there is no response.

Netskope Exhibit 1017

g. Click *OK*.



2. Edit the network interface:
   a. Go to *Network* > *Interfaces*.
   b. Double-click the interface that will receive the RADIUS accounting messages. The *Edit Interface* pane opens.
   c. In the *Administrative Access* section, select the *RADIUS Accounting* checkbox. This will open listening for port 1813 on this interface. The FortiGate will then be ready to receive RADIUS accounting messages.
   d. Click *OK*.



3. Create a local RSSO user group:
   a. Go to *User & Authentication* > *User Groups*.
   b. Click *Create New*.

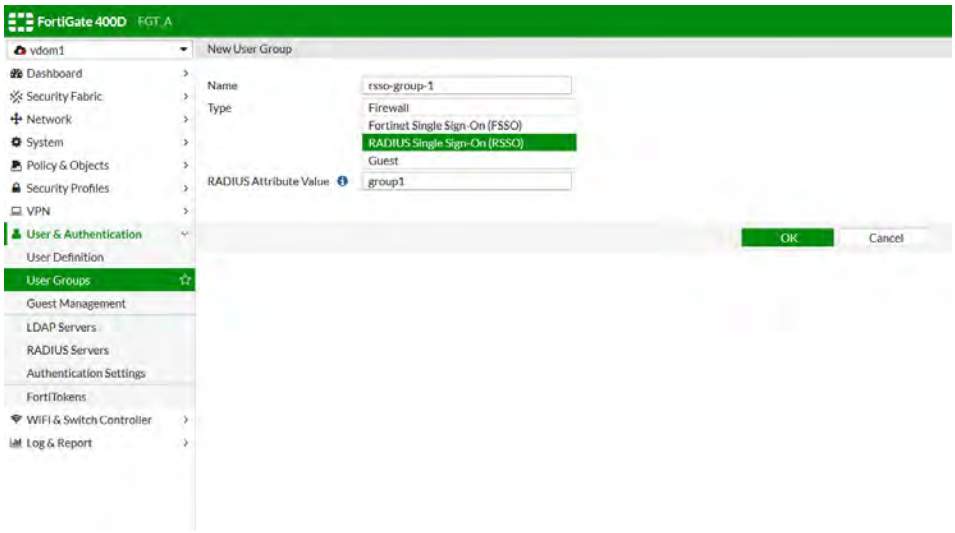
Netskope Exhibit 1017

c. Enter the group name.
d. For the *Type* field, click *RADIUS Single-Sign-ON (RSSO)*.
e. Enter a value for *RADIUS Attribute Value*.
This value by default is the class attribute. The FortiGate uses the content of this attribute in RADIUS accounting start messages to map a user to a FortiGate group, which then can be used in firewall policies.
In this example configuration, the FortiGate will only add a remote RADIUS user to the local firewall user list if the class attribute in the RADIUS accounting START message contains the value group1.

If your users are in multiple groups, you will need to add another local RSSO user group.

If the RADIUS attribute value used to map users to a local RSSO group is different than the RADIUS attribute in the RADIUS accounting messages forwarded by the server, you must change it in the CLI.

f. Click *OK*.



4. Edit the local RSSO agent to modify default options using the CLI.
For example, the default value for `rsso-endpoint-attribute` might work in common remote access scenarios where users are identified by their unique `Calling-Station-Id`, but in other scenarios the user name might be in a different attribute.

```
config user radius
    edit "Local RSSO Agent"
        set rsso-endpoint-attribute <attribute>
        set sso-attribute <attribute>
    next
end
```

5. Add the local RSSO user group to a firewall policy.

Verifying the RSSO configuration

Verification requires a working remote RADIUS server configured for RADIUS accounting forwarding and wireless or wired clients that use RADIUS for user authentication.

For a quick test, you can use one of the publicly available RADIUS test tools to send RADIUS accounting start and stop messages to the FortiGate. You can also use radclient.

To verify the RSSO configuration:

1. In radclient, enter the RADIUS attributes. These attributes are then executed with the FortiGate IP parameters (sends accounting messages to port 1813) and shared password you configured. `-x` is used for verbose output:

```
root@ControlPC:~# echo "Acct-Status-Type =Start,Framed-Ip-Address=10.1.100.185,User-
Name=test2,Acct-Session-Id=0211a4ef,Class=group1,Calling-Station-Id=00-0c-29-44-BE-B8" |
 radclient -x 10.1.100.1 acct 123456
Sending Accounting-Request of id 180 to 10.1.100.1 port 1813
        Acct-Status-Type = Start
        Framed-IP-Address = 10.1.100.185
        User-Name = "test2"
        Acct-Session-Id = "0211a4ef"
        Class = 0x67726f757031
        Calling-Station-Id = "00-0c-29-44-BE-B8"
rad_recv: Accounting-Response packet from host 10.1.100.1 port 1813, id=180, length=20
root@ControlPC:~#
```

2. Verify that the user is in the local firewall user list with the correct type (`rsso`) and local firewall group (`rsso-group1`):

```
# diagnose firewall auth l

10.1.100.185, test2
        type: rsso, id: 0, duration: 5, idled: 5
        flag(10): radius
        server: vdom1
        packets: in 0 out 0, bytes: in 0 out 0
        group_id: 3
        group_name: rsso-group-1

----- 1 listed, 0 filtered ------
```

Threat feeds

Threat feeds dynamically import an external block lists from an HTTP server in the form of a plain text file. Block lists can be used to enforce special security requirements, such as long term policies to always block access to certain websites, or short term requirements to block access to known compromised locations. The lists are dynamically imported, so that any changes are immediately imported by FortiOS.

There are four types of threat feeds:

FortiGuard Category	The file contains one URL per line. It is available as a *Remote Category* in Web Filter profiles and SSL inspection exemptions. Example:

	`http://example/com.url`
	`https://example.com/url`
	`http://example.com:8080/url`

IP Address	The file contains one IP/IP range/subnet per line. It is available as an *External IP Block List* in DNS Filter profiles, and as a *Source/Destination* in IPv4, IPv6, and proxy policies.
	Example:
	`192.168.2.100`
	`172.200.1.4/16`
	`172.16.1.2/24`
	`172.16.8.1-172.16.8.100`
	`2001:0db8::eade:27ff:fe04:9a01/120`
	`2001:0db8::eade:27ff:fe04:aa01-2001:0db8::eade:27ff:fe04:ab01`

Domain Name	The file contains one domain per line. Simple wildcards are supported. It is available as a *Remote Category* in DNS Filter profiles. See External resources for DNS filter on page 279 for more information.
	Example:
	`mail.*.example.com`
	`*-special.example.com`
	`www.*example.com`
	`example.com`

Malware Hash	The file contains one hash per line in the format `<hex hash> [optional hash description]`. Each line supports MD5, SHA1, and SHA256 hex hashes. It is automatically used for virus outbreak prevention on antivirus profiles with `external-blocklist` enabled.
	Note: For optimal performance, do not mix different hashes in the list. Only use one of MD5, SHA1, or SHA26.
	Example:
	`292b2e6bb027cd4ff4d24e338f5c48de`
	`dda37961870ce079defbf185eeeef905 Trojan-Ransom.Win32.Locky.abfl`
	`3fa86717650a17d075d856a41b3874265f8e9eab Trojan-Ransom.Win32.Locky.abfl`
	`c35f705df9e475305c0984b05991d444450809c35dd1d96106bb8e7128b9082f Trojan-Ransom.Win32.Locky.abfl`
	See External malware block list for antivirus on page 864 for an example.

External resources file format

File format requirements for an external resources file:

- The file is in plain text format with each URL list, IP address, and domain name occupying one line.
- The file is limited to 10 MB or 128 × 1024 (131072) entries, whichever limit is hit first.
- The entry limit also follows the table size limitation defined by CMDB per model.
- The external resources update period can be set to 1 minute, hourly, daily, weekly, or monthly (43200 min, 30 days).
- The external resources type as category (URL list) and domain (domain name list) share the category number range 192 to 221 (total of 30 categories).
- There is no duplicated entry validation for the external resources file (entry inside each file or inside different files).
- If the number of entries exceed the limit, a warning is displayed. Additional entries beyond the threshold will not be loaded.

For domain name list (type = domain):

- Simple wildcards are allowed in the domain name list, for example: *.test.com.
- IDN (international domain name) is supported.

For IP address list (type = address):

- The IP address can be a single IP address, subnet address, or address range. For example, 192.168.1.1, 192.168.10.0/24, or 192.168.100.1-192.168.100.254.
- The address can be an IPv4 or IPv6 address. An IPv6 address does not need to be in [ ] format.

To create a threat feed in the GUI:

1. Go to *Security Fabric* > *External Connectors*.
2. Click *Create New*.
3. In the *Thread Feeds* section, click on the required feed type.
4. Configure the connector settings:

Name	Enter a name for the threat feed connector.
URI of external resource	Enter the link to the external resource file. The file should be a plain text file with one entry on each line.
HTTP basic authentication	Enable/disable basic HTTP authentication. When enabled, enter the username and password in the requisite fields.
Refresh Rate	The time interval to refresh the external resource, in minutes (1 - 43200, default = 5).
Comments	Optionally, enter a description of the connector.
Status	Enable/disable the connector.

5. Click *OK*.

To create a threat feed in the CLI:

```
config system external-resource
    edit <name>
        set status {enable | disable}
        set type {category | address | domain | malware}
        set category <integer>
        set username <string>
        set password <string>
        set comments [comments]
        *set resource <resource-uri>
        *set refresh-rate <integer>
        set source-ip <string>
    next
end
```

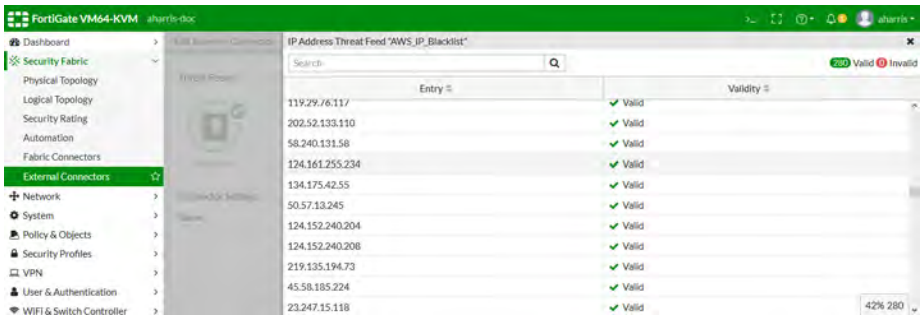
Parameters marked with a * are mandatory and must be filled in. Other parameters either have default values or are optional.

Netskope Exhibit 1017

Update history

To review the update history of a threat feed, go to *Security Fabric* > *External Connectors*, select a feed, and click *Edit*. The *Last Update* field shows the date and time that the feed was last updated.

Click *View Entries* to view the current entries in the list.

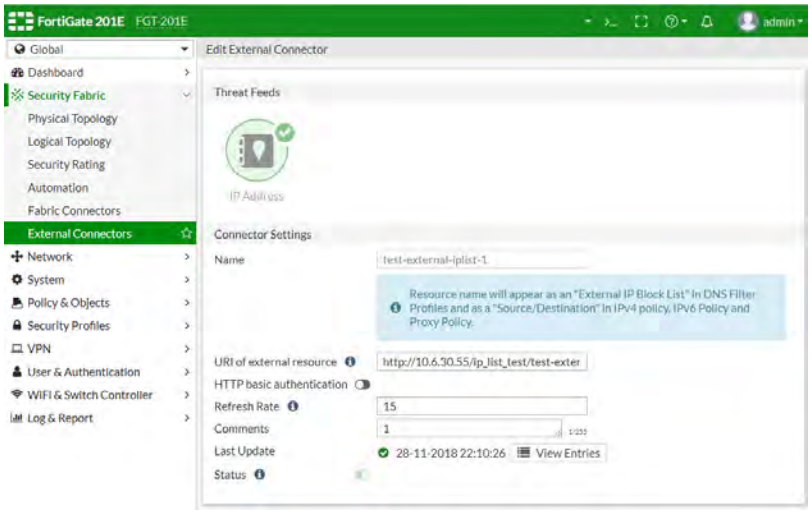


External Block List (Threat Feed) – Policy

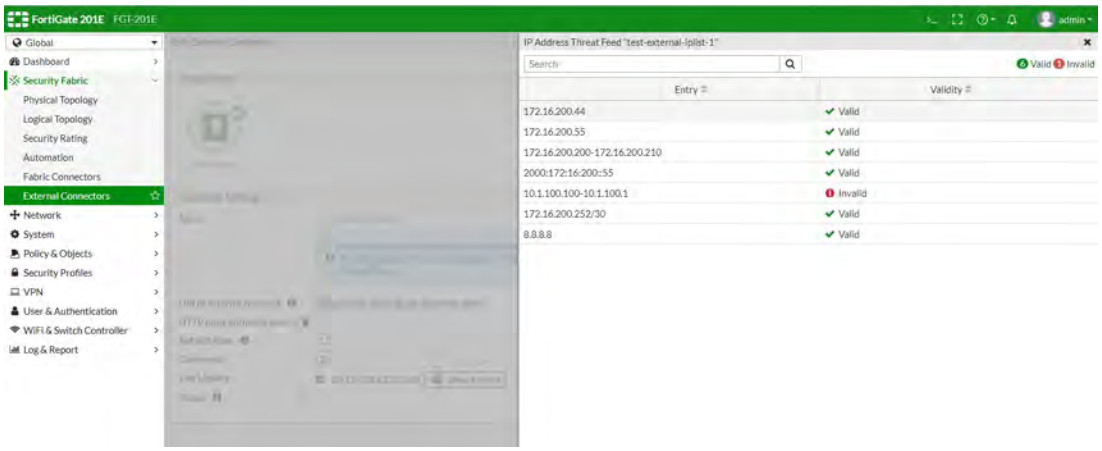
You can use the External Block List (Threat Feed) for web filtering and DNS. You can also use External Block List (Threat Feed) in firewall policies.

Sample configuration

In *Security Fabric* > *External Connectors* > *Threat Feeds* > *IP Address*, create or edit an external IP list object.



Click *View Entries* to see the external IP list.



To create an external iplist object using the CLI:

```
config system external-resource
   edit "test-external-iplist-1"
      set status enable
      set type address
      set username ''
      set password ENC
      set comments ''
      set resource "http://10.6.30.55/ip_list_test/test-external-iplist-2.txt"
      set refresh-rate 15
   next
end
```

To apply an external iplist object to the firewall policy using the CLI:

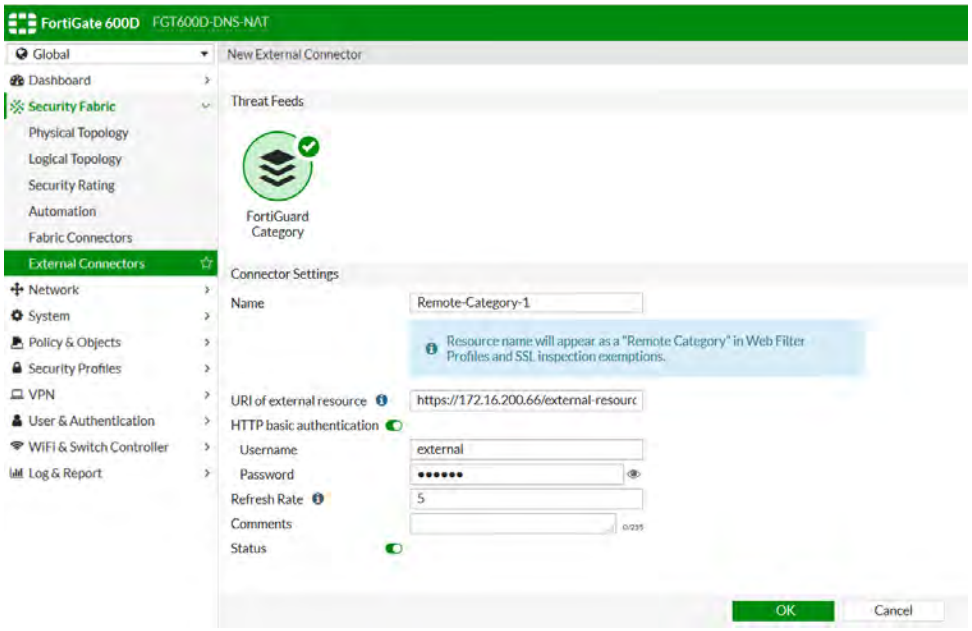
```
config firewall policy
   edit 1
      set name "policyid-1"
      set srcintf "wan2"
      set dstintf "wan1"
      set srcaddr "all"
      set dstaddr "test-external-iplist-1"
      set action accept
      set schedule "always"
      set service "ALL"
      set logtraffic all
      set auto-asic-offload disable
      set nat enable
   next
end
```

External Block List (Threat Feed) - Authentication

The external *Threat Feed* connector (block list retrieved by HTTPS) supports username and password authentication.

To enable username and password authentication:

1. Navigate to *Security Fabric* > *External Connectors*.
2. Edit an existing *Threat Feed* or create a new one by selecting *Create New*.
3. In *Connector Settings*, select the *HTTP basic authentication* toggle to enable the feature.
4. Enter a username and password.



5. Select *OK* to save your changes.

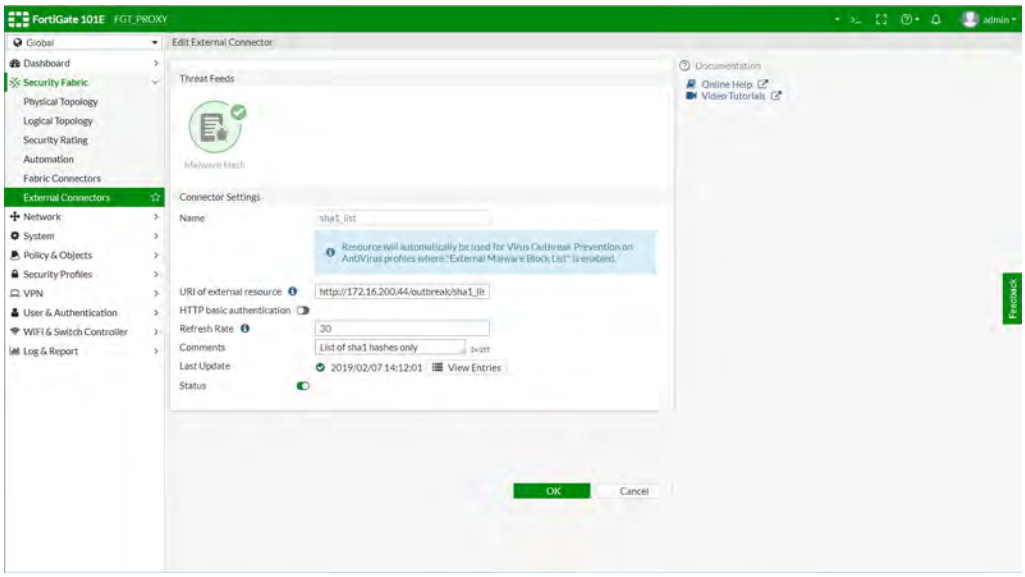
External Block List (Threat Feed) - File Hashes

The Malware Hash type of *Threat Feed* connector supports a list of file hashes that can be used as part of virus outbreak prevention.

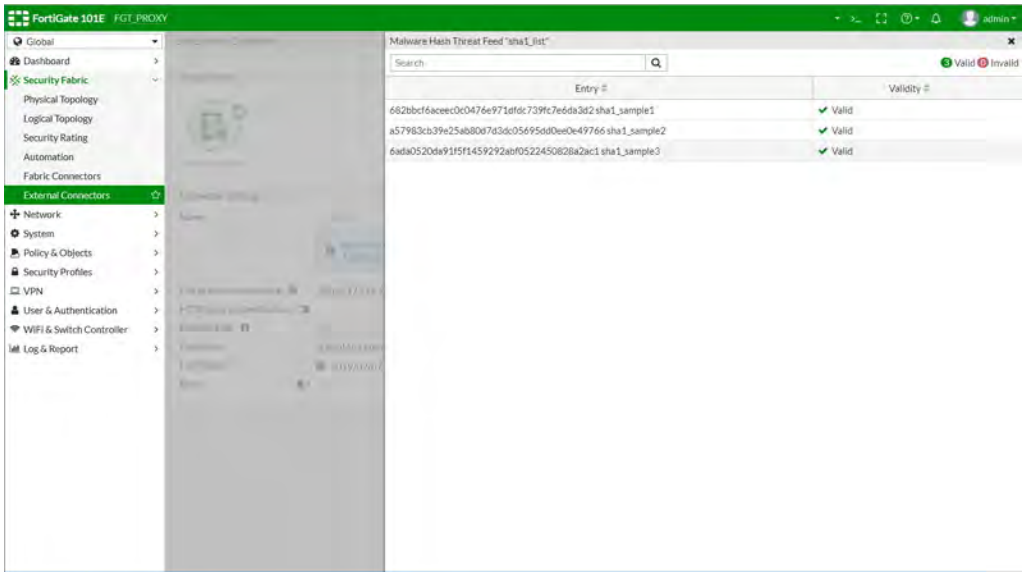
To configure Malware Hash:

1. Navigate to *Security Fabric* > *External Connectors* and click *Create New*.
2. In the *Threat Feeds* section, click *Malware Hash*.

3. To configure *Malware Hash*, fill in the *Connector Settings* section.



4. Beside the *Last Update* field, click *View Entries* to display the external Malware Hash list contents.



New Malware value for external-resource parameter in CLI

```
FGT_PROXY (external-resource) # edit sha1_list
new entry 'sha1_list' added

FGT_PROXY (sha1_list) # set type ?
category    FortiGuard category.
address     Firewall IP address.
domain      Domain Name.
malware     Malware hash.
```

To configure external Malware Hash list sources in CLI:

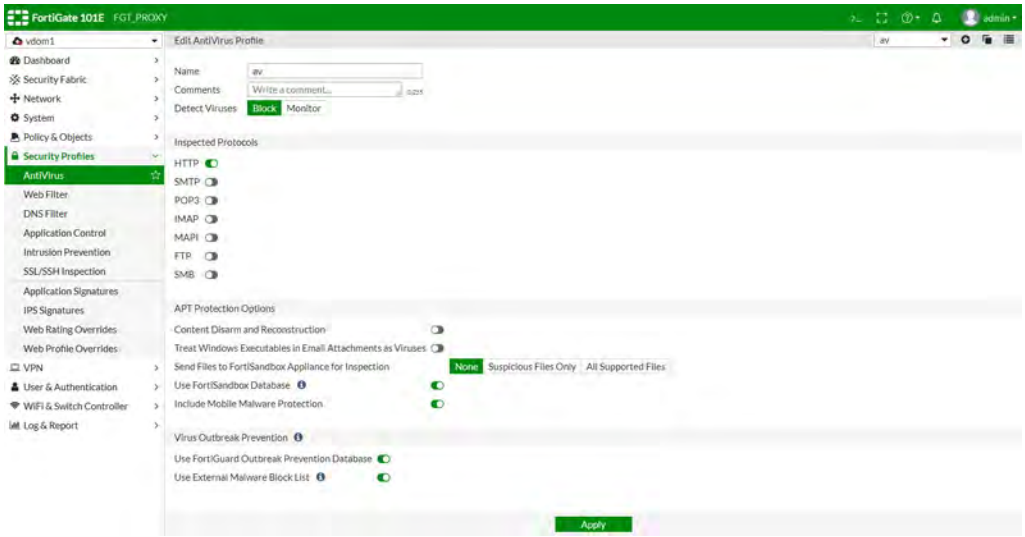
```
config global
    config system external-resource
        edit "md5_list"
            set type malware
            set comments "List of md5 hashes only"
            set resource "http://172.16.200.44/outbreak/md5_list"
            set refresh-rate 30
        next
        edit "sha1_list"
            set type malware
            set comments "List of sha1 hashes only"
            set resource "http://172.16.200.44/outbreak/sha1_list"
            set refresh-rate 30
        next
        edit "sha256_list"
            set type malware
            set comments "List of sha256 hashes only"
            set resource "http://172.16.200.44/outbreak/sha256_list"
            set refresh-rate 30
        next
    end

end
```

Update to antivirus profile

In *Security Profiles* > *AntiVirus*, the *Virus Outbreak Prevention* section allows you to enable the following options:

- Use Fortinet outbreak Prevention Database.
- Use External Malware Block List.



To configure virus outbreak prevention options in CLI:

You must first enable `outbreak-prevention` in the protocol and then enable `external-blocklist` under `outbreak-prevention`.

```
config antivirus profile
    edit "av"
        set analytics-db enable
        config http
            set options scan
            set outbreak-prevention full-archive
        end
        config ftp
            set options scan
            set outbreak-prevention files
        end
        config imap
            set options scan
            set outbreak-prevention full-archive
        end
        config pop3
            set options scan
            set outbreak-prevention full-archive
        end
        config smtp
            set options scan
            set outbreak-prevention files
        end
        config mapi
            set options scan
            set outbreak-prevention full-archive
        end
        config nntp
            set options scan
            set outbreak-prevention full-archive
        end
        config smb
            set options scan
            set outbreak-prevention full-archive
        end
        config outbreak-prevention
            set ftgd-service enable
            set external-blocklist enable
        end
    next
end
```

Update to utm-virus category logs

This feature adds the fields `filehash` and `filehashsrc` to outbreak prevention detection events.

Example of the utm-virus log generated when a file is detected by FortiGuard queried outbreak prevention:

```
2: date=2018-07-30 time=13:57:59 logid="0204008202" type="utm" subtype="virus"
eventtype="outbreak-prevention" level="warning" vd="root" evnttime=1532984279 msg="Blocked
by Virus Outbreak Prevention service." action="blocked" service="HTTP" sessionid=174777
srcip=192.168.101.20 dstip=172.16.67.148 srcport=37044 dstport=80 srcintf="lan"
srcintfrole="lan" dstintf="wan1" dstintfrole="wan" policyid=1 proto=6 direction="incoming"
filename="zhvo_test.com" checksum="583369a5" quarskip="No-skip"
virus="503e99fe40ee120c45bc9a30835e7256fff3e46a" dtype="File Hash"
filehash="503e99fe40ee120c45bc9a30835e7256fff3e46a"    filehashsrc="fortiguard"
```

```
url="http://172.16.67.148/zhvo_test.com" profile="mhash_test" agent="Firefox/43.0"
analyticssubmit="false" crscore=30 crlevel="high&#8220;
```

Example of the utm-virus log generated when a file is detected by External Malware Hash List outbreak prevention:

```
1: date=2018-07-30 time=13:59:41 logid="0207008212" type="utm" subtype="virus"
eventtype="malware-list" level="warning" vd="root" eventtime=1532984381 msg="Blocked by
local malware list." action="blocked" service="HTTP" sessionid=174963 srcip=192.168.101.20
dstip=172.16.67.148 srcport=37045 dstport=80 srcintf="lan" srcintfrole="lan" dstintf="wan1"
dstintfrole="wan" policyid=1 proto=6 direction="incoming" filename="mhash_block.com"
checksum="90f0cb57" quarskip="No-skip" virus="mhash_block.com" dtype="File Hash"
filehash="93bdd30bd381b018b9d1b89e8e6d8753"    filehashsrc="test_list"
url="http://172.16.67.148/mhash_block.com" profile="mhash_test" agent="Firefox/43.0"
analyticssubmit="false"
```

External resources for DNS filter

External resources provides the ability to dynamically import an external block list into an HTTP server. This feature enables the FortiGate to retrieve a dynamic URL, domain name, IP address, or malware hash list from an external HTTP server periodically. The FortiGate uses these external resources as the web filter's remote categories, DNS filter's remote categories, policy address objects, or antivirus profile's malware definitions. If external resources are updated, FortiGate objects are also updated dynamically.

External resource is divided into four types:

- URL list (type = category)
- Domain name list (type = domain)
- IP address list (type = address)
- Malware hash list (type = malware)

Remote categories and external IP block list

The DNS filter profile can use two types of external resources: *domain type* (domain name list) and *address type* (IP address list).

When a *domain type* external resource is configured, it is treated as a remote category in the DNS filter profile. If the domain name in DNS query matches the entry in this external resource file, it is treated as the remote category and follows the action configured for this category in DNS filter profile.

When an *address type* external resource is configured, it can be enabled as *external-ip-blocklist* in DNS filter profile. If a DNS resolved IP address in DNS response matches the entry in the *external-ip-blocklist*, this DNS query is blocked by the DNS filter.

For external resources file format and limits, see External resources file format on page 271.

Configuring external resources in the CLI

In the CLI, you can configure external resources files in an external HTTP server. Under global, configure the external resources file location and specify the resource type.

To configure external resources:

```
config system external-resource
   edit "Ext-Resource-Type-as-Domain-1"
```


Netskope Exhibit 1017

```
      set type domain
      set category 194
      set resource "http://172.16.200.66/external-resources/Ext-Resource-Type-as-Domain-
1.txt"
      set refresh-rate 1
   next
   edit "Ext-Resource-Type-as-Address-1"
      set status enable
      set type address
      set username ''
      set password
      set comments ''
      set resource "http://172.16.200.66/external-resources/Ext-Resource-Type-as-Address-
1.txt"
      set refresh-rate 1
   next
end
```

In each VDOM, the domain type external resource can be used in the DNS filter as remote category. In this example, the domain name list in the Ext-Resource-Type-as-Domain-1.txt file is treated as a remote category (category ID 194). The IP address list in the Ext-Resource-Type-as-Address-1.txt file can be applied in the DNS filter as an `external-ip-blocklist`. If the DNS resolved IP address matches any entry in the list in that file, the DNS query is blocked.

To configure the external IP block list and apply it to a policy:

```
config dnsfilter profile
   edit "default"
      set comment "Default dns filtering."
      config ftgd-dns
         config filters
            edit 1
               set category 194
               set action block
            next
            edit 2
               set category 12
            next
            edit 3
            next
         end
      end
      set block-botnet enable
      set external-ip-blocklist "Ext-Resource-Type-as-Address-1"
   next
end
```

```
config firewall policy
   edit 1
      set name "DNSFilter"
      set srcintf "port10"
      set dstintf "port9"
      set srcaddr "all"
      set dstaddr "all"
      set action accept
      set schedule "always"
      set service "ALL"
```

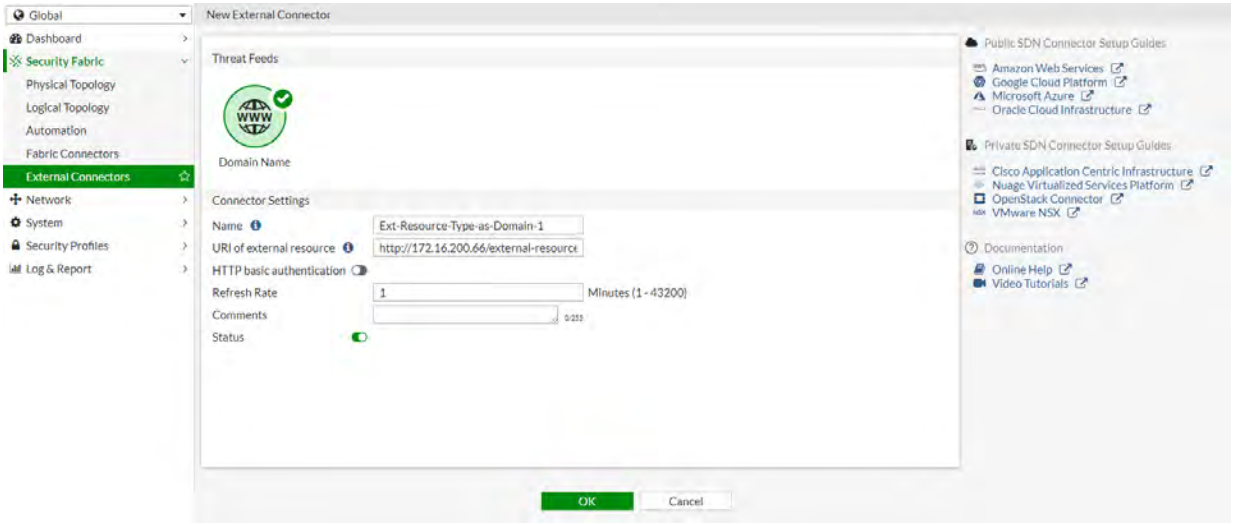

Netskope Exhibit 1017

```
        set utm-status enable
        set logtraffic all
        set dnsfilter-profile "default"
        set profile-protocol-options "protocol"
        set ssl-ssh-profile "protocols"
        set nat enable
    next
end
```

Configuring external resources in the GUI

To configure, edit, or view the entries for external resources in the GUI:

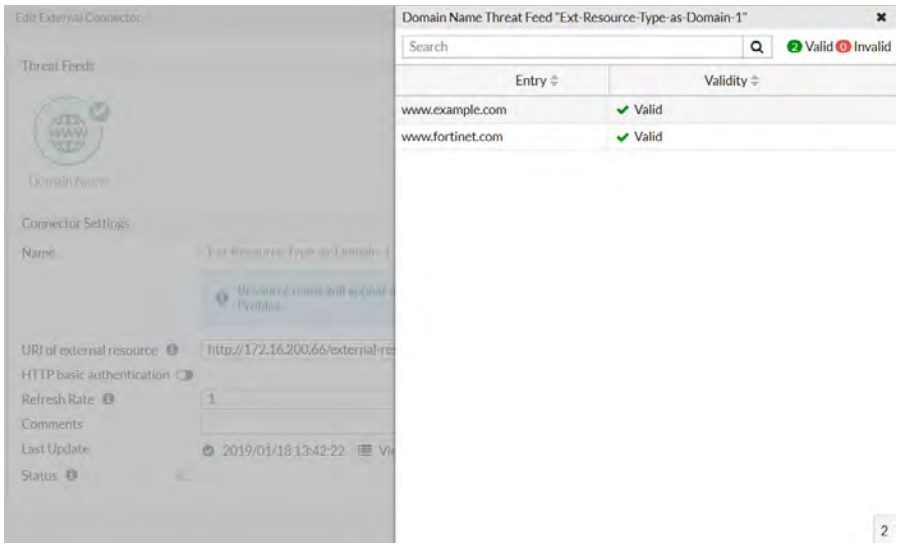
1. Go to *Global* > *Security Fabric* > *External Connectors*.
2. Click *Create New*.
3. In the *Threat Feeds* section, select *Domain Name* or *IP Address*.
4. Enter the Resource *Name*, URL, location of the resource file, resource authentication credentials, and *Refresh Rate*.



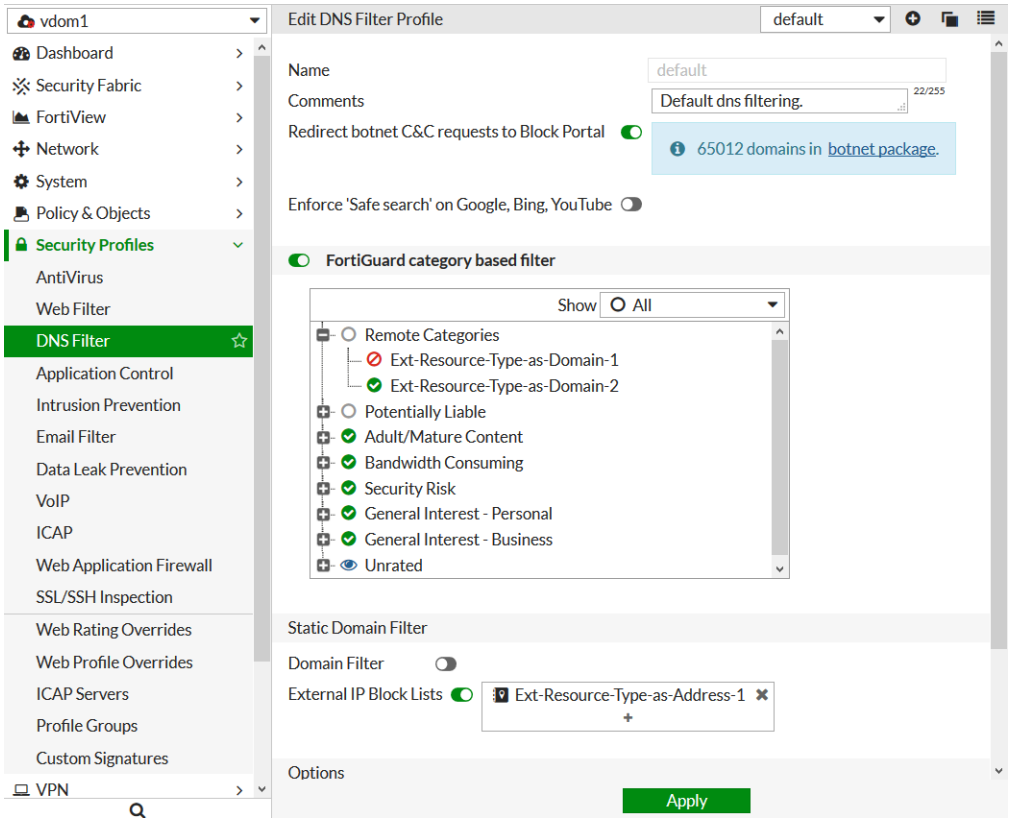
5. Click *OK*.
6. Double-click the *Threat Feeds Object* you just configured to open the *Edit* page

Netskope Exhibit 1017

7. Click *View Entries* to view the entry list in the external resources file.



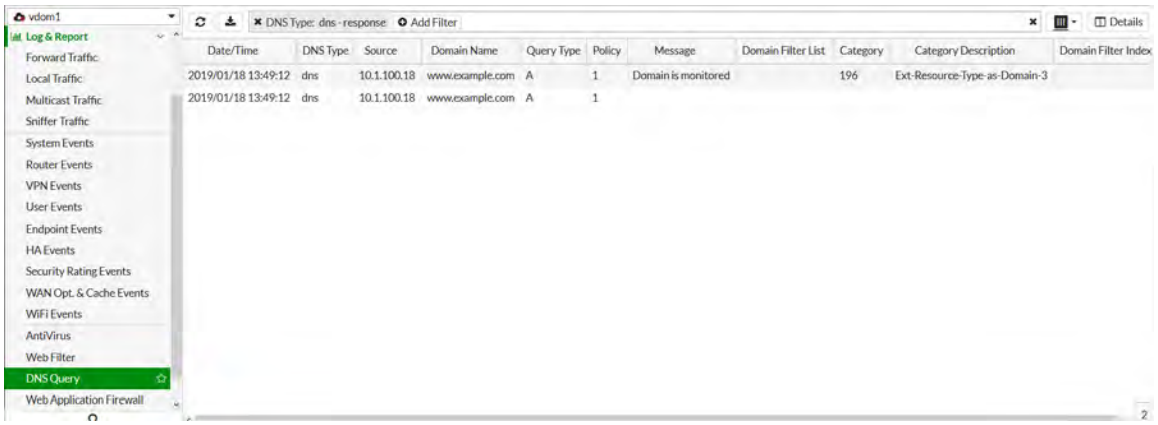
8. Go to *VDOM* > *Security Profiles* > *DNS Filter* and open a DNS filter profile. The configured external resources displays, and you can apply it in each DNS filter profile (remote category or external IP block lists).



Log sample

Remote categories

Go to *VDOM* > *Log & Report* > *DNS Query*. Some domains that match the remote category list are rated as remote category, overriding their original domain rating.



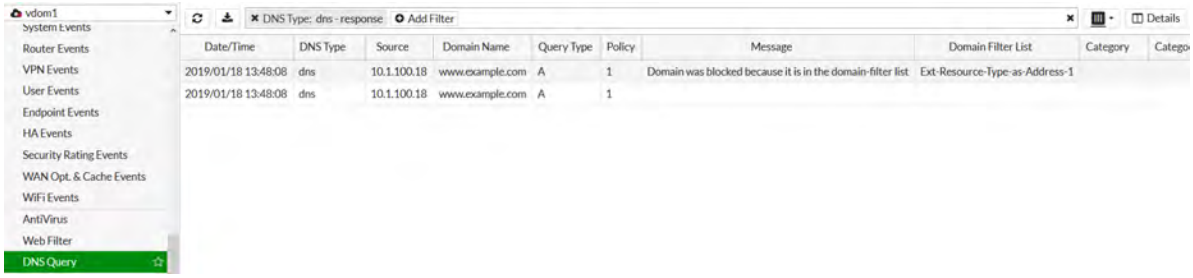
Log example:

```
1: date=2019-01-18 time=13:49:12 logid="1501054802" type="utm" subtype="dns" eventtype="dns-
response" level="notice" vd="vdom1" eventtime=1547848151 policyid=1 sessionid=82998
srcip=10.1.100.18 srcport=42985 srcintf="port10" srcintfrole="undefined" dstip=172.16.95.16
dstport=53 dstintf="port9" dstintfrole="undefined" proto=17 profile="default" xid=38234
qname="www.example.com" qtype="A" qtypeval=1 qclass="IN" ipaddr="93.184.216.34" msg="Domain
is monitored" action="pass" cat=196 catdesc="Ext-Resource-Type-as-Domain-3"

2: date=2019-01-18 time=13:49:12 logid="1500054000" type="utm" subtype="dns" eventtype="dns-
query" level="information" vd="vdom1" eventtime=1547848151 policyid=1 sessionid=82998
srcip=10.1.100.18 srcport=42985 srcintf="port10" srcintfrole="undefined" dstip=172.16.95.16
dstport=53 dstintf="port9" dstintfrole="undefined" proto=17 profile="default" xid=38234
qname="www.example.com" qtype="A" qtypeval=1 qclass="IN"
```

External IP block lists

Go to *VDOM* > *Log & Report* > *DNS Query*. If the DNS query resolved IP address matches the entry in the `external-ip-blocklist`, the DNS query is blocked.



Log example:

```
1: date=2019-01-18 time=13:50:53 logid="1501054400" type="utm" subtype="dns" eventtype="dns-response" level="warning" vd="vdom1" eventtime=1547848253 policyid=1 sessionid=83206 srcip=10.1.100.18 srcport=47281 srcintf="port10" srcintfrole="undefined" dstip=172.16.95.16 dstport=53 dstintf="port9" dstintfrole="undefined" proto=17 profile="default" xid=7501 qname="www.example.com" qtype="A" qtypeval=1 qclass="IN" msg="Domain was blocked because it is in the domain-filter list" action="redirect" domainfilteridx=0 domainfilterlist="Ext-Resource-Type-as-Address-1"

2: date=2019-01-18 time=13:50:53 logid="1500054000" type="utm" subtype="dns" eventtype="dns-query" level="information" vd="vdom1" eventtime=1547848253 policyid=1 sessionid=83206 srcip=10.1.100.18 srcport=47281 srcintf="port10" srcintfrole="undefined" dstip=172.16.95.16 dstport=53 dstintf="port9" dstintfrole="undefined" proto=17 profile="default" xid=7501 qname="www.example.com" qtype="A" qtypeval=1 qclass="IN"
```

Automation stitches

Automation stitches automate the activities between the different components in the Security Fabric, decreasing the response times to security events. Events from any source in the Security Fabric can be monitored, and action responses can be set up to any destination.



Automation stitches can also be used on FortiGate devices that are not part of a Security Fabric.

Automation stitches that use cloud-based actions, such as AWS Lambda and Azure Function, have the option to delay an action after the previous action is completed.

An automation stitch consists of two parts, the trigger and the actions. The trigger is the condition or event on the FortiGate that activates the action, for example, a specific log, or a failed log in attempt. The action is what the FortiGate does in response to the trigger.

Diagnose commands are available in the CLI to test, log, and display the history and settings of stitches.



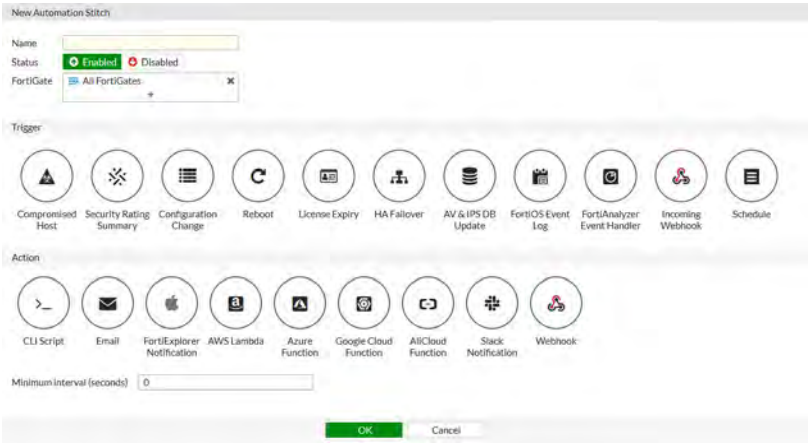
Automation stitches can only be created on the root FortiGate in a Security Fabric.

Creating automation stitches

To create an automation stitch, a trigger event and a response action or actions are selected. Automation stitches can also be tested after they are created.

To create an automation stitch in the GUI:

1. On the root FortiGate, go to *Security Fabric* > *Automation*.
2. Click *Create New*. The *New Automation Stitch* page opens.



3. Enter the following information:

Name	Enter a name for the automation stitch.
Status	Enable/disable the stitch.
FortiGate	Select the FortiGate device to apply the automation stitch to, or select *All FortiGates* to apply it to all of them.
Trigger	Select a trigger.
Action	Select and configure one or more actions.
Minimum interval (seconds)	Enter a minimum time interval during which notifications for the same trigger event will not be sent. After the time interval elapses, an alert is sent that includes the last event since the time interval elapsed.

4. Click *OK*.

To create an automation stitch in the CLI:

1. Create an automation trigger:

```
config system automation-trigger
    edit <automation-trigger-name>
        set trigger-type {event-based | scheduled}
        set event-type <option>
        set license-type <option>
        set ioc-level {medium | high}
        set logid <integer>
        set trigger-frequency {hourly | daily | weekly | monthly}
        set trigger-weekday <option>
        set trigger-day <integer>
        set trigger-hour <integer>
        set trigger-minute <integer>
        set faz-event-severity <string>
```

```
        set faz-event-tags <string>
    next
end
```

The available options will vary depending on the selected event type.

2. Create an automation action:

```
config system automation-action
    edit <name>
        set action-type <option>
        set email-to <names>
        set email-from <string>
        set email-subject <string>
        set message <string>
        set minimum-interval <integer>
        set delay <integer>
        set required {enable | disable}
        set aws-api-id <string>
        set aws-region <string>
        set aws-domain <string>
        set aws-api-stage <string>
        set aws-api-path <string>
        set aws-api-key <string>
        set azure-app <string>
        set azure-function <string>
        set azure-domain <string>
        set azure-function-authorization {anonymous | function | admin}
        set azure-api-key <string>
        set gcp-function-region <string>
        set gcp-project <string>
        set gcp-function-domain <string>
        set gcp-function <string>
        set alicloud-account-id <string>
        set alicloud-region <string>
        set alicloud-function-domain <string>
        set alicloud-version <string>
        set alicloud-service <string>
        set alicloud-function <string>
        set alicloud-function-authorization {anonymous | function}
        set alicloud-access-key-id <string>
        set alicloud-access-key-secret <string>
        set protocol {http | https}
        set method {post | put | get | patch | delete}
        set uri <string>
        set http-body <string>
        set port <integer>
        set headers <header>
        set script <string>
        set security-tag <string>
        set sdn-connector <connector_name>
    next
end
```



Enter `edit 0` to get the next available ID.

Netskope Exhibit 1017

3. Create an automation destination:

```
config system automation-destination
    edit <name>
        set type {fortigate | ha-cluster}
        set destination <serial numbers>
        set ha-group-id <integer>
    next
end
```

4. Create the automation stitch:

```
config system automation-stitch
    edit <automation-stitch-name>
        set status {enable | disable}
        set trigger <trigger-name>
        set action <action-name>
        set destination <serial-number>
    next
end
```

To test an automation stitch:

In the GUI, go to *Security Fabric* > *Automation*, right-click on the automation stitch and select *Test Automation Stitch*.

In the CLI, enter the following command:

```
diagnose automation test <stitch-name> <log>
```

Automation webhook stitches

The *Automation* menu contains eight webhook automation stitches, including an *Incoming Webhook Quarantine* trigger for API calls to the FortiGate, as well as a predefined *License Expired Notification* that replaces the existing license expiry alerts.

The automation stitches are available in new FortiGate installations by default. To install the stitches on an existing device, perform a factory reset.



Performing a factory reset will wipe the existing configurations from the FortiGate.

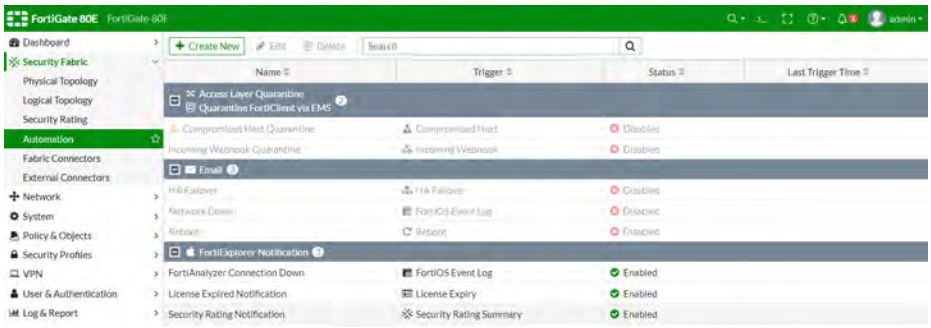
Before performing a factory reset, backup the existing configuration. Contact Fortinet support for additional assistance.

The following webhook stitches are included in the *Automation* menu:

- Compromised Host Quarantine
- Incoming Webhook quarantine
- HA Failover
- Network Down
- Reboot
- FortiAnalyzer Connection Down
- License Expired Notification
- Security rating Notification

To view the automation stitches in the GUI, go to *Security Fabric* > *Automation*.

Netskope Exhibit 1017



After the factory reset, the email alert feature will be removed from the GUI (*Log & Report* > *Email Alert Settings*), and replaced with the *Email* automation stitches.

You can continue using the email alert feature with the CLI console.

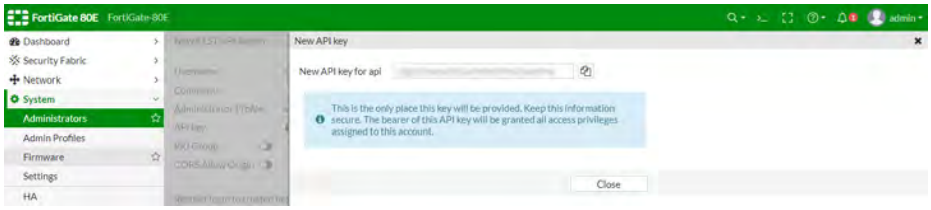
To configure the automation stitches in the CLI console, use the following commands:

```
config system automation-action
config system automation-trigger
config system automation-stitch
```

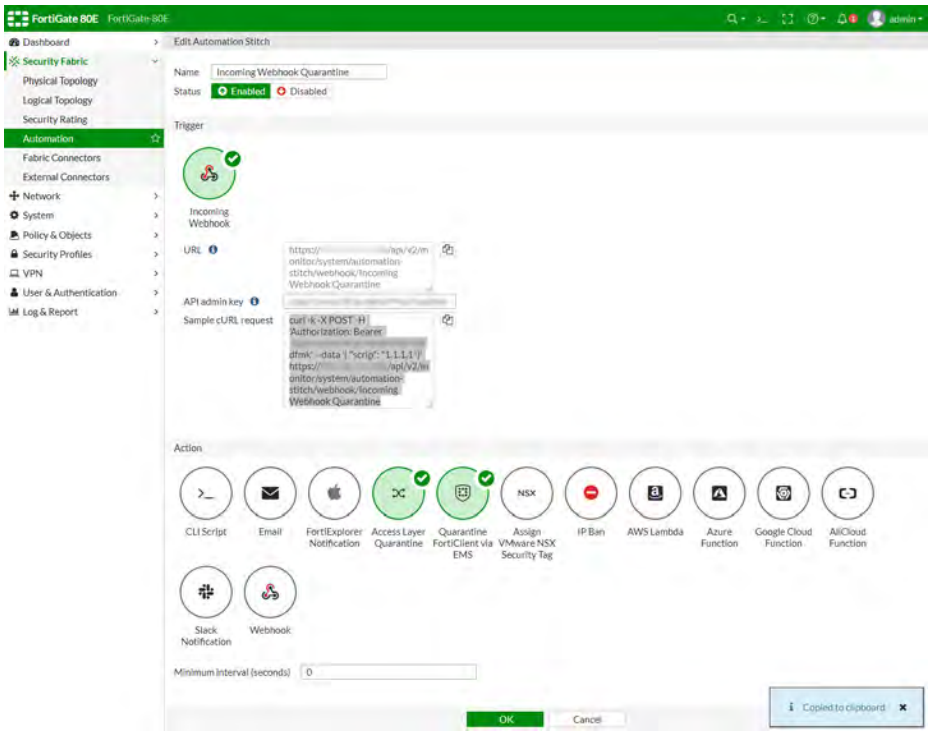
To view the configurations for the new automation stitches, see the CLI reference at the bottom of the page.

To trigger an Incoming Webhook Quarantine stitch with the GUI:

1. Create new API user.
   a. Go to *System* > *Administrators*.
   b. Click *Create New* > *REST API Admin*.
   c. Configure the *New REST API Admin* settings, and record the API key.



2. Get the sample cURL request.
   a. Go to *Security Fabric* > *Automation*.
   b. Under *Incoming Webhook*, right-click *Incoming Webhook Quarantine*, and select *Edit*.
   c. Click *Enabled*, to enable the rule.
   d. In the *API admin key* field, enter the API key you recorded in the previous step. A *Sample cURL request* is created.
   e. Copy the *Sample cURL request*.



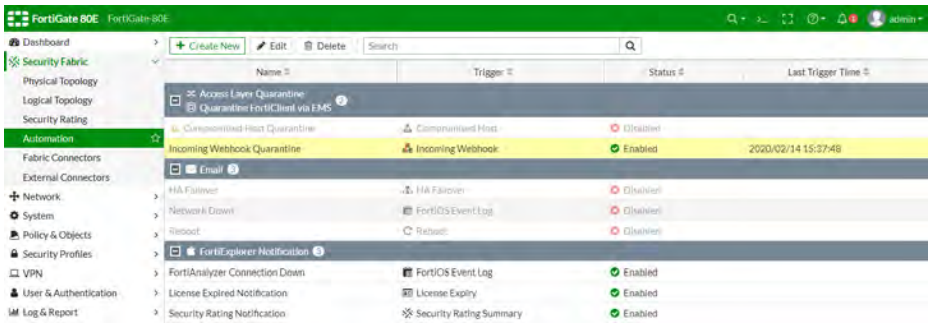
3. Execute the request:
   a. Edit the sample cURL you recorded in the previous step.
   b. Add parameters to the `data` field (`"mac"` and *"fctuid"*), and then execute the request.

```
root@pc:~# curl -k -X POST -H 'Authorization: Bearer
cfgtct1mmx3fQxr4khb994p7swdfmk' --data '{ "mac":"0c:0a:00:0c:ce:b0", "fctuid":
"0000BB0B0ABD0D00B0D0A0B0E0F0B00B"}'
https://172.16.116.226/api/v2/monitor/system/automation-
stitch/webhook/Incoming%20Webhook%20Quarantine
{
   "http_method":"POST",
   "status":"success",
   "http_status":200,
   "serial":"FGT00E0Q00000000",
   "version":"v6.4.0",
   "build":1545
```



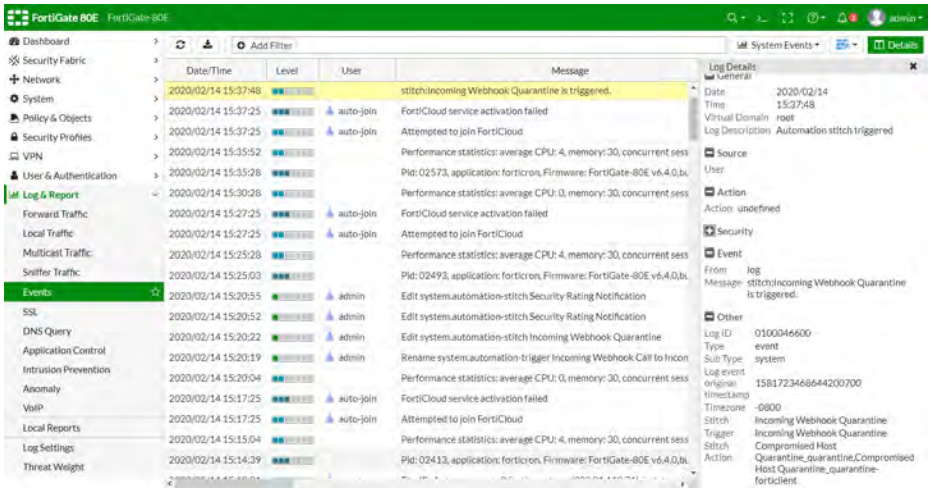
Encode the spaces in the automation-stitch name with `%20`. For example, `Incoming%20Webhook%20Quarantine`

The automation rule *Incoming Webhook Quarantine* is triggered.



The MAC address is quarantined in FortiGate and an event log is created.

The FortiClient UUID is quarantined by EMS on the server side.



To create an automated stitch with the CLI:

1. Create new API user and record the API key.

```
config system api-user
   edit "api"
      set api-key ENC SH00vqP0GKWKyZNz0FP0/jq00O0Ka/DHVEKdxUi+0kRDNKPpZppnnMk0KeunBI=
      set accprofile "api_profile"
      set vdom "root"
      config trusthost
         edit 1
            set ipv4-trusthost 10.6.30.0 200.200.200.0
         next
      end
   next
end
```

2. Configure the automation stitch, `Incoming Webhook Quarantine`.

```
config system automation-stitch
   edit "Incoming Webhook Quarantine"
      set status enable
      set trigger "Incoming Webhook Quarantine"
      set action "Compromised Host Quarantine_quarantine" "Compromised Host Quarantine_
          quarantine-forticlient"
   next
end
```

3. Add parameters in the `data` field (`"mac"` and `"fctuid"`), then execute the request on a device.

```
root@pc56:~# curl -k -X POST -H 'Authorization: Bearer
cfgtct1mmx0fQxr4khb000p70wdfmk' --data '{ "mac":"0c:0a:00:0c:ce:b0", "fctuid":
"3000BB0B0ABD0D00B0D0A0B0E0F0B00B"}'
https://100.10.100.200/api/v2/monitor/system/automation-
stitch/webhook/Incoming%20Webhook%20Quarantine
{
   "http_method":"POST",
   "status":"success",
   "http_status":200,
   "serial":"FGT80E0Q00000000",
   "version":"v6.4.0",
   "build":1545
```



Encode the spaces in the automation-stitch name with `%20`. For example, `Incoming%20Webhook%20Quarantine`

The automation rule "Incoming Webhook Quarantine" is triggered. The MAC address is quarantined in FortiGate, and an event log is created. The FortiClient UUID will be quarantined on the EMS server side.

```
config user quarantine
   config targets
      edit "0c:0a:00:0c:ce:b0"
         config macs
            edit 0c:0a:00:0c:ce:b0
               set description "Quarantined by automation stitch: Incoming Webhook
                    Quarantine"
            next
         end
      next
   end
end
date=2020-02-14 time=15:37:48 logid="0100046600" type="event" subtype="system"
level="notice" vd="root" eventtime=1581723468644200712 tz="-0800"
logdesc="Automation stitch triggered" stitch="Incoming Webhook Quarantine"
trigger="Incoming Webhook Quarantine" stitchaction="Compromised Host Quarantine_
quarantine,Compromised Host Quarantine_quarantine-forticlient" from="log"
msg="stitch:Incoming Webhook Quarantine is triggered."
```

CLI Reference

Network down

config system automation-action

```
config system automation-action
   edit "Network Down_email"
      set action-type email
      set email-from ''
      set email-subject "Network Down"
      set minimum-interval 0
      set delay 0
      set required disable
```


Netskope Exhibit 1017

```
      set message "%%log%%"
   next
end
```

config system automation-trigger

```
config system automation-trigger
   edit "Network Down"
      set trigger-type event-based
      set event-type event-log
      set logid 20099
      config fields
         edit 1
            set name "status"
            set value "DOWN"
         next
      end
   next
end
```

config system automation-stitch

```
config system automation-stitch
   edit "Network Down"
      set status disable
      set trigger "Network Down"
      set action "Network Down_email"
   next
end
```

HA failover

config system automation-action

```
config system automation-action
   edit "HA Failover_email"
      set action-type email
      set email-from ''
      set email-subject "HA Failover"
      set minimum-interval 0
      set delay 0
      set required disable
      set message "%%log%%"
   next
end
```

config system automation-trigger

```
config system automation-trigger
   edit "HA Failover"
      set trigger-type event-based
      set event-type ha-failover
   next
end
```


Netskope Exhibit 1017

config system automation-stitch

```
config system automation-stitch
   edit "HA Failover"
      set status disable
      set trigger "HA Failover"
      set action "HA Failover_email"
   next
end
```

Reboot

config system automation-action

```
config system automation-action
   edit "Reboot_email"
      set action-type email
      set email-from ''
      set email-subject "Reboot"
      set minimum-interval 0
      set delay 0
      set required disable
      set message "%%log%%"
   next
end
```

config system automation-trigger

```
config system automation-trigger
   edit "Reboot"
      set trigger-type event-based
      set event-type reboot
   next
end
```

config system automation-stitch

```
config system automation-stitch
   edit "Reboot"
      set status disable
      set trigger "Reboot"
      set action "Reboot_email"
   next
end
```

Connection down

config system automation-action

```
config system automation-action
   edit "FortiAnalyzer Connection Down_ios-notification"
      set action-type ios-notification
      set minimum-interval 0
      set delay 0
      set required disable
```


Netskope Exhibit 1017

```
   next
end
```

config system automation-trigger

```
config system automation-trigger
   edit "FortiAnalyzer Connection Down"
      set trigger-type event-based
      set event-type event-log
      set logid 22902
   next
end
```

config system automation-stitch

```
config system automation-stitch
   edit "FortiAnalyzer Connection Down"
      set status enable
      set trigger "FortiAnalyzer Connection Down"
      set action "FortiAnalyzer Connection Down_ios-notification"
   next
end
```

License expired

config system automation-action

```
config system automation-action
   edit "License Expired Notification_ios-notification"
      set action-type ios-notification
      set minimum-interval 0
      set delay 0
      set required disable
   next
end
```

config system automation-trigger

```
config system automation-trigger
   edit "License Expired Notification"
      set trigger-type event-based
      set event-type license-near-expiry
      set license-type any
   next
end
```

config system automation-stitch

```
config system automation-stitch
   edit "License Expired Notification"
      set status enable
      set trigger "License Expired Notification"
      set action "License Expired Notification_ios-notification"
   next
end
```


Netskope Exhibit 1017

Compromised host

config system automation-action

```
config system automation-action
   edit "Compromised Host Quarantine_quarantine"
      set action-type quarantine
      set minimum-interval 0
      set delay 0
      set required disable
   next
end
```

config system automation-trigger

```
config system automation-trigger
   edit "Compromised Host Quarantine"
      set trigger-type event-based
      set event-type ioc
      set ioc-level high
   next
end
```

config system automation-stitch

```
config system automation-stitch
   edit "Compromised Host Quarantine"
      set status disable
      set trigger "Compromised Host Quarantine"
      set action "Compromised Host Quarantine_quarantine" "Compromised Host Quarantine_
           quarantine-forticlient"
   next
end
```

Quarantine FortiClient

config system automation-action

```
config system automation-action
   edit "Compromised Host Quarantine_quarantine-forticlient"
      set action-type quarantine-forticlient
      set minimum-interval 0
      set delay 0
      set required disable
   next
end
```

config system automation-trigger

```
config system automation-trigger
   edit "Compromised Host Quarantine"
      set trigger-type event-based
      set event-type ioc
      set ioc-level high
   next
```


Netskope Exhibit 1017

```
end
```

config system automation-stitch

```
config system automation-stitch
   edit "Compromised Host Quarantine"
      set status disable
      set trigger "Compromised Host Quarantine"
      set action "Compromised Host Quarantine_quarantine" "Compromised Host Quarantine_
          quarantine-forticlient"
   next
end
```

Security rating

config system automation-action

```
config system automation-action
   edit "Security Rating Notification_ios-notification"
      set action-type ios-notification
      set minimum-interval 0
      set delay 0
      set required disable
   next
end
```

config system automation-trigger

```
config system automation-trigger
   edit "Security Rating Notification"
      set trigger-type event-based
      set event-type security-rating-summary
   next
end
```

config system automation-stitch

```
config system automation-stitch
   edit "Security Rating Notification"
      set status enable
      set trigger "Security Rating Notification"
      set action "Security Rating Notification_ios-notification"
   next
end
```

Chaining and delaying actions

Automation stitches that use cloud-based or webhook actions have the option to delay an action after the previous action is completed. The execution of the actions can be delayed by up to 3600 seconds (one hour).

To configure this option in the GUI, select a cloud-based action, then enter the required value, in seconds, in the action configuration's *Delay* field.

To configure a delay in the CLI, use the following command:

Netskope Exhibit 1017

```
config system automation-action
    edit <name>
        set action-type {aws-lambda | azure-function | google-cloud-function | alicloud-
function | webhook}
        set required {enable | disable}
        set delay <seconds>
    next
end
```

Triggers

The following table outlines the available automation stitch triggers:

Trigger	Description
Compromised Host	An Indicator of Compromise (IoC) is detected on a host endpoint. The threat level must be selected and can be *Medium* or *High*. If *Medium* is selected, both medium and high level threats are included. **Note:** Additional actions are available only for *Compromised Host* triggers: • Access Layer Quarantine • Quarantine FortiClient via EMS • Assign VMware NSX Security Tag • IP Ban
Security Rating Summary	A summary is available for a recently run Security Rating.
Configuration Change	A FortiGate configuration change has occurred.
Reboot	A FortiGate is rebooting.
Low memory	This option is only available in the CLI. Conserve mode due to low memory. See Execute a CLI script based on CPU and memory thresholds on page 328 for an example.
High CPU	This option is only available in the CLI. High CPU usage. See Execute a CLI script based on CPU and memory thresholds on page 328 for an example.
License Expiry	A FortiGuard license is expiring. The license type must be selected. Options include: • FortiCare Support • FortiGuard Web Filter • FortiGuard AntiSpam • FortiGuard AntiVirus • FortiGuard IPS • FortiGuard Management Service • FortiGate Cloud
HA Failover	An HA failover is occurring.
AV & IPS DB Update	The antivirus and IPS database is updating.

Trigger	Description
FortiOS Event Log	The specified FortiOS log has occurred. The event must be selected from the event list.
FortiAnalyzer Event Handler	The specified FortiAnalyzer event handler has occurred. See FortiAnalyzer event handler trigger on page 298 for details.
Schedule	A scheduled monthly, weekly, daily, or hourly trigger. Set to occur on a specific minute of an specific hour on a specific day.
FortiGate Cloud-Based IOC	IOC detection from the FortiGate Cloud IOC service. This option requires an IOC license, a web filter license, and FortiCloud logging must be enabled.

FortiAnalyzer event handler trigger

You can trigger automation stitches based on FortiAnalyzer event handlers. This allows you to define rules based on complex correlations across devices, log types, frequencies, and other criteria.

To set up a FortiAnalyzer event handler trigger:

1. Configure a FortiGate event handler on the FortiAnalyzer
2. Configure FortiAnalyzer logging on the FortiGate on page 299
3. Configure an automation stitch that is triggered by a FortiAnalyzer event handler on page 299

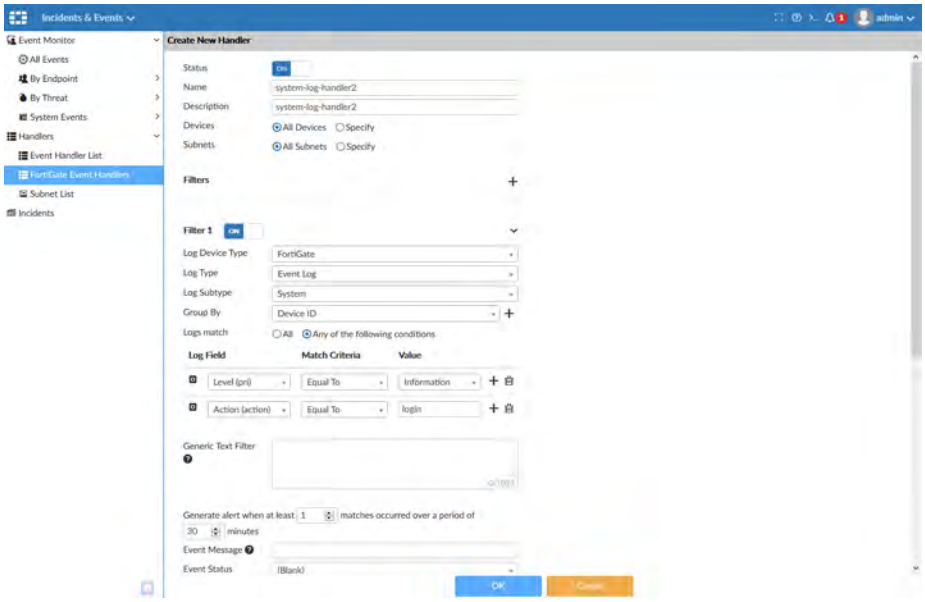
Configure a FortiGate event handler on the FortiAnalyzer

On the FortiAnalyzer, configure an event handler for the automation stitch. In this example, the event handler is triggered when an administrator logs in to the FortiGate.

Netskope Exhibit 1017

To configure an event handler on the FortiAnalyzer:

1. Go to *Incidents & Events* > *Handlers* > *FortiGate Event Handlers*.
2. Configure an event handler for the automation stitch.



3. Click *OK*.

Configure FortiAnalyzer logging on the FortiGate

See Configuring FortiAnalyzer on page 120 for more information.

To configure FortiAnalyzer logging in the GUI:

1. Go to *Security Fabric* > *Fabric Connectors* and double-click the *FortiAnalyzer Logging* card.
2. Click *Enabled* and configure the settings as needed.
3. Click *OK*.

To configure FortiAnalyzer logging in the CLI:

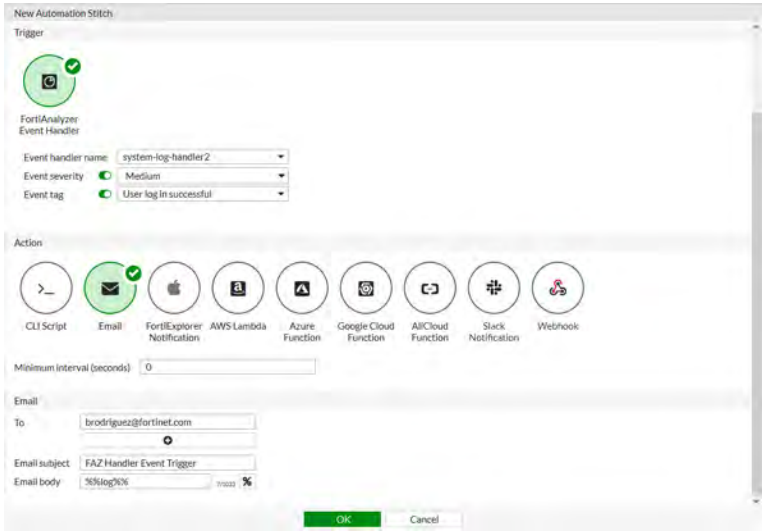
```
config log fortianalyzer setting
    set status enable
    set server "10.6.30.250"
    set serial "FL-4HET318900407"
    set upload-option realtime
    set reliable enable
end
```

Configure an automation stitch that is triggered by a FortiAnalyzer event handler

When a FortiAnalyzer event handler is triggered, it sends a notification to the FortiGate automation framework, which generates a log and triggers the automation stitch.

To configure an automation stitch that is triggered by a FortiAnalyzer event handler in the GUI:

1. Go to *Security Fabric* > *Automation*.
2. Click *Create New*.
3. In the *Trigger* section, select *FortiAnalyzer Event Handler*.
4. Set *Event handler name* to the event that was created on the FortiAnalyzer.
5. Set the *Event severity*, and select or create an *Event tag*.



6. In the *Action* section, select *Email* and configure the email recipient and message.
7. Click *OK*.

To configure an automation stitch that is triggered by a FortiAnalyzer event handler in the CLI:

1. Create an automation action:

```
config system automation-action
    edit "auto-faz-1_email"
        set action-type email
        set email-to "jnjssll@fortinet.com"
        set email-subject "CSF stitch alert"
        set message "User login FortiGate successfully."
    next
end
```

2. Create an automation trigger:

```
config system automation-trigger
    edit "auto-faz-1"
        set event-type faz-event
        set faz-event-name "system-log-handler2"
        set faz-event-severity "medium"
        set faz-event-tags "User login successfully"
    next
end
```

3. Create the automation stitch:

```
config system automation-stitch
    edit "auto-faz-1"
```

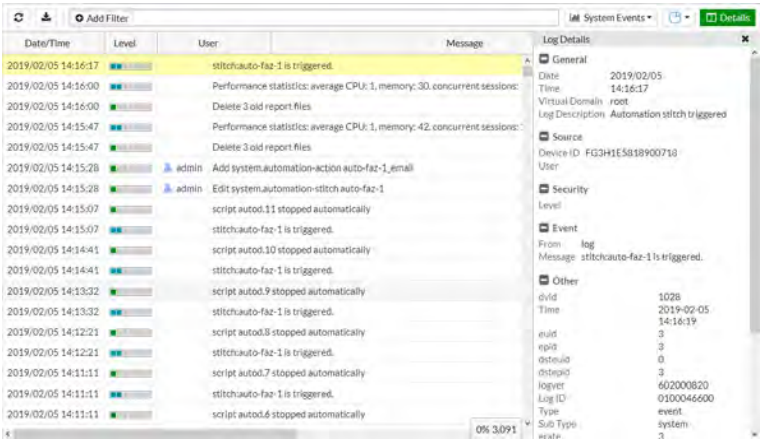

Netskope Exhibit 1017

```
        set trigger "auto-faz-1"
        set action "auto-faz-1_email"
    next
end
```

View the trigger event log

To see the trigger event log in the GUI:

1. Log in to the FortiGate.
   The FortiAnalyzer sends notification to the FortiGate automation framework, generates an event log on the FortiGate, and triggers the automation stitch.
2. Go to *Log & Report* > *Events* and select *System Events*.

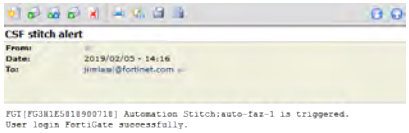


To see event logs in the CLI:

```
execute log display
    ...
    date=2019-02-05 time=14:16:17 logid="0100046600" type="event" subtype="system"
level="notice" vd="root" eventtime=1549404977 logdesc="Automation stitch triggered"
stitch="auto-faz-1" trigger="auto-faz-1" from="log" msg="stitch:auto-faz-1 is triggered."
    ...
```

Sample email

The email sent by the action will look similar to the following:



Actions

The following table outlines the available automation stitch actions. Multiple actions can be added and reorganized as needed by dragging and dropping.

Netskope Exhibit 1017

Action	Description
Alert	Generate a FortiOS dashboard alert. This option is only available in the CLI.
CLI Script	Run one or more CLI scripts. See CLI script action on page 303 for details. See Execute a CLI script based on CPU and memory thresholds on page 328 for an example.
Disable SSID	Disable the SSID interface. This option is only available in the CLI.
Email	Send a custom email message to the selected recipients. At least one recipient and an email subject must be specified. The email body can use parameters from logs or previous action results. Wrapping the parameter with %% will replace the expression with the JSON value for the parameter, for example: *%%results.source%%* is the source property from the previous action.
FortiExplorer Notification	Send push notifications to FortiExplorer. The FortiGate must be registered to FortiCare on the iOS App that will receive the notification.
Access Layer Quarantine	This option is only available for Compromised Host triggers. Impose a dynamic quarantine on multiple endpoints based on the access layer.
Quarantine FortiClient via EMS	This option is only available for Compromised Host triggers. Use FortiClient EMS to block all traffic from the source addresses that are flagged as compromised hosts. Quarantined devices are flagged on the Security Fabric topology views. Go to *Monitor* > *Quarantine Monitor* to view and manage quarantined IP addresses.
Assign VMware NSX Security Tag	This option is only available for Compromised Host triggers. If an endpoint instance in a VMware NSX environment is compromised, the configured security tag is assigned to the compromised endpoint. See NSX Quarantine action on page 305 for details.
IP Ban	This option is only available for Compromised Host triggers. Block all traffic from the source addresses flagged by the IoC. Go to *Monitor* > *Quarantine Monitor* to view and manage banned IP addresses.
AWS Lambda	Send log data to an integrated AWS service. See AWS Lambda action on page 307 for details.
Azure Function	Send log data to an Azure function. See Azure Function action on page 309 for details.
Google Cloud Function	Send log data to a Google Cloud function. See Google Cloud Function action on page 312 for details.
AliCloud Function	Send log data to an AliCloud function. See AliCloud Function action on page 314 for details.

Action	Description
Slack Notification	Send a notification to a Slack channel. See Slack notification action on page 316 for details.
Webhook	Send an HTTP request using a REST callback. See Webhook action on page 320 for details, and Slack integration webhook on page 325 and Microsoft Teams integration webhook on page 326 for examples.

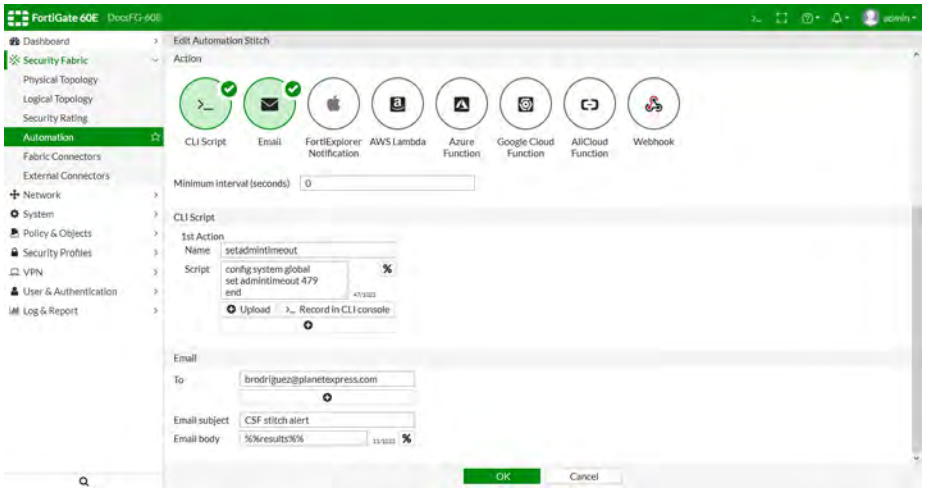
CLI script action

CLI scripts can be run when an automation stitch is triggered. The scripts can be manually entered, uploaded as a file, or recorded in the CLI console. The output of the script can be sent as an email action.

In this example, the script sets the idle timeout value to 479 minutes, and sends an email with the script output.

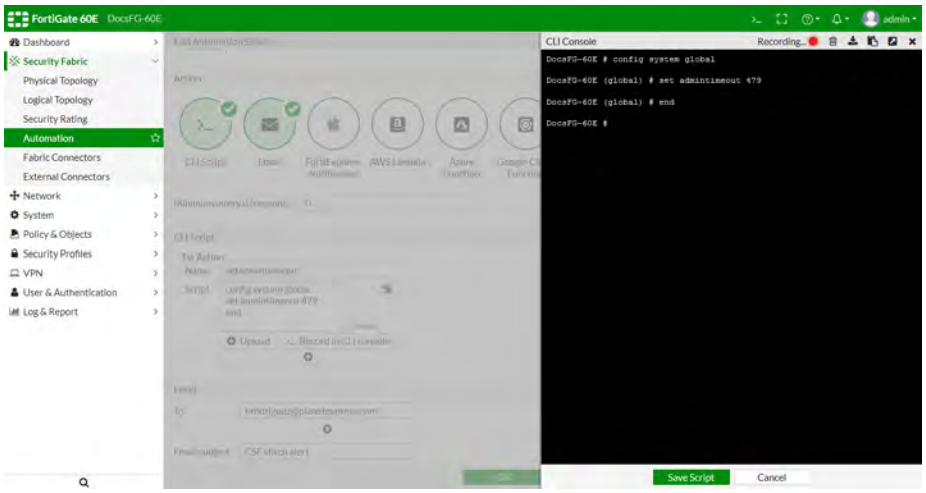
To configure a CLI script automation stitch in the GUI:

1. Go to *Security Fabric* > *Automation*.
2. Click *Create New*.
3. Enter a name for the stitch, and select the FortiGate devices that it will be applied to.
4. Select a trigger, such as *Security Rating Summary*.
5. Select *CLI Script* and *Email* actions.
6. Configure the CLI script:
   - To manually enter the script, type it into the *Script* field.



   - To upload a script file, click *Upload* and locate the file on your management computer.

- To record the script in the CLI console, click *>_Record in CLI console*, then enter the CLI commands.



7. Configure the email action.
8. Click *OK*.

To configure a CLI script automation stitch in the CLI:

1. Create an automation action:

```
config system automation-action
    edit "set admintimeout479"
        set action-type cli-script
        set minimum-interval 0
        set delay 0
        set required enable
        set script "config system global
            set admintimeout 479
            end"
    next
    edit "auto-cli-1_email"
        set action-type email
        set email-to "jnkssll@fortinet.com"
        set email-subject "CSF stitch alert"
        set message "%%results%%"
        set minimum-interval 0
    next
end
```

2. Create an automation trigger:

```
config system automation-trigger
    edit "auto-cli-1"
        set trigger-type event-based
        set event-type security-rating-summary
    next
end
```

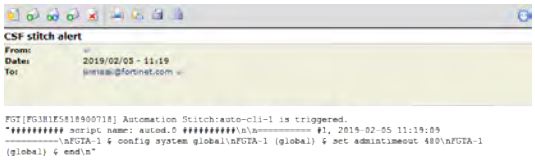
3. Create the automation stitch:

```
config system automation-stitch
    edit "auto-cli-1"
        set status enable
```

```
            set trigger "auto-cli-1"
            set action "set admintimeout479
        next
    end
```

Email sample

The email sent by the action will look similar to the following:



NSX Quarantine action

If an endpoint instance in a VMware NSX environment is compromised, this action will assign the configured security tag to the compromised endpoint.

This action is only available when the automation trigger is set to compromised host.

To set up the NSX quarantine action, you need to:

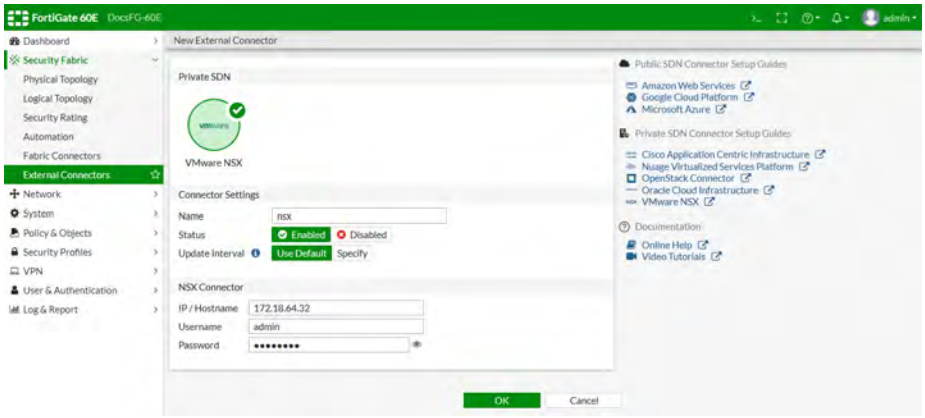
1. Configure a VMware NSX SDN connector
2. Configure an NSX security tag automation stitch
3. Configure FortiAnalyzer logging on the FortiGate

Configure a VMware NSX SDN connector

The FortiGate retrieves security tags from the VMware NSX server through the connector.

To configure a VMware NSX SDN connector in the GUI:

1. Go to *Security Fabric* > *External Connectors*.
2. Click *Create New*.
3. Select *VMware NSX*.
4. Configure the settings.



5. Click *OK*.

To configure a VMware NSX SDN connector in the CLI:

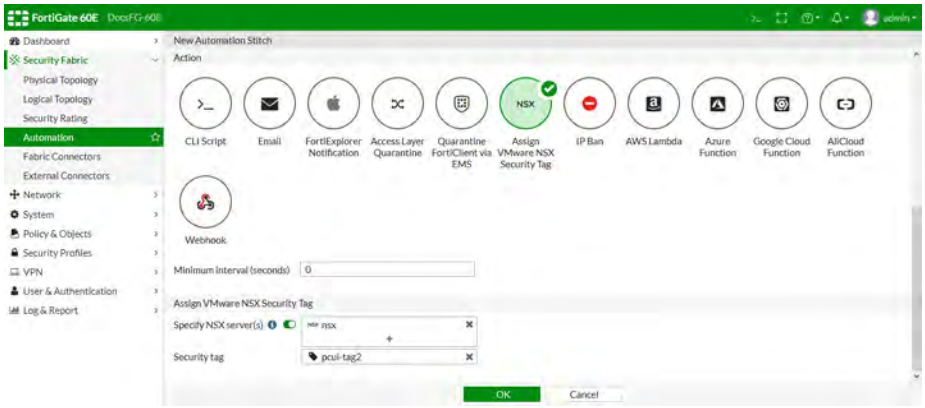
```
config system sdn-connector
    edit "nsx"
        set type nsx
        set server "172.18.64.32"
        set username "admin"
        set password xxxxxx
    next
end
```

Configure an NSX security tag automation stitch

Security tags are retrieved from the VMware NSX server through the NSX SDN connector.

To configure an NSX security tag automation stitch in the GUI:

1. Go to *Security Fabric* > *Automation*.
2. Click *Create New*.
3. In the *Trigger* section, select *Compromised Host*.
4. In the *Action* section, select *Assign VMware NSX Security Tag*.
5. Configure the settings.



6. Click *OK*.

To configure an NSX security tag automation stitch in the CLI:

1. Create an automation action:

```
config system automation-action
    edit "pcui-test_quarantine-nsx"
        set action-type quarantine-nsx
        set security-tag "pcui-tag2"
        set sdn-connector "nsx"
    next
end
```

2. Create an automation trigger:

```
config system automation-trigger
    edit "pcui-test"
        set ioc-level high
```

```
        next
    end
```

3. Create the automation stitch:

```
config system automation-stitch
    edit "pcui-test"
        set trigger "pcui-test"
        set action "pcui-test_quarantine-nsx"
    next
end
```

Configure FortiAnalyzer logging on the FortiGate

The FortiAnalyzer is used to send endpoint compromise notification to the FortiGate.

See Configuring FortiAnalyzer on page 120 for more information.

To configure FortiAnalyzer logging in the GUI:

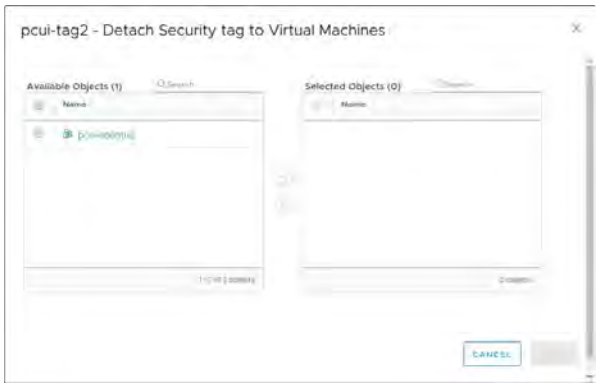
1. Go to *Security Fabric* > *Fabric Connectors* and double-click the *FortiAnalyzer Logging* card.
2. Click *Enabled* and configure the settings as needed.
3. Click *Apply*.

To configure FortiAnalyzer logging in the CLI:

```
config log fortianalyzer setting
    set status enable
    set server "172.18.64.234"
    set serial "FL-8HFT718900132"
    set upload-option realtime
    set reliable enable
end
```

When an endpoint instance is compromised

When an endpoint instance, such as *pcui-ubuntu2*, in the VMware NSX environment is compromised, the automation stitch is triggered. The FortiGate then assigns the configured security tag, *pcui-tag2* in this example, to the compromised NSX endpoint instance.

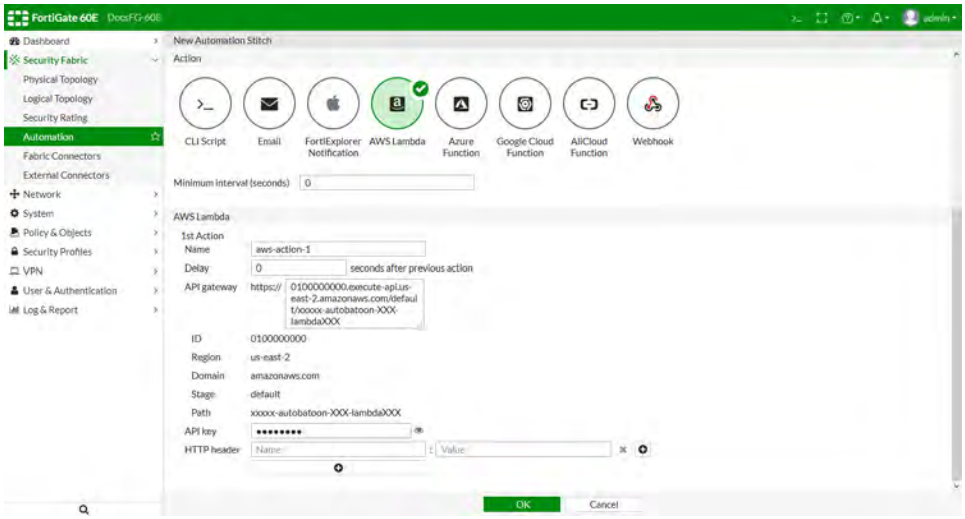


AWS Lambda action

AWS Lambda functions can be called when an automation stitch is triggered.

To configure an AWS Lambda function automation stitch in the GUI:

1. Go to *Security Fabric* > *Automation*.
2. Click *Create New*.
3. Enter a name for the stitch, and select the FortiGate devices that it will be applied to.
4. Select a trigger, such as *Security Rating Summary*.
5. Select *AWS Lambda Function* and configure its settings.



Name	The action name.
Delay	The amount of time after the previous action before this action executes, in seconds (0 - 3600, default = 0).
API gateway	The API gateway URL, in the format: `{restapi-id}.execute-api.{region}.{domain}/{stage}/{path}` The CLI must be used to manually enter the individual parameters.
API key	The API key configured in your API gateway.
HTTP header	The HTTP request header name and value. Multiple headers can be added.
+	Click to add another action. Actions can be reorganized as needed by dragging and dropping.
Name	The action name.
Delay	The amount of time after the previous action before this action executes, in seconds (0 - 3600, default = 0).

6. Click *OK*.

To configure an AWS Lambda function automation stitch in the CLI:

1. Create an automation action:

```
config system automation-action
    edit "aws-action-1"
        set action-type aws-lambda
        set aws-api-id "0100000000"
        set aws-region "us-east-2"
        set aws-api-stage "default"
        set aws-api-path "xxxxx-autobatoon-XXX-lambdaXXX"
        set aws-api-key *************
    next
end
```

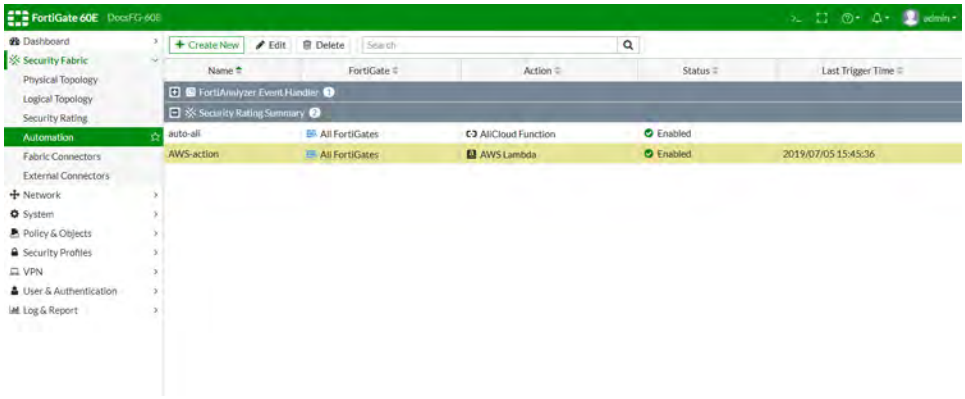
2. Create an automation trigger:

```
config system automation-trigger
    edit "auto-aws"
        set event-type security-rating-summary
    next
end
```

3. Create the automation stitch:

```
config system automation-stitch
    edit "auto-aws"
        set trigger "auto-aws"
        set action "aws-action-1"
    next
end
```

When the automation stitch is triggered, the FortiGate shows the stitch trigger time:



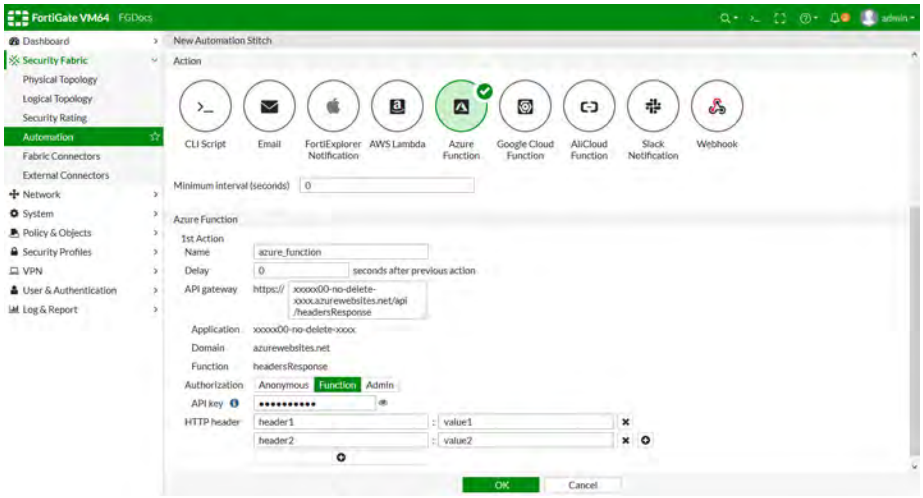
In AWS, the log shows that the function was called, executed, and finished.

Azure Function action

Azure functions can be called when an automation stitch is triggered.

To configure an Azure function automation stitch in the GUI:

1. Go to *Security Fabric* > *Automation*.
2. Click *Create New*.
3. Enter a name for the stitch, and select the FortiGate devices that it will be applied to.
4. Select a trigger, such as *Security Rating Summary*.
5. Select *Azure Function* and configure its settings.



Name	The action name.
Delay	The amount of time after the previous action before this action executes, in seconds (0 - 3600, default = 0).
API gateway	The API gateway URL, in the format: `{application}.{domain}/api/{function}` The CLI must be used to manually enter the individual parameters.
Authorization	The authorization level: *Anonymous*, *Function*, or *Admin*.
API key	The API key configured in your API gateway. This options is only available when *Authorization* is not *Anonymous*.
HTTP header	The HTTP request header name and value. Multiple headers can be added.
+	Click to add another action. Actions can be reorganized as needed by dragging and dropping.

6. Click *OK*.

To configure an Azure function automation stitch in the CLI:

1. Create an automation action:

```
config system automation-action
    edit "azure_function"
        set action-type azure-function
        set azure-app "xxxxx00-no-delete-xxxx"
        set azure-function "headersResponse"
        set azure-function-authorization function
```

```
        set azure-api-key **********
        set headers "header1:value1" "header2:value2"
    next
end
```

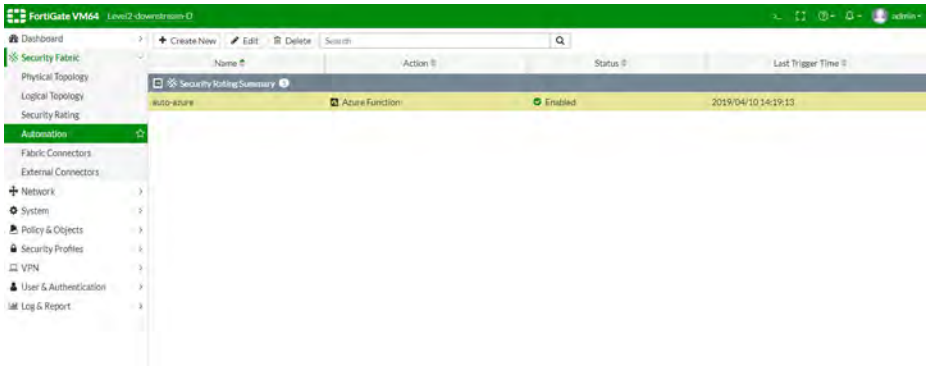
2. Create an automation trigger:

```
config system automation-trigger
    edit "auto-azure"
        set event-type security-rating-summary
    next
end
```

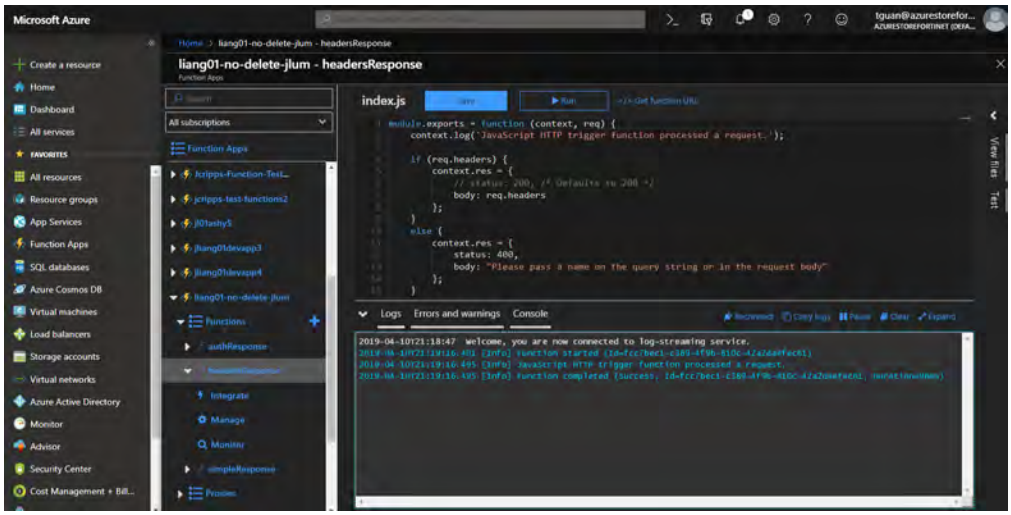
3. Create the automation stitch:

```
config system automation-stitch
    edit "auto-azure"
        set trigger "auto-azure"
        set action "azure_function"
    next
end
```

When the automation stitch is triggered, the FortiGate shows the stitch trigger time:



In Azure, the function log shows that the function was called, executed, and finished:

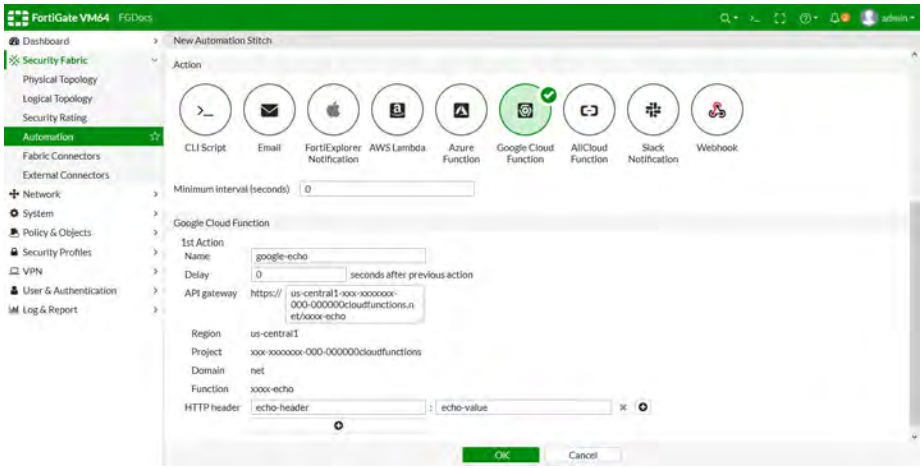


Google Cloud Function action

Google Cloud functions can be called when an automation stitch is triggered.

To configure a Google Cloud function automation stitch in the GUI:

1. Go to *Security Fabric* > *Automation*.
2. Click *Create New*.
3. Enter a name for the stitch, and select the FortiGate devices that it will be applied to.
4. Select a trigger, such as *Security Rating Summary*.
5. Select *Google Cloud Function* and configure its settings.



Name	The action name.
Delay	The amount of time after the previous action before this action executes, in seconds (0 - 3600, default = 0).
API gateway	The API gateway URL, in the format: `{region}-{project}{domain}/{function}` The CLI must be used to manually enter the individual parameters.
HTTP header	The HTTP request header name and value. Multiple headers can be added.
+	Click to add another action. Actions can be reorganized as needed by dragging and dropping.

6. Click *OK*.

To configure a Google Cloud function automation stitch in the CLI:

1. Create an automation action:

```
config system automation-action
    edit "google-echo"
        set action-type google-cloud-function
        set gcp-function-region "us-central1"
        set gcp-project "xxx-xxxxxxx-000-000000"
        set gcp-function-domain "cloudfunctions.net"
        set gcp-function "xxxx-echo"
```

```
        set headers "echo-header:echo-value"
    next
end
```

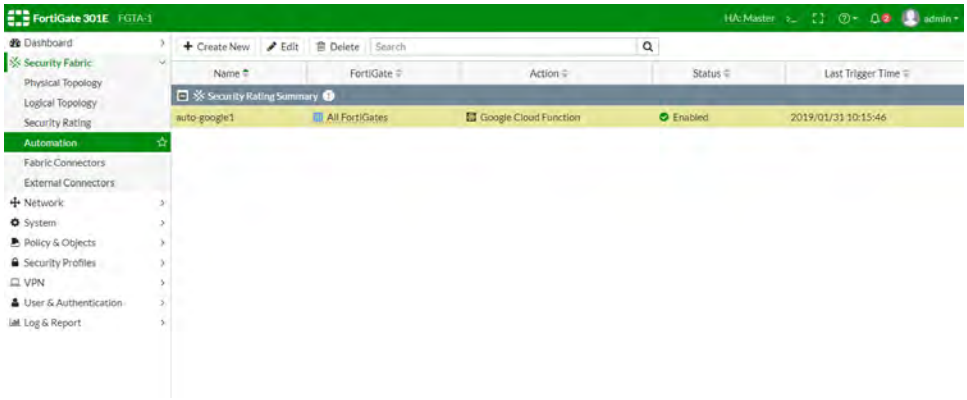
2. Create an automation trigger:

```
config system automation-trigger
    edit "auto-google1"
        set event-type security-rating-summary
    next
end
```

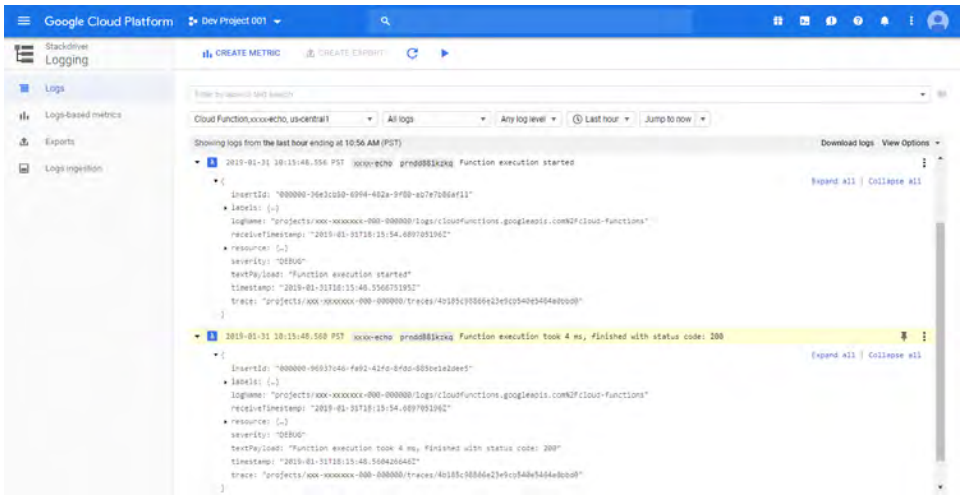
3. Create the automation stitch:

```
config system automation-stitch
    edit "auto-google1"
        set trigger "auto-google1"
        set action "google-echo"
    next
end
```

When the automation stitch is triggered, the FortiGate shows the stitch trigger time:



In Google Cloud, go to *Logs* to see the function log showing that the configured function was called, executed, and finished:

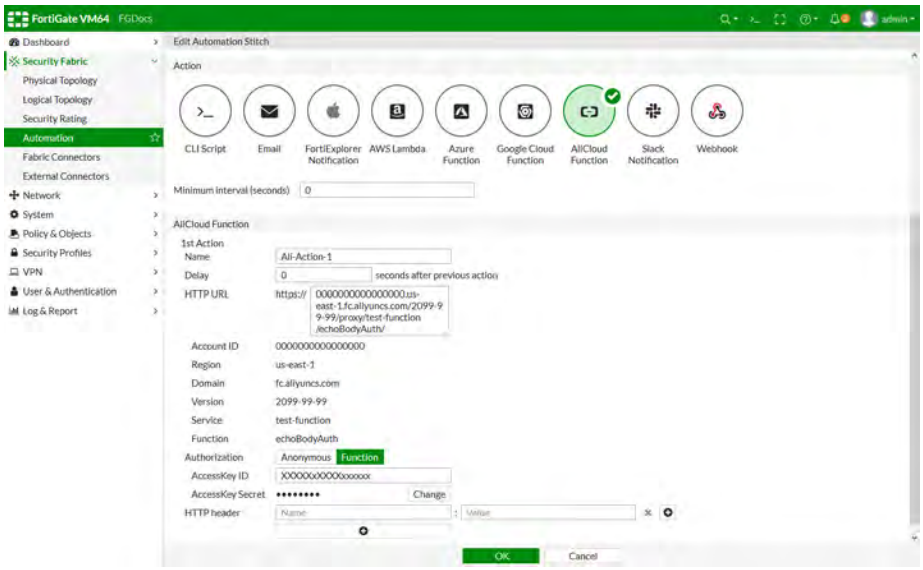


AliCloud Function action

AliCloud functions can be called when an automation stitch is triggered.

To configure an AliCloud function automation stitch in the GUI:

1. Go to *Security Fabric* > *Automation*.
2. Click *Create New*.
3. Enter a name for the stitch, and select the FortiGate devices that it will be applied to.
4. Select a trigger, such as *Security Rating Summary*.
5. Select *AliCloud Function* and configure its settings.



Name	The action name.
Delay	The amount of time after the previous action before this action executes, in seconds (0 - 3600, default = 0).
HTTP URL	The HTTP URL, in the format: `{account id}.{region}.{domain}/{version}/proxy/{service}/{function}` The CLI must be used to manually enter the individual parameters.
Authorization	The authorization level: *Anonymous*, or *Function*.
AccessKey ID	The access key ID This options is only available when *Authorization* is *Function*.
AccessKey Secret	The access key secret. This options is only available when *Authorization* is *Function*.
HTTP header	The HTTP request header name and value. Multiple headers can be added.
+	Click to add another action. Actions can be reorganized as needed by dragging and dropping.

6. Click *OK*.

To configure an AliCloud function automation stitch in the CLI:

1. Create an automation action:

```
config system automation-action
    edit "Ali-Action-1"
        set action-type alicloud-function
        set alicloud-account-id "0000000000000000"
        set alicloud-region "us-east-1"
        set alicloud-version "2099-99-99"
        set alicloud-service "test-function"
        set alicloud-function "echoBodyAuth"
        set alicloud-function-authorization function
        set alicloud-access-key-id "XXXXXxXXXXxxxxxx"
        set alicloud-access-key-secret xxxxxx
    next
end
```

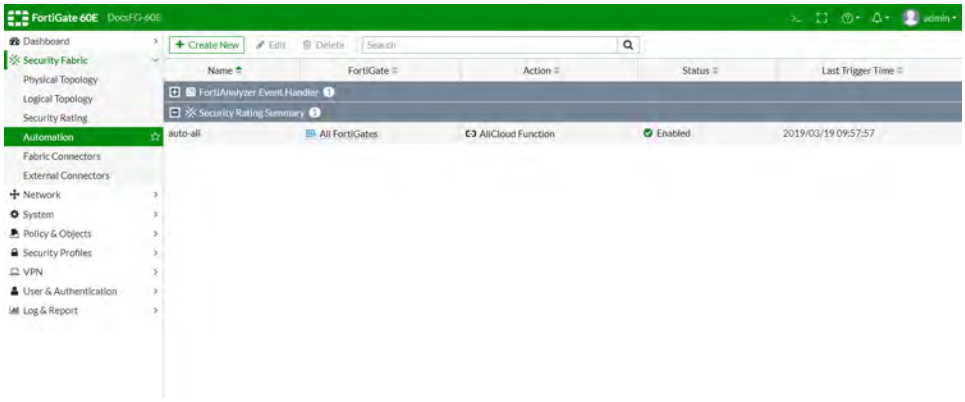
2. Create an automation trigger:

```
config system automation-trigger
    edit "auto-ali"
        set event-type security-rating-summary
    next
end
```

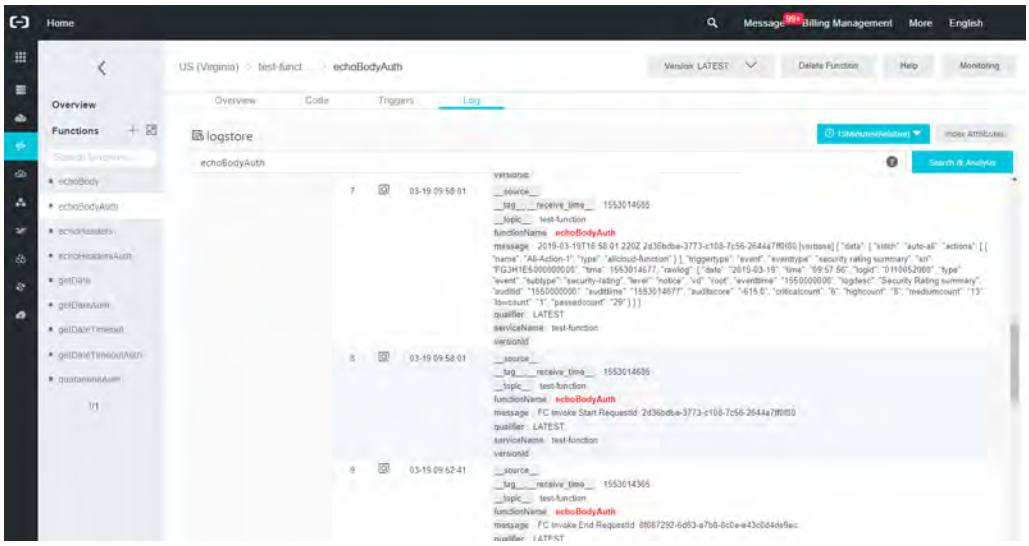
3. Create the automation stitch:

```
config system automation-stitch
    edit "auto-ali"
        set trigger "auto-ali"
        set action "Ali-Action-1"
    next
end
```

When the automation stitch is triggered, the FortiGate shows the stitch trigger time:



In AliCloud, the function log shows that the function was called, executed, and finished:

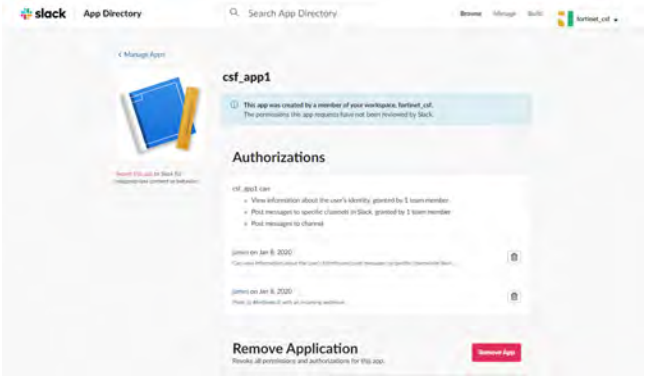


Slack notification action

To configure the automation stitch, create an incoming webhook in Slack, and then enter the webhook URL in the corresponding field of the notification action in FortiGate.

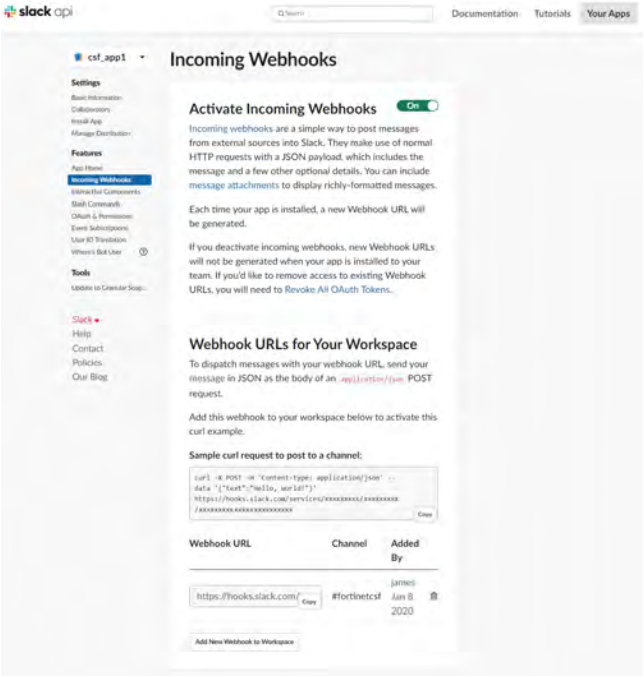
To create an incoming webhook in Slack:

1. Go to the Slack website, and create a workspace.
2. Create a Slack application for the workspace.



3. Add an *Incoming Webhook* to a channel in the workspace.

4. Activate the *Incoming Webhook*, and record the *Webhook URL*.

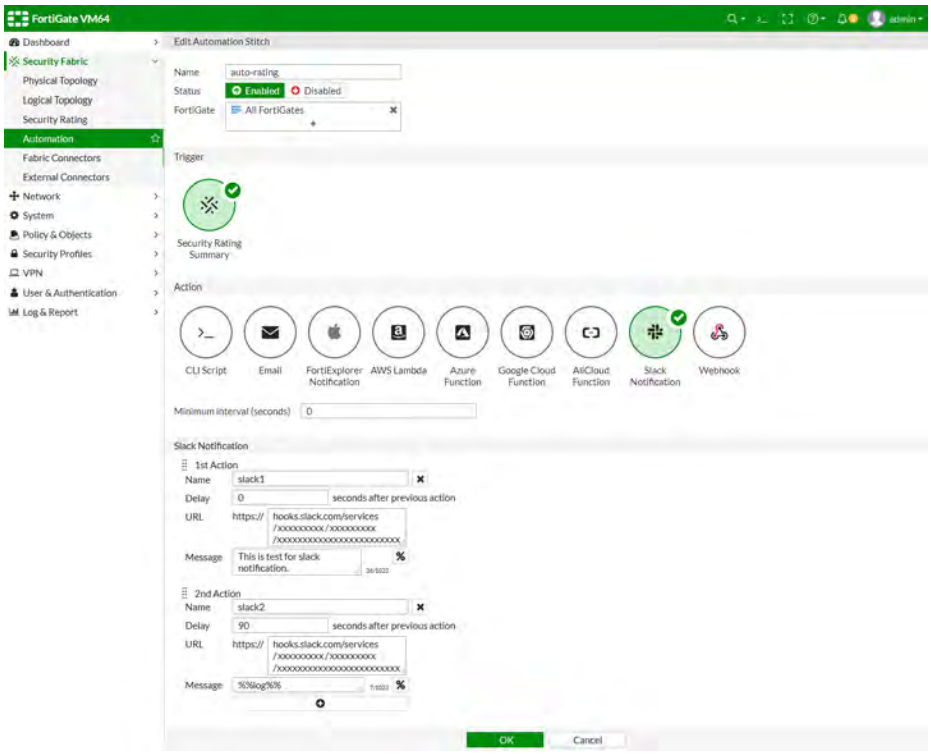


For information about using incoming webhooks in Slack, see https://api.slack.com/incoming-webhooks.

To configure a Slack notification in the GUI:

1. Go to *Security Fabric* > *Automation*.
2. Choose an automation stitch, and click *Edit*.

3. Select *Slack Notification*, and configure the notification settings.



Name	Enter a name for the notification.
Delay	Enter the number of seconds to delay the notification after the previous action is triggered.
URL	Enter the *Webhook URL* you recorded when you created the incoming webhook in Slack.
Message	Take one of the following actions: • Configure the message parameters. Click *%* to view a description of the available parameters. • Enter the message to display in the Slack channel.

4. (Optional) Click the plus (+) sign to add another action.
5. Click *OK*.
6. Run the automation stitch to trigger the action.

To configure a Slack notification in the CLI:

1. Add the webhook URL the Slack notification action:

```
config system automation-action
   edit "slack1"
      set action-type slack-notification
      set minimum-interval 0
      set delay 0
      set required disable
      set message "This is test for slack notification."
      set uri "hooks.slack.com/services/xxxxxxxxx/xxxxxxxxx/xxxxxxxxxxxxxxxxxxxxxxx"
   next
   edit "slack2"
      set action-type slack-notification
```

```
        set minimum-interval 0
        set delay 90
        set required disable
        set message "%%log%%"
        set uri "hooks.slack.com/services/xxxxxxxxx/xxxxxxxxx/xxxxxxxxxxxxxxxxxxxxxxxx"
    next
end
```

2. Create the trigger for the notification:

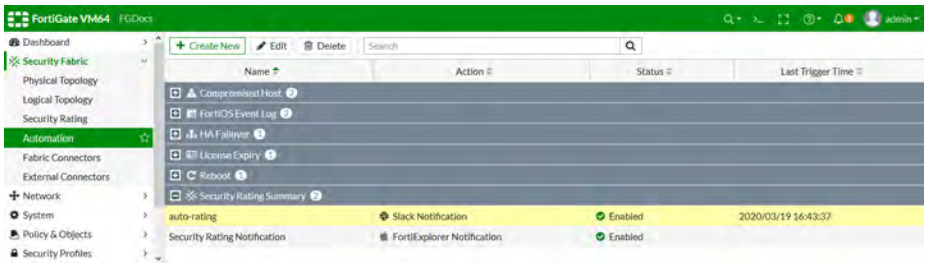
```
config system automation-trigger
    edit "auto-rating"
        set trigger-type event-based
        set event-type security-rating-summary
    next
end
```

3. Configure the action for the trigger:

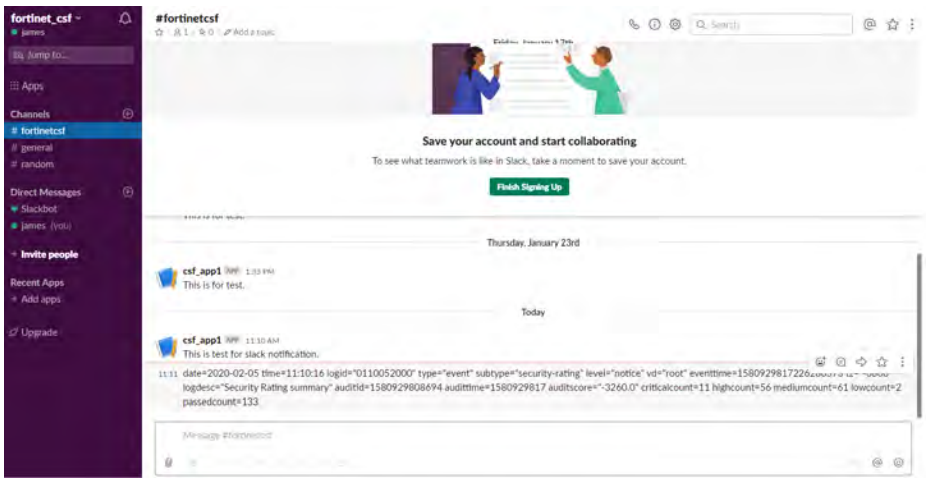
```
config system automation-stitch
    edit "auto-rating"
        set status enable
        set trigger "auto-rating"
        set action "slack1" "slack2"
    next
end
```

4. Trigger the notification.

The notification action is triggered in FortiGate.



The message you entered in the automation stitch is delivered to the Slack channel.



Webhook action

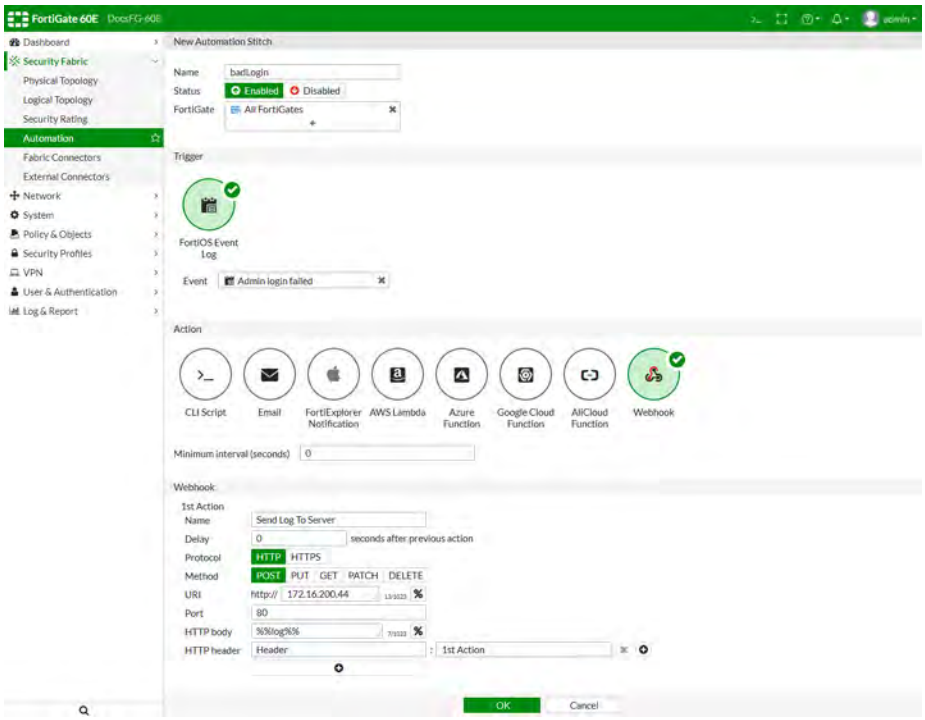
The webhook automation stitch action makes HTTP and HTTPS requests to a specified server, with custom headers, bodies, ports, and methods. It can be used to leverage the ubiquity of HTML requests and APIs to integrate with many other tools.

The URI and HTTP body can use parameters from logs or previous action results. Wrapping the parameter with %% will replace the expression with the JSON value for the parameter, for example: *%%results.source%%* is the source property from the previous action.

In this example, a specific log message (failed administrator log in attempt) triggers the FortiGate to send the contents of the log to a server. The server responds with a generic reply. This example assumes that the server is already configured and able to communicate with the FortiGate.

To configure the webhook automation stitch in the GUI:

1. Go to *Security Fabric* > *Automation*.
2. Click *Create New*.
3. Enter a name for the stitch, and select the FortiGate devices that it will be applied to.
4. Select the trigger *FortiOS Event Log*.
5. Set *Event* to *Admin login failed*.
6. Select *Webhook* and configure the settings:



Name	The action name.
Delay	The amount of time after the previous action before this action executes, in seconds (0 - 3600, default = 0).

Protocol	The request protocol to use: *HTTP* or *HTTPS*.
Method	The request method: *POST*, *PUT*, *GET*, *PATCH*, or *DELETE*.
URI	The request API URI.
Port	The protocol port.
HTTP body	The request body, if required, as a serialized JSON string. Use the parameter *%%log%%* to send the contents of the log from the trigger.
HTTP header	The HTTP request header name and value.
+	Click to add another action. Actions can be reorganized as needed by dragging and dropping.

7. Click *OK*.

To configure the webhook automation stitch in the CLI:

1. Create the automation action:

```
config system automation-action
    edit "Send Log To Server"
        set action-type webhook
        set uri "172.16.200.44"
        set http-body "%%log%%"
        set port 80
        set headers "Header:1st Action"
    next
end
```

2. Create an automation trigger:

```
config system automation-trigger
    edit "badLogin"
        set event-type event-log
        set logid 32002
    next
end
```

3. Create the automation stitch:

```
config system automation-stitch
    edit "badLogin"
        set trigger "badLogin"
        set action "Send Log To Server"
    next
end
```


Netskope Exhibit 1017

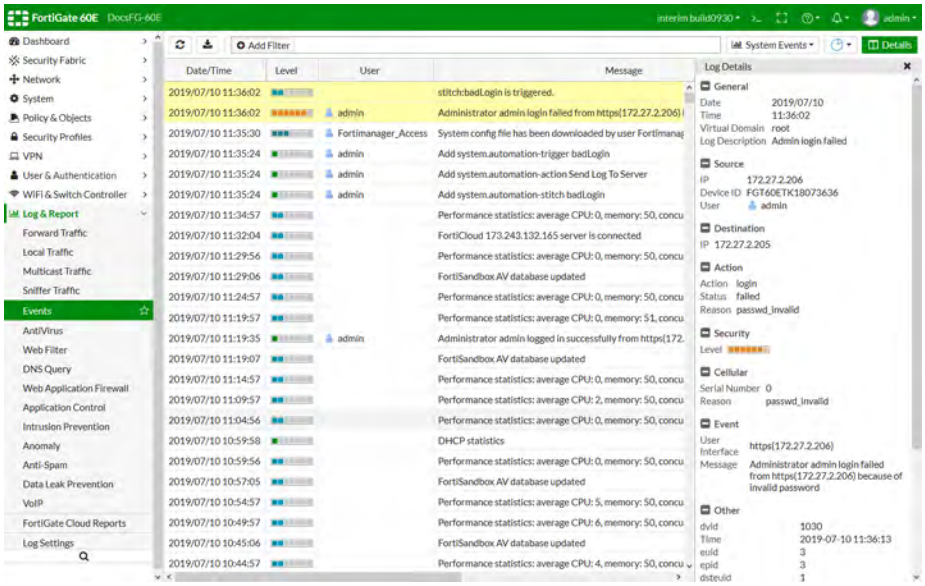
To test the automation stitch:

1. Attempt to log in to the FortiGate with an incorrect username or password.
2. On the server, check the log to see that its contents have been sent by the FortiGate.

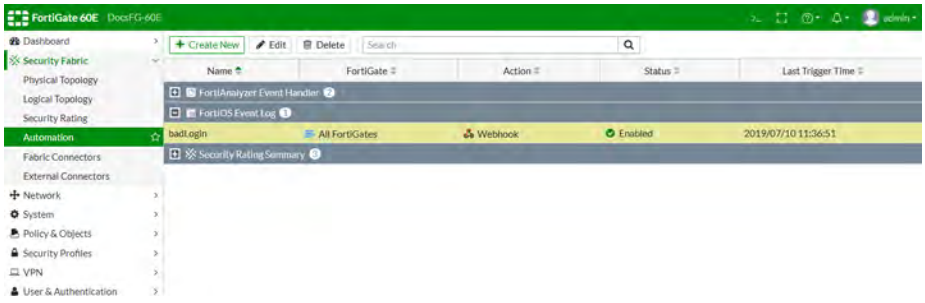


The body content is replaced with the log of the trigger.

3. On the FortiGate, go to *Log & Report* > *Events* and select *System Events* to confirm that the stitch was activated.



4. Go to *Security Fabric* > *Automation* to see the last time that the stitch was triggered.



Diagnose commands

- Enable log dumping:

```
# diagnose test application autod 1
```

```
autod dumped total:1 logs, num of logids:1
autod log dumping is enabled

vdom:root(0) logid:32002 len:408 log:
date=2019-05-30 time=17:41:03 logid="0100032002" type="event" subtype="system"
level="alert" vd="root" eventtime=1559263263858888451 tz="-0700" logdesc="Admin login
failed" sn="0" user="admin" ui="http(10.6.30.254)" method="http" srcip=10.6.30.254
dstip=10.6.30.5 action="login" status="failed" reason="passwd_invalid"
msg="Administrator admin login failed from http(10.6.30.254) because of invalid
password"
autod log dumping is disabled

autod logs dumping summary:
        logid:32002 count:1

autod dumped total:1 logs, num of logids:1
```

- Show automation settings:

```
# diagnose test application autod 2
csf: enabled   root:yes
total stitches activated: 2

stitch: badLogin
       destinations: all
       trigger: badLogin

       local hit: 6 relayed to: 6 relayed from: 6
       actions:
               Send Log To Server type:webhook interval:0
                       delay:0 required:no
                       proto:0 method:0 port:80
                       uri: 172.16.200.44
                       http body: %%log%%
                       headers:
                       0. Header:1st Action
```

- Show automation statistics:

```
# diagnose test application autod 3

stitch: badLogin

       local hit: 1 relayed to: 1 relayed from: 1
       last trigger:Wed Jul 10 12:14:14 2019
       last relay:Wed Jul 10 12:14:14 2019

       actions:
               Send Log To Server:
                       done: 1 relayed to: 1 relayed from: 1
                       last trigger:Wed Jul 10 12:14:14 2019
                       last relay:Wed Jul 10 12:14:14 2019

logid2stitch mapping:
id:32002  local hit: 3 relayed to: 3 relayed from: 3
       badLogin
```


Netskope Exhibit 1017

```
action run cfg&stats:
total:55 cur:0 done:55 drop:0
       email:
               flags:10
               stats: total:4 cur:0 done:4 drop:0
       ios-notification:
               flags:1
               stats: total:0 cur:0 done:0 drop:0
       alert:
               flags:0
               stats: total:0 cur:0 done:0 drop:0
       disable-ssid:
               flags:7
               stats: total:0 cur:0 done:0 drop:0
       quarantine:
               flags:7
               stats: total:0 cur:0 done:0 drop:0
       quarantine-forticlient:
               flags:4
               stats: total:0 cur:0 done:0 drop:0
       quarantine-nsx:
               flags:4
               stats: total:0 cur:0 done:0 drop:0
       ban-ip:
               flags:7
               stats: total:0 cur:0 done:0 drop:0
       aws-lambda:
               flags:11
               stats: total:21 cur:0 done:21 drop:0
       webhook:
               flags:11
               stats: total:6 cur:0 done:6 drop:0
       cli-script:
               flags:10
               stats: total:4 cur:0 done:4 drop:0
       azure-function:
               flags:11
               stats: total:0 cur:0 done:0 drop:0
       google-cloud-function:
               flags:11
               stats: total:0 cur:0 done:0 drop:0
       alicloud-function:
               flags:11
               stats: total:20 cur:0 done:20 drop:0
```

- Enable debug output and turn on automation debug messages for about 30 minutes:

```
# diagnose debug enable
# diagnose debug application autod -1
```

```
__auto_generate_generic_curl_request()-358: Generating generic automation CURL request
for action (Send Log To Server).
__auto_generate_generic_curl_request()-406: Generic automation CURL request POST data
for action (Send Log To Server):
date=2019-05-30 time=16:44:43 logid="0100032002" type="event" subtype="system"
level="alert" vd="root" eventtime=1559259884209355090 tz="-0700" logdesc="Admin login
failed" sn="0" user="admin" ui="http(10.6.30.254)" method="http" srcip=10.6.30.254
```


Netskope Exhibit 1017

```
dstip=10.6.30.5 action="login" status="failed" reason="passwd_invalid"
msg="Administrator admin login failed from http(10.6.30.254) because of invalid
password"

__auto_generic_curl_request_close()-512: Generic CURL request response body from
http://172.16.200.44:
{
  "userId": 1,
  "id": 1,
  "title": "Test Response",
  "body": "ABCDEFGHIJKLMNOPQRSTUVWXYZ"
}
```

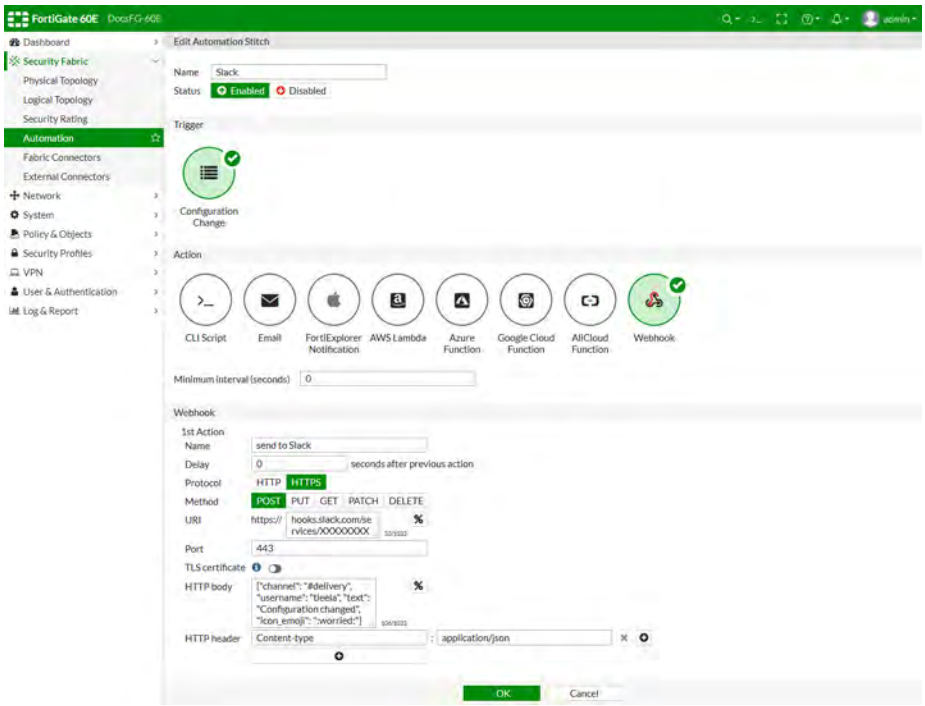
Slack integration webhook

A webhook can be created to post messages and notifications to Slack. For information about using incoming webhooks in Slack, see https://api.slack.com/incoming-webhooks.

In this example, a configuration change triggers the FortiGate to post a message to Slack.

To create a webhook automation stitch for Slack integration in the GUI:

1. Go to *Security Fabric* > *Automation*.
2. Click *Create New*.
3. Enter a name for the stitch.
4. Select the trigger *Configuration Change*.
5. Select *Webhook* and configure the settings:



6. Click *OK*.

To create a webhook automation stitch for Slack integration in the CLI:

1. Create the automation action:

```
config system automation-action
    edit "send to Slack"
        set action-type webhook
        set protocol https
        set uri "hooks.slack.com/services/XXXXXXXX"
        set http-body "{\"channel\": \"#delivery\", \"username\": \"tleela\", \"text\":
\"Configuration changed\", \"icon_emoji\": \":worried:\"}"
        set port 443
        set headers "Content-type:application/json"
    next
end
```

2. Create the automation trigger:

```
config system automation-trigger
    edit "config change"
        set event-type config-change
    next
end
```

3. Create the automation stitch:

```
config system automation-stitch
    edit "Slack"
        set trigger "config change"
        set action "send to Slack"
    next
end
```

Microsoft Teams integration webhook

A webhook can be created to post messages and notifications to Microsoft Teams.

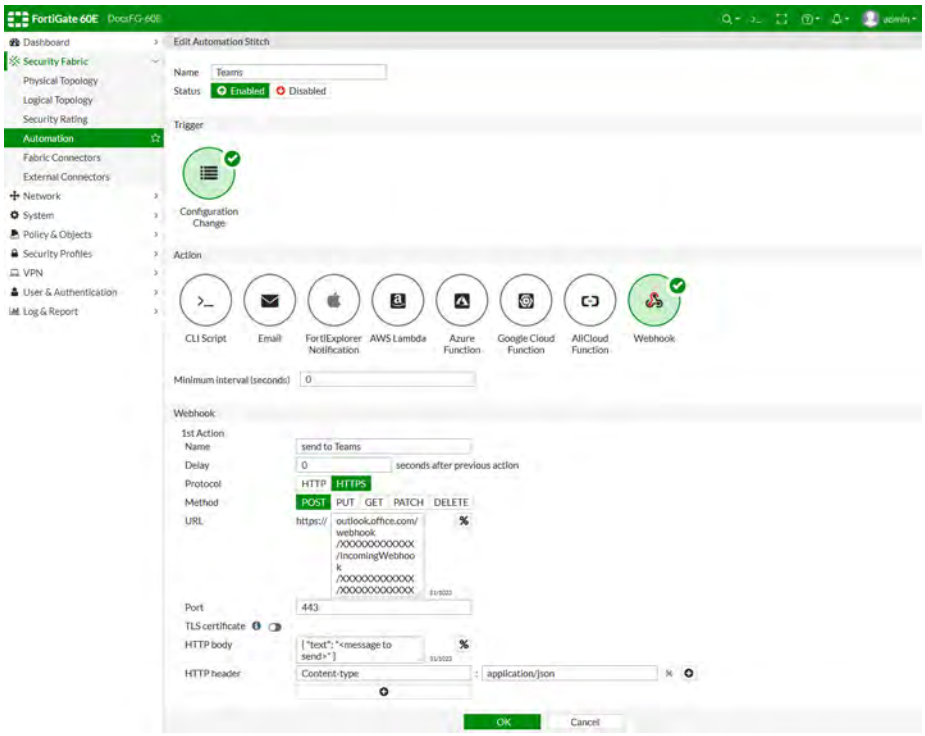
In this example, a configuration change triggers the FortiGate to post a message to Teams.

To create a webhook automation stitch for Teams integration in the GUI:

1. Create an incoming webhook in Teams. See https://docs.microsoft.com/en-us/microsoftteams/platform/webhooks-and-connectors/how-to/add-incoming-webhook for information.
2. Go to *Security Fabric* > *Automation*.
3. Click *Create New*.
4. Enter a name for the stitch.
5. Select the trigger *Configuration Change*.

Netskope Exhibit 1017

6. Select *Webhook* and configure the settings:



The *URI* is the URL from the incoming webhook created in Teams. The *HTTP body* can also contain log parameters.

7. Click *OK*.

To create a webhook automation stitch for Teams integration in the CLI:

1. Create an incoming webhook in Teams. See https://docs.microsoft.com/en-us/microsoftteams/platform/webhooks-and-connectors/how-to/add-incoming-webhook for information.
2. Create the automation action:

```
config system automation-action
    edit "send to Teams"
        set action-type webhook
        set protocol https
        set uri
"outlook.office.com/webhook/XXXXXXXXXXXX/IncomingWebhook/XXXXXXXXXXXX/XXXXXXXXXXXX"
        set http-body "{ \"text\": \"<message to send>\" }"
        set port 443
        set headers "Content-type:application/json"
    next
end
```

3. Create the automation trigger:

```
config system automation-trigger
    edit "Teams"
        set event-type config-change
    next
end
```

4. Create the automation stitch:

```
config system automation-stitch
    edit "Teams"
        set trigger "Teams"
        set action "send to Teams"
    next
end
```



For information about more advanced messages that can be configured and sent to the webhook, see https://docs.microsoft.com/en-us/microsoftteams/platform/webhooks-and-connectors/how-to/connectors-using

Execute a CLI script based on CPU and memory thresholds

Automation stitches can be created to run a CLI script and send an email message when CPU or memory usage exceeds specified thresholds.

In this example, two automation stitches are created that run a CLI script to collect debug information, and then email the results of the script to a specified email address when CPU usage threshold is exceeded or memory usage causes the FortiGate to enter conserve mode.



Automation stitches that use *High CPU* and *Conserve Mode* triggers can only be created in the CLI. Once create, they can be edited in the GUI.

To define CPU and memory usage thresholds:

```
config system global
    set cpu-use-threshold <percent>
    set memory-use-threshold-extreme <percent>
    set memory-use-threshold-green <percent>
    set memory-use-threshold-red <percent>
end
```

Where:

cpu-use-threshold	Threshold at which CPU usage is reported, in percent of total possible CPU utilization (default = 90).
memory-use-threshold-extreme	Threshold at which memory usage is considered extreme, and new sessions are dropped, in percent of total RAM (default = 95).
memory-use-threshold-green	Threshold at which memory usage forces the FortiGate to exit conserve mode, in percent of total RAM (default = 82).
memory-use-threshold-red	Threshold at which memory usage forces the FortiGate to enter conserve mode, in percent of total RAM (default = 88).

Configure the automation stitches

High CPU usage stitch

To create an automation stitch for high CPU usage:

1. Create an automation action to run a CLI script:

```
config system automation-action
    edit "high_cpu_debug"
        set action-type cli-script
        set required enable
        set script "diagnose debug cli 8
diagnose debug console timestamp enable
diagnose debug enable
diagnose debug crashlog read
get system performance status
get system session status
diagnose sys session full-stat
diagnose firewall iprope state
diagnose sys flash list
diagnose hardware sysinfo memory
diagnose hardware sysinfo slab
diagnose hardware sysinfo shm
diagnose hardware deviceinfo disk
get system arp
diagnose ip arp list
diagnose ip address list
get router info routing-table all
get router info kernel
diagnose ip rtcache list
diagnose sys top-summary
diagnose sys top 9 99"
    next
end
```

2. Create an automation action to send an email:

```
config system automation-action
    edit "auto_high_cpu_email"
        set action-type email
        set email-to "person@fortinet.com"
        set email-subject "CSF stitch alert: high_cpu"
        set message "%%results%%"
    next
end
```

3. Create an automation trigger:

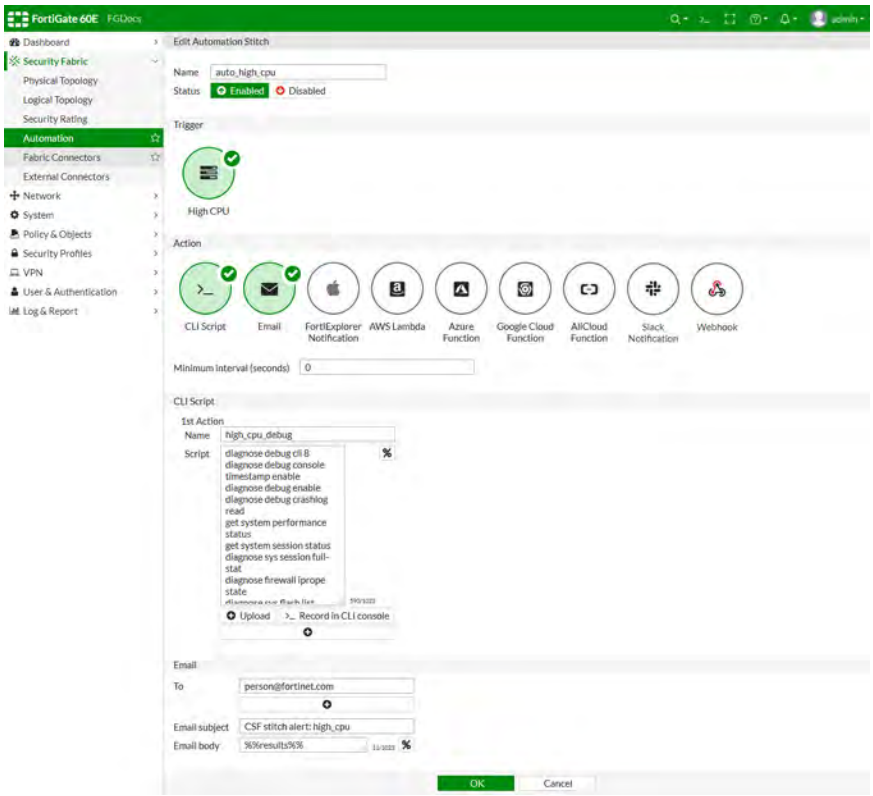
```
config system automation-trigger
    edit "auto_high_cpu"
        set event-type high-cpu
    next
end
```

4. Create an automation stitch:

```
config system automation-stitch
    edit "auto_high_cpu"
        set trigger "auto_high_cpu"
        set action "high_cpu_debug" "auto_high_cpu_email"
    next
end
```

To edit the automation stitch in the GUI:

1. Go to *Security Fabric* > *Automation*.
2. Double click the *auto_high_cpu* stitch.



3. Edit the stitch as required, then click *OK*.

High memory usage stitch

To create an automation stitch for high memory usage:

1. Create an automation action to run a CLI script:

```
config system automation-action
    edit "high_memory_debug"
        set action-type cli-script
        set required enable
        set script "diagnose debug cli 8
diagnose debug console timestamp enable
diagnose debug enable
```

```
diagnose debug crashlog read
get system performance status
get system session status
diagnose sys session full-stat
diagnose firewall iprope state
diagnose sys flash list
diagnose hardware sysinfo memory
diagnose hardware sysinfo slab
diagnose hardware sysinfo shm
diagnose hardware deviceinfo disk
get system arp
diagnose ip arp list
diagnose ip address list
get router info routing-table all
get router info kernel
diagnose ip rtcache list
diagnose sys top-summary
diagnose sys top 9 99"
    next
end
```

2. Create an automation action to send an email:

```
config system automation-action
    edit "auto_high_memory_email"
        set action-type email
        set email-to "person@fortinet.com"
        set email-subject "CSF stitch alert: high_memory"
        set message "%%results%%"
    next
end
```

3. Create an automation trigger:

```
config system automation-trigger
    edit "auto_high_memory"
        set event-type low-memory
    next
end
```

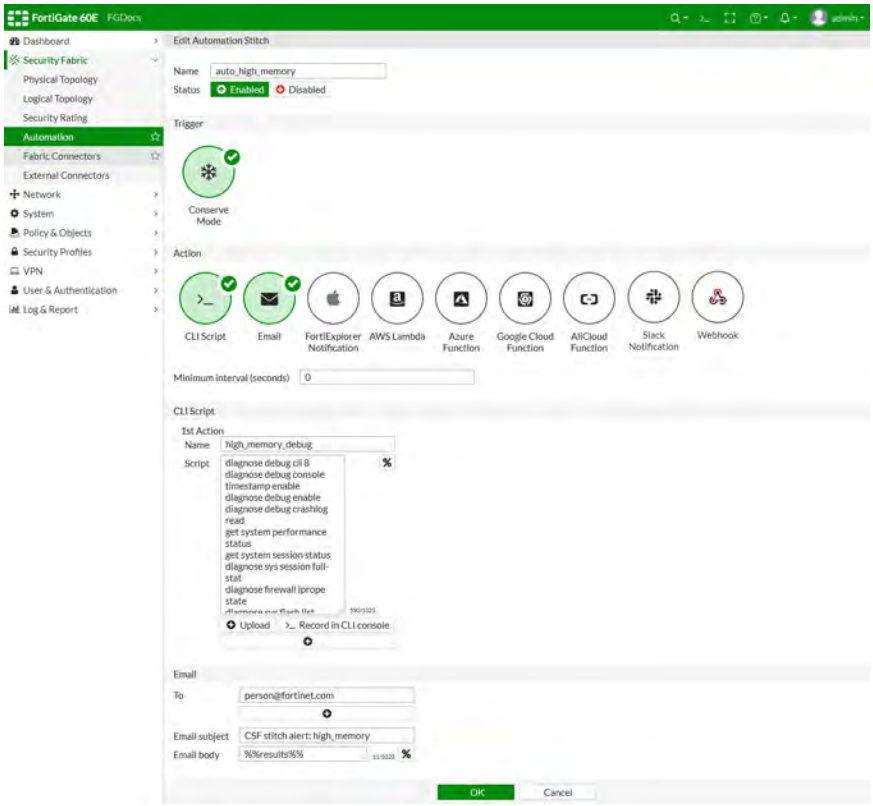
4. Create an automation stitch:

```
config system automation-stitch
    edit "auto_high_memory"
        set trigger "auto_high_memory"
        set action "high_memory_debug" "auto_high_memory_email"
    next
end
```


Netskope Exhibit 1017

To edit the automation stitch in the GUI:

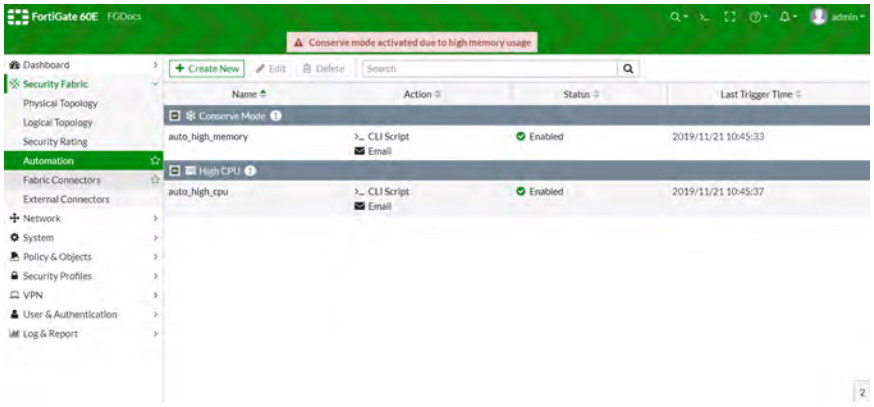
1. Go to *Security Fabric* > *Automation*.
2. Double click the *auto_high_memory* stitch.



3. Edit the stitch as required, then click *OK*.

Results

When FortiGate enters conserve mode due to the `memory-use-threshold-red` being exceeded, the GUI displays a notice, and the *auto_high_memory* automation stitch is triggered, causing the CLI script to run and the results of the script to be emailed to the specified address.



Here is an example of the email message:

```
CSF stitch alert: high_memory
noreply@notification.fortinet.net
Thu 11/21/2019 11:06 AM
James Li
FGT[FGVM16TM19000026] Automation Stitch:auto_high_memory is triggered.
########## script name: autod.47 ##########
========== #1, 2019-11-21 11:07:24 ==========
FGVM16TM19000026 $  diag deb cli 8
Debug messages will be on for 25 minutes.
FGVM16TM19000026 $  diag deb console timestamp enable
FGVM16TM19000026 $  diag deb enable
FGVM16TM19000026 $  diag deb crashlog read
1: 2019-08-08 11:35:25 the killed daemon is /bin/dhcpcd: status=0x0
2: 2019-08-08 17:52:47 the killed daemon is /bin/pyfcgid: status=0x0
3: 2019-08-23 11:32:31 from=license status=INVALID
4: 2019-08-23 11:32:32 from=license status=INVALID
5: 2019-11-21 09:53:31 from=license status=VALID
...
```

Troubleshooting

The following topics provide troubleshooting information for the Fortinet Security Fabric:

- Viewing a summary of all connected FortiGates in a Security Fabric on page 333
- Diagnosing automation stitches on page 336

Viewing a summary of all connected FortiGates in a Security Fabric

In downstream FortiGates, the `diagnose sys csf global` command shows a summary of all of the connected FortiGates in the Security Fabric.

To view a Security Fabric summary on a downstream FortiGate:

```
# diagnose sys csf global
Current vision:
[
  {
    "path":"FGVM01TM19000001",
    "mgmt_ip_str":"104.196.102.183",
    "mgmt_port":10403,
    "sync_mode":1,
    "saml_role":"identity-provider",
    "admin_port":443,
    "serial":"FGVM01TM19000001",
    "host_name":"admin-root",
    "firmware_version_major":6,
    "firmware_version_minor":2,
    "firmware_version_patch":0,
    "firmware_version_build":1010,
    "subtree_members":[
      {
```


Netskope Exhibit 1017

```
        "serial":"FGVM01TM19000002"
      },
      {
        "serial":"FGVM01TM19000003"
      },
      {
        "serial":"FGVM01TM19000004"
      },
      {
        "serial":"FGVM01TM19000005"
      }
    ]
  },
  {
    "path":"FGVM01TM19000001:FGVM01TM19000002",
    "mgmt_ip_str":"104.196.102.183",
    "mgmt_port":10423,
    "sync_mode":1,
    "saml_role":"service-provider",
    "admin_port":443,
    "serial":"FGVM01TM19000002",
    "host_name":"Branch_Office_01",
    "firmware_version_major":6,
    "firmware_version_minor":2,
    "firmware_version_patch":0,
    "firmware_version_build":1010,
    "upstream_intf":"Branch-HQ-A",
    "upstream_serial":"FGVM01TM19000001",
    "parent_serial":"FGVM01TM19000001",
    "parent_hostname":"admin-root",
    "upstream_status":"Authorized",
    "upstream_ip":22569994,
    "upstream_ip_str":"10.100.88.1",
    "subtree_members":[
    ],
    "is_discovered":true,
    "ip_str":"10.0.10.2",
    "downstream_intf":"To-HQ-A",
    "idx":1
  },
  {
    "path":"FGVM01TM19000001:FGVM01TM19000003",
    "mgmt_ip_str":"104.196.102.183",
    "mgmt_port":10407,
    "sync_mode":1,
    "saml_role":"service-provider",
    "admin_port":443,
    "serial":"FGVM01TM19000003",
    "host_name":"Enterprise_Second_Floor",
    "firmware_version_major":6,
    "firmware_version_minor":2,
    "firmware_version_patch":0,
    "firmware_version_build":1010,
    "upstream_intf":"port3",
    "upstream_serial":"FGVM01TM19000001",
    "parent_serial":"FGVM01TM19000001",
```


Netskope Exhibit 1017

```
    "parent_hostname":"admin-root",
    "upstream_status":"Authorized",
    "upstream_ip":22569994,
    "upstream_ip_str":"10.100.88.1",
    "subtree_members":[
    ],
    "is_discovered":true,
    "ip_str":"10.100.88.102",
    "downstream_intf":"port1",
    "idx":2
  },
  {
    "path":"FGVM01TM19000001:FGVM01TM19000004",
    "mgmt_ip_str":"104.196.102.183",
    "mgmt_port":10424,
    "sync_mode":1,
    "saml_role":"service-provider",
    "admin_port":443,
    "serial":"FGVM01TM19000004",
    "host_name":"Branch_Office_02",
    "firmware_version_major":6,
    "firmware_version_minor":2,
    "firmware_version_patch":0,
    "firmware_version_build":1010,
    "upstream_intf":"HQ-MPLS",
    "upstream_serial":"FGVM01TM19000001",
    "parent_serial":"FGVM01TM19000001",
    "parent_hostname":"admin-root",
    "upstream_status":"Authorized",
    "upstream_ip":22569994,
    "upstream_ip_str":"10.100.88.1",
    "subtree_members":[
    ],
    "is_discovered":true,
    "ip_str":"10.0.12.3",
    "downstream_intf":"To-HQ-MPLS",
    "idx":3
  },
  {
    "path":"FGVM01TM19000001:FGVM01TM19000005",
    "mgmt_ip_str":"104.196.102.183",
    "mgmt_port":10404,
    "sync_mode":1,
    "saml_role":"service-provider",
    "admin_port":443,
    "serial":"FGVM01TM19000005",
    "host_name":"Enterprise_First_Floor",
    "firmware_version_major":6,
    "firmware_version_minor":2,
    "firmware_version_patch":0,
    "firmware_version_build":1010,
    "upstream_intf":"port3",
    "upstream_serial":"FGVM01TM19000001",
    "parent_serial":"FGVM01TM19000001",
    "parent_hostname":"admin-root",
    "upstream_status":"Authorized",
```


Netskope Exhibit 1017

```
    "upstream_ip":22569994,
    "upstream_ip_str":"10.100.88.1",
    "subtree_members":[
    ],
    "is_discovered":true,
    "ip_str":"10.100.88.101",
    "downstream_intf":"port1",
    "idx":4
  }
]
```

Diagnosing automation stitches

Diagnose commands are available to:

- Test an automation stitch
- Enable or disable log dumping for automation stitches
- Display the settings of every automation stitch
- Display statistics on every automation stitch

To test an automation stitch:

```
diagnose automation test <automation-stitch-name>
```

Example:

```
# diagnose automation test HA-failover
automation test is done. stitch:HA-failover
```

To toggle log dumping:

```
diagnose test application autod 1
```

Examples:

```
# diagnose test application autod 1
autod log dumping is enabled
```

```
# diagnose test application autod 1
autod log dumping is disabled

autod logs dumping summary:
autod dumped total:7 logs, num of logids:4
```

To display the settings for all of the automation stitches:

```
diagnose test application autod 2
```

Example:

```
# diagnose test application autod 2
csf: enabled root:yes
total stitches activated: 3

stitch: Compromised-IP-Banned
        destinations: all
```

```
         trigger: Compromised-IP-Banned

       local hit: 0 relayed to: 0 relayed from: 0
        actions:
               Compromised-IP-Banned_ban-ip type:ban-ip interval:0

stitch: HA-failover
        destinations: HA-failover_ha-cluster_25;
        trigger: HA-failover

       local hit: 0 relayed to: 0 relayed from: 0
        actions:
               HA-failover_email type:email interval:0
               subject: HA Failover
               mailto:admin@example.com;

stitch: rebooot
       destinations: all
       trigger: reboot

       local hit: 0 relayed to: 0 relayed from: 0
       actions:
               action1 type:alicloud-function interval:0
                       delay:1 required:yes
                       Account ID: id
                       Region: region
                       Function domain: fc.aliyuncs.com
                       Version: versoin
                       Service name: serv
                       Function name: funcy
                       headers:
```

To display statistic on all of the automation stitches:

```
diagnose test application autod 3
```

Example:

```
stitch: Compromised-IP-Banned
       local hit: 0 relayed to: 0 relayed from: 0
       last trigger:Wed Dec 31 20:00:00 1969
       last relay:Wed Dec 31 20:00:00 1969
       actions:
               Compromised-IP-Banned_ban-ip:
                       done: 1 relayed to: 0 relayed from: 0
                       last trigger:Wed Dec 31 20:00:00 1969
                       last relay:

stitch: HA-failover
       local hit: 0 relayed to: 0 relayed from: 0
       last trigger:Thu May 24 11:35:22 2018
       last relay:Thu May 24 11:35:22 2018
       actions:
               HA-failover_email:
                       done: 1 relayed to: 1 relayed from: 1
                       last trigger:Thu May 24 11:35:22 2018
                       last relay:Thu May 24 11:35:22 2018
```


Netskope Exhibit 1017

```
stitch: rebooot
       local hit: 2 relayed to: 1 relayed from: 1
       last trigger:Fri May  3 13:30:56 2019
       last relay:Fri May  3 13:30:23 2019
       actions:
              action1
                     done: 1 relayed to: 0 relayed from: 0
                     last trigger:Fri May  3 13:30:56 2019
                     last relay:

logid2stitch mapping:
id:20103  local hit: 0 relayed to: 0 relayed from: 0
       License Expiry
       lambada

id:32138  local hit: 2 relayed to: 1 relayed from: 1
       Compromised-IP-Banned
       HA-failover
       rebooot

action run cfg&stats:
total:2 cur:0 done:1 drop:1
       email:
              flags:10
              stats: total:1 cur:0 done:1 drop:0
       ios-notification:
              flags:1
              stats: total:0 cur:0 done:0 drop:0
       alert:
              flags:0
              stats: total:0 cur:0 done:0 drop:0
       disable-ssid:
              flags:7
              stats: total:0 cur:0 done:0 drop:0
       quarantine:
              flags:7
              stats: total:0 cur:0 done:0 drop:0
       quarantine-forticlient:
              flags:4
              stats: total:0 cur:0 done:0 drop:0
       quarantine-nsx:
              flags:4
              stats: total:0 cur:0 done:0 drop:0
       ban-ip:
              flags:7
              stats: total:0 cur:0 done:0 drop:0
       aws-lambda:
              flags:11
              stats: total:0 cur:0 done:0 drop:0
       webhook:
              flags:11
              stats: total:0 cur:0 done:0 drop:0
       cli-script:
              flags:10
              stats: total:0 cur:0 done:0 drop:0
```


Netskope Exhibit 1017

```
azure-function:
        flags:11
        stats: total:1 cur:0 done:0 drop:1
google-cloud-function:
        flags:11
        stats: total:0 cur:0 done:0 drop:0
alicloud-function:
        flags:11
        stats: total:0 cur:0 done:0 drop:0
```


Netskope Exhibit 1017

Network

The following topics provide information about network settings:

- Interfaces on page 340
- DNS on page 372
- Explicit and transparent proxies on page 383
- SD-WAN on page 444
- DHCP server on page 428
- Static routes on page 433
- Policy routes on page 434
- RIP on page 434
- OSPF on page 434
- BGP on page 434
- Direct IP support for LTE/4G on page 435
- LLDP reception on page 438
- Route leaking between VRFs on page 441

Interfaces

Physical and virtual interfaces allow traffic to flow between internal networks, and between the internet and internal networks. FortiGate has options for setting up interfaces and groups of subnetworks that can scale as your organization grows. You can create and edit VLAN, EMAC-VLAN, switch interface, zones, and so on.

The following topics provide information about interfaces:

- Interface settings on page 341
- Aggregation and redundancy on page 344
- VLANs on page 346
- Enhanced MAC VLANs on page 352
- Inter-VDOM routing on page 355
- Software switch on page 360
- Zone on page 361
- Virtual Wire Pair on page 363
- Virtual switch support for FortiGate 300E series on page 365
- Failure detection for aggregate and redundant interfaces on page 367
- VLAN inside VXLAN on page 367
- Virtual Wire Pair with VXLAN on page 370

Netskope Exhibit 1017

Interface settings

Administrators can configure both physical and virtual FortiGate interfaces in *Network* > *Interfaces*. There are different options for configuring interfaces when FortiGate is in NAT mode or transparent mode.

To configure an interface in the GUI:

1. Go to *Network* > *Interfaces*.
2. Click *Create New* > *Interface*.
3. Configure the interface fields:

Interface Name	Physical interface names cannot be changed.
Alias	Enter an alternate name for a physical interface on the FortiGate unit. This field appears when you edit an existing physical interface. The alias does not appear in logs. The maximum length of the alias is 25 characters.
Type	The configuration type for the interface, such as VLAN or Software Switch.
Link Status	Indicates whether the interface is connected to a network or not (link status is up or down). This field is available when you edit an existing physical interface.
Interface	This field is available when *Type* is set to *VLAN*. Select the name of the physical interface that you want to add a VLAN interface to. Once created, the VLAN interface is listed below its physical interface in the *Interface* list. You cannot change the physical interface of a VLAN interface except when you add a new VLAN interface.
VLAN ID	This field is available when *Type* is set to *VLAN*. Enter the VLAN ID. The VLAN ID can be any number between 1 and 4094 and must match the VLAN ID added by the IEEE 802.1Q-compliant router or switch that is connected to the VLAN subinterface. The VLAN ID cannot be edited after the interface is added.
Virtual Domain	Select the virtual domain to add the interface to. Only administrator accounts with the *super_admin* profile can change the *Virtual Domain*.
Role	Set the role setting for the interface. Different settings will be shown or hidden when editing an interface depending on the role. • *LAN*: Used to connected to a local network of endpoints. It is default role for new interfaces. • *WAN*: Used to connected to the internet. When WAN is selected, the *Estimated bandwidth* setting is available, and the following settings are not: *DHCP server*, *Create address object matching subnet*, *Device detection*, *Security mode*, *One-arm sniffer*, *Dedicate to extension/fortiap modes*, and *Admission Control*.and will show Estimated Bandwidth settings. • *DMZ*: Used to connected to the DMZ. When selected, *DHCP server* and

Netskope Exhibit 1017

	Security mode are not available. • *Undefined*: The interface has no specific role. When selected, *Create address object matching subnet* is not available.
Interface Members	This section can has different formats depending on the *Type*: *Software Switch*: This field is read-only, and shows the interfaces that belong to the virtual interface of the software switch. *802.3ad Aggregate* or *Redundant Interface*: This field includes the available and selected interface lists.
Addressing mode	Select the addressing mode for the interface. • *Manual*: Add an IP address and netmask for the interface. If IPv6 configuration is enabled,you can add both an IPv4 and an IPv6 address. • *DHCP*: Get the interface IP address and other network settings from a DHCP server. • *PPPoE*: Get the interface IP address and other network settings from a PPPoE server. This option is only available on the low-end FortiGate models.
IP/Netmask	If *Addressing Mode* is set to *Manual*, enter an IPv4 address and subnet mask for the interface. FortiGate interfaces cannot have multiple IP addresses on the same subnet.
IPv6 Address/Prefix	If *Addressing Mode* is set to *Manual* and IPv6 support is enabled, enter an IPv6 address and subnet mask for the interface. A single interface can have an IPv4 address, IPv6 address, or both.
Create address object matching subnet	This option is available when *Role* is set to *LAN* or *DMZ*. Enable this option to automatically create an address object that matches the interface subnet.
Secondary IP Address	Add additional IPv4 addresses to this interface.
IPv4 Administrative Access	Select the types of administrative access permitted for IPv4 connections to this interface. See Configure administrative access to interfaces on page 343.
IPv6 Administrative Access	Select the types of administrative access permitted for IPv6 connections to this interface. See Configure administrative access to interfaces on page 343.
DHCP Server	Select to enable a DHCP server for the interface.
Device Detection	Enable/disable passively gathering device identity information about the devices on the network that are connected to this interface.
Security Mode	Enable/disable captive portal authentication for this interface. After enabling captive portal authentication, you can configure the authentication portal, user and group access, custom portal messages, exempt sources and destinations/services, and redirect after captive portal.
Outbound shaping profile	Enable/disable traffic shaping on the interface. This allows you to enforce bandwidth limits on individual interfaces.
Comments	Enter a description of the interface of up to 255 characters.
Status	Enable/disable the interface.

- *Enabled*: The interface is active and can accept network traffic.
- *Disabled*: The interface is not active and cannot accept traffic.

4. Click *OK*.

To configure an interface in the CLI:

```
config system interface
    edit "<Interface_Name>"
        set vdom "<VDOM_Name>"
        set mode static/dhcp/pppoe
        set ip <IP_address> <netmask>
        set security-mode {none | captive-portal}
        set egress-shaping-profile <Profile_name>
        set device-identification {enable | disable}
        set allowaccess ping https ssh http
        set secondary-IP enable
        config secondaryip
            edit 1
                set ip 9.1.1.2 255.255.255.0
                set allowaccess ping https ssh snmp http
            next
        end
    next
end
```

Configure administrative access to interfaces

You can configure the protocols that administrators can use to access interfaces on the FortiGate. This helps secure access to the FortiGate by restricting access to a limited number of protocols. It helps prevent users from accessing interfaces that you don't want them to access, such as public-facing ports.

As a best practice, you should configure administrative access when you're setting the IP address for a port.

To configure administrative access to interfaces in the GUI:

1. Go to *Network* > *Interfaces*.
2. Create or edit an interface.
3. In the *Administrative Access* section, select which protocols to enable for *IPv4* and *IPv6 Administrative Access*.

HTTPS	Allow secure HTTPS connections to the FortiGate GUI through this interface. If configured, this option is enabled automatically.
HTTP	Allow HTTP connections to the FortiGate GUI through this interface. This option can only be enabled if HTTPS is already enabled.
PING	The interface responds to pings. Use this setting to verify your installation and for testing.
FMG-Access	Allow FortiManager authorization automatically during the communication exchanges between FortiManager and FortiGate devices.
SSH	Allow SSH connections to the CLI through this interface.

SNMP	Allow a remote SNMP manager to request SNMP information by connecting to this interface.
FTM	Allow FortiToken Mobile Push (FTM) access.
RADIUS Accounting	Allow RADIUS accounting information on this interface.
Security Fabric Connection	Allow Security Fabric access. This enables FortiTelemetry and CAPWAP.

Aggregation and redundancy

Link aggregation (IEEE 802.3ad) enables you to bind two or more physical interfaces together to form an aggregated (combined) link. This new link has the bandwidth of all the links combined. If a link in the group fails, traffic is transferred automatically to the remaining interfaces. The only noticeable effect is reduced bandwidth.

This feature is similar to redundant interfaces. The major difference is a redundant interface group only uses one link at a time, where an aggregate link group uses the total bandwidth of the functioning links in the group, up to eight (or more).

An interface is available to be an aggregate interface if:

- It is a physical interface and not a VLAN interface or subinterface.
- It is not already part of an aggregate or redundant interface.
- It is in the same VDOM as the aggregated interface. Aggregate ports cannot span multiple VDOMs.
- It does not have an IP address and is not configured for DHCP or PPPoE.
- It is not referenced in any security policy, VIP, IP Pool, or multicast policy.
- It is not an HA heartbeat interface.
- It is not one of the FortiGate-5000 series backplane interfaces.

When an interface is included in an aggregate interface, it is not listed on the *Network* > *Interfaces* page. Interfaces still appear in the CLI although configuration for those interfaces do not take affect. You cannot configure the interface individually and it is not available for inclusion in security policies, VIPs, IP pools, or routing.

Sample configuration

This example creates an aggregate interface on a FortiGate-140D POE using ports 3-5 with an internal IP address of 10.1.1.123, as well as the administrative access to HTTPS and SSH.

To create an aggregate interface using the GUI:

1. Go to *Network* > *Interfaces* and select *Create New* > *Interface*.
2. For *Interface Name*, enter *Aggregate*.
3. For the Type, select *802.3ad Aggregate*.
4. In the physical *Interface Members*, click to add interfaces and select ports 4, 5, and 6.
5. For *Addressing mode*, select *Manual*.
6. For the IP address for the port, enter *10.1.1.123/24*.
7. For *Administrative Access*, select *HTTPS* and *SSH*.
8. Select *OK*.

Netskope Exhibit 1017

To create an aggregate interface using the CLI:

```
FG140P3G15800330 (aggregate) # show
config system interface
    edit "aggregate"
        set vdom "root"
        set ip 10.1.1.123 255.255.255.0
        set allowaccess ping https ssh snmp http fgfm radius-acct capwap ftm
        set type aggregate
        set member "port3" "port4" "port5"
        set device-identification enable
        set lldp-transmission enable
        set fortiheartbeat enable
        set role lan
        set snmp-index 45
    next
end
```

Redundancy

In a redundant interface, traffic only goes over one interface at any time. This differs from an aggregated interface where traffic goes over all interfaces for increased bandwidth. This difference means redundant interfaces can have more robust configurations with fewer possible points of failure. This is important in a fully-meshed HA configuration.

An interface is available to be in a redundant interface if:

- It is a physical interface and not a VLAN interface.
- It is not already part of an aggregated or redundant interface.
- It is in the same VDOM as the redundant interface.
- It does not have an IP address and is not configured for DHCP or PPPoE.
- It has no DHCP server or relay configured on it.
- It does not have any VLAN subinterfaces.
- It is not referenced in any security policy, VIP, or multicast policy.
- It is not monitored by HA.
- It is not one of the FortiGate-5000 series backplane interfaces.

When an interface is included in a redundant interface, it is not listed on the *Network* > *Interfaces* page. You cannot configure the interface individually and it is not available for inclusion in security policies, VIPs, or routing.

Sample configuration

To create a redundant interface using the GUI:

1. Go to *Network* > *Interfaces* and select *Create New* > *Interface*.
2. For *Interface Name*, enter *Redundant*.
3. For the Type, select *Redundant Interface*.
4. In the physical *Interface Members*, click to add interfaces and select ports 4, 5, and 6.
5. For *Addressing mode*, select *Manual*.
6. For the IP address for the port, enter *10.13.101.100/24*.
7. For *Administrative Access*, select *HTTPS* and *SSH*.
8. Select *OK*.

Netskope Exhibit 1017

To create a redundant interface using the CLI:

```
config system interface
    edit "red"
        set vdom "root"
        set ip 10.13.101.100 255.255.255.0
        set allowaccess https http
        set type redundant
        set member "port4" "port5" "port6"
        set device-identification enable
        set role lan
        set snmp-index 9
    next
end
```

VLANs

Virtual Local Area Networks (VLANs) multiply the capabilities of your FortiGate unit and can also provide added network security. VLANs use ID tags to logically separate devices on a network into smaller broadcast domains. These smaller domains forward packets only to devices that are part of that VLAN domain. This reduces traffic and increases network security.

VLANs in NAT mode

In NAT mode, the FortiGate unit functions as a layer-3 device. In this mode, the FortiGate unit controls the flow of packets between VLANs and can also remove VLAN tags from incoming VLAN packets. The FortiGate unit can also forward untagged packets to other networks such as the Internet.

In NAT mode, the FortiGate unit supports VLAN trunk links with IEEE 802.1Q-compliant switches or routers. The trunk link transports VLAN-tagged packets between physical subnets or networks. When you add VLAN subinterfaces to the FortiGate's physical interfaces, the VLANs have IDs that match the VLAN IDs of packets on the trunk link. The FortiGate unit directs packets with VLAN IDs to subinterfaces with matching IDs.

You can define VLAN subinterfaces on all FortiGate physical interfaces. However, if multiple virtual domains are configured on the FortiGate unit, you only have access to the physical interfaces on your virtual domain. The FortiGate unit can tag packets leaving on a VLAN subinterface. It can also remove VLAN tags from incoming packets and add a different VLAN tag to outgoing packets.

Normally in VLAN configurations, the FortiGate unit's internal interface is connected to a VLAN trunk, and the external interface connects to an Internet router that is not configured for VLANs. In this configuration, the FortiGate unit can apply different policies for traffic on each VLAN interface connected to the internal interface, which results in less network traffic and better security.

Sample topology

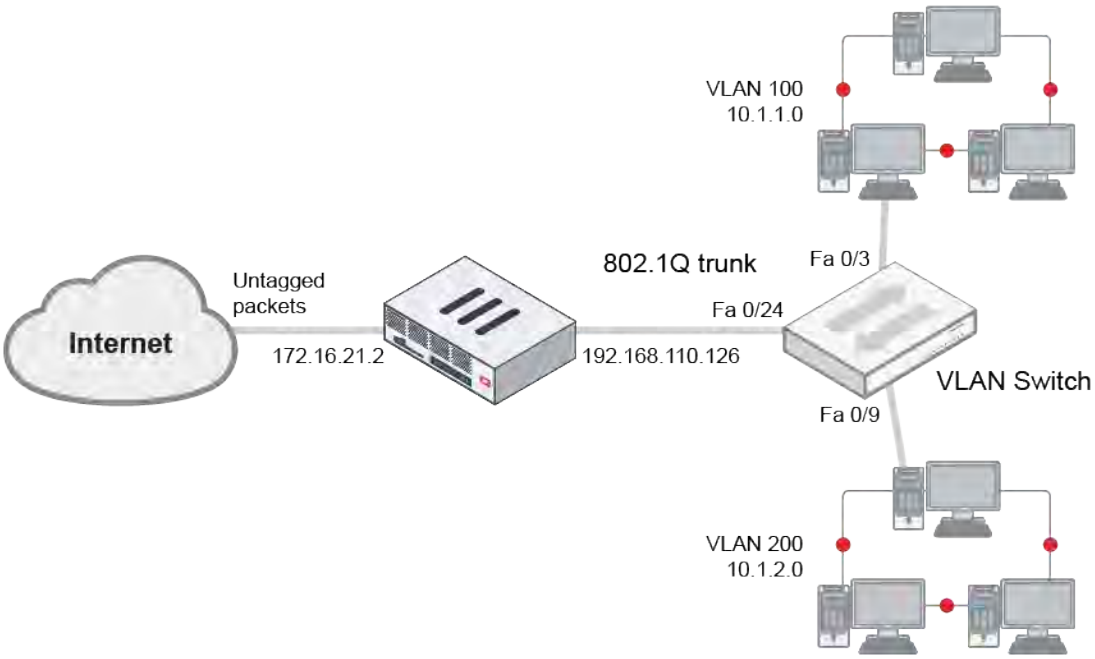
In this example, two different internal VLAN networks share one interface on the FortiGate unit and share the connection to the Internet. This example shows that two networks can have separate traffic streams while sharing a single interface. This configuration can apply to two departments in a single company or to different companies.

There are two different internal network VLANs in this example. VLAN_100 is on the 10.1.1.0/255.255.255.0 subnet, and VLAN_200 is on the 10.1.2.0/255.255.255.0 subnet. These VLANs are connected to the VLAN switch.

Netskope Exhibit 1017

The FortiGate internal interface connects to the VLAN switch through an 802.1Q trunk. The internal interface has an IP address of 192.168.110.126 and is configured with two VLAN subinterfaces (VLAN_100 and VLAN_200). The external interface has an IP address of 172.16.21.2 and connects to the Internet. The external interface has no VLAN subinterfaces.

When the VLAN switch receives packets from VLAN_100 and VLAN_200, it applies VLAN ID tags and forwards the packets of each VLAN both to local ports and to the FortiGate unit across the trunk link. The FortiGate unit has policies that allow traffic to flow between the VLANs, and from the VLANs to the external network.



Sample configuration

In this example, both the FortiGate unit and the Cisco 2950 switch are installed and connected and basic configuration has been completed. On the switch, you need access to the CLI to enter commands. No VDOMs are enabled in this example.

General configuration steps include:

1. Configure the external interface.
2. Add two VLAN subinterfaces to the internal network interface.
3. Add firewall addresses and address ranges for the internal and external networks.
4. Add security policies to allow:
   - the VLAN networks to access each other.
   - the VLAN networks to access the external network.

To configure the external interface:

```
config system interface
    edit external
        set mode static
        set ip 172.16.21.2 255.255.255.0
end
```

To add VLAN subinterfaces:

```
config system interface
    edit VLAN_100
        set vdom root
        set interface internal
        set type vlan
        set vlanid 100
        set mode static
        set ip 10.1.1.1 255.255.255.0
        set allowaccess https ping
    next
    edit VLAN_200
        set vdom root
        set interface internal
        set type vlan
        set vlanid 200
        set mode static
        set ip 10.1.2.1 255.255.255.0
        set allowaccess https ping
end
```

To add the firewall addresses:

```
config firewall address
    edit VLAN_100_Net
        set type ipmask
        set subnet 10.1.1.0 255.255.255.0
    next
    edit VLAN_200_Net
        set type ipmask
        set subnet 10.1.2.0 255.255.255.0
end
```

To add security policies:

Policies 1 and 2 do not need NAT enabled, but policies 3 and 4 do need NAT enabled.

```
config firewall policy
    edit 1
        set srcintf VLAN_100
        set srcaddr VLAN_100_Net
        set dstintf VLAN_200
        set dstaddr VLAN_200_Net
        set schedule always
        set service ALL
        set action accept
        set nat disable
        set status enable
    next
    edit 2
        set srcintf VLAN_200
        set srcaddr VLAN_200_Net
        set dstintf VLAN_100
        set dstaddr VLAN_100_Net
        set schedule always
```


Netskope Exhibit 1017

```
        set service ALL
        set action accept
        set nat disable
        set status enable
    next
    edit 3
        set srcintf VLAN_100
        set srcaddr VLAN_100_Net
        set dstintf external
        set dstaddr all
        set schedule always
        set service ALL
        set action accept
        set nat enable
        set status enable
    next
    edit 4
        set srcintf VLAN_200
        set srcaddr VLAN_200_Net
        set dstintf external
        set dstaddr all
        set schedule always
        set service ALL
        set action accept
        set nat enable
        set status enable
    end
```

VLANs in transparent mode

In transparent mode, the FortiGate unit behaves like a layer-2 bridge but can still provide services such as antivirus scanning, web filtering, spam filtering, and intrusion protection to traffic. Some limitations of transparent mode is that you cannot use SSL VPN, PPTP/L2TP VPN, DHCP server, or easily perform NAT on traffic. The limits in transparent mode apply to IEEE 802.1Q VLAN trunks passing through the unit.

You can insert the FortiGate unit operating in transparent mode into the VLAN trunk without making changes to your network. In a typical configuration, the FortiGate unit internal interface accepts VLAN packets on a VLAN trunk from a VLAN switch or router connected to internal network VLANs. The FortiGate external interface forwards VLAN-tagged packets through another VLAN trunk to an external VLAN switch or router and on to external networks such as the Internet. You can configure the unit to apply different policies for traffic on each VLAN in the trunk.

To pass VLAN traffic through the FortiGate unit, you add two VLAN subinterfaces with the same VLAN ID, one to the internal interface and the other to the external interface. You then create a security policy to permit packets to flow from the internal VLAN interface to the external VLAN interface. If required, create another security policy to permit packets to flow from the external VLAN interface to the internal VLAN interface. Typically in transparent mode, you do not permit packets to move between different VLANs. Network protection features such as spam filtering, web filtering, and anti-virus scanning, are applied through the UTM profiles specified in each security policy, enabling very detailed control over traffic.

When the FortiGate unit receives a VLAN-tagged packet on a physical interface, it directs the packet to the VLAN subinterface with the matching VLAN ID. The VLAN tag is removed from the packet and the FortiGate unit then applies security policies using the same method it uses for non-VLAN packets. If the packet exits the FortiGate unit through a VLAN subinterface, the VLAN ID for that subinterface is added to the packet and the packet is sent to the corresponding physical interface.

Netskope Exhibit 1017

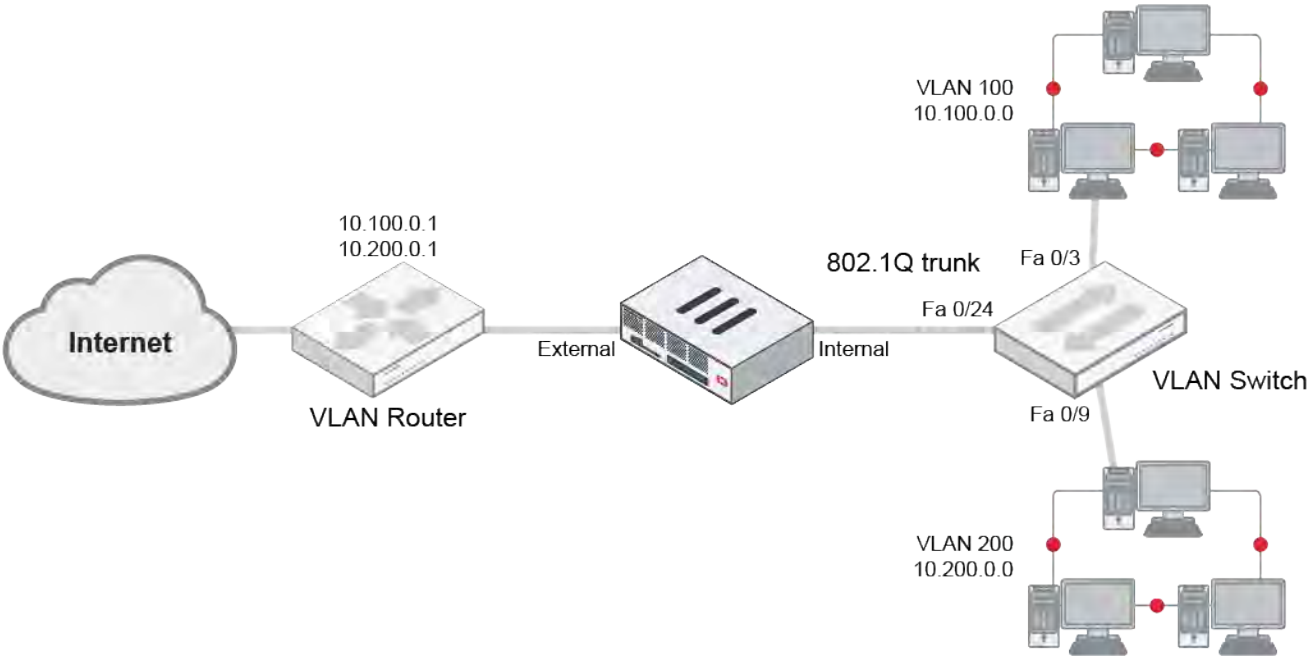
Sample topology

In this example, the FortiGate unit is operating in transparent mode and is configured with two VLANs: one with an ID of 100 and the other with ID 200. The internal and external physical interfaces each have two VLAN subinterfaces, one for VLAN_100 and one for VLAN_200.

The IP range for the internal VLAN_100 network is 10.100.0.0/255.255.0.0, and for the internal VLAN_200 network is 10.200.0.0/255.255.0.0.

The internal networks are connected to a Cisco 2950 VLAN switch which combines traffic from the two VLANs onto one in the FortiGate unit's internal interface. The VLAN traffic leaves the FortiGate unit on the external network interface, goes on to the VLAN switch, and on to the Internet. When the FortiGate units receives a tagged packet, it directs it from the incoming VLAN subinterface to the outgoing VLAN subinterface for that VLAN.

In this example, we create a VLAN subinterface on the internal interface and another one on the external interface, both with the same VLAN ID. Then we create security policies that allow packets to travel between the VLAN_100_int interface and the VLAN_100_ext interface. Two policies are required: one for each direction of traffic. The same is required between the VLAN_200_int interface and the VLAN_200_ext interface, for a total of four security policies.



Sample configuration

There are two main steps to configure your FortiGate unit to work with VLANs in transparent mode:

1. Add VLAN subinterfaces.
2. Add security policies.

You can also configure the protection profiles that manage antivirus scanning, web filtering, and spam filtering.

To add VLAN subinterfaces:

```
config system interface
    edit VLAN_100_int
        set type vlan
```

```
        set interface internal
        set vlanid 100
    next
    edit VLAN_100_ext
        set type vlan
        set interface external
        set vlanid 100
    next
    edit VLAN_200_int
        set type vlan
        set interface internal
        set vlanid 200
    next
    edit VLAN_200_ext
        set type vlan
        set interface external
        set vlanid 200
    end
```

To add security policies:

```
config firewall policy
    edit 1
        set srcintf VLAN_100_int
        set srcaddr all
        set dstintf VLAN_100_ext
        set dstaddr all
        set action accept
        set schedule always
        set service ALL
    next
    edit 2
        set srcintf VLAN_100_ext
        set srcaddr all
        set dstintf VLAN_100_int
        set dstaddr all
        set action accept
        set schedule always
        set service ALL
    next
    edit 3
        set srcintf VLAN_200_int
        set srcaddr all
        set dstintf VLAN_200_ext
        set dstaddr all
        set action accept
        set schedule always
        set service ALL
    next
    edit 4
        set srcintf VLAN_200_ext
        set srcaddr all
        set dstintf VLAN_200_int
        set dstaddr all
        set action accept
        set schedule always
```


Netskope Exhibit 1017

```
        set service ALL
end
```

Enhanced MAC VLANs

The Media Access Control (MAC) Virtual Local Area Network (VLAN) feature in Linux allows you to configure multiple virtual interfaces with different MAC addresses (and therefore different IP addresses) on a physical interface.

FortiGate implements an enhanced MAC VLAN consisting of a MAC VLAN with bridge functionality. Because each MAC VLAN has a unique MAC address, virtual IP addresses (VIPs) and IP pools are supported, and you can disable Source Network Address Translation (SNAT) in policies.

MAC VLAN cannot be used in a transparent mode virtual domain (VDOM). In a transparent mode VDOM, a packet leaves an interface with the MAC address of the original source instead of the interface's MAC address. FortiGate implements an enhanced version of MAC VLAN where it adds a MAC table in the MAC VLAN which learns the MAC addresses when traffic passes through.

If you configure a VLAN ID for an enhanced MAC VLAN, it won't join the switch of the underlying interface. When a packet is sent to this interface, a VLAN tag is inserted in the packet and the packet is sent to the driver of the underlying interface. When the underlying interface receives a packet, if the VLAN ID doesn't match, it won't deliver the packet to this enhanced MAC VLAN interface.



When using a VLAN ID, the ID and the underlying interface must be a unique pair, even if the belong to different VDOMs. This is because the underlying, physical interface uses the VLAN ID as the identifier to dispatch traffic among the VLAN and enhanced MAC VLAN interfaces.

If you use an interface in an enhanced MAC VLAN, do not use it for other purposes such as a management interface, HA heartbeat interface, or in Transparent VDOMs.

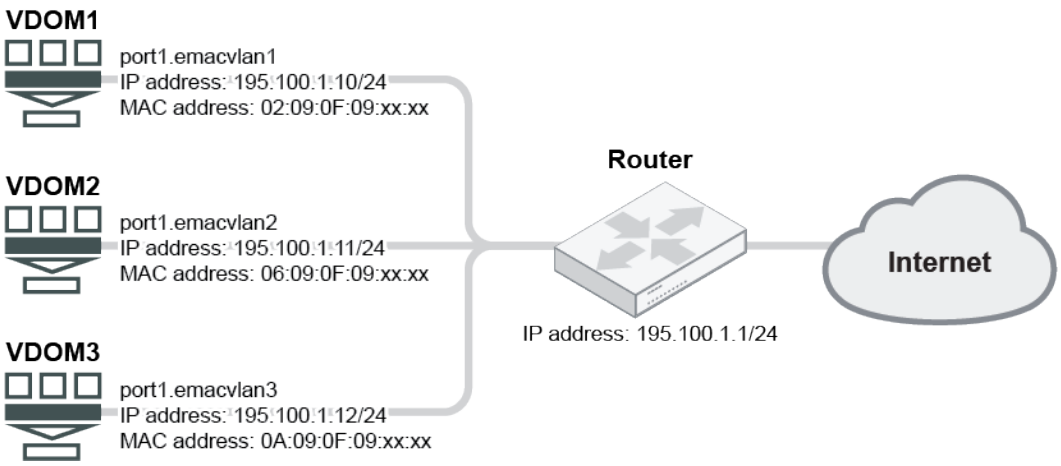
If a physical interface is used by an EMAC VLAN interface, you cannot use it in a Virtual Wire Pair.

In high availability (HA) configurations, enhanced MAC VLAN is treated as a physical interface. It's assigned a unique physical interface ID and the MAC table is synchronized with the secondary devices in the same HA cluster.

Example 1: Enhanced MAC VLAN configuration for multiple VDOMs that use the same interface or VLAN

In this example, a FortiGate is connected, through port 1 to a router that's connected to the Internet. Three VDOMs share the same interface (port 1) which connects to the same router that's connected to the Internet. Three enhanced MAC VLAN interfaces are configured on port 1 for the three VDOMs. The enhanced MAC VLAN interfaces are in the same IP subnet segment and each have unique MAC addresses.

The underlying interface (port 1) can be a physical interface, an aggregate interface, or a VLAN interface on a physical or aggregate interface.



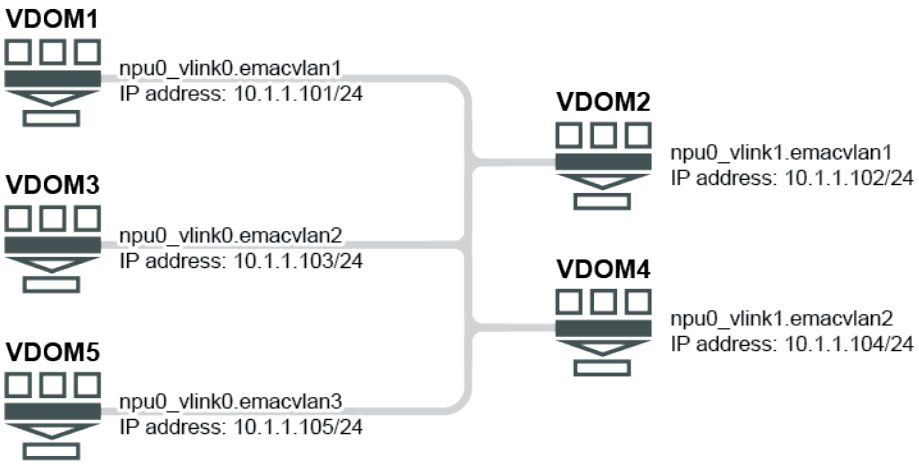
To configure enhanced MAC VLAN for this example in the CLI:

```
config system interface
    edit port1.emacvlan1
        set vdom VDOM1
        set type emac-vlan
        set interface port1
    next
    edit port 1.emacvlan2
        set vdom VDOM2
        set type emac-vlan
        set interface port1
    next
    edit port1.emacvlan3
        set vdom VDOM3
        set type emac-vlan
        set interface port1
    next
end
```

Example 2: Enhanced MAC VLAN configuration for shared VDOM links among multiple VDOMs

In this example, multiple VDOMs can connect to each other using enhanced MAC VLAN on network processing unit (NPU) virtual link (Vlink) interfaces.

FortiGate VDOM links (NPU-Vlink) are designed to be peer-to-peer connections and VLAN interfaces on NPU Vlink ports use the same MAC address. Connecting more than two VDOMs using NPU Vlinks and VLAN interfaces is not recommended.



To configure enhanced MAC VLAN for this example in the CLI:

```
config system interface
    edit npu0_vlink0.emacvlan1
        set vdom VDOM1
        set type emac-vlan
        set interface npu0_vlink0
    next
    edit npu0_vlink0.emacvlan2
        set vdom VDOM3
        set type emac-vlan
        set interface npu0_vlink0
    next
    edit npu0_vlink1.emacvlan1
        set vdom VDOM2
        set type emac-vlan
        set interface npu0_vlink1
    next
end
```

Example 3: Enhanced MAC VLAN configuration for unique MAC addresses for each VLAN interface on the same physical port

Some networks require a unique MAC address for each VLAN interface when the VLAN interfaces share the same physical port. In this case, the enhanced MAC VLAN interface is used the same way as normal VLAN interfaces.

To configure this, use the `set vlanid` command for the VLAN tag. The VLAN ID and interface must be a unique pair, even if they belong to different VDOMs.

To configure enhanced MAC VLAN:

```
config system interface
    edit <interface-name>
        set type emac-vlan
        set vlanid <VLAN-ID>
        set interface <physical-interface>
    next
end
```

Inter-VDOM routing

In the past, virtual domains (VDOMs) were separate from each other and there was no internal communication. Any communication between VDOMs involved traffic leaving on a physical interface belonging to one VDOM and re-entering the FortiGate unit on another physical interface belonging to another VDOM to be inspected by firewall policies in both directions.

Inter-VDOM routing changes this. With VDOM links, VDOMs can communicate internally without using additional physical interfaces.

Inter-VDOM routing is the communication between VDOMs. VDOM links are virtual interfaces that connect VDOMs. A VDOM link contains a pair of interfaces, each one connected to a VDOM and forming either end of the inter-VDOM connection.

When VDOMs are configured on your FortiGate unit, configuring inter-VDOM routing and VDOM-links is very much like creating a VLAN interface. VDOM-links are managed through the web-based manager or CLI. In the web-based manager, VDOM link interfaces are managed in the network interface list.



VDOM link does not support traffic offload. If you want to use traffic offload, use NPU-VDOM-LINK.

To configure a VDOM link in the GUI:

1. Go to *Network* > *Interfaces*.
2. Click *Create New* > *VDOM Link*.
3. Configure the fields including the *Name*, *Virtual Domain*, IP information, access levels, and other fields.



By default, VDOM links are created as point-to-point (ppp) links. If required, the link type can be changed in the CLI.

For example, when running OSPF in IPv6, a link-local address is required in order to communicate with OSPF neighbors. For a VDOM link to obtain a link-local address its type must be set to `ethernet`.

To configure a VDOM link in the CLI:

```
config system vdom-link
    edit "<vdom-link-name>"
        set type {ppp | ethernet}    <--Optional
    next
end
config system interface
    edit "<vdom-link-name0>"
        set vdom "<VDOM Name>"
        set type vdom-link
    next
    edit "<vdom-link-name1>"
        set vdom "<VDOM Name>"
        set type vdom-link
    next
end
```

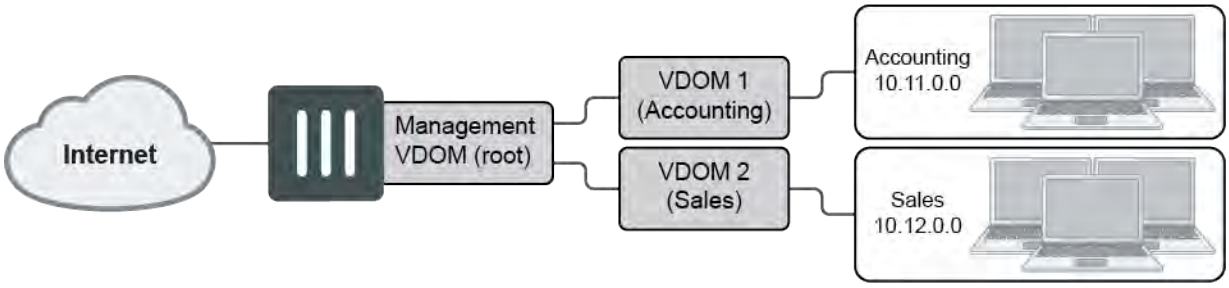
To delete a VDOM link in the GUI:

1. Go to *Network* > *Interfaces*.
2. Select a *VDOM Link* and click *Delete*.

To delete a VDOM link in the CLI:

```
config system vdom-link
    delete <VDOM-LINK-Name>
end
```

Sample configuration: Inter-VDOM routing



This example shows how to configure a FortiGate unit to use inter-VDOM routing.

Two departments of a company, Accounting and Sales, are connected to one FortiGate. The company uses a single ISP to connect to the Internet.

This example includes the following general steps. We recommend following the steps in the order below.

Create the VDOMs

To enable VDOMs and create the Sales and Accounting VDOMs:

```
config system global
    set vdom-mode multi-vdom
end
config system vdom
    edit Accounting
    next
    edit Sales
    next
end
```

Configure the physical interfaces

Next, configure the physical interfaces. This example uses three interfaces on the FortiGate unit: port2 (internal), port3 (DMZ), and port1 (external). Port2 and port3 interfaces each have a department's network connected. Port1 is for all traffic to and from the Internet and uses DHCP to configure its IP address, which is common with many ISPs.

To configure the interfaces:

```
config global
    config system interface
        edit port2
            set alias AccountingLocal
            set vdom Accounting
            set mode static
            set ip 172.100.1.1 255.255.0.0
            set allowaccess https ping ssh
            set description "The accounting dept internal interface"
        next
        edit port3
            set alias SalesLocal
            set vdom Sales
            set mode static
            set ip 192.168.1.1 255.255.0.0
            set allowaccess https ping ssh
            set description "The sales dept. internal interface"
        next
        edit port1
            set alias ManagementExternal
            set vdom root
            set mode DHCP
            set distance 5
            set gwdetect enable
            set dns-server-override enable
            set allowaccess https ssh snmp
            set description "The systemwide management interface."
        next
    end
end
```

Configure the VDOM links

To complete the connection between each VDOM and the management VDOM, add the two VDOM links. One pair is the Accounting – management link and the other is the Sales – management link.

When configuring inter-VDOM links, you do not have to assign IP addresses to the links unless you are using advanced features such as dynamic routing that require them. Not assigning IP addresses results in faster configuration and more available IP addresses on your networks.

To configure the Accounting and management VDOM link:

```
config global
    config system vdom-link
        edit AccountVlnk
        next
    end
    config system interface
        edit AccountVlnk0
            set vdom Accounting
            set ip 11.11.11.2 255.255.255.0
            set allowaccess https ping ssh
            set description "Accounting side of the VDOM link"
        next
```


Netskope Exhibit 1017

```
        edit AccountVlnk1
            set vdom root
            set ip 11.11.11.1 255.255.255.0
            set allowaccess https ping ssh
            set description "Management side of the VDOM link"
        next
    end
end
```

To configure the Sales and management VDOM link:

```
config global
    config system vdom-link
        edit SalesVlnk
        next
    end
    config system interface
        edit SalesVlnk0
            set vdom Sales
            set ip 12.12.12.2 255.255.255.0
            set allowaccess https ping ssh
            set description "Sales side of the VDOM link"
        next
        edit SalesVlnk1
            set vdom root
            set ip 12.12.12.1 255.255.255.0
            set allowaccess https ping ssh
            set description "Management side of the VDOM link"
        next
    end
end
```

Configure the firewall and security profile

With the VDOMs, physical interfaces, and VDOM links configured, the firewall must now be configured to allow the proper traffic. Firewalls are configured per-VDOM, and firewall objects and routes must be created for each VDOM separately.

To configure the firewall policies from AccountingLocal to Internet:

```
config vdom
    edit Accounting
        config firewall policy
            edit 1
                set name "Accounting-Local-to-Management"
                set srcintf port2
                set dstintf AccountVlnk
                set srcaddr all
                set dstaddr all
                set action accept
                set schedule always
                set service ALL
                set nat enable
            next
        end
```

```
    next
    edit root
        config firewall policy
            edit 2
                set name "Accounting-VDOM-to-Internet"
                set srcintf AccountVlnk
                set dstintf port1
                set srcaddr all
                set dstaddr all
                set action accept
                set schedule always
                set service ALL
                set nat enable
            next
        end
    next
end
```

To configure the firewall policies from SalesLocal to the Internet:

```
config vdom
    edit root
        config firewall policy
            edit 6
                set name "Sales-local-to-Management"
                set srcintf port2
                set srcaddr all
                set dstintf SalesVlnk
                set dstaddr all
                set schedule always
                set service ALL
                set action accept
                set logtraffic enable
            next
        end
    next
    edit Sales
        config firewall policy
            edit 7
                set name "Sales-VDOM-to-Internet"
                set srcintf SalesVlnk
                set srcaddr SalesManagement
                set dstintf external
                set dstaddr all
                set schedule always
                set service OfficeServices
                set action accept
                set logtraffic enable
            next
        end
    next
end
```

Test the configuration

When the inter-VDOM routing has been configured, test the configuration to confirm proper operation.

Netskope Exhibit 1017

Testing connectivity ensures that physical networking connections, FortiGate unit interface configurations, and firewall policies are properly configured.

The easiest way to test connectivity is to use the `ping` and `traceroute` command to confirm the connectivity of different routes on the network.

Test both from AccountingLocal to Internet and from SalesLocal to Internet.

Software switch

A software switch, or soft switch, is a virtual switch that is implemented at the software or firmware level and not at the hardware level. A software switch can be used to simplify communication between devices connected to different FortiGate interfaces. For example, using a software switch, you can place the FortiGate interface connected to an internal network on the same subnet as your wireless interfaces. Then devices on the internal network can communicate with devices on the wireless network without any additional configuration on the FortiGate unit, such as additional security policies.

A software switch can also be useful if you require more hardware ports for the switch on a FortiGate unit. For example, if your FortiGate unit has a 4-port switch, WAN1, WAN2, and DMZ interfaces, and you need one more port, you can create a soft switch that can include the four-port switch and the DMZ interface, all on the same subnet. These types of applications also apply to wireless interfaces, virtual wireless interfaces, and physical interfaces such as those in FortiWiFi and FortiAP units.

Similar to a hardware switch, a software switch functions like a single interface. A soft switch has one IP address and all the interfaces in the software switch are on the same subnet. Traffic between devices connected to each interface are not regulated by security policies, and traffic passing in and out of the switch are controlled by the same policy.

When setting up a software switch, consider the following:

- Ensure you have a back up of the configuration.
- Ensure you have at least one port or connection such as the console port to connect to the FortiGate unit. If you accidentally combine too many ports, you need a way to undo errors.
- The ports that you include must not have any link or relation to any other aspect of the FortiGate unit, such as DHCP servers, security policies, and so on.
- For increased security, you can create a captive portal for the switch to allow only specific user groups access to the resources connected to the switch.

To create a software switch in the GUI:

1. Go to *Network* > *Interfaces*.
2. Click *Create New* > *Interface*.
3. Set *Type* to *Software Switch*.
4. Configure the *Interface Name*, *Virtual Domain*, *Interface Members*, and other fields.

To create a software switch in the CLI:

```
config system switch-interface
    edit <switch-name>
        set type switch
        set member <interface_list>
    next
end
```


Netskope Exhibit 1017

```
config system interface
    edit <switch_name>
        set ip <ip_address>
        set allowaccess https ssh ping
    next
end
```

Sample configuration: software switch

For this example, the wireless interface (WiFi) needs to be on the same subnet as the DMZ1 interface to facilitate wireless syncing from an iPhone and a local computer. Because synching between two subnets is problematic, putting both interfaces on the same subnet the synching will work. The software switch will accomplish this.

1. Clear the interfaces and back up the configuration.
   a. Ensure the interfaces are not used for other security policy or for other use on the FortiGate unit.
   b. Check the WiFi and DMZ1 ports to ensure DHCP is not enabled on the interface and that there are no other dependencies on these interfaces.
   c. Save the current configuration so that if something doesn't work, recovery can be quick.
2. Merge the interfaces.
   Merge the WiFi port and DMZ1 port to create a software switch named `synchro` with an IP address of 10.10.21.12.
   Use the following CLI commands to create the switch, add the IP, and then set the administrative access for HTTPS, SSH and Ping.

```
config system switch-interface
    edit synchro
        set type switch
        set member dmz1 wifi
    next
end
config system interface
    edit synchro
        set ip 10.10.21.12
        set allowaccess https ssh ping
    next
end
```

   When the soft switch is set up, you now add security policies, DHCP servers, and any other configuration you normally do to configure interfaces on the FortiGate unit.

Zone

Zones are a group of one or more physical or virtual FortiGate interfaces that you can apply security policies to control inbound and outbound traffic. Grouping interfaces and VLAN subinterfaces into zones simplifies the creation of security policies where a number of network segments can use the same policy settings and protection profiles.

When you add a zone, you select the names of the interfaces and VLAN subinterfaces to add to the zone. Each interface still has its own address. Routing is still done between interfaces, that is, routing is not affected by zones. You can use security policies to control the flow of intra-zone traffic.

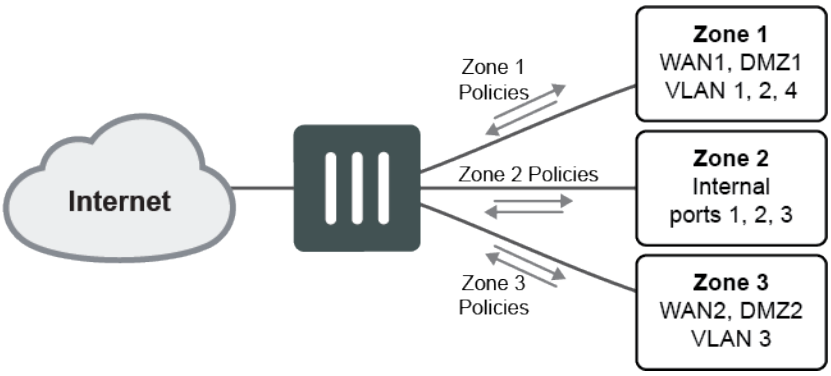
For example, in the sample configuration below, the network includes three separate groups of users representing different entities on the company network. While each group has its own set of ports and VLANs in each area, they can all use the same security policy and protection profiles to access the Internet. Rather than the administrator making nine

Netskope Exhibit 1017

separate security policies, he can make administration simpler by adding the required interfaces to a zone and creating three policies.

Sample configuration

You can configure policies for connections to and from a zone but not between interfaces in a zone. For this example, you can create a security policy to go between zone 1 and zone 3, but not between WAN2 and WAN1, or WAN1 and DMZ1.



To create a zone in the GUI:

1. Go to *Network* > *Interfaces*.

   If VDOMs are enabled, go to the VDOM to create a zone.

2. Click *Create New* > *Zone*.
3. Configure the *Name* and add the *Interface Members*.

To configure a zone to include the internal interface and a VLAN using the CLI:

```
config system zone
    edit Zone_1
        set interface internal VLAN_1
        set intrazone deny/allow
    next
end
```

Using zone in a firewall policy

To configure a firewall policy to allow any interface to access the Internet using the CLI:

```
config firewall policy
    edit 2
        set name "2"
        set srcintf "Zone_1"
        set dstintf "port15"
        set srcaddr "all"
        set dstaddr "all"
```

```
        set action accept
        set schedule "always"
        set service "ALL"
        set nat enable
    next
end
```

Intra-zone traffic

In the zone configuration you can `set intrazone deny` to prohibit the different interfaces in the same zone to talk to each other.

For example, if you have ten interfaces in your zone and the `intrazone` setting is `deny`. You now want to allow traffic between a very small number of networks on different interfaces that are part of the zone but you do not want to disable the intra-zone blocking.

In this example, the zone VLANs are defined as: `192.168.1.0/24`, `192.168.2.0/24`, ... `192.168.10.0/24`.

This policy allows traffic from 192.168.1.x to 192.168.2.x even though they are in the same zone and intra-zone blocking is enabled. The intra-zone blocking acts as a default deny rule and you have to specifically override it by creating a policy within the zone.

To enable intra-zone traffic, create the following policy:

Source Interface	Zone-name, e.g., `Vlans`
Source Address	192.168.1.0/24
Destination	Zone-name (same as Source Interface, i.e., `Vlans`)
Destination Address	192.168.2.0/24

Virtual Wire Pair

A virtual wire pair consists of two interfaces that do not have IP addressing and are treated like a transparent mode VDOM. All traffic received by one interface in the virtual wire pair can only be forwarded to the other interface, provided a virtual wire pair firewall policy allows this traffic. Traffic from other interfaces cannot be routed to the interfaces in a virtual wire pair. Redundant and 802.3ad aggregate (LACP) interfaces can be included in a virtual wire pair.

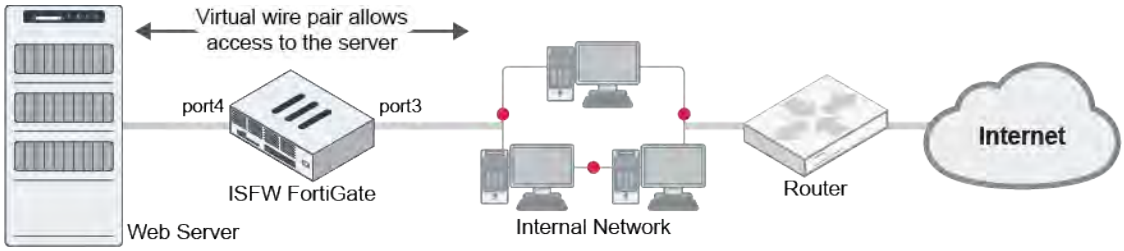
Virtual wire pairs are useful for a typical topology where MAC addresses do not behave normally. For example, port pairing can be used in a Direct Server Return (DSR) topology where the response MAC address pair may not match the request's MAC address pair.

Example

In this example, a virtual wire pair (port3 and port4) makes it easier to protect a web server that is behind a FortiGate operating as an Internal Segmentation Firewall (ISFW). Users on the internal network access the web server through the ISFW over the virtual wire pair.



Interfaces used in a virtual wire pair cannot be used to access the ISFW FortiGate. Before creating a virtual wire pair, make sure you have a different port configured to allow admin access using your preferred protocol.



To add a virtual wire pair using the CLI:

```
config system virtual-wire-pair
    edit "VWP-name"
        set member "port3" "port4"
        set wildcard-vlan disable
    next
end
```

To add a virtual wire pair using the GUI:

1. Go to *Network* > *Interfaces*.
2. Click *Create New* > *Virtual Wire Pair*.
3. Select the *Interface Members* to add to the virtual wire pair.
   These interfaces cannot be part of a switch, such as the default LAN/internal interface.
4. If required, enable *Wildcard VLAN* and set the *VLAN Filter*..
5. Click *OK*.

To create a virtual wire pair policy using the CLI:

```
config firewall policy
    edit 1
        set name "VWP-Policy"
        set srcintf "port3" "port4"
        set dstintf "port3" "port4"
        set srcaddr "all"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ALL"
        set utm-status enable
        set fsso disable
    next
end
```

To create a virtual wire pair policy using the GUI:

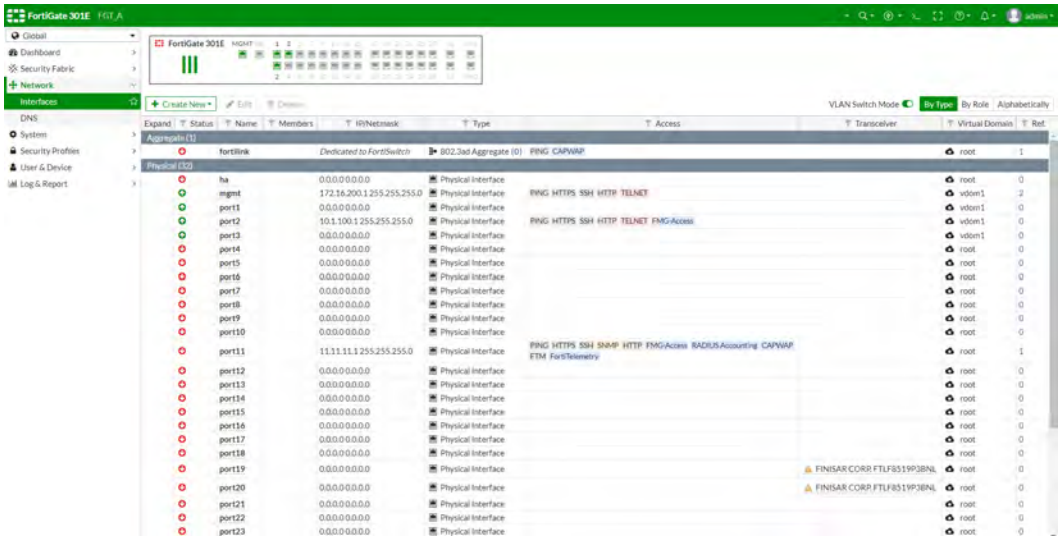
1. Go to *Policy & Objects* > *IPv4 Virtual Wire Pair Policy*.
2. Click *Create New*.
3. Select the direction that traffic is allowed to flow.
4. Configure the other fields.
5. Click *OK*.

Virtual switch support for FortiGate 300E series

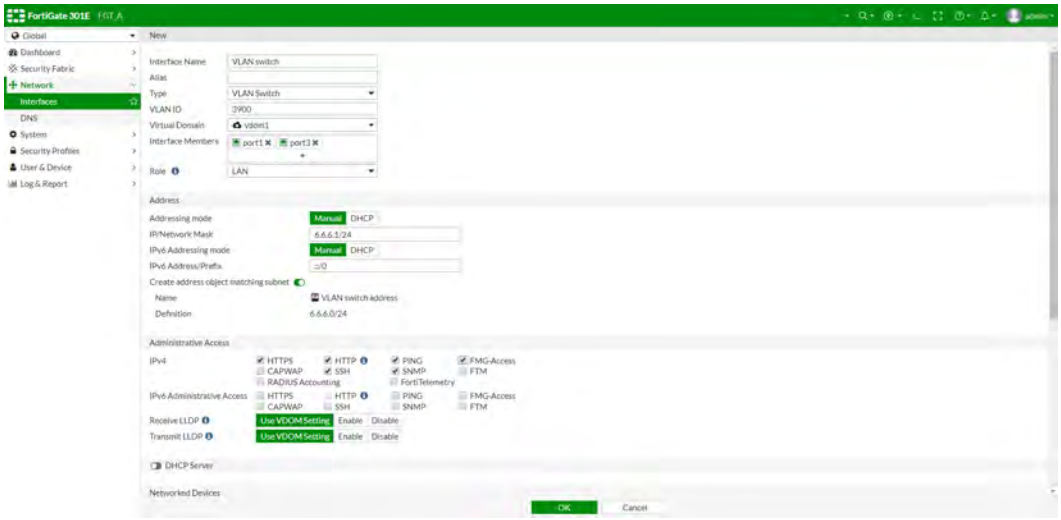
On the FortiGate 300E series, switch ports can be assigned to different VLANs.

To create a VLAN switch in the GUI:

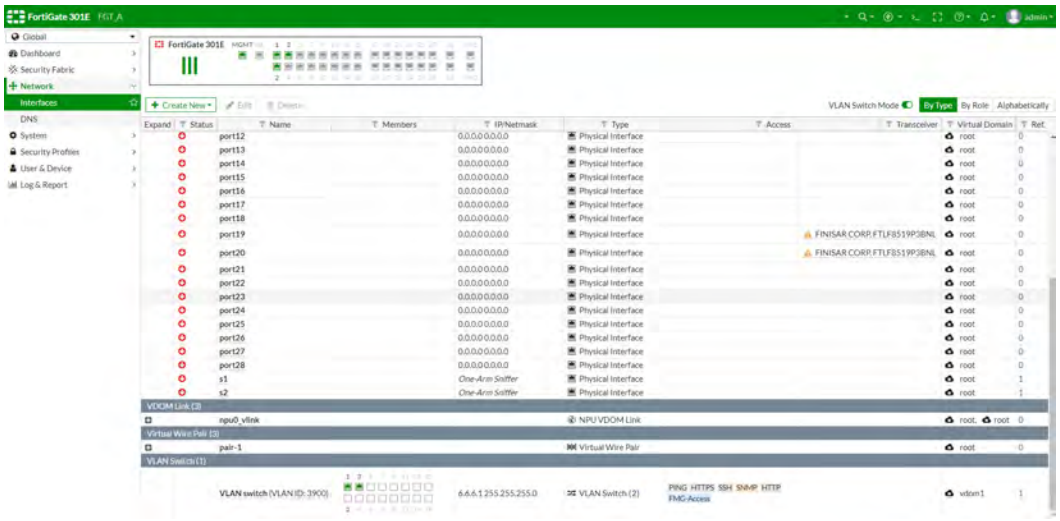
1. Go to *Network* > *Interfaces* and enable *VLAN Switch Mode*.



2. Click *Create New* > *Interface*.
3. Enter an interface name and configure the following:
   a. For *Type*, select *VLAN Switch*.
   b. (Optional) Enter a *VLAN ID* (range is 3900–3999).
   c. If applicable, select a *Virtual Domain*.
   d. Add the *Interface Members*.
   e. Configure the *Address* and *Administrative Access* settings as needed.
4. Click *OK*.



The new VLAN switch is visible in the interface table:



To create a VLAN switch in the CLI:

1. Enable VLAN switch mode:

```
config system global
    set virtual-switch-vlan enable
end
```

2. Create the VLAN switch. Optionally, you can assign an ID to the VLAN:
The default ID is `0`. You can use the default ID, or you can assign an ID to the VLAN (3900–3999).

```
config system virtual-switch
    edit "VLAN switch"
        set physical-switch "sw0"
        set vlan 3900
        config port
            edit "port1"
            next
            edit "port3"
            next
        end
    next
end
```

3. Configure the VLAN switch interface:

```
config system interface
    edit "VLAN switch"
        set vdom "vdom1"
        set ip 6.6.6.1 255.255.255.0
        set allowaccess ping https ssh snmp http fgfm
        set type hard-switch
        set snmp-index 15
    next
end
```

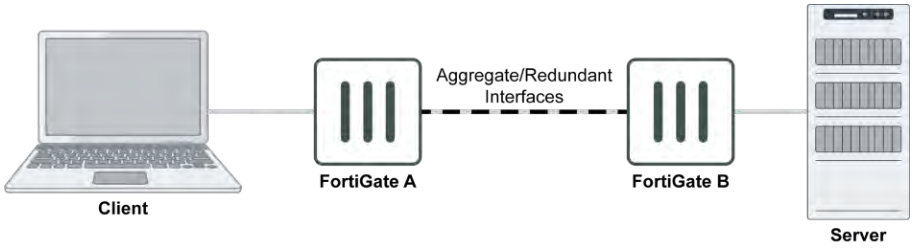
4. (Optional) Create a trunk interface:

```
config system interface
    edit port2
```

```
        set trunk enable
    next
end
```

Failure detection for aggregate and redundant interfaces

When an aggregate or redundant interface goes down, the corresponding fail-alert interface changes to down. When an aggregate or redundant interface comes up, the corresponding fail-alert interface changes to up.



Fail-detect for aggregate and redundant interfaces can be configured using the CLI.

To configure an aggregate interface so that port3 goes down with it:

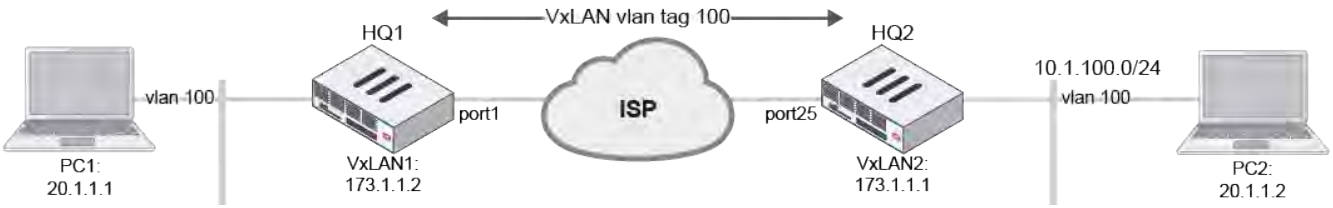
```
config system interface
   edit "agg1"
      set vdom "root"
      set fail-detect enable
      set fail-alert-method link-down
      set fail-alert-interfaces "port3"
      set type aggregate
      set member "port1" "port2"
   next
end
```

To configure a redundant interface so that port4 goes down with it:

```
config system interface
   edit "red1"
      set vdom "root"
      set fail-detect enable
      set fail-alert-method link-down
      set fail-alert-interfaces "port4"
      set type redundant
      set member "port1" "port2"
   next
end
```

VLAN inside VXLAN

VLANs can be assigned to VXLAN interfaces. In a data center network where VXLAN is used to create an L2 overlay network and for multitenant environments, a customer VLAN tag can be assigned to VXLAN interface. This allows the VLAN tag from VLAN traffic to be encapsulated within the VXLAN packet.



To configure VLAN inside VXLAN on HQ1:

1. Configure VXLAN:

```
config system vxlan
   edit "vxlan1"
      set interface port1
      set vni 1000
      set remote-ip 173.1.1.1
   next
end
```

2. Configure system interface:

```
config system interface
   edit vlan100
     set vdom root
     set vlanid 100
     set interface dmz
   next
   edit vxlan100
     set type vlan
     set vlanid 100
     set vdom root
     set interface vxlan1
   next
end
```

3. Configure software-switch:

```
config system switch-interface
   edit sw1
      set vdom root
      set member vlan100 vxlan100
      set intra-switch-policy implicit
   next
end
```

The default `intra-switch-policy implicit` behavior allows traffic between member interfaces within the switch. Therefore, it is not necessary to create firewall policies to allow this traffic.

Instead of creating a software-switch, it is possible to use a virtual-wire-pair as well. See Virtual Wire Pair with VXLAN on page 370.

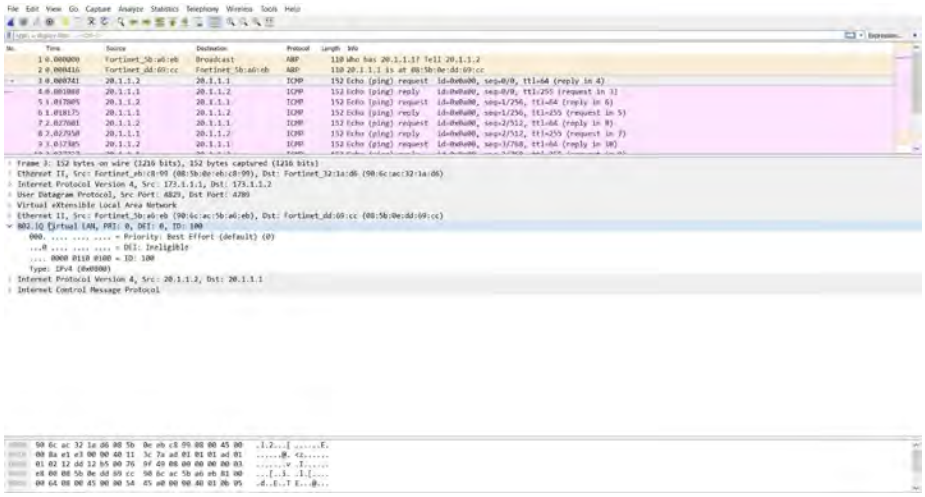
To configure VLAN inside VXLAN on HQ2:

1. Configure VXLAN:
```
config system vxlan
   edit "vxlan2"
      set interface port25
      set vni 1000
      set remote-ip 173.1.1.2
   next
end
```
2. Configure system interface:
```
config system interface
      edit vlan100
      set vdom root
      set vlanid 100
      set interface port20
   next
   edit vxlan100
      set type vlan
      set vlanid 100
      set vdom root
      set interface vxlan2
   next
end
```
3. Configure software-switch:
```
config system switch-interface
   edit sw1
      set vdom root
      set member vlan100 vxlan100
   next
end
```

To verify the configuration:

Ping PC1 from PC2.

The following is captured on HQ2:



This captures the VXLAN traffic between 172.1.1.1 and 172.1.1.2 with the VLAN 100 tag inside.

Virtual Wire Pair with VXLAN

Virtual wire pairs can be used with VXLAN interfaces.

In this examples, VXLAN interfaces are added between FortiGate HQ1 and FortiGate HQ2, a virtual wire pair is added in HQ1, and firewall policies are created on both HQ1 and HQ2.



To create VXLAN interface on HQ1:

```
config system interface
    edit "port11"
        set vdom "root"
        set ip 10.2.2.1 255.255.255.0
        set allowaccess ping https ssh snmp telnet
    next
end
config system vxlan
    edit "vxlan1"
        set interface "port11"
        set vni 1000
        set remote-ip "10.2.2.2"
    next
end
```

To create VXLAN interface on HQ2:

```
config system interface
    edit "port11"
        set vdom "root"
        set ip 10.2.2.2 255.255.255.0
        set allowaccess ping https ssh snmp http
    next
end
config system vxlan
    edit "vxlan1"
        set interface "port11"
        set vni 1000
        set remote-ip "10.2.2.1"
    next
end
config system interface
    edit "vxlan1"
        set vdom "root"
        set ip 10.1.100.2 255.255.255.0
        set allowaccess ping https ssh snmp
```

```
    next
end
```

To create a virtual wire pair on HQ1:

```
config system virtual-wire-pair
    edit "vwp1"
        set member "port10" "vxlan1"
    next
end
```

To create a firewall policy on HQ1:

```
config firewall policy
    edit 5
        set name "vxlan-policy"
        set srcintf "port10" "vxlan1"
        set dstintf "port10" "vxlan1"
        set srcaddr "all"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ALL"
        set utm-status enable
        set ssl-ssh-profile "certificate-inspection"
        set av-profile "default"
        set webfilter-profile "default"
        set dnsfilter-profile "default"
        set ips-sensor "default"
        set application-list "default"
        set fsso disable
    next
end
```

To create a firewall policy on HQ2:

```
config firewall policy
    edit 5
        set name "1"
        set srcintf "port13"
        set dstintf "vxlan1"
        set srcaddr "all"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ALL"
        set fsso disable
        set nat enable
    next
end
```


Netskope Exhibit 1017

DNS

Domain name system (DNS) is used by devices to locate websites by mapping a domain name to a website's IP address.

A FortiGate can serve different roles based on user requirements:

- A FortiGate can control what DNS server a network uses.
- A FortiGate can function as a DNS server.

FortiGuard Dynamic DNS (DDNS) allows a remote administrator to access a FortiGate's Internet-facing interface using a domain name that remains constant even when its IP address changes.

FortiOS supports DNS configuration for both IPv4 and IPv6 addressing. When a user requests a website, the FortiGate looks to the configured DNS servers to provide the IP address of the website in order to know which server to contact to complete the transaction.

The FortiGate queries the DNS servers whenever it needs to resolve a domain name into an IP address, such as for NTP or web servers defined by their domain names.

The following topics provide information about DNS:

- Important DNS CLI commands on page 372
- DNS domain list on page 373
- FortiGate DNS server on page 375
- DDNS on page 377
- DNS latency information on page 380
- DNS over TLS on page 382
- DNS troubleshooting on page 382

Important DNS CLI commands

DNS settings can be configured with the following CLI command:

```
config system dns
    set primary <ip_address>
    set secondary <ip_address>
    set dns-over-tls {enable | disable | enforce}
    set ssl-certificate <string>
    set domain <domains>
    set ip6-primary <ip6_address>
    set ip6-secondary <ip6_address>
    set timeout <integer>
    set retry <integer>
    set dns-cache-limit <integer>
    set dns-cache-ttl <integer>
    set cache-notfound-responses {enable | disable}
    set source-ip <class_ip>
end
```

For a FortiGate with multiple logical CPUs, you can set the DNS process number from 1 to the number of logical CPUs. The default DNS process number is 1.

```
config system global
    set dnsproxy-worker-count <integer>
end
```

dns-over-tls

DNS over TLS (DoT) is a security protocol for encrypting and wrapping DNS queries and answers via the Transport Layer Security (TLS) protocol. It can be enabled, disabled, or enforced:

- `disable`: Disable DNS over TLS (default).
- `enable`: Use TLS for DNS queries if TLS is available.
- `enforce`: Use only TLS for DNS queries. Does not fall back to unencrypted DNS queries if TLS is unavailable.

For more information, see DNS over TLS on page 382.

cache-notfound-responses

When enabled, any DNS requests that are returned with `NOT FOUND` can be stored in the cache. The DNS server is not asked to resolve the host name for `NOT FOUND` entries. By default, this option is disabled.

dns-cache-limit

Set the number of DNS entries that are stored in the cache (0 to 4294967295, default = 5000). Entries that remain in the cache provide a quicker response to requests than going out to the Internet to get the same information.

dns-cache-ttl

The duration that the DNS cache retains information, in seconds (60 to 86400 (1 day), default = 1800).

DNS domain list

You can configure up to eight domains in the DNS settings using the GUI or the CLI.

When a client requests a URL that does not include an FQDN, FortiOS resolves the URL by traversing through the DNS domain list and performing a query for each domain until the first match is found.

By default, FortiGate uses FortiGuard's DNS servers:

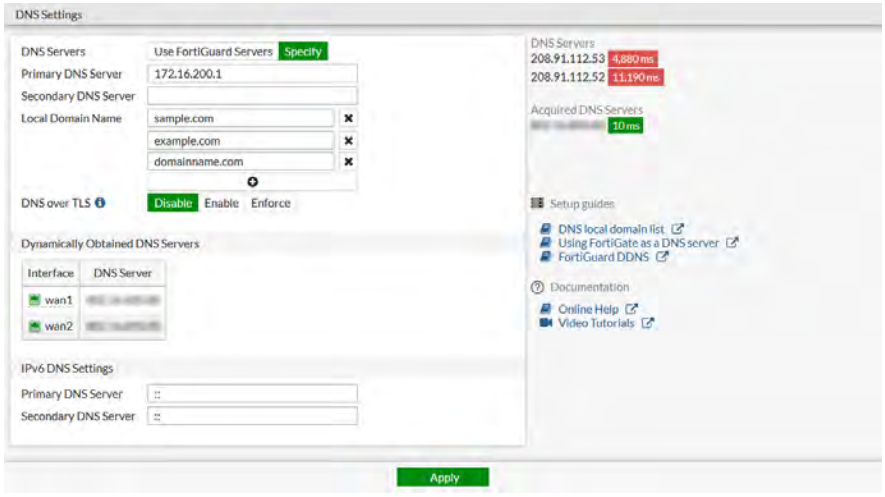
- Primary: 208.91.112.53
- Secondary: 208.91.112.52

You can also customize the DNS timeout time and the number of retry attempts.

To configure a DNS domain list in the GUI:

1. Go to *Network* > *DNS*.
2. Set *DNS Servers* to *Specify*.
3. Configure the primary and secondary DNS servers as needed.
4. In the *Local Domain Name* field, enter the first domain (`sample.com` in this example).

5. Click the + to add more domains (`example.com` and `domainname.com` in this example). You can enter up to eight domains.



6. Configure additional DNS settings as needed.
7. Click *Apply*.

To configure a DNS domain list in the CLI:

```
config system dns
    set primary 172.16.200.1
    set domain "sample.com" "example.com" "domainname.com"
end
```

Verify the DNS configuration

In the following example, the local DNS server has the entry for *host1* mapped to the FQDN of *host1.sample.com*, and the entry for *host2* is mapped to the FQDN of *host2.example.com*.

To verify that the DNS domain list is configured:

1. Open Command Prompt.
2. Enter `ping host1`.
   The system returns the following response:
   ```
   PING host1.sample.com (1.1.1.1): 56 data bytes
   ```
   As the request does not include an FQDN, FortiOS traverses the configured DNS domain list to find a match. Because *host1* is mapped to the *host1.sample.com*, FortiOS resolves *host1* to *sample.com*, the first entry in the domain list.
3. Enter `ping host2`.
   The system returns the following response:
   ```
   PING host2.example.com (2.2.2.2): 56 data bytes
   ```
   FortiOS traverses the domain list to find a match. It first queries sample.com, the first entry in the domain list, but does not find a match. It then queries the second entry in the domain list, example.com. Because *host2* is mapped to the FQDN of *host2.example.com*, FortiOS resolves *host2* to *example.com*.

DNS timeout and retry settings

The DNS timeout and retry settings can be customized using the CLI.

```
config system dns
    set timeout <integer>
    set retry <integer>
end
```

Variable	Description
timeout <integer>	The DNS query timeout interval, in seconds (1 - 10, default = 5).
retry <integer>	The number of times to retry the DNS query (0 - 5, default - 2).

FortiGate DNS server

You can create local DNS servers for your network. Depending on your requirements, you can either manually maintain your entries (primary DNS server), or use it to refer to an outside source (secondary DNS server).

A local, primary DNS server requires that you to manually add all URL and IP address combinations. Using a primary DNS server for local services can minimize inbound and outbound traffic, and access time. Making it authoritative is not recommended, because IP addresses can change, and maintaining the list can become labor intensive.

A secondary DNS server refers to an alternate source to obtain URL and IP address combinations. This is useful when there is a primary DNS server where the entry list is maintained.

In version 6.2 and later, FortiGate as a DNS server also supports TLS connections to a DNS client. See DNS over TLS on page 382 for details.

By default, DNS server options are not available in the FortiGate GUI.

To enable DNS server options in the GUI:

1. Go to *System* > *Feature Visibility*.
2. Enable *DNS Database* in the *Additional Features* section.
3. Click *Apply*.

Example configuration

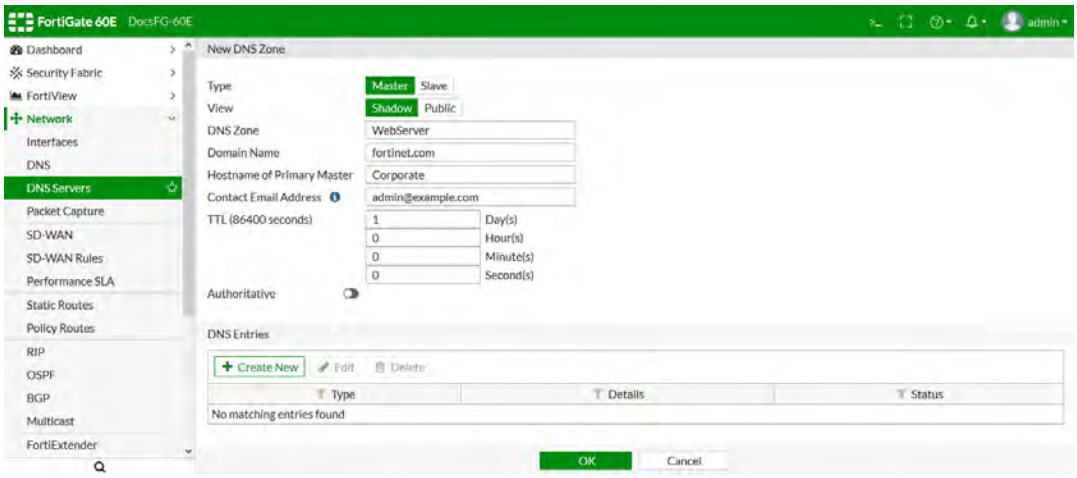
This section describes how to create an unauthoritative primary DNS server. The interface mode is recursive so that, if the request cannot be fulfilled, the external DNS servers will be queried.

To configure FortiGate as a primary DNS server in the GUI:

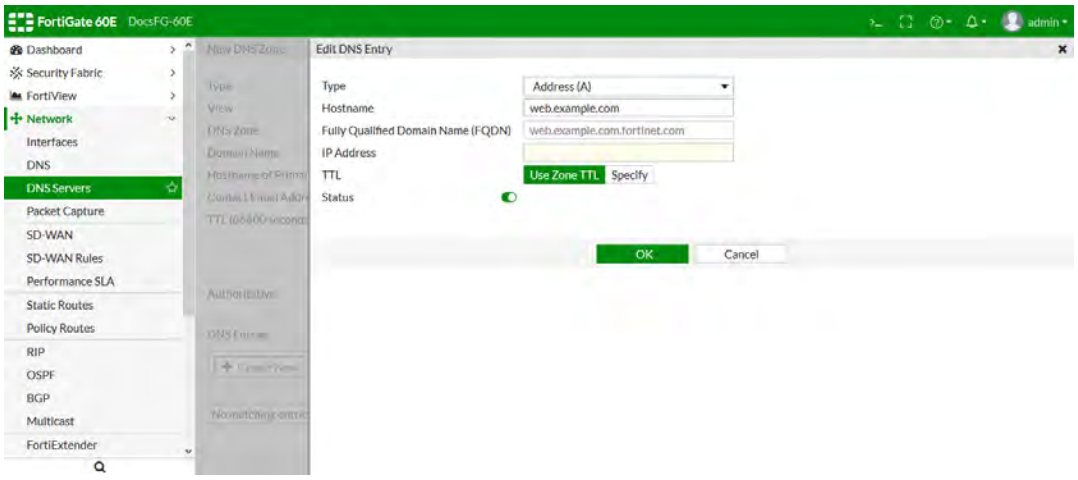
1. Go to *Network* > *DNS Servers*.
2. In the *DNS Database* table, click *Create New*.
3. Set *Type* to *Master*.
4. Set *View* to *Shadow*.
   The *View* setting controls the accessibility of the DNS server. If you select *Public*, external users can access or use the DNS server. If you select *Shadow*, only internal users can use it.

Netskope Exhibit 1017

5. Enter a *DNS Zone*, for example, *WebServer*.
6. Enter the *Domain Name* of the zone, for example, *fortinet.com*.
7. Enter the *Hostname* of the DNS server, for example, *Corporate*.
8. Enter the *Contact Email Address* for the administrator, for example, *admin@example.com*.
9. Disable *Authoritative*.

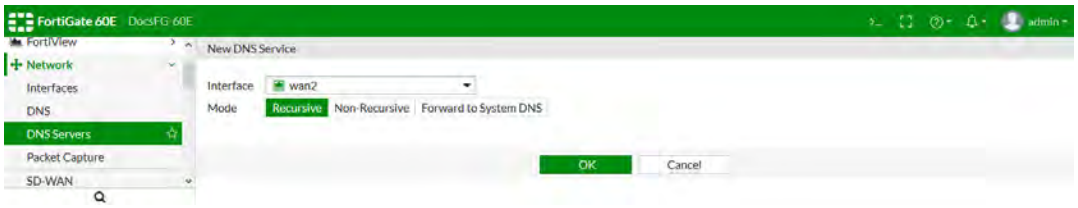


10. Add DNS entries:
    a. In the *DNS Entries* table, click *Create New*.
    b. Select a *Type*, for example *Address (A)*.
    c. Set the *Hostname*, for example *web.example.com*.



    d. Configure the remaining settings as needed. The options vary depending on the selected *Type*.
    e. Click *OK*.
11. Add more DNS entries as needed.
12. Click *OK*.
13. Enable DNS services on an interface:
    a. Go to *Network* > *DNS Servers*.
    b. In the *DNS Service on Interface* table, click *Create New*.
    c. Select the *Interface* for the DNS server, such as *wan2*.

d. Set the *Mode* to *Recursive*.



e. Click *OK*.

To configure FortiGate as a primary DNS server in the CLI:

```
config system dns-database
    edit WebServer
        set domain example.com
        set type master
        set view shadow
        set ttl 86400
        set primary-name corporate
        set contact admin@example.com
        set authoritative disable
        config dns-entry
            edit 1
                set hostname web.example.com
                set type A
                set ip 192.168.21.12
                set status enable
            next
        end
    next
end
```

```
config system dns-server
    edit wan1
        set mode recursive
    next
end
```

DDNS

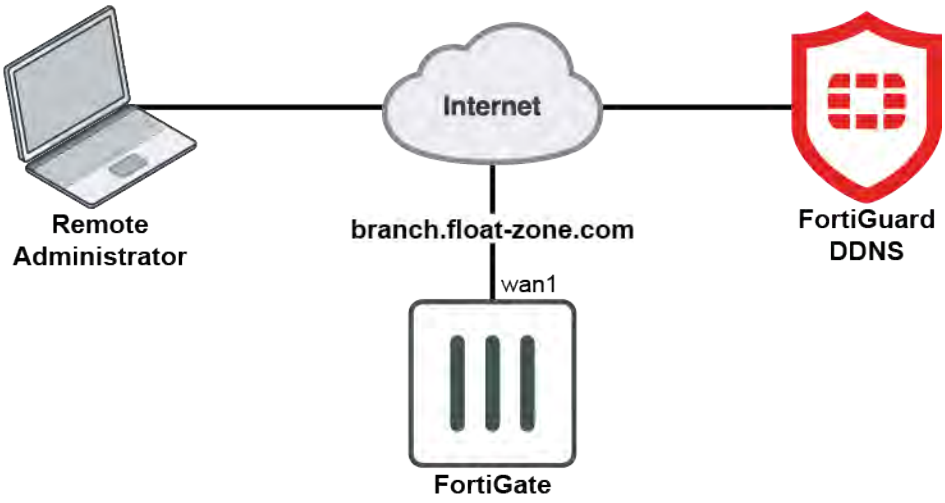
If your external IP address changes regularly and you have a static domain name, you can configure the external interface to use a dynamic DNS (DDNS) service. This ensures that external users and customers can always connect to your company firewall. If you have a FortiGuard subscription, you can use FortiGuard as the DDNS server.



- FortiGate does not support DDNS for pure TP mode.
- FortiGate models 1000D and higher do not support DDNS in the GUI.

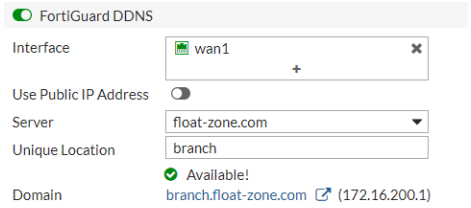
You can configure FortiGuard as the DDNS server using the GUI or CLI.

Sample topology



To configure FortiGuard as the DDNS server in the GUI:

1. Go to *Network* > *DNS*
2. Enable *FortiGuard DDNS*.
3. Select the *Interface* with the dynamic connection.
4. Select the *Server* that you have an account with.
5. Enter your *Unique Location*.



6. Click *Apply*.

To configure FortiGuard as the DDNS server in the CLI:

```
config system ddns
    edit <1>
        set ddns-server FortiGuardDDNS
        set ddns-domain "branch.float-zone.com"
        set monitor-interface "wan1"
    next
end
```

DDNS servers other than FortiGuard

If you do not have a FortiGuard subscription, or want to use a different DDNS server, you can configure a DDNS server for each interface. Only the first configure port appears in the GUI. The available commands vary depending on the selected DDNS server.

To configure DDNS servers other than FortiGuard in the CLI:

```
config system ddns
    edit <DDNS_ID>
        set monitor-interface <external_interface>
        set ddns-server <ddns_server_selection>
        ...
    next
end
```

Refresh DDNS IP addresses

You can configure FortiGate to refresh DDNS IP addresses. FortiGate periodically checks the DDNS server that is configured.

To configure FortiGate to refresh DDNS IP addresses using the CLI:

```
config system ddns
    edit <1>
        set ddns-server FortiGuardDDNS
        set use-public-ip enable
        set update-interval seconds
    next
end
```

Disable cleartext

When `clear-text` is disabled, FortiGate uses the SSL connection to send and receive (DDNS) updates.

To disable cleartext and set the SSL certificate using the CLI:

```
config system ddns
    edit <1>
        set clear-text disable
        set ssl-certificate <cert_name>
    next
end
```

DDNS update override

A DHCP server has an override command option that allows DHCP server communications to go through DDNS to perform updates for the DHCP client. This enforces a DDNS update of the A field every time even if the DHCP client does not request it. This allows support for the `allow`, `ignore`, and `deny client-updates` options.

To enable DDNS update override using the CLI:

```
config system dhcp server
    edit <0>
        set ddns-update enable
        set ddns-update_override enable
        set ddns-server-ip <ddns_server_ip>
```

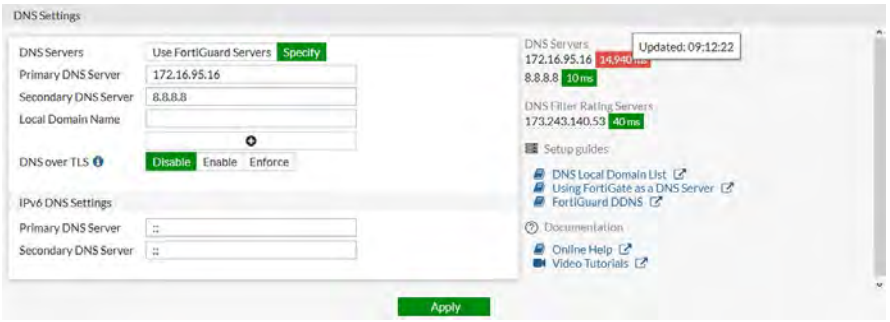

Netskope Exhibit 1017

```
        set ddns-zone <ddns_zone>
    next
end
```

DNS latency information

High latency in DNS traffic can result in an overall sluggish experience for end-users. In the *DNS Settings* pane, you can quickly identify DNS latency issues in your configuration.

Go to *Network* > *DNS* to view DNS latency information in the right side bar. If you use FortiGuard DNS, latency information for DNS, DNS filter, web filter, and outbreak prevention servers is also visible. Hover your pointer over a latency value to see when it was last updated.



To view DNS latency information using the CLI:

```
# diagnose test application dnsproxy 2
worker idx: 0
worker: count=1 idx=0
retry_interval=500 query_timeout=1495
DNS latency info:
vfid=0 server=2001::1 latency=1494 updated=73311
vfid=0 server=208.91.112.52 latency=1405 updated=2547
vfid=0 server=208.91.112.53 latency=19 updated=91
SDNS latency info:
vfid=0 server=173.243.138.221 latency=1 updated=707681
DNS_CACHE: alloc=35, hit=26
RATING_CACHE: alloc=1, hit=49
DNS UDP: req=66769 res=63438 fwd=83526 alloc=0 cmp=0 retrans=16855 to=3233
          cur=111 switched=8823467 num_switched=294 v6_cur=80 v6_switched=7689041 num_v6_
switched=6
         ftg_res=8 ftg_fwd=8 ftg_retrans=0
DNS TCP: req=0, res=0, fwd=0, retrans=0 alloc=0, to=0
FQDN: alloc=45 nl_write_cnt=9498 nl_send_cnt=21606 nl_cur_cnt=0
Botnet: searched=57 hit=0 filtered=57 false_positive=0
```

To view the latency from web filter and outbreak protection servers using the CLI:

```
# diagnose debug rating
Locale    : english

Service   : Web-filter
Status    : Enable
License   : Contract
```

```
Service  : Antispam
Status   : Disable

Service  : Virus Outbreak Prevention
Status   : Disable

--= Server List (Tue Jan 22 08:03:14 2019) --=

IP               Weight RTT Flags  TZ  Packets  Curr Lost Total Lost  Updated Time
173.243.138.194  10     0 DI       -8  700      0         2           Tue Jan 22 08:02:44
2019
173.243.138.195  10     0          -8  698      0         4           Tue Jan 22 08:02:44
2019
173.243.138.198  10     0          -8  698      0         4           Tue Jan 22 08:02:44
2019
173.243.138.196  10     0          -8  697      0         3           Tue Jan 22 08:02:44
2019
173.243.138.197  10     1          -8  694      0         0           Tue Jan 22 08:02:44
2019
96.45.33.64      10     22 D       -8  701      0         6           Tue Jan 22 08:02:44
2019
64.26.151.36     40     62         -5  704      0         10          Tue Jan 22 08:02:44
2019
64.26.151.35     40     62         -5  703      0         9           Tue Jan 22 08:02:44
2019
209.222.147.43   40     70 D       -5  696      0         1           Tue Jan 22 08:02:44
2019
66.117.56.42     40     70         -5  697      0         3           Tue Jan 22 08:02:44
2019
66.117.56.37     40     71         -5  702      0         9           Tue Jan 22 08:02:44
2019
65.210.95.239    40     74         -5  695      0         1           Tue Jan 22 08:02:44
2019
65.210.95.240    40     74         -5  695      0         1           Tue Jan 22 08:02:44
2019
45.75.200.88     90     142        0   706      0         12          Tue Jan 22 08:02:44
2019
45.75.200.87     90     155        0   714      0         20          Tue Jan 22 08:02:44
2019
45.75.200.85     90     156        0   711      0         17          Tue Jan 22 08:02:44
2019
45.75.200.86     90     159        0   704      0         10          Tue Jan 22 08:02:44
2019
62.209.40.72     100    157        1   701      0         7           Tue Jan 22 08:02:44
2019
62.209.40.74     100    173        1   705      0         11          Tue Jan 22 08:02:44
2019
62.209.40.73     100    173        1   699      0         5           Tue Jan 22 08:02:44
2019
121.111.236.179  180    138        9   706      0         12          Tue Jan 22 08:02:44
2019
121.111.236.180  180    138        9   704      0         10          Tue Jan 22 08:02:44
2019
```

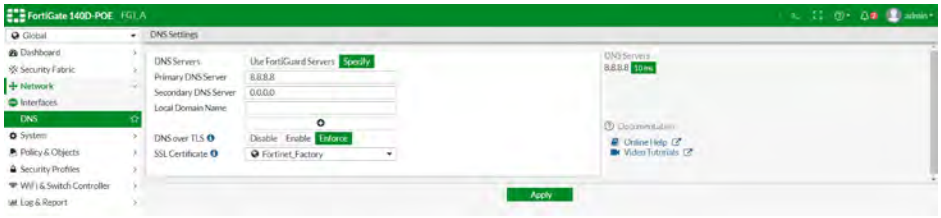

Netskope Exhibit 1017

DNS over TLS

DNS over TLS (DoT) is a security protocol for encrypting and wrapping DNS queries and answers via the TLS protocol. The goal of DNS over TLS is to increase user privacy and security by preventing eavesdropping and manipulation of DNS data via man-in-the-middle attacks. There is an option in the FortiOS DNS profile settings to enforce DoT for this added security.

To configure DoT in the GUI:

1. Go to *Network* > *DNS*. The *DNS Settings* pane opens.
2. For *DNS over TLS*, click *Enforce*.



3. Click *Apply*.

To configure DoT in the CLI:

```
config system dns
    set primary 8.8.8.8
    set dns-over-tls enforce
    set ssl-certificate "Fortinet_Factory"
end
```

DNS troubleshooting

The following diagnose command can be used to collect DNS debug information. If you do not specify worker ID, the default worker ID is `0`.

```
# diagnose test application dnsproxy
worker idx: 0
1. Clear DNS cache
2. Show stats
3. Dump DNS setting
4. Reload FQDN
5. Requery FQDN
6. Dump FQDN
7. Dump DNS cache
8. Dump DNS DB
9. Reload DNS DB
10. Dump secure DNS policy/profile
11. Dump Botnet domain
12. Reload Secure DNS setting
13. Show Hostname cache
14. Clear Hostname cache
15. Show SDNS rating cache
16. Clear SDNS rating cache
```

```
17. DNS debug bit mask
99. Restart dnsproxy worker
```

To view useful information about the ongoing DNS connection:

```
# diagnose test application dnsproxy 3
```

Important fields include:

`tls`	1 if the connection is TLS, 0 if the connection is not TLS.
`rt`	The round trip time of the DNS latency.
`probe`	The number of probes sent.

To dump the second DNS worker's cache:

```
# diagnose test application dnsproxy 7 1
```

To enable debug on the second worker:

```
# diagnose debug application dnsproxy -1 1
```

To enable debug on all workers by specifying -1 as worker ID:

```
# diagnose debug application dnsproxy -1 -1
```

Explicit and transparent proxies

This section contains instructions for configuring explicit and transparent proxies.

- Explicit web proxy on page 383
- Transparent proxy on page 388
- FTP proxy on page 386
- Proxy policy addresses on page 391
- Proxy policy security profiles on page 399
- Explicit proxy authentication on page 405
- Transparent web proxy forwarding on page 410
- Upstream proxy authentication in transparent proxy mode on page 411
- Multiple dynamic header count on page 413
- Restricted SaaS access (Office 365, G Suite, Dropbox) on page 416
- Explicit proxy and FortiSandbox Cloud on page 418
- Proxy chaining (web proxy forwarding servers) on page 421
- Agentless NTLM authentication for web proxy on page 425

Explicit web proxy

Explicit web proxy can be configured on FortiGate for proxying HTTP and HTTPS traffic.

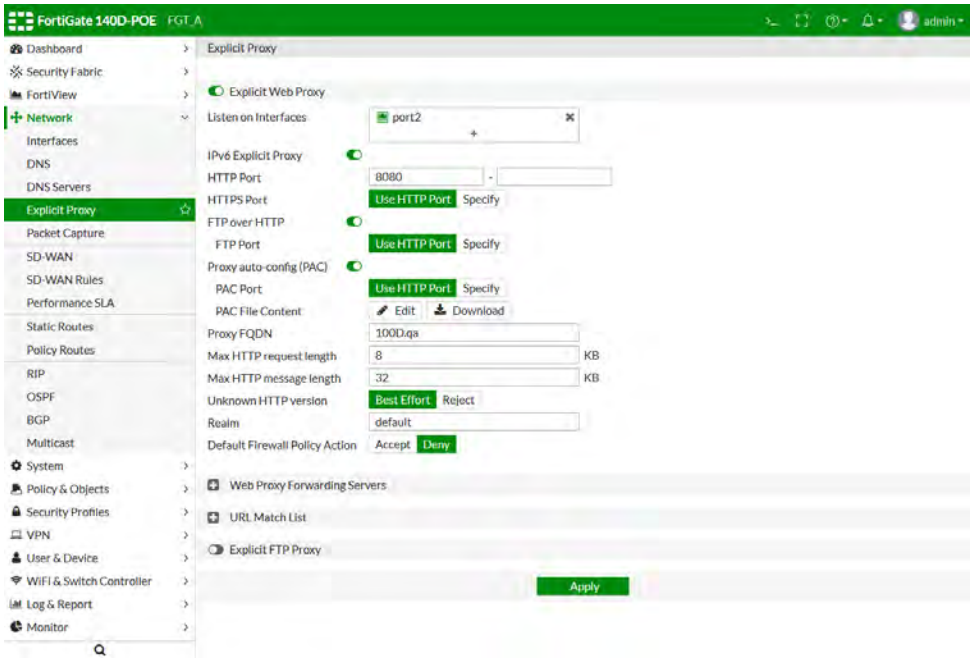
Netskope Exhibit 1017

To deploy explicit proxy, individual client browsers can be manually configured to send requests directly to the proxy, or they can be configured to download proxy configuration instructions from a Proxy Auto-Configuration (PAC) file.

When explicit proxy is configured on an interface, the interface IP address can be used by client browsers to forward requests directly to the FortiGate. FortiGate also supports PAC file configuration.

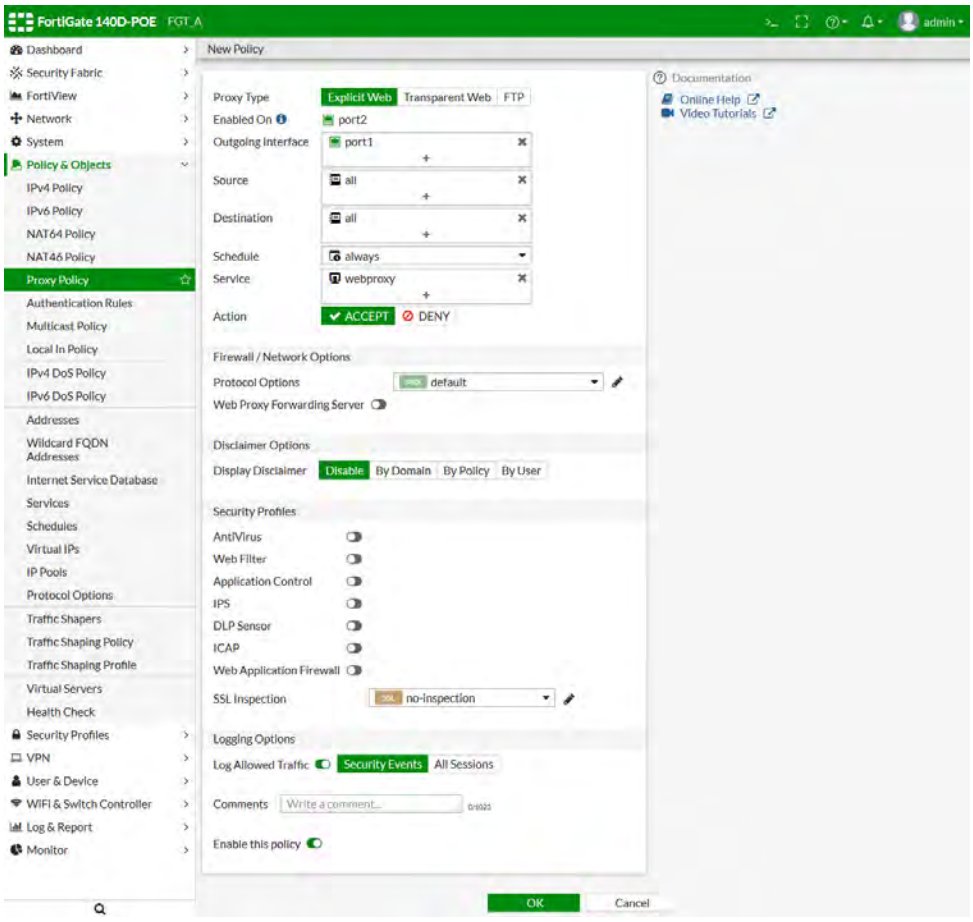
To configure explicit web proxy in the GUI:

1. Enable and configure explicit web proxy:
   a. Go to *Network* > *Explicit Proxy*.
   b. Enable *Explicit Web Proxy*.
   c. Select *port2* as the *Listen on Interfaces* and set the *HTTP Port* to *8080*.
   d. Configure the remaining settings as needed.



   e. Click *Apply*.
2. Create an explicit web proxy policy:
   a. Go to *Policy & Objects* > *Proxy Policy*.
   b. Click *Create New*.
   c. Set *Proxy Type* to *Explicit Web* and *Outgoing Interface* to *port1*.

d. Also set *Source* and *Destination* to *all*, *Schedule* to *always*, *Service* to *webproxy*, and *Action* to *ACCEPT*.



e. Click *OK* to create the policy.

This example creates a basic policy. If required, security profiles can be enabled, and deep SSL inspection can be selected to inspect HTTPS traffic.

---

3. Configure a client to use the FortiGate explicit proxy:
Set the FortiGate IP address as the proxy IP address in the browser, or use an automatic configuration script for the PAC file.

To configure explicit web proxy in the CLI:

1. Enable and configure explicit web proxy:

```
config web-proxy explicit
    set status enable
    set ftp-over-http enable
    set socks enable
    set http-incoming-port 8080
    set ipv6-status enable
    set unknown-http-version best-effort
end
config system interface
```

```
        edit "port2"
            set vdom "vdom1"
            set ip 10.1.100.1 255.255.255.0
            set allowaccess ping https ssh snmp http telnet
            set type physical
            set explicit-web-proxy enable
            set snmp-index 12
            end
        next
    end
```

2. Create an explicit web proxy policy:

```
config firewall proxy-policy
    edit 1
        set proxy explicit-web
        set dstintf "port1"
        set srcaddr "all"
        set dstaddr "all"
        set service "webproxy"
        set action accept
        set schedule "always"
        set logtraffic all
    next
end
```



This example creates a basic policy. If required, security profiles can be enabled, and deep SSL inspection can be selected to inspect HTTPS traffic.

3. Configure a client to use the FortiGate explicit web proxy:
   Set the FortiGate IP address as the proxy IP address in the browser, or use an automatic configuration script for the PAC file.

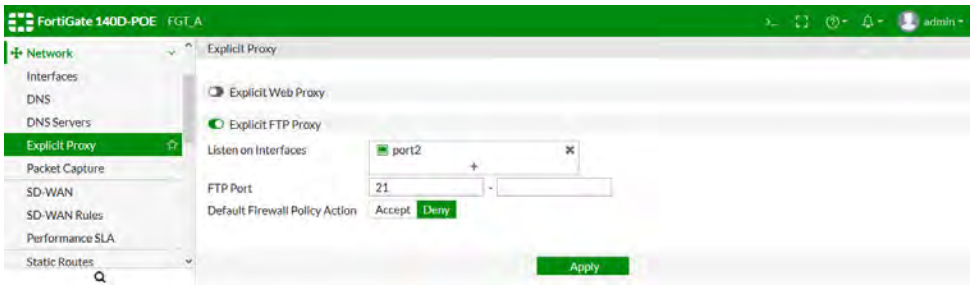
FTP proxy

FTP proxies can be configured on the FortiGate so that FTP traffic can be proxied. When the FortiGate is configured as an FTP proxy, FTP client applications should be configured to send FTP requests to the FortiGate.

To configure explicit FTP proxy in the GUI:

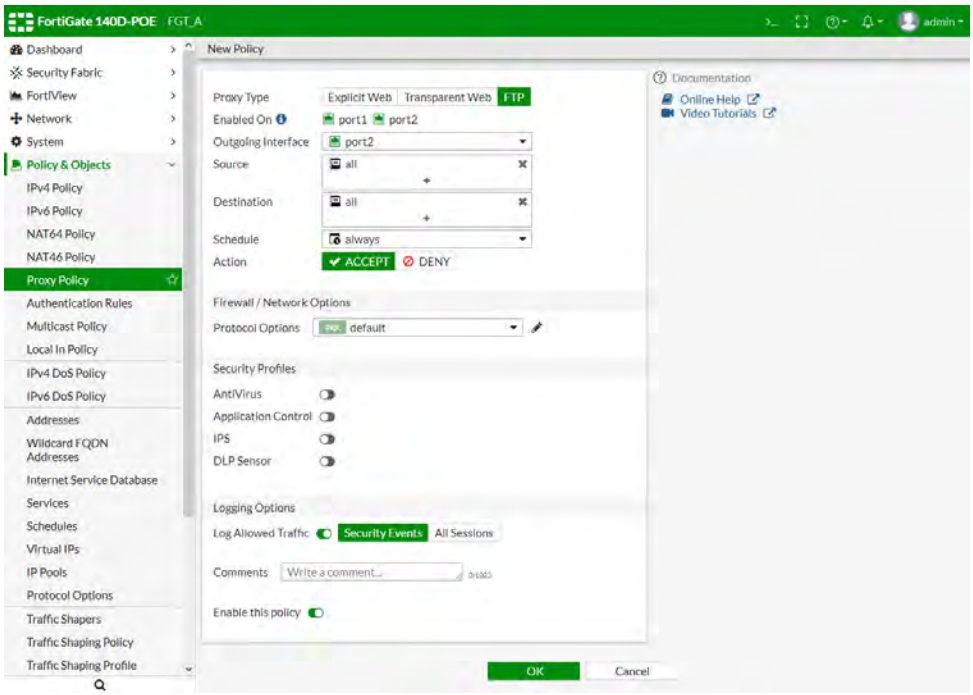
1. Enable and configure explicit FTP proxy:
   a. Go to *Network* > *Explicit Proxy*.
   b. Enable *Explicit FTP Proxy*.
   c. Select *port2* as the *Listen on Interfaces* and set the *HTTP Port* to *21*.

d. Configure the *Default Firewall Policy Action* as needed.



e. Click *Apply*.

2. Create an explicit FTP proxy policy:
   a. Go to *Policy & Objects* > *Proxy Policy*.
   b. Click *Create New*.
   c. Set *Proxy Type* to *FTP* and *Outgoing Interface* to *port1*.
   d. Also set *Source* and *Destination* to *all*, *Schedule* to *always*, and *Action* to *ACCEPT*.



e. Click *OK* to create the policy.



This example creates a basic policy. If required, security profiles can be enabled.

3. Configure the FTP client application to use the FortiGate IP address.

To configure explicit FTP proxy in the CLI:

1. Enable and configure explicit FTP proxy:

```
config ftp-proxy explicit
    set status enable
    set incoming-port 21
end
config system interface
    edit "port2"
        set vdom "vdom1"
        set ip 10.1.100.1 255.255.255.0
        set allowaccess ping https ssh snmp http telnet
        set type physical
        set explicit-ftp-proxy enable
        set snmp-index 12
        end
    next
end
```

2. Create an explicit FTP proxy policy:

```
config firewall proxy-policy
    edit 4
        set proxy ftp
        set dstintf "port2"
        set srcaddr "all"
        set dstaddr "all"
        set action accept
        set schedule "always"
    next
end
```



This example creates a basic policy. If required, security profiles can be enabled.

3. Configure the FTP client application to use the FortiGate IP address.

Transparent proxy

In a transparent proxy deployment, the user's client software, such as a browser, is unaware that it is communicating with a proxy.

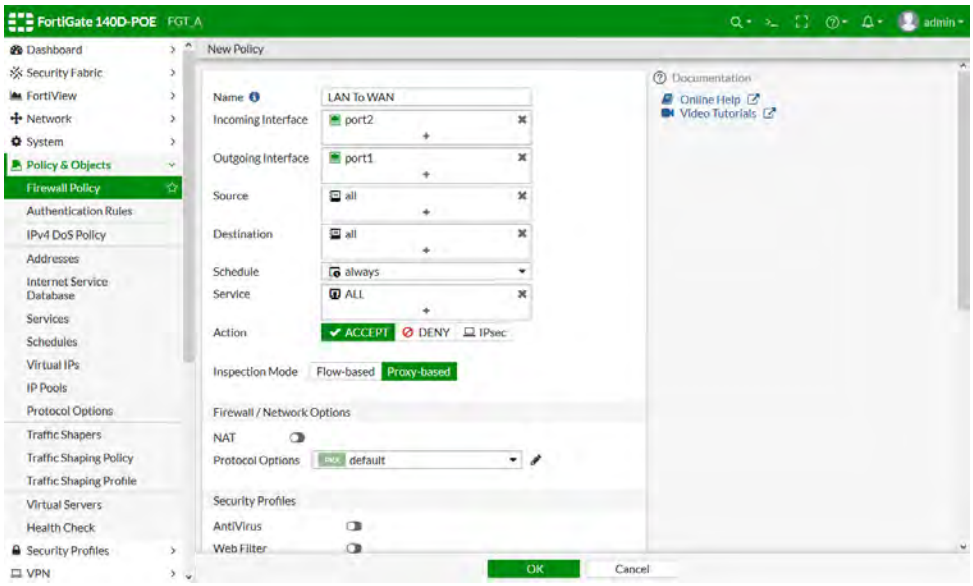
Users request Internet content as usual, without any special client configuration, and the proxy serves their requests. FortiGate also allows user to configure in transparent proxy mode.

To configure transparent proxy in the GUI:

1. Configure a regular firewall policy with HTTP redirect:
   a. Go to *Policy & Objects* > *Firewall Policy*.
   b. Click *Create New*.
   c. Name the policy appropriately, set the *Incoming Interface* to *port2*, and set the *Outgoing Interface* to *port1*.
   d. Also set *Source* and *Destination* to *all*, *Schedule* to *always*, *Service* to *ALL*, and *Action* to *ACCEPT*.

Netskope Exhibit 1017

e. Set *Inspection Mode* to *Proxy-based* and *SSL Inspection* to *deep-inspection*.



f. Configure the remaining settings as needed.

g. Click *OK*.



HTTP redirect can only be configured in the CLI. To redirect HTTPS traffic, deep inspection is required.

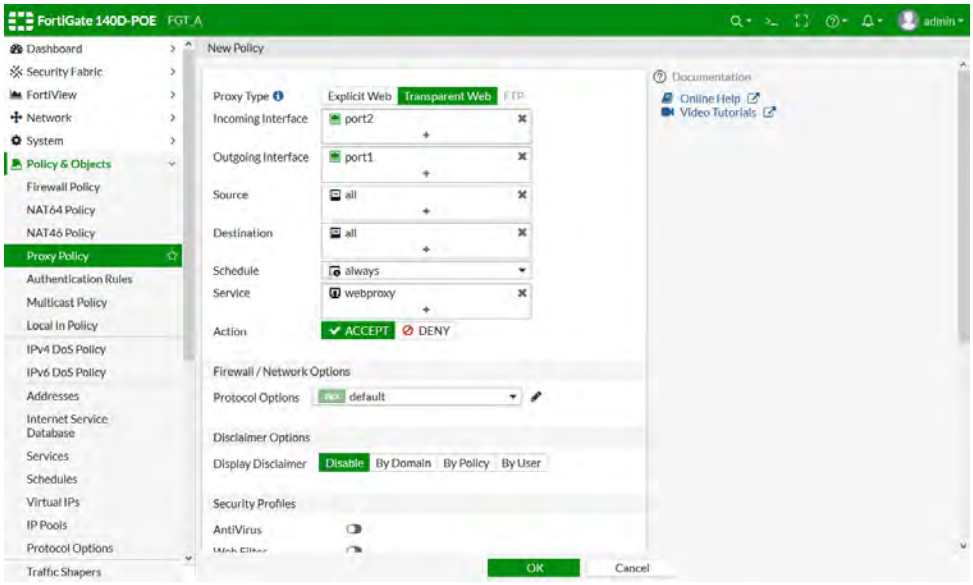
2. Configure a transparent proxy policy:

a. Go to *Policy & Objects* > *Proxy Policy*.

b. Click *Create New*.

c. Set *Proxy Type* to *Transparent Web*, set the *Incoming Interface* to *port2*, and set the *Outgoing Interface* to *port1*.

d. Also set *Source* and *Destination* to *all*, *Schedule*to *always*, *Service* to *webproxy*, and *Action* to *ACCEPT*.



e. Configure the remaining settings as needed.

f. Click *OK* to create the policy.



This example creates a basic policy. If required, security profiles can be enabled, and deep SSL inspection can be selected to inspect HTTPS traffic.

3. No special configure is required on the client to use FortiGate transparent proxy. As the client is using the FortiGate as its default gateway, requests will first hit the regular firewall policy, and then be redirected to the transparent proxy policy.

To configure transparent proxy in the CLI:

1. Configure a regular firewall policy with HTTP redirect:

```
config firewall policy
    edit 1
        set name "1"
        set srcintf "port2"
        set dstintf "port1"
        set srcaddr "all"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ALL"
        set inspection-mode proxy
        set http-policy-redirect enable
        set fsso disable
        set ssl-ssh-profile "deep-inspection"
        set nat enable
    next
end
```

2. Configure a transparent proxy policy:

```
config firewall proxy-policy
    edit 5
        set proxy transparent-web
        set srcintf "port2"
        set dstintf "port1"
        set srcaddr "all"
        set dstaddr "all"
        set service "webproxy"
        set action accept
        set schedule "always"
    next
end
```



This example creates a basic policy. If required, security profiles can be enabled, and deep SSL inspection can be selected to inspect HTTPS traffic.

3. No special configure is required on the client to use FortiGate transparent proxy. As the client is using the FortiGate as its default gateway, requests will first hit the regular firewall policy, and then be redirected to the transparent proxy policy.

Netskope Exhibit 1017

Proxy policy addresses

Proxy addresses are designed to be used only by proxy policies. The following address types are available:

- Host regex match on page 391
- URL pattern on page 392
- URL category on page 393
- HTTP method on page 394
- HTTP header on page 395
- User agent on page 396
- Advanced (source) on page 397
- Advanced (destination) on page 398

Fast policy match

The fast policy match function improves the performance of IPv4 explicit and transparent web proxies on FortiGate devices.

When enabled, after the proxy policies are configured, the FortiGate builds a fast searching table based on the different proxy policy matching criteria. When fast policy matching is disabled, web proxy traffic is compared to the policies one at a time from the beginning of the policy list.

Fast policy matching is enabled by default, and can be configured with the following CLI command:

```
config web-proxy global
    set fast-policy-match {enable | disable}
end
```

Host regex match

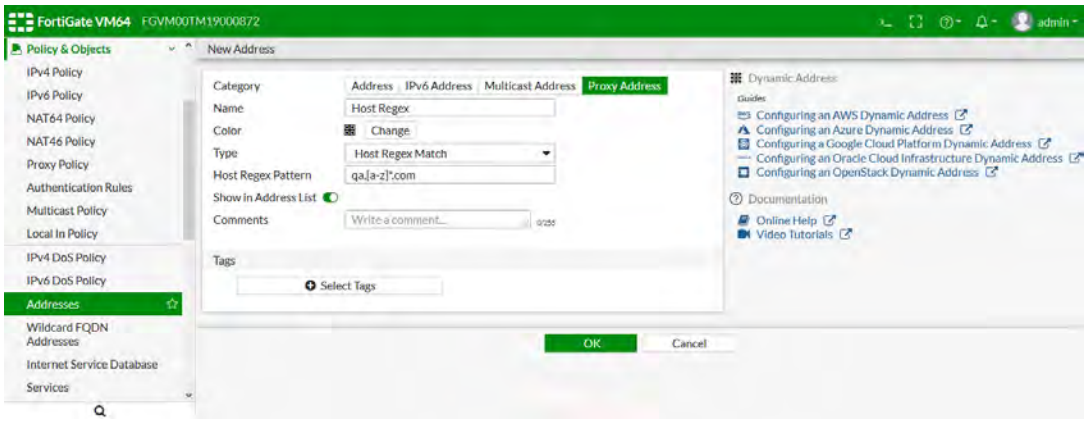
In this address type, a user can create a hostname as a regular expression. Once created, the hostname address can be selected as a destination of a proxy policy. This means that a policy will only allow or block requests that match the regular expression.

This example creates a host regex match address with the pattern *qa.[a-z]*.com*.

To create a host regex match address in the GUI:

1. Go to *Policy & Objects* > *Addresses*.
2. Click *Create New* > *Address*.
3. Set the following:
   - *Category* to *Proxy Address*,
   - *Name* to *Host Regex*,
   - *Type* to *Host Regex Match*, and
   - *Host Regex Pattern* to *qa.[a-z]*.com*.

Netskope Exhibit 1017



4. Click *OK*.

To create a host regex match address in the CLI:

```
config firewall proxy-address
    edit "Host Regex"
        set type host-regex
        set host-regex "qa.[a-z]*.com"
    next
end
```

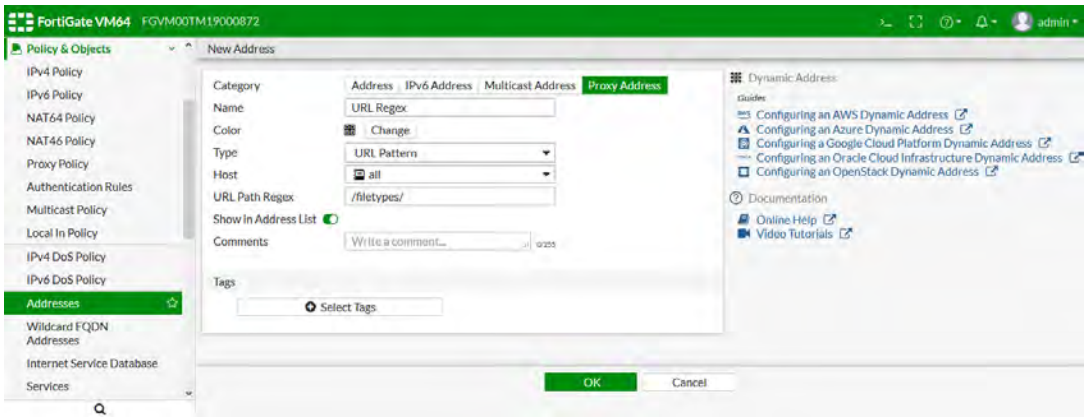
URL pattern

In this address type, a user can create a URL path as a regular expression. Once created, the path address can be selected as a destination of a proxy policy. This means that a policy will only allow or block requests that match the regular expression.

This example creates a URL pattern address with the pattern */filetypes/*.

To create a URL pattern address in the GUI:

1. Go to *Policy & Objects* > *Addresses*.
2. Click *Create New* > *Address*.
3. Set the following:
   - *Category* to *Proxy Address*,
   - *Name* to *URL Regex*,
   - *Type* to *URL Pattern*,
   - *Host* to *all*, and
   - *URL Path Regex* to */filetypes/*.



4. Click *OK*.

To create a URL pattern address in the CLI:

```
config firewall proxy-address
    edit "URL Regex"
        set type url
        set host "all"
        set path "/filetypes/"
    next
end
```

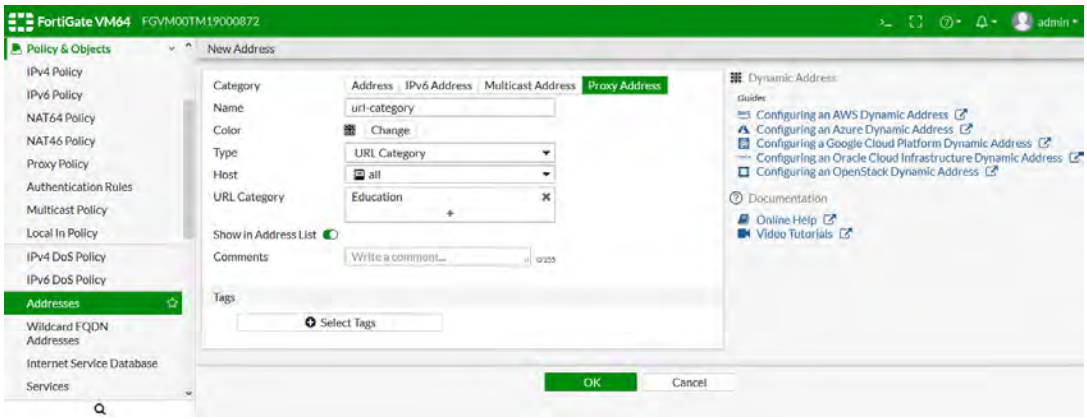
URL category

In this address type, a user can create a URL category based on a FortiGuard URL ID. Once created, the address can be selected as a destination of a proxy policy. This means that a policy will only allow or block requests that match the URL category.

The example creates a URL category address for URLs in the *Education* category. For more information about categories, see https://fortiguard.com/webfilter/categories.

To create a URL category address in the GUI:

1. Go to *Policy & Objects* > *Addresses*.
2. Click *Create New* > *Address*.
3. Set the following:
   - *Category* to *Proxy Address*,
   - *Name* to *url-category*,
   - *Type* to *URL Category*,
   - *Host* to *all*, and
   - *URL Category* to *Education*.



4. Click *OK*.

To create a URL category address in the CLI:

```
config firewall proxy-address
    edit "url-category"
        set type category
        set host "all"
        set category 30
    next
end
```

To see a list of all the categories and their numbers, when editing the address, enter `set category ?`.

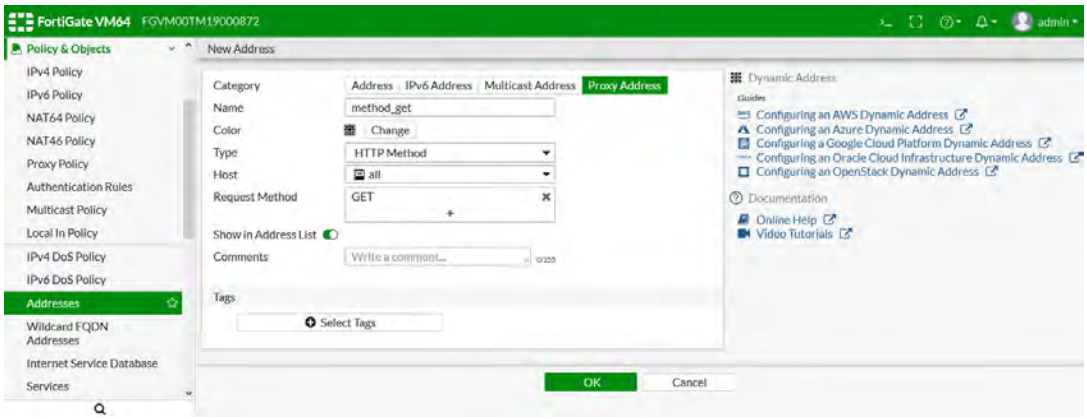
HTTP method

In this address type, a user can create an address based on the HTTP request methods that are used. Multiple method options are supported, including: *CONNECT*, *DELETE*, *GET*, *HEAD*, *OPTIONS*, *POST*, *PUT*, and *TRACE*. Once created, the address can be selected as a source of a proxy policy. This means that a policy will only allow or block requests that match the selected HTTP method.

The example creates a HTTP method address that uses the GET method.

To create a HTTP method address in the GUI:

1. Go to *Policy & Objects* > *Addresses*.
2. Click *Create New* > *Address*.
3. Set the following:
   - *Category* to *Proxy Address*,
   - *Name* to *method_get*,
   - *Type* to *HTTP Method*,
   - *Host* to *all*, and
   - *Request Method* to *GET*.



4. Click *OK*.

To create a HTTP method address in the CLI:

```
config firewall proxy-address
    edit "method_get"
        set type method
        set host "all"
        set method get
    next
end
```

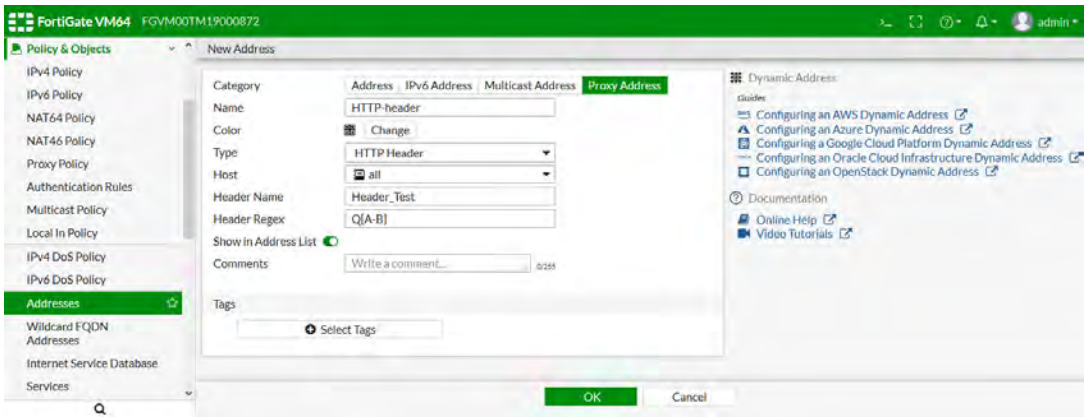
HTTP header

In this address type, a user can create a HTTP header as a regular expression. Once created, the header address can be selected as a source of a proxy policy. This means that a policy will only allow or block requests where the HTTP header matches the regular expression.

This example creates a HTTP header address with the pattern *Q[A-B]*.

To create a HTTP header address in the GUI:

1. Go to *Policy & Objects* > *Addresses*.
2. Click *Create New* > *Address*.
3. Set the following:
   - *Category* to *Proxy Address*,
   - *Name* to *HTTP-header*,
   - *Type* to *HTTP Header*,
   - *Host* to *all*,
   - *Header Name* to *Header_Test*, and
   - *Header Regex* to *Q[A-B]*.



4. Click *OK*.

To create a HTTP header address in the CLI:

```
config firewall proxy-address
    edit "method_get"
        set type header
        set host "all"
        set header-name "Header_Test"
        set header "Q[A-B]"
    next
end
```

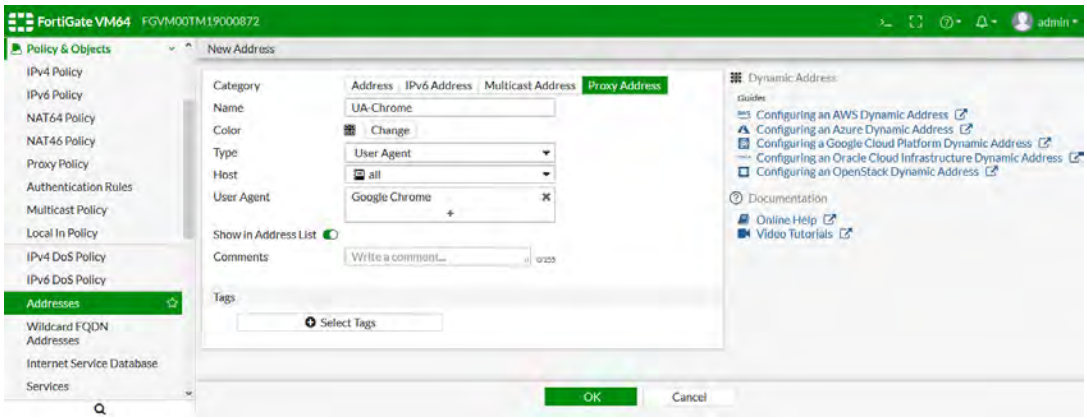
User agent

In this address type, a user can create an address based on the names of the browsers that are used as user agents. Multiple browsers are supported, such as Chrome, Firefox, Internet Explorer, and others. Once created, the address can be selected as a source of a proxy policy. This means that a policy will only allow or block requests from the specified user agent.

This example creates a user agent address for Google Chrome.

To create a user agent address in the GUI:

1. Go to *Policy & Objects* > *Addresses*.
2. Click *Create New* > *Address*.
3. Set the following:
   - *Category* to *Proxy Address*,
   - *Name* to *UA-Chrome*,
   - *Type* to *User Agent*,
   - *Host* to *all*, and
   - *User Agent* to *Google Chrome*.



4. Click *OK*.

To create a user agent address in the CLI:

```
config firewall proxy-address
    edit "UA-Chrome"
        set type ua
        set host "all"
        set ua chrome
    next
end
```

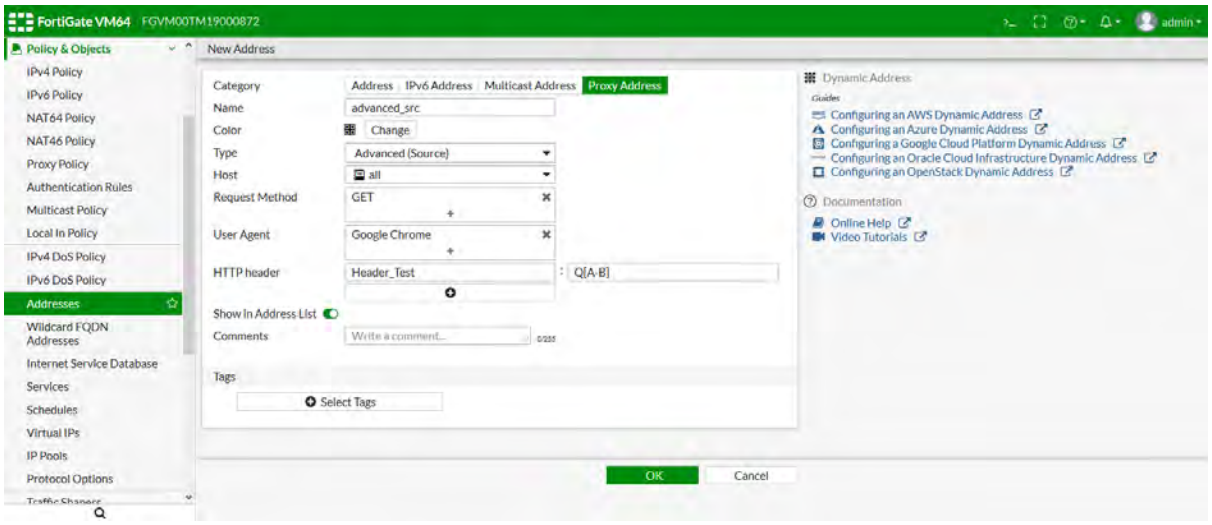
Advanced (source)

In this address type, a user can create an address based on multiple parameters, including HTTP method, User Agent, and HTTP header. Once created, the address can be selected as a source of a proxy policy. This means that a policy will only allow or block requests that match the selected address.

This example creates an address that uses the get method, a user agent for Google Chrome, and an HTTP header with the pattern *Q[A-B]*.

To create an advanced (source) address in the GUI:

1. Go to *Policy & Objects* > *Addresses*.
2. Click *Create New* > *Address*.
3. Set the following:
   - *Category* to *Proxy Address*,
   - *Name* to *advanced_src*,
   - *Type* to *Advanced (Source)*,
   - *Host* to *all*,
   - *Request Method* to *GET*,
   - *User Agent* to *Google Chrome*, and
   - *HTTP header* to *Header_Test* : *Q[A-B]*.



4. Click *OK*.

To create an advanced (source) address in the CLI:

```
config firewall proxy-address
    edit "advance_src"
        set type src-advanced
        set host "all"
        set method get
        set ua chrome
        config header-group
            edit 1
                set header-name "Header_Test"
                set header "Q[A-B]"
            next
        end
    next
end
```

Advanced (destination)

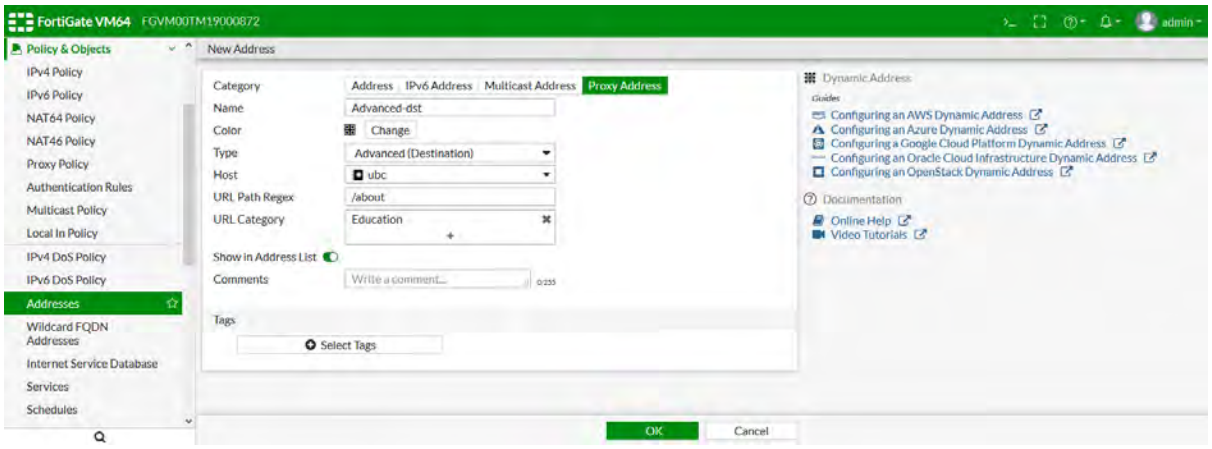
In this address type, a user can create an address based on URL pattern and URL category parameters. Once created, the address can be selected as a destination of a proxy policy. This means that a policy will only allow or block requests that match the selected address.

This example creates an address with the URL pattern */about* that are in the *Education* category. For more information about categories, see https://fortiguard.com/webfilter/categories.

To create an advanced (destination) address in the GUI:

1. Go to *Policy & Objects* > *Addresses*.
2. Click *Create New* > *Address*.
3. Set the following:
   - *Category* to *Proxy Address*,
   - *Name* to *Advanced-dst*,

- *Type* to *Advanced (Destination)*,
- *Host* to *all*,
- *URL Path Regex* to */about*, and
- *URL Category* to *Education*.



4. Click *OK*.

To create an advanced (destination) address in the CLI:

```
config firewall proxy-address
    edit "Advanced-dst"
        set type dst-advanced
        set host "ubc"
        set path "/about"
        set category 30
    next
end
```

Proxy policy security profiles

Web proxy policies support most security profile types.

Security profiles must be created before they can be used in a policy, see Security Profiles on page 851 for information.

Explicit web proxy policy

The security profiles supported by explicit web proxy policies are:

- *AntiVirus*
- *Web Filter*
- *Application Control*
- *IPS*
- *DLP Sensor*
- *ICAP*

- *Web Application Firewall*
- *SSL Inspection*

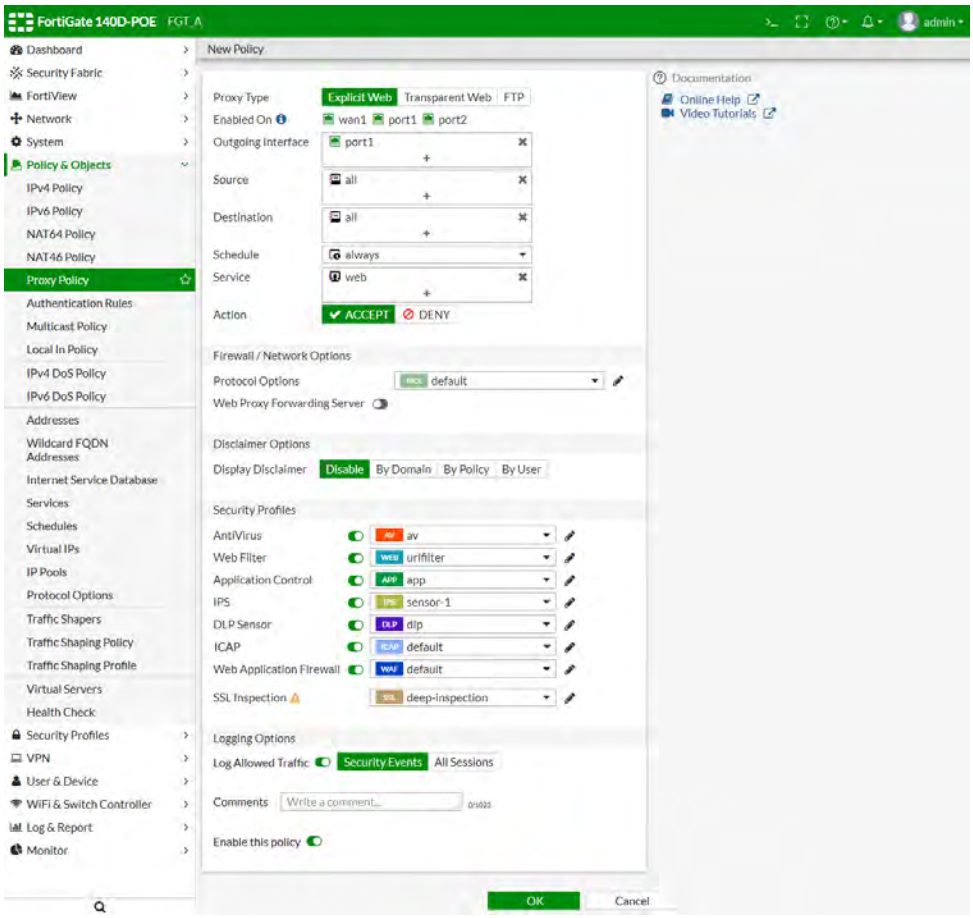
To configure security profiles on an explicit web proxy policy in the GUI:

1. Go to *Policy & Objects* > *Proxy Policy*.
2. Click *Create New*.
3. Set the following:

Proxy Type	Explicit Web
Outgoing Interface	port1
Source	all
Destination	all
Schedule	always
Service	webproxy
Action	ACCEPT

4. In the *Firewall / Network Options* section, set *Protocol Options* to *default*.
5. In the *Security Profiles* section, make the following selections (for this example, these profiles have all already been created):

AntiVirus	av
Web Filter	urlfiler
Application Control	app
IPS	Sensor-1
DLP Sensor	dlp
ICAP	default
Web Application Firewall	default
SSL Inspection	deep-inspection



6. Click *OK* to create the policy.

To configure security profiles on an explicit web proxy policy in the CLI:

```
config firewall proxy-policy
    edit 1
        set proxy explicit-web
        set dstintf "port1"
        set srcaddr "all"
        set dstaddr "all"
        set service "web"
        set action accept
        set schedule "always"
        set utm-status enable
        set av-profile "av"
        set webfilter-profile "urlfilter"
        set dlp-sensor "dlp"
        set ips-sensor "sensor-1"
        set application-list "app"
        set icap-profile "default"
        set waf-profile "default"
        set ssl-ssh-profile "deep-inspection"
    next
end
```

Transparent proxy

The security profiles supported by transparent proxy policies are:

- *AntiVirus*
- *Web Filter*
- *Application Control*
- *IPS*
- *DLP Sensor*
- *ICAP*
- *Web Application Firewall*
- *SSL Inspection*

To configure security profiles on a transparent proxy policy in the GUI:

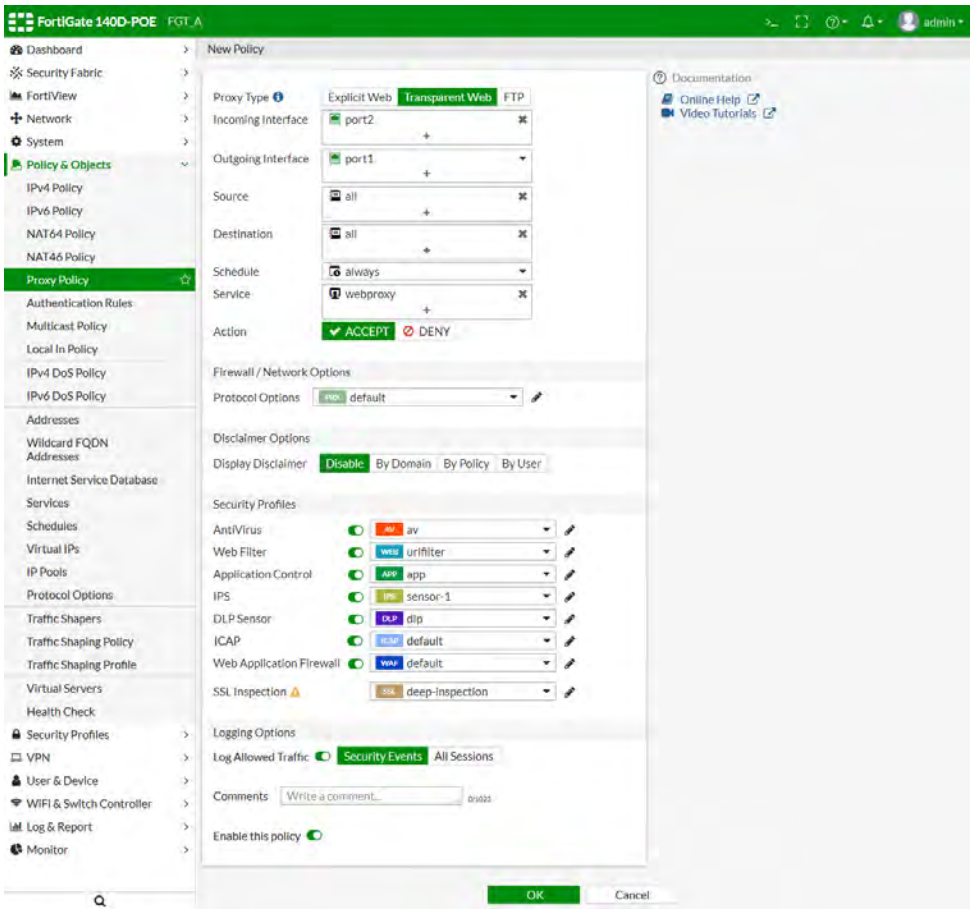
1. Go to *Policy & Objects* > *Proxy Policy*.
2. Click *Create New*.
3. Set the following:

Proxy Type	Explicit Web
Incoming Interfae	port2
Outgoing Interface	port1
Source	all
Destination	all
Schedule	always
Service	webproxy
Action	ACCEPT

4. In the *Firewall / Network Options* section, set *Protocol Options* to *default*.
5. In the *Security Profiles* section, make the following selections (for this example, these profiles have all already been created):

AntiVirus	av
Web Filter	urlfiler
Application Control	app
IPS	Sensor-1
DLP Sensor	dlp
ICAP	default
Web Application Firewall	default
SSL Inspection	deep-inspection

Netskope Exhibit 1017



6. Click *OK* to create the policy.

To configure security profiles on a transparent proxy policy in the CLI:

```
config firewall proxy-policy
    edit 2
        set proxy transparent-web
        set srcintf "port2"
        set dstintf "port1"
        set srcaddr "all"
        set dstaddr "all"
        set service "webproxy"
        set action accept
        set schedule "always"
        set utm-status enable
        set av-profile "av"
        set webfilter-profile "urlfilter"
        set dlp-sensor "dlp"
        set ips-sensor "sensor-1"
        set application-list "app"
        set icap-profile "default"
        set waf-profile "default"
        set ssl-ssh-profile "certificate-inspection"
    next
end
```

FTP proxy

The security profiles supported by FTP proxy policies are:

- *AntiVirus*
- *Application Control*
- *IPS*
- *DLP Sensor*

To configure security profiles on an FTP proxy policy in the GUI:

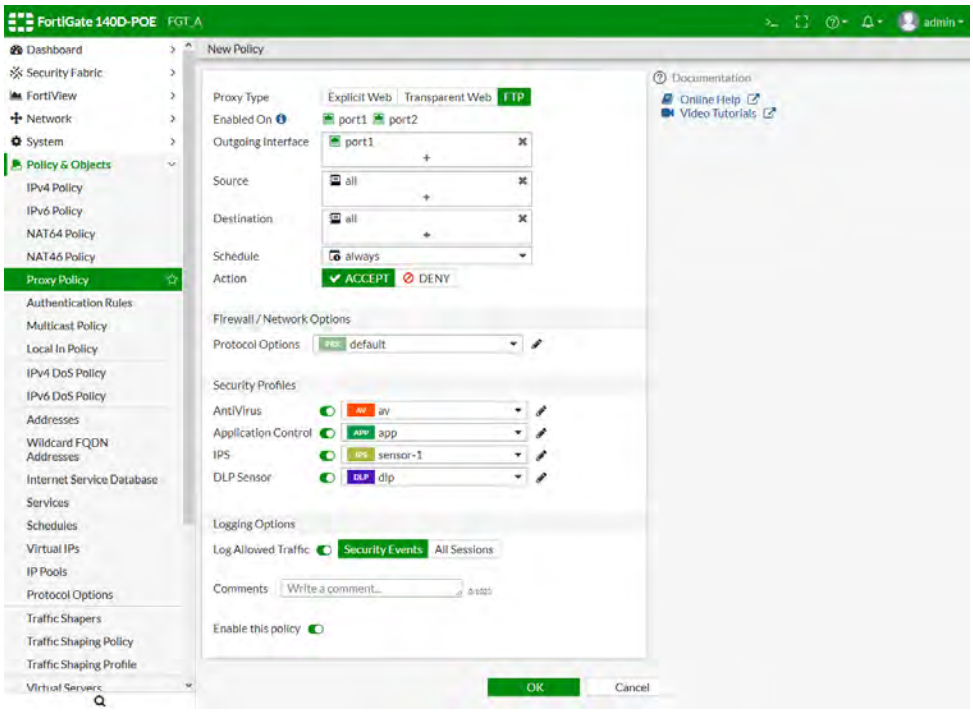
1. Go to *Policy & Objects* > *Proxy Policy*.
2. Click *Create New*.
3. Set the following:

Proxy Type	FTP
Outgoing Interface	port1
Source	all
Destination	all
Schedule	always
Action	ACCEPT

4. In the *Firewall / Network Options* section, set *Protocol Options* to *default*.
5. In the *Security Profiles* section, make the following selections (for this example, these profiles have all already been created):

AntiVirus	av
Application Control	app
IPS	Sensor-1
DLP Sensor	dlp

Netskope Exhibit 1017



6. Click *OK* to create the policy.

To configure security profiles on an FTP proxy policy in the CLI:

```
config firewall proxy-policy
    edit 3
        set proxy ftp
        set dstintf "port1"
        set srcaddr "all"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set utm-status enable
        set av-profile "av"
        set dlp-sensor "dlp"
        set ips-sensor "sensor-1"
        set application-list "app"
    next
end
```

Explicit proxy authentication

FortiGate supports multiple authentication methods. This topic explains using an external authentication server with Kerberos as the primary and NTLM as the fallback.

To configure Explicit Proxy with authentication:

1. Enable and configure the explicit proxy on page 406.
2. Configure the authentication server and create user groups on page 406.
3. Create an authentication scheme and rules on page 408.

4. Create an explicit proxy policy and assign a user group to the policy on page 409.
5. Verify the configuration on page 410.

Enable and configure the explicit proxy

To enable and configure explicit web proxy in the GUI:

1. Go to *Network* > *Explicit Proxy*.
2. Enable *Explicit Web Proxy*.
3. Select *port2* as the *Listen on Interfaces* and set the *HTTP Port* to *8080*.
4. Configure the remaining settings as needed.
5. Click *Apply*.

To enable and configure explicit web proxy in the CLI:

```
config web-proxy explicit
    set status enable
    set ftp-over-http enable
    set socks enable
    set http-incoming-port 8080
    set ipv6-status enable
    set unknown-http-version best-effort
end
config system interface
    edit "port2"
        set vdom "vdom1"
        set ip 10.1.100.1 255.255.255.0
        set allowaccess ping https ssh snmp http telnet
        set type physical
        set explicit-web-proxy enable
        set snmp-index 12
        end
    next
end
```

Configure the authentication server and create user groups

Since we are using an external authentication server with Kerberos authentication as the primary and NTLM as the fallback, Kerberos authentication is configured first and then FSSO NTLM authentication is configured.

For successful authorization, the FortiGate checks if user belongs to one of the groups that is permitted in the security policy.

To configure an authentication server and create user groups in the GUI:

1. Configure Kerberos authentication:
   a. Go to *User & Authentication* > *LDAP Servers*.
   b. Click *Create New*.

Netskope Exhibit 1017

c. Set the following:

Name	ldap-kerberos
Server IP	172.18.62.220
Server Port	389
Common Name Identifier	cn
Distinguished Name	dc=fortinetqa,dc=local

d. Click *OK*

2. Define Kerberos as an authentication service. This option is only available in the CLI.
3. Configure FSSO NTLM authentication:
FSSO NTLM authentication is supported in a Windows AD network. FSSO can also provide NTLM authentication service to the FortiGate unit. When a user makes a request that requires authentication, the FortiGate initiates NTLM negotiation with the client browser, but does not process the NTLM packets itself. Instead, it forwards all the NTLM packets to the FSSO service for processing.
   a. Go to *Security Fabric* > *External Connectors*.
   b. Click *Create New* and select *Fortinet Single Sign-On Agent* from the *Endpoint/Identity* category.
   c. Set the *Name* to *FSSO*, *Primary FSSO Agent* to *172.16.200.220*, and enter a password.
   d. Click *OK*.
4. Create a user group for Kerberos authentication:
   a. Go to *User & Authentication* > *User Groups*.
   b. Click *Create New*.
   c. Set the *Name* to *Ldap-Group*, and *Type* to *Firewall*.
   d. In the *Remote Groups* table, click *Add*, and set the *Remote Server* to the previously created *ldap-kerberos* server.
   e. Click *OK*.
5. Create a user group for NTLM authentication:
   a. Go to *User & Authentication* > *User Groups*.
   b. Click *Create New*.
   c. Set the *Name* to *NTLM-FSSO-Group*, *Type* to *Fortinet Single Sign-On (FSSO)*, and add *FORTINETQA/FSSO* as a member.
   d. Click *OK*.

To configure an authentication server and create user groups in the CLI:

1. Configure Kerberos authentication:

```
config user ldap
    edit "ldap-kerberos"
        set server "172.18.62.220"
        set cnid "cn"
        set dn "dc=fortinetqa,dc=local"
        set type regular
        set username "CN=root,CN=Users,DC=fortinetqa,DC=local"
        set password **********
    next
end
```


Netskope Exhibit 1017

2. Define Kerberos as an authentication service:

```
config user krb-keytab
    edit "http_service"
        set pac-data disable
        set principal "HTTP/FGT.FORTINETQA.LOCAL@FORTINETQA.LOCAL"
        set ldap-server "ldap-kerberos"
        set keytab
"BQIAAABFAAIAEEZPUlRJTkVUUUEuTE9DQUwABEhUVFAAFEZHVC5GT1JUSU5FVFFBLkxPQ0FMAAAAAQAAAAA
EAAEACKLCMonpitnVAAAARQACABBGT1JUSU5FVFFBLkxPQ0FMAARIVFRQABRGR1QuRk9SVElORVRRQS5MT0N
BTAAAAAEAAAAABAADAAiiwjKJ6YrZ1QAAAE0AAgAQRk9SVElORVRRQS5MT0NBTAAESFRUUAAURkdULkZPUlR
JTkVUUUEuTE9DQUwAAAABAAAAAAQAFwAQUHo9uqR9cSkzyxdzKCEXdwAAAF0AAgAQRk9SVElORVRRQS5MT0N
BTAAESFRUUAAURkdULkZPUlRJTkVUUUEuTE9DQUwAAAABAAAAAAQAEgAgzee854Aq1HhQiKJZvV4tL2Poy7h
MIARQpK8MCB//BIAAAABNAAIAEEZPUlRJTkVUUUEuTE9DQUwABEhUVFAAFEZHVC5GT1JUSU5FVFFBLkxPQ0F
MAAAAAQAAAAAEABEAEG49vHEiiBghr63Z/lnwYrU="
    next
end
```

3. Configure FSSO NTLM authentication:

```
config user fsso
    edit "1"
        set server "172.18.62.220"
        set password **********
    next
end
```

4. Create a user group for Kerberos authentication:

```
config user group
    edit "Ldap-Group"
        set member "ldap" "ldap-kerberos"
    next
end
```

5. Create a user group for NTLM authentication:

```
config user group
    edit "NTLM-FSSO-Group"
        set group-type fsso-service
        set member "FORTINETQA/FSSO"
    next
end
```

Create an authentication scheme and rules

Explicit proxy authentication is managed by authentication schemes and rules. An authentication scheme must be created first, and then the authentication rule.

To create an authentication scheme and rules in the GUI:

1. Create an authentication scheme:
   a. Go to *Policy & Objects* > *Authentication Rules*.
   b. Click *Create New* > *Authentication Schemes*.
   c. Set the *Name* to *Auth-scheme-Negotiate* and select *Negotiate* as the *Method*.
   d. Click *OK*.

Netskope Exhibit 1017

2. Create an authentication rule:
   a. Go to *Policy & Objects* > *Authentication Rules*.
   b. Click *Create New* > *Authentication Rules*.
   c. Set the *Name* to *Auth-Rule*, *Source Address* to *all*, and *Protocol* to *HTTP*.
   d. Enable *Authentication Scheme*, and select the just created *Auth-scheme-Negotiate* scheme.
   e. Click *OK*.

To create an authentication scheme and rules in the CLI:

1. Create an authentication scheme:

```
config authentication scheme
    edit "Auth-scheme-Negotiate"
        set method negotiate       <<< Accepts both Kerberos and NTLM as fallback
    next
end
```

2. Create an authentication rule:

```
config authentication rule
    edit "Auth-Rule"
        set status enable
        set protocol http
        set srcaddr "all"
        set ip-based enable
        set active-auth-method "Auth-scheme-Negotiate"
        set comments "Testing"
    next
end
```

Create an explicit proxy policy and assign a user group to the policy

To create an explicit proxy policy and assign a user group to it in the GUI:

1. Go to *Policy & Object* > *Proxy Policy*.
2. Click *Create New*.
3. Set *Proxy Type* to *Explicit Web* and *Outgoing Interface* to *port1*.
4. Set *Source* to *all*, and the just created user groups *NTLM-FSSO-Group* and *Ldap-Group*.
5. Also set *Destination* to *all*, *Schedule* to *always*, *Service* to *webproxy*, and *Action* to *ACCEPT*.
6. Click *OK*.

To create an explicit proxy policy and assign a user group to it in the CLI:

```
config firewall proxy-policy
    edit 1
        set proxy explicit-web
        set dstintf "port1"
        set srcaddr "all"
        set dstaddr "all"
        set service "web"
        set action accept
        set schedule "always"
        set logtraffic all
```

```
            set groups "NTLM-FSSO-Group" "Ldap-Group"
            set av-profile "av"
            set ssl-ssh-profile "deep-custom"
        next
    end
```

Verify the configuration

Log in using a domain and system that would be authenticated using the Kerberos server, then enter the `diagnose wad user list` CLI command to verify:

```
# diagnose wad user list
ID: 8, IP: 10.1.100.71, VDOM: vdom1
  user name    : test1@FORTINETQA.LOCAL
  duration     : 389
  auth_type    : IP
  auth_method : Negotiate
  pol_id       : 1
  g_id         : 1
  user_based   : 0
  expire       : no
  LAN:
    bytes_in=4862 bytes_out=11893
  WAN:
    bytes_in=7844 bytes_out=1023
```

Log in using a system that is not part of the domain. The NTLM fallback server should be used:

```
# diagnose wad user list
ID: 2, IP: 10.1.100.202, VDOM: vdom1
  user name    : TEST31@FORTINETQA
  duration     : 7
  auth_type    : IP
  auth_method : NTLM
  pol_id       : 1
  g_id         : 5
  user_based   : 0
  expire       : no
  LAN:
    bytes_in=6156 bytes_out=16149
  WAN:
    bytes_in=7618 bytes_out=1917
```

Transparent web proxy forwarding

In FortiOS, there is an option to enable proxy forwarding for transparent web proxy policies and regular firewall policies for HTTP and HTTPS.

In previous versions of FortiOS, you could forward proxy traffic to another proxy server (proxy chaining) with explicit proxy. Now, you can forward web traffic to the upstream proxy without having to reconfigure your browsers or publish a proxy auto-reconfiguration (PAC) file.

Once configured, the FortiGate forwards traffic generated by a client to the upstream proxy. The upstream proxy then forwards it to the server.

Netskope Exhibit 1017

To enable proxy forwarding using the CLI:

1. Configure the web proxy forwarding server:

```
config web-proxy forward-server
    edit "PC_03"
        set ip 172.16.200.46
        set healthcheck enable
        set monitor "http://www.google.ca"
    next
end
```

2. Append the web proxy forwarding server to a firewall policy:

```
config firewall policy
    edit 1
        set name "LAN to WAN"
        set srcintf "port2"
        set dstintf "port1"
        set srcaddr "all"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ALL"
        set utm-status enable
        set logtraffic all
        set webproxy-forward-server "PC_03"
        set fsso disable
        set av-profile "av"
        set ssl-ssh-profile "deep-custom"
        set nat enable
    next
end
```

Upstream proxy authentication in transparent proxy mode

A downstream proxy FortiGate that needs to be authenticated by the upstream web proxy can use the basic authentication method to send its username and password, in the base64 format, to the upstream web proxy for authentication. If the authentication succeeds, web traffic that is forwarded from the downstream proxy FortiGate to the upstream proxy can be accepted and forwarded to its destinations.

In this example, a school has a FortiGate acting as a downstream proxy that is configured with firewall policies for each user group (students and staff). In each policy, a forwarding server is configured to forward the web traffic to the upstream web proxy.

The username and password that the upstream web proxy uses to authenticate the downstream proxy are configured on the forwarding server, and are sent to the upstream web proxy with the forwarded HTTP requests.

	Username	Password
student.proxy.local:8080	students	ABC123
staff.proxy.local:8081	staff	123456

On the downstream FortiGate, configure forwarding servers with the usernames and passwords for authentication on the upstream web proxy, then apply those servers to firewall policies for transparent proxy. For explicit web proxy, the forwarding servers can be applied to proxy policies.

When the transparent proxy is configured, clients can access websites without configuring a web proxy in their browser. The downstream proxy sends the username and password to the upstream proxy with forwarded HTTP requests to be authenticated.

To configure the forwarding server on the downstream FortiGate:

```
config web-proxy forward-server
    edit "Student_Upstream_WebProxy"
        set addr-type fqdn
        set fqdn "student.proxy.local"
        set port 8080
        set username "student"
        set password ABC123
    next
    edit "Staff_Upstream_WebProxy"
        set addr-type fqdn
        set fqdn "staff.proxy.local"
        set port 8081
        set username "staff"
        set password 123456
    next
end
```

To configure firewall policies for transparent proxy:

```
config firewall policy
    edit 1
        set srcintf "Vlan_Student"
        set dstintf "port9"
        set srcaddr "Student_Subnet"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ALL"
        set utm-status enable
        set inspection-mode proxy
        set ssl-ssh-profile "deep-inspection"
        set av-profile "av"
        set webproxy-forward-server "Student_Upstream_WebProxy"
        set nat enable
    next
    edit 2
        set srcintf "Vlan_Staff"
        set dstintf "port9"
        set srcaddr "Staff_Subnet"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ALL"
        set utm-status enable
        set inspection-mode proxy
        set ssl-ssh-profile "deep-inspection"
```


Netskope Exhibit 1017

```
        set av-profile "av"
        set webproxy-forward-server "Staff_Upstream_WebProxy"
        set nat enable
    next
end
```

Multiple dynamic header count

Multiple dynamic headers are supported for web proxy profiles, as well as Base64 encoding and the append/new options.

Administrators only have to select the dynamic header in the profile. The FortiGate will automatically display the corresponding static value. For example, if the administrator selects the `$client-ip` header, the FortiGate will display the actual client IP address.

The supported headers are:

`$client-ip`	Client IP address
`$user`	Authentication user name
`$domain`	User domain name
`$local_grp`	Firewall group name
`$remote_grp`	Group name from authentication server
`$proxy_name`	Proxy realm name

To configure dynamic headers using the CLI:

Since authentication is required, FSSO NTLM authentication is configured in this example.

1. Configure LDAP:

```
config user ldap
    edit "ldap-kerberos"
        set server "172.18.62.220"
        set cnid "cn"a
        set dn "dc=fortinetqa,dc=local"
        set type regular
        set username "CN=root,CN=Users,DC=fortinetqa,DC=local"
        set password **********
    next
end
```

2. Configure FSSO:

```
config user fsso
    edit "1"
        set server "172.18.62.220"
        set password **********
    next
end
```

3. Configure a user group:

```
config user group
    edit "NTLM-FSSO"
        set group-type fsso-service
        set member "FORTINETQA/FSSO"
    next
end
```

4. Configure an authentication scheme:

```
config authentication scheme
    edit "au-sch-ntlm"
        set method ntlm
    next
end
```

5. Configure an authentication rule:

```
config authentication rule
    edit "au-rule-fsso"
        set srcaddr "all"
        set active-auth-method "au-sch-ntlm"
    next
end
```

6. Create a web proxy profile that adds a new dynamic and custom `Via` header:

```
config web-proxy profile
    edit "test"
        set log-header-change enable
        config headers
            edit 1
                set name "client-ip"
                set content "$client-ip"
            next
            edit 2
                set name "Proxy-Name"
                set content "$proxy_name"
            next
            edit 3
                set name "user"
                set content "$user"
            next
            edit 4
                set name "domain"
                set content "$domain"
            next
            edit 5
                set name "local_grp"
                set content "$local_grp"
            next
            edit 6
                set name "remote_grp"
                set content "$remote_grp"
            next
            edit 7
                set name "Via"
                set content "Fortigate-Proxy"
```


Netskope Exhibit 1017

```
            next
        end
    next
end
```

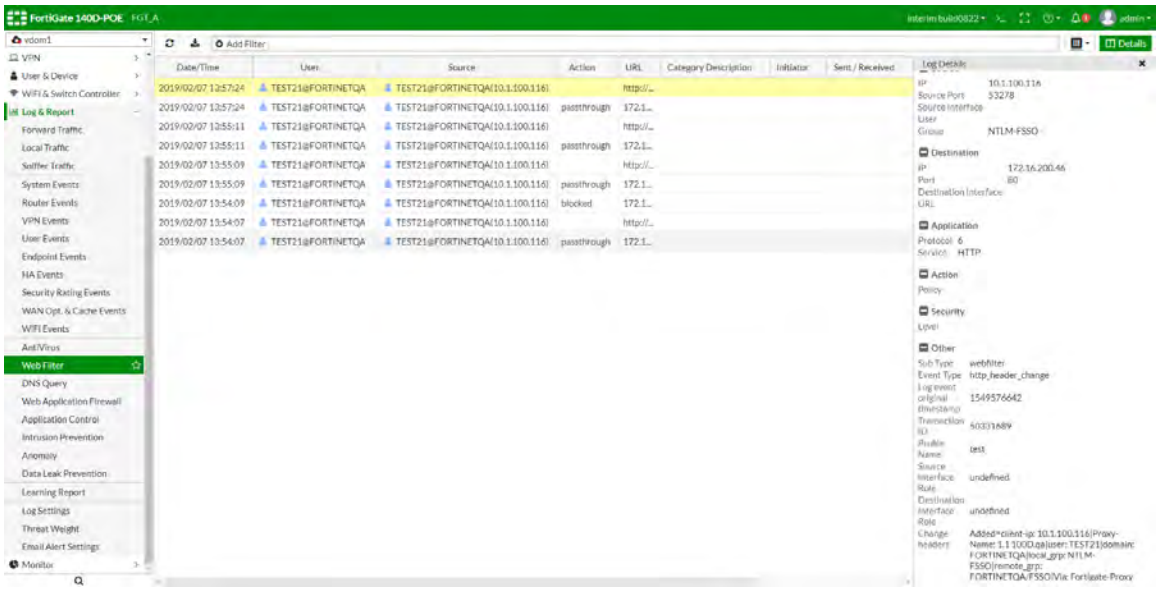
7. In the proxy policy, append the web proxy profile created in the previous step:

```
config firewall proxy-policy
    edit 1
        set proxy explicit-web
        set dstintf "port1"
        set srcaddr "all"
        set dstaddr "all"
        set service "web"
        set action accept
        set schedule "always"
        set logtraffic all
        set groups "NTLM-FSSO"
        set webproxy-profile "test"
        set utm-status enable
        set av-profile "av"
        set webfilter-profile "content"
        set ssl-ssh-profile "deep-custom"
    next
end
```

8. Once traffic is being generated from the client, look at the web filter logs to verify that it is working.
The corresponding values for all the added header fields displays in the *Change headers* section at the bottom of the *Log Details* pane.

```
1: date=2019-02-07 time=13:57:24 logid="0344013632" type="utm" subtype="webfilter"
eventtype="http_header_change" level="notice" vd="vdom1" eventtime=1549576642 policyid=1
transid=50331689 sessionid=1712788383 user="TEST21@FORTINETQA" group="NTLM-FSSO"
profile="test" srcip=10.1.100.116 srcport=53278 dstip=172.16.200.46 dstport=80
srcintf="port2" srcintfrole="undefined" dstintf="port1" dstintfrole="undefined" proto=6
service="HTTP" url="http://172.16.200.46/" agent="curl/7.22.0" chgheaders="Added=client-
ip: 10.1.100.116|Proxy-Name: 1.1 100D.qa|user: TEST21|domain: FORTINETQA|local_grp:
NTLM-FSSO|remote_grp: FORTINETQA/FSSO|Via: Fortigate-Proxy"
```


Netskope Exhibit 1017



Restricted SaaS access (Office 365, G Suite, Dropbox)

With the web proxy profile, you can specify access permissions for Microsoft Office 365, Google G Suite, and Dropbox. You can insert vendor-defined headers that restrict access to the specific accounts. You can also insert custom headers for any destination.

You can configure the web proxy profile with the required headers for the specific destinations, and then directly apply it to a policy to control the header's insertion.

To implement Office 365 tenant restriction, G Suite account access control, and Dropbox network access control:

1. Configure a web proxy profile according to the vendors' specifications:
   a. Define the traffic destination (service provider).
   b. Define the header name, defined by the service provider.
   c. Define the value that will be inserted into the traffic, defined by your settings.
2. Apply the web proxy profile to a policy.

The following example creates a web proxy profile for Office 365, G Suite, and Dropbox access control.

Due to vendors' changing requirements, this example may no longer comply with the vendors' official guidelines.

To create a web proxy profile for access control using the CLI:

1. Configure the web proxy profile:

```
config web-proxy profile
   edit "SaaS-Tenant-Restriction"
        set header-client-ip pass
```

```
        set header-via-request pass
        set header-via-response pass
        set header-x-forwarded-for pass
        set header-front-end-https pass
        set header-x-authenticated-user pass
        set header-x-authenticated-groups pass
        set strip-encoding disable
        set log-header-change disable
        config headers
            edit 1
                set name "Restrict-Access-To-Tenants"  <---header name defined by
Office365 spec. input EXACTLY as it is
                set dstaddr "Microsoft Office 365" <----built-in destination address for
Office365
                set action add-to-request
                set base64-encoding disable
                set add-option new
                set protocol https http
                set content "contoso.onmicrosoft.com,fabrikam.onmicrosoft.com"  <----
your tenants restriction configuration
            next
            edit 2
                set name "Restrict-Access-Context"  <----header name defined by
Office365 spec. input EXACTLY as it is
                set dstaddr "Microsoft Office 365"  <----build-in destination address
for Office365
                set action add-to-request
                set base64-encoding disable
                set add-option new
                set protocol https http
                set content "456ff232-35l2-5h23-b3b3-3236w0826f3d" <----your directory
ID can find in Azure portal
            next
            edit 3
                set name "X-GooGApps-Allowed-Domains"  <----header name defined by
Google G suite.
                set dstaddr "G Suite"  <---- built-in G Suite destination address
                set action add-to-request
                set base64-encoding disable
                set add-option new
                set protocol https http
                set content "abcd.com"    <----your domain restriction when you create G
Suite account
            next

            edit 4
                set name "X-Dropbox-allowed-Team-Ids" <----header defined by Dropbox
                set dstaddr "wildcard.dropbox.com"  <----build-in destination address
for Dropbox
                set action add-to-request
                set base64-encoding disable
                set add-option new
                set protocol https http
                set content "dbmid:FDFSVF-DFSDF"  <----your team-Id in Dropbox
            next
        end
```


Netskope Exhibit 1017

```
    next
end
```

2. Apply the web proxy profile to a firewall policy:

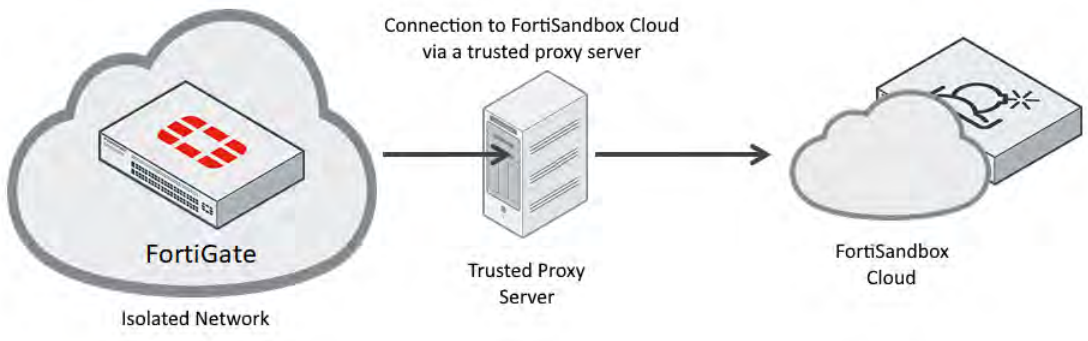
```
config firewall policy
     edit 1
                set name "WF"
                set srcintf "port10" "wifi"
                set dstintf "port9"
                set srcaddr "all"
                set dstaddr "all"
                set action accept
                set schedule "always"
                set service "ALL"
                set webproxy-profile "SaaS-Tenant-Restriction"
                set utm-status enable
                set utm-inspection-mode proxy
                set logtraffic all
                set webfilter-profile "blocktest2"
          set application-list "g-default"
                set profile-protocol-options "protocol"
                set ssl-ssh-profile "protocols"
                set nat enable
        next
end
```

References

- Office 365: Use tenant restrictions to manage access to SaaS cloud applications
- G Suite: Block access to consumer accounts
- Dropbox: Network control

Explicit proxy and FortiSandbox Cloud

Explicit proxy connections can leverage FortiSandbox Cloud for advanced threat scanning and updates. This allows FortiGates behind isolated networks to connect to FortiCloud services.



To configure FortiGuard services to communicate with an explicit proxy server:

```
config system fortiguard
    set proxy-server-ip 172.16.200.44
    set proxy-server-port 3128
    set proxy-username "test1"
    set proxy-password ENC
Y0+KTg9UsILkv8+nDe+Pe3VlnlaHUMzLkfAXLATknW/xm/Xv7EdZHTnua1djM+waZA1vxCh8LV7Ci4sEhj/PABSTShSt
xskEn3E1+CjxviwVSljgF6AD+zJZF/+4jkspq+PogZT3LVO68+kqsPdU4rikuy1BbnsbZcPxC/MJyuIx7343bdKYqp+I
UprQUR2wf8tiMg==
end
```

To verify the explicit proxy connection to FortiSandbox Cloud:

```
# diagnose debug application forticldd -1
Debug messages will be on for 30 minutes.
# diagnose debug enable
[2942] fds_handle_request: Received cmd 23 from pid-2526, len 0
[40] fds_queue_task: req-23 is added to Cloud-sandbox-controller
[178] fds_svr_default_task_xmit: try to get IPs for Cloud-sandbox-controller
[239] fds_resolv_addr: resolve aptctrl1.fortinet.com
[169] fds_get_addr: name=aptctrl1.fortinet.com, id=32, cb=0x2bc089
[101] dns_parse_resp: DNS aptctrl1.fortinet.com -&gt; 172.16.102.21
[227] fds_resolv_cb: IP-1: 172.16.102.21
[665] fds_ctx_set_addr: server: 172.16.102.21:443
[129] fds_svr_default_pickup_server: Cloud-sandbox-controller: 172.16.102.21:443
[587] fds_https_start_server: server: 172.16.102.21:443
[579] ssl_new: SSL object is created
[117] https_create: proxy server 172.16.200.44 port:3128
[519] fds_https_connect: https_connect(172.16.102.21) is established.
[261] fds_svr_default_on_established: Cloud-sandbox-controller has connected to
ip=172.16.102.21
[268] fds_svr_default_on_established: server-Cloud-sandbox-controller handles cmd-23
[102] fds_pack_objects: number of objects: 1
[75] fds_print_msg: FCPC: len=109
[81] fds_print_msg: Protocol=2.0
[81] fds_print_msg: Command=RegionList
[81] fds_print_msg: Firmware=FG101E-FW-6.02-0917
[81] fds_print_msg: SerialNumber=FG101E4Q17002429
[81] fds_print_msg: TimeZone=-7
[75] fds_print_msg: http req: len=248
[81] fds_print_msg: POST https://172.16.102.21:443/FCPService HTTP/1.1
[81] fds_print_msg: User-Agent: Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.1)
[81] fds_print_msg: Host: 172.16.102.21:443
[81] fds_print_msg: Cache-Control: no-cache
[81] fds_print_msg: Connection: close
[81] fds_print_msg: Content-Type: application/octet-stream
[81] fds_print_msg: Content-Length: 301
[524] fds_https_connect: http request to 172.16.102.21: header=248, ext=301.
[257] fds_https_send: sent 248 bytes: pos=0, len=248
[265] fds_https_send: 172.16.102.21: sent 248 byte header, now send 301-byte body
[257] fds_https_send: sent 301 bytes: pos=0, len=301
[273] fds_https_send: sent the entire request to server: 172.16.102.21:443
[309] fds_https_recv: read 413 bytes: pos=413, buf_len=2048
[332] fds_https_recv: received the header from server: 172.16.102.21:443, [HTTP/1.1 200
Content-Type: application/octet-stream
```


Netskope Exhibit 1017

```
Content-Length: 279
Date: Thu, 20 Jun 2019 16:41:11 GMT
Connection: close]
[396] fds_https_recv: Do memmove buf_len=279, pos=279
[406] fds_https_recv: server: 172.16.102.21:443, buf_len=279, pos=279
[453] fds_https_recv: received a packet from server-172.16.102.21:443: sz=279, objs=1
[194] __ssl_data_ctx_free: Done
[839] ssl_free: Done
[830] ssl_disconnect: Shutdown
[481] fds_https_recv: obj-0: type=FCPR, len=87
[294] fds_svr_default_on_response: server-Cloud-sandbox-controller handles cmd-23
[75] fds_print_msg: fcpr:  len=83
[81] fds_print_msg: Protocol=2.0
[81] fds_print_msg: Response=202
[81] fds_print_msg: ResponseItem=Region:Europe,Global,Japan,US
[81] fds_print_msg: existing:Japan
[3220] aptctrl_region_res: Got rsp: Region:Europe,Global,Japan,US
[3222] aptctrl_region_res: Got rsp: Region existing:Japan
[439] fds_send_reply: Sending 28 bytes data.
[395] fds_free_tsk: cmd=23; req.noreply=1
# [136] fds_on_sys_fds_change: trace
[2942] fds_handle_request: Received cmd 22 from pid-170, len 0
[40] fds_queue_task: req-22 is added to Cloud-sandbox-controller
[587] fds_https_start_server: server: 172.16.102.21:443
[579] ssl_new: SSL object is created
[117] https_create: proxy server 172.16.200.44 port:3128
[519] fds_https_connect: https_connect(172.16.102.21) is established.
[261] fds_svr_default_on_established: Cloud-sandbox-controller has connected to
ip=172.16.102.21
[268] fds_svr_default_on_established: server-Cloud-sandbox-controller handles cmd-22
[102] fds_pack_objects: number of objects: 1
[75] fds_print_msg: FCPC: len=146
[81] fds_print_msg: Protocol=2.0
[81] fds_print_msg: Command=UpdateAPT
[81] fds_print_msg: Firmware=FG101E-FW-6.02-0917
[81] fds_print_msg: SerialNumber=FG101E4Q17002429
[81] fds_print_msg: TimeZone=-7
[81] fds_print_msg: TimeZoneInMin=-420
[81] fds_print_msg: DataItem=Region:US
[75] fds_print_msg: http req: len=248
[81] fds_print_msg: POST https://172.16.102.21:443/FCPService HTTP/1.1
[81] fds_print_msg: User-Agent: Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.1)
[81] fds_print_msg: Host: 172.16.102.21:443
[81] fds_print_msg: Cache-Control: no-cache
[81] fds_print_msg: Connection: close
[81] fds_print_msg: Content-Type: application/octet-stream
[81] fds_print_msg: Content-Length: 338
[524] fds_https_connect: http request to 172.16.102.21: header=248, ext=338.
[257] fds_https_send: sent 248 bytes: pos=0, len=248
[265] fds_https_send: 172.16.102.21: sent 248 byte header, now send 338-byte body
[257] fds_https_send: sent 338 bytes: pos=0, len=338
[273] fds_https_send: sent the entire request to server: 172.16.102.21:443
[309] fds_https_recv: read 456 bytes: pos=456, buf_len=2048
[332] fds_https_recv: received the header from server: 172.16.102.21:443, [HTTP/1.1 200
Content-Type: application/octet-stream
Content-Length: 322
```


Netskope Exhibit 1017

```
Date: Thu, 20 Jun 2019 16:41:16 GMT
Connection: close]
[396] fds_https_recv: Do memmove buf_len=322, pos=322
[406] fds_https_recv: server: 172.16.102.21:443, buf_len=322, pos=322
[453] fds_https_recv: received a packet from server-172.16.102.21:443: sz=322, objs=1
[194] __ssl_data_ctx_free: Done
[839] ssl_free: Done
[830] ssl_disconnect: Shutdown
[481] fds_https_recv: obj-0: type=FCPR, len=130
[294] fds_svr_default_on_response: server-Cloud-sandbox-controller handles cmd-22
[75] fds_print_msg: fcpr:  len=126
[81] fds_print_msg: Protocol=2.0
[81] fds_print_msg: Response=202
[81] fds_print_msg: ResponseItem=Server1:172.16.102.51:514
[81] fds_print_msg: Server2:172.16.102.52:514
[81] fds_print_msg: Contract:20210215
[81] fds_print_msg: NextRequest:86400
[615] parse_apt_contract_time_str: The APTContract is valid to Mon Feb 15 23:59:59 2021
[616] parse_apt_contract_time_str: FGT current local time is Thu Jun 20 09:41:16 2019
[3289] aptctrl_update_res: Got rsp: APT=172.16.102.51:514 APTAlter=172.16.102.52:514 next-
upd=86400
[395] fds_free_tsk: cmd=22; req.noreply=1
```

Proxy chaining (web proxy forwarding servers)

For the explicit web proxy you can configure web proxy forwarding servers to use proxy chaining to redirect web proxy sessions to other proxy servers. Proxy chaining can be used to forward web proxy sessions from the FortiGate unit to one or more other proxy servers on your network or on a remote network. You can use proxy chaining to integrate the FortiGate explicit web proxy with a web proxy solution that you already have in place.

A FortiGate unit can forward sessions to most web proxy servers including a remote FortiGate unit with the explicit web proxy enabled. No special configuration of the explicit web proxy on the remote FortiGate unit is required.

You can deploy the explicit web proxy with proxy chaining in an enterprise environment consisting of small satellite offices and a main office. If each office has a FortiGate unit, users at each of the satellite offices can use their local FortiGate unit as an explicit web proxy server. The satellite office FortiGate units can forward explicit web proxy sessions to an explicit web proxy server at the central office. From here the sessions can connect to web servers on the Internet.

FortiGate proxy chaining does not support web proxies in the proxy chain authenticating each other.

The following examples assume explicit web proxy has been enabled.

To enable explicit web proxy in the GUI:

1. Go to *System* > *Feature Visibility*.
2. In the *Security Features* column, enable *Explicit Proxy*.
3. Configure the explicit web proxy settings. See Explicit web proxy on page 383.

To add a web proxy forwarding server in the GUI:

1. Go to *Network* > *Explicit Proxy*. The *Explicit Proxy* page opens.
2. In the *Web Proxy Forwarding Servers* section, click *Create New*.

Netskope Exhibit 1017

3. Configure the server settings and click *OK*.

Name	Enter the name of the forwarding server.
Proxy Address Type	Select the type of IP address of the forwarding server. A forwarding server can have an *FQDN* or *IP* address.
Proxy Address	Enter the IP address of the forwarding server.
Port	Enter the port number on which the proxy receives connections. Traffic leaving the FortiGate explicit web proxy for this server has its destination port number changed to this number.
Server Down Action	Select the action the explicit web proxy will take if the forwarding server is down. • *Block*: Blocks the traffic if the remote server is down. • *Use Original Server*: Forwards the traffic from the FortiGate to its destination as if no forwarding server is configured.
Health Monitor	Select to enable health check monitoring.
Health Check Monitor Site	Enter the address of a remote site.

Example

The following example adds a web proxy forwarding server named `fwd-srv` at address `proxy.example.com` and port `8080`.

To add a web proxy forwarding server in the CLI:

```
config web-proxy forward-server
   edit fwd-srv
      set addr-type fqdn
      set fqdn proxy.example.com
      set port 8080
   end
```

Web proxy forwarding server monitoring and health checking

By default, a FortiGate unit monitors a web proxy forwarding server by forwarding a connection to the remote server every 10 seconds. The remote server is assumed to be down if it does not respond to the connection. FortiGate continues checking the server. The server is assumed to be back up when the server sends a response. If you enable health checking, the FortiGate unit attempts to get a response from a web server every 10 seconds by connecting through the remote forwarding server.

You can configure health checking for each remote server and specify a different website to check for each one.

If the remote server is found to be down you can configure the FortiGate unit to block sessions until the server comes back up or to allow sessions to connect to their destination, bypassing the remote forwarding server. You cannot configure the FortiGate unit to fail over to another remote forwarding server.

To configure proxy server monitor and health checking in the GUI:

1. Go to *Network* > *Explicit Proxy*. The *Explicit Proxy* page opens.
2. In the *Web Proxy Forwarding Servers* section, select a server and click *Editi*.

Netskope Exhibit 1017

3. Configure the *Server Down Action* and *Health Monitor* settings.

Server Down Action	Select the action the explicit web proxy will take if the forwarding server is down. • *Block*: Blocks the traffic if the remote server is down. • *Use Original Server*: Forwards the traffic from the FortiGate to its destination as if no forwarding server configured.
Health Monitor	Select to enable health check monitoring.
Health Check Monitor Site	Enter the address of a remote site.

4. Click *OK*.

Example

The following example enables health checking for a web proxy forwarding server and sets the server down option to bypass the forwarding server if it is down.

To configure proxy server monitor and health checking in the CLI:

```
config web-proxy forward-server
   edit fwd-srv
      set healthcheck enable
      set monitor http://example.com
      set server-down-option pass
   end
```

Grouping forwarding servers and load balancing traffic to the servers

You can add multiple web proxy forwarding servers to a forwarding server group and then add the server group to an explicit web proxy policy instead of adding a single server. Forwarding server groups are created from the FortiGate CLI but can be added to policies from the web-based manager (or from the CLI).

When you create a forwarding server group you can select a load balancing method to control how sessions are load balanced to the forwarding servers in the server group. Two load balancing methods are available:

- *Weighted* load balancing sends more sessions to the servers with higher weights. You can configure the weight for each server when you add it to the group.
- *Least-session* load balancing sends new sessions to the forwarding server that is processing the fewest sessions.

When you create a forwarding server group you can also enable *affinity*. Enable affinity to have requests from the same client processed by the same server. This can reduce delays caused by using multiple servers for a single multi-step client operation. Affinity takes precedence over load balancing.

You can also configure the behavior of the group if all of the servers in the group are down. You can select to block traffic or you can select to have the traffic pass through the FortiGate explicit proxy directly to its destination instead of being sent to one of the forwarding servers.

Example

The following example adds a forwarding server group that uses weighted load balancing to load balance traffic to three forwarding servers. Server weights are configured to send most traffic to `server2`. The group has `affinity` enabled and blocks traffic if all of the forward servers are down.

Netskope Exhibit 1017

To configure load balancing in the CLI:

```
config web-proxy forward-server
   edit server_1
      set ip 172.20.120.12
      set port 8080
   next
   edit server_2
      set ip 172.20.120.13
      set port 8000
   next
   edit server_3
      set ip 172.20.120.14
      set port 8090
   next
   end
config web-proxy forward-server-group
   edit New-fwd-group
      set affinity enable
      set ldb-method weighted
      set group-down-option block
         config server-list
            edit server_1
               set weight 10
            next
            edit server_2
               set weight 40
            next
            edit server_3
               set weight 10
            next
   end
```

Adding proxy chaining to an explicit web proxy policy

You can enable proxy chaining for web proxy sessions by adding a web proxy forwarding server or server group to an explicit web proxy policy. In a policy you can select one web proxy forwarding server or server group. All explicit web proxy traffic accepted by this security policy is forwarded to the specified web proxy forwarding server or server group.

To add an explicit web proxy forwarding server in the GUI:

1. Go to *Policy & Objects* > *Proxy Policy* and click *Create New*.
2. Configure the policy settings.
3.

Proxy Type	Explicit Web
Outgoing Interface	wan1
Source	Internal_subnet
Destination	all
Schedule	always
Service	webproxy
Action	Accept

4. Enable *Web Proxy Forwarding Server* and select the forwarding server, (for example,*fwd-srv*).
5. Click *OK*.

Example

The following example adds a security policy that allows all users on the `10.31.101.0` subnet to use the explicit web proxy for connections through the `wan1` interface to the Internet. The policy forwards web proxy sessions to a remote forwarding server named *fwd-srv*.

To add an explicit web proxy forwarding server in the CLI:

```
config firewall proxy-policy
   edit o
      set proxy explicit-web
      set dstintf "wan1"
      set srcaddr "Internal_subnet"
      set dstaddr "all"
      set service "webproxy"
      set action accept
      set schedule "always"
      set webproxy-forward-server "fwd-srv"
   end
```

Agentless NTLM authentication for web proxy

In FortiOS, agentless Windows NT LAN Manager (NTLM) authentication includes support for the following items:

- Multiple servers
- Individual users

You can use multiple domain controller servers for the agentless NTLM. They can be used for load balancing and high service stability.

You can also use user-based matching in groups for Kerberos and agentless NTLM. In these scenarios, FortiOS matches the user's group information from an LDAP server.

To support multiple domain controllers for agentless NTLM using the CLI:

1. Configure an LDAP server:
```
config user ldap
   edit "ldap-kerberos"
      set server "172.18.62.177"
      set cnid "cn"
      set dn "dc=fortinetqa,dc=local"
      set type regular
      set username "CN=root,CN=Users,DC=fortinetqa,DC=local"
      set password **********
   next
end
```
2. Configure multiple domain controllers:
```
config user domain-controller
   edit "dc1"
      set ip-address 172.18.62.177
      config extra-server
```

```
        edit 1
          set ip-address 172.18.62.220
        next
      end
      set ldap-server "ldap-kerberos"
  next
end
```

3. Create an authentication scheme and rule:

```
config authentication scheme
  edit "au-ntlm"
    set method ntlm
    set domain-controller "dc1"
  next
end
config authentication rule
  edit "ru-ntlm"
    set srcaddr "all"
    set ip-based disable
    set active-auth-method "au-ntlm"
  next
end
```

4. In the proxy policy, append the user group for authorization:

```
config firewall proxy-policy
  edit 1
    set proxy explicit-web
    set dstintf "port1"
    set srcaddr "all"
    set dstaddr "all"
    set service "web"
    set action accept
    set schedule "always"
    set groups "ldap-group"
    set utm-status enable
    set av-profile "av"
    set ssl-ssh-profile "deep-custom"
  next
end
```

   This configuration uses a round-robin method. When the first user logs in, the FortiGate sends the authentication request to the first domain controller. Later when another user logs in, the FortiGate sends the authentication request to another domain controller.

5. Verify the behavior after the user successfully logs in:

```
# diagnose wad user list
ID: 1825, IP: 10.1.100.71, VDOM: vdom1
  user name : test1
  duration : 497
  auth_type : Session
  auth_method : NTLM
  pol_id : 1 g_id : 5
  user_based : 0 e
  xpire : 103
  LAN:
    bytes_in=2167 bytes_out=7657
  WAN:
    bytes_in=3718 bytes_out=270
```


Netskope Exhibit 1017

To support individual users for agentless NTLM using the CLI:

1. Configure an LDAP server:
```
config user ldap
   edit "ldap-kerberos"
      set server "172.18.62.177"
      set cnid "cn"
      set dn "dc=fortinetqa,dc=local"
      set type regular
      set username "CN=root,CN=Users,DC=fortinetqa,DC=local"
      set password **********
   next
end
```
2. Configure the user group and allow user-based matching:
```
config user group
   edit "ldap-group"
      set member "ldap" "ldap-kerberos"
      config match
         edit 1
            set server-name "ldap-kerberos"
            set group-name "test1"
         next
      end
   next
end
```
3. Create an authentication scheme and rule:
```
config authentication scheme
   edit "au-ntlm"
      set method ntlm
      set domain-controller "dc1"
   next
end
config authentication rule
   edit "ru-ntlm"
      set srcaddr "all"
      set ip-based disable
      set active-auth-method "au-ntlm"
   next
end
```
4. In the proxy policy, append the user group for authorization:
```
config firewall proxy-policy
   edit 1
      set proxy explicit-web
      set dstintf "port1"
      set srcaddr "all"
      set dstaddr "all"
      set service "web"
      set action accept
      set schedule "always"
      set groups "ldap-group"
      set utm-status enable
      set av-profile "av"
      set ssl-ssh-profile "deep-custom"
   next
end
```


Netskope Exhibit 1017

This implementation lets you configure a single user instead of a whole group. The FortiGate will now allow the user named `test1`.

To verify the configuration using the CLI:

```
diagnose wad user list
   ID: 1827, IP: 10.1.15.25, VDOM: vdom1
   user name : test1
   duration : 161
   auth_type : Session
   auth_method : NTLM
   pol_id : 1
   g_id : 5
   user_based : 0
   expire : 439
   LAN:
      bytes_in=1309 bytes_out=4410
   WAN:
      bytes_in=2145 bytes_out=544
```

DHCP server

A DHCP server provides an address from a defined address range to a client on the network, when requested.

You can configure one or more DHCP servers on any FortiGate interface. A DHCP server dynamically assigns IP addresses to hosts on the network connected to the interface. The host computers must be configured to obtain their IP addresses using DHCP.

You can configure a FortiGate interface as a DHCP relay. The interface forwards DHCP requests from DHCP clients to an external DHCP server and returns the responses to the DHCP clients. The DHCP server must have appropriate routing so that its response packets to the DHCP clients arrive at the unit.

For more information about options, see:

- DHCP options on page 429
- IP address assignment with relay agent information option on page 430
- DHCP client options on page 432

Configure DHCP on the FortiGate

To add a DHCP server on the GUI:

1. Go to *Network* > *Interfaces*.
2. Edit an interface.
3. Enable the *DHCP Server* option and configure the settings.

To add a DHCP server on the CLI:

```
config system dhcp server
    edit 1
```


Netskope Exhibit 1017

```
        set dns-service default
        set default-gateway 192.168.1.2
        set netmask 255.255.255.0
        set interface "port1"
        config ip-range
            edit 1
                set start-ip 192.168.1.1
                set end-ip 192.168.1.1
            next
            edit 2
                set start-ip 192.168.1.3
                set end-ip 192.168.1.254
            next
        end
        set timezone-option default
        set tftp-server "172.16.1.2"
    next
end
```

DHCP options

When adding a DHCP server, you can include DHCP codes and options. The DHCP options are BOOTP vendor information fields that provide additional vendor-independent configuration parameters to manage the DHCP server. For example, you might need to configure a FortiGate DHCP server that gives out a separate option as well as an IP address, such as an environment that needs to support PXE boot with Windows images.

The option numbers and codes are specific to the application. The documentation for the application indicates the values to use. Option codes are represented in a option value/HEX value pairs. The option is a value between 1 and 255.

You can add up to three DHCP code/option pairs per DHCP server.

For detailed information about DHCP options, see RFC 2132, DHCP Options and BOOTP Vendor Extensions.

To configure option 252 with value http://192.168.1.1/wpad.dat using the CLI:

```
config system dhcp server
    edit <server_entry_number>
        set option1 252 687474703a2f2f3139322e3136382e312e312f777061642e646174
    next
end
```

Option 82

The DHCP relay agent information option (option 82 in RFC 3046) helps protect the FortiGate against attacks such as spoofing (forging) of IP addresses and MAC addresses, and DHCP IP address starvation.

This option is disabled by default. However, when `dhcp-relay-service` is enabled, `dhcp-relay-agent-option` becomes enabled.

To configure the DHCP relay agent option using the CLI:

```
config system interface
    edit <interface>
```

```
        set vdom root
        set dhcp-relay-service enable
        set dhcp-relay-ip <ip>
        set dhcp-relay-agent-option enable
        set vlanid <id>
    next
end
```

See IP address assignment with relay agent information option on page 430 for an example.

Option 42

This option specifies a list of the NTP servers available to the client by IP address.

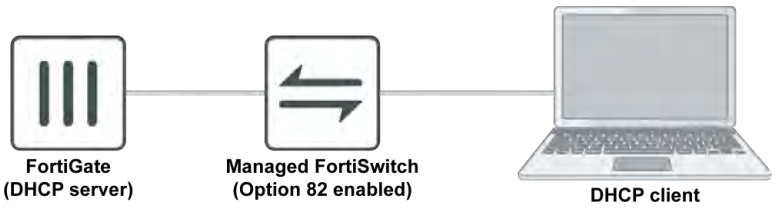
```
config system dhcp server
    edit 2
        set ntp-service {local | default | specify}
        set ntp-server1 <class_ip>
        set ntp-server2 <class_ip>
        set ntp-server3 <class_ip>
    next
end
```

The NTP service options include:

- `local`: The IP address of the interface that the DHCP server is added to becomes the client's NTP server IP address.
- `default`: Clients are assigned the FortiGate's configured NTP servers.
- `specify`: Specify up to three NTP servers in the DHCP server configuration.

IP address assignment with relay agent information option

Option 82 (DHCP relay information option) helps protect the FortiGate against attacks such as spoofing (or forging) of IP and MAC addresses, and DHCP IP address starvation.



The following CLI variables are included in the `config system dhcp server` > `config reserved-address` command:

`circuit-id-type {hex | string}`	DHCP option type; hex or string (default).
`circuit-id <value>`	Option 82 circuit ID of the client that will get the reserved IP address. Format: *vlan-mod-port* • vlan: VLAN ID (2 bytes) • mod: 1 = snoop, 0 = relay (1 byte) • port: port number (1 byte)

`remote-id-type {hex | string}`	DHCP option type; hex or string (default).
`remote-id <value>`	Option 82 remote ID of the client that will get the reserved IP address. Format: the MAC address of the client.
`type {mac | option82}`	The DHCP reserved address type; mac (default) or option82.

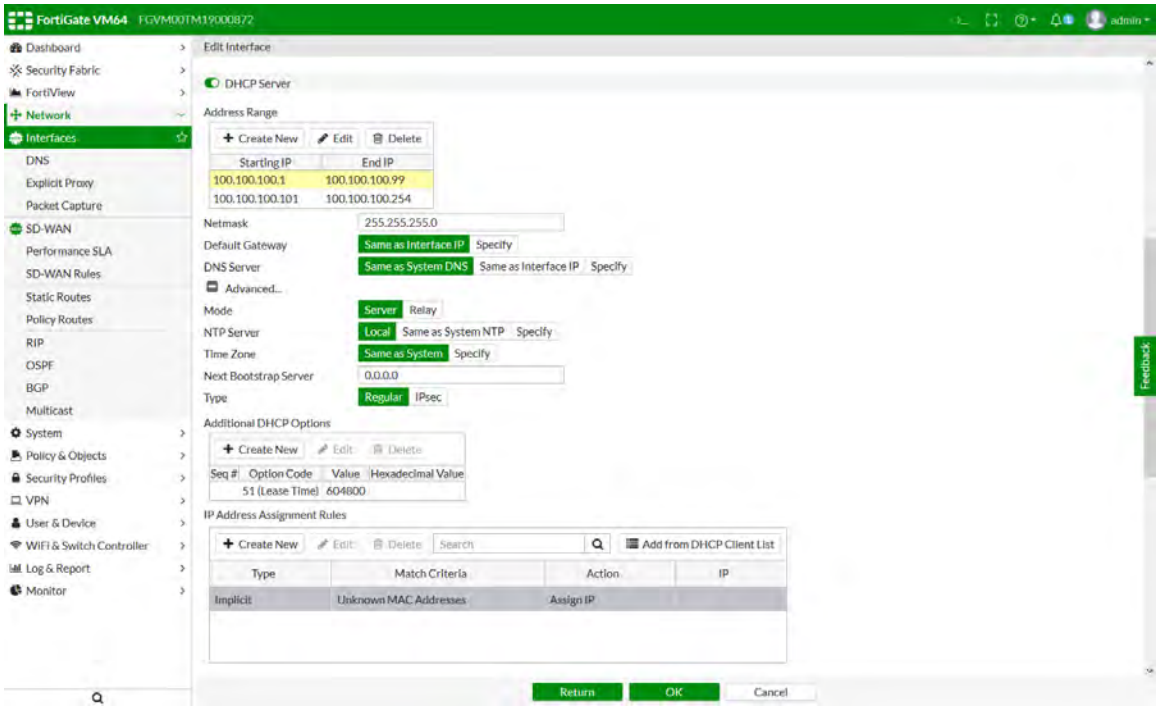
To create an IP address assignment rule using option 82 in the GUI:

1. Go to *Network* > *Interfaces*.
2. Edit an existing port, or create a new one.



The port *Role* must be *LAN* or *Undefined*.

3. Enable *DHCP Server*.
4. Configure the address ranges and other settings as needed.
5. Click + to expand the *Advanced* options.
6. In the *IP Address Assignment Rules* table, click *Create New*.

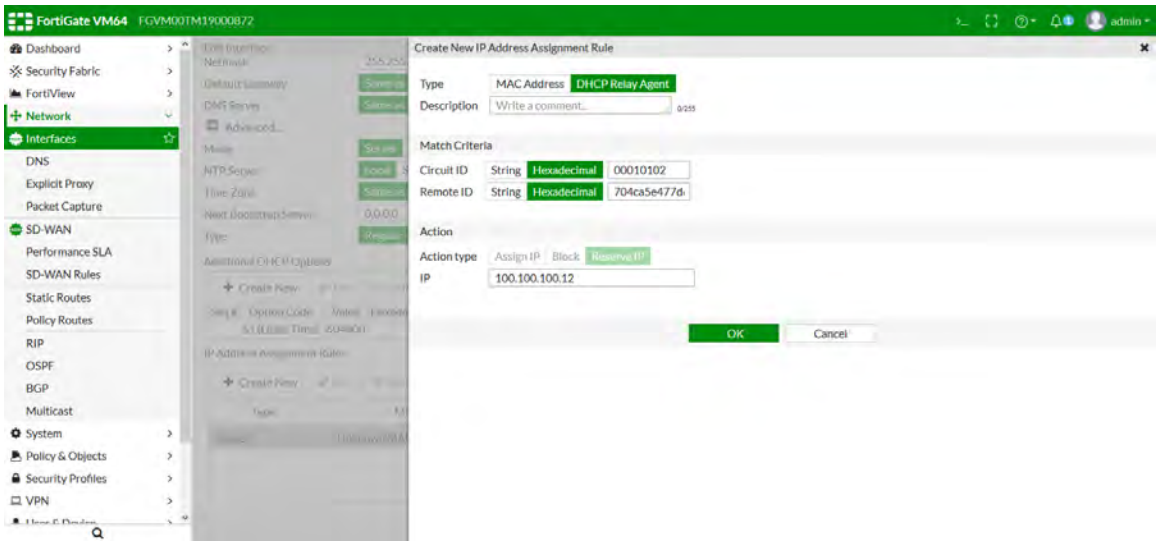


The *Create New IP Address Assignment Rule* pane opens.

7. Configure the new rule:
   a. For the *Type*, select *DHCP Relay Agent*.
   b. Enter the *Circuit ID* and *Remote ID*.
   c. Enter the *IP* address that will be reserved.

Netskope Exhibit 1017

8. Click *OK*.



To create an IP address assignment rule using option 82 with the CLI:

```
config system dhcp server
    edit 1
        set netmask 255.255.255.0
        set interface "port4"
        config ip-range
            edit 1
                set start-ip 100.100.100.1
                set end-ip 100.100.100.99
            next
            edit 2
                set start-ip 100.100.100.101
                set end-ip 100.100.100.254
            next
        end
        config reserved-address
            edit 1
                set type option82
                set ip 100.100.100.12
                set circuit-id-type hex
                set circuit-id "00010102"
                set remote-id-type hex
                set remote-id "704ca5e477d6"
            next
        end
    next
end
```

DHCP client options

When an interface is in DHCP addressing mode, DHCP client options can be configured in the CLI. For example, a vendor class identifier (usually DCHP client option 60) can be specified so that a request can be matched by a specific DHCP offer.

Multiple options can be configured, but any options not recognized by the DHCP server are discarded.

To configure client option 60 - vendor class identifier:

```
config system interface
    edit port1
        set vdom vdom1
        set mode dhcp
        config client-options
            edit 1
                set code 60
                set type hex
                set value aabbccdd
            next
        end
        set type physical
        set snmp-index 4
    next
end
```

Variable	Description
`code <integer>`	DHCP client option code (0 - 255, default = 0). See Dynamic Host Configuration Protocol (DHCP) and Bootstrap Protocol (BOOTP) Parameters for a list of possible options.
`type {hex | string | ip | fqdn}`	DHCP client option type (default = hex).
`value <string>`	DHCP client option value.
`ip <ip>`	DHCP client option IP address. This option is only available when `type` is `ip`.

Static routes

The following topics include information about static routes:

- Deploying the Security Fabric on page 158
- Security Fabric over IPsec VPN on page 170
- Viewing and controlling network risks via topology view on page 156
- Adding a static route on page 446
- Configure VDOM-A on page 618
- Configure VDOM-A on page 627
- Adding IPsec aggregate members in the GUI on page 1163
- IPsec VPN in an HA environment on page 1158
- IPsec VPN to Azure with virtual network gateway on page 1101
- FortiGate as dialup client on page 1116
- ADVPN with BGP as the routing protocol on page 1219
- ADVPN with OSPF as the routing protocol on page 1229
- ADVPN with RIP as the routing protocol on page 1237

Netskope Exhibit 1017

- Basic site-to-site VPN with pre-shared key on page 1071
- Site-to-site VPN with digital certificate on page 1076
- Tunneled Internet browsing on page 1144
- FortiGate multiple connector support on page 1580

Policy routes

The following topics include information about policy routes:

- IPsec aggregate for redundancy and traffic load-balancing on page 1167
- Using BGP tags with SD-WAN rules on page 495

RIP

The following topics include information about Routing Information Protocol (RIP):

- ADVPN with RIP as the routing protocol on page 1237

OSPF

The following topics include information about Open Shortest Path First (OSPF):

- OSPF with IPsec VPN for network redundancy on page 1150
- IPsec aggregate for redundancy and traffic load-balancing on page 1167
- ADVPN with OSPF as the routing protocol on page 1229

BGP

The following topics include information about Border Gateway Protocol (BGP):

- Using BGP tags with SD-WAN rules on page 495
- BGP multiple path support on page 498
- Controlling traffic with BGP route mapping and service rules on page 501
- Applying BGP route-map to multiple BGP neighbors on page 507
- IBGP and EBGP support in VRF on page 513
- ADVPN and shortcut paths on page 516
- SD-WAN related diagnose commands on page 573
- IKEv2 IPsec site-to-site VPN to an AWS VPN gateway on page 1094

Netskope Exhibit 1017

- ADVPN with BGP as the routing protocol on page 1219
- Configuring RADIUS SSO authentication on page 1420

Direct IP support for LTE/4G

Direct IP is a public IP address that is assigned to a computing device, which allows the device to directly access the internet.

When an LTE modem is enabled in FortiOS, a DHCP interface is created. As a result, the FortiGate can acquire direct IP (which includes IP, DNS, and gateway) from the LTE network carrier.

Since some LTE modems require users to input the access point name (APN) for the LTE network, the LTE modem configuration allows you to set the APN.



LTE modems can only be enabled by using the CLI.

To enable direct IP support using the CLI:

1. Enable the LTE modem:

```
config system lte-modem
  set status enable
end
```

2. Check that the LTE interface was created:

```
config system interface
  edit "wwan"
    set vdom "root"
    set mode dhcp
    set status down
    set distance 1
    set type physical
    set snmp-index 23
  next
end
```

Shortly after the LTE modem joins its carrier network, `wwan` is enabled and granted direct IP:

```
# config system interface
(interface) # edit wwan
(wwan) # get
name                : wwan
....
ip                  : 100.112.75.43 255.255.255.248
....
status              : up
....
defaultgw           : enable
DHCP Gateway        : 100.112.75.41
```

```
Lease Expires       : Thu Feb 21 19:33:27 2019
dns-server-override : enable
Acquired DNS1       : 184.151.118.254
Acquired DNS2       : 70.28.245.227
....
```

PCs can reach the internet via the following firewall policy:

```
config firewall policy
....
  edit 5
    set name "LTE"
    set srcintf "port9"
    set dstintf "wwan"
    set srcaddr "all"
    set dstaddr "all"
    set action accept
    set schedule "always"
    set service "ALL"
    set utm-status enable
    set fsso disable
    set nat enable
  next
end
```

Sample LTE interface

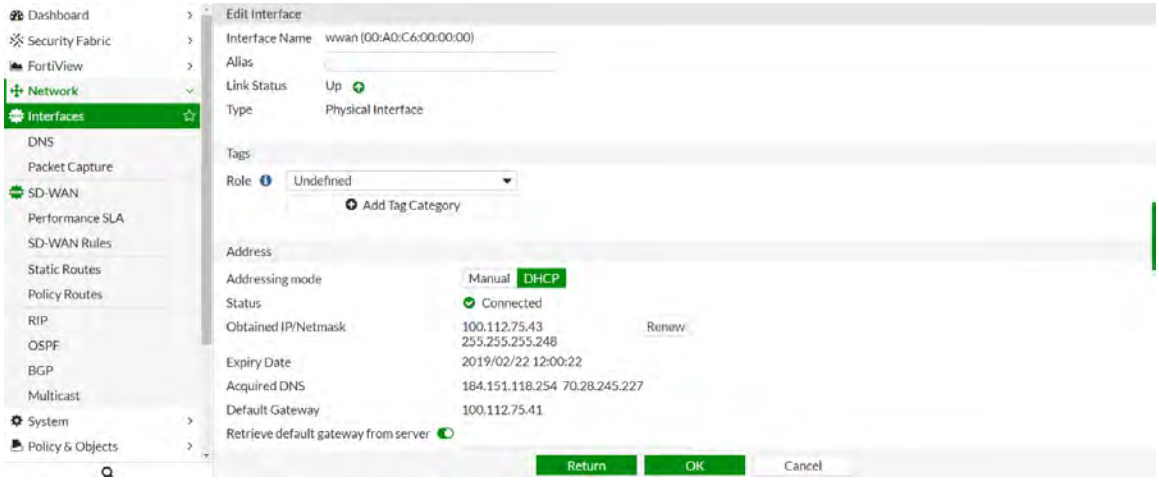
When an LTE modem is enabled, you can view the LTE interface in the GUI and check the acquired IP, DNS, and gateway.

To view the LTE interface in the GUI:

1. Go to *Network* > *Interfaces*.
2. Double-click the LTE interface to view the properties.
3. Look in the *Address* section to view the:
   a. *Obtained IP*
   b. *Acquired DNS*
   c. *Default Gateway*

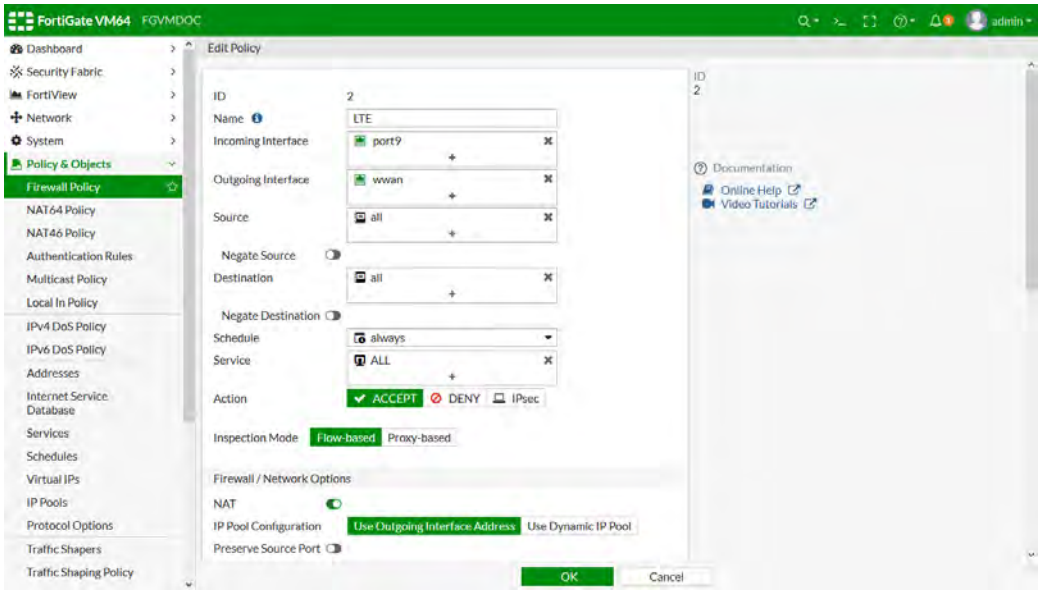
Netskope Exhibit 1017

4. Click *Return*.



To configure the firewall policy that uses the LTE interface:

1. Go to *Policy & Objects* > *Firewall Policy*.
2. Double-click the LTE policy. The *Edit Policy* pane opens.
3. In the *Outgoing Interface* field, select the interface (`wwan` in this example).
4. Configure the rest of the policy as needed.



5. Click *OK*.

Limitations

- Most LTE modems have a preset APN in their SIM card. Therefore, the APN does not need to be set in the FortiOS configuration. In cases where the internet cannot be accessed, consult with your carrier and set the APN in the LTE modem configuration (for example, inet.bell.ca):

```
config system lte-modem
  set status enable
  set apn "inet.bell.ca"
end
```

- Some models, such as the FortiGate 30E-3G4G, have built-in LTE modems. In this scenario, the LTE modem is enabled by default. The firewall policy via the LTE interface is also created by default. Once you plug in a SIM card, your network devices can connect to the internet.

  Sample FortiGate 30E-3G4G default configuration:

```
config system lte-modem
  set status enable
  set extra-init ''
  set manual-handover disable
  set force-wireless-profile 0
  set authtype none
  set apn ''
  set modem-port 255
  set network-type auto
  set auto-connect disable
  set gpsd-enabled disable
  set data-usage-tracking disable
  set gps-port 255
end

config firewall policy
....
  edit 3
    set srcintf "internal"
    set dstintf "wwan"
    set srcaddr "all"
    set dstaddr "all"
    set action accept
    set schedule "always"
    set service "ALL"
    set nat enable
  next
end
```

LLDP reception

Natively, device detection can scan LLDP as a source for device identification. However, the FortiGate does not read or store the full information. Enabling LLDP reception allows the FortiGate to receive and store LLDP messages, learn about active neighbors, and makes the LLDP information available via the CLI, REST API, and SNMP.

You will need to enable `device-identification` at the interface level, and then `lldp-reception` can be enabled on three levels: globally, per VDOM, or per interface.

To configure device identification on an interface:

```
config system interface
    edit <port>
```


Netskope Exhibit 1017

```
        set device-identification enable
    next
end
```

To configure LLDP reception globally:

```
config system global
    set lldp-reception enable
end
```

To configure LLDP reception per VDOM:

```
config system setting
    set lldp-reception enable
end
```

To configure LLDP reception per interface:

```
config system interface
    edit <port>
        set lldp-reception enable
    next
end
```

To view the LLDP information in the GUI:

1. Go to *Dashboard* > *Users & Devices*.
2. Hover over the *Device Inventory* widget, and click *Expand to Full Screen*.



To view the received LLDP information in the CLI:

```
# diagnose user device list
    hosts
      vd root/0 44:0a:a0:0a:0a:0a gen 3 req S/2
        created 10290s gen 1 seen 0s port3 gen 1
        ip 172.22.22.22 src lldp
        type 20 'Other Network Device' src lldp id 155 gen 2
        os 'Artist EOS ' version '4.20.4' src lldp id 155
        host 'artist' src lldp
```

To view additional information about LLDP neighbors and ports:

```
# diagnose lldprx neighbor {summary | details | clear}
# diagnose lldprx port {details | summary | neighbor | filter}
# diagnose lldprx port neighbor {summary | details}
```

Note that the port index in the output corresponds to the port index from the following command:

```
# diagnose netlink interface list port2 port3 | grep index
      if=port2 family=00 type=1 index=4 mtu=1500 link=0 master=0
      if=port3 family=00 type=1 index=5 mtu=1500 link=0 master=0
```

To view the received LLDP information in the REST API:


```
{
  "http_method":"GET",
  "results":[
    {
      "mac":"90:9c:9c:c9:c9:90",
      "chassis_id":"90:9C:9C:C9:C9:90",
      "port":19,
      "port_id":"port12",
      "port_desc":"port12",
      "system_name":"S124DN3W00000000",
      "system_desc":"FortiSwitch-124D v3.6.6,build0416,180515 (GA)",
      "ttl":120,
      "addresses":[
        {
          "type":"ipv4",
          "address":"192.168.1.99"
        }
      ]
    }
  ],
  "vdom":"root",
  "path":"network",
  "name":"lldp",
  "action":"neighbors",
  "status":"success",
  "serial":"FG201E4Q00000000",
  "version":"v6.2.0",
  "build":866
}
```

```
{
  "http_method":"GET",
  "results":[
    {
      "name":"port1",
      "rx":320,
      "neighbors":1
    }
  ],
  "vdom":"root",
  "path":"network",
  "name":"lldp",
  "action":"ports",
  "mkey":"port1",
  "status":"success",
  "serial":"FG201E4Q00000000",
  "version":"v6.2.0",
  "build":866
}
```


Netskope Exhibit 1017

Route leaking between VRFs

This feature provides generic route leaking capabilities between locally defined VRFs (VRF-lite). If VRF leaking is not configured, VRFs are isolated.

In this example, interface *npu0_vlink0* belongs to VRF 10 and is used to leak 1.2.2.2/32 from VRF10 to VRF20, and interface *npu0_vlink1* belongs to VRF 20 and is used to leak 172.28.1.0/24 from VRF20 to VRF10. So, VRF10 can see 172.28.1.0/24, and VRF20 can see 1.2.2.2/32.

To configure VRF leaking:

1. Configure the prefix list and route map to filter what will be leaked:

```
config router prefix-list
    edit "1"
        config rule
            edit 1
                set prefix 1.2.2.2 255.255.255.255
            next
        end
    next
    edit "2"
        config rule
            edit 1
                set prefix 172.28.1.0 255.255.255.0
            next
        end
    next
end

config router route-map
    edit "from10"
        config rule
            edit 1
                set match-ip-address "1"
            next
        end
    next
    edit "from20"
        config rule
            edit 1
                set match-ip-address "2"
            next
        end
    next
end
```

2. Configure the VDOM link interfaces for the leaking and routing:

```
config system interface
    edit "npu0_vlink0"
        set vdom "root"
        set vrf 10
        set ip 172.16.201.1 255.255.255.0
        set allowaccess ping https ssh snmp http
    next
```


Netskope Exhibit 1017

```
    edit "npu0_vlink1"
        set vdom "root"
        set vrf 20
        set ip 172.16.201.2 255.255.255.0
        set allowaccess ping https ssh snmp http telnet
    next
end
```

3. Configure the BGP VRF leak:

```
config router bgp
    set as 44
    set router-id 4.4.4.4
    config neighbor
        edit "172.16.200.1"
            set soft-reconfiguration enable
            set remote-as 11
            set update-source "port1"
        next
        edit "172.16.202.1"
            set soft-reconfiguration enable
            set remote-as 22
            set update-source "port3"
        next
    end
    config vrf-leak
        edit "10"
            config target
                edit "20"
                    set route-map "from10"
                    set interface "npu0_vlink0"
                next
            end
        next
        edit "20"
            config target
                edit "10"
                    set route-map "from20"
                    set interface "npu0_vlink1"
                next
            end
        next
    end
end
```

4. Confirm that the filtered routed leaked as expected:

```
# get router info routing-table all
Routing table for VRF=10
B       1.1.1.1/32 [20/0] via 172.16.200.1, port1, 01:03:16
B       1.2.2.2/32 [20/0] via 172.16.200.1, port1, 01:03:16
B       172.28.1.0/24 [20/0] via 172.16.201.2, npu0_vlink0, 00:00:17
<<<<<<<<<<<<<<Leaked into VRF10 from VRF20

Routing table for VRF=20
B       1.2.2.2/32 [20/0] via 172.16.201.1, npu0_vlink1, 00:00:15   <<<<<<<<<<<<<<Leaked
into VRF 20 from VRF10
```


Netskope Exhibit 1017

```
B       172.28.1.0/24 [20/0] via 172.16.202.1, port3, 01:03:16
B       172.28.2.0/24 [20/0] via 172.16.202.1, port3, 01:03:16
```


Netskope Exhibit 1017

SD-WAN

SD-WAN is a software-defined approach to managing Wide-Area Networks (WAN). It allows you to offload internet-bound traffic, meaning that private WAN services remain available for real-time and mission critical applications. This added flexibility improves traffic flow and reduces pressure on the network.

SD-WAN platforms create hybrid networks that integrate broadband and other network services into the corporate WAN while maintaining the performance and security of real-time and sensitive applications.

SD-WAN with Application Aware Routing can measure and monitor the performance of multiple services in a hybrid network. It uses application routing to offer more granular control of where and when an application uses a specific service, allowing better use of the overall network.

Some of the key benefits of SD-WAN include:

- Reduced cost with transport independence across MPLS, 3G/4G LTE, and others.
- Improve business application performance thanks to increased availability and agility.
- Optimized user experience and efficiency with SaaS and public cloud applications.

SD-WAN has 3 objects:

- **SD-WAN interface**
  Also called members, SD-WAN interfaces are the ports and interfaces that are used to run traffic. At least one interface must be configured for SD-WAN to function; up to 255 member interfaces can be configured. See Configuring the SD-WAN interface on page 445.
- **Performance SLA**
  Also called health-check, performance SLAs are used to monitor member interface link quality, and to detect link failures. They can be used to remove routes, and to reroute traffic when an SD-WAN member cannot detect the server. They can also be used in SD-WAN rules to select the preferred member interface for forwarding traffic. See Performance SLA - link monitoring on page 455.
- **SD-WAN rule**
  Also called service, SD-WAN rules are used to control path selection. Specific traffic can be dynamically sent to the best link, or use a specific route. There are five modes:
  - *auto*: Assign interfaces a priority based on quality.
  - *manual*: Assign interfaces a priority manually.
  - *priority*: Assign interfaces a priority based on the link-cost-factor quality of the interface.
  - *sla*: Assign interfaces a priority based on selected SLA settings.
  - *load-balance*: Distribute traffic among all available links based on the load balance algorithm.

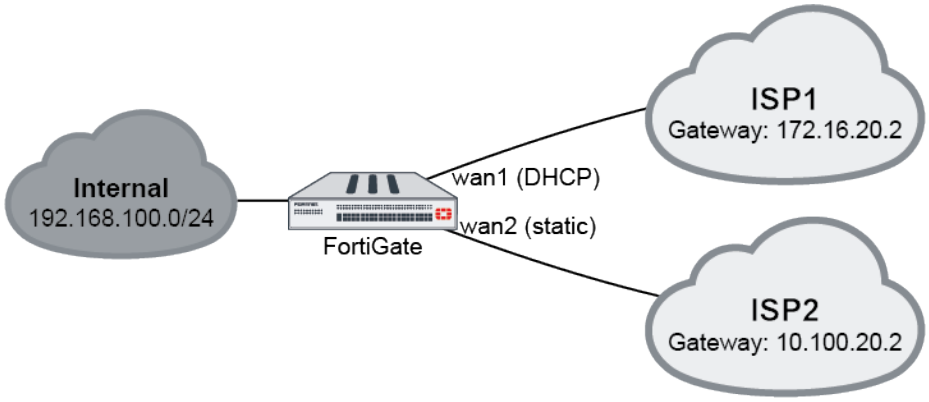
  See SD-WAN rules - best quality on page 464, SD-WAN rules - lowest cost (SLA) on page 467, and SD-WAN rules - maximize bandwidth (SLA) on page 470.

SD-WAN quick start

This section provides an example of how to start using SD-WAN for load balancing and redundancy.

Netskope Exhibit 1017

In this example, two ISP internet connections, wan1 (DHCP) and wan2 (static), use SD-WAN to balance traffic between them at 50% each.



1. Configuring the SD-WAN interface on page 445
2. Adding a static route on page 446
3. Selecting the implicit SD-WAN algorithm on page 446
4. Configuring firewall policies for SD-WAN on page 447
5. Link monitoring and failover on page 448
6. Results on page 449
7. Configuring SD-WAN in the CLI on page 453

Configuring the SD-WAN interface

First, SD-WAN must be enabled and member interfaces must be selected. The selected FortiGate interfaces can be of any type (physical, aggregate, VLAN, IPsec, and others), but must be removed from any other configurations on the FortiGate.

In this step, two interfaces are configured and then selected as SD-WAN member interfaces. This example uses a mix of static and dynamic IP addresses; your deployment could also use only one or the other.

Once the SD-WAN interface is configured, it is referenced as *SD-WAN* in the GUI for static routes and firewall policies, and `virtual-wan-link` can be enabled in the CLI.

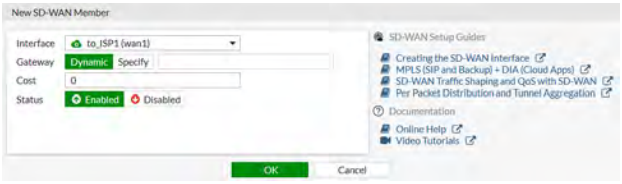
To configure SD-WAN members:

1. Configure the wan1 and wan2 interfaces. See Interface settings on page 341 for details.
   a. Set the wan1 interface *Addressing mode* to *DHCP* and *Distance* to *10*.

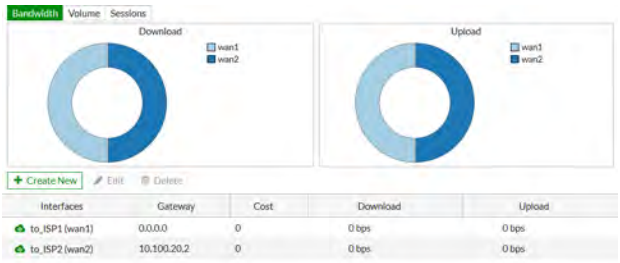
   By default, a DHCP interface has a distance of 5, and a static route has a distance of 10. It is important to account for this when configuring your SD-WAN for 50/50 load balancing by setting the DHCP interface's distance to 10.

   b. Set the wan2 interface *IP/Netmask* to *10.100.20.1 255.255.255.0*.
2. Go to *Network* > *SD-WAN Interface*. Routing for each SD-WAN interface is defined here.
3. Click *Create New*.
4. Set the *Interface* to *wan1*.

5. As wan1 uses DHCP, set *Gateway* to *Dynamic*.
   If IPv6 visibility is enabled in the GUI, an IPv6 gateway can also be added for each member. See Feature visibility on page 694 for details.
6. Leave *Cost* as *0*.
   The *Cost* field is used by the Lowest Cost (SLA) strategy. The link with the lowest cost is chosen to pass traffic. The lowest possible *Cost* is *0*.



7. Set *Status* to *Enable*, and click *OK*.
8. Repeat the above steps for wan2, setting *Gateway* to the ISP's gateway: *10.100.20.2*.



Adding a static route

After you create an SD-WAN interface, FortiGate adds a virtual interface for SD-WAN to the interface list that can be used to create routes.

You must configure a default route for the SD-WAN interface. The default gateways for each SD-WAN member interface do not need to be defined in the static routes table. FortiGate will decide which route or routes are preferred using Equal Cost Multi-Path (ECMP) based on distance and priority.

To create a static route for SD-WAN:

1. Go to *Network* > *Static Routes*.
2. Click *Create New*. The *New Static Route* page opens.
3. Set *Destination* to *Subnet*, and leave the IP address and subnet mask as *0.0.0.0/0.0.0.0*.
4. From the *Interface* drop-down list, select *SD-WAN*.



5. Ensure that *Status* is *Enabled*.
6. Click *OK*.

Selecting the implicit SD-WAN algorithm

SD-WAN rules define specific routing options to route traffic to an SD-WAN member.

If no routing rules are defined, the default *Implicit* rule is used. It can be configured to use one of five different load balancing algorithms. See Implicit rule on page 461 for more details and examples.

This example shows four methods to equally balance traffic between the two WAN connections. Go to *Network* > *SD-WAN Rules* and edit the *sd-wan* rule to select the method that is appropriate for your requirements.

- *Source IP* (CLI command: `source-ip-based`):
  Select this option to balance traffic equally between the SD-WAN members according to a hash algorithm based on the source IP addresses.
- *Session* (`weight-based`):
  Select this option to balance traffic equally between the SD-WAN members by the session numbers ratio among its members. Use weight 50 for each of the 2 members.
- *Source-Destination IP* (`source-dest-ip-based`):
  Select this option to balance traffic equally between the SD-WAN members according to a hash algorithm based on the source and destination IP addresses.
- *Volume* (`measured-volume-based`):
  Select this option to balance traffic equally between the SD-WAN members according to the bandwidth ratio among its members.

Configuring firewall policies for SD-WAN

After you create an SD-WAN interface, FortiGate adds a virtual interface for SD-WAN to the interface list that can be used to create firewall policies.

You must configure a policy that allows traffic from your organization's internal network to the SD-WAN interface (`virtual-wan-link` in the CLI). You do not need to configure policies for each individual SD-WAN member interface because policies configured with the SD-WAN interface apply to all SD-WAN interface members.

To create a firewall policy for SD-WAN:

1. Go to *Policy & Objects* > *Firewall Policy*.
2. Click *Create New*. The *New Policy* page opens.

Netskope Exhibit 1017

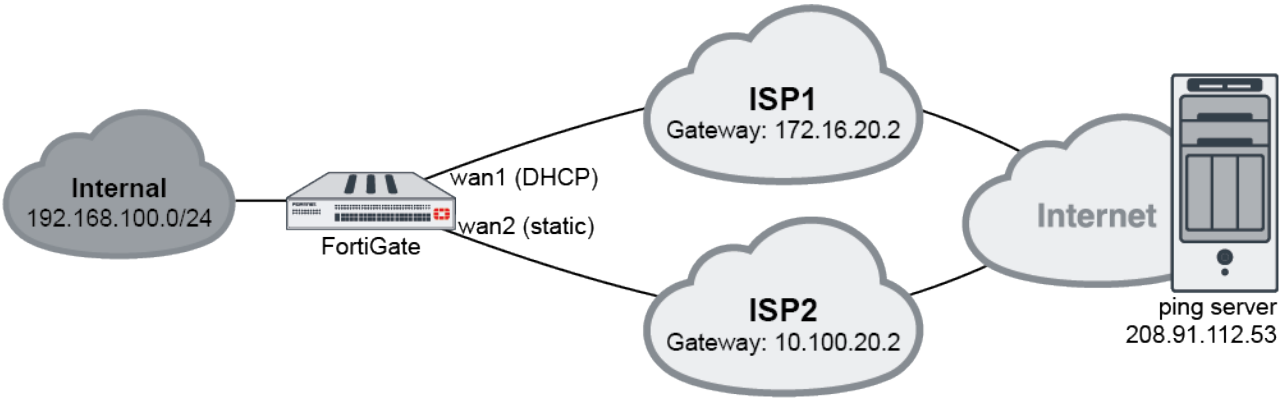
3. Configure the following:

Name	Enter a name for the policy.
Incoming Interface	*internal*
Outgoing Interface	*SD-WAN*
Source	*all*
Destination	*all*
Schedule	*always*
Service	*ALL*
Action	*ACCEPT*
Firewall / Network Options	Enable *NAT* and set *IP Pool Configuration* to *Use Outgoing Interface Address*.
Security Profiles	Apply profiles as required.
Logging Options	Enable *Log Allowed Traffic* and select *All Sessions*. This allows you to verify results later.

4. Enable the policy, then click *OK*.

Link monitoring and failover

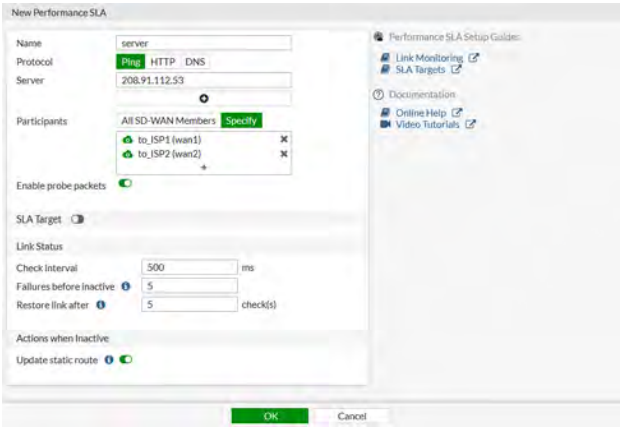
Performance SLA link monitoring measures the health of links that are connected to SD-WAN member interfaces by sending probing signals through each link to a server, and then measuring the link quality based on latency, jitter, and packet loss. If a link is broken, the routes on that link are removed and traffic is routed through other links. When the link is working again, the routes are re-enabled. This prevents traffic being sent to a broken link and lost.



In this example, the detection server IP address is 208.91.112.53. A performance SLA is created so that, if ping fails per the metrics defined, the routes to that interface are removed and traffic is detoured to the other interface. The ping protocol is used, but other protocols could also be selected as required.

To configure a performance SLA:

1. Go to *Network* > *Performance SLA*.
2. Click *Create New*. The *New Performance SLA* page opens.
3. Enter a name for the SLA and set *Protocol* to *Ping*.
4. In the *Server* field, enter the detection server IP address (208.91.112.53 in this example).
5. In the *Participants* field, select *Specify* and add wan1 and wan2.



   SLA targets are not required for link monitoring.
6. Configure the required metrics in *Link Status*.
7. Ensure that *Update static route* is enabled. This disables static routes for the inactive interface and restores routes on recovery.
8. Click *OK*.

Results

The following GUI pages show the function of the SD-WAN and can be used to confirm that it is setup and running correctly:

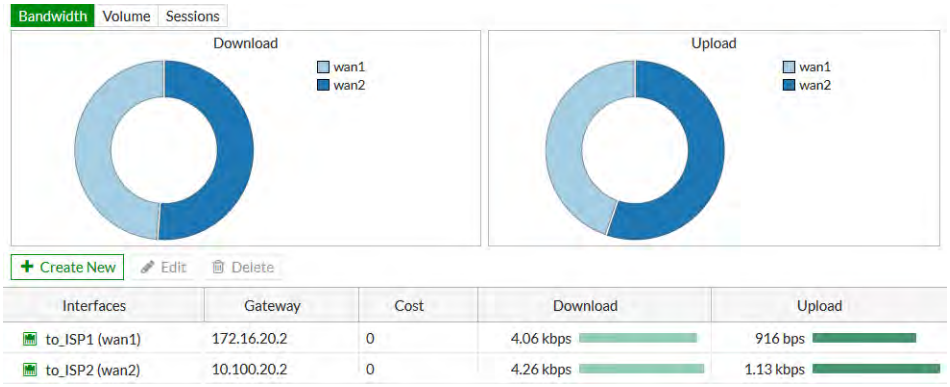
- Interface usage on page 449
- Routing table on page 452
- Firewall policy on page 453

Interface usage

Go to *Network* > *SD-WAN Interfaces* to review the SD-WAN interfaces' usage.

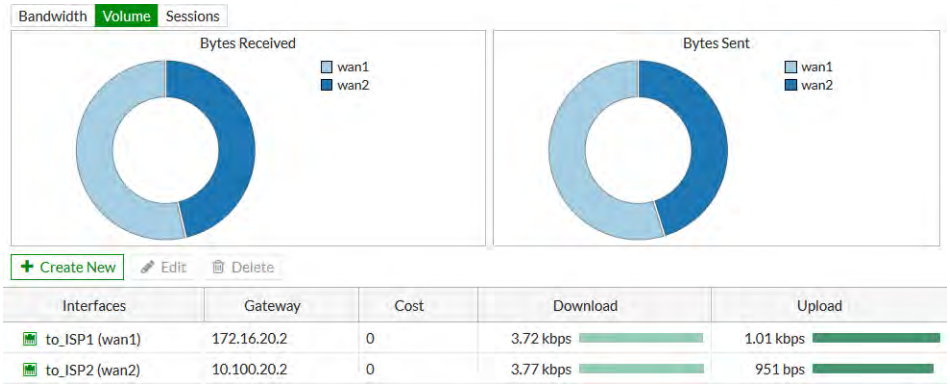
Bandwidth

Select *Bandwidth* to view the amount of downloaded and uploaded data for each interface.



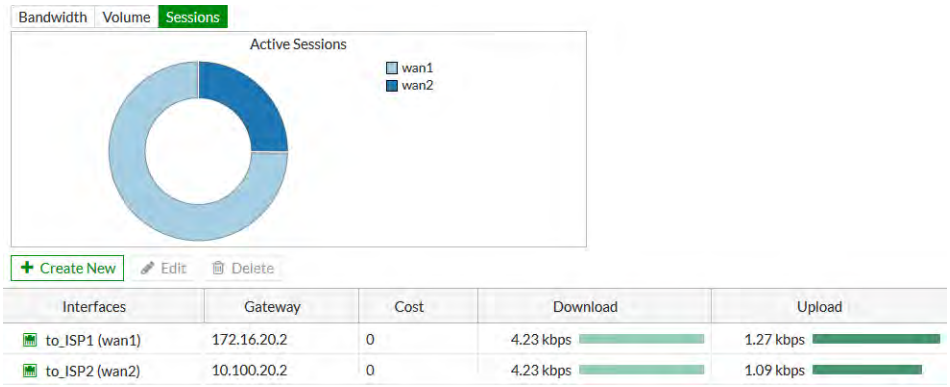
Volume

Select *Volume* to see donut charts of the received and sent bytes on the interfaces.



Sessions

Select *Sessions* to see a donut chart of the number of active sessions on each interface.



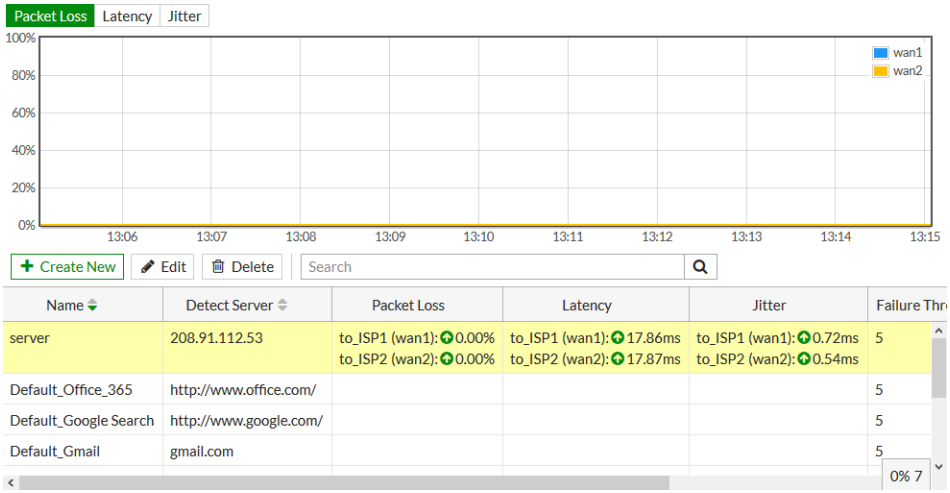
Netskope Exhibit 1017

Performance SLA

Go to *Network > Performance SLA* and select the SLA from the table (*server* in this example) to view the packet loss, latency, and jitter on each SD-WAN member in the health check server.

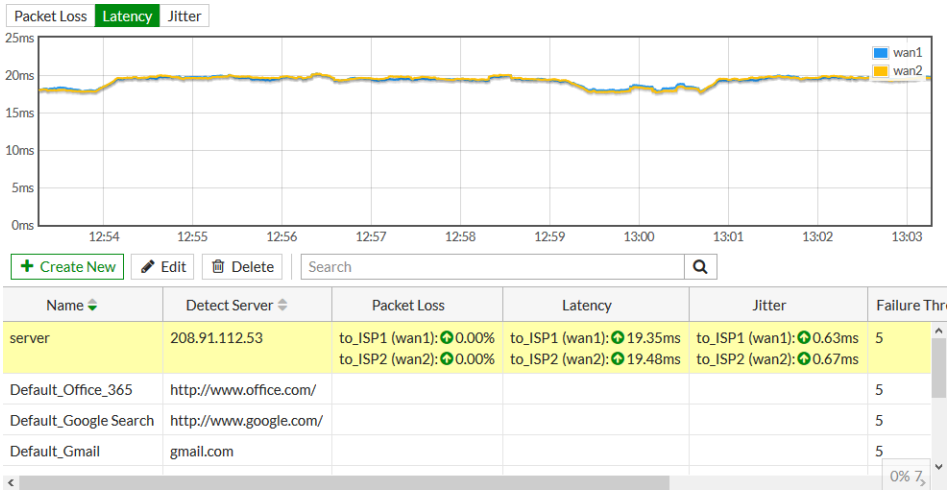
Packet loss

Select *Packet Loss* to see the percentage of packets lost for each member.



Latency

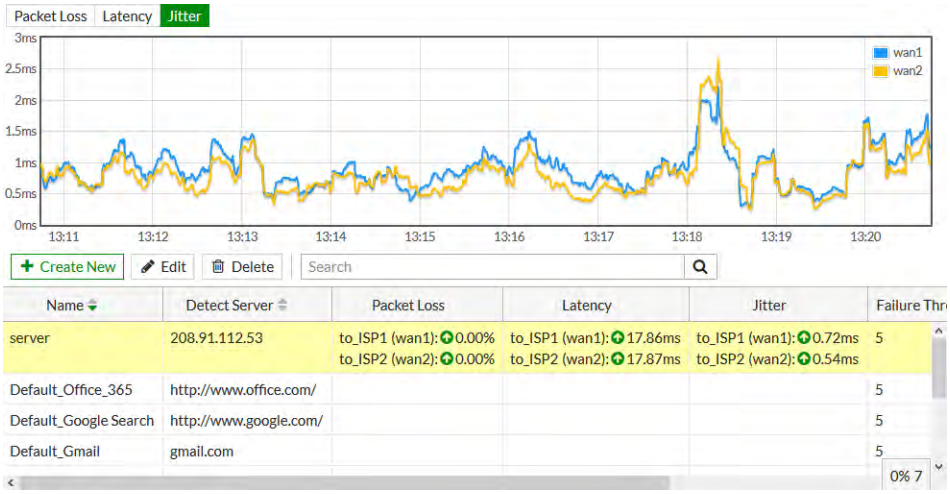
Select *Latency* to see the current latency, in milliseconds, for each member.



Netskope Exhibit 1017

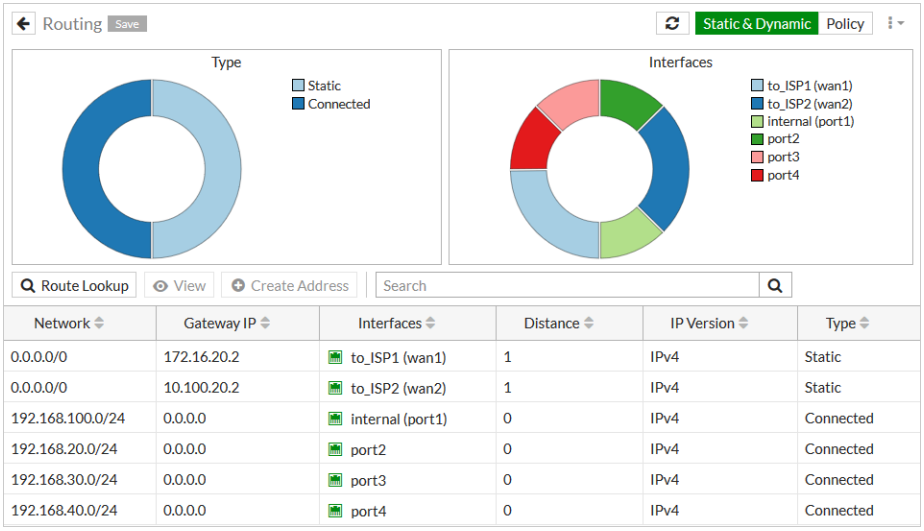
Jitter

Select *Jitter* to see the jitter, in milliseconds, for each member.



Routing table

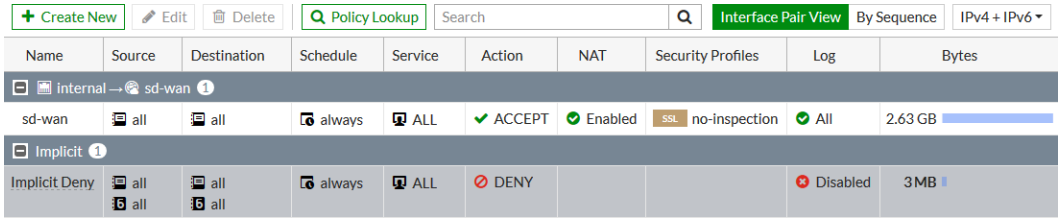
Go to *Dashboard* > *Network* and expand the *Static & Dynamic Routing* widget to review all static and dynamic routes. For more information about the widget, see Static & Dynamic Routing Monitor on page 63.



Netskope Exhibit 1017

Firewall policy

Go to *Policy & Objects* > *Firewall Policy* to review the SD-WAN policy.



Configuring SD-WAN in the CLI

This example can be entirely configured using the CLI.

To configure SD-WAN in the CLI:

1. Configure the wan1 and wan2 interfaces:

```
config system interface
    edit "wan1"
        set alias to_ISP1
        set mode dhcp
        set distance 10
    next
    edit "wan2"
        set alias to_ISP2
        set ip 10.100.20.1 255.255.255.0
    next
end
```

2. Enable SD-WAN and add the interfaces as members:

```
config system virtual-wan-link
    set status enable
    config members
        edit 1
            set interface "wan1"
        next
        edit 2
            set interface "wan2"
            set gateway 10.100.20.2
        next
    end
end
```

3. Create a static route for SD-WAN:

```
config router static
    edit 1
        set virtual-wan-link enable
    next
end
```

4. Select the implicit SD-WAN algorithm:

```
config system virtual-wan-link
    set load-balance-mode {source-ip-based | weight-based | source-dest-ip-based |
measured-volume-based}
end
```

5. Create a firewall policy for SD-WAN:

```
config firewall policy
    edit <policy_id>
        set name <policy_name>
        set srcintf internal
        set dstintf virtual-wan-link
        set srcaddr all
        set dstaddr all
        set action accept
        set schedule always
        set service ALL
        set utm-status enable
        set ssl-ssh-profile <profile_name>
        set av-profile <profile_name>
        set webfilter-profile <profile_name>
        set dnsfilter-profile <profile_name>
        set emailfilter-profile <profile_name>
        set ips_sensor <sensor_name>
        set application-list <app_list>
        set voip-profile <profile_name>
        set logtraffic all
        set nat enable
        set status enable
    next
end
```

6. Configure a performance SLA:

```
config system virtual-wan-link
    config health-check
        edit "server"
            set server "208.91.112.53"
            set update-static-route enable
            set members 1 2
        next
    end
end
```

Results

To view the routing table:

```
# get router info routing-table all

Routing table for VRF=0
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
```


Netskope Exhibit 1017

```
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default

S*      0.0.0.0/0 [1/0] via 172.16.20.2, wan1
                  [1/0] via 10.100.20.2, wan2
C       10.100.20.0/24 is directly connected, wan2
C       172.16.20.2/24 is directly connected, wan1
C       192.168.0.0/24 is directly connected, internal
```

To diagnose the Performance SLA status:

```
FGT # diagnose sys virtual-wan-link health-check
Health Check(server):
Seq(1): state(alive), packet-loss(0.000%) latency(15.247), jitter(5.231) sla_map=0x0
Seq(2): state(alive), packet-loss(0.000%) latency(13.621), jitter(6.905) sla_map=0x0
```

WAN path control

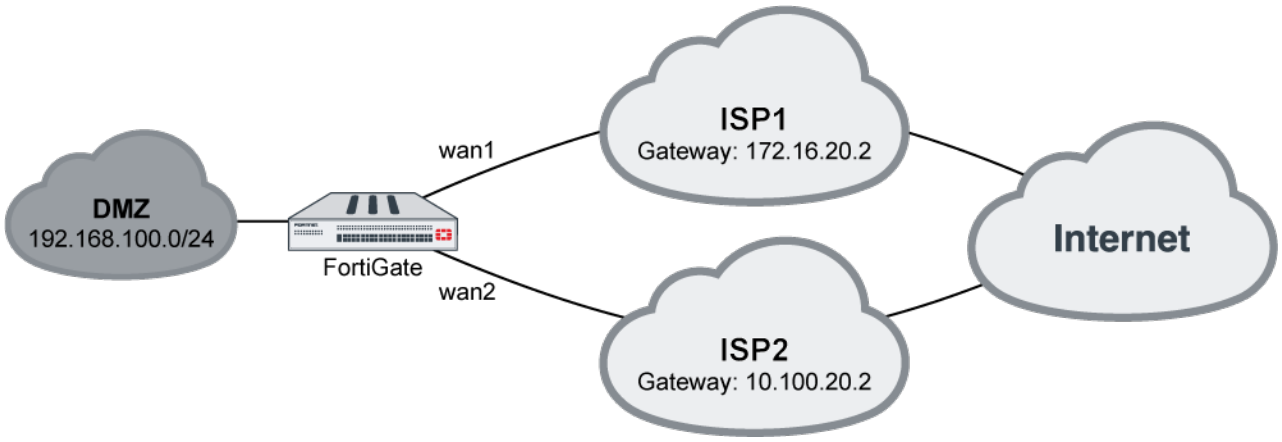
The following topics provide instructions on configuring WAN path control:

- Performance SLA - link monitoring on page 455
- Performance SLA - SLA targets on page 457
- Factory default health checks on page 458
- Implicit rule on page 461
- SD-WAN rules - best quality on page 464
- SD-WAN rules - lowest cost (SLA) on page 467
- SD-WAN rules - maximize bandwidth (SLA) on page 470
- Application steering using SD-WAN rules on page 472
- SD-WAN traffic shaping and QoS on page 484

Performance SLA - link monitoring

Performance SLA link monitoring measures the health of links that are connected to SD-WAN member interfaces by sending probing signals through each link to a server and measuring the link quality based on latency, jitter, and packet loss. If a link is broken, the routes on that link are removed, and traffic is routed through other links. When the link is working again, the routes are reenabbled. This prevents traffic being sent to a broken link and lost.

Netskope Exhibit 1017



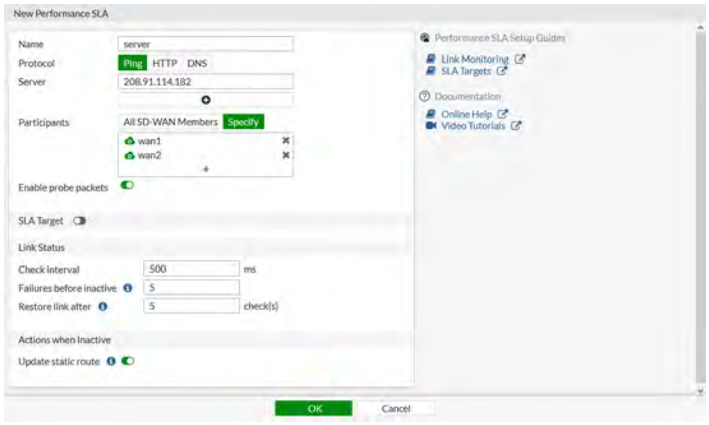
In this example:

- Interfaces wan1 and wan2 connect to the internet through separate ISPs
- The detection server IP address is 208.91.114.182

A performance SLA is created so that, if one link fails, its routes are removed and traffic is detoured to the other link.

To configure a Performance SLA using the GUI:

1. On the FortiGate, add wan1 and wan2 as SD-WAN members, then add a policy and static route. See SD-WAN quick start on page 444 for details.
2. Go to *Network* > *Performance SLA*.
3. Click *Create New*. The *Performance SLA* page opens.
4. Enter a name for the SLA and select a protocol.
5. In the *Server* field, enter the detection server IP address (208.91.114.182 in this example).
6. In the *Participants* field, select both wan1 and wan2.



7. Configured the remaining settings as needed, then click *OK*.

To configure a Performance SLA using the CLI:

```
config system virtual-wan-link
    config health-check
        edit "server"
            set server "208.91.114.182"
            set update-static-route enable
```

```
            set members 1 2
        next
    end
end
```

To diagnose the Performance SLA status:

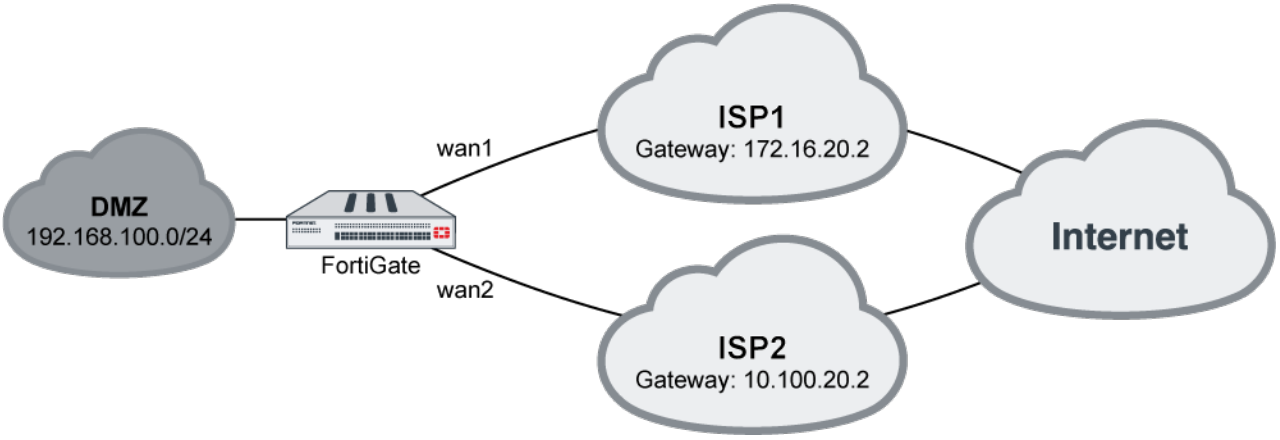
```
FGT # diagnose sys virtual-wan-link health-check
Health Check(server):
Seq(1): state(alive), packet-loss(0.000%) latency(15.247), jitter(5.231) sla_map=0x0
Seq(2): state(alive), packet-loss(0.000%) latency(13.621), jitter(6.905) sla_map=0x0
```

Performance SLA - SLA targets

SLA targets are a set of constraints that are used in SD-WAN rules to control the paths that traffic take.

The available constraints are:

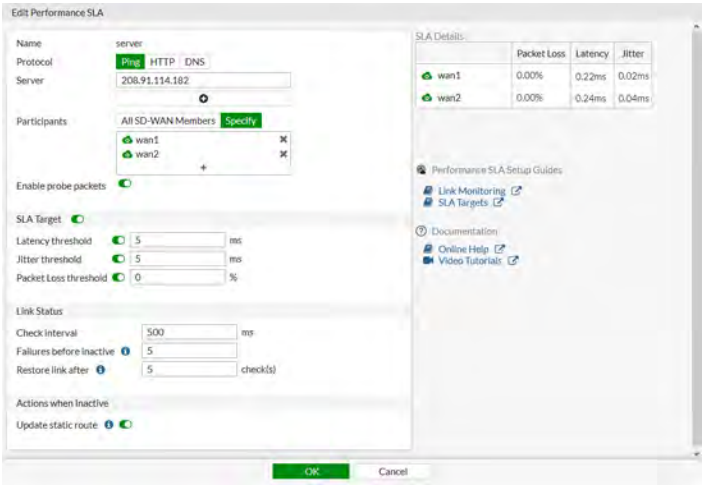
- *Latency threshold*: Latency for SLA to make decision, in milliseconds (0 - 10000000, default = 5).
- *Jitter threshold*: Jitter for SLA to make decision, in milliseconds (0 - 10000000, default = 5).
- *Packet loss threshold*: Packet loss for SLA to make decision, in percentage (0 - 100, default = 0).



To configure Performance SLA targets using the GUI:

1. On the FortiGate, add wan1 and wan2 as SD-WAN members, then add a policy and static route. See SD-WAN quick start on page 444 for details.
2. Go to *Network* > *Performance SLA*.
3. Create a new Performance SLA or edit an existing one. See Performance SLA - link monitoring on page 455.

4. Enable *SLA Targets*and configure the constraints. To add multiple SLA targets, use the CLI.



5. Configured the remaining settings as needed, then click *OK*.

To configure Performance SLA targets using the GUI:

```
config system virtual-wan-link
    config health-check
        edit "server"
            set server "208.91.114.182"
            set members 1 2
            config sla
                edit 1
                    set link-cost-factor latency jitter packet-loss
                    set latency-threshold 10
                    set jitter-threshold 10
                    set packetloss-threshold 1
                next
            end
        next
    end
end
```

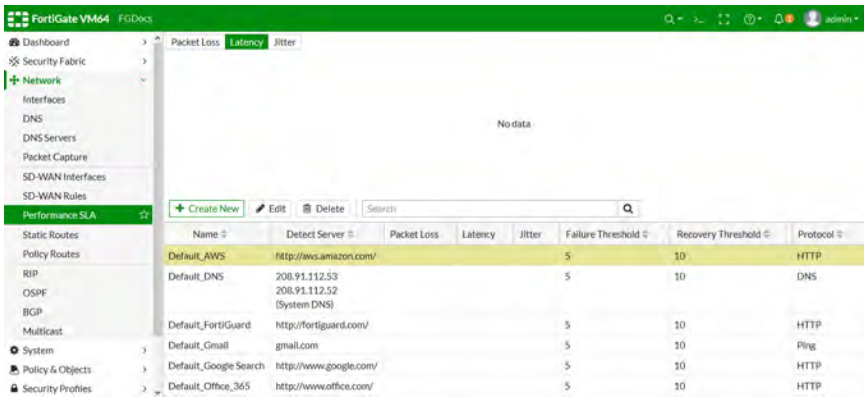
The `link-cost-factor` variable is used to select which constraints are enabled.

Factory default health checks

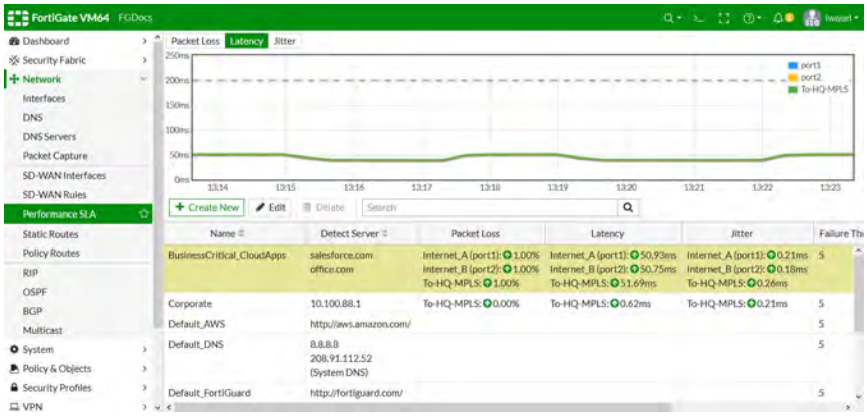
There are six predefined performance SLA profiles for newly created VDOMs or factory reset FortiGate devices:

- AWS
- System DNS
- FortiGuard
- Gmail
- Google Search
- Office 365

You can view and configure the SLA profiles in *Network* > *Performance SLA*.



After configuring a health check, you will be able to view packet loss, latency, and jitter data for the SLA profiles. If a value is colored red, it means that it failed to meet the SLA requirements.



To configure the performance SLA profiles in the CLI:

```
config system virtual-wan-link
    config health-check
        edit "Default_AWS"
            set server "aws.amazon.com"
            set protocol http
            set interval 1000
            set recoverytime 10
            config sla
                edit 1
                    set latency-threshold 250
                    set jitter-threshold 50
                    set packetloss-threshold 5
                next
            end
        next
        edit "Default_DNS"
            set system-dns enable
            set interval 1000
            set probe-timeout 1000
            set recoverytime 10
            config sla
                edit 1
```

```
            set latency-threshold 250
            set jitter-threshold 50
            set packetloss-threshold 5
        next
    end
next
edit "Default_FortiGuard"
    set server "fortiguard.com"
    set protocol http
    set interval 1000
    set recoverytime 10
    config sla
        edit 1
            set latency-threshold 250
            set jitter-threshold 50
            set packetloss-threshold 5
        next
    end
next
edit "Default_Gmail"
    set server "gmail.com"
    set interval 1000
    set recoverytime 10
    config sla
        edit 1
            set latency-threshold 250
            set jitter-threshold 50
            set packetloss-threshold 2
        next
    end
next
edit "Default_Google Search"
    set server "www.google.com"
    set protocol http
    set interval 1000
    set recoverytime 10
    config sla
        edit 1
            set latency-threshold 250
            set jitter-threshold 50
            set packetloss-threshold 5
        next
    end
next
edit "Default_Office_365"
    set server "www.office.com"
    set protocol http
    set interval 1000
    set recoverytime 10
    config sla
        edit 1
            set latency-threshold 250
            set jitter-threshold 50
            set packetloss-threshold 5
        next
    end
```


Netskope Exhibit 1017

```
        next
    end
end
```

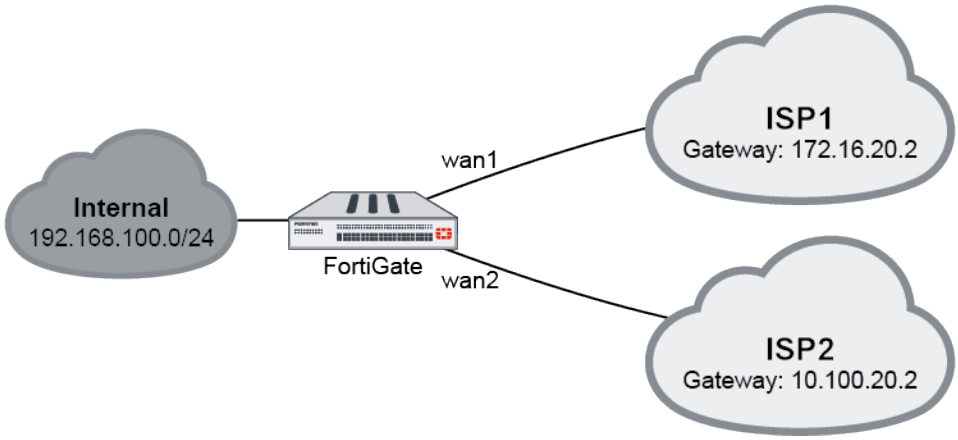
Implicit rule

SD-WAN supports five types of implicit rules (load-balance mode):

- Source IP (CLI command: `source-ip-based`): SD-WAN will load balance the traffic equally among its members according to a hash algorithm based on the source IP addresses.
- Session (`weight-based`): SD-WAN will load balance the traffic according to the session numbers ratio among its members.
- Spillover (`usage-based`): SD-WAN will use the first member until the bandwidth reaches its limit, then use the second, and so on.
- Source-Destination IP (`source-dest-ip-based`): SD-WAN will load balance the traffic equally among its members according to a hash algorithm based on both the source and destination IP addresses.
- Volume (`measured-volume-based`): SD-WAN will load balance the traffic according to the bandwidth ratio among its members.

Examples

The following four examples demonstrate how to use the implicit rules (load-balance mode).



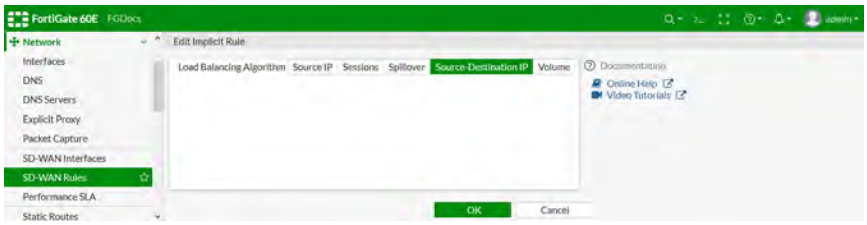
Example 1

Outgoing traffic is equally balanced between wan1 and wan2, using *source-ip-based* or *source-dest-ip-based* mode.

Using the GUI:

1. On the FortiGate, enable SD-WAN and add wan1 and wan2 as SD-WAN members, then add a policy and static route. See SD-WAN quick start on page 444 for details.
2. Go to *Network* > *SD-WAN Rules*.
3. Edit the *sd-wan* rule (the last default rule).

4. For the *Load Balancing Algorithm*, select either *Source IP* or *Source-Destination IP*.



5. Click *OK*.

Using the CLI:

1. Enable SD-WAN and add wan1 and wan2 as SD-WAN members, then add a policy and static route. See SD-WAN quick start on page 444 for details.
2. Set the load balancing algorithm:
   Source IP based:

```
config system virtual-wan-link
    set load-balance-mode source-ip-based
end
```

   Source-Destination IP based:

```
config system virtual-wan-link
    set load-balance-mode source-dest-ip-based
end
```

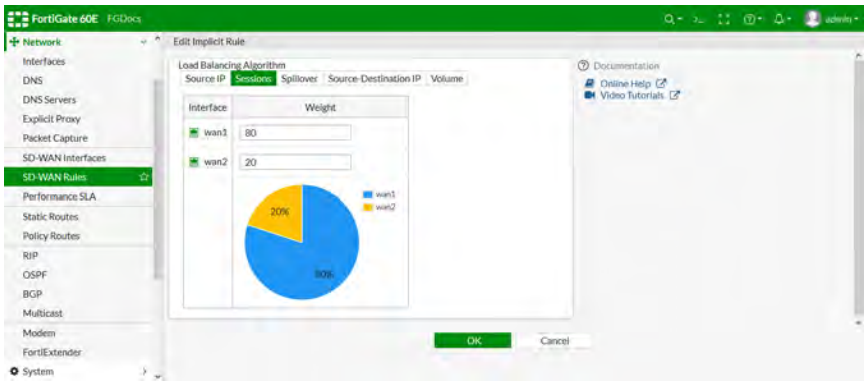
Example 2

Outgoing traffic is balanced between wan1 and wan2 with a customized ratio, using *weight-based* mode: wan1 runs 80% of the sessions, and wan2 runs 20% of the sessions.

Sessions with the same source and destination IP addresses (`src-ip` and `dst-ip`) will be forwarded to the same path, but will still be considered in later session ratio calculations.

Using the GUI:

1. Go to *Network* > *SD-WAN Rules*.
2. Edit the *sd-wan* rule (the last default rule).
3. For the *Load Balancing Algorithm*, select *Sessions*.
4. Enter 80 in the *wan1* field, and 20 in the *wan2* field.



Netskope Exhibit 1017

5. Click *OK*.

Using the CLI:

```
config system virtual-wan-link
    set load-balance-mode weight-based
    config members
        edit 1
            set interface "wan1"
            set weight 80
        next
        edit 2
            set interface "wan2"
            set weight 20
        next
    end
end
```

Example 3

Outgoing traffic is balanced between wan1 and wan2 with a customized ratio, using *measured-volume-based* mode: wan1 runs 80% of the volume, and wan2 runs 20% of the volume.

Using the GUI:

1. Go to *Network* > *SD-WAN Rules*.
2. Edit the *sd-wan* rule (the last default rule).
3. For the *Load Balancing Algorithm*, select *Volume*.
4. Enter 80 in the *wan1* field, and 20 in the *wan2* field.
5. Click *OK*.

Using the CLI:

```
config system virtual-wan-link
    set load-balance-mode measured-volume-based
    config members
        edit 1
            set interface "wan1"
            set volume-ratio 80
        next
        edit 2
            set interface "wan2"
            set volume-ratio 20
        next
    end
end
```

Example 4

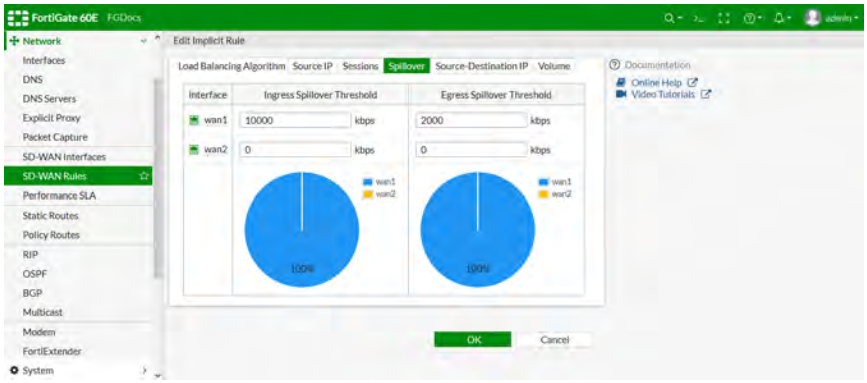
Load balancing can be used to reduce costs when internet connections are charged at different rates. For example, if wan2 charges based on volume usage and wan1 charges a fixed monthly fee, we can use wan1 at its maximum bandwidth, and use wan2 for overflow.

Netskope Exhibit 1017

In this example, wan1's bandwidth is 10Mbps down and 2Mbps up. Traffic will use wan1 until it reaches its spillover limit, then it will start to use wan2. Note that `auto-asic-offload` must be disabled in the firewall policy.

Using the GUI:

1. On the FortiGate, enable SD-WAN and add wan1 and wan2 as SD-WAN members, then add a policy and static route. See SD-WAN quick start on page 444 for details.
2. Go to *Network* > *SD-WAN Rules*.
3. Edit the *sd-wan* rule (the last default rule).
4. For the *Load Balancing Algorithm*, select *Spillover*.
5. Enter 10000 in the *wan1 Ingress Spillover Threshold* field, and 2000 in the *wan1 Egress Spillover Threshold* field.



6. Click *OK*.

Using the CLI:

```
config system virtual-wan-link
    set load-balance-mode usage-based
    config members
        edit 1
            set interface "wan1"
            set spillover-threshold 2000
            set ingress-spillover-threshold 10000
        next
    end
end
```

SD-WAN rules - best quality

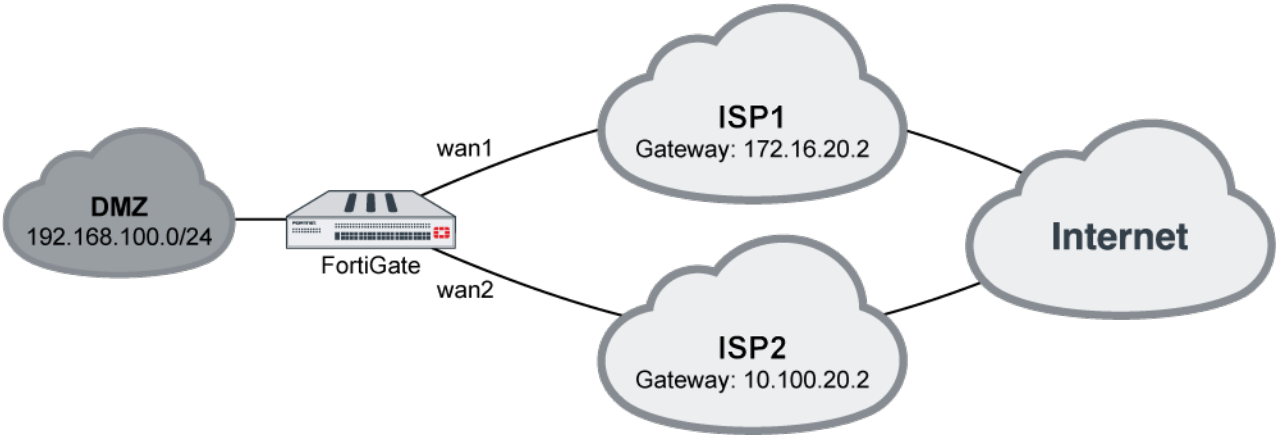
SD-WAN rules are used to control how sessions are distributed to SD-WAN members. Rules can be configured in one of five modes:

- `auto`: Interfaces are assigned a priority based on quality.
- Manual (`manual`): Interfaces are manually assigned a priority.
- Best Quality (`priority`): Interface are assigned a priority based on the link-cost-factor of the interface.
- Lowest Cost (SLA) (`sla`): Interfaces are assigned a priority based on selected SLA settings. See SD-WAN rules - lowest cost (SLA) on page 467.
- Maximize Bandwith (SLA) (`load-balance`): Traffic is distributed among all available links based on the selected load balancing algorithm. See SD-WAN rules - maximize bandwidth (SLA) on page 470.

Netskope Exhibit 1017

When using *Best Quality* mode, SD-WAN will choose the best link to forward traffic by comparing the *link-cost-factor*, selected from one of the following:

GUI	CLI	Description
Latency	`latency`	Select a link based on latency.
Jitter	`jitter`	Select a link based on jitter.
Packet Loss	`packet-loss`	Select a link based on packet loss.
Downstream	`inbandwidth`	Select a link based on available bandwidth of incoming traffic.
Upstream	`outbandwidth`	Select a link based on available bandwidth of outgoing traffic.
Bandwidth	`bibandwidth`	Select a link based on available bandwidth of bidirectional traffic.
custom-profile-1	`custom-profile-1`	Select link based on customized profile. If selected, set the following weights: • packet-loss-weight: Coefficient of packet-loss. • latency-weight: Coefficient of latency. • jitter-weight: Coefficient of jitter. • bandwidth-weight: Coefficient of reciprocal of available bidirectional bandwidth.

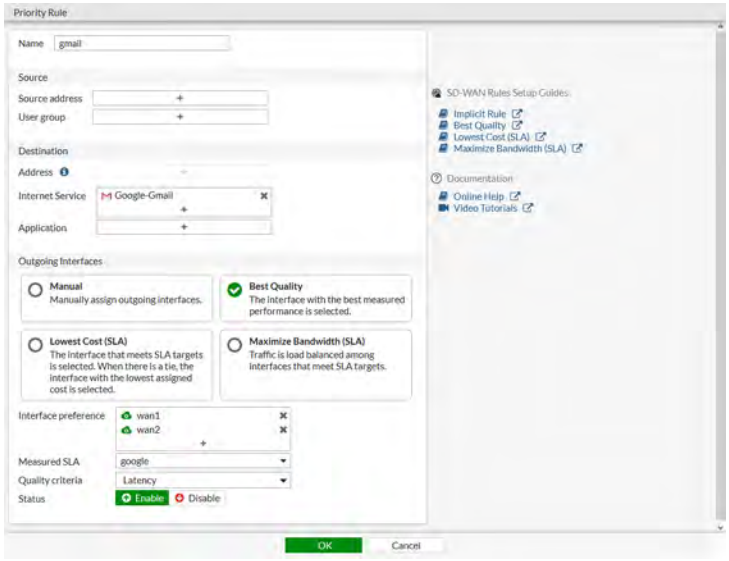


In this example, your wan1 and wan2 SD-WAN interfaces connect to two ISPs that both go to the public internet, and you want Gmail services to use the link with the least latency.

To configure an SD-WAN rule to use Best Quality:

1. On the FortiGate, add wan1 and wan2 as SD-WAN members, then add a policy and static route. See SD-WAN quick start on page 444 for details.
2. Create a new Performance SLA named *google*. See Performance SLA - link monitoring on page 455.
3. Go to *Network* > *SD-WAN Rules*.
4. Click *Create New*. The *Priority Rule* page opens.
5. Enter a name for the rule, such as *gmail*.

6. Configure the following settings:



Field	Setting
Internet Service	Google-Gmail
Strategy	Best Quality
Interface preference	wan1 and wan2
Measured SLA	google (created in step 2).
Quality criteria	Latency

7. Click *OK* to create the rule.

To configure an SD-WAN rule to use priority:

```
config system virtual-wan-link
    config health-check
        edit "google"
            set server "google.com"
            set members 1 2
        next
    end
    config service
        edit 1
            set name "gmail"
            set mode priority
            set internet-service enable
            set internet-service-id 65646
            set health-check "google"
            set link-cost-factor latency
            set priority-members 1 2
        next
    end
end
```

To diagnose the Performance SLA status:

```
FGT # diagnose sys virtual-wan-link health-check google
Health Check(google):
Seq(1): state(alive), packet-loss(0.000%) latency(14.563), jitter(4.334) sla_map=0x0
Seq(2): state(alive), packet-loss(0.000%) latency(12.633), jitter(6.265) sla_map=0x0

FGT # diagnose sys virtual-wan-link service 1
Service(1):

    TOS(0x0/0x0), protocol(0: 1->65535), Mode(priority), link-cost-facotr(latency), link-
cost-threshold(10), health-check(google) Members:

        1: Seq_num(2), alive, latency: 12.633, selected
        2: Seq_num(1), alive, latency: 14.563, selected

    Internet Service: Google-Gmail(65646)
```

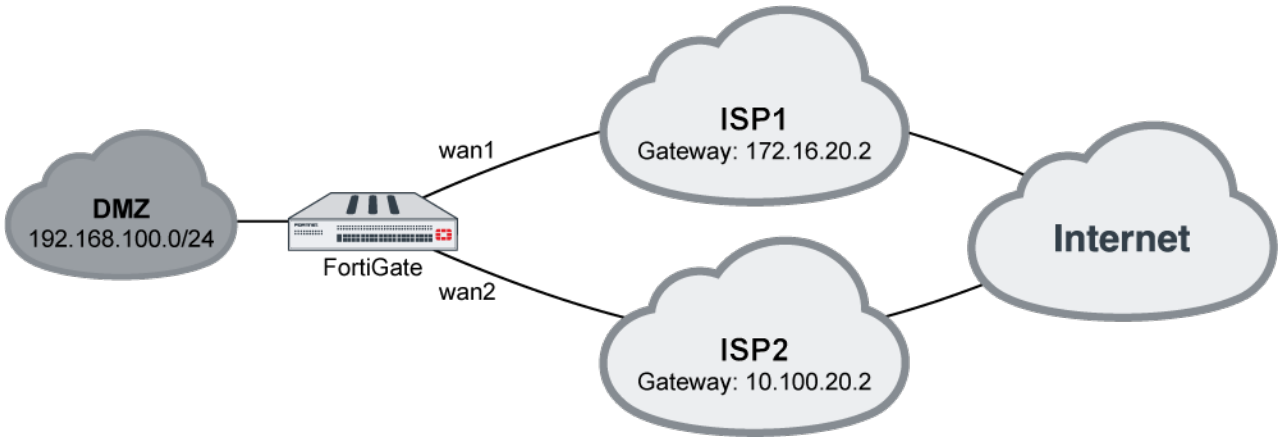
As wan2 has a smaller latency, SD-WAN will put Seq_num(2) on top of Seq_num(1) and wan2 will be used to forward Gmail traffic.

SD-WAN rules - lowest cost (SLA)

SD-WAN rules are used to control how sessions are distributed to SD-WAN members. Rules can be configured in one of five modes:

- `auto`: Interfaces are assigned a priority based on quality.
- Manual (`manual`): Interfaces are manually assigned a priority.
- Best Quality (`priority`): Interface are assigned a priority based on the link-cost-factor of the interface. See SD-WAN rules - best quality on page 464.
- Lowest Cost (SLA) (`sla`): Interfaces are assigned a priority based on selected SLA settings.
- Maximize Bandwidth (SLA) (`load-balance`): Traffic is distributed among all available links based on the selected load balancing algorithm. See SD-WAN rules - maximize bandwidth (SLA) on page 470.

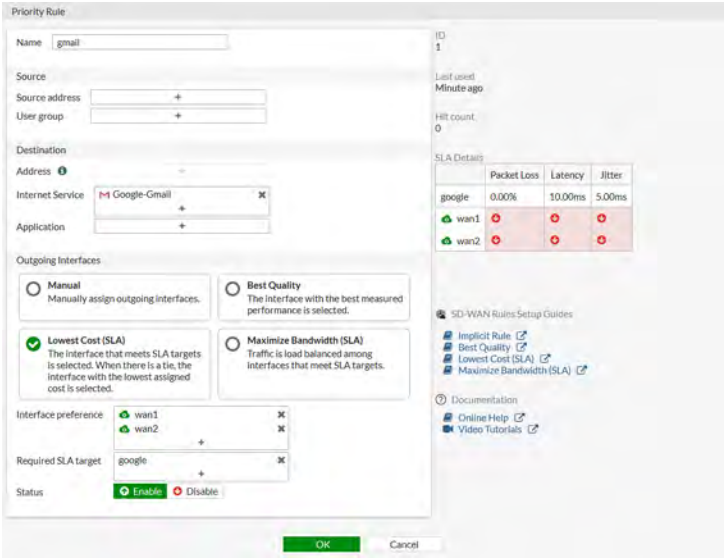
When using *Lowest Cost (SLA)* mode (`sla` in the CLI), SD-WAN will choose the lowest cost link that satisfies SLA to forward traffic. The lowest possible cost is `0`. If multiple eligible links have the same cost, the *Interface preference* order will be used to select a link.



In this example, your wan1 and wan2 SD-WAN interfaces connect to two ISPs that both go to the public internet. The cost of wan2 is less than that of wan1. You want to configure Gmail services to use the lowest cost interface, but the link quality must meet a standard of latency: 10ms, and jitter: 5ms.

To configure an SD-WAN rule to use Lowest Cost (SLA):

1. On the FortiGate, add wan1 and wan2 as SD-WAN members, then add a policy and static route. See SD-WAN quick start on page 444 for details.
2. Create a new Performance SLA named *google* that includes an SLA Target with *Latency threshold* = 10ms and *Jitter threshold* = 5ms. See Performance SLA - link monitoring on page 455.
3. Go to *Network* > *SD-WAN Rules*.
4. Click *Create New*. The *Priority Rule* page opens.
5. Enter a name for the rule, such as *gmail*.
6. Configure the following settings:



Field	Setting
Internet Service	Google-Gmail
Strategy	Lowest Cost (SLA)
Interface preference	wan1 and wan2
Required SLA target	google (created in step 2).

7. Click *OK* to create the rule.

To configure an SD-WAN rule to use SLA:

```
config system virtual-wan-link
    config members
        edit 1
            set interface "wan1"
            set cost 10
        next
        edit 2
```

```
            set interface "wan2"
            set cost 5
        next
    end
    config health-check
        edit "google"
            set server "google.com"
            set members 1 2
            config sla
                edit 1
                    set latency-threshold 10
                    set jitter-threshold 5
                next
            end
        next
    end
    config service
        edit 1
            set name "gmail"
            set mode sla
            set internet-service enable
            set internet-service-id 65646
            config sla
                edit "google"
                    set id 1
                next
            end
            set priority-members 1 2
        next
    end
end
```

To diagnose the Performance SLA status:

```
FGT # diagnose sys virtual-wan-link health-check google
Health Check(google):
Seq(1): state(alive), packet-loss(0.000%) latency(14.563), jitter(4.334) sla_map=0x0
Seq(2): state(alive), packet-loss(0.000%) latency(12.633), jitter(6.265) sla_map=0x0

FGT # diagnose sys virtual-wan-link service 1
Service(1): Address Mode(IPV4) flags=0x0

    TOS(0x0/0x0), Protocol(0: 1->65535), Mode(sla)
    Members:<<BR>>

        1: Seq_num(2), alive, sla(0x1), cfg_order(1), selected
        2: Seq_num(1), alive, sla(0x1), cfg_order(0), selected

    Internet Service: Google.Gmail(65646)
```

When both wan1 and wan2 meet the SLA requirements, Gmail traffic will only use wan2. If only wan1 meets the SLA requirements, Gmail traffic will only use wan1, even though it has a higher cost. If neither interface meets the requirements, wan2 will be used.

If both interface had the same cost and both met the SLA requirements, the first link configured in `set priority-members` would be used.

Netskope Exhibit 1017

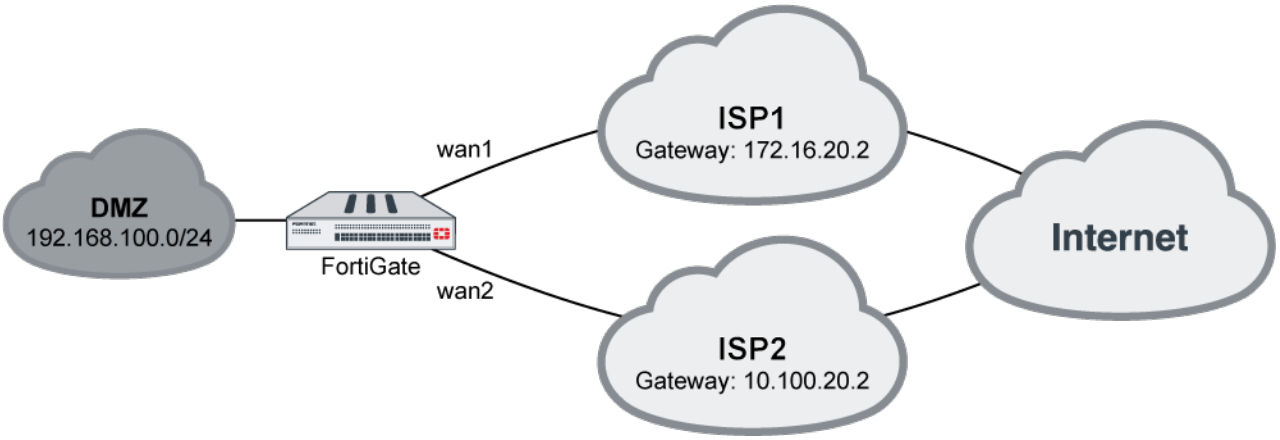
SD-WAN rules - maximize bandwidth (SLA)

SD-WAN rules are used to control how sessions are distributed to SD-WAN members. Rules can be configured in one of five modes:

- `auto`: Interfaces are assigned a priority based on quality.
- Manual (`manual`): Interfaces are manually assigned a priority.
- Best Quality (`priority`): Interface are assigned a priority based on the link-cost-factor of the interface. See SD-WAN rules - best quality on page 464.
- Lowest Cost (SLA) (`sla`): Interfaces are assigned a priority based on selected SLA settings. See SD-WAN rules - lowest cost (SLA) on page 467.
- Maximize Bandwidth (SLA) (`load-balance`): Traffic is distributed among all available links based on the selected load balancing algorithm.

When using *Maximize Bandwidth* mode (`load-balance` in the CLI), SD-WAN will all of the links that satisfies SLA to forward traffic based on a round-robin load balancing algorithm.

ADVPN is not supported in this mode.

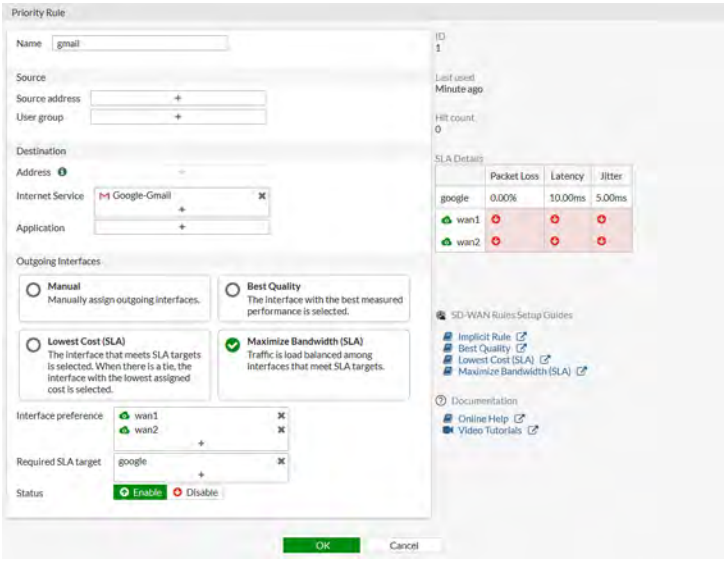


In this example, your wan1 and wan2 SD-WAN interfaces connect to two ISPs that both go to the public internet. You want to configure Gmail services to use both of the interface, but the link quality must meet a standard of latency: 10ms, and jitter: 5ms. This can maximize the bandwidth usage.

To configure an SD-WAN rule to use Maximize Bandwidth (SLA):

1. On the FortiGate, add wan1 and wan2 as SD-WAN members, then add a policy and static route. See SD-WAN quick start on page 444 for details.
2. Create a new Performance SLA named *google* that includes an SLA Target 1 with *Latency threshold* = 10ms and *Jitter threshold* = 5ms. See Performance SLA - link monitoring on page 455.
3. Go to *Network* > *SD-WAN Rules*.
4. Click *Create New*. The *Priority Rule* page opens.
5. Enter a name for the rule, such as *gmail*.

6. Configure the following settings:



Field	Setting
Internet Service	Google-Gmail
Strategy	Maximize Bandwidth (SLA)
Interface preference	wan1 and wan2
Required SLA target	google (created in step 2).

7. Click *OK* to create the rule.

To configure an SD-WAN rule to use SLA:

```
config system virtual-wan-link
    config health-check
        edit "google"
            set server "google.com"
            set members 1 2
            config sla
                edit 1
                    set latency-threshold 10
                    set jitter-threshold 5
                next
            end
        next
    end
    config service
        edit 1
            set name "gmail"
            set mode load-balance
            set internet-service enable
            set internet-service-id 65646
            config sla
                edit "google"
                    set id 1
```


Netskope Exhibit 1017

```
                next
            end
            set priority-members 1 2
        next
    end
end
```

To diagnose the performance SLA status:

```
FGT # diagnose sys virtual-wan-link health-check google
Health Check(google):
Seq(1): state(alive), packet-loss(0.000%) latency(14.563), jitter(4.334) sla_map=0x0
Seq(2): state(alive), packet-loss(0.000%) latency(12.633), jitter(6.265) sla_map=0x0

FGT # diagnose sys virtual-wan-link service 1
Service(1): Address Mode(IPV4) flags=0x0

    TOS(0x0/0x0), Protocol(0: 1->65535), Mode(load-balance)
    Members:<<BR>>

        1: Seq_num(1), alive, sla(0x1), num of pass(1), selected
        2: Seq_num(2), alive, sla(0x1), num of pass(1), selected

    Internet Service: Google.Gmail(65646)
```

When both wan1 and wan2 meet the SLA requirements, Gmail traffic will use both wan1 and wan2. If only one of the interfaces meets the SLA requirements, Gmail traffic will only use that interface.

If neither interface meets the requirements but health-check is still alive, then wan1 and wan2 tie. The traffic will try to balance between wan1 and wan2, using both interfaces to forward traffic.

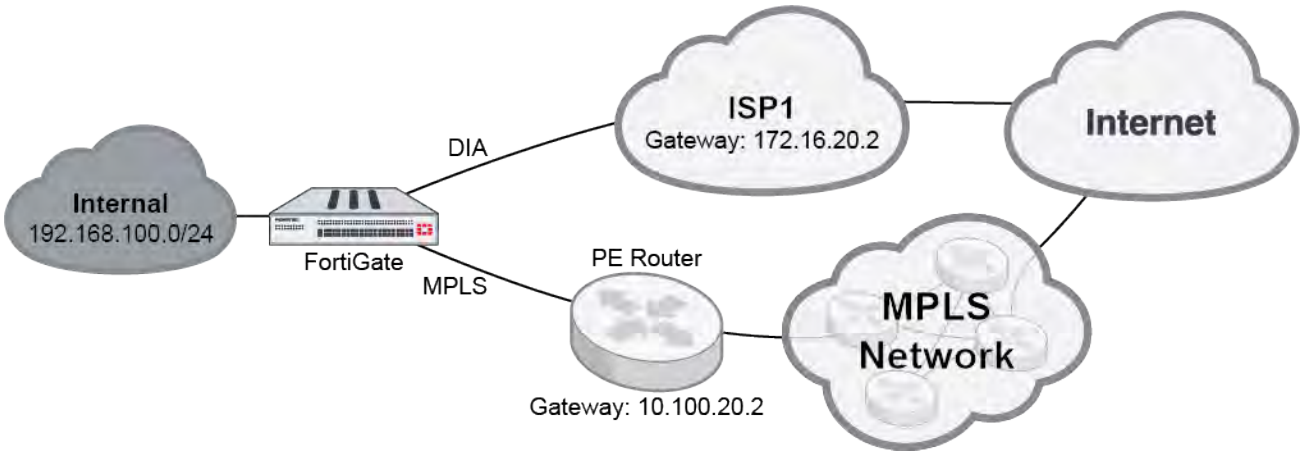
Application steering using SD-WAN rules

This topic covers how to use application steering in a topology with multiple WAN links. The following examples illustrate how to use different strategies to perform application steering to accommodate different business needs:

- Static application steering with a manual strategy on page 472
- Dynamic application steering with lowest cost and best quality strategies on page 476

Static application steering with a manual strategy

This example covers a typical usage scenario where the SD-WAN has two members: MPLS and DIA. DIA is primarily used for direct internet access to internet applications, such as Office365, Google applications, Amazon, and Dropbox. MPLS is primarily used for SIP, and works as a backup when DIA is not working.

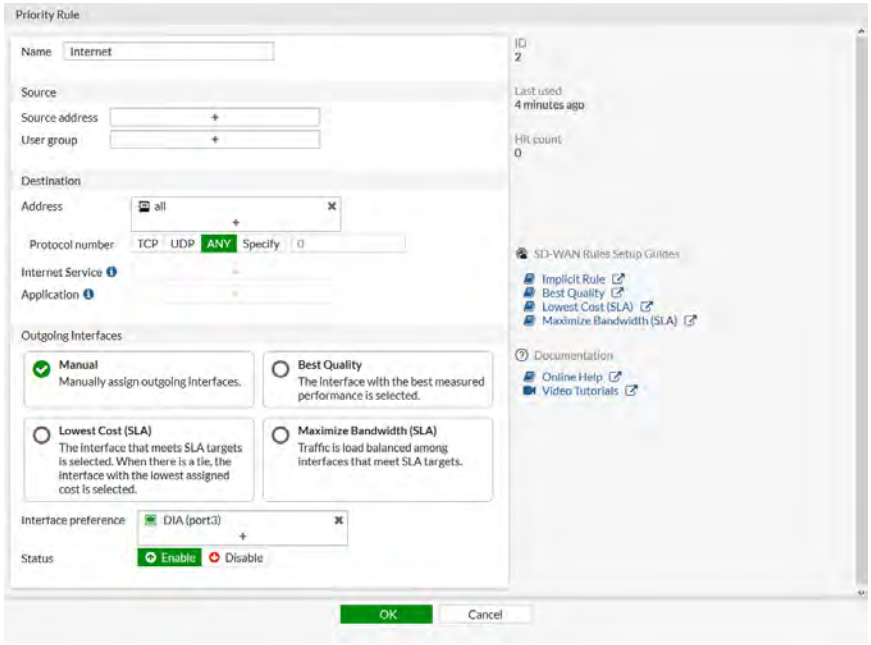


This example configures all SIP traffic to use MPLS while all other traffic uses DIA. If DIA is not working, the traffic will use MPLS.

To configure an SD-WAN rule to use SIP and DIA in the GUI:

1. Add port1 (DIA) and port2 (MPLS) as SD-WAN members, and configure a static route. See Configuring the SD-WAN interface on page 445 for details.
2. Create a firewall policy with an *Application Control* profile configured. See Configuring firewall policies for SD-WAN on page 447 for details.
3. Go to *Network* > *SD-WAN Rules*.
4. Click *Create New*. The *Priority Rule* page opens.
5. Enter a name for the rule, such as *SIP*.
6. Click the *Application* field and select the applicable SIP applications from the *Select Entries* panel.
7. Under *Outgoing Interfaces*, select *Manual*.
8. For *Interface preference*, select *MPLS*.
9. Click *OK*.
10. Click *Create New* to create another rule.
11. Enter a name for the rule, such as *Internet*.
12. Click the *Address* field and select *all* from the panel.
13. Under *Outgoing Interfaces*, select *Manual*.

14. For *Interface preference*, select *DIA*.



15. Click *OK*.

To configure the firewall policy using the CLI:

```
config firewall policy
    edit 1
        set name "1"
        set srcintf "dmz"
        set dstintf "virtual-wan-link"
        set srcaddr "all"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ALL"
        set utm-status enable
        set fsso disable
        set application-list "default"
        set ssl-ssh-profile "certificate-inspection"
        set nat enable
    next
end
```

To configure an SD-WAN rule to use SIP and DIA using the CLI:

```
config system virtual-wan-link
    set status enable
    config members
        edit 1
            set interface "MPLS"
        next
        edit 2
            set interface "DIA"
        next
```

```
    end
    config service
        edit 1
            set name "SIP"
            set internet-service enable
            set internet-service-app-ctrl 34640 152305677 38938 26180 26179 30251
            set priority-members 2
        next
        edit 2
            set name "Internet"
            set dst "all"
            set priority-members 1
        next
    end
end
```

All SIP traffic uses MPLS. All other traffic goes to DIA. If DIA is broken, the traffic uses MPLS. If you use VPN instead of MPLS to run SIP traffic, you must configure a VPN interface, for example vpn1, and then replace member 1 from MPLS to vpn1 for SD-WAN member.

To use the diagnose command to check performance SLA status using the CLI:

```
# diagnose sys virtual-wan-link service 1

Service(1): Address Mode(IPV4) flags=0x0

TOS(0x0/0x0), Protocol(0: 1->65535), Mode(manual)
Members:<<BR>>

1: Seq_num(1), alive, selected

Internet Service: SIP(4294836224 34640) SIP.Method(4294836225 152305677) SIP.Via.NAT
(4294836226 38938) SIP_Media.Type.Application(4294836227 26180) SIP_Message(4294836228
26179) SIP_Voice(4294836229 30251)
```

```
# diagnose sys virtual-wan-link service 2

Service(2): Address Mode(IPV4) flags=0x0

TOS(0x0/0x0), Protocol(0: 1->65535), Mode(manual)
Members:<<BR>>

1: Seq_num(2), alive, selected

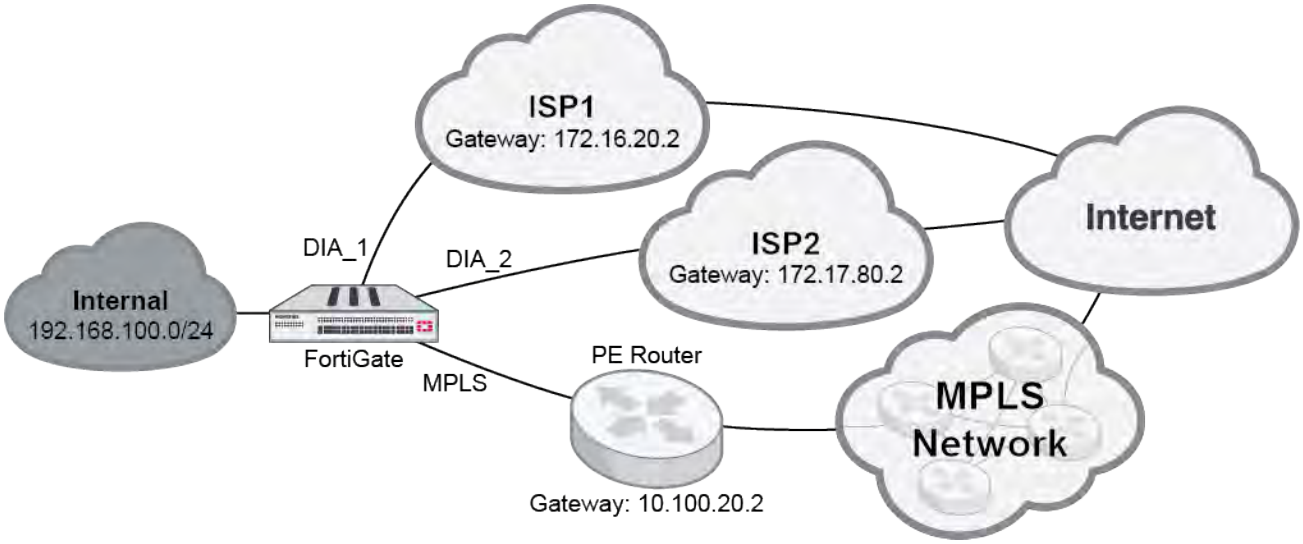
Dst address: 0.0.0.0-255.255.255.255
```

```
# diagnose sys virtual-wan-link internet-service-app-ctrl-list
Ctrl application(SIP 34640):Internet Service ID(4294836224)
Ctrl application(SIP.Method 152305677):Internet Service ID(4294836225)
Ctrl application(SIP.Via.NAT 38938):Internet Service ID(4294836226)
Ctrl application(SIP_Media.Type.Application 26180):Internet Service ID(4294836227)
Ctrl application(SIP_Message 26179):Internet Service ID(4294836228)
Ctrl application(SIP_Voice 30251):Internet Service ID(4294836229)
```


Netskope Exhibit 1017

Dynamic application steering with lowest cost and best quality strategies

In this example, the SD-WAN has three members: two ISPs (DIA_1 and DIA_2) that are used for access to internet applications, and an MPLS link that is used exclusively as a backup for business critical applications.



Business applications, such as Office365, Google, Dropbox, and SIP, use the *Lowest Cost (SLA)* strategy to provide application steering, and traffic falls back to MPLS only if both ISP1 and ISP2 are down. Non-business applications, such as Facebook and Youtube, use the *Best Quality* strategy to choose between the ISPs.

To configure the SD-WAN members, static route, and firewall policy in the GUI:

1. Add port1 (DIA_1), port2 (DIA_2), and port3 (MPLS) as SD-WAN members. Set the cost of DIA_1 and DIA_2 to *0*, and MPLS to *20*. See Configuring the SD-WAN interface on page 445 for details.



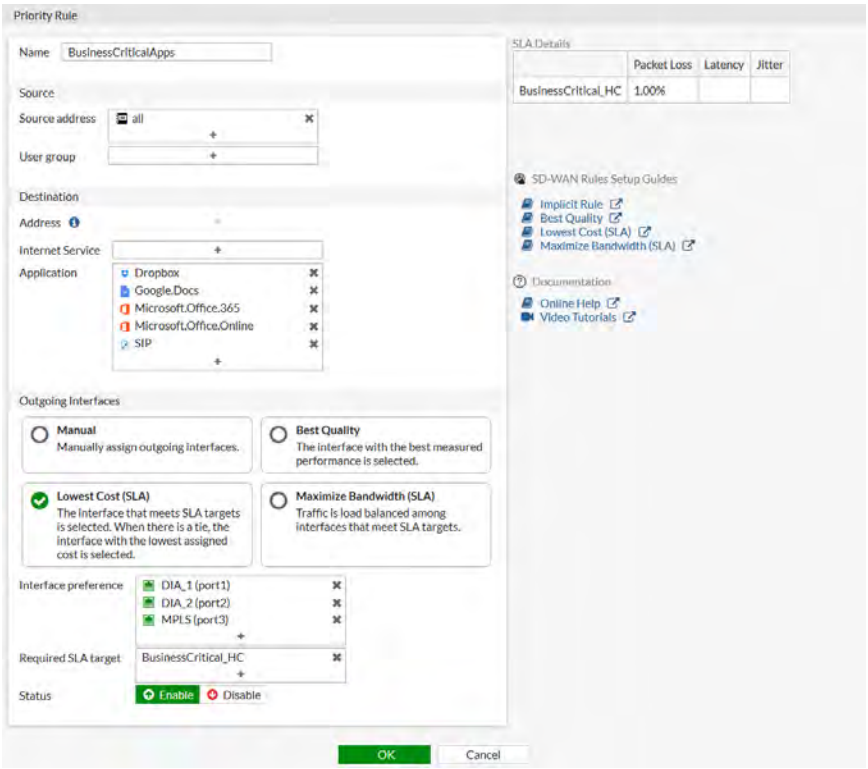
2. Configure a static route. See Adding a static route on page 446 for details.
3. Create a firewall policy to allow traffic out on SD-WAN, with an *Application Control* profile configured. See Configuring firewall policies for SD-WAN on page 447 for details.

To configure the SD-WAN rule and performance SLA checks for business critical application in the GUI:

1. Go to *Network* > *SD-WAN Rules*, and click *Create New*.
2. Set the name to *BusinessCriticalApps*.
   This rule will steer your business critical traffic to the appropriate link based on the *Lowest Cost (SLA)*.
3. Set *Source address* to *all*.
4. Under *Destination*, set *Application* to your required applications. In this example: *Microsoft.Office.365*, *Microsoft.Office.Online*, *Google.Docs*, *Dropbox*, and *SIP*.

Netskope Exhibit 1017

5. Under *Outgoing Interfaces*, select *Lowest Cost (SLA)*.
   The lowest cost is defined in the SD-WAN member interface settings (see Configuring the SD-WAN interface on page 445). The lowest possible cost is *0*, which represents the most preferred link. In this example, DIA_1 and DIA_2 both have a cost of *0*, while MPLS has a cost of *20* because it is used for backup.
6. In *Interface preference*, add the interfaces in order of preference when the cost of the links is tied. In this example, DIA_1, DIA_2, then MPLS.
   MPLS will always be chosen last, because it has the highest cost. DIA_1 and DIA_2 have the same cost, so an interface is selected based on their order in the *Interface preference* list.
7. Set *Required SLA target* to ensure that only links that pass your SLA target are chosen in this SD-WAN rule:
   a. Click in the *Required SLA target* field.
   b. In the *Select Entries* pane, click *Create*. The *New Performace SLA* pane opens.
   c. Set *Name* to *BusinessCritical_HC*.
      This health check is used for business critical applications in your SD-WAN rule.
   d. Leave *Protocol* set to *Ping*, and add up to two servers, such as *office.com* and *google.com*.
   e. Set *Participants* to *Specify*, and add all three interfaces: DIA_1, DIA_2, and MPLS.
   f. Enable *SLA Target*.
      The attributes in your target determine the quality of your link. The SLA target of each link is compared when determining which link to use based on the lowest cost. Links that meet the SLA target are preferred over links that fail, and move to the next step of selection based on cost. If no links meet the SLA target, then they all move to the next step.
      In this example, disable *Latency threshold* and *Jitter threshold*, and set *Packet loss threshold* to *1*.
   g. Click *OK*.
   h. Select the new performance SLA to set it as the *Required SLA target*.
   When multiple SLA targets are added, you can choose which target to use in the SD-WAN rule.



8. Click *OK* to create the SD-WAN rule.

Netskope Exhibit 1017

To configure the SD-WAN rule and performance SLA checks for non-business critical application in the GUI:

1. Go to *Network* > *SD-WAN Rules*, and click *Create New*.
2. Set the name to *NonBusinessCriticalApps*.
   This rule will steer your non-business critical traffic to the appropriate link based on the *Best Quality*. No SLA target must be met, as the best link is selected based on the configured quality criteria and interface preference order.
3. Set *Source address* to *all*.
4. Under *Destination*, set *Application* to your required applications. In this example: *Facebook*, and *Youtube*.
5. Under *Outgoing Interfaces*, select *Best Quality*.
6. In *Interface preference*, add the interfaces in order of preference.
   By default, a more preferred link has an advantage of 10% over a less preferred link. For example, when latency is used, the preferred link's calculated latency = real latency / (1+10%).

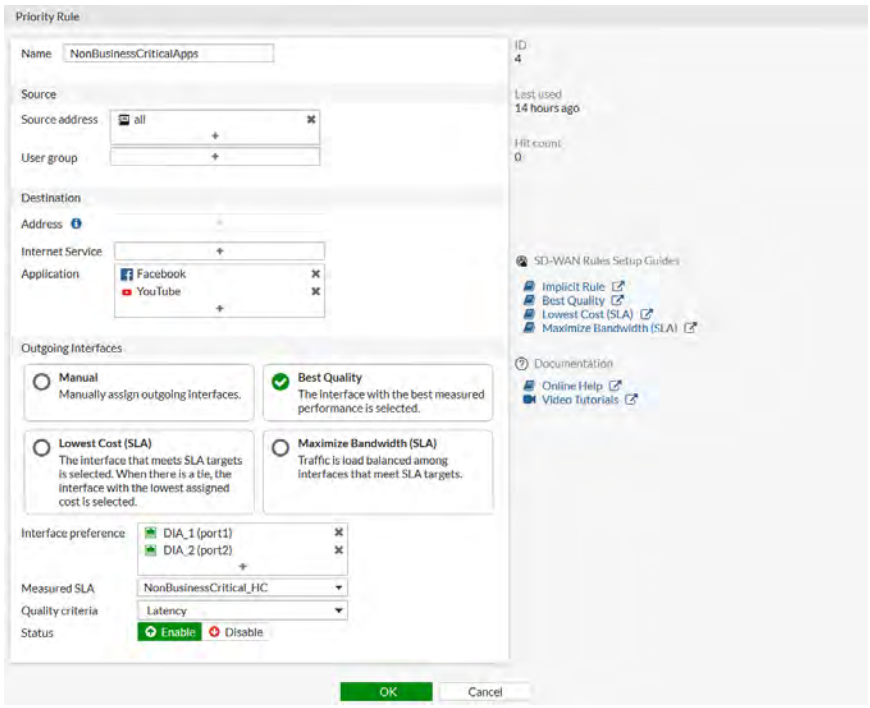
   

   The preferred link advantage can be customized in the CLI when the mode is `priority` (*Best Quality*) or `auto`:

   ```
   config system virtual-wan-link
       config service
           edit <id>
               set link-cost-threshold <integer>
           next
       end
   end
   ```

7. Create and apply a new performance SLA profile:
   a. Click in the *Measured SLA* field.
   b. In the drop-down list, click *Create*. The *New Performace SLA* pane opens.
   c. Set *Name* to *NonBusinessCritical_HC*.
      This health check is used for non-business critical applications in your SD-WAN rule.
   d. Leave *Protocol* set to *Ping*, and add up to two servers, such as *youtube.com* and *facebook.com*.
   e. Set *Participants* to *Specify*, and add the DIA_1 and DIA_2 interfaces. In this example, MPLS is not used for non-business critical applications.
   f. Leave *SLA Target* disabled.
   g. Click *OK*.
   h. Select the new performance SLA from the list to set it as the *Measured SLA*.
8. Set *Quality criteria* as required. In this example, *Latency* is selected.
   For bandwidth related criteria, such as *Downstream*, *Upstream*, and *Bandwidth* (bi-directional), the selection is based on available bandwidth. An estimated bandwidth should be configured on the interface to provide a baseline, maximum available bandwidth.

Netskope Exhibit 1017



9. Click *OK* to create the SD-WAN rule.

To configure the SD-WAN members, static route, and firewall policy in the CLI:

1. Configure the interfaces:

```
config system interface
    edit "port1"
        set ip <class_ip&net_netmask>
        set alias "DIA_1"
        set role wan
    next
    edit "port2"
        set ip <class_ip&net_netmask>
        set alias "DIA_2"
        set role wan
    next
    edit "port3"
        set ip <class_ip&net_netmask>
        set alias "MPLS"
        set role wan
    next
end
```

2. Configure the SD-WAN members:

```
config system virtual-wan-link
    set status enable
    config members
        edit 1
            set interface "port1"
            set gateway 172.16.20.2
        next
```

```
        edit 2
            set interface "port2"
            set gateway 172.17.80.2
        next
        edit 3
            set interface "port3"
            set gateway 10.100.20.2
            set cost 20
        next
    end
end
```

To configure the SD-WAN rule and performance SLA checks for business critical application in the CLI:

1. Configure the *BusinessCriticalApps_HC* health-check:

```
config system virtual-wan-link
    config health-check
        edit "BusinessCriticalApps_HC"
            set server "office.com" "google.com"
            set members 1 2 3
            config sla
                edit 1
                    set link-cost-factor packet-loss
                    set packetloss-threshold 1
                next
            end
        next
    end
end
```

2. Configure the *BusinessCriticalApps* service to use *Lowest Cost (SLA)*:

```
config system virtual-wan-link
    config service
        edit 1
            set name "BusinessCriticalApps"
            set mode sla
            set src "all"
            set internet-service enable
            set internet-service-app-ctrl 17459 16541 33182 16177 34640
            config sla
                edit "BusinessCriticalApps_HC"
                    set id 1
                next
            end
            set priority-members 1 2 3
        next
    end
end
```


Netskope Exhibit 1017

To configure the SD-WAN rule and performance SLA checks for non-business critical application in the CLI:

1. Configure the *nonBusinessCriticalApps_HC* health-check:

```
config system virtual-wan-link
    config health-check
        edit "NonBusinessCriticalApps_HC"
            set server "youtube.com" "facebook.com"
            set members 1 2
        next
    end
end
```

2. Configure the *BusinessCriticalApps* service to use *Lowest Cost (SLA)*:

```
config system virtual-wan-link
    config service
        edit 4
            set name "NonBusinessCriticalApps"
            set mode priority
            set src "all"
            set internet-service enable
            set internet-service-app-ctrl 15832 31077
            set health-check "NonBusinessCriticalApps_HC"
            set priority-members 1 2
        next
    end
end
```

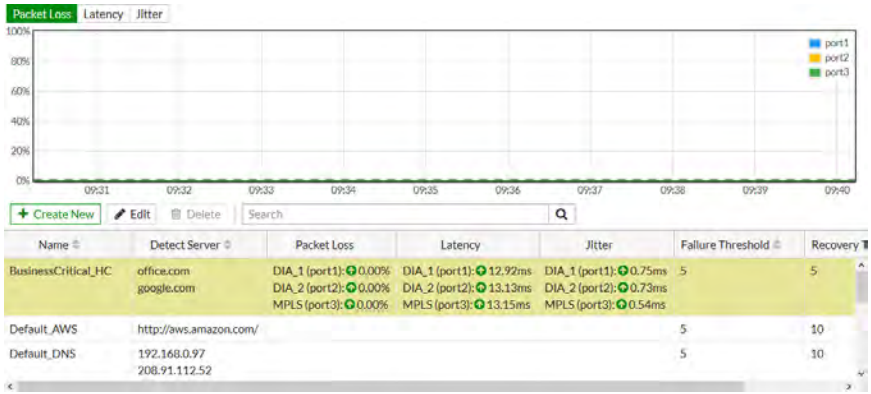
Verification

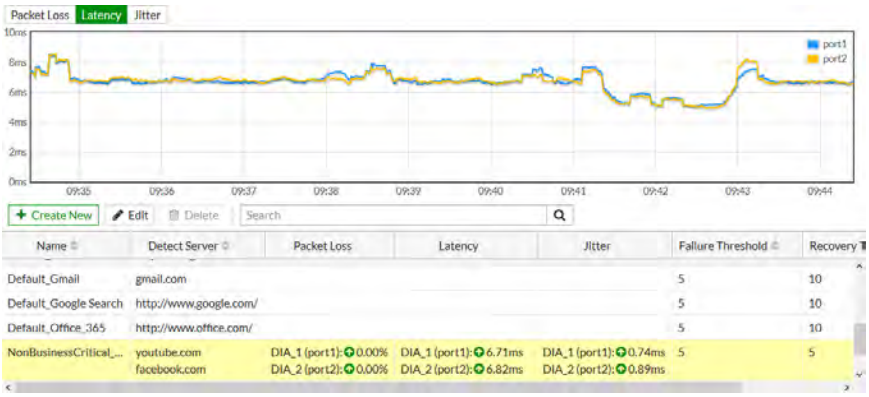
Check the following GUI pages, and run the following CLI commands to confirm that your traffic is being steered by the SD-WAN rules.

Health checks

To verify the status of each of the health checks in the GUI:

1. Go to *Network* > *Performance SLA* and select each of the health checks from the list.





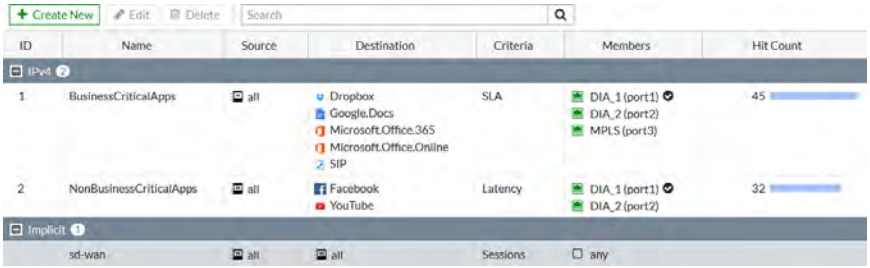
To verify the status of each of the health checks in the CLI:

```
# diagnose sys virtual-wan-link health-check
Health Check(BusinessCritical_HC):
Seq(1 port1): state(alive), packet-loss(0.000%) latency(12.884), jitter(0.919) sla_map=0x1
Seq(2 port2): state(alive), packet-loss(0.000%) latency(13.018), jitter(0.723) sla_map=0x1
Seq(3 port3): state(alive), packet-loss(0.000%) latency(13.018), jitter(0.923) sla_map=0x1
Health Check(NonBusinessCritical_HC):
Seq(1 port1): state(alive), packet-loss(0.000%) latency(6.888), jitter(0.953) sla_map=0x0
Seq(2 port2): state(alive), packet-loss(0.000%) latency(6.805), jitter(0.830) sla_map=0x0
```

Rule members and hit count

To verify the active members and hit count of the SD-WAN rule in the GUI:

1. Go to *Network* > *SD-WAN Rules*.



   The interface that is currently selected by the rule has a checkmark next to its name in the *Members* column. Hover the cursor over the checkmark to open a tooltip that gives the reason why that member is selected. If multiple members are selected, only the highest ranked member is highlighted (unless the mode is *Maximize Bandwidth (SLA)*).

To verify the active members and hit count of the SD-WAN rule in the CLI:

```
# diagnose sys virtual-wan-link service

Service(3): Address Mode(IPV4) flags=0x0
  Gen(13), TOS(0x0/0x0), Protocol(0: 1->65535), Mode(sla), sla-compare-order
  Members:
    1: Seq_num(1 port1), alive, sla(0x1), cfg_order(0), cost(0), selected
    2: Seq_num(2 port2), alive, sla(0x1), cfg_order(1), cost(0), selected
    3: Seq_num(3 port3), alive, sla(0x1), cfg_order(2), cost(20), selected
```


Netskope Exhibit 1017

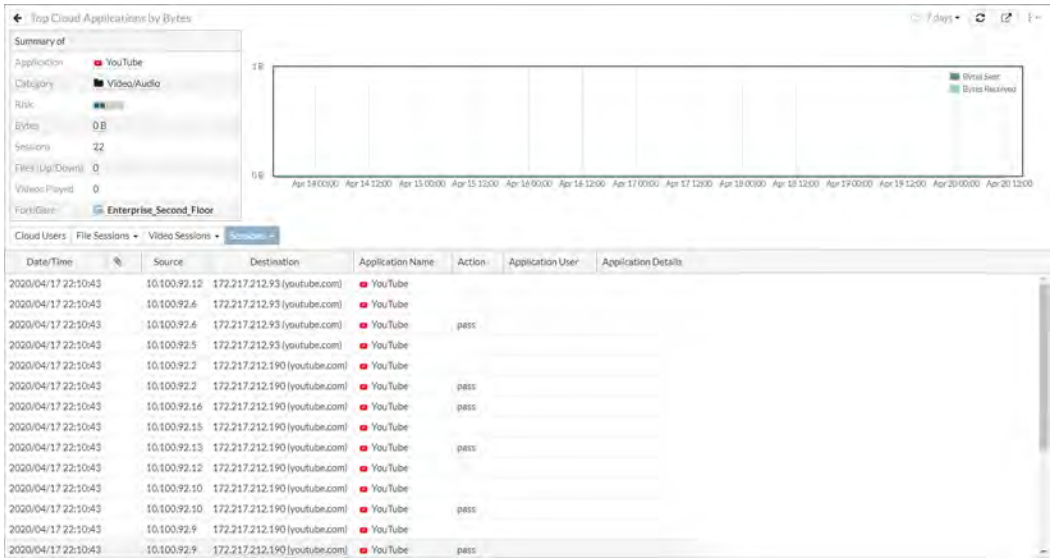
```
  Internet Service: Dropbox(4294836727,0,0,0 17459) Google.Docs(4294836992,0,0,0 16541)
Microsoft.Office.365(4294837472,0,0,0 33182) Microsoft.Office.Online(4294837475,0,0,0 16177)
SIP(4294837918,0,0,0 34640)
  Src address:
       0.0.0.0-255.255.255.255

Service(4): Address Mode(IPV4) flags=0x0
  Gen(211), TOS(0x0/0x0), Protocol(0: 1->65535), Mode(priority), link-cost-factor(latency),
link-cost-threshold(10), heath-check(NonBusinessCritical_HC)
  Members:
    1: Seq_num(1 port1), alive, latency: 5.712, selected
    2: Seq_num(2 port2), alive, latency: 5.511, selected
  Internet Service: Facebook(4294836806,0,0,0 15832) YouTube(4294838537,0,0,0 31077)
  Src address:
       0.0.0.0-255.255.255.255
```

Applications and sessions

To verify sessions in FortiView:

1. Go to a dashboard and add the *Top Cloud Applications by Bytes* widget. See Cloud application view on page 96 for details.
2. Drill down on an application, such as *YouTube*, then select the *Sessions* tab.



To verify applications identified by Application Control in SD-WAN:

```
# diagnose sys virtual-wan-link internet-service-app-ctrl-list

Facebook(15832 4294836697): 31.13.67.20 6 443 Fri April 17 22:33:39 2020
Facebook(15832 4294836697): 31.13.67.35 6 443 Fri April 17 22:33:41 2020
Facebook(15832 4294836697): 31.13.70.36 6 443 Fri April 17 22:36:41 2020
Facebook(15832 4294836697): 157.240.11.22 6 443 Fri April 17 22:36:42 2020
Facebook(15832 4294836697): 157.240.11.35 6 443 Fri April 17 22:36:41 2020
YouTube(31077 4294838227): 172.217.24.150 6 443 Fri April 17 22:32:16 2020
YouTube(31077 4294838227): 172.217.25.78 6 443 Fri April 17 22:32:16 2020
YouTube(31077 4294838227): 216.58.220.129 6 443 Fri April 17 22:32:34 2020
```


Netskope Exhibit 1017

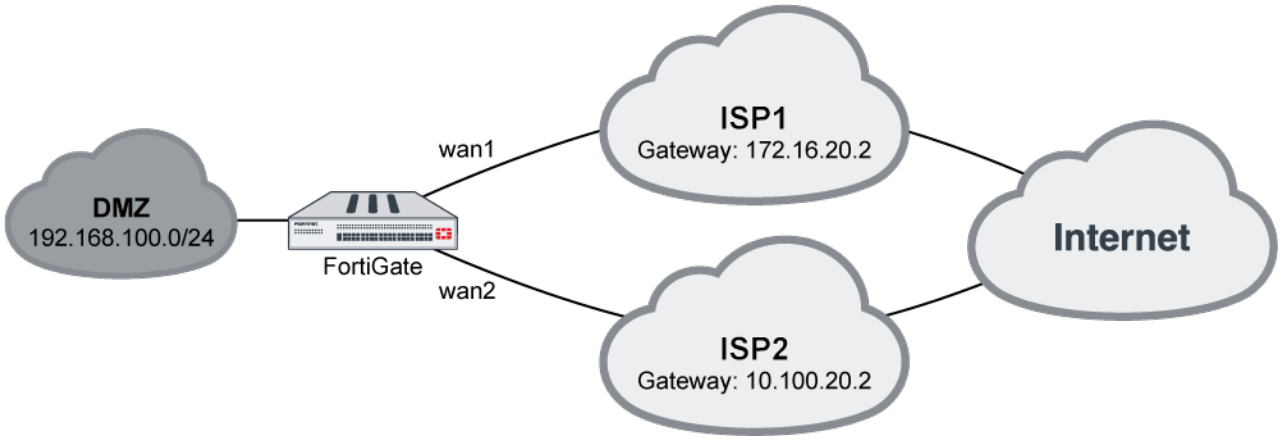
SD-WAN traffic shaping and QoS

Use a traffic shaper in a firewall shaping policy to control traffic flow. You can use it to control maximum and guaranteed bandwidth, or put certain traffic to one of the three different traffic priorities: high, medium, or low.

An advanced shaping policy can classify traffic into 30 groups. Use a shaping profile to define the percentage of the interface bandwidth that is allocated to each group. Each group of traffic is shaped to the assigned speed limit based on the outgoing bandwidth limit configured on the interface.

For more information, see the online help on shared policy traffic shaping and interface-based traffic shaping.

Sample topology



Sample configuration

This example shows a typical customer usage where the customer's SD-WAN has two member: wan1 and wan2 and each is 10Mb/s.

An overview of the procedures to configure SD-WAN traffic shaping and QoS with SD-WAN includes:

1. Give HTTP/HTTPS traffic high priority and give FTP low priority so that if there are conflicts, FortiGate will forward HTTP/HTTPS traffic first.
2. Even though FTP has low priority, configure FortiGate to give it a 1Mb/s guaranteed bandwidth on each SD-WAN member so that if there is no FTP traffic, other traffic can use all the bandwidth. If there is heavy FTP traffic, it can still be guaranteed a 1Mb/s bandwidth.
3. Traffic going to specific destinations such as a VOIP server uses wan1 to forward, and SD-WAN forwards with an Expedited Forwarding (EF) DSCP tag 101110.

To configure SD-WAN traffic shaping and QoS with SD-WAN in the GUI:

1. On the FortiGate, add wan1 and wan2 as SD-WAN members, then add a policy and static route.
   See SD-WAN quick start on page 444.
2. When you add a firewall policy, enable *Application Control*.
3. Go to *Policy & Objects* > *Traffic Shapers* and edit *low-priority*.
   a. Enable *Guaranteed Bandwidth* and set it to *1000* kbps.

4. Go to *Policy & Objects* > *Traffic Shaping Policy* and click *Create New*.
   a. Name the traffic shaping policy, for example, *HTTP-HTTPS*.
   b. Click the *Source* box and select *all*.
   c. Click the *Destination* box and select *all*.
   d. Click the *Service* box and select *HTTP* and *HTTPS*.
   e. Click the *Outgoing Interface* box and select *SD-WAN*.
   f. Enable both *Shared Shaper* and *Reverse Shaper* and select *high-priority* for both options.
   g. Click *OK*.
5. Go to *Policy & Objects* > *Traffic Shaping Policy* and click *Create New*.
   a. Name the traffic shaping policy, for example, *FTP*.
   b. Click the *Source* box and select *all*.
   c. Click the *Destination* box and select *all*.
   d. Click the *Service* box and select *FTP*, *FTP_GET*, and *FTP_PUT*.
   e. Click the *Outgoing Interface* box and select *SD-WAN*.
   f. Enable both *Shared Shaper* and *Reverse Shaper* and select *low-priority* for both options.
   g. Click *OK*
6. Go to *Network* > *SD-WAN Rules* and click *Create New*.
   a. Enter a name for the rule, such as *Internet*.
   b. In the *Destination* section, click the *Address* field and select the VOIP server you created in the firewall address.
   c. Under *Outgoing Interfaces*, select *Manual*.
   d. For *Interface preference*, select *wan1*.
   e. Click *OK*.
7. Use CLI commands to modify DSCP settings. See the DSCP CLI commands below.

To configure the firewall policy using the CLI:

```
config firewall policy
    edit 1
        set name "1"
        set srcintf "dmz"
        set dstintf ""virtual-wan-link""
        set srcaddr "all"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ALL"
        set nat enable
    next
end
```

To configure the firewall traffic shaper priority using the CLI:

```
config firewall shaper traffic-shaper
    edit "high-priority"
        set maximum-bandwidth 1048576
        set per-policy enable
    next
    edit "low-priority"
        set guaranteed-bandwidth 1000
```


Netskope Exhibit 1017

```
        set maximum-bandwidth 1048576
        set priority low
        set per-policy enable
    next
end
```

To configure the firewall traffic shaping policy using the CLI:

```
config firewall shaping-policy
    edit 1
        set name "http-https"
        set service "HTTP" "HTTPS"
        set dstintf "virtual-wan-link"
        set traffic-shaper "high-priority"
        set traffic-shaper-reverse "high-priority"
        set srcaddr "all"
        set dstaddr "all"
    next
    edit 2
        set name "FTP"
        set service "FTP" "FTP_GET" "FTP_PUT"
        set dstintf "virtual-wan-link"
        set traffic-shaper "low-priority"
        set traffic-shaper-reverse "low-priority"
        set srcaddr "all"
        set dstaddr "all"
    next
end
```

To configure SD-WAN traffic shaping and QoS with SD-WAN in the CLI:

```
config system virtual-wan-link
    set status enable
    config members
        edit 1
            set interface "wan1"
            set gateway x.x.x.x
        next
        edit 2
            set interface "wan2"
            set gateway x.x.x.x
        next
    end
    config service
        edit 1
            set name "SIP"
            set member 1
            set dst "voip-server"
            set dscp-forward enable
            set dscp-forward-tag 101110
        next
    end
end
```


Netskope Exhibit 1017

To use the diagnose command to check if specific traffic is attached to the correct traffic shaper:

```
# diagnose firewall iprope list 100015

policy index=1 uuid_idx=0 action=accept
flag (0):
shapers: orig=high-priority(2/0/134217728) reply=high-priority(2/0/134217728)
cos_fwd=0  cos_rev=0
group=00100015 av=00000000 au=00000000 split=00000000
host=0 chk_client_info=0x0 app_list=0 ips_view=0
misc=0 dd_type=0 dd_mode=0
zone(1): 0 -> zone(2): 36 38
source(1): 0.0.0.0-255.255.255.255, uuid_idx=6,
dest(1): 0.0.0.0-255.255.255.255, uuid_idx=6,
service(2):
        [6:0x0:0/(1,65535)->(80,80)] helper:auto
        [6:0x0:0/(1,65535)->(443,443)] helper:auto

policy index=2 uuid_idx=0 action=accept
flag (0):
shapers: orig=low-priority(4/128000/134217728) reply=low-priority(4/128000/134217728)
cos_fwd=0  cos_rev=0
group=00100015 av=00000000 au=00000000 split=00000000
host=0 chk_client_info=0x0 app_list=0 ips_view=0
misc=0 dd_type=0 dd_mode=0
zone(1): 0 -> zone(2): 36 38
source(1): 0.0.0.0-255.255.255.255, uuid_idx=6,
dest(1): 0.0.0.0-255.255.255.255, uuid_idx=6,
service(3):
        [6:0x0:0/(1,65535)->(21,21)] helper:auto
        [6:0x0:0/(1,65535)->(21,21)] helper:auto
        [6:0x0:0/(1,65535)->(21,21)] helper:auto

FGT_A (root) #
```

To use the diagnose command to check if the correct traffic shaper is applied to the session:

```
# diagnose sys session list
session info: proto=6 proto_state=01 duration=11 expire=3599 timeout=3600 flags=00000000
sockflag=00000000 sockport=0 av_idx=0 use=5
origin-shaper=low-priority prio=4 guarantee 128000Bps max 1280000Bps traffic 1050Bps drops
0B
reply-shaper=
per_ip_shaper=
class_id=0 shaping_policy_id=2 ha_id=0 policy_dir=0 tunnel=/ helper=ftp vlan_cos=0/255
state=may_dirty npu npd os mif route_preserve
statistic(bytes/packets/allow_err): org=868/15/1 reply=752/10/1 tuples=2
tx speed(Bps/kbps): 76/0 rx speed(Bps/kbps): 66/0
orgin->sink: org pre->post, reply pre->post dev=39->38/38->39 gwy=172.16.200.55/0.0.0.0
hook=post dir=org act=snat 10.1.100.11:58241->172.16.200.55:21(172.16.200.1:58241)
hook=pre dir=reply act=dnat 172.16.200.55:21->172.16.200.1:58241(10.1.100.11:58241)
pos/(before,after) 0/(0,0), 0/(0,0)
misc=0 policy_id=1 auth_info=0 chk_client_info=0 vd=4
serial=0003255f tos=ff/ff app_list=0 app=0 url_cat=0
rpdb_link_id = 00000000
dd_type=0 dd_mode=0
```


Netskope Exhibit 1017

```
npu_state=0x100000
npu info: flag=0x00/0x00, offload=0/0, ips_offload=0/0, epid=0/0, ipid=0/0,
vlan=0x0000/0x0000
vlifid=0/0, vtag_in=0x0000/0x0000 in_npu=0/0, out_npu=0/0, fwd_en=0/0, qid=0/0
no_ofld_reason:  offload-denied helper
total session 1
```

To use the diagnose command to check the status of a shared traffic shaper:

```
 # diagnose firewall shaper traffic-shaper list

name high-priority
maximum-bandwidth 131072 KB/sec
guaranteed-bandwidth 0 KB/sec
current-bandwidth 0 B/sec
priority 2
tos ff
packets dropped 0
bytes dropped 0

name low-priority
maximum-bandwidth 131072 KB/sec
guaranteed-bandwidth 125 KB/sec
current-bandwidth 0 B/sec
priority 4
tos ff
packets dropped 0
bytes dropped 0

name high-priority
maximum-bandwidth 131072 KB/sec
guaranteed-bandwidth 0 KB/sec
current-bandwidth 0 B/sec
priority 2
policy 1
tos ff
packets dropped 0
bytes dropped 0

name low-priority
maximum-bandwidth 131072 KB/sec
guaranteed-bandwidth 125 KB/sec
current-bandwidth 0 B/sec
priority 4
policy 2
tos ff
packets dropped 0
bytes dropped 0
```

Advanced configuration

The following topics provide instructions on SD-WAN advanced configuration:

- Self-originating traffic on page 489
- SDN dynamic connector addresses in SD-WAN rules on page 489
- Forward error correction on VPN overlay networks on page 492
- Using BGP tags with SD-WAN rules on page 495
- BGP multiple path support on page 498
- Controlling traffic with BGP route mapping and service rules on page 501
- Applying BGP route-map to multiple BGP neighbors on page 507
- IBGP and EBGP support in VRF on page 513
- ADVPN and shortcut paths on page 516
- SD-WAN monitor on ADVPN shortcuts on page 528
- SD-WAN integration with OCVPN on page 529
- DSCP matching (shaping) on page 537
- SD-WAN health check packet DSCP marker support on page 541
- Dual VPN tunnel wizard on page 541
- Internet service customization on page 542
- SD-WAN with FGCP HA on page 543

See also Per packet distribution and tunnel aggregation on page 1173.

Self-originating traffic



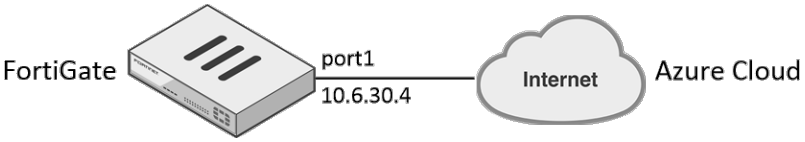
This topic applies to FortiOS 6.4.0. In other versions, self-originating (local-out) traffic behaves differently.

Self-originating traffic, such as Syslog, FortiAnalyzer logging, FortiGuard services, remote authentication, and others, relies on routing table lookups to determine the egress interface that is used to initiate the connection. Policy routes generated by SD-WAN rules do not apply to this traffic.

SDN dynamic connector addresses in SD-WAN rules

SDN dynamic connector addresses can be used in SD-WAN rules. FortiGate supports both public (AWS, Azure, GCP, OCI, AliCloud) and private (Kubernetes, VMware ESXi and NSX, OpenStack, ACI, Nuage) SDN connectors.

The configuration procedure for all of the supported SDN connector types is the same. This example uses an Azure public SDN connector.



There are four steps to create and use an SDN connector address in an SD-WAN rule:

1. Configure the FortiGate IP address and network gateway so that it can reach the Internet.
2. Create an Azure SDN connector.

3. Create a firewall address to associate with the configured SDN connector.
4. Use the firewall address in an SD-WAN service rule.

To create an Azure SDN connector:

1. Go to *Security Fabric* > *External Connectors*.
2. Click *Create New*.
3. In the *Public SDN* section, click *Microsoft Azure*.
4. Enter the following:

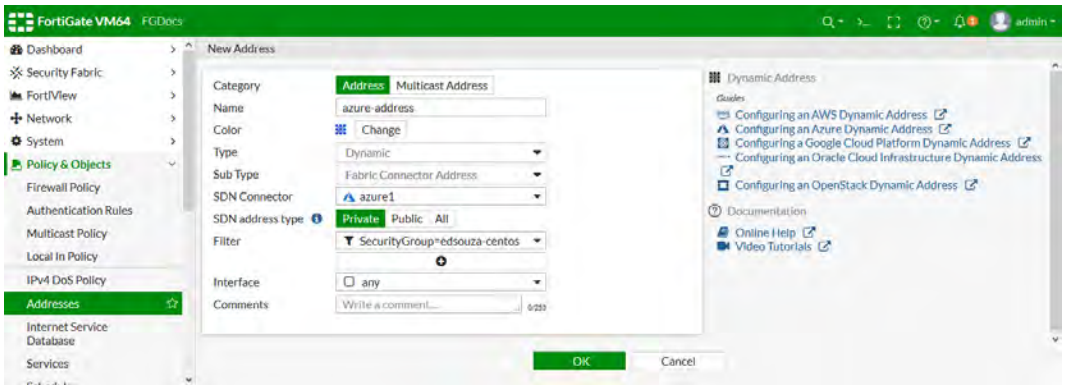
Name	azure1
Status	Enabled
Update Interval	Use Default
Server region	Global
Directory ID	942b80cd-1b14-42a1-8dcf-4b21dece61ba
Application ID	14dbd5c5-307e-4ea4-8133-68738141feb1
Client secret	xxxxxx
Resource path	disabled

5. Click *OK*.

To create a firewall address to associate with the configured SDN connector:

1. Go to *Policy & Objects* > *Addresses*.
2. Click *Create New* > *Address*.
3. Enter the following:

Category	Address
Name	azure-address
Type	Dynamic
Sub Type	Fabric Connector Address
SDN Connector	azure1
SDN address type	Private
Filter	SecurityGroup=edsouza-centos
Interface	Any



4. Click *OK*.

To use the firewall address in an SD-WAN service rule:

1. Go to *Network* > *SD-WAN Rules*.
2. Click *Create New*.
3. Set the *Name* to *Azure1*.
4. For the *Destination Address* select *azure-address*.
5. Configure the remaining settings as needed. See WAN path control on page 455 for details.
6. Click *OK*.

Diagnostics

Use the following CLI commands to check the status of and troubleshoot the connector.

To see the status of the SDN connector:

```
# diagnose sys sdn status

    SDN Connector       Type        Status          Updating        Last update
    ------------------------------------------------------------------------------------
-
    azure1              azure       connected       no              n/a
```

To debug the SDN connector to resolve the firewall address:

```
# diagnose debug application azd -1
    Debug messages will be on for 30 minutes.

    ...
    azd sdn connector azure1 start updating IP addresses
    azd checking firewall address object azure-address-1, vd 0
     IP address change, new list:
      10.18.0.4
      10.18.0.12
      ...
      ...
# diagnose sys virtual-wan-link service
```

```
Service(2): Address Mode(IPV4) flags=0x0
  TOS(0x0/0x0), Protocol(0: 1->65535), Mode(manual)
  Service role: standalone
  Member sub interface:
  Members:
    1: Seq_num(1), alive, selected
  Dst address:
        10.18.0.4 - 10.18.0.4
        10.18.0.12 - 10.18.0.12
         ... ...
         ... ...
         ... ...
```

Forward error correction on VPN overlay networks

This topic shows an SD-WAN with forward error correction (FEC) on VPN overlay networks. FEC is a technique used to control and correct errors in data transmission by sending redundant data across the VPN. It uses six parameters in IPsec phase1/phase1-interface settings:

`fec-ingress`	Enable/disable Forward Error Correction for ingress IPsec traffic (default = disable).
`fec-egress`	Enable/disable Forward Error Correction for egress IPsec traffic (default = disable).
`fec-base`	The number of base Forward Error Correction packets (1 - 100, default = 20).
`fec-redundant`	The number of redundant Forward Error Correction packets (1 - 100, default = 10).
`fec-send-timeout`	The time before sending Forward Error Correction packets, in milliseconds (1 - 1000, default = 8).
`fec-receive-timeout`	The time before dropping Forward Error Correction packets, in milliseconds (1 - 1000, default = 5000).

For every `fec-base` number of sent packets, the tunnel will send `fec-redundant` number of redundant packets.

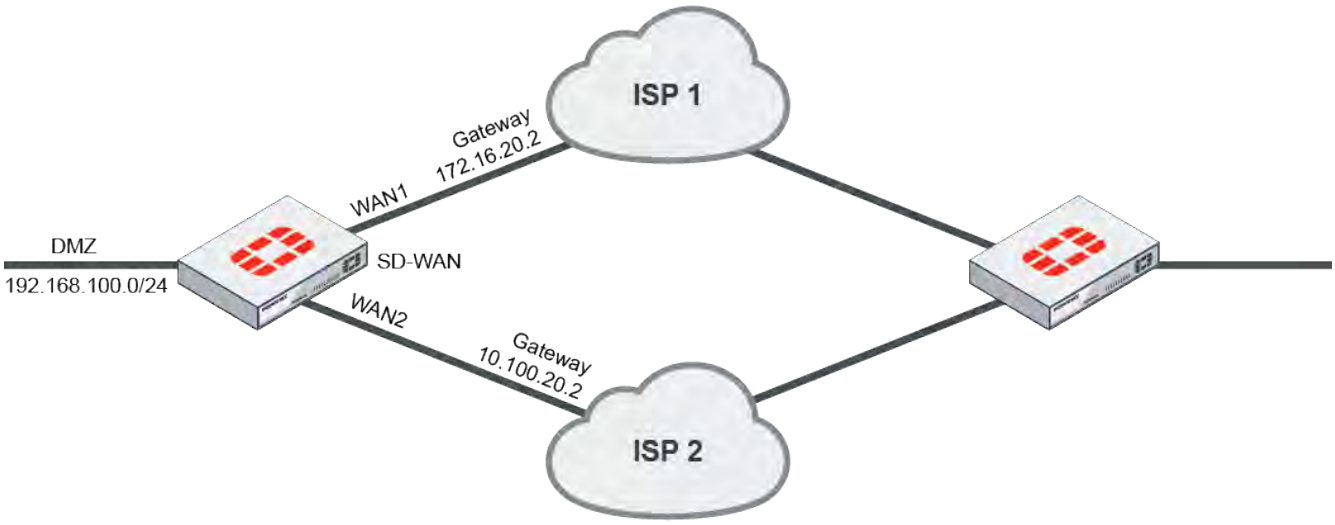


If your FortiGate is NPU capable, disable `npu-offload` in your phase1 configurations:

```
config vpn ipsec phase1-interface
    edit <name>
        set npu-offload disable
    next
end
```

Example

For example, a customer has two ISP connections, wan1 and wan2. Using these two connections, create two IPsec VPN interfaces as SD-WAN members. Configure FEC on each VPN interface to lower packet loss ratio by re-transmitting the packets using its backend algorithm.



To configure IPsec VPN:

```
config vpn ipsec phase1-interface
    edit "vd1-p1"
        set interface "wan1"
        set peertype any
        set net-device disable
        set proposal aes256-sha256
        set dhgrp 14
        set remote-gw 172.16.201.2
        set psksecret ftnt1234
        set fec-egress enable
        set fec-send-timeout 8
        set fec-base 20
        set fec-redundant 10
        set fec-ingress enable
        set fec-receive-timeout 5000
    next
    edit "vd1-p2"
        set interface "wan2"
        set peertype any
        set net-device disable
        set proposal aes256-sha256
        set dhgrp 14
        set remote-gw 172.16.202.2
        set psksecret ftnt1234
        set fec-egress enable
        set fec-send-timeout 8
        set fec-base 20
        set fec-redundant 10
        set fec-ingress enable
        set fec-receive-timeout 5000
    next
end
config vpn ipsec phase2-interface
    edit "vd1-p1"
        set phase1name "vd1-p1"
    next
```

```
    edit "vd1-p2"
        set phase1name "vd1-p2"
    next
end
```

To configure the interface:

```
config system interface
    edit "vd1-p1"
        set ip 172.16.211.1 255.255.255.255
        set remote-ip 172.16.211.2 255.255.255.255
    next
    edit "vd1-p2"
        set ip 172.16.212.1 255.255.255.255
        set remote-ip 172.16.212.2 255.255.255.255
    next
end
```

To configure the firewall policy:

```
config firewall policy
    edit 1
        set name "1"
        set srcintf "dmz"
        set dstintf ""virtual-wan-link""
        set srcaddr "all"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ALL"
        set nat enable
    next
end
```

To configure SD-WAN:

```
config system virtual-wan-link
    set status enable
    config members
        edit 1
            set interface "vd1-p1"
            set gateway 172.16.211.2
        next
        edit 1
            set interface "vd2-p2"
            set gateway 172.16.212.2
        next
    end
end
```

To use the diagnose command to check VPN FEC status:

```
# diagnose vpn tunnel list
list all ipsec tunnel in vd 0
-----------------------------------------------------
```

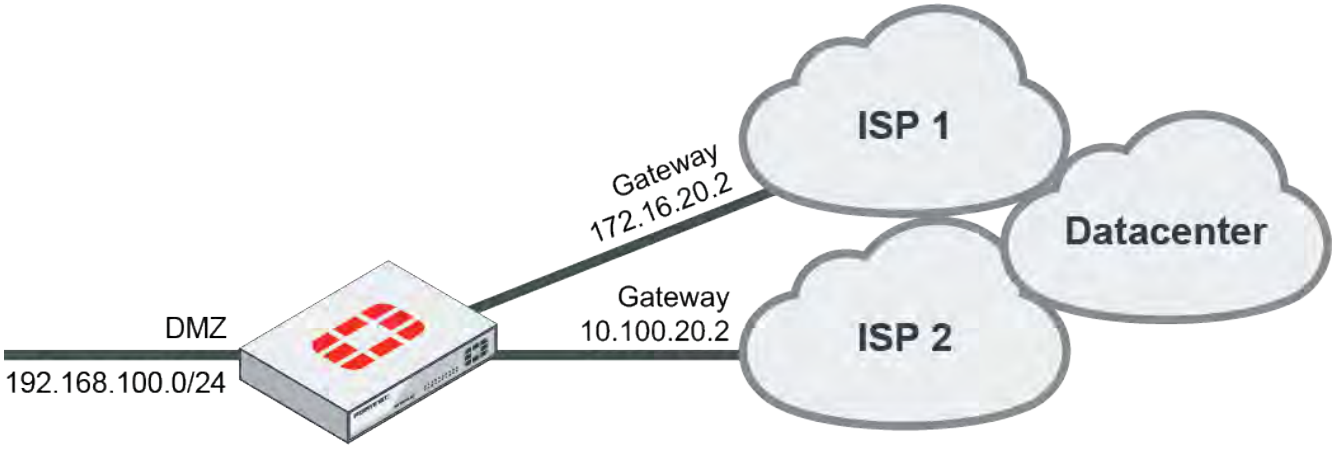

Netskope Exhibit 1017

```
name=vd1 ver=1 serial=1 172.16.200.1:0->172.16.200.2:0
bound_if=11 lgwy=static/1 tun=intf/0 mode=auto/1 encap=none/3600 options[0e10]=create_dev
frag-rfc fec-egress fec-ingress  accept_traffic=1

proxyid_num=1 child_num=0 refcnt=11 ilast=8 olast=8 ad=/0
stat: rxp=0 txp=0 rxb=0 txb=0
dpd: mode=on-demand on=1 idle=20000ms retry=3 count=0 seqno=0
natt: mode=none draft=0 interval=0 remote_port=0
fec-egress: base=20 redundant=10 remote_port=50000     <<<<<<<<<<<<<<<<<<<<<<
fec-ingress: base=20 redundant=10                      <<<<<<<<<<<<<<<<<<<<<<
proxyid=demo proto=0 sa=1 ref=2 serial=1
  src: 0:10.1.100.0/255.255.255.0:0
  dst: 0:173.1.1.0/255.255.255.0:0
  SA:  ref=3 options=10226 type=00 soft=0 mtu=1390 expire=42897/0B replaywin=2048
       seqno=1 esn=0 replaywin_lastseq=00000000 itn=0 qat=0
  life: type=01 bytes=0/0 timeout=42899/43200
  dec: spi=181f4f81 esp=aes key=16 6e8fedf2a77691ffdbf3270484cb2555
       ah=sha1 key=20 f92bcf841239d15d30b36b695f78eaef3fad05c4
  enc: spi=0ce10190 esp=aes key=16 2d684fb19cbae533249c8b5683937329
       ah=sha1 key=20 ba7333f89cd34cf75966bd9ffa72030115919213
  dec:pkts/bytes=0/0, enc:pkts/bytes=0/0
```

Using BGP tags with SD-WAN rules

SD-WAN rules can use Border Gateway Protocol (BGP) learned routes as dynamic destinations.



In this example, a customer has two ISP connections, wan1 and wan2. wan1 is used primarily for direct access to internet applications, and wan2 is used primarily for traffic to the customer's data center.

The customer could create an SD-WAN rule using the data center's IP address range as the destination to force that traffic to use wan2, but the data center's IP range is not static. Instead, a BGP tag can be used.

For this example, wan2's BGP neighbor advertises the data center's network range with a community number of 30:5.

This example assumes that SD-WAN is enable on the FortiGate, wan1 and wan2 are added as SD-WAN members, and a policy and static route have been created. See SD-WAN quick start on page 444 for details.

To configure BGP tags with SD-WAN rules:

1. Configure the community list:

```
config router community-list
    edit "30:5"
        config rule
            edit 1
                set action permit
                set match "30:5"
            next
        end
    next
end
```

2. Configure the route map:

```
config router route-map
    edit "comm1"
        config rule
            edit 1
                set match-community "30:5"
                set set-route-tag 15
            next
        end
    next
end
```

3. Configure BGP:

```
config router bgp
    set as xxxxx
    set router-id xxxx
    config neighbor
        edit "10.100.20.2"
            set soft-reconfiguration enable
            set remote-as xxxxx
            set route-map-in "comm1"
        next
    end
end
```

4. Configure a firewall policy:

```
config firewall policy
    edit 1
        set name "1"
        set srcintf "dmz"
        set dstintf ""virtual-wan-link""
        set srcaddr "all"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ALL"
        set nat enable
    next
end
```


Netskope Exhibit 1017

5. Edit the SD-WAN configuration:

```
config system virtual-wan-link
    set status enable
    config members
        edit 1
            set interface "wan1"
            set gateway 172.16.20.2
        next
        edit 2
            set interface "wan2"
        next
    end
    config service
        edit 1
            set name "DataCenter"
            set mode manual
            set route-tag 15
            set members 2
        next
    end
end
```

Troubleshooting BGP tags with SD-WAN rules

Check the network community

Use the `get router info bgp network` command to check the network community:

```
# get router info bgp network
BGP table version is 5, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network Next Hop Metric LocPrf Weight RouteTag Path
*> 0.0.0.0/0 10.100.1.5 32768 0 ?
*> 1.1.1.1/32 0.0.0.0 32768 0 ?
*> 10.1.100.0/24 172.16.203.2 32768 0 ?
*> 10.100.1.0/30 0.0.0.0 32768 0 ?
*> 10.100.1.4/30 0.0.0.0 32768 0 ?
*> 10.100.1.248/29 0.0.0.0 32768 0 ?
*> 10.100.10.0/24 10.100.1.5 202 10000 15 20 e
*> 172.16.200.0/24 0.0.0.0 32768 0 ?
*> 172.16.200.200/32
                    0.0.0.0 32768 0 ?
*> 172.16.201.0/24 172.16.200.4 32768 0 ?
*> 172.16.203.0/24 0.0.0.0 32768 0 ?
*> 172.16.204.0/24 172.16.200.4 32768 0 ?
*> 172.16.205.0/24 0.0.0.0 32768 0 ?
*> 172.16.206.0/24 0.0.0.0 32768 0 ?
*> 172.16.207.1/32 0.0.0.0 32768 0 ?
*> 172.16.207.2/32 0.0.0.0 32768 0 ?
*> 172.16.212.1/32 0.0.0.0 32768 0 ?
*> 172.16.212.2/32 0.0.0.0 32768 0 ?
*> 172.17.200.200/32
```


Netskope Exhibit 1017

```
                    0.0.0.0 32768 0 ?
*> 172.27.1.0/24 0.0.0.0 32768 0 ?
*> 172.27.2.0/24 0.0.0.0 32768 0 ?
*> 172.27.5.0/24 0.0.0.0 32768 0 ?
*> 172.27.6.0/24 0.0.0.0 32768 0 ?
*> 172.27.7.0/24 0.0.0.0 32768 0 ?
*> 172.27.8.0/24 0.0.0.0 32768 0 ?
*> 172.29.1.0/24 0.0.0.0 32768 0 ?
*> 172.29.2.0/24 0.0.0.0 32768 0 ?
*> 192.168.1.0 0.0.0.0 32768 0 ?

Total number of prefixes 28

# get router info bgp network 10.100.11.0
BGP routing table entry for 10.100.10.0/24
Paths: (2 available, best 1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
   172.10.22.2
  20
    10.100.20.2 from 10.100.20.2 (6.6.6.6)
      Origin EGP metric 200, localpref 100, weight 10000, valid, external, best
      Community: 30:5 <<<<==========================
      Last update: Wen Mar 20 18:45:17 2019
```

Check dynamic BGP addresses

Use the `get router info route-map-address` command to check dynamic BGP addresses:

```
# get router info route-map-address
Extend-tag: 15, interface(wan2:16)
        10.100.11.0/255.255.255.0
```

Check dynamic BGP addresses used in policy routes

Use the `diagnose firewall proute list` command to check dynamic BGP addresses used in policy routes:

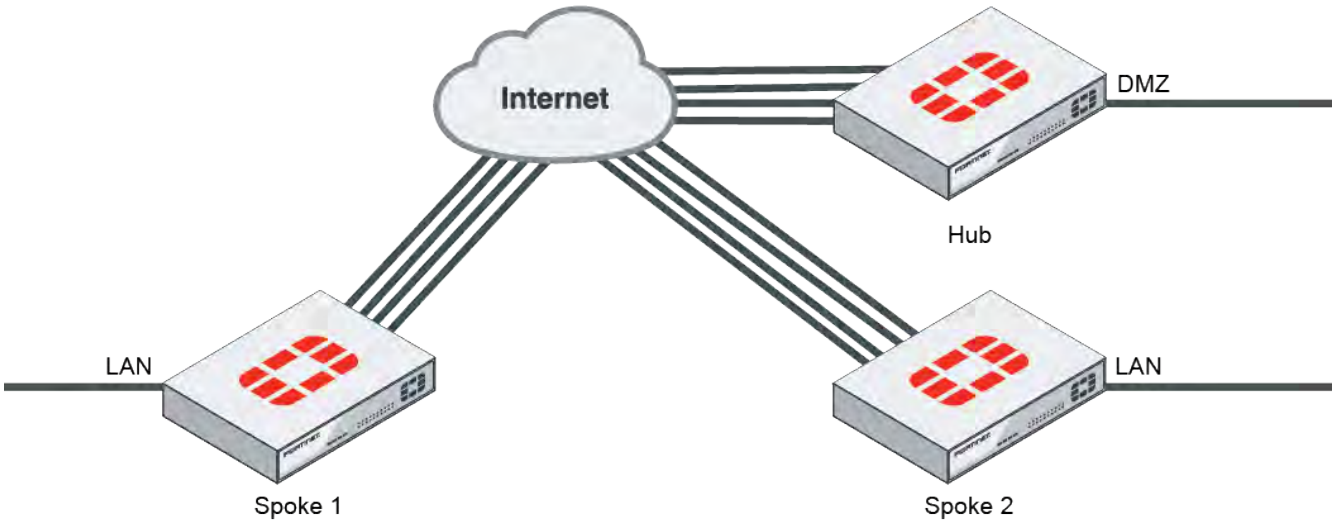
```
# diagnose firewall proute list
list route policy info(vf=root):

id=4278779905 vwl_service=1(DataCenter) flags=0x0 tos=0x00 tos_mask=0x00 protocol=0
sport=0:65535 iif=0 dport=1-65535 oif=16
source wildcard(1): 0.0.0.0/0.0.0.0
destination wildcard(1): 10.100.11.0/255.255.255.0
```

BGP multiple path support

BGP supports multiple paths, allowing an ADVPN to advertise multiple paths. This allows BGP to extend and keep additional network paths according to RFC 7911.

In this example, Spoke1 and Spoke2 each have four VPN tunnels that are connected to the Hub with ADVPN. The Spoke-Hub has established four BGP neighbors on all four tunnels.



Spoke 1 and Spoke 2 can learn four different routes from each other.

To configure the hub:

```
config router bgp
    set as 65505
    set router-id 11.11.11.11
    set ibgp-multipath enable
    set additional-path enable
    set additional-path-select 4
    config neighbor-group
        edit "gr1"
            set capability-default-originate enable
            set remote-as 65505
            set additional-path both
            set adv-additional-path 4
            set route-reflector-client enable
        next
    end
    config neighbor-range
        edit 1
            set prefix 10.10.0.0 255.255.0.0
            set neighbor-group "gr1"
        next
    end
    config network
        edit 12
            set prefix 11.11.11.11 255.255.255.255
        next
    end
end
```

To configure a spoke:

```
config router bgp
    set as 65505
    set router-id 2.2.2.2
    set ibgp-multipath enable
```

```
        set additional-path enable
        set additional-path-select 4
        config neighbor
            edit "10.10.100.254"
                set soft-reconfiguration enable
                set remote-as 65505
                set additional-path both
                set adv-additional-path 4
            next
            edit "10.10.200.254"
                set soft-reconfiguration enable
                set remote-as 65505
                set additional-path both
                set adv-additional-path 4
            next
            edit "10.10.203.254"
                set soft-reconfiguration enable
                set remote-as 65505
                set additional-path both
                set adv-additional-path 4
            next
            edit "10.10.204.254"
                set soft-reconfiguration enable
                set remote-as 65505
                set additional-path both
                set adv-additional-path 4
            next
        end
        config network
            edit 3
                set prefix 22.1.1.0 255.255.255.0
            next
        end
    end
```

To view the BGP routing table on a spoke:

```
Spoke1 # get router info routing-table bgp
    Routing table for VRF=0
    B*      0.0.0.0/0 [200/0] via 10.10.200.254, vd2-2, 03:57:26
            [200/0] via 10.10.203.254, vd2-3, 03:57:26
            [200/0] via 10.10.204.254, vd2-4, 03:57:26
            [200/0] via 10.10.100.254, vd2-1, 03:57:26
    B       1.1.1.1/32 [200/0] via 11.1.1.1 (recursive via 12.1.1.1), 03:57:51
            [200/0] via 11.1.1.1 (recursive via 12.1.1.1), 03:57:51
            [200/0] via 11.1.1.1 (recursive via 12.1.1.1), 03:57:51
            [200/0] via 11.1.1.1 (recursive via 12.1.1.1), 03:57:51
    B       11.11.11.11/32 [200/0] via 10.10.200.254, vd2-2, 03:57:51
            [200/0] via 10.10.203.254, vd2-3, 03:57:51
            [200/0] via 10.10.204.254, vd2-4, 03:57:51
            [200/0] via 10.10.100.254, vd2-1, 03:57:51
    B       33.1.1.0/24 [200/0] via 10.10.204.3, vd2-4, 03:57:26
            [200/0] via 10.10.203.3, vd2-3, 03:57:26
            [200/0] via 10.10.200.3, vd2-2, 03:57:26
            [200/0] via 10.10.100.3, vd2-1, 03:57:26
            [200/0] via 10.10.204.3, vd2-4, 03:57:26
```


Netskope Exhibit 1017

```
[200/0] via 10.10.203.3, vd2-3, 03:57:26
[200/0] via 10.10.200.3, vd2-2, 03:57:26
[200/0] via 10.10.100.3, vd2-1, 03:57:26
[200/0] via 10.10.204.3, vd2-4, 03:57:26
[200/0] via 10.10.203.3, vd2-3, 03:57:26
[200/0] via 10.10.200.3, vd2-2, 03:57:26
[200/0] via 10.10.100.3, vd2-1, 03:57:26
[200/0] via 10.10.204.3, vd2-4, 03:57:26
[200/0] via 10.10.203.3, vd2-3, 03:57:26
[200/0] via 10.10.200.3, vd2-2, 03:57:26
[200/0] via 10.10.100.3, vd2-1, 03:57:26
```

Controlling traffic with BGP route mapping and service rules

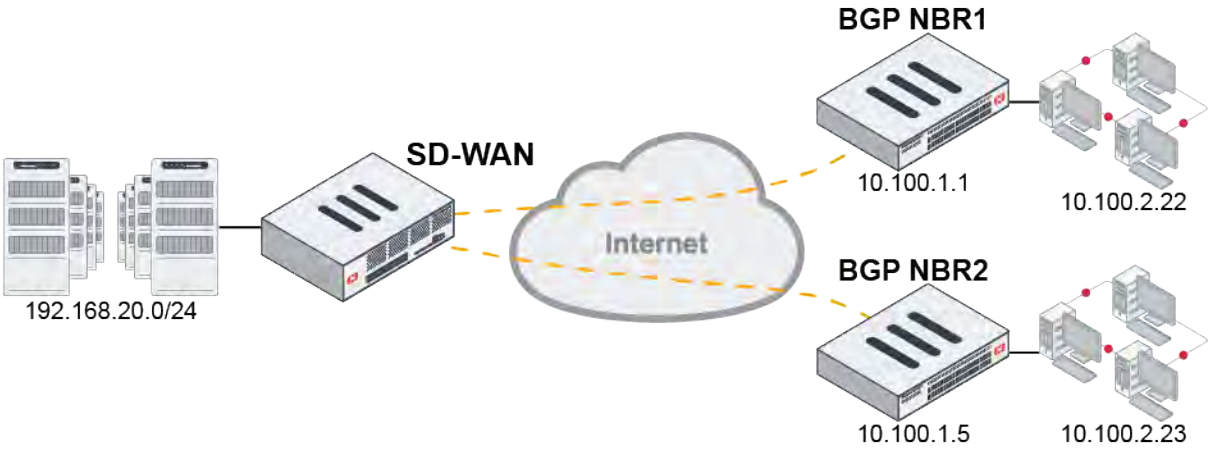
SD-WAN allows you to select different outbound WAN links based on performance SLAs. It is important that BGP neighbors are aware of these settings, and changes to them.

BGP can adapt to changes in SD-WAN link SLAs in the following ways:

- Applying different route-maps based on the SD-WAN's health checks. For example, different BGP community strings can be advertised to BGP neighbors when SLAs are not met.
- Traffic can be selectively forwarded based on the active BGP neighbor. If the SD-WAN service's role matches the active SD-WAN neighbor, the service is enabled. If there is no match, then the service is disabled.

Example

In this topology, a branch FortiGate has two SD-WAN gateways serving as the primary and secondary gateways. The gateways reside in different datacenters, but have a full mesh network between them.



This example shows how route-maps and service rules are selected based on performance SLAs and the member that is currently active. Traffic flows through the primary gateway unless the neighbor's health check is outside of its SLA. If that happens, traffic routes to the secondary gateway.

BGP NBR1 is the primary neighbor and BGP NBR2 is the secondary neighbor.

The branch FortiGate's wan1 and wan2 interfaces are members of the SD-WAN. When the SD-WAN neighbor status is primary, it will advertise community *20:1* to BGP NBR1 and *20:5* to BGP NBR2. When the SD-WAN neighbor status is secondary, it will advertise *20:5* to BGP NBR1 and *20:2* to BGP NBR2.

Only one of the primary or secondary neighbors can be active at one time. The SD-WAN neighbor status is used to decide which neighbor is selected:

- **Primary**: The primary neighbor takes precedence if its SLAs are met.
- **Secondary**: If the primary neighbor's SLAs are not met, the secondary neighbor becomes active if its SLAs are met.
- **Standalone**: If neither the primary or secondary neighbor's SLAs are met, the SD-WAN neighbor status becomes standalone.

Route map

SD-WAN is configured to let BGP advertise different communities when the SLA status changes. When the SLA is missed, it triggers BGP to advertise a different community to its BGP neighbor based on its route-map. The BGP neighbors can use the received community string to select the best path to reach the branch.

To configure BGP route-maps and neighbors:

1. Configure an access for the routes to be matched:

```
config router access-list
    edit "net192"
        config rule
            edit 1
                set prefix 192.168.20.0 255.255.255.0
            next
        end
    next
end
```

2. Configure the primary neighbor's preferred route-map:

```
config router route-map
    edit "comm1"
        config rule
            edit 1
                set match-ip-address "net192"
                set set-community "20:1"
            next
        end
    next
end
```

3. Configure the secondary neighbor's preferred route-map:

```
config router route-map
    edit "comm2"
        config rule
            edit 1
                set match-ip-address "net192"
                set set-community "20:2"
            next
        end
    next
end
```


Netskope Exhibit 1017

4. Configure the failed route-map:

```
config router route-map
    edit "comm5"
        config rule
            edit 1
                set match-ip-address "net192"
                set set-community "20:5"
            next
        end
    next
end
```

5. Configure BGP neighbors:

```
config router bgp
    set as 65412
    set router-id 1.1.1.1
    set ibgp-multipath enable
    config neighbor
        edit "10.100.1.1"
            set soft-reconfiguration enable
            set remote-as 20
            set route-map-out "comm5"
            set route-map-out-preferable "comm1"
        next
        edit "10.100.1.5"
            set soft-reconfiguration enable
            set remote-as 20
            set route-map-out "comm5"
            set route-map-out-preferable "comm2"
        next
    end
end
```

When SLAs are met, `route-map-out-preferable` is used. When SLAs are missed, `route-map-out` is used.

To configure SD-WAN:

1. Configure the SD-WAN members:

```
config system virtual-wan-link
    set status enable
    config members
        edit 1
            set interface "port1"
        next
        edit 2
            set interface "port2"
        next
    end
end
```

2. Configure health checks for each member:

```
config system virtual-wan-link
    config health-check
        edit "ping"
```


Netskope Exhibit 1017

```
            set server "10.100.2.22"
            set members 1
            config sla
                edit 1
                    set link-cost-factor packet-loss
                    set packetloss-threshold 1
                next
            end
        next
        edit "ping2"
            set server "10.100.2.23"
            set members 2
            config sla
                edit 1
                    set link-cost-factor packet-loss
                    set packetloss-threshold 1
                next
            end
        next
    end
end
```

3. Configure the SD-WAN neighbors and assign them a role and the health checks used to determine if the neighbor meets the SLA:
   SD-WAN neighbors can only be configured in the CLI.

```
config system virtual-wan-link
    config neighbor
        edit "10.100.1.1"
            set member 1
            set role primary
            set health-check "ping"
            set sla-id 1
        next
        edit "10.100.1.5"
            set member 2
            set role secondary
            set health-check "ping2"
            set sla-id 1
        next
    end
end
```

Service rules

Create SD-WAN service rules to direct traffic to the primary neighbor when its SLAs are met, and to the secondary neighbor when the primary neighbor's SLAs are missed.

To configure the SD-WAN service rules:

```
config system virtual-wan-link
    config service
        edit 1
            set name "Primary-Out"
            set role primary
            set dst "all"
```


Netskope Exhibit 1017

```
            set src "all"
            set priority-members 1
        next
        edit 2
            set name "Secondary-Out"
            set role secondary
            set dst "all"
            set src "all"
            set priority-members 2
        next
    end
end
```



If neither the primary nor secondary neighbors are active, the SD-WAN neighbor status becomes standalone. Only service rules with `standalone-action` enabled will continue to pass traffic. This option is disabled by default.

Verification

To verify when the primary neighbor is passing traffic:

1. Verify the health check status:

```
FortiGate-Branch # diagnose sys virtual-wan-link health-check
Health Check(ping):
Seq(1 port1): state(alive), packet-loss(0.000%) latency(0.569), jitter(0.061) sla_
map=0x1
Health Check(ping2):
Seq(2 port2): state(alive), packet-loss(0.000%) latency(3.916), jitter(2.373) sla_
map=0x1
```

2. Verify SD-WAN neighbor status:

```
FortiGate-Branch # diagnose sys virtual-wan-link neighbor
SD-WAN neighbor status: hold-down(disable), hold-down-time(0), hold_boot_time(0)
        Selected role(primary) last_secondary_select_time/current_time in seconds 0/572
Neighbor(10.100.1.1): member(1) role(primary)
        Health-check(ping:1)  sla-pass selected alive
Neighbor(10.100.1.5): member(2) role(secondary)
        Health-check(ping2:1)  sla-pass alive
```

3. Verify service rules status:

```
FortiGate-Branch # diagnose sys virtual-wan-link service

Service(1): Address Mode(IPV4) flags=0x0
  Gen(3), TOS(0x0/0x0), Protocol(0: 1->65535), Mode(manual)
  Service role: primary
  Members:
    1: Seq_num(1 port1), alive, selected
  Src address:
        0.0.0.0-255.255.255.255

  Dst address:
        0.0.0.0-255.255.255.255
```


Netskope Exhibit 1017

```
Service(2): Address Mode(IPV4) flags=0x0
  Gen(6), TOS(0x0/0x0), Protocol(0: 1->65535), Mode(manual)
  Service role: secondary, disabled by unselected.
  Members:
    1: Seq_num(2 port2), alive, selected
  Src address:
        0.0.0.0-255.255.255.255

  Dst address:
        0.0.0.0-255.255.255.255
```

4. Verify neighbor routers:
   a. Primary neighbor router:

```
FGT-NBR1 # get router info bgp network 192.168.20.0
BGP routing table entry for 192.168.20.0/24
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer
  64512
    10.100.1.2 from 10.100.1.2 (192.168.122.98)
      Origin IGP metric 0, localpref 100, valid, external, best
      Community: 20:1
      Last update: Thu Apr 30 13:41:40 2020
```

   b. Secondary neighbor router:

```
FGT-NBR2 # get router info bgp network 192.168.20.0
VRF 0 BGP routing table entry for 192.168.20.0/24
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer
  Original VRF 0
  64512
    10.100.1.6 from 10.100.1.6 (192.168.122.98)
      Origin IGP metric 0, localpref 100, valid, external, best
      Community: 20:5
      Last update: Thu Apr 30 13:41:39 2020
```

To verify when the secondary neighbor is passing traffic:

1. Verify the health check status:

```
FortiGate-Branch # diagnose sys virtual-wan-link health-check
Health Check(ping):
Seq(1 port1): state(dead), packet-loss(54.000%) sla_map=0x0
Health Check(ping2):
Seq(2 port2): state(alive), packet-loss(0.000%) latency(4.339), jitter(3.701) sla_
map=0x1
```

2. Verify SD-WAN neighbor status:

```
FortiGate-Branch # diagnose sys virtual-wan-link neighbor
SD-WAN neighbor status: hold-down(disable), hold-down-time(0), hold_boot_time(0)
        Selected role(secondary) last_secondary_select_time/current_time in seconds
936/936
Neighbor(10.100.1.1): member(1) role(primary)
        Health-check(ping:1)  sla-fail dead
Neighbor(10.100.1.5): member(2) role(secondary)
        Health-check(ping2:1)  sla-pass selected alive
```


Netskope Exhibit 1017

3. Verify service rules status:

```
FortiGate-Branch # diagnose sys virtual-wan-link service

Service(1): Address Mode(IPV4) flags=0x0
  Gen(4), TOS(0x0/0x0), Protocol(0: 1->65535), Mode(manual)
  Service role: primary, disabled by unselected.
  Members:
    1: Seq_num(1 port1), alive, selected
  Src address:
        0.0.0.0-255.255.255.255

  Dst address:
        0.0.0.0-255.255.255.255

Service(2): Address Mode(IPV4) flags=0x0
  Gen(7), TOS(0x0/0x0), Protocol(0: 1->65535), Mode(manual)
  Service role: secondary
  Members:
    1: Seq_num(2 port2), alive, selected
  Src address:
        0.0.0.0-255.255.255.255

  Dst address:
        0.0.0.0-255.255.255.255
```

4. Verify neighbor routers:
   a. Primary neighbor router:

```
FGT-NBR1 # get router info bgp network 192.168.20.0
BGP routing table entry for 192.168.20.0/24
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer
  64512
    10.100.1.2 from 10.100.1.2 (192.168.122.98)
      Origin IGP metric 0, localpref 100, valid, external, best
      Community: 20:5
      Last update: Thu Apr 30 15:41:58 2020
```

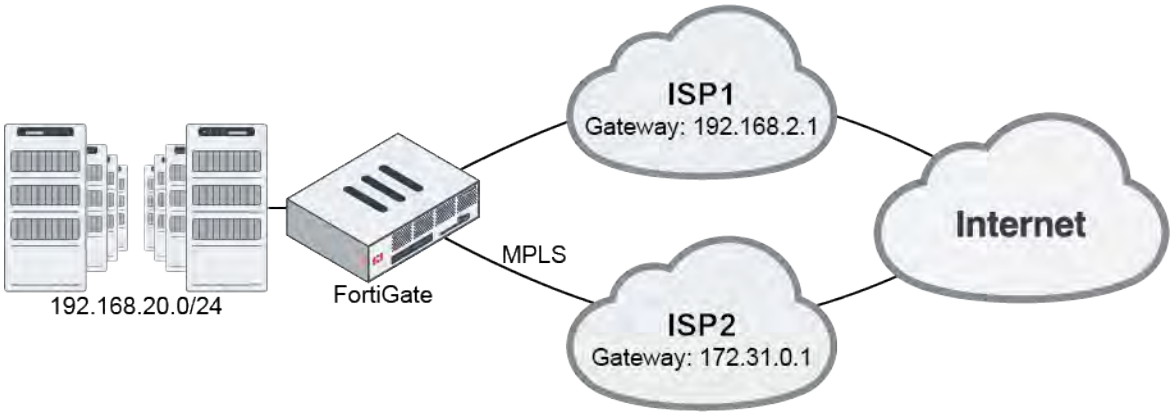
   b. Secondary neighbor router:

```
FGT-NBR2 # get router info bgp network 192.168.20.0
VRF 0 BGP routing table entry for 192.168.20.0/24
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer
  Original VRF 0
  64512
    10.100.1.6 from 10.100.1.6 (192.168.122.98)
      Origin IGP metric 0, localpref 100, valid, external, best
      Community: 20:2
      Last update: Thu Apr 30 15:42:07 2020
```

Applying BGP route-map to multiple BGP neighbors

Controlling traffic with BGP route mapping and service rules explained how BGP can apply different route-maps to the primary and secondary SD-WAN neighbors based on SLA health checks.

In this example, SD-WAN neighbors that are not bound to primary and secondary roles are configured.



The FortiGate has multiple SD-WAN links and has formed BGP neighbors with both ISPs.

ISP1 is used primarily for outbound traffic, and has an SD-WAN service rule using the lowest cost algorithm applied to it. When SLAs for ISP1 are not met, it will fail over to the MPLS line.

Inbound traffic is allowed by both WAN links, with each WAN advertising a community string when SLAs are met. When SLAs are not met, the WAN links advertise a different community string.

This example uses two SD-WAN links. The topology can be expanded to include more links as needed.

To configure BGP route-maps and neighbors:

1. Configure an access list for routes to be matched:

```
config router access-list
    edit "net192"
        config rule
            edit 1
                set prefix 192.168.20.0 255.255.255.0
            next
        end
    next
end
```

2. Configure route-maps for neighbor ISP1:

```
config router route-map
    edit "comm1"
        config rule
            edit 1
                set match-ip-address "net192"
                set set-community "64511:1"
            next
        end
    next
    edit "comm-fail1"
        config rule
            edit 1
                set match-ip-address "net192"
                set set-community "64511:5"
            next
```


Netskope Exhibit 1017

```
        end
    next
end
```

3. Configure route-maps for neighbor ISP2:

```
config router route-map
    edit "comm2"
        config rule
            edit 1
                set match-ip-address "net192"
                set set-community "64522:1"
            next
        end
    next
    edit "comm-fail2"
        config rule
            edit 1
                set match-ip-address "net192"
                set set-community "64522:5"
            next
        end
    next
end
```

4. Configure the BGP neighbors:

```
config router bgp
    set as 64512
    set keepalive-timer 1
    set holdtime-timer 3
    config neighbor
        edit "192.168.2.1"
            set soft-reconfiguration enable
            set remote-as 64511
            set route-map-out "comm-fail1"
            set route-map-out-preferable "comm1"
        next
        edit "172.31.0.1"
            set soft-reconfiguration enable
            set remote-as 64522
            set route-map-out "comm-fail2"
            set route-map-out-preferable "comm2"
        next
    end
    config network
        edit 1
            set prefix 192.168.20.0 255.255.255.0
        next
    end
end
```

When SLAs are met, `route-map-out-preferable` is used. When SLAs are missed, `route-map-out` is used.

Netskope Exhibit 1017

To configure SD-WAN:

1. Configure the SD-WAN members:

```
config system virtual-wan-link
    set status enable
    config members
        edit 1
            set interface "port1"
                  set gateway 192.168.2.1
        next
        edit 2
            set interface "MPLS"
            set cost 20
        next
    end
end
```

2. Configure the health checks that must be met:

```
config system virtual-wan-link
    config health-check
        edit "pingserver"
            set server "8.8.8.8"
            set members 2 1
            config sla
                edit 1
                    set link-cost-factor packet-loss
                    set packetloss-threshold 2
                next
            end
        next
    end
end
```

3. Configure the SD-WAN neighbors and assign them a role and the health checks used to determine if the neighbor meets the SLA:
   When no role is defined, the default role, `standalone`, is used.

```
config system virtual-wan-link
    config neighbor
        edit "192.168.2.1"
            set member 1
            set health-check "pingserver"
            set sla-id 1
        next
        edit "172.31.0.1"
            set member 2
            set health-check "pingserver"
            set sla-id 1
        next
    end
end
```


Netskope Exhibit 1017

Service rules

Create SD-WAN service rules to direct traffic to the SD-WAN links based on the lowest cost algorithm The same SLA health check and criteria that are used for the SD-WAN neighbor are used for this SD-WAN service rule.

When no roles are defined in the service rule, the default role, `standalone`, is used.

To configure the SD-WAN service rule:

```
config system virtual-wan-link
    config service
        edit 1
            set name "OutboundAll"
            set mode sla
            set dst "all"
            set src "all"
            config sla
                edit "pingserver"
                    set id 1
                next
            end
            set priority-members 1 2
        next
    end
end
```

Verification

To verify that when both SLAs are met, port1 is selected due to its lower cost:

1. Verify the health check status:

```
FortiGate-Branch # diagnose sys virtual-wan-link health-check
Health Check(pingserver):
Seq(2 MPLS): state(alive), packet-loss(0.000%) latency(24.709), jitter(14.996) sla_
map=0x1
Seq(1 port1): state(alive), packet-loss(0.000%) latency(28.771), jitter(14.840) sla_
map=0x1
```

2. Verify SD-WAN neighbor status:

```
FortiGate-Branch # diagnose sys virtual-wan-link neighbor
Neighbor(192.168.2.1): member(1) role(standalone)
        Health-check(pingserver:1)  sla-pass selected alive
Neighbor(172.31.0.1): member(2) role(standalone)
        Health-check(pingserver:1)  sla-pass selected alive
```

3. Verify service rules status:
   Because the service role is `standalone`, it matches both neighbors. The mode (SLA) determines that port1 is lower cost.

```
FortiGate-Branch # diagnose sys virtual-wan-link service

Service(1): Address Mode(IPV4) flags=0x0
  Gen(1), TOS(0x0/0x0), Protocol(0: 1->65535), Mode(sla), sla-compare-order
  Service role: standalone
  Members:
```


Netskope Exhibit 1017

```
    1: Seq_num(1 port1), alive, sla(0x1), cfg_order(0), cost(0), selected
    2: Seq_num(2 MPLS), alive, sla(0x1), cfg_order(1), cost(20), selected
  Src address:
        0.0.0.0-255.255.255.255

  Dst address:
        0.0.0.0-255.255.255.255
```

4. Verify neighbor routers:
   a. Primary neighbor router:

```
FGT-NBR1 # get router info bgp network 192.168.20.0
BGP routing table entry for 192.168.20.0/24
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer
  64512
    192.168.2.5 from 192.168.2.5 (192.168.122.98)
      Origin IGP metric 0, localpref 100, valid, external, best
      Community: 64511:1
      Last update: Thu Apr 30 23:59:05 2020
```

   b. Secondary neighbor router:

```
FGT-NBR2 # get router info bgp network 192.168.20.0
VRF 0 BGP routing table entry for 192.168.20.0/24
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer
  Original VRF 0
  64512
    172.31.0.2 from 172.31.0.2 (192.168.122.98)
      Origin IGP metric 0, localpref 100, valid, external, best
      Community: 64522:1
      Last update: Fri May  1 00:11:28 2020
```

To verify that when neighbor ISP1 misses SLAs, MPLS is selected and BGP advertises a different community string for ISP1:

1. Verify the health check status:

```
FortiGate-Branch # diagnose sys virtual-wan-link health-check
Health Check(pingserver):
Seq(2 MPLS): state(alive), packet-loss(0.000%) latency(25.637), jitter(17.820) sla_
map=0x1
Seq(1 port1): state(dead), packet-loss(16.000%) sla_map=0x0
```

2. Verify SD-WAN neighbor status:

```
FortiGate-Branch # diagnose sys virtual-wan-link neighbor
Neighbor(192.168.2.1): member(1) role(standalone)
        Health-check(pingserver:1)  sla-fail dead
Neighbor(172.31.0.1): member(2) role(standalone)
        Health-check(pingserver:1)  sla-pass selected alive
```

3. Verify service rules status:
   As SLA failed for neighbor ISP1, MPLS is preferred.

```
FortiGate-Branch # diagnose sys virtual-wan-link service
```


Netskope Exhibit 1017

```
Service(1): Address Mode(IPV4) flags=0x0
  Gen(3), TOS(0x0/0x0), Protocol(0: 1->65535), Mode(sla), sla-compare-order
  Service role: standalone
  Members:
    1: Seq_num(2 MPLS), alive, sla(0x1), cfg_order(1), cost(20), selected
    2: Seq_num(1 port1), dead, sla(0x0), cfg_order(0), cost(0)
  Src address:
        0.0.0.0-255.255.255.255

  Dst address:
        0.0.0.0-255.255.255.255
```

4. Verify neighbor routers:
   The community received on ISP1 is updated.
   a. Primary neighbor router:

```
FGT-NBR1 # get router info bgp network 192.168.20.0
BGP routing table entry for 192.168.20.0/24
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer
  64512
    192.168.2.5 from 192.168.2.5 (192.168.122.98)
      Origin IGP metric 0, localpref 100, valid, external, best
      Community: 64511:5
      Last update: Fri May  1 00:33:26 2020
```

   b. Secondary neighbor router:

```
FGT-NBR2 # get router info bgp network 192.168.20.0
VRF 0 BGP routing table entry for 192.168.20.0/24
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer
  Original VRF 0
  64512
    172.31.0.2 from 172.31.0.2 (192.168.122.98)
      Origin IGP metric 0, localpref 100, valid, external, best
      Community: 64522:1
      Last update: Fri May  1 00:22:42 2020
```

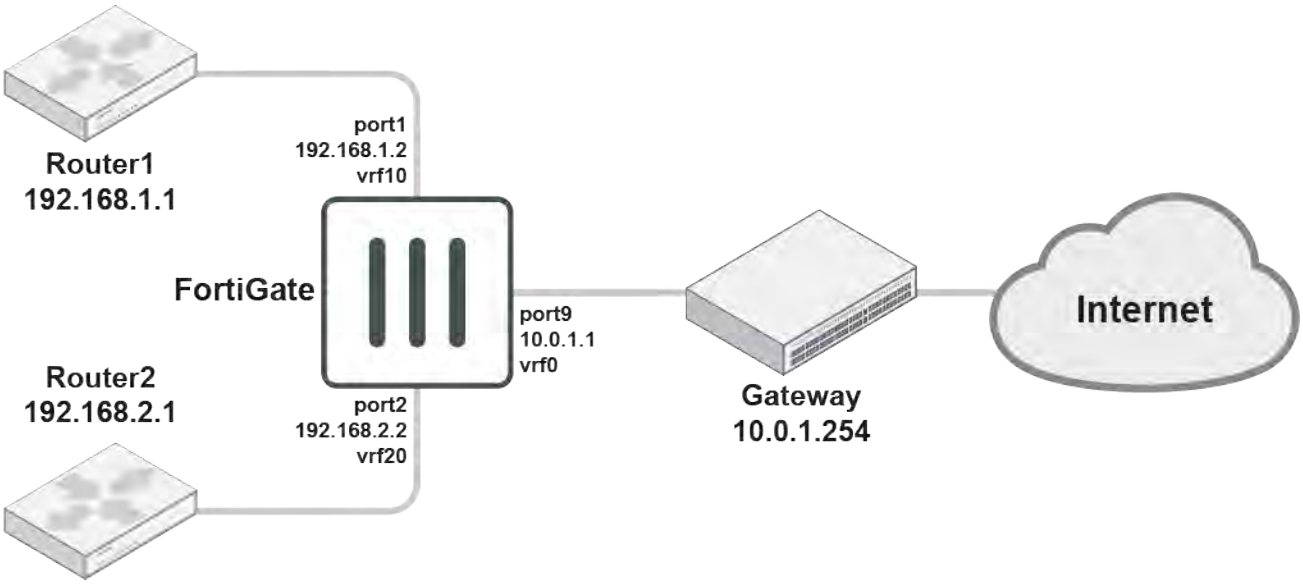
IBGP and EBGP support in VRF

Support is included for internal and external border gateway protocols (IBGP and EBGP) in virtual routing and forwarding (VRF).

FortiGate can establish neighbor connections with other FortiGates or routers, and the learned routes are put into different VRF tables according to the neighbor's settings.

This example uses the following topology:

Netskope Exhibit 1017



- BGP routes learned from the Router1 neighbor are put into vrf10.
- BGP routes learned from the Router2 neighbor are put into vrf20.

To configure this example:

```
config system interface
    edit port1
        set vrf 10
    next
    edit port2
        set vrf 20
    next
end
```

```
config router bgp
    config neighbor
        edit "192.168.1.1"
            set update-source port1
        next
        edit "192.168.2.1"
            set interface port2
        next
    end
end
```

Results

Using the above topology:

- Both Router1 and Router2 establish OSPF and BGP neighbor with the FortiGate.
- Router1 advertises 10.10.1.0/24 into OSPF and 10.10.2.0/24 into BGP.
- Router2 advertises 20.20.1.0/24 into OSPF and 20.20.2.0/24 into BGP.

When port1 and port2 have not set VRF, all of the routing is in VRF=0:

```
# get router info routing-table all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default

Routing table for VRF=0
S*      0.0.0.0/0 [5/0] via 10.0.1.254, port9
C       10.0.1.0/24 is directly connected, port9
O       10.10.1.0/24 [110/10] via 192.168.1.1, port1, 00:18:31
B       10.10.2.0/24 [20/200] via 192.168.1.1, port1, 00:01:31
O       20.20.1.0/22 [110/10] via 192.168.2.1, port2, 00:19:05
B       20.20.2.0/24 [20/200] via 192.168.2.1, port2, 00:01:31
C       192.168.1.0/24 is directly connected, port1
C       192.168.2.0/24 is directly connected, port2
```

After VRF is set for BGP, BGP routes are added to the VRF tables along with OSPF and connected routes:

```
# get router info routing-table all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default

Routing table for VRF=0
S*      0.0.0.0/0 [5/0] via 10.0.1.254, port9
C       10.0.1.0/24 is directly connected, port9

Routing table for VRF=10
O       10.10.1.0/24 [110/10] via 192.168.1.1, port1, 00:18:31
B       10.10.2.0/24 [20/200] via 192.168.1.1, port1, 00:01:31
C       192.168.1.0/24 is directly connected, port1

Routing table for VRF=20
O       20.20.1.0/22 [110/10] via 192.168.2.1, port2, 00:19:05
B       20.20.2.0/24 [20/200] via 192.168.2.1, port2, 00:01:31
C       192.168.2.0/24 is directly connected, port2
```

BGP neighbor groups

This feature is also supported in the BGP neighbor groups. For example:

```
config router bgp
    config neighbor-group
        edit "FGT"
            set update-source "port1"
        next
    end
    config neighbor-range
        edit 1
            set prefix 172.16.201.0 255.255.255.0
            set neighbor-group "FGT"
        next
```


Netskope Exhibit 1017

```
    end
end
```

Note that the `set interface` command is not supported.

ADVPN and shortcut paths

This topic provides an example of how to use SD-WAN and ADVPN together.

ADVPN (Auto Discovery VPN) is an IPsec technology that allows a traditional hub-and-spoke VPN's spokes to establish dynamic, on-demand, direct tunnels between each other to avoid routing through the topology's hub device. The primary advantage is that it provides full meshing capabilities to a standard hub-and-spoke topology. This greatly reduces the provisioning effort for full spoke-to-spoke low delay reachability, and addresses the scalability issues associated with very large fully meshed VPN networks.

If a customer's head office and branch offices all have two or more internet connections, they can build a dual-hub ADVPN network. Combined with SD-WAN technology, the customer can load-balance traffic to other offices on multiple dynamic tunnels, control specific traffic using specific connections, or choose better performance connections dynamically.

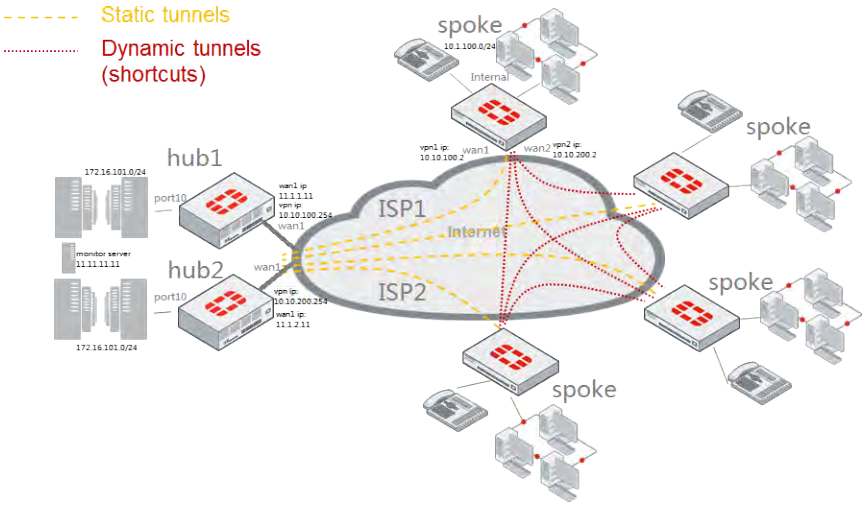


SD-WAN load-balance mode rules (or services) do not support ADVPN members. Other modes' rules, such as SLA and priority, support ADVPN members.

This topic covers three parts:

1. Configure dual-hub ADVPN with multiple branches.
2. Configure BGP to exchange routing information among hubs and spokes.
3. Configure SD-WAN on spoke to do load-balancing and control traffic.

Configuration example



A typical ADVPN configuration with SD-WAN usually has two hubs, and each spoke connects to two ISPs and establishes VPN tunnels with both hubs.

This example shows a hub-and-spoke configuration using two hubs and one spoke:

- Hub1 and Hub2 both use wan1 to connect to the ISPs and port10 to connect to internal network.
- Spoke1 uses wan1 to connect to ISP1 and wan2 to connect to ISP2.
- wan1 sets up VPN to hub1.
- wan2 sets up VPN to hub2.

The SD-WAN is configured on the spoke. It uses the two VPN interfaces as members and two rules to control traffic to headquarters or other spokes using ADVPN VPN interfaces. You can create more rules if required.

For this example:

- Use SD-WAN member 1 (via ISP1) and its dynamic shortcuts for financial department traffic if member 1 meets SLA requirements. If it doesn't meet SLA requirements, it will use SD-WAN member 2 (via ISP2).
- Use SD-WAN member 2 (via ISP2) and its dynamic shortcuts for engineering department traffic.
- Load balance other traffic going to hubs and other spokes between these two members.
- Set up all other traffic to go with their original ISP connection. All other traffic does not go through SD-WAN.
- Set up basic network configuration to let all hubs and spokes connect to their ISPs and the Internet.

Hub internal network	172.16.101.0/24
Spoke1 internal network	10.1.100.0/24
ADVPN 1 network	10.10.100.0/24
ADVPN 2 network	10.10.200.0/24
Hub1 wan1 IP	11.1.1.11
Hub2 wan1 IP	11.1.2.11
Hub1 VPN IP	10.10.100.254
Hub2 VPN IP	10.10.200.254
Spoke1 to hub1 VPN IP	10.10.100.2
Spoke1 to hub2 VPN IP	10.10.200.2
Ping server in Headquarters	11.11.11.11
Internal subnet of spoke1	22.1.1.0/24
Internal subnet of spoke2	33.1.1.0/24
Firewall addresses	Configure hub_subnets and spoke_subnets before using in policies. These can be customized.

The GUI does not support some ADVPN related options, such as auto-discovery-sender, auto-discovery-receiver, auto-discovery-forwarder, and IBGP neighbor-group setting, so this example only provides CLI configuration commands.

Hub1 sample configuration

To configure the IPsec phase1 and phase2 interface:

```
config vpn ipsec phase1-interface
    edit "hub-phase1"
        set type dynamic
```


Netskope Exhibit 1017

```
        set interface "wan1"
        set peertype any
        set net-device disable
        set proposal aes128-sha256 aes256-sha256 3des-sha256 aes128-sha1 aes256-sha1 3des-
sha1
        set add-route disable
        set dpd on-idle
        set auto-discovery-sender enable
        set tunnel-search nexthop
        set psksecret sample
        set dpd-retryinterval 5
    next
end
config vpn ipsec phase2-interface
    edit "hub-phase2"
        set phase1name "hub-phase1"
        set proposal aes128-sha1 aes256-sha1 3des-sha1 aes128-sha256 aes256-sha256 3des-
sha256
    next
end
```

To configure the VPN interface and BGP:

```
config system interface
    edit "hub-phase1"
        set ip 10.10.100.254 255.255.255.255
        set remote-ip 10.10.100.253 255.255.255.0
    next
end
config router bgp
    set as 65505
    config neighbor-group
        edit "advpn"
            set link-down-failover enable
            set remote-as 65505
            set route-reflector-client enable
        next
    end
    config neighbor-range
        edit 1
            set prefix 10.10.100.0 255.255.255.0
            set neighbor-group "advpn"
        next
    end
    config network
        edit 1
            set prefix 172.16.101.0 255.255.255.0
        next
        edit 2
            set prefix 11.11.11.0 255.255.255.0
         next
    end
end
```


Netskope Exhibit 1017

To configure the firewall policy:

```
config firewall policy
    edit 1
        set name "spoke2hub"
        set srcintf "hub-phase1"
        set dstintf "port10"
        set srcaddr "spoke_subnets"
        set dstaddr "hub_subnets"
        set action accept
        set schedule "always"
        set service "ALL"
        set comments "allow traffic from spokes to headquater"
    next
    edit 2
        set name "spoke2spoke"
        set srcintf "hub-phase1"
        set dstintf "hub-phase1"
        set srcaddr "spoke_subnets"
        set dstaddr "spoke_subnets"
        set action accept
        set schedule "always"
        set service "ALL"
        set comments "allow traffic from spokes to spokes"
    next
    edit 3
        set name "internal2spoke"
        set srcintf "port10"
        set dstintf "hub-phase1"
        set srcaddr "hub_subnets"
        set dstaddr "spoke_subnets"
        set action accept
        set schedule "always"
        set service "ALL"
        set comments "allow traffic from headquater to spokes"
    next
end
```

Hub2 sample configuration

Hub2 configuration is the same as hub1 except the wan1 IP address, VPN interface IP address, and BGP neighbor-range prefix.

To configure the IPsec phase1 and phase2 interface:

```
config vpn ipsec phase1-interface
    edit "hub-phase1"
        set type dynamic
        set interface "wan1"
        set peertype any
        set net-device disable
        set proposal aes128-sha256 aes256-sha256 3des-sha256 aes128-sha1 aes256-sha1 3des-
sha1
        set add-route disable
        set dpd on-idle
        set auto-discovery-sender enable
```


Netskope Exhibit 1017

```
        set tunnel-search nexthop
        set psksecret sample
        set dpd-retryinterval 5
    next
end
config vpn ipsec phase2-interface
    edit "hub-phase2"
        set phase1name "hub-phase1"
        set proposal aes128-sha1 aes256-sha1 3des-sha1 aes128-sha256 aes256-sha256 3des-
sha256
    next
end
```

To configure the VPN interface and BGP:

```
config system interface
    edit "hub-phase1"
        set ip 10.10.200.254 255.255.255.255
        set remote-ip 10.10.200.253 255.255.255.0
    next
end
config router bgp
    set as 65505
    config neighbor-group
        edit "advpn"
            set link-down-failover enable
            set remote-as 65505
            set route-reflector-client enable
        next
    end
    config neighbor-range
        edit 1
            set prefix 10.10.200.0 255.255.255.0
            set neighbor-group "advpn"
        next
    end
    config network
        edit 1
            set prefix 172.16.101.0 255.255.255.0
        next
        edit 2
            set prefix 11.11.11.0 255.255.255.0
         next
    end
end
```

To configure the firewall policy:

```
config firewall policy
    edit 1
        set name "spoke2hub"
        set srcintf "hub-phase1"
        set dstintf "port10"
        set srcaddr "spoke_subnets"
        set dstaddr "hub_subnets"
        set action accept
```


Netskope Exhibit 1017

```
        set schedule "always"
        set service "ALL"
        set comments "allow traffic from spokes to headquater"
    next
    edit 2
        set name "spoke2spoke"
        set srcintf "hub-phase1"
        set dstintf "hub-phase1"
        set srcaddr "spoke_subnets"
        set dstaddr "spoke_subnets"
        set action accept
        set schedule "always"
        set service "ALL"
        set comments "allow traffic from spokes to spokes"
    next
    edit 3
        set name "internal2spoke"
        set srcintf "port10"
        set dstintf "hub-phase1"
        set srcaddr "hub_subnets"
        set dstaddr "spoke_subnets"
        set action accept
        set schedule "always"
        set service "ALL"
        set comments "allow traffic from headquater to spokes"
    next
end
```

Spoke1 sample configuration

To configure the IPsec phase1 and phase2 interface:

```
config vpn ipsec phase1-interface
    edit "spoke1-phase1"
        set interface "wan1"
        set peertype any
        set net-device enable
        set proposal aes128-sha256 aes256-sha256 aes128-sha1 aes256-sha1
        set add-route disable
        set dpd on-idle
        set auto-discovery-receiver enable
        set remote-gw 11.1.1.11
        set psksecret sample
        set dpd-retryinterval 5
    next
    edit "spoke1-2-phase1"
        set interface "wan2"
        set peertype any
        set net-device enable
        set proposal aes128-sha256 aes256-sha256 aes128-sha1 aes256-sha1
        set add-route disable
        set dpd on-idle
        set auto-discovery-receiver enable
        set remote-gw 11.1.2.11
        set psksecret sample
        set dpd-retryinterval 5
```


Netskope Exhibit 1017

```
    next
end
config vpn ipsec phase2-interface
    edit "spoke1-phase2"
        set phase1name "spoke1-phase1"
        set proposal aes128-sha1 aes256-sha1 aes128-sha256 aes256-sha256 aes128gcm aes256gcm
chacha20poly1305
        set auto-negotiate enable
    next
    edit "spoke1-2-phase2"
        set phase1name "spoke1-2-phase1"
        set proposal aes128-sha1 aes256-sha1 aes128-sha256 aes256-sha256 aes128gcm aes256gcm
chacha20poly1305
        set auto-negotiate enable
    next
end
```

To configure the VPN interface and BGP:

```
config system interface
    edit "spoke1-phase1"
        set ip 10.10.100.2 255.255.255.255
        set remote-ip 10.10.100.254 255.255.255.0
    next
    edit "spoke1-2-phase1"
        set ip 10.10.200.2 255.255.255.255
        set remote-ip 10.10.200.254 255.255.255.0
    next
end
config router bgp
    set as 65505
    config neighbor
        edit "10.10.100.254"
            set advertisement-interval 1
            set link-down-failover enable
            set remote-as 65505
        next
        edit "10.10.200.254"
            set advertisement-interval 1
            set link-down-failover enable
            set remote-as 65505
        next
    end
    config network
        edit 1
            set prefix 10.1.100.0 255.255.255.0
        next
    end
end
```

To configure SD-WAN:

```
config system virtual-wan-link
    set status enable
    config members
        edit 1
```


Netskope Exhibit 1017

```
            set interface "spoke1-phase1"
        next
        edit 2
            set interface "spoke1-2-phase1"
        next
    end
    config health-check
        edit "ping"
            set server "11.11.11.11"
            set members 1 2
            config sla
                edit 1
                    set latency-threshold 200
                    set jitter-threshold 50
                    set packetloss-threshold 5
                next
            end
        next
    end
    config service
        edit 1
            set mode sla
            set dst "finacial-department"
            config sla
                edit "ping"
                    set id 1
                next
            end
            set priority-member 1 2
        next
        edit 2
            set member 2
            set dst "engineering-department"
        next
    end
end
```

To configure the firewall policy:

```
config firewall policy
    edit 1
        set name "outbound_advpn"
        set srcintf "internal"
        set dstintf "virtual-wan-link"
        set srcaddr "spoke_subnets"
        set dstaddr "spoke_subnets" "hub_subnets"
        set action accept
        set schedule "always"
        set service "ALL"
        set comments "allow internal traffic going out to headquater and other spokes"
    next
    edit 2
        set name "inbound_advpn"
        set srcintf "virtual-wan-link"
        set dstintf "internal"
        set srcaddr "spoke_subnets" "hub_subnets"
```


Netskope Exhibit 1017

```
        set dstaddr "spoke_subnets"
        set action accept
        set schedule "always"
        set service "ALL"
        set comments "allow headquater and other spokes traffic coming in"
    next
end
```

Troubleshooting ADVPN and shortcut paths

Before spoke vs spoke shortcut VPN is established

Use the following CLI commands to check status before spoke vs spoke shortcut VPN is established.

```
# get router info bgp summary
BGP router identifier 2.2.2.2, local AS number 65505
BGP table version is 13
3 BGP AS-PATH entries
0 BGP community entries

Neighbor        V         AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.10.100.254   4      65505    3286    3270       11    0    0 00:02:15        5
10.10.200.254   4      65505    3365    3319       12    0    0 00:02:14        5

Total number of neighbors 2
```

```
# get router info routing-table bgp

Routing table for VRF=0
B*      0.0.0.0/0 [200/0] via 10.10.200.254, spoke1-2-phase1, 00:00:58
                  [200/0] via 10.10.100.254, spoke1-phase1, 00:00:58
B       1.1.1.1/32 [200/0] via 11.1.1.1 (recursive via 12.1.1.1), 00:01:29
                   [200/0] via 11.1.1.1 (recursive via 12.1.1.1), 00:01:29
B       11.11.11.0/24 [200/0] via 10.10.200.254, spoke1-2-phase1, 00:01:29
                      [200/0] via 10.10.100.254, spoke1-phase1, 00:01:29
B       33.1.1.0/24 [200/0] via 10.10.200.3, spoke1-2-phase1, 00:00:58
                    [200/0] via 10.10.100.3, spoke1-phase1, 00:00:58
                    [200/0] via 10.10.200.3, spoke1-2-phase1, 00:00:58
                    [200/0] via 10.10.100.3, spoke1-phase1, 00:00:58
```

```
# diagnose vpn tunnel list
list all ipsec tunnel in vd 3
------------------------------------------------------
name=spoke1-phase1 ver=1 serial=5 12.1.1.2:0->11.1.1.11:0 dst_mtu=15324
bound_if=48 lgwy=static/1 tun=intf/0 mode=auto/1 encap=none/536 options[0218]=npu create_dev
frag-rfc  accept_traffic=1

proxyid_num=1 child_num=0 refcnt=22 ilast=0 olast=0 ad=r/2
stat: rxp=1 txp=185 rxb=16428 txb=11111
dpd: mode=on-demand on=1 idle=20000ms retry=3 count=0 seqno=4
natt: mode=none draft=0 interval=0 remote_port=0
proxyid=spoke1 proto=0 sa=1 ref=4 serial=1 auto-negotiate adr
  src: 0:0.0.0.0/0.0.0.0:0
  dst: 0:0.0.0.0/0.0.0.0:0
  SA:  ref=6 options=1a227 type=00 soft=0 mtu=15262 expire=42820/0B replaywin=2048
       seqno=ba esn=0 replaywin_lastseq=00000002 itn=0 qat=0
  life: type=01 bytes=0/0 timeout=42903/43200
```


Netskope Exhibit 1017

```
  dec: spi=03e01a2a esp=aes key=16 56e673f0df05186aa657f55cbb631c13
       ah=sha1 key=20 b0d50597d9bed763c42469461b03da8041f87e88
  enc: spi=2ead61bc esp=aes key=16 fe0ccd4a3ec19fe6d520c437eb6b8897
       ah=sha1 key=20 e3e669bd6df41b88eadaacba66463706f26fb53a
  dec:pkts/bytes=1/16368, enc:pkts/bytes=185/22360
  npu_flag=03 npu_rgwy=11.1.1.11 npu_lgwy=12.1.1.2 npu_selid=0 dec_npuid=1 enc_npuid=1
-----------------------------------------------------
name=spoke1-2-phase1 ver=1 serial=6 112.1.1.2:0->11.1.2.11:0 dst_mtu=15324
bound_if=90 lgwy=static/1 tun=intf/0 mode=auto/1 encap=none/536 options[0218]=npu create_dev
frag-rfc  accept_traffic=1

proxyid_num=1 child_num=0 refcnt=21 ilast=0 olast=0 ad=r/2
stat: rxp=1 txp=186 rxb=16498 txb=11163
dpd: mode=on-demand on=1 idle=20000ms retry=3 count=0 seqno=74
natt: mode=none draft=0 interval=0 remote_port=0
proxyid=spoke1-2 proto=0 sa=1 ref=4 serial=1 auto-negotiate adr
  src: 0:0.0.0.0/0.0.0.0:0
  dst: 0:0.0.0.0/0.0.0.0:0
  SA:  ref=6 options=1a227 type=00 soft=0 mtu=15262 expire=42818/0B replaywin=2048
       seqno=bb esn=0 replaywin_lastseq=00000002 itn=0 qat=0
  life: type=01 bytes=0/0 timeout=42901/43200
  dec: spi=03e01a2b esp=aes key=16 fe49f5042a5ad236250bf53312db1346
       ah=sha1 key=20 5dbb15c8cbc046c284bb1c6425dac2b3e15bec85
  enc: spi=2ead61bd esp=aes key=16 d6d97be52c3cccb9e88f28a9db64ac46
       ah=sha1 key=20 e20916ae6ea2295c2fbd5cbc8b8f5dd8b17f52f1
  dec:pkts/bytes=1/16438, enc:pkts/bytes=186/22480
  npu_flag=03 npu_rgwy=11.1.2.11 npu_lgwy=112.1.1.2 npu_selid=1 dec_npuid=1 enc_npuid=1
```

diagnose sys virtual-wan-link service

```
Service(1): Address Mode(IPV4) flags=0x0
  TOS(0x0/0x0), Protocol(0: 1->65535), Mode(sla)
  Member sub interface:
  Members:
    1: Seq_num(1), alive, sla(0x1), cfg_order(0), cost(0), selected
    2: Seq_num(2), alive, sla(0x1), cfg_order(1), cost(0), selected
  Dst address: 33.1.1.1-33.1.1.100

Service(2): Address Mode(IPV4) flags=0x0
  TOS(0x0/0x0), Protocol(0: 1->65535), Mode(manual)
  Member sub interface:
  Members:
    1: Seq_num(2), alive, selected
  Dst address: 33.1.1.101-33.1.1.200
```

diagnose firewall proute list

```
list route policy info(vf=vd2):

id=2132869121 vwl_service=1 vwl_mbr_seq=1 dscp_tag=0xff 0xff flags=0x0 tos=0x00 tos_
mask=0x00 protocol=0 sport=0:65535 iif=0 dport=1-65535 oif=70 oif=71
destination(1): 33.1.1.1-33.1.1.100
source wildcard(1): 0.0.0.0/0.0.0.0

id=2132869122 vwl_service=2 vwl_mbr_seq=2 dscp_tag=0xff 0xff flags=0x0 tos=0x00 tos_
mask=0x00 protocol=0 sport=0:65535 iif=0 dport=1-65535 oif=71
destination(1): 33.1.1.101-33.1.1.200
source wildcard(1): 0.0.0.0/0.0.0.0
```


Netskope Exhibit 1017

After spoke vs spoke shortcut VPN is established

Use the following CLI commands to check status after spoke vs spoke shortcut VPN is established.

```
# get router info routing-table bgp

Routing table for VRF=0
B*      0.0.0.0/0 [200/0] via 10.10.200.254, spoke1-2-phase1, 00:01:33
                  [200/0] via 10.10.100.254, spoke1-phase1, 00:01:33
B       1.1.1.1/32 [200/0] via 11.1.1.1 (recursive via 12.1.1.1), 00:02:04
                   [200/0] via 11.1.1.1 (recursive via 12.1.1.1), 00:02:04
B       11.11.11.0/24 [200/0] via 10.10.200.254, spoke1-2-phase1, 00:02:04
                      [200/0] via 10.10.100.254, spoke1-phase1, 00:02:04
B       33.1.1.0/24 [200/0] via 10.10.200.3, spoke1-2-phase1_0, 00:01:33
                    [200/0] via 10.10.100.3, spoke1-phase1_0, 00:01:33
                    [200/0] via 10.10.200.3, spoke1-2-phase1_0, 00:01:33
                    [200/0] via 10.10.100.3, spoke1-phase1_0, 00:01:33
```

```
# diagnose sys virtual-wan-link service

Service(1): Address Mode(IPV4) flags=0x0
  TOS(0x0/0x0), Protocol(0: 1->65535), Mode(sla)
  Member sub interface:
    1: seq_num(1), interface(spoke1-phase1):
       1: spoke1-phase1_0(111)
    2: seq_num(2), interface(spoke1-2-phase1):
       1: spoke1-2-phase1_0(113)
  Members:
    1: Seq_num(1), alive, sla(0x1), cfg_order(0), cost(0), selected
    2: Seq_num(2), alive, sla(0x1), cfg_order(1), cost(0), selected
  Dst address: 33.1.1.1-33.1.1.100

Service(2): Address Mode(IPV4) flags=0x0
  TOS(0x0/0x0), Protocol(0: 1->65535), Mode(manual)
  Member sub interface:
    1: seq_num(2), interface(spoke1-2-phase1):
       1: spoke1-2-phase1_0(113)
  Members:
    1: Seq_num(2), alive, selected
  Dst address: 33.1.1.101-33.1.1.200
```

```
# diagnose vpn tunnel list
list all ipsec tunnel in vd 3
------------------------------------------------------
name=spoke1-phase1 ver=1 serial=5 12.1.1.2:0->11.1.1.11:0 dst_mtu=15324
bound_if=48 lgwy=static/1 tun=intf/0 mode=auto/1 encap=none/536 options[0218]=npu create_dev
frag-rfc  accept_traffic=1

proxyid_num=1 child_num=1 refcnt=20 ilast=0 olast=0 ad=r/2
stat: rxp=1 txp=759 rxb=16428 txb=48627
dpd: mode=on-demand on=1 idle=20000ms retry=3 count=0 seqno=4
natt: mode=none draft=0 interval=0 remote_port=0
proxyid=vd2-1 proto=0 sa=1 ref=5 serial=1 auto-negotiate adr
  src: 0:0.0.0.0/0.0.0.0:0
  dst: 0:0.0.0.0/0.0.0.0:0
  SA:  ref=6 options=1a227 type=00 soft=0 mtu=15262 expire=42536/0B replaywin=2048
       seqno=2f8 esn=0 replaywin_lastseq=00000002 itn=0 qat=0
  life: type=01 bytes=0/0 timeout=42898/43200
```


Netskope Exhibit 1017

```
  dec: spi=03e01a42 esp=aes key=16 1f131bda108d33909d49fc2778bd08bb
       ah=sha1 key=20 14131d3f0da9b741a2fd13d530b0553aa1f58983
  enc: spi=2ead61d8 esp=aes key=16 81ed24d5cd7bb59f4a80dceb5a560e1f
       ah=sha1 key=20 d2ccc2f3223ce16514e75f672cd88c4b4f48b681
  dec:pkts/bytes=1/16360, enc:pkts/bytes=759/94434
  npu_flag=03 npu_rgwy=11.1.1.11 npu_lgwy=12.1.1.2 npu_selid=0 dec_npuid=1 enc_npuid=1
-----------------------------------------------------
name=spoke1-2-phase1 ver=1 serial=6 112.1.1.2:0->11.1.2.11:0 dst_mtu=15324
bound_if=90 lgwy=static/1 tun=intf/0 mode=auto/1 encap=none/536 options[0218]=npu create_dev
frag-rfc  accept_traffic=1

proxyid_num=1 child_num=1 refcnt=19 ilast=0 olast=0 ad=r/2
stat: rxp=1 txp=756 rxb=16450 txb=48460
dpd: mode=on-demand on=1 idle=20000ms retry=3 count=0 seqno=74
natt: mode=none draft=0 interval=0 remote_port=0
proxyid=vd2-2 proto=0 sa=1 ref=5 serial=1 auto-negotiate adr
  src: 0:0.0.0.0/0.0.0.0:0
  dst: 0:0.0.0.0/0.0.0.0:0
  SA:  ref=6 options=1a227 type=00 soft=0 mtu=15262 expire=42538/0B replaywin=2048
       seqno=2f5 esn=0 replaywin_lastseq=00000002 itn=0 qat=0
  life: type=01 bytes=0/0 timeout=42900/43200
  dec: spi=03e01a43 esp=aes key=16 7fc87561369f88b56d08bfda769eb45b
       ah=sha1 key=20 0ed554ef231c5ac16dc2e71d1907d7347dda33d6
  enc: spi=2ead61d9 esp=aes key=16 00286687aa1762e7d8216881d6720ef3
       ah=sha1 key=20 59d5eec6299ebcf038c190860774e2833074d7c3
  dec:pkts/bytes=1/16382, enc:pkts/bytes=756/94058
  npu_flag=03 npu_rgwy=11.1.2.11 npu_lgwy=112.1.1.2 npu_selid=1 dec_npuid=1 enc_npuid=1
-----------------------------------------------------
name=spoke1-phase1_0 ver=1 serial=55 12.1.1.2:0->13.1.1.3:0 dst_mtu=15324
bound_if=48 lgwy=static/1 tun=intf/0 mode=dial_inst/3 encap=none/728 options[02d8]=npu
create_dev no-sysctl rgwy-chg frag-rfc  accept_traffic=1

 parent=vd2-1 index=0
proxyid_num=1 child_num=0 refcnt=18 ilast=8 olast=8 ad=r/2
stat: rxp=0 txp=0 rxb=0 txb=0
dpd: mode=on-demand on=1 idle=20000ms retry=3 count=0 seqno=0
natt: mode=none draft=0 interval=0 remote_port=0
proxyid=vd2-1 proto=0 sa=1 ref=2 serial=1 auto-negotiate adr
  src: 0:0.0.0.0/0.0.0.0:0
  dst: 0:0.0.0.0/0.0.0.0:0
  SA:  ref=3 options=1a227 type=00 soft=0 mtu=15262 expire=42893/0B replaywin=2048
       seqno=1 esn=0 replaywin_lastseq=00000000 itn=0 qat=0
  life: type=01 bytes=0/0 timeout=42901/43200
  dec: spi=03e01a44 esp=aes key=16 c3b77a98e3002220e2373b73af14df6e
       ah=sha1 key=20 d18d107c248564933874f60999d6082fd7a78948
  enc: spi=864f6dba esp=aes key=16 eb6181806ccb9bac37931f9eadd4d5eb
       ah=sha1 key=20 ab788f7a372877a5603c4ede1be89a592fc21873
  dec:pkts/bytes=0/0, enc:pkts/bytes=0/0
  npu_flag=00 npu_rgwy=13.1.1.3 npu_lgwy=12.1.1.2 npu_selid=51 dec_npuid=0 enc_npuid=0
-----------------------------------------------------
name=spoke1-2-phase1_0 ver=1 serial=57 112.1.1.2:0->113.1.1.3:0 dst_mtu=15324
bound_if=90 lgwy=static/1 tun=intf/0 mode=dial_inst/3 encap=none/728 options[02d8]=npu
create_dev no-sysctl rgwy-chg frag-rfc  accept_traffic=1

 parent=vd2-2 index=0
proxyid_num=1 child_num=0 refcnt=17 ilast=5 olast=5 ad=r/2
```


Netskope Exhibit 1017

```
stat: rxp=0 txp=0 rxb=0 txb=0
dpd: mode=on-demand on=1 idle=20000ms retry=3 count=0 seqno=0
natt: mode=none draft=0 interval=0 remote_port=0
proxyid=vd2-2 proto=0 sa=1 ref=3 serial=1 auto-negotiate adr
  src: 0:0.0.0.0/0.0.0.0:0
  dst: 0:0.0.0.0/0.0.0.0:0
  SA:  ref=3 options=1a227 type=00 soft=0 mtu=15262 expire=42900/0B replaywin=2048
       seqno=1 esn=0 replaywin_lastseq=00000000 itn=0 qat=0
  life: type=01 bytes=0/0 timeout=42901/43200
  dec: spi=03e01a45 esp=aes key=16 0beb519ed9f800e8b4c0aa4e1df7da35
       ah=sha1 key=20 bc9f38db5296cce4208a69f1cc8a9f7ef4803c37
  enc: spi=864f6dbb esp=aes key=16 1d26e3556afcdb9f8e3e33b563b44228
       ah=sha1 key=20 564d05ef6f7437e1fd0a88d5fee7b6567f9d387e
  dec:pkts/bytes=0/0, enc:pkts/bytes=0/0
  npu_flag=00 npu_rgwy=113.1.1.3 npu_lgwy=112.1.1.2 npu_selid=53 dec_npuid=0 enc_npuid=0
```

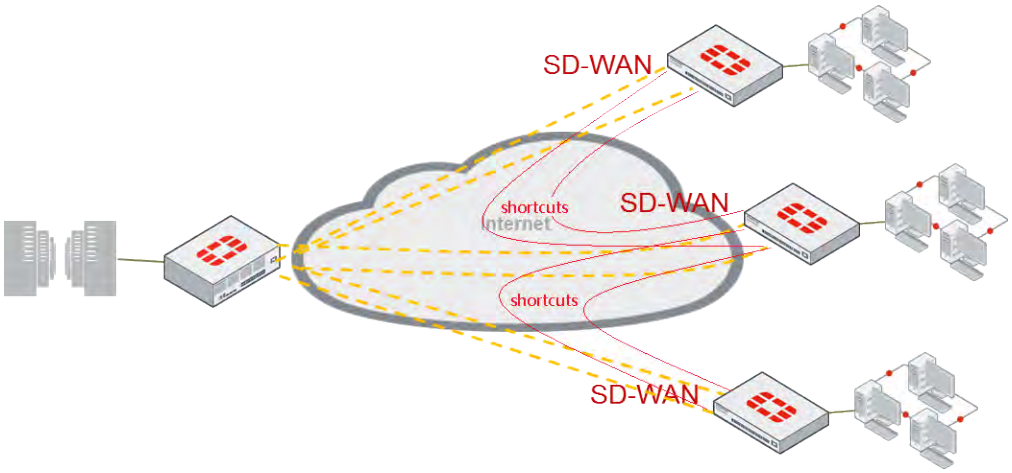
```
# diagnose firewall proute list
list route policy info(vf=vd2):

id=2132869121 vwl_service=1 vwl_mbr_seq=1 dscp_tag=0xff 0xff flags=0x0 tos=0x00 tos_
mask=0x00 protocol=0 sport=0:65535 iif=0 dport=1-65535 oif=111 oif=70 oif=113 oif=71
destination(1): 33.1.1.1-33.1.1.100
source wildcard(1): 0.0.0.0/0.0.0.0

id=2132869122 vwl_service=2 vwl_mbr_seq=2 dscp_tag=0xff 0xff flags=0x0 tos=0x00 tos_
mask=0x00 protocol=0 sport=0:65535 iif=0 dport=1-65535 oif=113 oif=71
destination(1): 33.1.1.101-33.1.1.200
source wildcard(1): 0.0.0.0/0.0.0.0
```

SD-WAN monitor on ADVPN shortcuts

SD-WAN monitors ADVPN shortcut link quality by dynamically creating link monitors for each ADVPN link. The dynamic link monitor on the spoke will use ICMP probes and the IP address of the gateway as the monitored server. These ICMP probes will not be counted as actual user traffic that keeps the spoke-to-spoke tunnel alive.



- When no shortcut is established:

  ```
  # diagnose sys virtual-wan-link health-check
  Health Check(ping):
  ```

```
Seq(1 tunnel-1): state(alive), packet-loss(0.000%) latency(0.038), jitter(0.006) sla_
map=0x3
Seq(2 tunnel-2): state(alive), packet-loss(0.000%) latency(0.035), jitter(0.004) sla_
map=0x3
```

- When one shortcut is established:

```
# diagnose sys virtual-wan-link health-check
Health Check(ping):
Seq(1 tunnel-1): state(alive), packet-loss(0.000%) latency(0.039), jitter(0.003) sla_
map=0x3
Seq(1 tunnel-1_0): state(alive), packet-loss(0.000%) latency(0.060), jitter(0.023) sla_
map=0x3
Seq(2 tunnel-2): state(alive), packet-loss(0.000%) latency(0.035), jitter(0.002) sla_
map=0x3
```

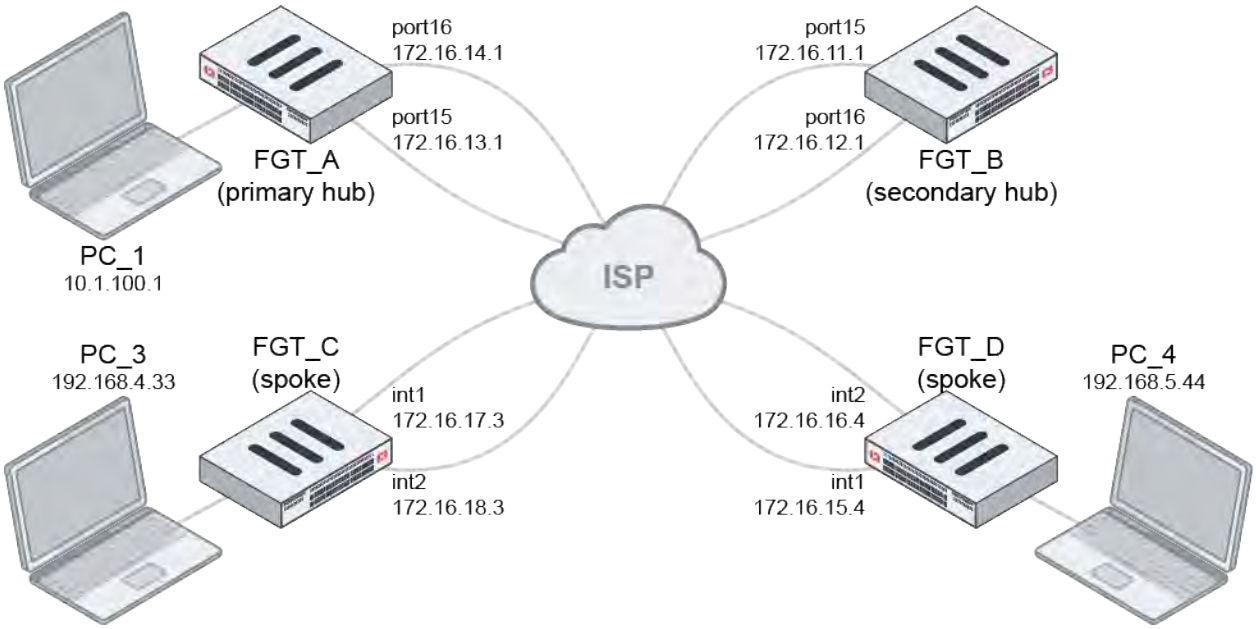
- When more than one shortcut is established:

```
# diagnose sys virtual-wan-link health-check
Health Check(ping):
Seq(1 tunnel-1): state(alive), packet-loss(0.000%) latency(0.036), jitter(0.004) sla_
map=0x3
Seq(1 tunnel-1_0): state(alive), packet-loss(0.000%) latency(0.041), jitter(0.009) sla_
map=0x3
Seq(2 tunnel-2): state(alive), packet-loss(0.000%) latency(0.030), jitter(0.005) sla_
map=0x3
Seq(2 tunnel-2_0): state(alive), packet-loss(0.000%) latency(0.031), jitter(0.004) sla_
map=0x3
```

SD-WAN integration with OCVPN

OCVPN has the capability to enable SD-WAN in order to dynamically add its tunnel interfaces as SD-WAN members. Users can configure SD-WAN health checks and service rules to direct traffic over the OCVPN tunnels.

The following example uses a dual hub and spoke topology. Each hub and spoke has two WAN link connections to the ISP. The spokes generate two IPsec tunnels to each hub (four tunnels in total). BGP neighbors are established over each tunnel and routes from the hubs and other spokes learned from all neighbors, which forms an ECMP scenario. All tunnels are placed as SD-WAN members, so traffic can be distributed across tunnels based on the configured SD-WAN service rules.



To integrate SD-WAN with OCVPN in the GUI:

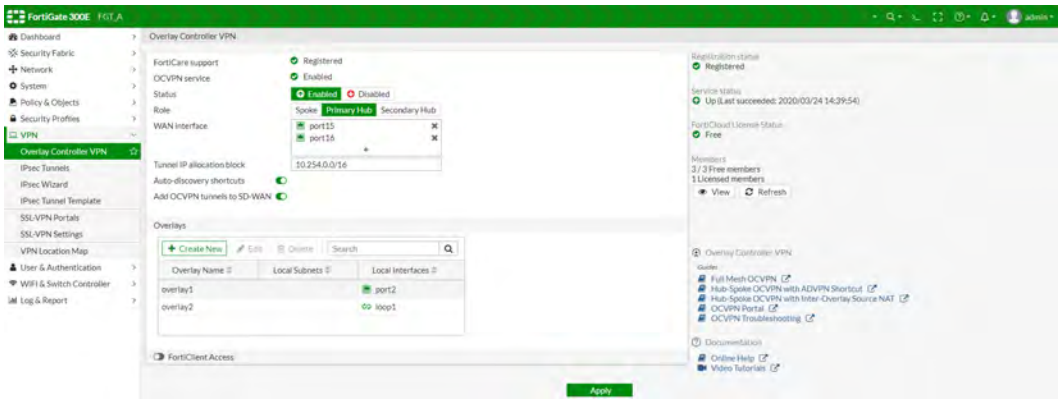
1. Configure the primary hub:
   a. Go to *VPN* > *Overlay Controller VPN* and set the *Status* to *Enable*.
   b. For *Role*, select *Primary Hub*.
   c. Enter the WAN interfaces (*port15* and *port16*) and tunnel IP allocation block (*10.254.0.0/16*).

   The WAN interface is position sensitive, meaning a tunnel will be created with the first position interface on the hub to the first position interface on the spoke, and so on. In this example, FGT_A (primary hub) will create two tunnels with FGT_C (spoke):
   - FGT_A port15 <==> FGT_C internal1
   - FGT_A port16 <==> FGT_C internal2

   d. Enable *Auto-discovery shortcuts*.
   e. Enable *Add OCVPN tunnels to SD-WAN*. The IPsec tunnels will be added automatically to the SD-WAN members if SD-WAN is enabled.
2. Configure the overlays on the primary hub:
   a. In the *Overlays* section, click *Create New*.
   b. Enter a name and add the local interface (*port2*). Note the overlay is either based on local subnets or local interfaces, but not both.
   By default, inter-overlay traffic is not enabled. Toggle *Allow traffic from other overlays* to enable it.
   c. Click *OK* and repeat these steps to create the second overlay (*loop1*).



d. Click *Apply*.



3. Configure the secondary hub with the same settings as the primary hub.

4. Configure the spoke:

a. Go to *VPN* > *Overlay Controller VPN* and set the *Status* to *Enable*.

b. For *Role*, select *Spoke*.

c. Enter the WAN interfaces (*internal1* and *internal2*).

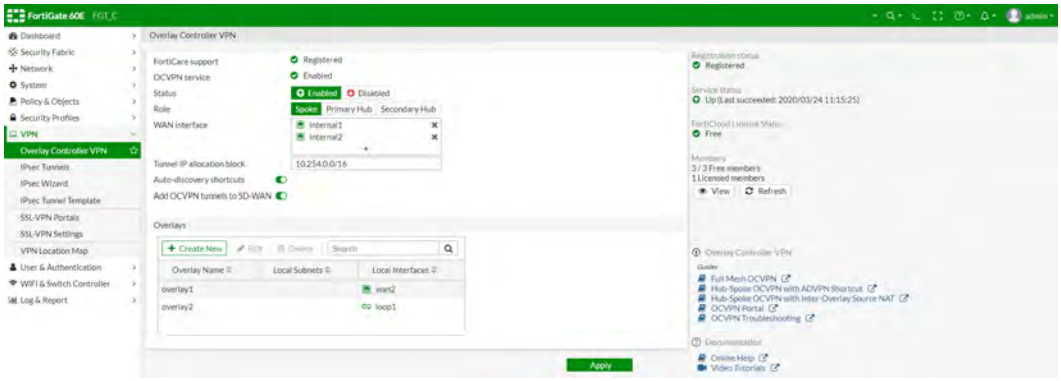
d. Enable *Auto-discovery shortcuts*.

e. Enable *Add OCVPN tunnels to SD-WAN*. The IPsec tunnels will be added automatically to the SD-WAN members if SD-WAN is enabled.

f. Configure the overlays.

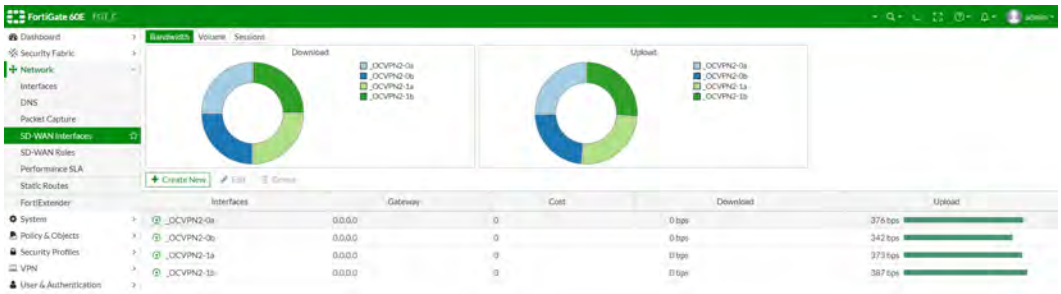
The overlay names on the spokes must match the hub for the traffic to be allowed through the same overlay.

g. Click *Apply*.



5. Configure the other spoke with the same settings.

6. On a spoke, go to *Network* > *SD-WAN Interfaces* to view the configuration generated by OCVPN.



Firewall policies will be automatically generated by OCVPN between the local interfaces and the SD-WAN interface. Each policy will define the proper local and remote networks for its source and destination addresses.

To integrate SD-WAN with OCVPN in the CLI:

1. Configure the primary hub:

```
config vpn ocvpn
    set role primary-hub
    set sdwan enable
    set wan-interface "port15" "port16"
    set ip-allocation-block 10.254.0.0 255.255.0.0
    config overlays
        edit "overlay1"
            config subnets
                edit 1
                    set type interface
                    set interface "port2"
                next
            end
        next
        edit "overlay2"
            config subnets
                edit 1
                    set type interface
                    set interface "loop1"
                next
            end
        next
    end
end
```

2. Configure the secondary hub with the same settings as the primary hub.
3. Configure the spoke:

```
config vpn ocvpn
    set status enable
    set sdwan enable
    set wan-interface "internal1" "internal2"
    config overlays
        edit "overlay1"
            config subnets
                edit 1
                    set type interface
                    set interface "wan2"
```

```
                next
            end
        next
        edit "overlay2"
            config subnets
                edit 1
                    set type interface
                    set interface "loop1"
                next
            end
        next
    end
end
```

4. Configure the other spoke with the same settings.
5. Configure SD-WAN:

```
config system virtual-wan-link
    set status enable
    config members
        edit 1
            set interface "_OCVPN2-0a"
        next
        edit 2
            set interface "_OCVPN2-0b"
        next
        edit 3
            set interface "_OCVPN2-1a"
        next
        edit 4
            set interface "_OCVPN2-1b"
        next
    end
end
```

Firewall policies will be automatically generated by OCVPN between the local interfaces and the SD-WAN interface. Each policy will define the proper local and remote networks for its source and destination addresses.

To verify the integration is working after the ADVPN shortcut is triggered:

1. Check the routing table on the spoke:

```
FGT_C # get router info routing-table all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default

Routing table for VRF=0
S*      0.0.0.0/0 [10/0] via 172.16.17.2, internal1
                  [10/0] via 172.16.18.2, internal2
B       10.1.100.0/24 [200/0] via 10.254.7.254, _OCVPN2-0a, 00:10:24
                      [200/0] via 10.254.15.254, _OCVPN2-0b, 00:10:24
B       10.1.200.0/24 [200/0] via 10.254.7.254, _OCVPN2-0a, 00:10:24
                      [200/0] via 10.254.15.254, _OCVPN2-0b, 00:10:24
```


Netskope Exhibit 1017

```
B       10.2.100.0/24 [200/0] via 10.254.71.254, _OCVPN2-1a, 00:10:15
                      [200/0] via 10.254.79.254, _OCVPN2-1b, 00:10:15
B       10.2.200.0/24 [200/0] via 10.254.71.254, _OCVPN2-1a, 00:10:15
                      [200/0] via 10.254.79.254, _OCVPN2-1b, 00:10:15
B       10.254.0.0/16 [200/0] via 10.254.7.254, _OCVPN2-0a, 00:10:15
                      [200/0] via 10.254.15.254, _OCVPN2-0b, 00:10:15
                      [200/0] via 10.254.71.254, _OCVPN2-1a, 00:10:15
                      [200/0] via 10.254.79.254, _OCVPN2-1b, 00:10:15
C       10.254.0.0/21 is directly connected, _OCVPN2-0a
C       10.254.0.1/32 is directly connected, _OCVPN2-0a
C       10.254.8.0/21 is directly connected, _OCVPN2-0b
C       10.254.8.1/32 is directly connected, _OCVPN2-0b
C       10.254.64.0/21 is directly connected, _OCVPN2-1a
C       10.254.64.1/32 is directly connected, _OCVPN2-1b_0 <==shortcut tunnel
C       10.254.64.2/32 is directly connected, _OCVPN2-1a
C       10.254.72.0/21 is directly connected, _OCVPN2-1b
C       10.254.72.2/32 is directly connected, _OCVPN2-1b
                       is directly connected, _OCVPN2-1b_0
C       172.16.17.0/24 is directly connected, internal1
C       172.16.18.0/24 is directly connected, internal2
C       172.16.200.0/24 is directly connected, wan1
C       192.168.1.0/24 is directly connected, internal
C       192.168.4.0/24 is directly connected, wan2
B       192.168.5.0/24 [200/0] via 10.254.0.2, _OCVPN2-0a, 00:00:10
                       [200/0] via 10.254.8.2, _OCVPN2-0b, 00:00:10
                       [200/0] via 10.254.0.2, _OCVPN2-0a, 00:00:10
                       [200/0] via 10.254.8.2, _OCVPN2-0b, 00:00:10
                       [200/0] via 10.254.64.1, _OCVPN2-1b_0, 00:00:10
                       [200/0] via 10.254.72.1, _OCVPN2-1b, 00:00:10
                       [200/0] via 10.254.64.1, _OCVPN2-1b_0, 00:00:10
                       [200/0] via 10.254.72.1, _OCVPN2-1b, 00:00:10
C       192.168.44.0/24 is directly connected, loop1
B       192.168.55.0/24 [200/0] via 10.254.0.2, _OCVPN2-0a, 00:00:10
                        [200/0] via 10.254.8.2, _OCVPN2-0b, 00:00:10
                        [200/0] via 10.254.0.2, _OCVPN2-0a, 00:00:10
                        [200/0] via 10.254.8.2, _OCVPN2-0b, 00:00:10
                        [200/0] via 10.254.64.1, _OCVPN2-1b_0, 00:00:10
                        [200/0] via 10.254.72.1, _OCVPN2-1b, 00:00:10
                        [200/0] via 10.254.64.1, _OCVPN2-1b_0, 00:00:10
                        [200/0] via 10.254.72.1, _OCVPN2-1b, 00:00:10
```

2. Check the VPN tunnel state:

```
FGT_C # diagnose vpn tunnel list

list all ipsec tunnel in vd 0
------------------------------------------------------
name=_OCVPN2-1b_0 ver=2 serial=1c 172.16.18.3:0->172.16.15.4:0 dst_mtu=1500
bound_if=9 lgwy=static/1 tun=intf/0 mode=dial_inst/3 encap=none/728 options[02d8]=npu
create_dev no-sysctl rgwy-chg frag-rfc  accept_traffic=1 overlay_id=4

 parent=_OCVPN2-1b index=0
proxyid_num=1 child_num=0 refcnt=15 ilast=0 olast=0 ad=r/2
stat: rxp=641 txp=1025 rxb=16436 txb=16446
dpd: mode=on-idle on=1 idle=20000ms retry=3 count=0 seqno=0
natt: mode=none draft=0 interval=0 remote_port=0
proxyid=_OCVPN2-1b proto=0 sa=1 ref=3 serial=1 auto-negotiate adr
```


Netskope Exhibit 1017

```
  src: 0:0.0.0.0/0.0.0.0:0
  dst: 0:0.0.0.0/0.0.0.0:0
  SA:  ref=6 options=1a227 type=00 soft=0 mtu=1438 expire=42650/0B replaywin=1024
       seqno=407 esn=0 replaywin_lastseq=00000280 itn=0 qat=0 hash_search_len=1
  life: type=01 bytes=0/0 timeout=43186/43200
  dec: spi=90f03d9d esp=aes key=16 6cb33685bbc67d5c85488e0176ecf7b0
       ah=sha1 key=20 7d11b3babe62c840bf444b7b1f637b4324722a71
  enc: spi=7bc94bda esp=aes key=16 b4d8fc731d411eb24448b4077a5872ca
       ah=sha1 key=20 b724064d827304a6d80385ed4914461108b7312f
  dec:pkts/bytes=641/16368, enc:pkts/bytes=2053/123426
  npu_flag=03 npu_rgwy=172.16.15.4 npu_lgwy=172.16.18.3 npu_selid=1f dec_npuid=1 enc_
npuid=1
-----------------------------------------------------
name=_OCVPN2-0a ver=2 serial=18 172.16.17.3:0->172.16.13.1:0 dst_mtu=1500
bound_if=8 lgwy=static/1 tun=intf/0 mode=auto/1 encap=none/536 options[0218]=npu create_
dev frag-rfc  accept_traffic=1 overlay_id=1

proxyid_num=1 child_num=0 refcnt=20 ilast=0 olast=0 ad=r/2
stat: rxp=1665 txp=2922 rxb=278598 txb=70241
dpd: mode=on-idle on=1 idle=20000ms retry=3 count=0 seqno=7
natt: mode=none draft=0 interval=0 remote_port=0
proxyid=_OCVPN2-0a proto=0 sa=1 ref=4 serial=1 auto-negotiate adr
  src: 0:0.0.0.0/0.0.0.0:0
  dst: 0:0.0.0.0/0.0.0.0:0
  SA:  ref=6 options=1a227 type=00 soft=0 mtu=1438 expire=41599/0B replaywin=1024
       seqno=890 esn=0 replaywin_lastseq=00000680 itn=0 qat=0 hash_search_len=1
  life: type=01 bytes=0/0 timeout=42899/43200
  dec: spi=90f03d95 esp=aes key=16 a6ffcc197bb1b46ec745d0b595cdd69a
       ah=sha1 key=20 8007c134e41edf282f95daf9c9033d688ef05ccc
  enc: spi=a1bf21bf esp=aes key=16 ead05be389b0dec222f969e2f9c46b1d
       ah=sha1 key=20 b04105d34d4b0e61b018f2e60591f9b1510783bb
  dec:pkts/bytes=1665/278538, enc:pkts/bytes=4237/265074
  npu_flag=03 npu_rgwy=172.16.13.1 npu_lgwy=172.16.17.3 npu_selid=1b dec_npuid=1 enc_
npuid=1
-----------------------------------------------------
name=_OCVPN2-1a ver=2 serial=1a 172.16.17.3:0->172.16.11.1:0 dst_mtu=1500
bound_if=8 lgwy=static/1 tun=intf/0 mode=auto/1 encap=none/536 options[0218]=npu create_
dev frag-rfc  accept_traffic=1 overlay_id=3

proxyid_num=1 child_num=0 refcnt=17 ilast=0 olast=0 ad=r/2
stat: rxp=1 txp=2913 rxb=16376 txb=69642
dpd: mode=on-idle on=1 idle=20000ms retry=3 count=0 seqno=5
natt: mode=none draft=0 interval=0 remote_port=0
proxyid=_OCVPN2-1a proto=0 sa=1 ref=28 serial=1 auto-negotiate adr
  src: 0:0.0.0.0/0.0.0.0:0
  dst: 0:0.0.0.0/0.0.0.0:0
  SA:  ref=6 options=1a227 type=00 soft=0 mtu=1438 expire=41653/0B replaywin=1024
       seqno=887 esn=0 replaywin_lastseq=00000002 itn=0 qat=0 hash_search_len=1
  life: type=01 bytes=0/0 timeout=42900/43200
  dec: spi=90f03d9b esp=aes key=16 ee03f5b0f617a26c6177e91d60abf90b
       ah=sha1 key=20 f60cbbc4ebbd6d0327d23137da707b7ab2dc49e6
  enc: spi=a543a7d3 esp=aes key=16 1d37efab13a5c0347b582b2198b15cb8
       ah=sha1 key=20 427ee4c82bac6f26f0bcabfe04328c7f57ce682e
  dec:pkts/bytes=1/16316, enc:pkts/bytes=4229/264036
  npu_flag=03 npu_rgwy=172.16.11.1 npu_lgwy=172.16.17.3 npu_selid=1d dec_npuid=1 enc_
npuid=1
```


Netskope Exhibit 1017

```
---------------------------------------------------
name=_OCVPN2-0b ver=2 serial=19 172.16.18.3:0->172.16.14.1:0 dst_mtu=1500
bound_if=9 lgwy=static/1 tun=intf/0 mode=auto/1 encap=none/536 options[0218]=npu create_
dev frag-rfc  accept_traffic=1 overlay_id=2

proxyid_num=1 child_num=0 refcnt=20 ilast=0 olast=0 ad=r/2
stat: rxp=1665 txp=2917 rxb=278576 txb=69755
dpd: mode=on-idle on=1 idle=20000ms retry=3 count=0 seqno=7
natt: mode=none draft=0 interval=0 remote_port=0
proxyid=_OCVPN2-0b proto=0 sa=1 ref=4 serial=1 auto-negotiate adr
  src: 0:0.0.0.0/0.0.0.0:0
  dst: 0:0.0.0.0/0.0.0.0:0
  SA:  ref=6 options=1a227 type=00 soft=0 mtu=1438 expire=41599/0B replaywin=1024
       seqno=88b esn=0 replaywin_lastseq=00000680 itn=0 qat=0 hash_search_len=1
  life: type=01 bytes=0/0 timeout=42899/43200
  dec: spi=90f03d96 esp=aes key=16 9d7eb233c1d095b30796c3711d53f2fd
       ah=sha1 key=20 d8feacd42b5e0ba8b5e38647b2f2734c94644bd1
  enc: spi=a1bf21c0 esp=aes key=16 d2c0984bf86dc504c5475230b24034f0
       ah=sha1 key=20 3946e4033e1f42b0d9a843b94448f56fd5b57bee
  dec:pkts/bytes=1665/278516, enc:pkts/bytes=4233/264411
  npu_flag=03 npu_rgwy=172.16.14.1 npu_lgwy=172.16.18.3 npu_selid=1c dec_npuid=1 enc_
npuid=1
---------------------------------------------------
name=_OCVPN2-1b ver=2 serial=1b 172.16.18.3:0->172.16.12.1:0 dst_mtu=1500
bound_if=9 lgwy=static/1 tun=intf/0 mode=auto/1 encap=none/536 options[0218]=npu create_
dev frag-rfc  accept_traffic=1 overlay_id=4

proxyid_num=1 child_num=1 refcnt=19 ilast=1 olast=0 ad=r/2
stat: rxp=1 txp=2922 rxb=16430 txb=70173
dpd: mode=on-idle on=1 idle=20000ms retry=3 count=0 seqno=4
natt: mode=none draft=0 interval=0 remote_port=0
proxyid=_OCVPN2-1b proto=0 sa=1 ref=28 serial=1 auto-negotiate adr
  src: 0:0.0.0.0/0.0.0.0:0
  dst: 0:0.0.0.0/0.0.0.0:0
  SA:  ref=6 options=1a227 type=00 soft=0 mtu=1438 expire=41656/0B replaywin=1024
       seqno=890 esn=0 replaywin_lastseq=00000002 itn=0 qat=0 hash_search_len=1
  life: type=01 bytes=0/0 timeout=42903/43200
  dec: spi=90f03d9c esp=aes key=16 a655767c1ed6cff4575857eb3981ad81
       ah=sha1 key=20 bfc2bccd7103a201be2641d4c6147d437d2c3f70
  enc: spi=a543a7d4 esp=aes key=16 7221b814e483165b01edfdc8260d261a
       ah=sha1 key=20 d54819643c2f1b20da2aea4282d50a1f1bc1d72a
  dec:pkts/bytes=1/16370, enc:pkts/bytes=4238/265164
  npu_flag=03 npu_rgwy=172.16.12.1 npu_lgwy=172.16.18.3 npu_selid=1e dec_npuid=1 enc_
npuid=1
```

3. Check the SD-WAN state:

```
FGT_C # diagnose sys virtual-wan-link health-check
Health Check(Default_DNS):
Health Check(Default_Office_365):
Health Check(Default_Gmail):
Health Check(Default_AWS):
Health Check(Default_Google Search):
Health Check(Default_FortiGuard):
Health Check(ocvpn):
Seq(1 _OCVPN2-0a): state(alive), packet-loss(0.000%) latency(0.364), jitter(0.028) sla_
map=0x0
```


Netskope Exhibit 1017

```
Seq(2 _OCVPN2-0b): state(alive), packet-loss(0.000%) latency(0.287), jitter(0.026) sla_
map=0x0
Seq(3 _OCVPN2-1a): state(dead), packet-loss(100.000%) sla_map=0x0
Seq(4 _OCVPN2-1b): state(dead), packet-loss(100.000%) sla_map=0x0
Seq(4 _OCVPN2-1b_0): state(alive), packet-loss(0.000%) latency(0.289), jitter(0.029)
sla_map=0x0
```

DSCP matching (shaping)

This feature has three parts:

- DSCP matching in firewall policies
- DSCP matching in firewall shaping policies
- DSCP marking in firewall shaping policies

DSCP matching in firewall policies

Traffic is allowed or blocked according to the Differentiated Services Code Point (DSCP) values in the incoming packets.

The following CLI variables are available in the `config firewall policy` command:

`tos-mask <mask_value>`	Non-zero bit positions are used for comparison. Zero bit positions are ignored (default = 0x00). This variable replaces the `dscp-match` variable.
`tos <tos_value>`	Type of Service (ToC) value that is used for comparison (default = 0x00). This variable is only available when `tos-mask` is not zero. This variable replaces the `dscp-value` variable.
`tos-negate {enable | disable}`	Enable/disable negated ToS match (default = disable). This variable is only available when `tos-mask` is not zero. This variable replaces the `dscp-negate` variable.

DSCP matching in firewall shaping policies

Shaping is applied to the session or not according to the DSCP values in the incoming packets. The same logic and commands as in firewall policies are used.

DSCP marking in firewall shaping policies

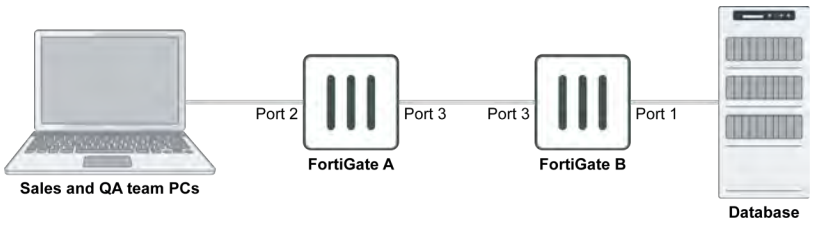
Traffic is allowed or blocked according to the DSCP values in the incoming packets. DSCP marking in firewall shaping policies uses the same logic and commands as in firewall policy and traffic-shaper.

When DSCP marking on `firewall shaper traffic-shaper`, `firewall shaping-policy`, and `firewall policy` all apply to the same session, `shaping-policy` overrides `policy`, and `shaper traffic-shaper` overrides both `shaping-policy` and `policy`.

The following CLI variables in `config firewall policy` are used to mark the packets:

`diffserv-forward {enable | disable}`	Enable/disable changing a packet's DiffServ values to the value specified in `diffservcode-forward` (default = disable).
`diffservcode-forward <dscp_value>`	The value that packet's DiffServ is set to (default = 000000). This variable is only available when `diffserv-forward` is enabled.
`diffserv-reverse {enable | disable}`	Enable/disable changing a packet's reverse (reply) DiffServ values to the value specified in `diffservcode-rev` (default = disable).
`diffservcode-rev <dscp_value>`	The value that packet's reverse (reply) DiffServ is set to (default = 000000). This variable is only available when `diffserv-rev` is enabled.

Examples



Example 1

FortiGate A marks traffic from the sales and QA teams with different DSCP values. FortiGate B does DSCP matching, allowing only the sales team to access the database.

1. Configure FortiGate A:

```
config firewall policy
    edit 1
        set srcintf "port2"
        set dstintf "port3"
        set srcaddr "QA"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ALL"
        set diffserv-forward enable
        set diffservcode-forward 110000
        set nat enable
    next
    edit 5
        set srcintf "port2"
        set dstintf "port3"
        set srcaddr "Sales"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ALL"
        set diffserv-forward enable
        set diffservcode-forward 111011
        set nat enable
```

```
    next
end
```

2. Configure FortiGate B:

```
config firewall policy
    edit 2
        set srcintf "port3"
        set dstintf "port1"
        set srcaddr "all"
        set dstaddr "Database"
        set action accept
        set schedule "always"
        set service "ALL"
        set tos-mask 0xf0
        set tos 0xe0
        set fsso disable
        set nat enable
    next
end
```

Example 2

FortiGate A marks traffic from the sales and QA teams with different DSCP values. FortiGate B uses a firewall shaping policy to do the DSCP matching, limiting the connection speed of the sales team to the database to 10MB/s.

1. Configure FortiGate A:

```
config firewall policy
    edit 1
        set srcintf "port2"
        set dstintf "port3"
        set srcaddr "QA"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ALL"
        set diffserv-forward enable
        set diffservcode-forward 110000
        set nat enable
    next
    edit 5
        set srcintf "port2"
        set dstintf "port3"
        set srcaddr "Sales"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ALL"
        set diffserv-forward enable
        set diffservcode-forward 111011
        set nat enable
    next
end
```


Netskope Exhibit 1017

2. Configure FortiGate B:

```
config firewall policy
    edit 2
        set srcintf "port3"
        set dstintf "port1"
        set srcaddr "all"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ALL"
        set nat enable
    next
end
config firewall shaper traffic-shaper
    edit "10MB/s"
        set guaranteed-bandwidth 60000
        set maximum-bandwidth 80000
    next
end
config firewall shaping-policy
    edit 1
        set service "ALL"
        set dstintf "port1"
        set tos-mask 0xf0
        set tos 0xe0
        set traffic-shaper "10MB/s"
        set srcaddr "all"
        set dstaddr "all"
    next
end
```

Example 3

FortiGate A has a traffic shaping policy to mark traffic from the QA team with a DSCP value of 100000, while reverse traffic is marked with 000011.

1. Configure FortiGate A:

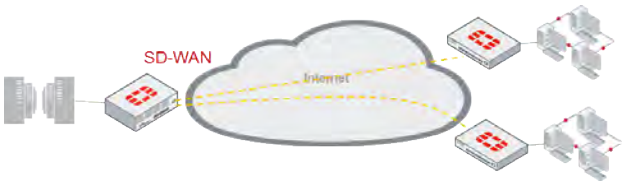
```
config firewall shaping-policy
    edit 1
        set name "QA Team 50MB"
        set service "ALL"
        set dstintf "port3"
        set traffic-shaper "50MB/s"
        set traffic-shaper-reverse "50MB/s"
        set diffserv-forward enable
        set diffserv-reverse enable
        set srcaddr "QA"
        set dstaddr "all"
        set diffservcode-forward 100000
        set diffservcode-rev 000011
    next
end
```


Netskope Exhibit 1017

SD-WAN health check packet DSCP marker support

SD-WAN health check probe packets support Differentiated Services Code Point (DSCP) markers for accurate evaluation of the link performance for high priority applications by upstream devices.

When the SD-WAN health check packet is sent out, the DSCP can be set with a CLI command.



To mark health-check packets with DSCP:

```
config system virtual-wan-link
   config health-check
      edit <name>
         set diffservcode <6 bits binary, range 000000-111111>
      next
   end
end
```

Dual VPN tunnel wizard

This wizard is used to automatically set up multiple VPN tunnels to the same destination over multiple outgoing interfaces. This includes automatically configuring IPsec, routing, and firewall settings, avoiding cumbersome and error-prone configuration steps.

To create a new SD-WAN VPN interface using the tunnel wizard:

1. Go to *Network* > *SD-WAN Interfaces* and click *Create New*.
2. In the *Interface* drop-down, click *+VPN*. The *Create IPsec VPN for SD-WAN members* pane opens.

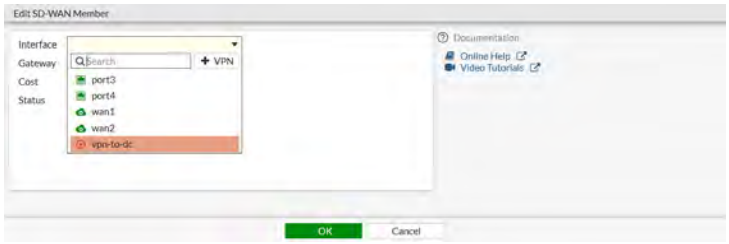


Netskope Exhibit 1017

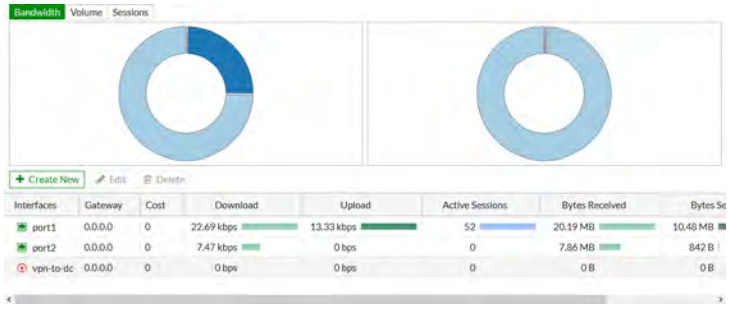
3. Enter the required information, then click *Create*.



4. Click *Close* to return to the SD-WAN page.
   The newly created VPN interface will be highlighted in the *Interface* drop-down list.



5. Select the VPN interface to add it as an SD-WAN member, then click *OK*.



Internet service customization

Internet Service Database (ISDB) entries can be tuned for their environments by adding custom ports and port ranges.

To add a custom port range:

```
config firewall internet-service-addition
    edit 65646
        set comment "Add custom port-range:tcp/8080-8090 into 65646"
        config entry
            edit 1
                set protocol 6
                config port-range
                    edit 1
                        set start-port 8080
                        set end-port 8090
                    next
                end
            next
        end
    next
end
```

```
Warning: Configuration will only be applied after rebooting or using the 'execute internet-
service refresh' command.

# execute internet-service refresh
Internet Service database is refreshed.
```

To verify that the change was applied:

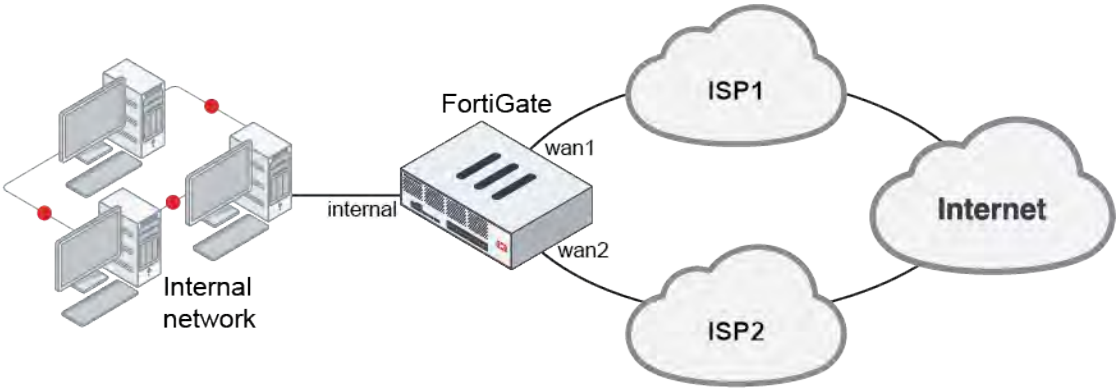
```
# diagnose internet-service info FG-traffic 6 8080 2.20.183.160
Internet Service: 65646(Google.Gmail)
```

SD-WAN with FGCP HA

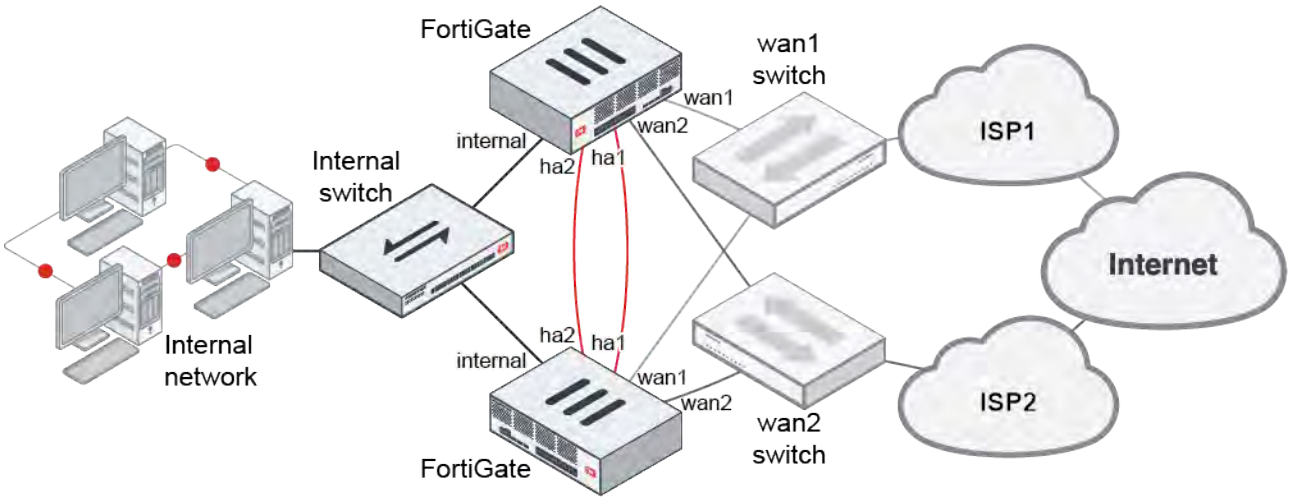
This example shows how to convert a standalone FortiGate SD-WAN solution to a FGCP HA cluster with full-mesh WAN set up. This configuration allows you to load balance your internet traffic between multiple ISP links. It also provides redundancy for your internet connection if your primary ISP in unavailable, or if one of the FortiGates in the HA cluster fails.

This example assumes that a standalone FortiGate has already been configured for SD-WAN by following the SD-WAN quick start on page 444.

Standalone FortiGate:



FGCP HA cluster:



The following devices are required to convert the topology to HA:

- A second FortiGate that is the same model running the same firmware version.
- Two switches for connecting each FortiGate's WAN interface to the corresponding ISP modem.

Before you begin:

- Ensure that the licenses and subscriptions on both HA members match.
- Ensure that there are one or more ports reserved for HA heartbeat.
- Ensure you have physical access to both HA members.

Enabling HA and re-cabling the WAN interfaces will cause network interruptions. This procedure should be performed during a maintenance window.

Configuring the standalone FortiGate for HA

After running the following commands, the FortiGate negotiates to establish an HA cluster. You might temporarily lose connectivity with the FortiGate as FGCP negotiations take place and the MAC addresses of the FortiGate interfaces are changed to HA virtual MAC addresses.

This configurations sets the HA mode to active-passive.

The ha1 and ha2 interfaces are configured as the heartbeat interfaces, with priorities set to 200 and 100 respectively. Setting different priorities for the heartbeat interfaces is a best practice, but is not required.

If you have more than one cluster on the same network, each cluster should have a different group ID. Changing the group ID changes the cluster interface's virtual MAC addresses. If the group IP causes a MAC address conflict on your network, select a different group ID.

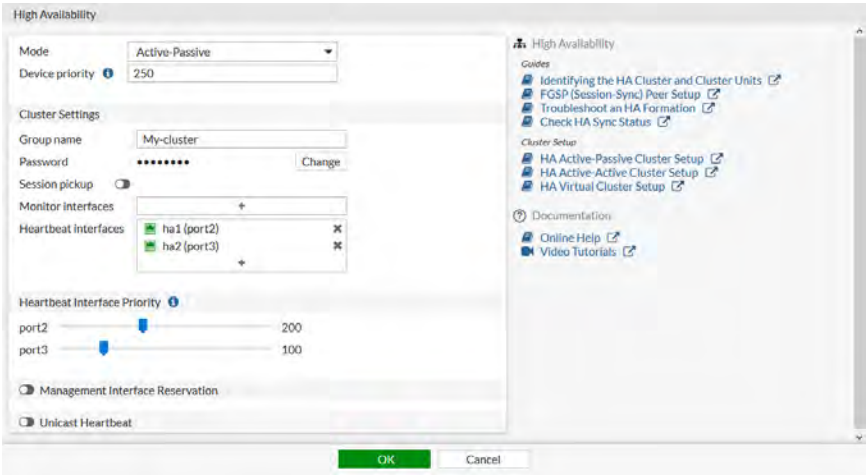
Enabling override and increasing the device priority means that this FortiGate always becomes the primary unit.

To configure the standalone FortiGate for HA in the GUI:

1. Go to *System* > *Settings* and change the *Host name* so that the FortiGate can be easily identified as the primary unit.
2. Go to *System* > *HA* and configure the following options:

Mode	Active-Passive
Device priority	250
Group name	My-cluster
Password	<password>
Heartbeat interfaces	ha1 and ha2
Heartbeat Interface Priority	**port2** (ha1): 200 **port3** (ha2): 100

Override and the group ID can only be configured from the CLI.



3. Click *OK*.
   Connectivity with the FortiGate will temporarily be lost.

To configure the standalone FortiGate for HA in the CLI:

1. Change the host name so that the FortiGate can be easily identified:

```
config system global
    set hostname primary_FG
end
```

2. Configure HA:

```
config system ha
    set mode a-p
    set group-id 100
    set group-name My-cluster
```

```
    set password <password>
    set priority 250
    set override enable
    set hbdev ha1 200 ha2 100
end
```



If HA mode does not start after running the above steps, ensure that none of the FortiGate's interfaces use DHCP or PPPoE addressing.

Configuring the secondary FortiGate for HA

The secondary FortiGate must be the same model and running the same firmware version as the primary FortiGate. The HA settings are the same as the for the primary unit, except the secondary device has a lower priority and override is not enabled.



It is best practice to reset the FortiGate to factory default settings prior to configuring HA. This reduces the chance of synchronization problems.

```
# execute factoryreset
This operation will reset the system to factory default!
Do you want to continue? (y/n) y
```

This is unnecessary if the device is new from the factory.

To configure the secondary FortiGate for HA in the GUI:

1. Go to *System* > *Settings* and change the *Host name* so that the FortiGate can be easily identified as the backup unit.
2. Go to *System* > *HA* and configure the options the same as for the primary FortiGate, except with a lower priority:

Mode	Active-Passive
Device priority	128
Group name	My-cluster
Password	<password>
Heartbeat interfaces	ha1 and ha2
Heartbeat Interface Priority	**port2** (ha1): 200 **port3** (ha2): 100

3. Click *OK*.

To configure the secondary FortiGate for HA in the CLI:

1. Change the host name so that the secondary FortiGate can be easily identified:

```
config system global
    set hostname secondary_FG
end
```

2. Configure HA:

```
config system ha
    set mode a-p
    set group-id 100
    set group-name My-cluster
    set password <password>
    set priority 128
    set hbdev ha1 200 ha2 100
end
```

Connecting the heartbeat interfaces between the FortiGates

To connect and check the heartbeat interfaces:

1. Connect the heartbeat interfaces ha1 and ha2 between the primary and secondary FortiGate.
   a. An HA primary is selected. Because the primary FortiGate has a higher priority and override enabled, it assumes the role of HA primary.
   b. The secondary FortiGate synchronizes its configuration from the primary device.
2. Verify that the checksums match between the primary and secondary FortiGates:

   ```
   # diagnose sys ha checksum cluster
   ```

   If all of the cluster members have identical checksums, then their configurations are synchronized. If the checksums are not the same, wait for a few minutes, then repeat the command. Some parts of the configuration might take a significant amount of time to synchronize (tens of minutes).

Connecting other traffic interfaces

After the device configurations are synchronized, you can connect the rest of the traffic interfaces. Making these connections will disrupt traffic as cables are disconnected and reconnected.

Switches must be used between the cluster and the ISPs, and between the cluster and the internal network, as shown in the topology diagram.

Checking cluster operations

The *HA Status* dashboard widget shows the synchronization status. Hover over the host names of each FortiGate in the widget to verify that they are synchronized and have the same checksum.

To view more information about the cluster status, including the number of sessions passing through the cluster members, go to *System* > *HA*.

See Check HA sync status on page 643 for more information.

Results

1. Browse the internet on a computer in the internal network.
2. Go to *Network* > *SD-WAN Interfaces* to see the bandwidth, volume, and sessions for traffic on the SD-WAN interfaces. See Results on page 449 for details.

Netskope Exhibit 1017

3. Go to *Dashboard* > *Network*, and expand the *SD-WAN* widget to see information about each interface, such as the number of sessions and the bit rate.



Testing HA failover

All traffic should currently be flowing through the primary FortiGate. If it becomes unavailable, traffic fails over to the secondary FortiGate. When the primary FortiGate rejoins the cluster, the secondary FortiGate continues to operate as the primary FortiGate.

To test this, ping a reliable IP address from a computer in the internal network, and then power off the primary FortiGate.

There will be a momentary pause in the ping results until traffic diverts to the backup FortiGate, allowing the ping traffic to continue:

```
64 bytes from 184.25.76.114: icmp_seq=69 ttl=52 time=8.719 ms\
64 bytes from 184.25.76.114: icmp_seq=70 ttl=52 time=8.822 ms\
64 bytes from 184.25.76.114: icmp_seq=74 ttl=52 time=8.901 ms\
Request timeout for icmp_seq 75\
64 bytes from 184.25.76.114: icmp_seq=76 ttl=52 time=8.860 ms\
64 bytes from 184.25.76.114: icmp_seq=77 ttl=52 time=9.174 ms\
64 bytes from 184.25.76.114: icmp_seq=83 ttl=52 time=8.639 ms}
```



If you are using port monitoring, you can also unplug the primary FortiGate's internet facing interface to test failover.

After the secondary FortiGate becomes the primary, you can log into the cluster using the same IP address as before the fail over. If the primary FortiGate is powered off, you will be logged into the backup FortiGate. Check the host name to verify what device you have logged into. The FortiGate continues to operate in HA mode, and if you restart the primary FortiGate, it will rejoin the cluster and act as the backup FortiGate. Traffic is not disrupted when the restarted FortiGate rejoins the cluster.

You can also use the CLI to force an HA failover. See Force HA failover for testing and demonstrations on page 661 for information.

Testing ISP failover

To test a failover of the redundant internet configuration, you need to simulate a failed internet connection to one of the ports. You can do this by disconnecting power from the wan1 switch, or by disconnecting the wan1 interfaces of both FortiGates from ISP1.

After disconnecting, verify that users still have internet access

- Go to *Dashboard* > *Network*, and expand the *SD-WAN* widget. The *Upload* and *Download* columns for wan1 show that traffic is not going through that interface.



- Go to *Network* > *SD-WAN Interfaces*. The Bandwidth, Volume, and Sessions tabs show that traffic is entirely diverted to wan2.

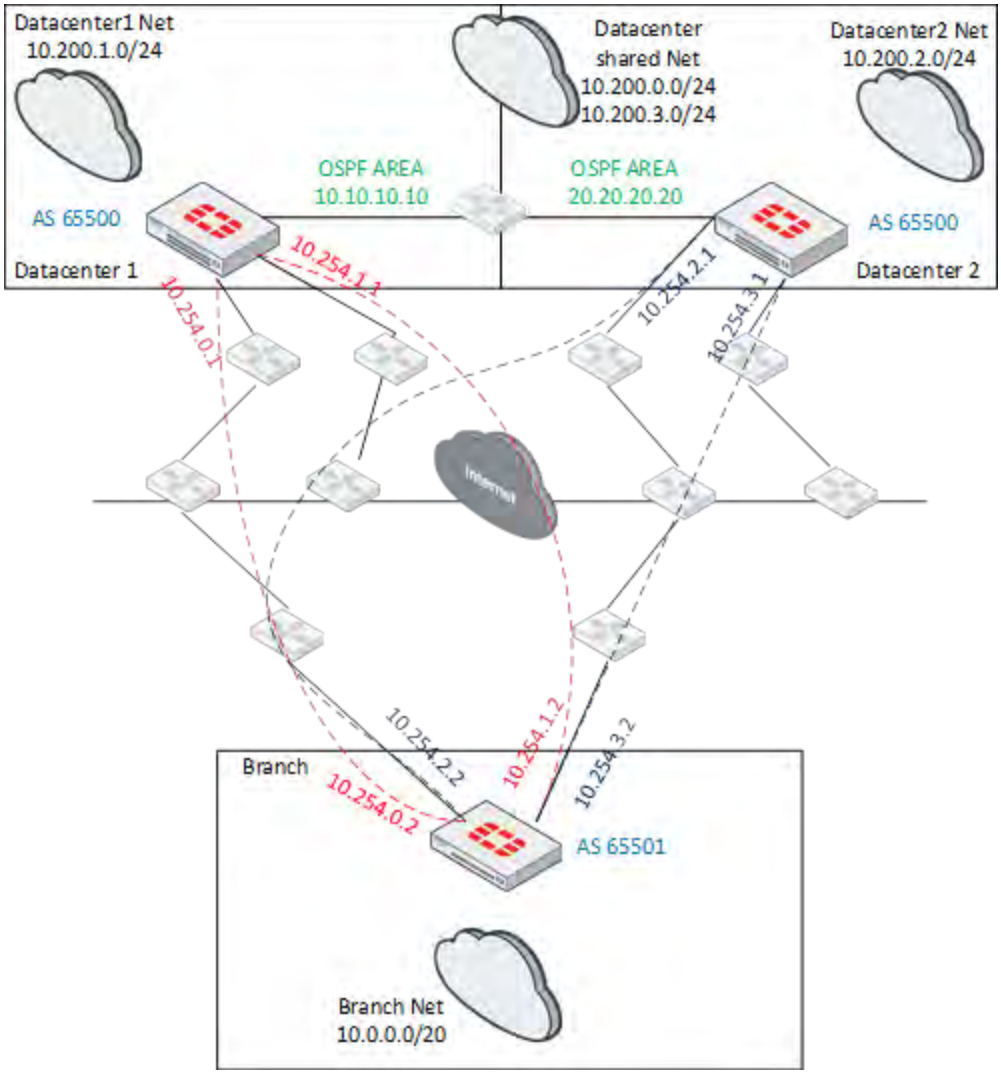


Users on the network should not notice the wan1 failure. If you are using the wan1 gateway IP address to connect to the administrator dashboard, it will appear as though you are still connecting through wan1.

After verifying a successful failover, reestablish the connection to ISP1.

Hub and spoke SD-WAN deployment example

This topology diagram shows an overview of the network that is configured in this example:



Datacenter configuration

The datacenter is configured to support:

- Zero touch provisioning of new spokes
- Point to multipoint VPN
- Central management of access with the datacenter firewall
- Dynamic peering, to share routing information between branches and the datacenter
- VDOM compatibility, with inter-VDOM links for isolation and segmentation

To configure the datacenter, complete the following steps:

1. Configure dial-up (dynamic) VPN
2. Configure VPN interfaces
3. Configure loopback interface
4. Configure BGP
5. Firewall policies
6. Configure a black hole route

Configure dial-up (dynamic) VPN

Dial-up, or dynamic, VPNs are used to facilitate zero touch provisioning of new spokes to establish VPN connections to the hub FortiGate.

The `exchange-interface-ip` option is enabled to allow the exchange of IPsec interface IP addresses. This allows a point to multipoint connection to the hub FortiGate.

The `add-route` option is disabled to allow multiple dial-up tunnels to be established to the same host that is advertising the same network. This dynamic network discovery is facilitated by the BGP configuration; see Configure BGP on page 552 for details.

Wildcard security associations are defined for the phase2 interface because routing is used to determine if traffic is subject to encryption and transmission through the IPsec VPN tunnel. The phase1 interface name must be 11 characters or less.

A dynamic VPN configuration must be defined for each interface that connects to the internet.

To configure the IPsec phase1 interfaces:

```
config vpn ipsec phase1-interface
    edit "vpn-isp-a"
        set type dynamic
        set interface "port2"
        set peertype any
        set exchange-interface-ip enable
        set proposal aes256-sha256
        set add-route disable
        set dhgrp 5
        set net-device enable
        set psksecret ********
    next
    edit "vpn-isp-b"
```


Netskope Exhibit 1017

```
        set type dynamic
        set interface "port3"
        set peertype any
        set exchange-interface-ip enable
        set proposal aes256-sha256
        set add-route disable
        set dhgrp 5
        set net-device enable
        set psksecret ********
    next
end
```

To configure the IPsec phase2 interfaces:

```
config vpn ipsec phase2-interface
    edit "vpn-isp-a_p2"
        set phase1name "vpn-isp-a"
        set proposal aes256-sha256
        set pfs disable
        set replay disable
    next
    edit "vpn-isp-b_p2"
        set phase1name "vpn-isp-b"
        set proposal aes256-sha256
        set pfs disable
        set replay disable
    next
end
```

Configure VPN interfaces

To establish the BGP session, IP addresses must be assigned to the tunnel interfaces that BGP will use to peer.

The hub IP address is set to the address that the tunnels connect to. The remote IP address is set to highest unused IP address that is part of the tunnel network. This establishes two connected routes directly back to the branch FortiGate in the hub FortiGate's routing table.

Ping is allowed on the virtual interface to confirm that a point to point tunnel has been established between the hub and branch FortiGates.

To define IP addresses for VPN interfaces:

```
config system interface
    edit "vpn-isp-a"
        set vdom "root"
        set ip 10.254.0.1 255.255.255.255
        set allowaccess ping
        set type tunnel
        set remote-ip 10.254.0.254 255.255.255.0
        set interface "port2"
    next
    edit "vpn-isp-b"
        set vdom "root"
        set ip 10.254.1.1 255.255.255.255
        set allowaccess ping
```

```
        set type tunnel
        set remote-ip 10.254.1.254 255.255.255.0
        set interface "port3"
    next
end
```

Configure loopback interface

A loopback interface must be defined on the hub FortiGate to be used as a common probe point for the FortiGates that are using SD-WAN. The FortiGates send a probe packet from each of their SD-WAN member interfaces so that they can determine the best route according to their policies. Ping is allowed so that it can be used for measurements.

To configure the loopback interface on the hub FortiGate:

```
config system interface
    edit "loopback_0"
        set vdom "root"
        set ip 10.255.255.1 255.255.255.255
        set allowaccess ping
        set type loopback
    next
end
```

Configure BGP

Network route discovery is facilitated by BGP.

EBGP is used to prevent the redistribution of routes that are in the same Autonomous System (AS) number as the host. It is also required to influence route selection on the branches with AS-Path prepending. EBGP multipath is enabled so that the hub FortiGate can dynamically discover multiple paths for networks that are advertised at the branches.

The neighbor range and group settings are configured to allow peering relationships to be established without defining each individual peer. Connecting branches have their tunnel interfaces configured within the range of the BGP peer.

In order to facilitate the fastest route failovers, configure the following timers to their lowest levels: `scan-time`, `advertisement-interval`, `keep-alive-timer`, and `holdtime-timer`.

To configure BGP on the hub FortiGate:

```
config router bgp
    set as 65500
    set router-id 10.10.0.1
    set ebgp-multipath enable
    set graceful-restart enable
    config neighbor-group
        edit "branch-peers-1"
            set soft-reconfiguration enable
            set remote-as 65501
        next
        edit "branch-peers-2"
            set soft-reconfiguration enable
            set remote-as 65501
        next
```

```
    end
    config neighbor-range
        edit 1
            set prefix 10.254.0.0 255.255.255.0
            set neighbor-group "branch-peers-1"
        next
        edit 2
            set prefix 10.254.1.0 255.255.255.0
            set neighbor-group "branch-peers-2"
        next
    end
    config network
        edit 1
            set prefix 10.200.1.0 255.255.255.0
        next
        edit 2
            set prefix 10.200.0.0 255.255.255.0
        next
        edit 3
            set prefix 10.200.3.0 255.255.255.0
        next
    end
end
```

Firewall policies

Centralized access is controlled from the hub FortiGate using Firewall policies. In addition to layer three and four inspection, security policies can be used in the policies for layer seven traffic inspection.

It is best practice to only allow the networks and services that are required for communication through the firewall. The following rules are the minimum that must be configured to allow SD-WAN to function:

Source Interface	Destination Interface	Source Address	Destination Address	Action	Schedule	Service	Comments
<vpn interfaces>	<internal Interface>	<branch tunnel IP addresses>	<hub FortiGate internal interface>	Accept	Always	ICMP	Allow health checks to the hub FortiGate
<vpn interfaces>	<internal Interface>	<branch networks>	<datacenter networks>	Accept	Always	<allowed services>	Allow traffic from branch networks

For this example, a simple policy that allows all traffic is configured.

To configure a firewall policy:

```
config firewall policy
    edit 1
        set name "Allow All"
        set srcintf "any"
        set dstintf "any"
        set srcaddr "all"
```


Netskope Exhibit 1017

```
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ALL"
    next
end
```

Configure a black hole route

If there is a temporary loss of connectivity to the branch routes, it is best practice to send the traffic that is destined for those networks into a black hole until connectivity is restored.

To configure a black hole route for branch networks:

```
config router static
    edit 6
        set dst 10.0.0.0/14
        set distance 254
        set blackhole enable
    next
end
```

Branch configuration

The branches are configured to support:

- Client side SD-WAN with intelligent load balancing based on link quality
- Easy to create configuration templates for quick spoke deployment
- Split tunnel deployment for local internet access
- VDOM compatibility, with inter-VDOM links for isolation and segmentation

To configure a branch, complete the following steps:

1. Configure VPN to the hub
2. Configure VPN interfaces
3. Configure BGP
4. Configure SD-WAN
5. Firewall configuration

Configure VPN to the hub

The branch uses a normal site-to-site VPN configuration.

Wildcard security associations are define in the phase2 configuration because dynamic routing with BGP determines what traffic must traverse the VPN tunnel for encryption/transmission.

To make sure that the VPN is established, `auto-negotiate` is enabled.

To configure the IPsec phase1 interfaces:

```
config vpn ipsec phase1-interface
    edit "vpn_dc1-1"
        set interface "port2"
        set peertype any
        set exchange-interface-ip enable
        set proposal aes256-sha256
        set dhgrp 5
        set remote-gw 172.16.0.78
        set psksecret ********
    next
    edit "vpn_dc1-2"
        set interface "port3"
        set peertype any
        set exchange-interface-ip enable
        set proposal aes256-sha256
        set dhgrp 5
        set remote-gw 172.16.0.82
        set psksecret ********
    next
end
```

To configure the IPsec phase2 interfaces:

```
config vpn ipsec phase2-interface
    edit "vpn_dc1-1_p2"
        set phase1name "vpn_dc1-1"
        set proposal aes256-sha256
        set pfs disable
        set replay disable
        set auto-negotiate enable
    next
    edit "vpn_dc1-2_p2"
        set phase1name "vpn_dc1-2"
        set proposal aes256-sha256
        set pfs disable
        set replay disable
        set auto-negotiate enable
    next
end
```

Configure VPN interfaces

The branch must define its local tunnel interface IP address, and the remote tunnel interface IP address of the datacenter FortiGate, to establish the point to multipoint VPN.

To define IP addressses for VPN interfaces:

```
config system interface
    edit "vpn_dc1-1"
        set vdom "root"
        set ip 10.255.0.2 255.255.255.255
        set allowaccess ping
```


Netskope Exhibit 1017

```
        set type tunnel
        set remote-ip 10.255.0.1 255.255.255.255
        set interface "port2"
    next
    edit "vpn_dc1-2"
        set vdom "root"
        set ip 10.255.1.2 255.255.255.255
        set allowaccess ping
        set type tunnel
        set remote-ip 10.255.1.1 255.255.255.255
        set interface "port3"
    next
end
```

Configure BGP

BGP enables learning dynamic routes from the datacenter. The BGP configuration is normal, with the definition of the datacenter FortiGate tunnel IP addresses set as BGP peers.

Routes that have the same network mask, administrative distance, priority, and AS length are automatically considered for SD-WAN when the interfaces that those routes are on are added to the SD-WAN interface group.

In order to facilitate the fastest route failovers, configure the following timers to their lowest levels: `scan-time`, `advertisement-interval`, `keep-alive-timer`, and `holdtime-timer`.

The `distance-external` option might need to be configured if you need routes that are learned from BGP to take precedence over static routes.

To configure BGP on the branch FortiGate:

```
config router bgp
    set as 65501
    set router-id 10.254.0.2
    set ebgp-multipath enable
    config neighbor
        edit "10.254.0.1"
            set soft-reconfiguration enable
            set remote-as 65500
        next
        edit "10.254.1.1"
            set soft-reconfiguration enable
            set remote-as 65500
        next
    end
end
```

Configure SD-WAN

SD-WAN configuration is required to load balance based on the quality of the links. It can be configured to select the best link based on characteristics such as jitter, packet loss, and latency. A policy route is created by the FortiGate to select the best link based on the defined criteria.

For SD-WAN interfaces, or members, the peer is defined to reference the BGP neighbor that is tied to that specific interface.

Netskope Exhibit 1017

The health check is the ping server that gathers the link characteristics used for link selection. It is recommended that the minimum `failtime` be set to `2`.

The service definition defines the criteria for the policy routes. It can match based on the following characteristics:

- Protocol
- Destination Address
- Source Address
- Identity Based Group
- Internet Service Definition
- Source Port
- Destination Port
- Destination Route Tag

To dynamically determine the networks of the policy routes, routes that are learned from a BGP neighbor are matched against a route map, and a tag is defined for the matching routes. The service rules learn the networks based on these tags, instead of defining objects based on the learned addresses' network prefixes . See Dynamic definition of SD-WAN routes on page 559 for details on configuring the FortiGate to use the destination tags for the SD-WAN service definition.

To define the SD-WAN member interfaces:

```
config system virtual-wan-link
    set status enable
    config members
        edit 1
            set interface "vpn_dc1-1"
        next
        edit 2
            set interface "vpn_dc1-2"
        next
    end
end
```

To define the SD-WAN health checks:

```
config system virtual-wan-link
    config health-check
        edit "datacenter1"
            set server "10.200.1.1"
            set interval 1
            set failtime 2
            set recoverytime 10
        next
    end
end
```

To define the SD-WAN service rules:

```
config system virtual-wan-link
    config service
        edit 1
            set mode priority
            set dst n-corporate
            set health-check "datacenter1"
```

```
            set priority-members 1 2
        next
    end
end
```

Firewall configuration

Centralized access is controlled from the hub FortiGate using Firewall policies. In addition to layer three and four inspection, security policies can be used in the policies for layer seven traffic inspection.

It is best practice to only allow the networks and services that are required for communication through the firewall. The following rules are the minimum that must be configured to allow SD-WAN to function:

Source Interface	Destination Interface	Source Address	Destination Address	Action	Schedule	Service	Comments
<internal interface>	<virtual wan link>	<branch networks>	<datacenter networks>	Accept	Always	<allowed services>	Allow traffic from branch to datacenter
<virtual wan link>	<internal Interface>	<datacenter networks>	<branch networks>	Accept	Always	<allowed services>	Allow traffic from datacenter to branch

For this example, a simple policy that allows all traffic is configured.

To configure a firewall policy:

```
config firewall policy
    edit 1
        set name "Allow All"
        set srcintf "any"
        set dstintf "any"
        set srcaddr "all"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ALL"
    next
end
```

Validation

The following commands can be used to validate the connections on the datacenter and branches.

Datacenter

Routing table:

```
get router info routing-table all
```

VPN establishment:

```
diagnose vpn ike gateway list
```

Branch

SD-WAN validation:

```
diagnose sys virtual-wan-link member
diagnose sys virtual-wan-link service
diagnose sys virtual-wan-link health-check
```

Routing table:

```
get router info routing-table all
get router info route-map-address
get router info bgp route-map <route-map-name>
```

VPN establishment:

```
diagnose vpn ike gateway list
```

Dynamic definition of SD-WAN routes

Dynamic definitions of SD-WAN routes alleviate administrators from needing to know the destination of the traffic that is being load balanced, which, in an environment where routes are constantly added and removed, required a significant amount of administrative overhead.

The FortiGate can be configured to apply a route map to a BGP neighbor, and tag the routes that are learned from that neighbor with the `set-route-tag` command. After those routes are assigned a tag ID in the route map, the ID can be referenced in the SD-WAN rule.

To define the route map to apply to the BGP neighbor:

```
config router route-map
    edit "map-comm1"
        config rule
            edit 1
                set match-origin igp
                set set-route-tag 12
            next
            edit 2
                set match-ip-address "pf-all-in"
                set set-route-tag 11
            next
        end
    next
end
```

To apply the route map to the BGP neighbor:

```
config router bgp
    config neighbor
        edit "10.254.0.1"
            set route-map-in "map-comm1"
        next
    end
end
```

To reference tagged routes in an SD-WAN rule:

```
config system virtual-wan-link
    config service
        edit 1
            set mode priority
            set dst-tag 11
            set health-check "datacenter1"
            set priority-members 1 2
        next
    end
end
```

Adding another datacenter

Datacenter FortiGates should be configured to establish an OSPF neighbor relationship with the internal core router. This allows the dynamic redistribution of routes to the branches that are receiving updates from the datacenter FortiGates.

To ensure the fastest failover with OSPF, the following timers are set to their minimum levels: `spf-timers`, `hello-interval`, `dead-interval`.

Bi-directional forwarding is enabled to allow the fastest convergence time if there is a failure with a peering neighbor.

To configure OSPF:

```
config router ospf
    set router-id 10.10.10.10
    set spf-timers 0 1
    set distribute-list-in "pf-datacenter2-tunnel"
    set restart-mode graceful-restart
    config area
        edit 10.10.10.10
        next
    end
    config ospf-interface
        edit "port5"
            set interface "port5"
            set dead-interval 3
            set hello-interval 1
            set bfd enable
        next
    end
    config network
```


Netskope Exhibit 1017

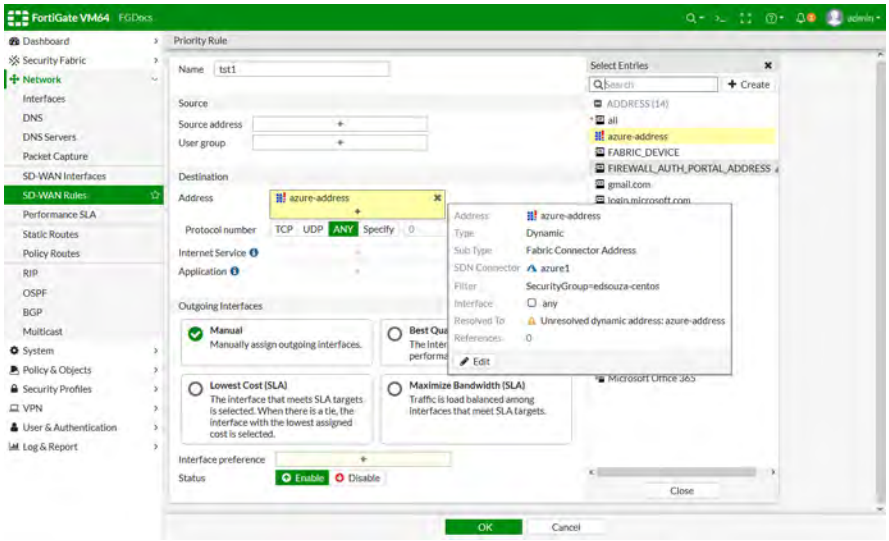
```
        edit 1
            set prefix 192.168.100.0 255.255.255.252
            set area 10.10.10.10
        next
    end
    config redistribute "connected"
        set status enable
        set routemap "redistribute-branch-tunnel"
    end
    config redistribute "static"
    end
    config redistribute "rip"
    end
    config redistribute "bgp"
        set status enable
        set routemap "redistribute-branch-networks"
    end
    config redistribute "isis"
    end
end
```

Dynamic connector addresses in SD-WAN policies

Dynamic Fabric Connector addresses can be used in SD-WAN policies.

To add a dynamic connector address to an SD-WAN rule:

1. Go to *Network* > *SD-WAN Rules*.
2. Click *Create New*.
3. In the *Name* field, enter the rule name.
4. In the *Destination* > *Address* field, click the +. The *Select Entries* pane will appear.
5. In the *Select Entries* pane, select the dynamic connector.

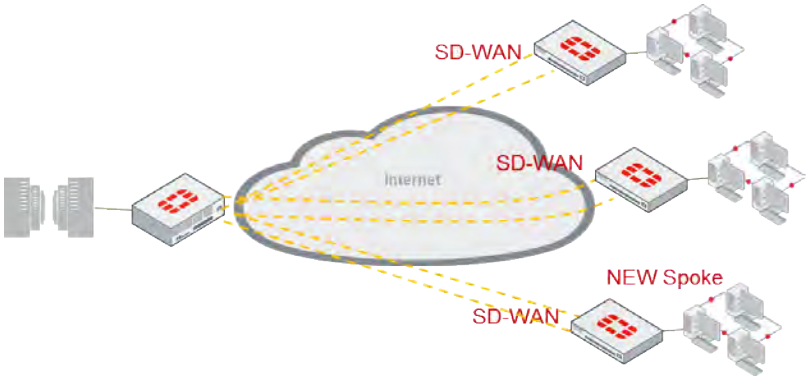


6. Configure the remaining settings as required.
7. Click *OK* to save the rule.

SD-WAN configuration portability

When configuring SD-WAN, adding interfaces to members is optional. This allows a configuration to be copied directly from one device to another, without requiring the devices to have interfaces with the same names.

After the configuration is pasted to the new device, add the interfaces to the new device to make it fully functional.



To copy the SD-WAN configuration from an existing spoke to a new spoke:

1. Copy the configuration from the configured spoke:

```
config system virtual-wan-link
    set status enable
    config members
        edit 1
            set interface "_OCVPN3-0.0"
        next
        edit 2
            set interface "_OCVPN3-0.1"
        next
    end
    config health-check
        edit "office"
            set server "office365.com"
            set protocol http
            set sla-fail-log-period 300
            set sla-pass-log-period 300
            set members 2 1
            config sla
                edit 1
                    set latency-threshold 300
                    set jitter-threshold 200
                next
                edit 2
                    set link-cost-factor latency
                    set latency-threshold 20
```

```
                next
            end
        next
        ...
    end
    config service
        edit 2
            set name "Office365"
            set mode sla
            set internet-service enable
            set internet-service-app-ctrl 327782
            config sla
                edit "office"
                    set id 1
                next
            end
            set priority-members 2 1
        next
        ...
    end
end
```

2. Paste the configuration onto the new spoke:

```
config system virtual-wan-link
    set status enable
    config members
        edit 1
        next
        edit 2
        next
    end
    config health-check
        edit "office"
            set server "office365.com"
            set protocol http
            set sla-fail-log-period 300
            set sla-pass-log-period 300
            set members 2 1
            config sla
                edit 1
                    set latency-threshold 300
                    set jitter-threshold 200
                next
                edit 2
                    set link-cost-factor latency
                    set latency-threshold 20
                next
            end
        next
        ...
    end
    config service
        edit 2
            set name "Office365"
            set mode sla
```


Netskope Exhibit 1017

```
            set internet-service enable
            set internet-service-app-ctrl 327782
            config sla
                edit "office"
                    set id 1
                next
            end
            set priority-members 2 1
        next
        ...
    end
end
```

The member interfaces are not copied over. Already configured interfaces are not unset. The member is disabled until an interface is configured.

3. Configure the member interfaces on the new spoke:

```
config system virtual-wan-link
    config members
        edit 1
            set interface "_OCVPN4-0.0"
        next
        edit 2
            set interface "_OCVPN4-0.1"
        next
    end
end
```

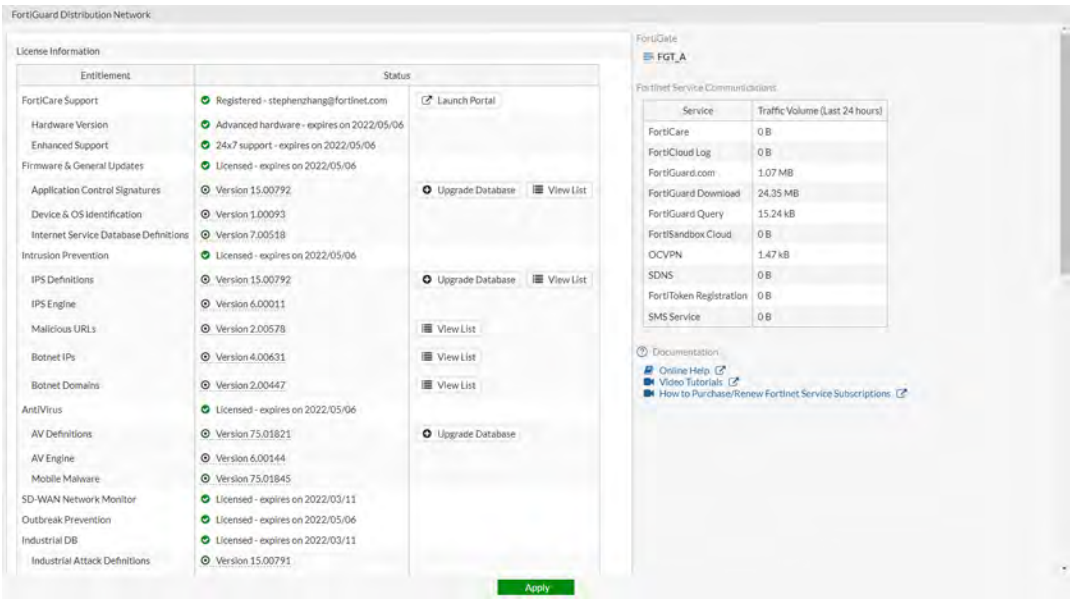
After the interfaces are configured, the new spoke will function like the other spokes.

Interface speedtest

An interface speedtest can be performed on WAN interfaces in the GUI. The results of the test can be added to the interface's *Estimated bandwidth*. An SD-WAN Network Monitor license is required.

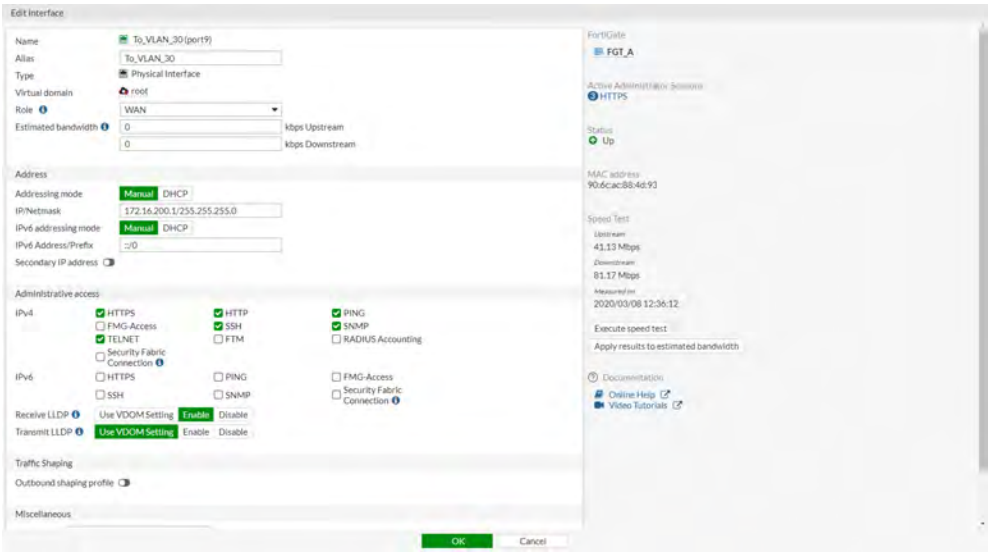
The *License* widget and the *System* > *FortiGuard* page display the SD-WAN Network Monitor license status.



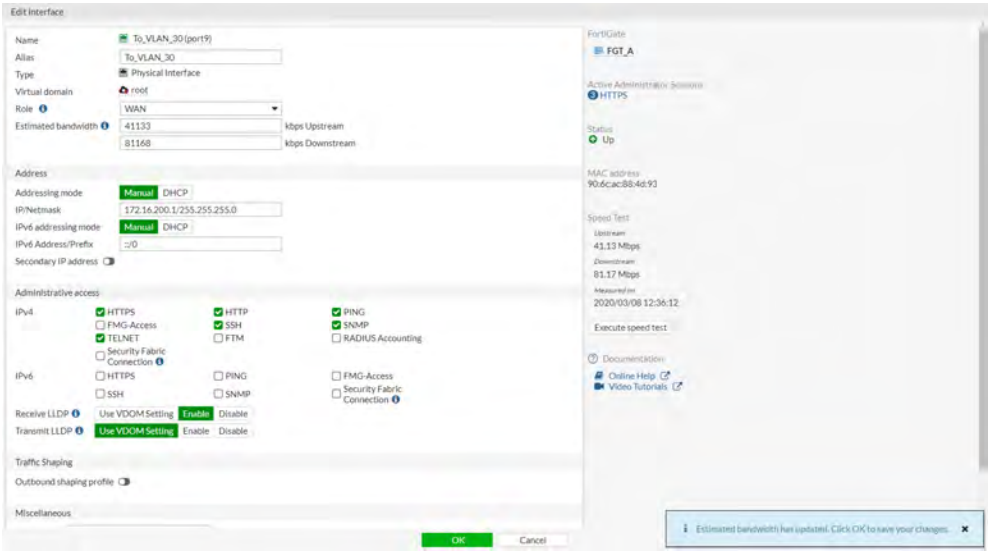


To run an interface speedtest in the GUI:

1. Go to *Network* > *Interfaces*.
2. Edit a WAN interface. The interfaces can be grouped by role using the grouping dropdown on the right side of the toolbar.
3. Click *Execute speed test* in the right pane.



4. When the test completes, click *Apply results to estimated bandwidth*.



The speedtest results are used to populate the *Estimated bandwidth* fields.

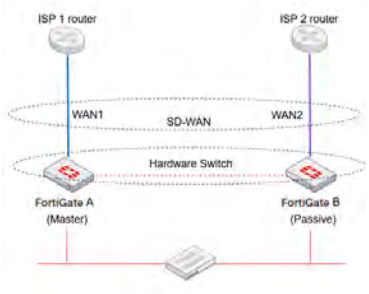
5. Click *OK*.

The FortiGate must be connected to FortiGuard, and able to reach either the AWS or Google speedtest servers.

Configuring SD-WAN in an HA cluster using internal hardware switches

This setup is not fully compliant with a regular HA configuration. Failover is also unnecessary. Flipping can occur if failover is configured using a ping server interface.

Two FortiGates with internal hardware switches can be configured as an active-active (A-A) HA pair. In the following topology, both FortiGates forward traffic through internal switches connected to service providers. SD-WAN is configured on all upstream interfaces and overlays.



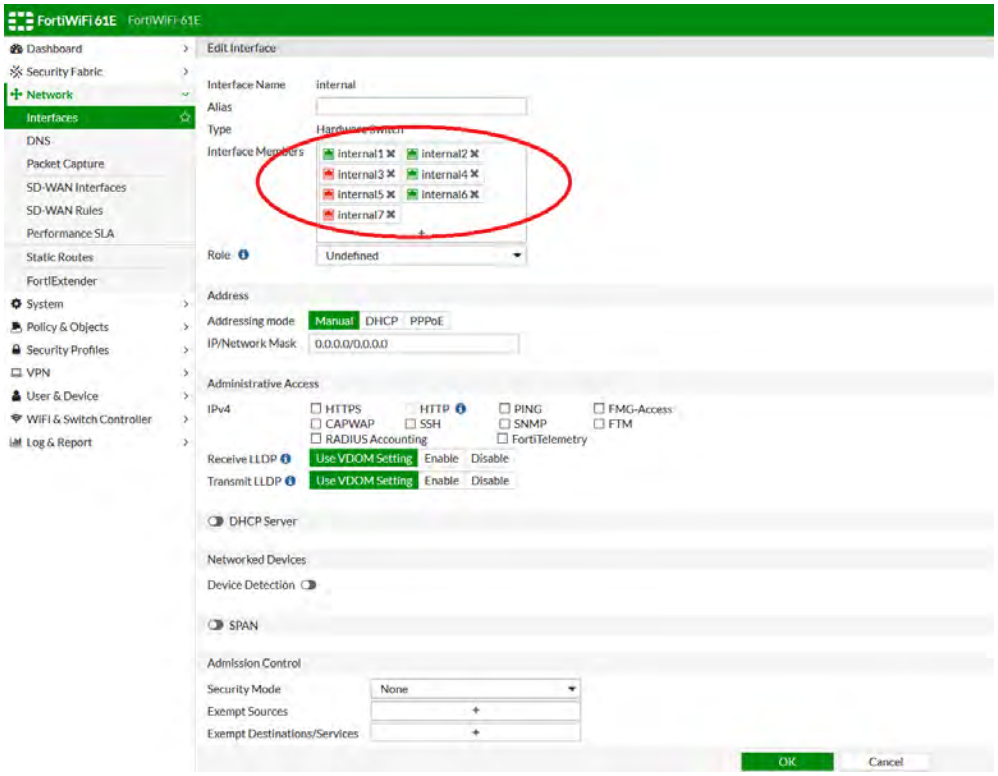
This setup is not fully compliant with a regular HA configuration. In a regular HA configuration, two logical switches are used: one for incoming traffic and one for outgoing traffic. In this example, only incoming traffic has a switch while outgoing traffic uses the internal switch. This means that if FortiGate A loses power, ISP 1 will not be available.

Traffic will flow either through ISP 1 directly or through ISP 2 via a connection between FortiGate A and FortiGate B's internal switch interface. FortiGate A decides how traffic will go through ISP 1 or ISP 2 based on SD-WAN rules. If ISP 1 is not available, then traffic will go through ISP 2.

Failover is unnecessary in this setup. Because SD-WAN will automatically failover traffic to the accessible ISP, traffic is not blocked so there is no network downtime. Also, the hardware switch interface cannot be monitored as an HA interface. If HA failover is required, a ping server must be used. The ping server monitor interface has to be configured under HA settings. This failover setup, however, results in flipping. FortiGate B will act as the primary after failover and traffic will still flow to the available ISP. Since the broken link to one of the ISPs still fails, HA will start flipping until the link is back up.

To configure the HA A-A cluster with internal hardware switches:

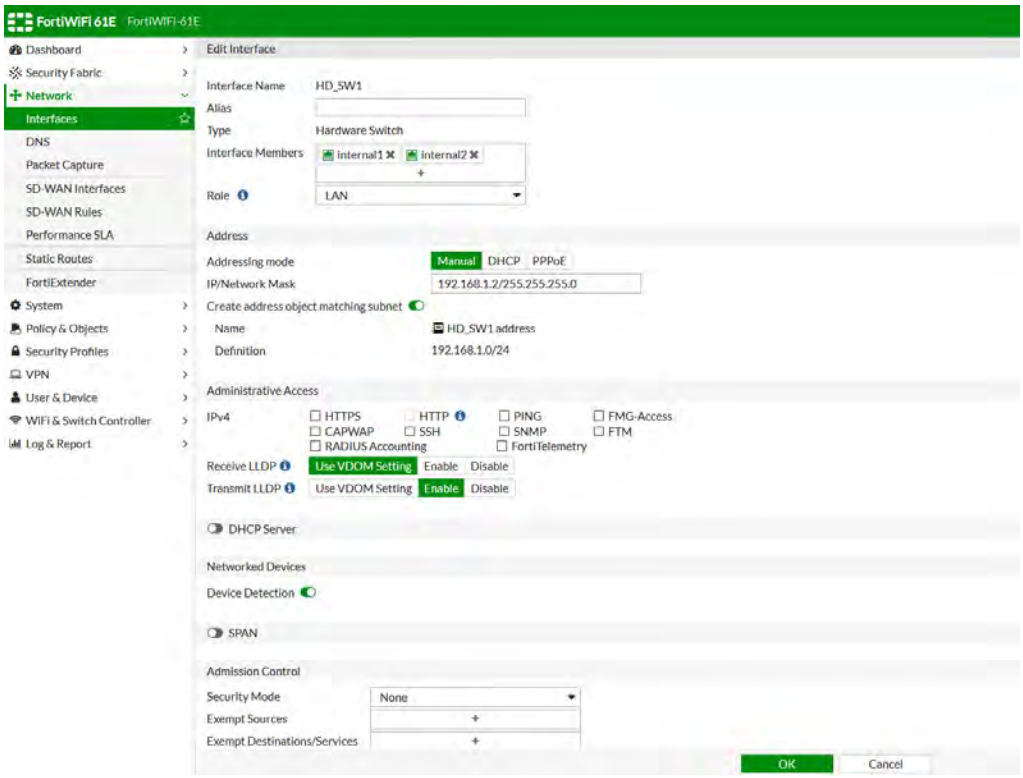
1. Configure two FortiGates with internal switches in an active-active HA cluster (follow the steps in HA active-active cluster setup on page 648), starting by connecting the heartbeat interface.
2. On the primary FortiGate, remove the existing interface members:
   a. Go to *Network* > *Interfaces*.
   b. In the *LAN* section, double-click the *internal* interface to edit it.
   c. In the *Interface Members* box, remove all the interfaces.



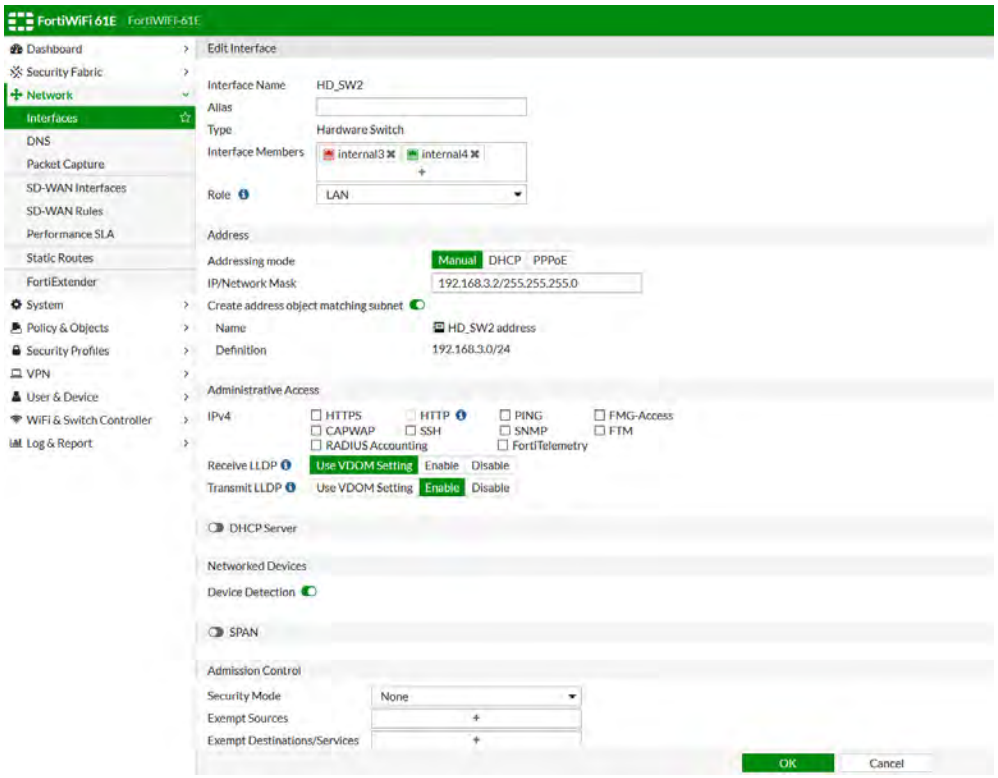
   d. Click *OK*.
3. On the primary FortiGate, configure the hardware switch interfaces for the two ISPs:
   a. Go to *Network* > *Interfaces*.
   b. Click *Create New* > *Interface*.

Netskope Exhibit 1017

c. Enter a name (*HD_SW1*).
d. For type, select *Hardware Switch*.
e. For *Interface Members*, add two interfaces (*internal1* and *internal2*).
f. Configure the other settings as needed.
g. Click *OK*.



h. Repeat these steps to create a second hardware switch interface (*HD_SW2*) with two interface members (*internal3* and *internal4*).

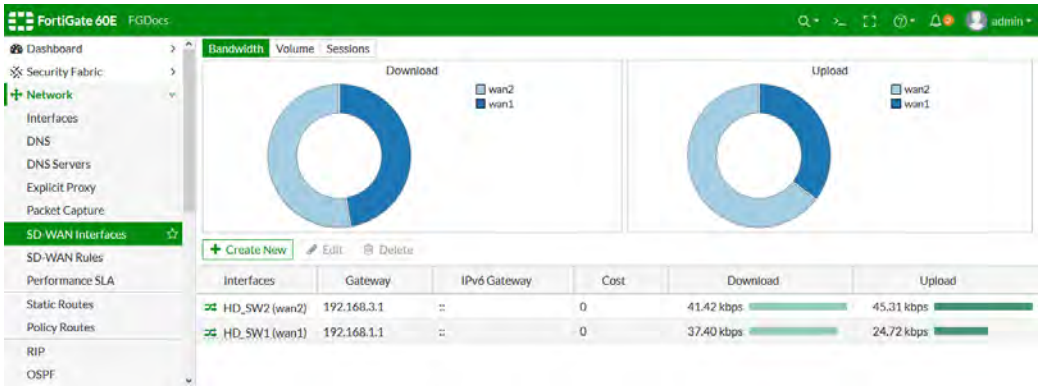


4. On the primary FortiGate, set up SD-WAN:



The primary FortiGate makes all the SD-WAN decisions.

   a. Go to *Network* > *SD-WAN Interfaces*.
   b. Click *Create New*.
   c. In the *Interface* dropdown, select *HD_SW1* and enter the *Gateway* address.
   d. Click *OK*.
   e. Repeat these steps to add the second interface (*HD_SW2*).
   f. Click *Apply*.



5. Connect the devices as shown in the topology:
   a. Connect the incoming interface to the internal switch on both FortiGates.
   b. On FortiGate A, connect ISP 1 to HD_SW1.
   c. On FortiGate B, connect ISP 2 to HD_SW2.
   d. For HD_SW1, connect FortiGate A directly to B.
   e. For HD_SW2, connect FortiGate A directly to B.



The default implicit rule load-balancing algorithm for SD-WAN is the source IP address. For more information about rule types and configurations, see Implicit rule on page 461.

Troubleshooting SD-WAN

The following topics provide instructions on SD-WAN troubleshooting:

- Understanding SD-WAN related logs on page 570
- SD-WAN related diagnose commands on page 573
- SLA logging on page 577
- SLA monitoring using the REST API on page 579
- SD-WAN bandwidth monitoring service on page 581

Understanding SD-WAN related logs

This topic lists the SD-WAN related logs and explains when the logs will be triggered.

Health-check detects a failure:

- When health-check detects a failure, it will record a log:

```
34: date=2019-03-23 time=17:26:06 logid="0100022921" type="event" subtype="system"
level="critical" vd="root" eventtime=1553387165 logdesc="Routing information changed"
name="test" interface="R150" status="down" msg="Static route on interface R150 may be
removed by health-check test. Route:  (10.100.1.2->10.100.2.22 ping-down)"
```

- When health-check detects a recovery, it will record a log:

```
32: date=2019-03-23 time=17:26:54 logid="0100022921" type="event" subtype="system"
level="critical" vd="root" eventtime=1553387214 logdesc="Routing information changed"
name="test" interface="R150" status="up" msg="Static route on interface R150 may be
added by health-check test. Route:  (10.100.1.2->10.100.2.22 ping-up)"
```

Health-check has an SLA target and detects SLA qualification changes:

- When health-check has an SLA target and detects SLA changes, and changes to fail:

```
5: date=2019-04-11 time=11:48:39 logid="0100022923" type="event" subtype="system"
level="notice" vd="root" eventtime=1555008519816639290 logdesc="Virtual WAN Link status"
msg="SD-WAN Health Check(ping) SLA(1): number of pass members changes from 2 to 1."
```


Netskope Exhibit 1017

- When health-check has an SLA target and detects SLA changes, and changes to pass:

```
2: date=2019-04-11 time=11:49:46 logid="0100022923" type="event" subtype="system"
level="notice" vd="root" eventtime=1555008586149038471 logdesc="Virtual WAN Link status"
msg="SD-WAN Health Check(ping) SLA(1): number of pass members changes from 1 to 2."
```

SD-WAN calculates a link's session/bandwidth over/under its ratio and stops/resumes traffic:

- When SD-WAN calculates a link's session/bandwidth over its configured ratio and stops forwarding traffic:

```
3: date=2019-04-10 time=17:15:40 logid="0100022924" type="event" subtype="system"
level="notice" vd="root" eventtime=1554941740185866628 logdesc="Virtual WAN Link volume
status" interface="R160" msg="The member(3) enters into conservative status with limited
ablity to receive new sessions for too much traffic."
```

- When SD-WAN calculates a link's session/bandwidth according to its ratio and resumes forwarding traffic:

```
1: date=2019-04-10 time=17:20:39 logid="0100022924" type="event" subtype="system"
level="notice" vd="root" eventtime=1554942040196041728 logdesc="Virtual WAN Link volume
status" interface="R160" msg="The member(3) resume normal status to receive new sessions
for internal adjustment."
```

The SLA mode service rule's SLA qualified member changes:

- When the SLA mode service rule's SLA qualified member changes. In this example R150 fails the SLA check, but is still alive:

```
14: date=2019-03-23 time=17:44:12 logid="0100022923" type="event" subtype="system"
level="notice" vd="root" eventtime=1553388252 logdesc="Virtual WAN Link status"
msg="Service2() prioritized by SLA will be redirected in seq-num order 2(R160) 1(R150)."
15: date=2019-03-23 time=17:44:12 logid="0100022923" type="event" subtype="system"
level="notice" vd="root" eventtime=1553388252 logdesc="Virtual WAN Link status"
interface="R150" msg="The member1(R150) SLA order changed from 1 to 2. "
16: date=2019-03-23 time=17:44:12 logid="0100022923" type="event" subtype="system"
level="notice" vd="root" eventtime=1553388252 logdesc="Virtual WAN Link status"
interface="R160" msg="The member2(R160) SLA order changed from 2 to 1. "
```

- When the SLA mode service rule's SLA qualified member changes. In this example R150 changes from fail to pass:

```
1: date=2019-03-23 time=17:46:05 logid="0100022923" type="event" subtype="system"
level="notice" vd="root" eventtime=1553388365 logdesc="Virtual WAN Link status"
msg="Service2() prioritized by SLA will be redirected in seq-num order 1(R150) 2(R160)."
2: date=2019-03-23 time=17:46:05 logid="0100022923" type="event" subtype="system"
level="notice" vd="root" eventtime=1553388365 logdesc="Virtual WAN Link status"
interface="R160" msg="The member2(R160) SLA order changed from 1 to 2. "
3: date=2019-03-23 time=17:46:05 logid="0100022923" type="event" subtype="system"
level="notice" vd="root" eventtime=1553388365 logdesc="Virtual WAN Link status"
interface="R150" msg="The member1(R150) SLA order changed from 2 to 1. "
```

The priority mode service rule member's link status changes:

- When priority mode service rule member's link status changes. In this example R150 changes to better than R160, and both are still alive:

```
1: date=2019-03-23 time=17:33:23 logid="0100022923" type="event" subtype="system"
level="notice" vd="root" eventtime=1553387603 logdesc="Virtual WAN Link status"
msg="Service2() prioritized by packet-loss will be redirected in seq-num order 1(R150) 2
(R160)."
```


Netskope Exhibit 1017

```
2: date=2019-03-23 time=17:33:23 logid="0100022923" type="event" subtype="system"
level="notice" vd="root" eventtime=1553387603 logdesc="Virtual WAN Link status"
interface="R160" msg="The member2(R160) link quality packet-loss order changed from 1 to
2. "
3: date=2019-03-23 time=17:33:23 logid="0100022923" type="event" subtype="system"
level="notice" vd="root" eventtime=1553387603 logdesc="Virtual WAN Link status"
interface="R150" msg="The member1(R150) link quality packet-loss order changed from 2 to
1. "
```

- When priority mode service rule member's link status changes. In this example R160 changes to better than R150, and both are still alive:

```
6: date=2019-03-23 time=17:32:01 logid="0100022923" type="event" subtype="system"
level="notice" vd="root" eventtime=1553387520 logdesc="Virtual WAN Link status"
msg="Service2() prioritized by packet-loss will be redirected in seq-num order 2(R160) 1
(R150)."
7: date=2019-03-23 time=17:32:01 logid="0100022923" type="event" subtype="system"
level="notice" vd="root" eventtime=1553387520 logdesc="Virtual WAN Link status"
interface="R150" msg="The member1(R150) link quality packet-loss order changed from 1 to
2. "
8: date=2019-03-23 time=17:32:01 logid="0100022923" type="event" subtype="system"
level="notice" vd="root" eventtime=1553387520 logdesc="Virtual WAN Link status"
interface="R160" msg="The member2(R160) link quality packet-loss order changed from 2 to
1. "
```

SD-WAN member is used in service and it fails the health-check:

- When SD-WAN member fails the health-check, it will stop forwarding traffic:

```
6: date=2019-04-11 time=13:33:21 logid="0100022923" type="event" subtype="system"
level="notice" vd="root" eventtime=1555014801844089814 logdesc="Virtual WAN Link status"
interface="R160" msg="The member2(R160) link is unreachable or miss threshold. Stop
forwarding traffic. "
```

- When SD-WAN member passes the health-check again, it will resume forwarding logs:

```
2: date=2019-04-11 time=13:33:36 logid="0100022923" type="event" subtype="system"
level="notice" vd="root" eventtime=1555014815914643626 logdesc="Virtual WAN Link status"
interface="R160" msg="The member2(R160) link is available. Start forwarding traffic. "
```

Load-balance mode service rule's SLA qualified member changes:

- When load-balance mode service rule's SLA qualified member changes. In this example R150 changes to not meet SLA:

```
2: date=2019-04-11 time=14:11:16 logid="0100022923" type="event" subtype="system"
level="notice" vd="root" eventtime=1555017075926510687 logdesc="Virtual WAN Link status"
msg="Service1(rule2) will be load balanced among members  2(R160) with available
routing."
3: date=2019-04-11 time=14:11:16 logid="0100022923" type="event" subtype="system"
level="notice" vd="root" eventtime=1555017075926508676 logdesc="Virtual WAN Link status"
interface="R150" msg="The member1(R150) SLA order changed from 1 to 2. "
4: date=2019-04-11 time=14:11:16 logid="0100022923" type="event" subtype="system"
level="notice" vd="root" eventtime=1555017075926507182 logdesc="Virtual WAN Link status"
interface="R160" msg="The member2(R160) SLA order changed from 2 to 1. "
```

- When load-balance mode service rule's SLA qualified member changes. In this example R150 changes to meet SLA:

Netskope Exhibit 1017

```
1: date=2019-04-11 time=14:33:23 logid="0100022923" type="event" subtype="system"
level="notice" vd="root" eventtime=1555017075926510668 logdesc="Virtual WAN Link status"
msg="Service1(rule2) will be load balanced among members 1(R150) 2(R160) with available
routing."
2: date=2019-03-23 time=14:33:23 logid="0100022923" type="event" subtype="system"
level="notice" vd="root" eventtime=1553387603592651068 logdesc="Virtual WAN Link status"
interface="R160" msg="The member2(R160) link quality packet-loss order changed from 1 to
2. "
3: date=2019-03-23 time=14:33:23 logid="0100022923" type="event" subtype="system"
level="notice" vd="root" eventtime=1553387603592651068 logdesc="Virtual WAN Link status"
interface="R150" msg="The member1(R150) link quality packet-loss order changed from 2 to
1. "
```

SLA link status logs, generated with interval sla-fail-log-period or sla-pass-log-period:

- When SLA fails, SLA link status logs will be generated with interval sla-fail-log-period:

```
7: date=2019-03-23 time=17:45:54 logid="0100022925" type="event" subtype="system"
level="notice" vd="root" eventtime=1553388352 logdesc="Link monitor SLA information"
name="test" interface="R150" status="up" msg="Latency: 0.016, jitter: 0.002, packet
loss: 21.000%, inbandwidth: 0Mbps, outbandwidth: 200Mbps, bibandwidth: 200Mbps, sla_map:
0x0"
```

- When SLA passes, SLA link status logs will be generated with interval sla-pass-log-period:

```
5: date=2019-03-23 time=17:46:05 logid="0100022925" type="event" subtype="system"
level="information" vd="root" eventtime=1553388363 logdesc="Link monitor SLA
information" name="test" interface="R150" status="up" msg="Latency: 0.017, jitter:
0.003, packet loss: 0.000%, inbandwidth: 0Mbps, outbandwidth: 200Mbps, bibandwidth:
200Mbps, sla_map: 0x1"
```

SD-WAN related diagnose commands

This topic lists the SD-WAN related diagnose commands and related output.

To check SD-WAN health-check status:

```
FGT # diagnose sys virtual-wan-link health-check
Health Check(server):
Seq(1): state(alive), packet-loss(0.000%) latency(15.247), jitter(5.231) sla_map=0x0
Seq(2): state(alive), packet-loss(0.000%) latency(13.621), jitter(6.905) sla_map=0x0

FGT  # diagnose sys virtual-wan-link health-check
Health Check(ping):
Seq(1): state(alive), packet-loss(0.000%) latency(0.683), jitter(0.082) sla_map=0x0
Seq(2): state(dead), packet-loss(100.000%) sla_map=0x0

FGT # diagnose sys virtual-wan-link health-check google
Health Check(google):
Seq(1): state(alive), packet-loss(0.000%) latency(14.563), jitter(4.334) sla_map=0x0
Seq(2): state(alive), packet-loss(0.000%) latency(12.633), jitter(6.265) sla_map=0x0
```

To check SD-WAN member status:

- When SD-WAN load-balance mode is *source-ip-based*/*source-dest-ip-based*.

```
FGT # diagnose sys virtual-wan-link member
Member(1): interface: port13, gateway: 10.100.1.1 2004:10:100:1::1, priority: 0, weight:
0
Member(2): interface: port15, gateway: 10.100.1.5 2004:10:100:1::5, priority: 0, weight:
0
```

- When SD-WAN load-balance mode is *weight-based*.

```
FGT # diagnose sys virtual-wan-link  member
Member(1): interface: port13, gateway: 10.100.1.1 2004:10:100:1::1, priority: 0, weight:
33
Member(2): interface: port15, gateway: 10.100.1.5 2004:10:100:1::5, priority: 0, weight:
66
```

- When SD-WAN load-balance mode is *measured-volume-based*.
  - Both members are under volume and still have room:

```
FGT # diagnose sys virtual-wan-link member
Member(1): interface: port13, gateway: 10.100.1.1 2004:10:100:1::1, priority: 0,
weight: 33
  Config volume ratio: 33, last reading: 8211734579B, volume room 33MB
Member(2): interface: port15, gateway: 10.100.1.5 2004:10:100:1::5, priority: 0,
weight: 66
  Config volume ratio: 66, last reading: 24548159B, volume room 66MB
```

  - Some members are overloaded and some still have room:

```
FGT # diagnose sys virtual-wan-link member
Member(1): interface: port1, gateway: 10.10.0.2, priority: 0, weight: 0
  Config volume ratio: 10, last reading: 10297221000B, overload volume 1433MB
Member(2): interface: port2, gateway: 10.11.0.2, priority: 0, weight: 38
  Config volume ratio: 50, last reading: 45944239916B, volume room 38MB
```

- When SD-WAN load balance mode is *usage-based*/*spillover*.
  - When no spillover occurs:

```
FGT # diagnose sys virtual-wan-link member
Member(1): interface: port13, gateway: 10.100.1.1 2004:10:100:1::1, priority: 0,
weight: 255
  Egress-spillover-threshold: 400kbit/s, ingress-spillover-threshold: 300kbit/s
  Egress-overbps=0, ingress-overbps=0
Member(2): interface: port15, gateway: 10.100.1.5 2004:10:100:1::5, priority: 0,
weight: 254
  Egress-spillover-threshold: 0kbit/s, ingress-spillover-threshold: 0kbit/s
  Egress-overbps=0, ingress-overbps=0
```

  - When member has reached limit and spillover occurs:

```
FGT # diagnose sys virtual-wan-link member
Member(1): interface: port13, gateway: 10.100.1.1 2004:10:100:1::1, priority: 0,
weight: 255
  Egress-spillover-threshold: 400kbit/s, ingress-spillover-threshold: 300kbit/s
  Egress-overbps=1, ingress-overbps=1
Member(2): interface: port15, gateway: 10.100.1.5 2004:10:100:1::5, priority: 0,
weight: 254
```


Netskope Exhibit 1017

```
Egress-spillover-threshold: 0kbit/s, ingress-spillover-threshold: 0kbit/s
Egress-overbps=0, ingress-overbps=0
```

- You can also use the `diagnose netlink dstmac list` command to check if you are over the limit.

```
FGT # diagnose netlink dstmac list port13
dev=port13 mac=08:5b:0e:ca:94:9d rx_tcp_mss=0 tx_tcp_mss=0 egress_overspill_
threshold=51200 egress_bytes=103710 egress_over_bps=1 ingress_overspill_
threshold=38400 ingress_bytes=76816 ingress_over_bps=1 sampler_rate=0
```

To check SD-WAN service rules status:

- *Manual mode* service rules.

```
FGT # diagnose sys virtual-wan-link  service
Service(1): Address Mode(IPV4) flags=0x0
  TOS(0x0/0x0), Protocol(0: 1->65535), Mode(manual)
  Members:
    1: Seq_num(2), alive, selected
  Dst address: 10.100.21.0-10.100.21.255
```

- *Auto mode* service rules.

```
FGT # diagnose sys virtual-wan-link service
Service(1): Address Mode(IPV4) flags=0x0
  TOS(0x0/0x0), Protocol(0: 1->65535), Mode(auto), link-cost-factor(latency), link-cost-
threshold(10), health-check(ping)
  Members:
    1: Seq_num(2), alive, latency: 0.011
    2: Seq_num(1), alive, latency: 0.018, selected
  Dst address: 10.100.21.0-10.100.21.255
```

- *Priority mode* service rules.

```
FGT # diagnose sys virtual-wan-link  service
Service(1): Address Mode(IPV4) flags=0x0
  TOS(0x0/0x0), Protocol(0: 1->65535), Mode(priority), link-cost-factor(latency), link-
cost-threshold(10), health-check(ping)
  Members:
    1: Seq_num(2), alive, latency: 0.011, selected
    2: Seq_num(1), alive, latency: 0.017, selected
  Dst address: 10.100.21.0-10.100.21.255
```

- *Load-balance* mode service rules.

```
FGT # diagnose sys virtual-wan-link  service
Service(1): Address Mode(IPV4) flags=0x0
  TOS(0x0/0x0), Protocol(0: 1->65535), Mode(load-balance)
  Members:
    1: Seq_num(1), alive, sla(0x1), num of pass(1), selected
    2: Seq_num(2), alive, sla(0x1), num of pass(1), selected
  Dst address: 10.100.21.0-10.100.21.255
```

- *SLA mode* service rules.

```
FGT # diagnose sys virtual-wan-link service
Service(1): Address Mode(IPV4) flags=0x0
  TOS(0x0/0x0), Protocol(0: 1->65535), Mode(sla)
  Members:
    1: Seq_num(1), alive, sla(0x1), cfg_order(0), cost(0), selected
```


Netskope Exhibit 1017

```
      2: Seq_num(2), alive, sla(0x1), cfg_order(1), cost(0), selected
    Dst address: 10.100.21.0-10.100.21.255
```

To check interface logs from the past 15 minutes:

```
FGT (root) # diagnose sys virtual-wan-link intf-sla-log R150
Timestamp: Fri Apr 12 11:08:36 2019, used inbandwidth: 0bps, used outbandwidth: 0bps, used
bibandwidth: 0bps, tx bytes: 860bytes, rx bytes: 1794bytes.
Timestamp: Fri Apr 12 11:08:46 2019, used inbandwidth: 1761bps, used outbandwidth: 1710bps,
used bibandwidth: 3471bps, tx bytes: 2998bytes, rx bytes: 3996bytes.
Timestamp: Fri Apr 12 11:08:56 2019, used inbandwidth: 2452bps, used outbandwidth: 2566bps,
used bibandwidth: 5018bps, tx bytes: 7275bytes, rx bytes: 7926bytes.
Timestamp: Fri Apr 12 11:09:06 2019, used inbandwidth: 2470bps, used outbandwidth: 3473bps,
used bibandwidth: 5943bps, tx bytes: 13886bytes, rx bytes: 11059bytes.
Timestamp: Fri Apr 12 11:09:16 2019, used inbandwidth: 2433bps, used outbandwidth: 3417bps,
used bibandwidth: 5850bps, tx bytes: 17946bytes, rx bytes: 13960bytes.
Timestamp: Fri Apr 12 11:09:26 2019, used inbandwidth: 2450bps, used outbandwidth: 3457bps,
used bibandwidth: 5907bps, tx bytes: 22468bytes, rx bytes: 17107bytes.
```

To check SLA logs in the past 15 minutes:

```
FGT (root) # diagnose sys virtual-wan-link sla-log ping 1
Timestamp: Fri Apr 12 11:09:27 2019, vdom root, health-check ping, interface: R150, status:
up, latency: 0.014, jitter: 0.003, packet loss: 16.000%.
Timestamp: Fri Apr 12 11:09:28 2019, vdom root, health-check ping, interface: R150, status:
up, latency: 0.015, jitter: 0.003, packet loss: 15.000%.
Timestamp: Fri Apr 12 11:09:28 2019, vdom root, health-check ping, interface: R150, status:
up, latency: 0.014, jitter: 0.003, packet loss: 14.000%.
Timestamp: Fri Apr 12 11:09:29 2019, vdom root, health-check ping, interface: R150, status:
up, latency: 0.015, jitter: 0.003, packet loss: 13.000%.
```

To check Application Control used in SD-WAN and the matching IP addresses:

```
FGT # diagnose sys virtual-wan-link internet-service-app-ctrl-list
Ctrl application(Microsoft.Authentication 41475):Internet Service ID(4294836224)
        Protocol(6), Port(443)
        Address(2): 104.42.72.21 131.253.61.96
Ctrl application(Microsoft.CDN 41470):Internet Service ID(4294836225)
Ctrl application(Microsoft.Lync 28554):Internet Service ID(4294836226)
Ctrl application(Microsoft.Office.365 33182):Internet Service ID(4294836227)
Ctrl application(Microsoft.Office.365.Portal 41468):Internet Service ID(4294836228)
Ctrl application(Microsoft.Office.Online 16177):Internet Service ID(4294836229)
Ctrl application(Microsoft.OneNote 40175):Internet Service ID(4294836230)
Ctrl application(Microsoft.Portal 41469):Internet Service ID(4294836231)
        Protocol(6), Port(443)
        Address(8): 23.58.134.172 131.253.33.200 23.58.135.29 204.79.197.200 64.4.54.254
23.59.156.241 13.77.170.218 13.107.22.200
Ctrl application(Microsoft.Sharepoint 16190):Internet Service ID(4294836232)
Ctrl application(Microsoft.Sway 41516):Internet Service ID(4294836233)
Ctrl application(Microsoft.Tenant.Namespace 41471):Internet Service ID(4294836234)
```

To check IPsec aggregate interface when SD-WAN uses the per-packet distribution feature:

```
# diagnose sys ipsec-aggregate list
agg1 algo=L3 member=2 run_tally=2
```


Netskope Exhibit 1017

```
members:
        vd1-p1
        vd1-p2
```

To check BGP learned routes and determine if they are used in SD-WAN service:

```
FGT # get router info bgp network
FGT # get router info bgp network 10.100.11.0
BGP routing table entry for 10.100.10.0/24
Paths: (2 available, best 1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
   172.10.22.2
  20
    10.100.20.2 from 10.100.20.2 (6.6.6.6)
      Origin EGP metric 200, localpref 100, weight 10000, valid, external, best
      Community: 30:5
      Last update: Wen Mar 20 18:45:17 2019
FGT # get router info route-map-address
Extend-tag: 15, interface(wan2:16)
        10.100.11.0/255.255.255.0

 FGT # diagnose firewall proute list
list route policy info(vf=root):

id=4278779905 vwl_service=1(DataCenter) flags=0x0 tos=0x00 tos_mask=0x00 protocol=0
sport=0:65535 iif=0 dport=1-65535 oif=16
source wildcard(1): 0.0.0.0/0.0.0.0
destination wildcard(1): 10.100.11.0/255.255.255.0
```

SLA logging

The features adds an SD-WAN daemon function to keep a short, 10 minute history of SLA that can be viewed in the CLI.

Performance SLA results related to interface selection, session failover, and other information, can be logged. These logs can then be used for long-term monitoring of traffic issues at remote sites, and for reports and views in FortiAnalyzer.

The time intervals that Performance SLA fail and pass logs are generated in can be configured.

To configure the fail and pass logs' generation time interval:

```
config system virtual-wan-link
    config health-check
        edit "ping"
            set sla-fail-log-period 30
            set sla-pass-log-period 60
        next
    end
end
```

To view the 10 minute Performance SLA link status history:

```
FGT_A (root) # diagnose sys virtual-wan-link sla-log ping 1
Timestamp: Thu Feb 28 10:58:24 2019, vdom root, health-check ping, interface: R150, status:
```


Netskope Exhibit 1017

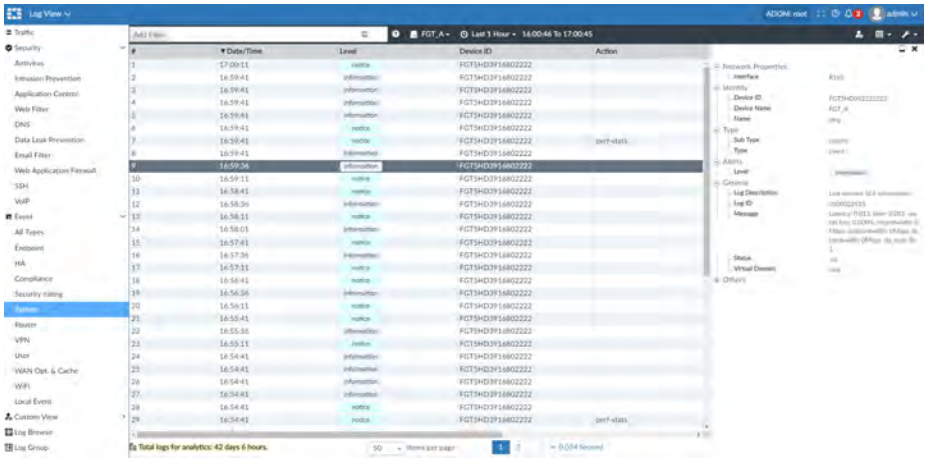
```
up, latency: 0.000, jitter: 0.000, packet loss: 0.000%.
Timestamp: Thu Feb 28 10:58:24 2019, vdom root, health-check ping, interface: R150, status:
up, latency: 0.097, jitter: 0.000, packet loss: 0.000%.
Timestamp: Thu Feb 28 10:58:25 2019, vdom root, health-check ping, interface: R150, status:
up, latency: 0.058, jitter: 0.040, packet loss: 0.000%.
Timestamp: Thu Feb 28 10:58:25 2019, vdom root, health-check ping, interface: R150, status:
up, latency: 0.044, jitter: 0.026, packet loss: 0.000%.
... ...
```

SLA pass logs

The FortiGate generates Performance SLA logs at the specified pass log interval (`sla-pass-log-period`) when SLA passes.

```
3: date=2019-02-28 time=11:53:26 logid="0100022925" type="event" subtype="system"
level="information" vd="root" eventtime=1551383604 logdesc="Link monitor SLA information"
name="ping" interface="R160" status="up" msg="Latency: 0.013, jitter: 0.001, packet loss:
0.000%, inbandwidth: 0Mbps, outbandwidth: 0Mbps, bibandwidth: 0Mbps, sla_map: 0x1"
7: date=2019-02-28 time=11:52:26 logid="0100022925" type="event" subtype="system"
level="information" vd="root" eventtime=1551383545 logdesc="Link monitor SLA information"
name="ping" interface="R160" status="up" msg="Latency: 0.013, jitter: 0.002, packet loss:
0.000%, inbandwidth: 0Mbps, outbandwidth: 0Mbps, bibandwidth: 0Mbps, sla_map: 0x1"
```

In the FortiAnalyzer GUI:

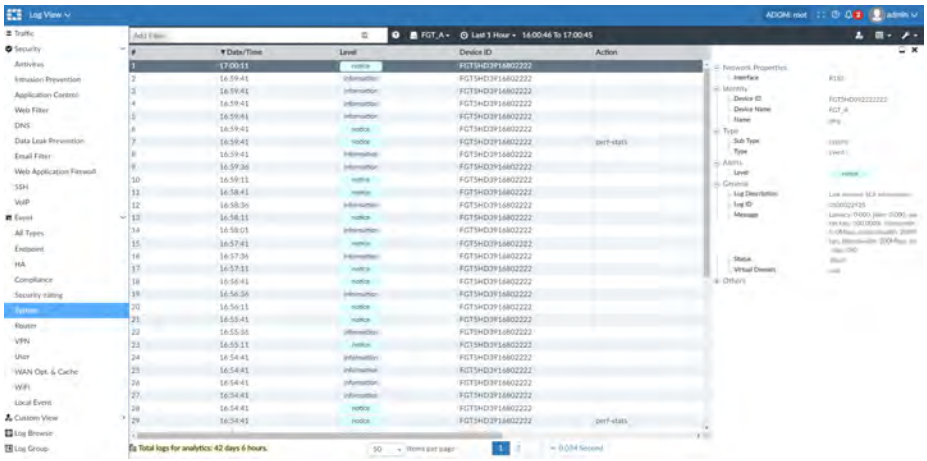


SLA fail logs

The FortiGate generates Performance SLA logs at the specified fail log interval (`sla-fail-log-period`) when SLA fails.

```
6: date=2019-02-28 time=11:52:32 logid="0100022925" type="event" subtype="system"
level="notice" vd="root" eventtime=1551383552 logdesc="Link monitor SLA information"
name="ping" interface="R150" status="down" msg="Latency: 0.000, jitter: 0.000, packet loss:
100.000%, inbandwidth: 0Mbps, outbandwidth: 200Mbps, bibandwidth: 200Mbps, sla_map: 0x0"
8: date=2019-02-28 time=11:52:02 logid="0100022925" type="event" subtype="system"
level="notice" vd="root" eventtime=1551383522 logdesc="Link monitor SLA information"
name="ping" interface="R150" status="down" msg="Latency: 0.000, jitter: 0.000, packet loss:
100.000%, inbandwidth: 0Mbps, outbandwidth: 200Mbps, bibandwidth: 200Mbps, sla_map: 0x0"
```


Netskope Exhibit 1017



SLA monitoring using the REST API

SLA log information and interface SLA information can be monitored using the REST API. This feature is also be used by FortiManager as part of its detailed SLA monitoring and drill-down features.

Interface log command example:

```
https://172.172.172.9/api/v2/monitor/virtual-wan/interface-log
{
  "http_method":"GET",
  "results":[
    {
      "interface":"port13",
      "logs":[
        {
          "timestamp":1547087168,
          "tx_bandwidth":3447,
          "rx_bandwidth":3457,
          "bi_bandwidth":6904,
          "tx_bytes":748875,
          "rx_bytes":708799,
          "egress_queue":[
          ]
        },
        {
          "timestamp":1547087178,
          "tx_bandwidth":3364,
          "rx_bandwidth":3400,
          "bi_bandwidth":6764,
          "tx_bytes":753789,
          "rx_bytes":712835,
          "egress_queue":[
          ]
        },
....
....
```

SLA log command example:

```
https://172.172.172.9/api/v2/monitor/virtual-wan/sla-log
{
  "http_method":"GET",
  "results":[
    {
      "name":"ping",
      "interface":"port13",
      "logs":[
        {
          "timestamp":1547087204,
          "link":"up",
          "latency":0.686433,
          "jitter":0.063400,
          "packetloss":0.000000
        },
        {
          "timestamp":1547087205,
          "link":"up",
          "latency":0.688433,
          "jitter":0.063133,
          "packetloss":0.000000
        },
        {
          "timestamp":1547087206,
          "link":"up",
          "latency":0.688300,
          "jitter":0.065267,
          "packetloss":0.000000
        },
....
....
```

CLI diagnose commands:

```
# diagnose sys virtual-wan-link sla-log ping 1
    Timestamp: Wed Jan 9 18:35:11 2019, vdom root, health-check ping, interface: port13,
status: up, latency: 0.698, jitter: 0.073, packet loss: 0.000%.
    Timestamp: Wed Jan 9 18:35:12 2019, vdom root, health-check ping, interface: port13,
status: up, latency: 0.704, jitter: 0.073, packet loss: 0.000%.
    Timestamp: Wed Jan 9 18:35:13 2019, vdom root, health-check ping, interface: port13,
status: up, latency: 0.709, jitter: 0.073, packet loss: 0.000%.
    Timestamp: Wed Jan 9 18:35:14 2019, vdom root, health-check ping, interface: port13,
status: up, latency: 0.707, jitter: 0.066, packet loss: 0.000%.
    Timestamp: Wed Jan 9 18:35:15 2019, vdom root, health-check ping, interface: port13,
status: up, latency: 0.710, jitter: 0.061, packet loss: 0.000%.
    Timestamp: Wed Jan 9 18:35:16 2019, vdom root, health-check ping, interface: port13,
status: up, latency: 0.707, jitter: 0.055, packet loss: 0.000%.
    Timestamp: Wed Jan 9 18:35:17 2019, vdom root, health-check ping, interface: port13,
status: up, latency: 0.703, jitter: 0.055, packet loss: 0.000%.

# diagnose sys virtual-wan-link intf-sla-log port13
    Timestamp: Wed Jan 9 18:33:49 2019, used inbandwidth: 3208bps, used outbandwidth:
3453bps, used bibandwidth: 6661bps, tx bytes: 947234bytes, rx bytes: 898622bytes.
    Timestamp: Wed Jan 9 18:33:59 2019, used inbandwidth: 3317bps, used outbandwidth:
```


Netskope Exhibit 1017

```
3450bps, used bibandwidth: 6767bps, tx bytes: 951284bytes, rx bytes: 902937bytes.
    Timestamp: Wed Jan 9 18:34:09 2019, used inbandwidth: 3302bps, used outbandwidth:
3389bps, used bibandwidth: 6691bps, tx bytes: 956268bytes, rx bytes: 907114bytes.
    Timestamp: Wed Jan 9 18:34:19 2019, used inbandwidth: 3279bps, used outbandwidth:
3352bps, used bibandwidth: 6631bps, tx bytes: 958920bytes, rx bytes: 910793bytes.
    Timestamp: Wed Jan 9 18:34:29 2019, used inbandwidth: 3233bps, used outbandwidth:
3371bps, used bibandwidth: 6604bps, tx bytes: 964374bytes, rx bytes: 914854bytes.
    Timestamp: Wed Jan 9 18:34:39 2019, used inbandwidth: 3235bps, used outbandwidth:
3362bps, used bibandwidth: 6597bps, tx bytes: 968250bytes, rx bytes: 918846bytes.
    Timestamp: Wed Jan 9 18:34:49 2019, used inbandwidth: 3165bps, used outbandwidth:
3362bps, used bibandwidth: 6527bps, tx bytes: 972298bytes, rx bytes: 922724bytes.
    Timestamp: Wed Jan 9 18:34:59 2019, used inbandwidth: 3184bps, used outbandwidth:
3362bps, used bibandwidth: 6546bps, tx bytes: 977282bytes, rx bytes: 927019bytes.
```

SD-WAN bandwidth monitoring service

The bandwidth measuring tool is used to detect true upload and download speeds. Bandwidth tests can be run on demand or automated using a script, and can be useful when configuring SD-WAN SLA and rules to balance SD-WAN traffic.

The speed test tool requires a valid SD-WAN Bandwidth Monitoring Service license.

The speed test tool is compatible with iperf3.6 with SSL support. It can test the upload bandwidth to the FortiGate Cloud speed test service. It can initiate the server connection and send download requests to the server. The tool can be run up to 10 times a day .

FortiGate downloads the speed test server list. The list expires after 24 hours. One of the speed test servers is selected, based on user input. The speed test runs, testing upload and download speeds. The test results are shown in the command terminal.

To download the speed test server list:

```
# execute speed-test-server download
Download completed.
```

To check the speed test server list:

```
# execute speed-test-server list
AWS_West valid
        Host: 34.210.67.183 5204 fortinet
        Host: 34.210.67.183 5205 fortinet
        Host: 34.210.67.183 5206 fortinet
        Host: 34.210.67.183 5207 fortinet
Google_West valid
        Host: 35.197.55.210 5204 fortinet
        Host: 35.197.55.210 5205 fortinet
        Host: 35.197.55.210 5206 fortinet
        Host: 35.197.55.210 5207 fortinet
        Host: 35.230.2.124 5204 fortinet
        Host: 35.230.2.124 5205 fortinet
        Host: 35.230.2.124 5206 fortinet
        Host: 35.230.2.124 5207 fortinet
        Host: 35.197.18.234 5204 fortinet
        Host: 35.197.18.234 5205 fortinet
```


Netskope Exhibit 1017

```
        Host: 35.197.18.234 5206 fortinet
        Host: 35.197.18.234 5207 fortinet
```

To run the speed test:

You can run the speed test without specifying a server. The system will automatically choose one server from the list and run the speed test.

```
# execute speed-test auto
The license is valid to run speed test.
Speed test quota for 2/1 is 9
current vdom=root
Run in uploading mode.
Connecting to host 35.230.2.124, port 5206
[ 16] local 172.16.78.185 port 2475 connected to 35.230.2.124 port 5206
[ ID] Interval Transfer Bitrate Retr Cwnd
[ 16] 0.00-1.01 sec 11.0 MBytes 91.4 Mbits/sec 0 486 KBytes
[ 16] 1.01-2.00 sec 11.6 MBytes 98.4 Mbits/sec 0 790 KBytes
[ 16] 2.00-3.01 sec 11.0 MBytes 91.6 Mbits/sec 15 543 KBytes
[ 16] 3.01-4.01 sec 11.2 MBytes 94.2 Mbits/sec 1 421 KBytes
[ 16] 4.01-5.01 sec 11.2 MBytes 93.5 Mbits/sec 0 461 KBytes
- - - - - - - - - - - - - - - - - - - - - - - - -
[ ID] Interval Transfer Bitrate Retr
[ 16] 0.00-5.01 sec 56.1 MBytes 93.8 Mbits/sec 16 sender
[ 16] 0.00-5.06 sec 55.8 MBytes 92.6 Mbits/sec receiver

speed test Done.
Run in reverse downloading mode!
Connecting to host 35.230.2.124, port 5206
Reverse mode, remote host 35.230.2.124 is sending
[ 16] local 172.16.78.185 port 2477 connected to 35.230.2.124 port 5206
[ ID] Interval Transfer Bitrate
[ 16] 0.00-1.00 sec 10.9 MBytes 91.4 Mbits/sec
[ 16] 1.00-2.00 sec 11.2 MBytes 93.9 Mbits/sec
[ 16] 2.00-3.00 sec 11.2 MBytes 94.0 Mbits/sec
[ 16] 3.00-4.00 sec 11.2 MBytes 93.9 Mbits/sec
[ 16] 4.00-5.00 sec 10.9 MBytes 91.1 Mbits/sec
- - - - - - - - - - - - - - - - - - - - - - - - -
[ ID] Interval Transfer Bitrate Retr
[ 16] 0.00-5.03 sec 57.5 MBytes 95.9 Mbits/sec 40 sender
[ 16] 0.00-5.00 sec 55.4 MBytes 92.9 Mbits/sec receiver

speed test Done
```

To run the speed test on a server farm or data center:

```
# execute speed-test auto AWS_West
The license is valid to run speed test.
Speed test quota for 2/1 is 8
current vdom=root
Run in uploading mode.
Connecting to host 34.210.67.183, port 5205
```


Netskope Exhibit 1017

To run the speed test on a local interface when there are multiple valid routes:

```
# execute speed-test port1 Google_West
The license is valid to run speed test.
Speed test quota for 2/1 is 6
bind to local ip 172.16.78.202
current vdom=root
Specified interface port1 does not comply with default outgoing interface port2 in routing
table!
Force to use the specified interface!
Run in uploading mode.
Connecting to host 35.197.18.234, port 5205
[ 11] local 172.16.78.202 port 20852 connected to 35.197.18.234 port 5205
[ ID] Interval Transfer Bitrate Retr Cwnd
[ 11] 0.00-1.01 sec 10.7 MBytes 89.0 Mbits/sec 0 392 KBytes
[ 11] 1.01-2.01 sec 10.5 MBytes 88.5 Mbits/sec 1 379 KBytes
[ 11] 2.01-3.01 sec 11.3 MBytes 94.5 Mbits/sec 0 437 KBytes
[ 11] 3.01-4.01 sec 11.2 MBytes 94.3 Mbits/sec 0 478 KBytes
[ 11] 4.01-5.00 sec 11.3 MBytes 95.2 Mbits/sec 0 503 KBytes
- - - - - - - - - - - - - - - - - - - - - - - - -
[ ID] Interval Transfer Bitrate Retr
[ 11] 0.00-5.00 sec 55.1 MBytes 92.3 Mbits/sec 1 sender
[ 11] 0.00-5.04 sec 54.5 MBytes 90.7 Mbits/sec receiver

speed test Done.
Run in reverse downloading mode!
Connecting to host 35.197.18.234, port 5205
Reverse mode, remote host 35.197.18.234 is sending
[ 11] local 172.16.78.202 port 20853 connected to 35.197.18.234 port 5205
[ ID] Interval Transfer Bitrate
[ 11] 0.00-1.00 sec 10.9 MBytes 91.1 Mbits/sec
[ 11] 1.00-2.00 sec 11.2 MBytes 94.0 Mbits/sec
[ 11] 2.00-3.00 sec 11.2 MBytes 94.0 Mbits/sec
[ 11] 3.00-4.00 sec 11.2 MBytes 94.0 Mbits/sec
[ 11] 4.00-5.00 sec 11.2 MBytes 94.0 Mbits/sec
- - - - - - - - - - - - - - - - - - - - - - - - -
[ ID] Interval Transfer Bitrate Retr
[ 11] 0.00-5.03 sec 57.4 MBytes 95.8 Mbits/sec 33 sender
[ 11] 0.00-5.00 sec 55.7 MBytes 93.4 Mbits/sec receiver

speed test Done.
```

To add a script to run a speed test automatically once every 24 hours:

```
config system auto-script
    edit "speedtest"
        set interval 86400
        set repeat 0
        set start auto
        set script "
execute speed-test-server download
execute speed-test"
    next
end
```


Netskope Exhibit 1017

To view the results of the speed test script:

```
execute auto-script result speedtest
```


Netskope Exhibit 1017

System

This topic contains information about FortiGate administration and system configuration that you can do after installing the FortiGate in your network.

Basic system settings

Administrators

By default, FortiGate has an administrator account with the username *admin* and no password. See Administrators on page 587 for more information.

Administrator profiles

An administrator profile defines what the administrator can see and do on the FortiGate. See Administrator profiles on page 587 for more information.

Password policy

Set up a password policy to enforce password criteria and change frequency. See Password policy on page 591 for more information.

Interfaces

Physical and virtual interface allow traffic to flow between internal networks, and between the internet and internal networks. See Interfaces on page 340 for more information.

Advanced system settings

SNMP

The simple network management protocol (SNMP) allows you to monitor hardware on your network. See SNMP on page 665 for more information.

DHCP server

You can configure one or more DHCP servers on any FortiGate interface. See DHCP server on page 428 for more information.

VDOM

You can use virtual domains (VDOMs) to divide a FortiGate into multiple virtual devices that function independently. See Virtual Domains on page 610 for more information.

High availability

You can configure multiple FortiGate devices, including private and public cloud VMs, in HA mode. See High Availability on page 631 for more information.

Certificates

You can manage certificates on the FortiGate. See Certificates on page 695 for more information.

Operating modes

A FortiGate or VDOM (in multi-vdom mode) can operate in either NAT/Route mode or Transparent mode.

NAT/Route mode

The FortiGate or VDOM is installed as a gateway between two networks, such as a private network and the internet. This allows the FortiGate to hide the IP addresses on the private network using NAT. NAT/Route mode can also be used when several ISPs are used for redundant internet connections.

By default, new VDOMs are set to NAT/Route operation mode.

See Configure VDOM-A on page 618 for more information.

Transparent mode

The FortiGate or VDOM is installed between the internal network and the router. The FortiGate does not changes any IP addresses, and only applies security scanning to traffic. When you add a FortiGate that is in transparent mode to a network, it only needs to be provided with a management IP address.

Transparent mode is primarily used when increased network protection is needed without changing the network configuration.

See Configure VDOM-A on page 627 for more information.

To change the operating mode of a FortiGate or VDOM:

```
config system settings
    set opmode {nat | transparent}
end
```

Administrators

By default, FortiGate has an administrator account with the username *admin* and no password. To prevent unauthorized access to the FortiGate, this account must be protected with a password. Additional administrators can be added for various functions, each with a unique username, password, and set of access privileges.

The following topics provide information about administrators:

- Administrator profiles on page 587
- Add a local administrator on page 589
- Remote authentication for administrators on page 589
- Password policy on page 591
- Admin profile option for diagnose access on page 592

Administrator profiles

Administrator profiles define what the administrator can do when logged into the FortiGate. When you set up an administrator account, you also assign an administrator profile which dictates what the administrator sees. Depending on the nature of the administrator's work, access level or seniority, you can allow them to view and configure as much or as little as is required.

By default, the FortiGate has an *admin* administrator account that uses the *super_admin* profile.

super_admin profile

This profile has access to all components of FortiOS, including the ability to add and remove other system administrators. For certain administrative functions, such as backing up and restoring the configuration, super_admin access is required. To ensure that there is always a method to administer the FortiGate, the super_admin profile can't be deleted or modified.



Lower level administrator profiles can't backup or restore the FortiOS configuration.

The super_admin profile is used by the default admin account. It is recommended that you add a password and rename this account once you have set up your FortiGate. In order to rename the default account, a second admin account is required.

Creating customized profiles

To create a profile in the GUI:

1. Go to *System* > *Admin Profiles*.
2. Select *Create New*.
3. Configure the following settings:
   - Name.
   - Access permissions.

- Override idle timeout.

4. Select *OK*.

To create a profile in the CLI:

```
config system accprofile
    edit "sample"
        set secfabgrp read-write
        set ftviewgrp read-write
        set authgrp read-write
        set sysgrp read-write
        set netgrp read-write
        set loggrp read-write
        set fwgrp read-write
        set vpngrp read-write
        set utmgrp read-write
        set wanoptgrp read-write
        set wifi read-write
    next
end
```

Edit profiles

To edit a profile in the GUI:

1. Go to *System* > *Admin Profiles*.
2. Choose the profile to be edited and select *Edit*.
3. Select *OK* to save any changes made.

To edit a profile in the CLI:

```
config system accprofile
    edit "sample"
        set secfabgrp read
    next
end
```

Delete profiles

To delete a profile in the GUI:

1. Go to *System* > *Admin Profiles*.
2. Choose the profile to be deleted and select *Delete*.
3. Select *OK*.

To delete a profile in the CLI:

```
config system accprofile
    delete "sample"
end
```

Add a local administrator

By default, FortiGate has one super admin named `admin`. You can create more administrator accounts with different privileges.

To create an administrator account in the GUI:

1. Go to *System* > *Administrators*.
2. Select *Create New* > *Administrator*.
3. Specify the *Username*.



Do not use the characters `< > ( ) # " '` in the administrator username.

Using these characters in an administrator username might have a cross site scripting (XSS) vulnerability.

4. Set *Type* to *Local User*.
5. Set the password and other fields.
6. Click *OK*.

To create an administrator account in the CLI:

```
config system admin
    edit <admin_name>
        set accprofile <profile_name>
        set vdom <vdom_name>
        set password <password for this admin>
    next
end
```

Remote authentication for administrators

Administrators can use remote authentication, such as LDAP, to connect to the FortiGate.

Setting up remote authentication for administrators includes the following steps:

1. Configure the LDAP server on page 589
2. Add the LDAP server to a user group on page 590
3. Configure the administrator account on page 590

Configure the LDAP server

To configure the LDAP server in the GUI:

1. Go to *User & Authentication* > *LDAP Servers* and select *Create New*.
2. Enter the server *Name*, *Server IP* address or *Name*.
3. Enter the *Common Name Identifier* and *Distinguished Name*.
4. Set the *Bind Type* to *Regular* and enter the *Username* and *Password*.
5. Click *OK*.

To configure the LDAP server in the CLI:

```
config user ldap
  edit <ldap_server_name>
    set server <server_ip>
    set cnid "cn"
    set dn "dc=XYZ,dc=fortinet,dc=COM"
    set type regular
    set username "cn=Administrator,dc=XYA, dc=COM"
    set password <password>
  next
end
```

Add the LDAP server to a user group

After configuring the LDAP server, create a user group that includes that LDAP server.

To create a user group in the GUI:

1. Go to *User & Authentication* > *User Groups* and select *Create New*.
2. Enter a *Name* for the group.
3. In the *Remote groups* section, select *Create New*.
4. Select the *Remote Server* from the dropdown list.
5. Click *OK*.

To create a user group in the CLI:

```
config user group
  edit <Group_name>
    set member "ldap_server_name"
  next
end
```

Configure the administrator account

After configuring the LDAP server and adding it to a user group, create a new administrator. For this administrator, instead of entering a password, use the new user group and the wildcard option for authentication.

To create an administrator in the GUI:

1. Go to *System* > *Administrators*.
2. Select *Create New* > *Administrator*.
3. Specify the *Username*.
4. Set *Type* to *Match a user on a remote server group*.
5. In *Remote User Group*, select the user group you created.
6. Select *Wildcard*.
   The Wildcard option allows LDAP users to connect as this administrator.
7. Select an *Administrator Profile*.
8. Click *OK*.

Netskope Exhibit 1017

To create an administrator in the CLI:

```
config system admin
  edit <admin_name>
    set remote-auth enable
    set accprofile super_admin
    set wild card enable
    set remote-group ldap
end
```

Other methods of administrator authentication

Administrator accounts can use different methods for authentication, including RADIUS, TACACS+, and PKI.

RADIUS authentication for administrators

To use a RADIUS server to authenticate administrators, you must:

- Configure the FortiGate to access the RADIUS server.
- Create the RADIUS user group.
- Configure an administrator to authenticate with a RADIUS server.

TACACS+ authentication for administrators

To use a TACACS+ server to authenticate administrators, you must:

- Configure the FortiGate to access the TACACS+ server.
- Create a TACACS+ user group.
- Configure an administrator to authenticate with a TACACS+ server.

PKI certificate authentication for administrators

To use PKI authentication for an administrator, you must:

- Configure a PKI user.
- Create a PKI user group.
- Configure an administrator to authenticate with a PKI certificate.

Password policy

Brute force password software can launch more than just dictionary attacks. It can discover common passwords where a letter is replaced by a number. For example, if `p4ssw0rd` is used as a password, it can be cracked.

Using secure passwords is vital for preventing unauthorized access to your FortiGate. When changing the password, consider the following to ensure better security:

- Do not use passwords that are obvious, such as the company name, administrator names, or other obvious words or phrases.
- Use numbers in place of letters, for example: `passw0rd`.
- Administrator passwords can be up to 64 characters.
- Include a mixture of numbers, symbols, and upper and lower case letters.

Netskope Exhibit 1017

- Use multiple words together, or possibly even a sentence, for example: `correcthorsebatterystaple`.
- Use a password generator.
- Change the password regularly and always make the new password unique and not a variation of the existing password. for example, do not change from `password` to `password1`.
- Make note of the password and store it in a safe place away from the management computer, in case you forget it; or ensure at least two people know the password in the event one person becomes unavailable. Alternatively, have two different admin logins.

FortiGate allows you to create a password policy for administrators and IPsec pre-shared keys. With this policy, you can enforce regular changes and specific criteria for a password policy, including:

- Minimum length between 8 and 64 characters.
- If the password must contain uppercase (A, B, C) and/or lowercase (a, b, c) characters.
- If the password must contain numbers (1, 2, 3).
- If the password must contain special or non-alphanumeric characters (!, @, #, $, %, ^, &, *, (, and)).
- Where the password applies (admin or IPsec or both).
- The duration of the password before a new one must be specified.

If you add a password policy or change the requirements on an existing policy, the next time that administrator logs into the FortiGate, the administrator is prompted to update the password to meet the new requirements before proceeding to log in.

For information about setting passwords, see Default administrator password on page 602.

To create a system password policy the GUI:

1. Go to *System* > *Settings*.
2. In the *Password Policy* section, change the *Password scope* to *Admin*, *IPsec*, or *Both*.
3. Configure the password policy options.
4. Click *Apply*.

To create a system password policy the CLI:

```
config system password-policy
    set status {enable | disable}
    set apply-to {admin-password | ipsec-preshared-key}
    set minimum-length <8-128>
    set min-lower-case-letter <0-128>
    set min-upper-case-letter <0-128>
    set min-non-alphanumeric <0-128>
    set min-number <0-128>
    set change-4-characters {enable | disable}
    set expire-status {enable | disable}
    set expire-day <1-999>
    set reuse-password {enable | disable}
end
```

Admin profile option for diagnose access

The `system-diagnostics` command in an administrator profile can be used to control access to diagnose commands for global and VDOM level administrators.

To block an administrator's access to diagnose commands:

1. Create an admin profile that cannot access diagnose commands:

```
config system accprofile
    edit "nodiagnose"
        ...
        set system-diagnostics disable
    next
end
```

2. Apply the profile to an administrator:

```
config system admin
    edit "nodiag"
        set accprofile "nodiagnose"
        set vdom "root"
        set password ********
    next
end
```

3. Log in as that administrator and confirm that they cannot access diagnose commands:

```
$ ?
config     Configure object.
get        Get dynamic and system information.
show       Show configuration.
execute    Execute static commands.
alias      Execute alias commands.
exit       Exit the CLI.
```

Firmware

Fortinet periodically updates the FortiGate firmware to include new features and resolve important issues. After you have registered your FortiGate unit, firmware updates can be downloaded from the Fortinet Customer Service & Support website.



Always back up the current configuration before installing new firmware. See Configuration backups on page 108.

Before you install any new firmware, follow the below steps:

1. Review the Release Notes for a new firmware release.
2. Review the Supported Upgrade Paths.
3. Download a copy of the currently installed firmware, in case you need to revert to it. See Downloading a firmware image on page 595 and Downgrading to a previous firmware version on page 599 for details.
4. Have a plan in place in case there is a critical failure, such as the FortiGate not coming back online after the update. This could include having console access to the device (Connecting to the CLI on page 24), ensuring that you TFTP server is working (Installing firmware from system reboot on page 600), and preparing a USB drive (Restoring from a USB drive on page 601).
5. Backup the current configuration, including local certificates. SeeConfiguration backups on page 108 for details.

6. Test the new firmware until you are satisfied that it applies to your configuration. See Testing a firmware version on page 597 and Controlled upgrade on page 602 for details.

Installing new firmware without reviewing release notes or testing the firmware may result in changes to settings and unexpected issues.



Only FortiGate admin users and administrators whose access profiles contain system read and write privileges can change the FortiGate firmware.

Firmware upgrade notifications

FortiGates with a firmware upgrade license that are connected to FortiGuard display upgrade notifications in the setup window, banner, and FortiGuard menu. The firmware notifications are enabled by default.

To configure firmware notifications in the CLI:

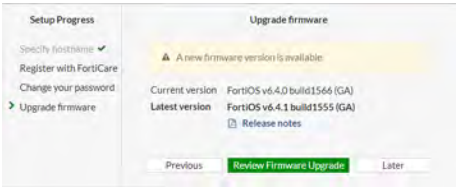
```
config system global
    set gui-firmware-upgrade-setup-warning {enable | disable}
end
```

To use the firmware upgrade notifications in the GUI:

1. When you log in to FortiGate, the *FortiGate Setup* window includes an *Upgrade firmware* step. Click *Begin*.

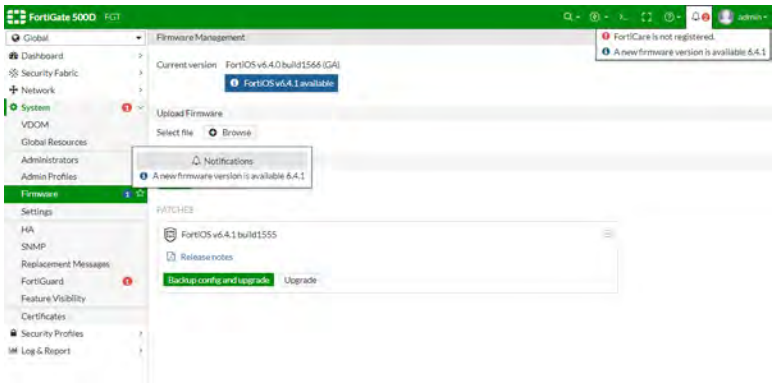


2. Follow the steps in the *Setup Progress*, then click *Review Firmware Upgrade*.



The *System* > *Firmware* page opens.

3. Notifications appear below the *Notification* icon in the banner, and beside *Firmware* in the tree menu.



Downloading a firmware image

Firmware images for all FortiGate units are available on the Fortinet Customer Service & Support website.

To download firmware:

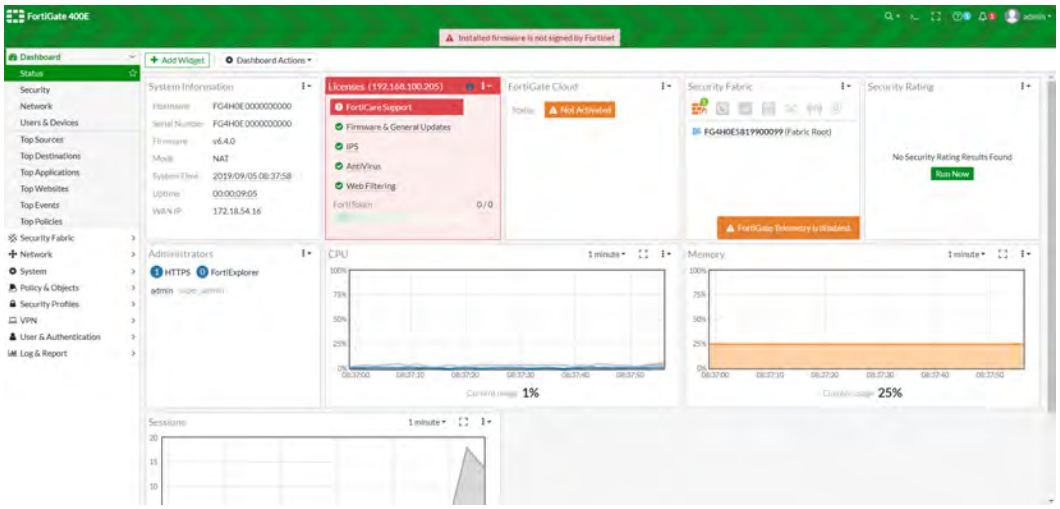
1. Log into the support site with your user name and password.
2. Go to *Download* > *Firmware Images*.
   A list of Release Notes is shown. If you have not already done so, download and review the Release Notes for the firmware version that you are upgrading your FortiGate unit to.
3. Select the *Download* tab.
4. Navigate to the folder for the firmware version that you are upgrading to.
5. Find your device model from the list. FortiWiFi devices have file names that start with *FWF*.
6. Click *HTTPS* in the far right column to download the firmware image to your computer.

Firmware can also be downloaded using FTP, but as FTP is not an encrypted file transferring protocol, HTTPS downloading is recommended.

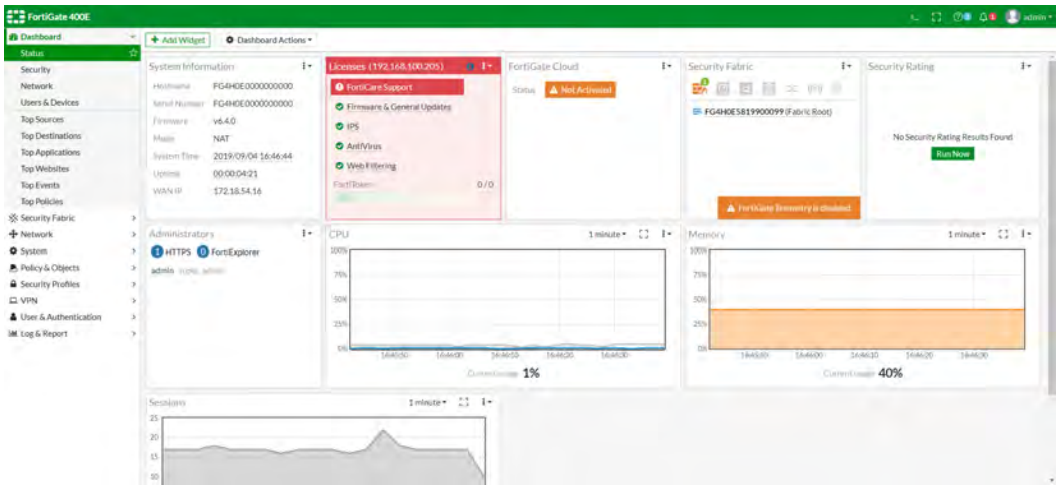
FortiOS image signing and verification

Official FortiOS firmware images are signed by the Fortinet CA. The BIOS checks the validity of an image when it is uploaded to the device. If the image is not signed by the Fortinet CA, a warning message is shown in the GUI.

Unsigned image:



Signed image:



This functionality is implemented on the following devices:

FortiGate 40F	FortiWiFi 41F	FortiGate 100F	FortiGate 3400E
FortiGate 40F-3G4G	FortiWiFi 41F-3G4G	FortiGate 101F	FortiGate 3401E
FortiGate 41F	FortiGate 60F	FortiGate 400E	FortiGate 3600E
FortiGate 41F-3G4G	FortiGate 61F	FortiGate 401E	FortiGate 3601E
FortiWiFi 40F	FortiWiFi 60F	FortiGate 600E	
FortiWiFi 40F-3G4G	FortiWiFi 61F	FortiGate 601E	

Testing a firmware version

The integrity of firmware images downloaded from Fortinet's support portal can be verified using a file checksum. A file checksum that does not match the expected value indicates a corrupt file. The corruption could be caused by errors in transfer or by file modification. A list of expected checksum values for each build of released code is available on Fortinet's support portal.

Image integrity is also verified when the FortiGate is booting up. This integrity check is done through a cyclic redundancy check (CRC). If the CRC fails, the FortiGate unit will encounter an error during the boot process.

Firmware images are signed and the signature is attached to the code as it is built. When upgrading an image, the running OS will generate a signature and compare it with the signature attached to the image. If the signatures do not match, the new OS will not load.

Testing before installation

FortiOS lets you test a new firmware image by installing the firmware image from a system reboot and saving it to system memory. After completing this procedure, the FortiGate unit operates using the new firmware image with the current configuration. The new firmware image is not permanently installed. The next time the FortiGate unit restarts, it operates with the originally installed firmware image using the current configuration. If the new firmware image operates successfully, you can install it permanently using the procedure explained in Upgrading the firmware.

For this procedure, you must install a TFTP server that you can connect to from the FortiGate internal interface. The TFTP server should be on the same subnet as the internal interface.

To test the new firmware version:

1. Connect to the CLI using an RJ-45 to USB (or DB-9) or null modem cable.
2. Ensure that the TFTP server is running.
3. Copy the new firmware image file to the root directory on the TFTP server.
4. Ensure that the FortiGate unit can connect to the TFTP server using the `execute ping` command.
5. Restart the FortiGate unit: `execute reboot`. The following message is shown:
   ```
   This operation will reboot the system!
   Do you want to continue? (y/n)
   ```
6. Type `y`. As the FortiGate unit starts, a series of system startup messages appears.
7. When the following messages appears:
   ```
   Press any key to display configuration menu.....
   ```
   Immediately press any key to interrupt the system startup.

   You have only three seconds to press any key. If you do not press a key during this time, the FortiGate will reboot, and you will have to log in and repeat the `execute reboot` command.

   If you successfully interrupt the startup process, the following messages appears:
   ```
   [G]: Get firmware image from TFTP server.
   [F]: Format boot device.
   [B]: Boot with backup firmware and set as default
   [C]: Configuration and information
   [Q]: Quit menu and continue to boot with default firmware.
   [H]: Display this list of options.
   Enter G, F, Q, or H:
   ```
8. Type *G* to get the new firmware image from the TFTP server. The following message appears: `Enter TFTP server address [192.168.1.168]:`

9. Type the address of the TFTP server, then press *Enter*. The following message appears: `Enter Local Address [192.168.1.188]:`
10. Type the IP address of the FortiGate unit to connect to the TFTP server.



The IP address must be on the same network as the TFTP server.
Make sure that you do not enter the IP address of another device on this network.

The following message appears:

```
Enter File Name [image.out]:
```

11. Enter the firmware image file name then press *Enter*. The TFTP server uploads the firmware image file to the FortiGate unit and the following message appears:

```
Save as Default firmware/Backup firmware/Run image without saving: [D/B/R]
```

12. Type *R*. The FortiGate image is installed to system memory and the FortiGate unit starts running the new firmware image, but with its current configuration.

Test the new firmware image as required. When done testing, reboot the FortiGate unit, and the it will resume using the firmware that was running before you installed the test firmware.

Upgrading the firmware

Installing a new firmware image replaces the current antivirus and attack definitions, along with the definitions included with the firmware release that is being installing. After you install new firmware, make sure that the antivirus and attack definitions are up to date.



Back up your configuration before making any firmware changes.

To upgrade the firmware in the GUI:

1. Log into the FortiGate GUI as the admin administrative user.
2. Go to *System* > *Firmware*.
3. Under *Upload Firmware*, click *Browse* and locate the previously downloaded firmware image file (see Downloading a firmware image on page 595).
4. Click *Backup config and upgrade*.
   The FortiGate unit backs up the current configuration to the management computer, uploads the firmware image file, upgrades to the new firmware version, and restarts. This process takes a few minutes.

To upgrade the firmware in the CLI:

1. Make sure that the TFTP server is running.
2. Copy the new firmware image file to the root directory of the TFTP server.
3. Log into the CLI.
4. Ping the TFTP server to ensure that the FortiGate can connect to it:

```
execute ping <tftp_ipv4>
```

5. Enter the following command to copy the firmware image from the TFTP server to the FortiGate unit:

```
execute restore image tftp <filename> <tftp_ipv4>
```

The FortiGate unit responds with the message:

```
This operation will replace the current firmware version!
Do you want to continue? (y/n)
```

6. Type `y`. The FortiGate unit uploads the firmware image file, upgrades to the new firmware version, and restarts. This process takes a few minutes.
7. Reconnect to the CLI.
8. Update the antivirus and attack definitions:
   ```
   execute update-now
   ```

Downgrading to a previous firmware version



Downgrading the firmware is not recommended.

This procedure downgrades the FortiGate to a previous firmware version. The backup configuration might not be able to be restored after downgrading.

To downgrade to a previous firmware version in the GUI:

1. Log into the FortiGate GUI as the admin administrative user.
2. Go to *System* > *Firmware*.
3. Under *Upload Firmware*, click *Browse* and locate the previously downloaded firmware image file (see Downloading a firmware image on page 595).
4. Click *Confirm version downgrade*.
5. Click *Backup config and downgrade*.
   The FortiGate unit backs up the current configuration to the management computer, uploads the firmware image file, upgrades to the new firmware version, and restarts. This process takes a few minutes.

To downgrade to a previous firmware version in the CLI:

1. Make sure that the TFTP server is running.
2. Copy the new firmware image file to the root directory of the TFTP server.
3. Log into the CLI.
4. Ping the TFTP server to ensure that the FortiGate can connect to it:
   ```
   execute ping <tftp_ipv4>
   ```
5. Enter the following command to copy the firmware image from the TFTP server to the FortiGate unit:
   ```
   execute restore image tftp <filename> <tftp_ipv4>
   ```
   The FortiGate unit responds with the message:
   ```
   This operation will replace the current firmware version!
   Do you want to continue? (y/n)
   ```
6. Type `y`. The FortiGate unit uploads the firmware image file, then a message similar to the following is shown:
   ```
   Get image from tftp server OK.
   Check image OK.
   This operation will downgrade the current firmware version!
   Do you want to continue? (y/n)
   ```
7. Type `y`. The FortiGate unit downgrades to the old firmware version and restarts. This process takes a few minutes.
8. Reconnect to the CLI.

9. Update the antivirus and attack definitions:
   ```
   execute update-now
   ```

Installing firmware from system reboot

In the event that the firmware upgrade does not load properly and the FortiGate unit will not boot, or continuously reboots, it is best to perform a fresh install of the firmware from a reboot using the CLI. If configured, the firmware can also be automatically installed from a USB drive; see Restoring from a USB drive on page 601 for details.

This procedure installs a firmware image and resets the FortiGate unit to factory default settings. You can use this procedure to upgrade to a new firmware version, revert to an older firmware version, or re-install the current firmware.

To use this procedure, you must connect to the CLI using the FortiGate console port and a RJ-45 to USB (or DB-9), or null modem cable. You must also install a TFTP server that you can connect to from the FortiGate internal interface. The TFTP server should be on the same subnet as the internal interface.

Before beginning this procedure, ensure that you backup the FortiGate unit configuration. See Configuration backups on page 108 for details. If you are reverting to a previous FortiOS version, you might not be able to restore the previous configuration from the backup configuration file.

Installing firmware replaces your current antivirus and attack definitions, along with the definitions included with the firmware release you are installing. After you install new firmware, make sure that antivirus and attack definitions are up to date.

To install firmware from a system reboot:

1. Connect to the CLI using the RJ-45 to USB (or DB-9) or null modem cable.
2. Ensure that the TFTP server is running.
3. Copy the new firmware image file to the root directory of the TFTP server.
4. Ensure that the FortiGate unit can connect to the TFTP server using the `execute ping` command.
5. Restart the FortiGate unit: `execute reboot`. The following message is shown:
   ```
   This operation will reboot the system!
   Do you want to continue? (y/n)
   ```
6. Type `y`. As the FortiGate unit starts, a series of system startup messages appears.
7. When the following messages appears:
   ```
   Press any key to display configuration menu.....
   ```
   Immediately press any key to interrupt the system startup.

   You have only three seconds to press any key. If you do not press a key during this time, the FortiGate will reboot, and you will have to log in and repeat the `execute reboot` command.

   If you successfully interrupt the startup process, the following messages appears:

   ```
   [C]: Configure TFTP parameters.
   [R]: Review TFTP parameters.
   [T]: Initiate TFTP firmware transfer.
   [F]: Format boot device.
   [I]: System information.
   [B]: Boot with backup firmware and set as default.
   [Q]: Quit menu and continue to boot.
   [H]: Display this list of options.

   Enter C,R,T,F,I,B,Q,or H:
   ```


Netskope Exhibit 1017

8. If necessary, type `C` to configure the TFTP parameters, then type `Q` to return to the previous menu:

```
[P]: Set firmware download port.
[D]: Set DHCP mode.
[I]: Set local IP address.
[S]: Set local subnet mask.
[G]: Set local gateway.
[V]: Set local VLAN ID.
[T]: Set remote TFTP server IP address.
[F]: Set firmware file name.
[E]: Reset TFTP parameters to factory defaults.
[R]: Review TFTP parameters.
[N]: Diagnose networking(ping).
[Q]: Quit this menu.
[H]: Display this list of options.

Enter P,D,I,S,G,V,T,F,E,R,N,Q,or H:
```



The IP address must be on the same network as the TFTP server.
Make sure that you do not enter the IP address of another device on this network.

9. Type `T` get the new firmware image from the TFTP server.
   The FortiGate unit loads the firmware.
10. Save the firmware as the default (`D`) or backup (`B`) firmware image, or run the image without saving it (`R`).
    The FortiGate unit installs the new firmware image and restarts. The installation might take a few minutes to complete.

Restoring from a USB drive

The FortiGate firmware can be manually restored from a USB drive, or installed automatically from a USB drive after a reboot.

To restore the firmware from a USB drive:

1. Copy the firmware file to the root directory on the USB drive.
2. Connect the USB drive to the USB port of the FortiGate device.
3. Connect to the FortiGate CLI using the RJ-45 to USB (or DB-9) or null modem cable.
4. Enter the following command:
   ```
   execute restore image usb <filename>
   ```
   The FortiGate unit responds with the following message:
   ```
   This operation will replace the current firmware version! Do you want to continue?
   (y/n)
   ```
5. Type `y`. The FortiGate unit restores the firmware and restarts. This process takes a few minutes.
6. Update the antivirus and attack definitions:
   ```
   execute update-now
   ```

To install firmware automatically from a USB drive:

1. Go to *System* > *Settings*.
2. In the *Start Up Settings* section, enable *Detect firmware* and enter the name of the firmware file.

Netskope Exhibit 1017

3. Copy the firmware file to the root directory on the USB drive.
4. Connect the USB drive to the USB port of the FortiGate device.
5. Reboot the FortiGate device.

Controlled upgrade

Using a controlled upgrade, you can upload a new version of the FortiOS firmware to a separate partition in the FortiGate memory for later upgrade. The FortiGate unit can be configured so that when it is rebooted, it will automatically load the new firmware. Using this option, you can stage multiple FortiGate units to upgrade simultaneously using FortiManager or a script.

To load the firmware for later installation:

```
execute restore secondary-image {ftp | tftp | usb} <filename_str>
```

To set the FortiGate unit so that when it reboots, the new firmware is loaded:

```
execute set-next-reboot {primary | secondary}
```

where `{primary | secondary}` is the partition with the preloaded firmware.

Settings

The default administrator password should be configured immediately after the FortiGate is installed, see Default administrator password on page 602.

After that, there are several system settings that should also be configured in *System* > *Settings*:

- Default administrator password on page 602
- Changing the host name on page 603
- Setting the system time on page 603
- Configuring ports on page 607
- Setting the idle timeout time on page 608
- Setting the password policy on page 608
- Changing the view settings on page 608
- Setting the administrator password retries and lockout time on page 609

Default administrator password

By default, your FortiGate has an administrator account set up with the username `admin` and no password. In order to prevent unauthorized access to the FortiGate, it is highly recommended that you add a password to this account.



In FortiOS 6.2.1 and later, adding a password to the *admin* administrator is mandatory. You will be prompted to configured it the first time you log in to the FortiGate using that account, after a factory reset, and after a new image installation.

To change the default password in the GUI:

1. Go to *System* > *Administrators*.
2. Edit the *admin* account.
3. Click *Change Password*.
4. If applicable, enter the current password in the *Old Password* field.
5. Enter a password in the *New Password* field, then enter it again in the *Confirm Password* field.
6. Click *OK*.

To change the default password in the CLI:

```
config system admin
    edit admin
        set password <password>
    next
end
```



It is also recommended that you change the user name of this account; however, since you cannot change the user name of an account that is currently in use, a second administrator account must be created in order to do this.

Changing the host name

The FortiGate host name is shown in the *Hostname* field in the *System Information* widget on a dashboard, as the command prompt in the CLI, as the SNMP system name, as the device name on FortiGate Cloud, and other places. If the FortiGate is in an HA cluster, use a unique host name to distinguish it from the other devices in the cluster.

An administrator requires *System* > *Configuration* read/write access to edit the host name. See Administrator profiles on page 587 for details.

To change the host name in the GUI:

1. Go to *System* > *Settings*.
2. In the *Host name* field, enter a new name.
3. Click *Apply*.

To change the host name in the CLI:

```
config system global
    set hostname <hostname>
end
```

Setting the system time

You can either manually set the FortiOS system time, or configure the device to automatically keep its system time correct by synchronizing with a Network Time Protocol (NTP) server.

Daylight savings time is enabled by default, and can only be configured in the CLI.



For many features to work, including scheduling, logging, and SSL-dependent features, the FortiOS system time must be accurate.

To configure the date and time in the GUI:

1. Go to *System* > *Settings*.
2. In the *System Time* section, configure the following settings to either manually set the time or use an NTP server:

Time Zone		Select a time zone from the list. This should be the time zone that the FortiGate is in.
Set Time		Select to either *Synchronize with an NTP Server*, or use *Manual settings*.
	Synchronize with an NTP Server	To use an NTP server other than FortiGuard, the CLI must be used. In the *Sync interval* field, enter how often, in minutes, that the device synchronizes its time with the NTP server.
	Manual settings	Manually enter the *Date*, *Hour* (in 24-hour format), *Minute*, and *Second* in their fields.
Setup device as local NTP server		Enable to configure the FortiGate as a local NTP server. In the *Listen on Interfaces* field, set the interface or interfaces that the FortiGate will listen for NTP requests on.

3. Click *Apply*.

To configure the date and time in the CLI:

1. Configure the timezone and daylight savings time:

```
config system global
    set timezone <integer>
    set dst {enable | disable}
end
```

2. Either manually configure the date and time, or configure an NTP server:

   Manual:

```
execute date <yyyy-mm-dd>
execute time <hh:mm:ss>
```

   NTP server:

```
config system ntp
    set ntpsync enable
    set type {fortiguard | custom}
    set syncinterval <integer>
    set source-ip <ip_address>
    set source-ip6 <ip6_address>
    set server-mode {enable | disable}
    set interface <interface>
    set authentication {enable | disable}
    set key-type {MD5 | SHA1}
```


Netskope Exhibit 1017

```
        set key <password>
        set key-id <integer>
        config ntpserver
            edit <server_id>
                set server <ip_address or hostname>
                set ntpv3 {enable | disable}
                set authentication {enable | disable}
                set key <password>
                set key-id <integer>
            next
        end
    end
```

SHA-1 authentication support (for NTPv4)

SHA-1 authentication support allows the NTP client to verify that severs are known and trusted and not intruders masquerading (accidentally or intentionally) as legitimate servers. In cryptography, SHA-1 is a cryptographic hash algorithmic function.



SHA-1 authentication support is only available for NTP clients, not NTP servers.

To configure authentication on a FortiGate NTP client:

```
config system ntp
    set ntpsync enable
    set type custom
    set syncinterval 1
    config ntpserver
        edit "883502"
            set server "10.1.100.11"
            set authentication enable
            set key
ENCi9NmcqsV3xBJvOkgIL3lFxA8mnNs2XKfB7spOQoUw4cm8FOOP0nrCbqx6rJ+om95+hVUHpaVZmepdd4KznPlAHNiu
liPgPOk
            set key-id 1
        next
    end
end
```

Command	Description
`authentication <enable | disable>`	Enable/disable MD5/SHA1 authentication (default = disable).
`key <passwd>`	Key for MD5/SHA1 authentication. Enter a password value.
`key-id <integer>`	Key ID for authentication. Enter an integer value from 0 to 4294967295.

To confirm that NTP authentication is set up correctly:

```
# diagnose sys ntp status
synchronized: yes, ntpsync: enabled, server-mode: disabled
ipv4 server(10.1.100.11) 10.1.100.11 -- reachable(0xff) S:4 T:6 selected
server-version=4, stratum=3
```

If NTP authentication is set up correctly, the server version is equal to 4.

PTPv2

Precision time protocol (PTP) is used to synchronize network clocks. It is best suited to situations where time accuracy is of the utmost importance, as it supports accuracy in the sub-microsecond range. Conversely, NTP accuracy is in the range of milliseconds or tens of milliseconds.

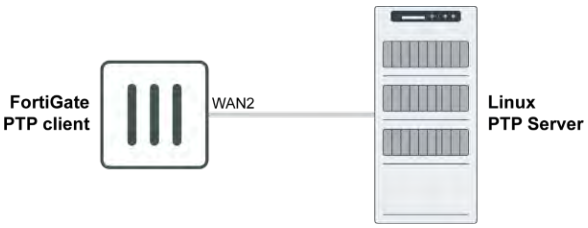
The following CLI commands are available:

```
config system ptp
  set status {enable | disable}
  set mode {multicast | hybrid}
  set delay-mechanism {E2E | P2P}
  set request-interval <integer>
  set interface <interface>
end
```

Command	Description
status {enable | disable}	Enable or disable the FortiGate system time by synchronizing with a PTP server (default = disable).
mode {multicast | hybrid}	Use multicast or hybrid transmission (default = multicast).
delay-mechanism {E2E | P2P}	Use end-to-end (E2E) or peer-to-peer (P2P) delay detection (default = E2E).
request-interval <integer>	The logarithmic mean interval between the delay request messages sent by the client to the server in seconds (default = 1).
interface <interface>	The interface that the PTP client will reply through.

Sample configuration

This example uses the following topology:



To configure a FortiGate to act as a PTP client that synchronizes itself with a Linux PTP server:

1. Enable debug messages:
   ```
   diagnose debug application ptpd -1
   ```

This command will provide details to debug the PTP communication with the server.

2. Check the system date:
```
execute date
   current date is: 2019-01-01
```
3. Configure PTP in global mode:
```
config system ptp
   set status enable
   set interface wan2
end
```
4. Check the system date again after synchronization with the PTP server:
```
execute date
   current date is: 2019-01-14
```

Configuring ports

To improve security, the default ports for administrative connections to the FortiGate can be changed. Port numbers must be unique. If a conflict exists with a particular port, a warning message is shown.

When connecting to the FortiGate after a port has been changed, the port number be included, for example: `https://192.168.1.99:100`.

To configure the ports in the GUI:

1. Go to *System* > *Settings*.
2. In the *Administration Settings* section, set the HTTP, HTTPS, SSH, and Telnet ports.
3. Enable *Redirect to HTTPS* to prevent HTTP from being used by administrators.
4. Click *Apply*.

To configure the ports in the CLI:

```
config system global
    set admin-port <port>
    set admin-sport <port>
    set admin-https-redirect {enable | disable}
    set admin-ssh-port <port>
    set admin-telnet-port <port>
end
```

Custom default service port range

The default service port range can be customized using the following CLI command:

```
config system global
   set default-service-source-port <port range>
end
```

Where `<port range>` is the new default service port range, that can have a minimum value of 0 and a maximum value up to 65535. The default value is 1 to 65535.

Netskope Exhibit 1017



This change effects the TCP/UDP protocol.

Setting the idle timeout time

The idle timeout period is the amount of time that an administrator will stay logged in to the GUI without any activity. This is to prevent someone from accessing the FortiGate if the management PC is left unattended. By default, it is set to five minutes.



A setting of higher than 15 minutes will have a negative effect on a security rating score. See Security rating on page 195 for more information.

To change the idle timeout in the GUI:

1. Go to *System* > *Settings*.
2. In the *Administration Settings* section, set the *Idle timeout* to up to 480 minutes.
3. Click *Apply*.

To change the idle timeout in the CLI:

```
config system global
    set admintimeout <integer>
end
```

Setting the password policy

A password policy can be created for administrators and IPsec pre-shared keys. See Password policy on page 591 for information.

Changing the view settings

The view settings change the look and language of the FortiOS GUI.

To change the view settings in the GUI:

1. Go to *System* > *Settings*.
2. In the *View Settings* section, configure the following settings:

Language	Set the GUI language: *English*, *French*, *Spanish*, *Portuguese*, *Japanese*, *Traditional Chinese*, *Simplifies Chinese*, *Korean*.
Lines per page	Set the number of lines per page, from 20 to 100.

Theme	Set the theme color: *Green*, *Red*, *Blue*, *Melongene*, or *Mariner*.
Date/Time Display	Set the date and time to display using the FortiGate's or the browser's timezone.
NGFW Mode	Set the NGFW mode to either *Profile-based* (default) or *Policy-based*. If *Policy-based* is selected, the *SSL/SSH Inspection* profile must be selected.

3. Click *Apply*.

To change the view settings in the CLI:

```
config system global
    set language {english | french | spanish | portuguese | japanese | trach | simch |
korean}
    set gui-lines-per-page <integer>
    set gui-theme {green | red | blue | melongene | mariner}
    set gui-date-time-source {system | browser}
end
config system settings
    set ngfw-mode {profile-based | policy-based}
    set ssl-ssh-profile {certificate-inspection | custom-deep-inspection | deep-inspection |
no-inspection}
end
```

Setting the administrator password retries and lockout time

By default, the number password retry attempts is set to three, allowing the administrator a maximum of three attempts at logging in to their account before they are locked out for a set amount of time (by default, 60 seconds).

The number of attempts and the default wait time before the administrator can try to enter a password again can be configured using the CLI.

A maximum of ten retry attempts can be configured, and the lockout period can be 1 to 2147483647 seconds (over 68 years). The higher the retry attempts, the higher the risk that someone might be able to guess the password.

To configure the lockout options:

```
config system global
    set admin-lockout-threshold <failed_attempts>
    set admin-lockout-duration <seconds>
end
```

Example:

To set the number of retry attempts to 1, and the lockout time to 5 minutes, enter the following commands:

```
config system global
    set admin-lockout-threshold 1
    set admin-lockout-duration 300
end
```



If the time span between the first failed log in attempt and the lockout threshold failed attempt is less than lockout time, the lockout will be triggered.

Virtual Domains

Virtual Domains (VDOMs) are used to divide a FortiGate into two or more virtual units that function independently. VDOMs can provide separate security policies and, in NAT mode, completely separate configurations for routing and VPN services for each connected network.

There are two VDOM modes:

- Split-task VDOM mode: One VDOM is used only for management, and the other is used to manage traffic. See Split-task VDOM mode on page 611.
- Multi VDOM mode: Multiple VDOMs can be created and managed as independent units. See Multi VDOM mode on page 615.

By default, most FortiGate units support 10 VDOMs, and many FortiGate models support purchasing a license key to increase the maximum number.

Global settings are configured outside of a VDOM. They effect the entire FortiGate, and include settings such as interfaces, firmware, DNS, some logging and sandboxing options, and others. Global settings should only be changed by top level administrators.



Enable the following to prevent accidentally creating VDOMs in the CLI:

```
config system global
    set edit-vdom-prompt enable
end
```

The FortiGate displays a prompt to confirm before the VDOM is created.

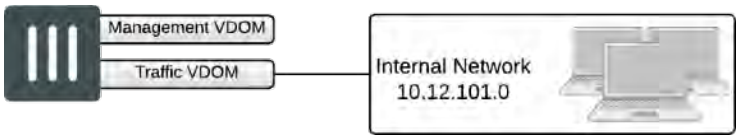
Switching VDOM modes

Current VDOM mode	New VDOM mode	Rule
No VDOM	Split-task VDOM	Allowed
Split-task VDOM	No VDOM	Allowed
No VDOM	Multi VDOM	Allowed only if the FortiGate is not a member of a Security Fabric. See Configuring the root FortiGate and downstream FortiGates on page 115 for more information.
Multi VDOM	No VDOM	Allowed

Current VDOM mode	New VDOM mode	Rule
Split-task VDOM	Multi VDOM	Allowed only if the FortiGate is not a member of a Security Fabric. See Configuring the root FortiGate and downstream FortiGates on page 115 for more information.
Multi VDOM	Split-task VDOM	Not Allowed. User must first switch to *No VDOM*

Split-task VDOM mode

In split-task VDOM mode, the FortiGate has two VDOMs: the management VDOM (*root*) and the traffic VDOM (*FG-traffic*).



The management VDOM is used to manage the FortiGate, and cannot be used to process traffic.

The following GUI sections are available when in the management VDOM:

- The Status dashboard
- Security Fabric topology and settings (read-only, except for *HTTP Service* settings)
- Interface and static route configuration
- FortiClient configuration
- Replacement messages
- Certificates
- System events
- Log and email alert settings
- Threat weight definitions

The traffic VDOM provides separate security policies, and is used to process all network traffic.

The following GUI sections are available when in the traffic VDOM:

- The Status, Top Usage LAN/DMZ, and Security dashboards
- Security Fabric topology, settings (read-only, except for *HTTP Service* settings), and External Connectors (*Endpoint/Identity* connectors only)
- FortiView
- Interface configuration
- Packet capture
- SD-WAN, SD-WAN Rules, and Performance SLA
- Static and policy routes
- RIP, OSPF, BGP, and Multicast
- Replacement messages
- Feature visibility
- Tags
- Certificates

- Policies and objects
- Security profiles
- VPNs
- User and device authentication
- Wifi and switch controller
- Logging
- Monitoring

Split-task VDOM mode is not available on all FortiGate models. The Fortinet Security Fabric supports split-task VDOM mode.

Enable split-task VDOM mode

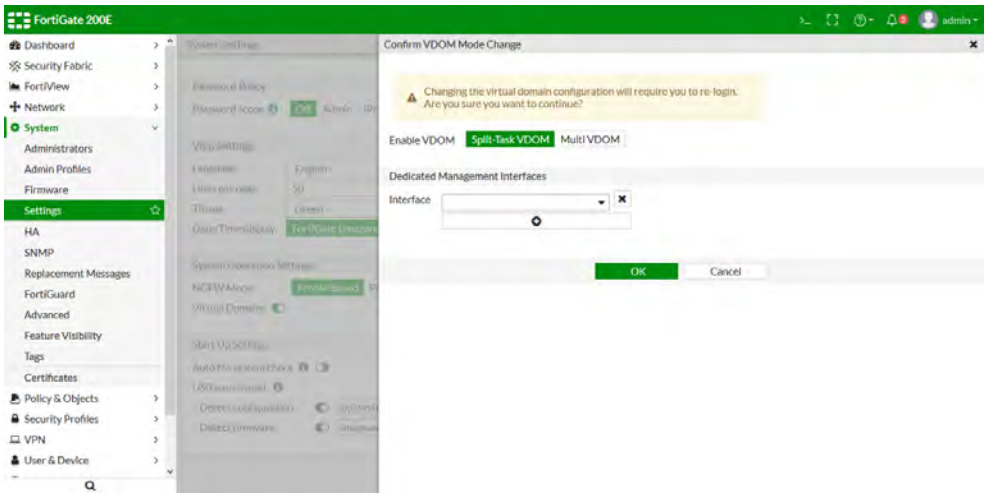
Split-task VDOM mode can be enabled in the GUI or CLI. Enabling it does not require a reboot, but does log you out of the FortiGate.

When split-task VDOM mode is enabled, all current management configuration is assigned to the *root* VDOM, and all non-management settings, such as firewall policies and security profiles, are deleted.

On FortiGate 60 series models and lower, VDOMs can only be enabled using the CLI.

To enable split-task VDOM mode in the GUI:

1. On the FortiGate, go to *System* > *Settings*.
2. In the *System Operation Settings* section, enable *Virtual Domains*.



3. Select *Split-Task VDOM* for the VDOM mode.
4. Select a *Dedicated Management Interface* from the *Interface* list. This interface is used to access the management VDOM, and cannot be used in firewall policies.
5. Click *OK*.

To enable split-task VDOM mode with the CLI:

```
config system global
   set vdom-mode split-vdom
end
```

Assign interfaces to a VDOM

An interface can only be assigned to one of the VDOMs. When split-task VDOM mode is enabled, all interfaces are assigned to the *root* VDOM. To use an interface in a policy, it must first be assigned to the traffic VDOM.

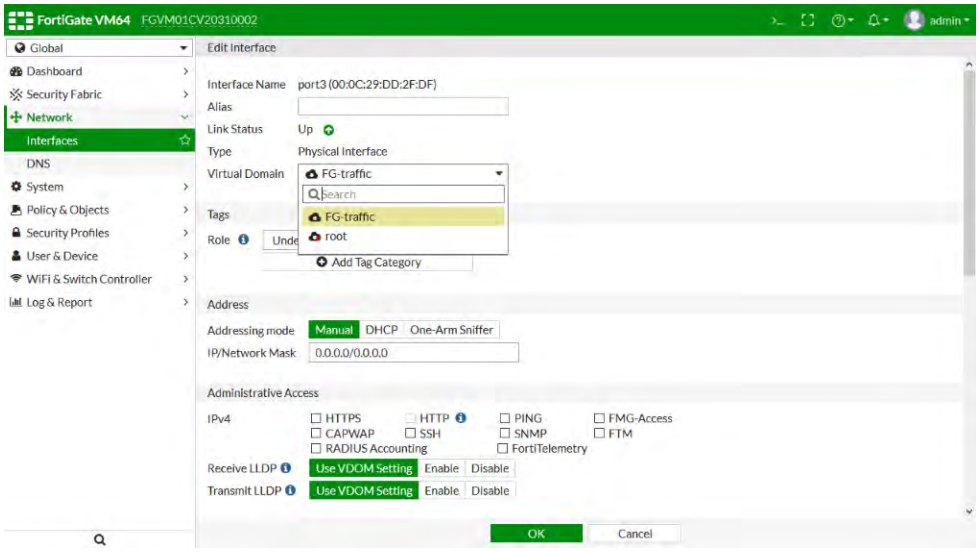
An interface cannot be moved if it is referenced in an existing configuration.



In the GUI, the interface list *Ref.* column shows if the interface is referenced in an existing configuration, and allows you to quickly access and edit those references.

To assign an interface to a VDOM in the GUI:

1. On the FortiGate, go to *Global* > *Network* > *Interfaces*.
2. Edit the interface that will be assigned to a VDOM.
3. Select the VDOM that the interface will be assigned to from the *Virtual Domain* list.



4. Click *OK*.

To assign an interface to a VDOM using the CLI:

```
config global
   config system interface
      edit <interface>
         set vdom <VDOM_name>
      next
   end
end
```

Create per-VDOM administrators

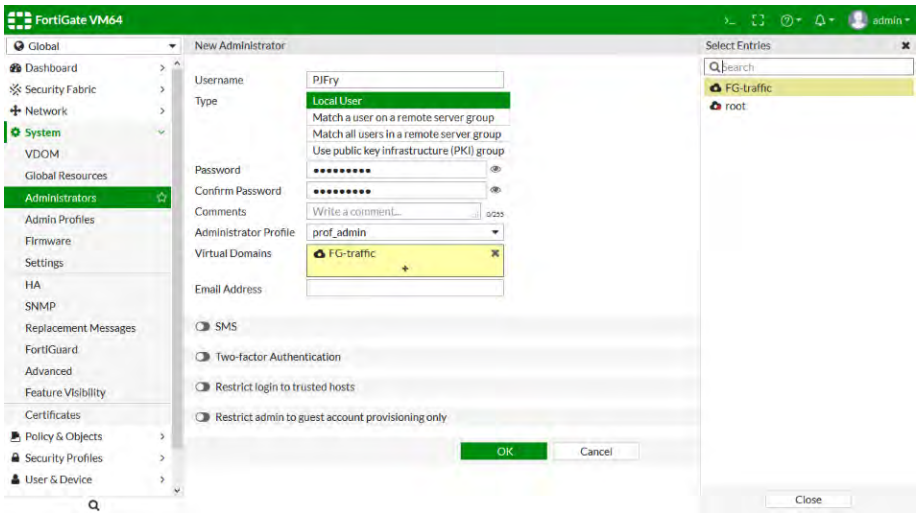
Per-VDOM administrators can be created that can access only the management or traffic VDOM. These administrators must use either the *prof_admin* administrator profile, or a custom profile.

A per-VDOM administrator can only access the FortiGate through a network interface that is assigned to the VDOM that they are assigned to. The interface must also be configured to allow management access. They can also connect to the FortiGate using the console port.

To assign an administrator to multiple VDOMs, they must be created at the global level. When creating an administrator at the VDOM level, the *super_admin* administrator profile cannot be used.

To create a per-VDOM administrator in the GUI:

1. On the FortiGate, connect to the management VDOM.
2. Go to *Global* > *System* > *Administrators* and click *Create New* > *Administrator*.
3. Fill in the required information, setting the *Type* as *Local User*.
4. In the *Virtual Domains* field, add the VDOM that the administrator will be assigned to, and if necessary, remove the other VDOM from the list.



5. Click *OK*.

To create a per-VDOM administrator using the CLI:

```
config global
   config system admin
      edit <name>
         set vdom <VDOM_name>
         set password <password>
         set accprofile <admin_profile>
         ...
      next
   end
end
```

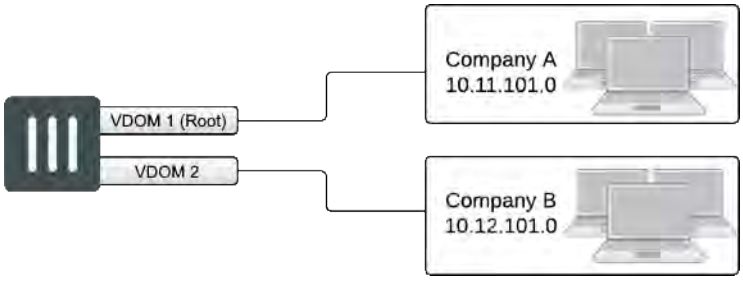
Multi VDOM mode

In multi VDOM mode, the FortiGate can have multiple VDOMs that function as independent units. One VDOM is used to manage global settings. The root VDOM cannot be deleted, and remains in the configuration even if it is not processing any traffic.

Multi VDOM mode isn't available on all FortiGate models. The Fortinet Security Fabric does not support multi VDOM mode.

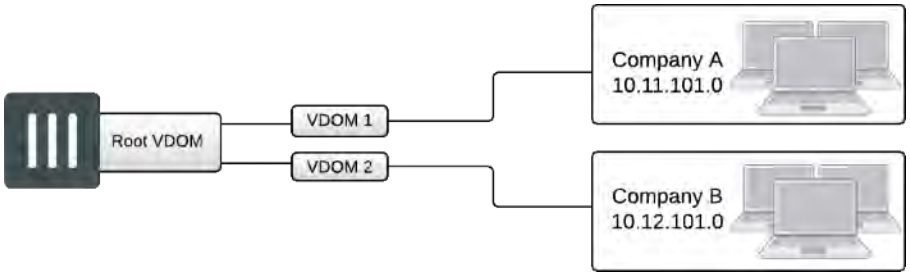
There are three main configuration types in multi VDOM mode:

Independent VDOMs:



Multiple, completely separate VDOMs are created. Any VDOM can be the management VDOM, as long as it has Internet access. There are no inter-VDOM links, and each VDOM is independently managed.

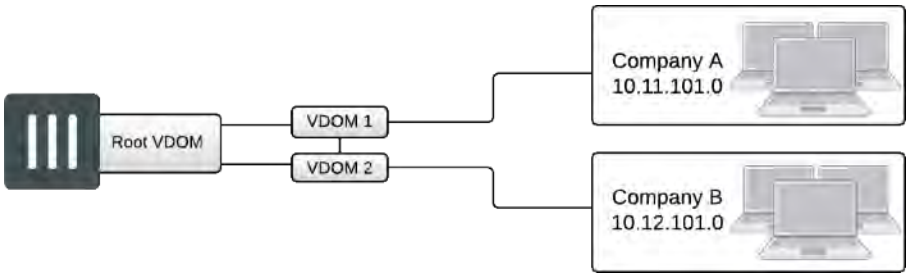
Management VDOM:



A management VDOM is located between the other VDOMs and the Internet, and the other VDOMs connect to the management VDOM with inter-VDOM links. The management VDOM has complete control over Internet access, including the types of traffic that are allowed in both directions. This can improve security, as there is only one point of ingress and egress.

There is no communication between the other VDOMs.

Meshed VDOMs:



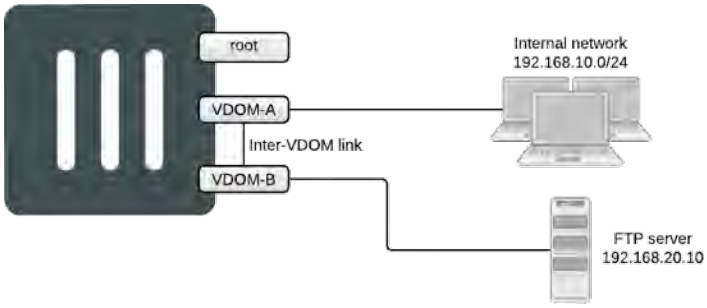
VDOMs can communicate with inter-VDOM links. In full-mesh configurations, all the VDOMs are interconnected. In partial-mesh configurations, only some of the VDOMs are interconnected.

In this configuration, proper security must be achieved by using firewall policies and ensuring secure account access for administrators and users.

Multi VDOM configuration examples

The following examples show how to configure per-VDOM settings, such as operation mode, routing, and security policies, in a network that includes the following VDOMs:

- VDOM-A: allows the internal network to access the Internet.
- VDOM-B: allows external connections to an FTP server.
- root: the management VDOM.



You can use VDOMs in either NAT or transparent mode on the same FortiGate. By default, VDOMs operate in NAT mode.

For both examples, multi VDOM mode must be enabled, and VDOM-A and VDOM-B must be created.

Enable multi VDOM mode

Multi VDOM mode can be enabled in the GUI or CLI. Enabling it does not require a reboot, but does log you out of the device. The current configuration is assigned to the *root* VDOM.



On FortiGate 60 series models and lower, VDOMs can only be enabled using the CLI.

To enable multi VDOM mode in the GUI:

1. On the FortiGate, go to *System* > *Settings*.
2. In the *System Operation Settings* section, enable *Virtual Domains*.
3. Select *Multi VDOM* for the VDOM mode.
4. Click *OK*.

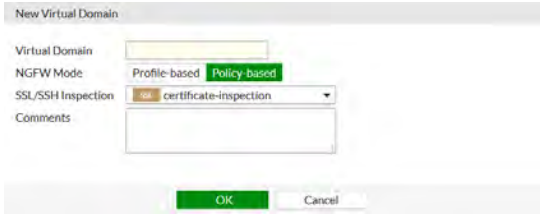
To enable multi VDOM mode with the CLI:

```
config system global
   set vdom-mode multi-vdom
end
```

Create the VDOMs

To create the VDOMs in the GUI:

1. In the *Global* VDOM, go to *System* > *VDOM*, and click *Create New*. The *New Virtual Domain* page opens.



2. In the *Virtual Domain* field, enter *VDOM-A*.
3. If required, set the *NGFW Mode*. If the *NGFW Mode* is *Policy-based*, select an *SSL/SSH Inspection* from the list.
4. Optionally, enter a comment.
5. Click *OK* to create the VDOM.
6. Repeat the above steps for *VDOM-B*.

To create the VDOMs with the CLI:

```
config vdom
   edit <VDOM-A>
      next
   edit <VDOM-B>
      next
   end
end
```

NAT mode

In this example, both VDOM-A and VDOM-B use NAT mode. A VDOM link is created that allows users on the internal network to access the FTP server.

This configuration requires the following steps:

1. Configure VDOM-A on page 618
2. Configure VDOM-B on page 620
3. Configure the VDOM link on page 622

Configure VDOM-A

VDOM-A allows connections from devices on the internal network to the Internet. WAN 1 and port 1 are assigned to this VDOM.

The per-VDOM configuration for VDOM-A includes the following:

- A firewall address for the internal network
- A static route to the ISP gateway
- A security policy allowing the internal network to access the Internet

All procedures in this section require you to connect to VDOM-A, either using a global or per-VDOM administrator account.

To add the firewall addresses in the GUI:

1. Go to *Policy & Objects* > *Addresses* and create a new address.
2. Enter the following information:

Name	internal-network
Type	Subnet
Subnet / IP Range	192.168.10.0/255.255.255.0
Interface	port1
Show in Address List	enabled

To add the firewall addresses with the CLI:

```
config vdom
  edit VDOM-A
    config firewall address
      edit internal-network
        set associated-interface port1
        set subnet 192.168.10.0 255.255.255.0
      next
    end
  next
end
```

To add a default route in the GUI:

1. Go to *Network* > *Static Routes* and create a new route.
2. Enter the following information:

Destination	Subnet
IP address	0.0.0.0/0.0.0.0
Gateway	172.20.201.7
Interface	wan1
Distance	10

To add a default route with the CLI:

```
config vdom
   edit VDOM-A
      config router static
         edit 0
            set gateway 172.20.201.7
            set device wan1
         next
      end
   next
end
```

To add the security policy in the GUI:

1. Connect to VDOM-A.
2. Go to *Policy & Objects* > *Firewall Policy* and create a new policy.
3. Enter the following information:

Name	VDOM-A-Internet
Incoming Interface	port1
Outgoing Interface	wan1
Source	internal-network
Destination	all
Schedule	always
Service	ALL
Action	ACCEPT
NAT	enabled

To add the security policy with the CLI:

```
config vdom
   edit VDOM-A
      config firewall policy
         edit 0
            set name VDOM-A-Internet
            set srcintf port1
            set dstintf wan1
            set srcaddr internal-network
            set dstaddr all
            set action accept
            set schedule always
            set service ALL
            set nat enable
         next
      end
   next
end
```


Netskope Exhibit 1017

Configure VDOM-B

VDOM-B allows external connections to reach an internal FTP server. WAN 2 and port 2 are assigned to this VDOM.

The per-VDOM configuration for VDOM-B includes the following:

- A firewall address for the FTP server
- A virtual IP address for the FTP server
- A static route to the ISP gateway
- A security policy allowing external traffic to reach the FTP server

All procedures in this section require you to connect to VDOM-B, either using a global or per-VDOM administrator account.

To add the firewall addresses in the GUI:

1. Go to *Policy & Objects* > *Addresses* and create a new address.
2. Enter the following information:

Address Name	FTP-server
Type	Subnet
Subnet / IP Range	192.168.20.10/32
Interface	port2
Show in Address List	enabled

To add the firewall addresses with the CLI:

```
config vdom
   edit VDOM-B
      config firewall address
         edit FTP-server
            set associated-interface port2
            set subnet 192.168.20.10 255.255.255.255
         next
      end
   next
end
```

To add the virtual IP address in the GUI:

1. Go to *Policy & Objects* > *Virtual IPs* and create a new virtual IP address.
2. Enter the following information:

Name	FTP-server-VIP
Interface	wan2
External IP Address/Range	172.25.177.42
Internal IP Address/Range	192.168.20.10

To add the virtual IP address with the CLI:

```
config firewall vip
   edit FTP-server-VIP
      set extip 172.25.177.42
      set extintf wan2
      set mappedip 192.168.20.10
   next
end
```

To add a default route in the GUI:

1. Go to *Network* > *Static Routes* and create a new route.
2. Enter the following information:

Destination	Subnet
IP address	0.0.0.0/0.0.0.0
Gateway	172.20.10.10
Interface	wan2
Distance	10

To add a default route with the CLI:

```
config vdom
   edit VDOM-B
      config router static
         edit 0
            set device wan2
            set gateway 172.20.10.10
         next
      end
   next
end
```

To add the security policy in the GUI:

1. Go to *Policy & Objects* > *Firewall Policy* and create a new policy.
2. Enter the following information:

Name	Access-server
Incoming Interface	wan2
Outgoing Interface	port2
Source	all
Destination	FTP-server-VIP
Schedule	always
Service	FTP

Netskope Exhibit 1017

Action	ACCEPT
NAT	enabled

To add the security policy with the CLI:

```
config vdom
   edit VDOM-B
      config firewall policy
         edit 0
            set name Access-server
            set srcintf wan2
            set dstintf port2
            set srcaddr all
            set dstaddr FTP-server-VIP
            set action accept
            set schedule always
            set service FTP
            set nat enable
         next
      end
   next
end
```

Configure the VDOM link

The VDOM link allows connections from VDOM-A to VDOM-B. This allows users on the internal network to access the FTP server through the FortiGate.

The configuration for the VDOM link includes the following:

- The VDOM link interface
- Firewall addresses for the FTP server on VDOM-A and for the internal network on VDOM-B
- Static routes for the FTP server on VDOM-A and for the internal network on VDOM-B
- Policies allowing traffic using the VDOM link

All procedures in this section require you to connect to the global VDOM using a global administrator account.

To add the VDOM link in the GUI:

1. Connect to root.
2. Go to *Global* > *Network* > *Interfaces* and select *Create New* > *VDOM link.*
3. Enter the following information:

Name		VDOM-link
Interface 0		
	Virtual Domain	VDOM-A
	IP/Netmask	0.0.0.0/0.0.0.0
Interface 1		

Virtual Domain	VDOM-B
IP/Netmask	0.0.0.0/0.0.0.0

To add the VDOM link with the CLI:

```
config global
   config system vdom-link
      edit vlink
   end
   config system interface
      edit VDOM-link0
         set vdom VDOM-A
         set ip 0.0.0.0 0.0.0.0
      next
      edit VDOM-link1
         set vdom VDOM-B
         set ip 0.0.0.0 0.0.0.0
      next
   end
end
```

To add the firewall address on VDOM-A in the GUI:

1. Connect to VDOM-A.
2. Go to *Policy & Objects* > *Addresses* and create a new address.
3. Enter the following information:

Address Name	FTP-server
Type	Subnet
Subnet / IP Range	192.168.20.10/32
Interface	VDOM-link0
Show in Address List	enabled
Static Route Configuration	enabled

To add the firewall addresses on VDOM-A with the CLI:

```
config vdom
   edit VDOM-B
      config firewall address
         edit FTP-server
            set associated-interface VDOM-link0
            set allow-routing enable
            set subnet 192.168.20.10 255.255.255.255
         next
      end
   next
end
```


Netskope Exhibit 1017

To add the static route on VDOM-A in the GUI:

1. Connect to VDOM-A.
2. Go to *Network* > *Static Routes* and create a new route.
3. Enter the following information:

Destination	Named Address
Named Address	FTP-server
Gateway	0.0.0.0
Interface	VDOM-link0

To add the static route on VDOM-A with the CLI:

```
config vdom
   edit VDOM-A
      config router static
         edit 0
            set device VDOM-link0
            set dstaddr FTP-server
         next
      end
   next
end
```

To add the security policy on VDOM-A in the GUI:

1. Connect to VDOM-A.
2. Go to *Policy & Objects* > *Firewall Policy* and create a new policy.
3. Enter the following information:

Name	Access-FTP-server
Incoming Interface	port1
Outgoing Interface	VDOM-link0
Source	internal-network
Destination	FTP-server
Schedule	always
Service	FTP
Action	ACCEPT
NAT	disabled

To add the security policy on VDOM-A with the CLI:

```
config vdom
   edit VDOM-A
      config firewall policy
         edit 0
```


Netskope Exhibit 1017

```
            set name Access-FTP-server
            set srcintf port1
            set dstintf VDOM-link0
            set srcaddr internal-network
            set dstaddr FTP-server
            set action accept
            set schedule always
            set service FTP
        next
      end
   next
end
```

To add the firewall address on VDOM-B in the GUI:

1. Connect to VDOM-B.
2. Go to *Policy & Objects* > *Addresses* and create a new address.
3. Enter the following information:

Address Name	internal-network
Type	Subnet
Subnet / IP Range	192.168.10.0/24
Interface	VDOM-link1
Show in Address List	enabled
Static Route Configuration	enabled

To add the firewall addresses on VDOM-B with the CLI:

```
config vdom
   edit VDOM-B
      config firewall address
         edit internal-network
            set associated-interface VDOM-link1
            set allow-routing enable
            set subnet 192.168.10.0 255.255.255.0
         next
      end
   next
end
```

To add the static route on VDOM-B in the GUI:

1. Connect to VDOM-B.
2. Go to *Network* > *Static Routes* and create a new route.
3. Enter the following information:

Destination	Named Address
Named Address	internal-network

Gateway	0.0.0.0
Interface	VDOM-link1

To add the static route on VDOM-B with the CLI:

```
config vdom
   edit VDOM-B
      config router static
         edit 0
            set device VDOM-link1
            set dstaddr internal-network
         next
      end
   next
end
```

To add the security policy on VDOM-B in the GUI:

1. Connect to VDOM-B.
2. Go to *Policy & Objects* > *Firewall Policy* and create a new policy.
3. Enter the following information:

Name	Internal-server-access
Incoming Interface	VDOM-link1
Outgoing Interface	port2
Source	internal-network
Destination	FTP-server
Schedule	always
Service	FTP
Action	ACCEPT
NAT	disabled

To add the security policy on VDOM-B with the CLI:

```
config vdom
   edit VDOM-B
      config firewall policy
         edit 0
            set name Internal-server-access
            set srcintf VDOM-link1
            set dstintf port2
            set srcaddr internal-network
            set dstaddr FTP-server
            set action accept
            set schedule always
            set service FTP
         next
      end
```


Netskope Exhibit 1017

```
    next
end
```

NAT and transparent mode

In this example, VDOM-A uses NAT mode and VDOM-B uses transparent mode.

This configuration requires the following steps:

1. Configure VDOM-A on page 627
2. Configure VDOM-B on page 629

Configure VDOM-A

VDOM-A allows connections from devices on the internal network to the Internet. WAN 1 and port 1 are assigned to this VDOM.

The per-VDOM configuration for VDOM-A includes the following:

- A firewall address for the internal network
- A static route to the ISP gateway
- A security policy allowing the internal network to access the Internet

All procedures in this section require you to connect to VDOM-A, either using a global or per-VDOM administrator account.

To add the firewall addresses in the GUI:

1. Go to *Policy & Objects* > *Addresses* and create a new address.
2. Enter the following information:

Name	internal-network
Type	Subnet
Subnet / IP Range	192.168.10.0/24
Interface	port1
Show in Address List	enabled

To add the firewall addresses with the CLI:

```
config vdom
   edit VDOM-A
      config firewall address
         edit internal-network
            set associated-interface port1
            set subnet 192.168.10.0 255.255.255.0
         next
      end
   next
end
```

To add a default route in the GUI:

1. Go to *Network* > *Static Routes* and create a new route.
2. Enter the following information:

Destination	Subnet
IP address	0.0.0.0/0.0.0.0
Gateway	172.20.201.7
Interface	wan1
Distance	10

To add a default route with the CLI:

```
config vdom
   edit VDOM-A
      config router static
         edit 0
            set gateway 172.20.201.7
            set device wan1
         next
      end
   next
end
```

To add the security policy in the GUI:

1. Connect to VDOM-A.
2. Go to *Policy & Objects* > *Firewall Policy* and create a new policy.
3. Enter the following information:

Name	VDOM-A-Internet
Incoming Interface	port1
Outgoing Interface	wan1
Source	internal-network
Destination	all
Schedule	always
Service	ALL
Action	ACCEPT
NAT	enabled

To add the security policy with the CLI:

```
config vdom
   edit VDOM-A
      config firewall policy
         edit 0
```


Netskope Exhibit 1017

```
            set name VDOM-A-Internet
            set srcintf port1
            set dstintf wan1
            set srcaddr internal-network
            set dstaddr all
            set action accept
            set schedule always
            set service ALL
            set nat enable
         next
      end
   next
end
```

Configure VDOM-B

VDOM-B allows external connections to reach an internal FTP server. WAN 2 and port 2 are assigned to this VDOM.

The per-VDOM configuration for VDOM-B includes the following:

- A firewall address for the FTP server
- A static route to the ISP gateway
- A security policy allowing external traffic to reach the FTP server

All procedures in this section require you to connect to VDOM-B, either using a global or per-VDOM administrator account.

To add the firewall addresses in the GUI:

1. Go to *Policy & Objects* > *Addresses* and create a new address.
2. Enter the following information:

Address Name	FTP-server
Type	Subnet
Subnet / IP Range	172.25.177.42/32
Interface	port2
Show in Address List	enabled

To add the firewall addresses with the CLI:

```
config vdom
   edit VDOM-B
      config firewall address
         edit FTP-server
            set associated-interface port2
            set subnet 172.25.177.42 255.255.255.255
         next
      end
   next
end
```

To add a default route in the GUI:

1. Go to *Network* > *Routing Table* and create a new route.
2. Enter the following information:

Destination	Subnet
IP address	0.0.0.0/0.0.0.0
Gateway	172.20.10.10

To add a default route with the CLI:

```
config vdom
   edit VDOM-B
      config router static
         edit 0
            set gateway 172.20.10.10
         next
      end
      next
   end
```

To add the security policy in the GUI:

1. Connect to VDOM-B.
2. Go to *Policy & Objects* > *Firewall Policy* and create a new policy.
3. Enter the following information:

Name	Access-server
Incoming Interface	wan2
Outgoing Interface	port2
Source	all
Destination	FTP-server
Schedule	always
Service	FTP
Action	ACCEPT

To add the security policy with the CLI:

```
   config vdom
      edit VDOM-B
         config firewall policy
            edit 0
               set name Access-server
               set srcintf wan2
               set dstintf port2
               set srcaddr all
               set dstaddr FTP-server-VIP
               set action accept
```

```
            set schedule always
            set service FTP
        next
    end
  next
end
```

High Availability

The following sections provide instructions on configuring High Availability (HA):

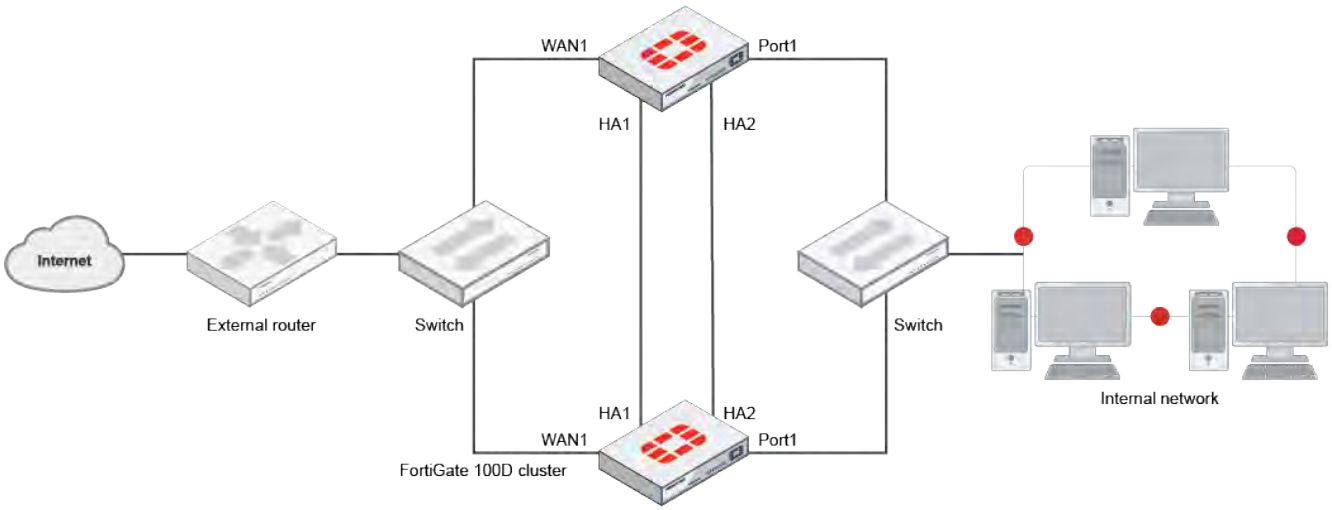
- Introduction to the FGCP cluster on page 631
- Failover protection on page 632
- FGSP (session synchronization) peer setup on page 634
- UTM inspection on asymmetric traffic in FGSP on page 635
- UTM inspection on asymmetric traffic on L3 on page 637
- Encryption for L3 on asymmetric traffic in FGSP on page 639
- Synchronizing sessions between FGCP clusters on page 639
- Using standalone configuration synchronization on page 641
- Troubleshoot an HA formation on page 643
- Check HA sync status on page 643
- Disabling stateful SCTP inspection on page 644
- Upgrading FortiGates in an HA cluster on page 645
- HA cluster setup examples on page 646
- Routing data over the HA management interface on page 655
- Override FortiAnalyzer and syslog server settings on page 657
- Force HA failover for testing and demonstrations on page 661
- Querying autoscale clusters for FortiGate VM on page 664

Introduction to the FGCP cluster

High availability (HA) is usually required in a system where there is high demand for little downtime. There are usually hot-swaps, backup routes, or standby backup units and as soon as the active entity fails, backup entities will start functioning. This results in minimal interruption for the users.

The FortiGate Clustering Protocol (FGCP) is a proprietary HA solution whereby FortiGates can find other member FortiGates to negotiate and create a cluster. A FortiGate HA cluster consists of at least two FortiGates (members) configured for HA operation. All FortiGates in the cluster must be the same model and have the same firmware installed. Cluster members must also have the same hardware configuration (such as the same number of hard disks). All cluster members share the same configurations except for their host name and priority in the HA settings. The cluster works like a device but always has a hot backup device.

Netskope Exhibit 1017



Critical cluster components

The following are critical components in an HA cluster:

- Heartbeat connections: members will use this to communicate with each other. In general, a two-member cluster is most common. We recommend double back-to-back heartbeat connections.
- Identical connections for internal and external interfaces: as demonstrated in the topology, we recommend similar connections from each member to the switches for the cluster to function properly.

General operation

The following are best practices for general cluster operation:

- Ensure that heartbeat communication is present.
- Enable the session synchronization option in daily operation (see FGSP (session synchronization) peer setup on page 634).
- Monitor traffic flowing in and out of the interfaces.

Failover

FGCP provides failover protection in the following scenarios:

- The active device loses power.
- A monitored interface loses a connection.

After failover occurs, the user will not notice any difference, except that the active device has changed. See Failover protection on page 632 for more information.

Failover protection

The FortiGate Clustering Protocol (FGCP) provides failover protection, meaning that a cluster can provide FortiGate services even when one of the devices in the cluster encounters a problem that would result in the complete loss of connectivity for a stand-alone FortiGate unit. Failover protection provides a backup mechanism that can be used to reduce the risk of unexpected downtime, especially in mission-critical environments.

FGCP supports failover protection in three ways:

1. Link failover maintains traffic flow if a link fails.
2. If a device loses power, it automatically fails over to a backup unit with minimal impact on the network.
3. Optionally, if an SSD fails, it can automatically fail over to a backup unit.

When session-pickup is enabled in the HA settings, existing TCP session are kept, and users on the network are not impacted by downtime as the traffic can be passed without reestablishing the sessions.

When and how the failover happens

1. Link fails

Before triggering a failover when a link fails, the administrator must ensure that monitor interfaces are configured. Normally, the internal interface that connects to the internal network, and an outgoing interface for traffic to the internet or outside the network, should be monitored. Any of those links going down will trigger a failover.

2. Loss of power for active unit.

When an active (primary) unit loses power, a backup (secondary) unit automatically becomes the primary, and the impact on traffic is minimal. There are no settings for this kind of fail over.

3. SSD failure

HA failover can be triggered by an SSD failure.

To enable an SSD failure triggering HA fail over:

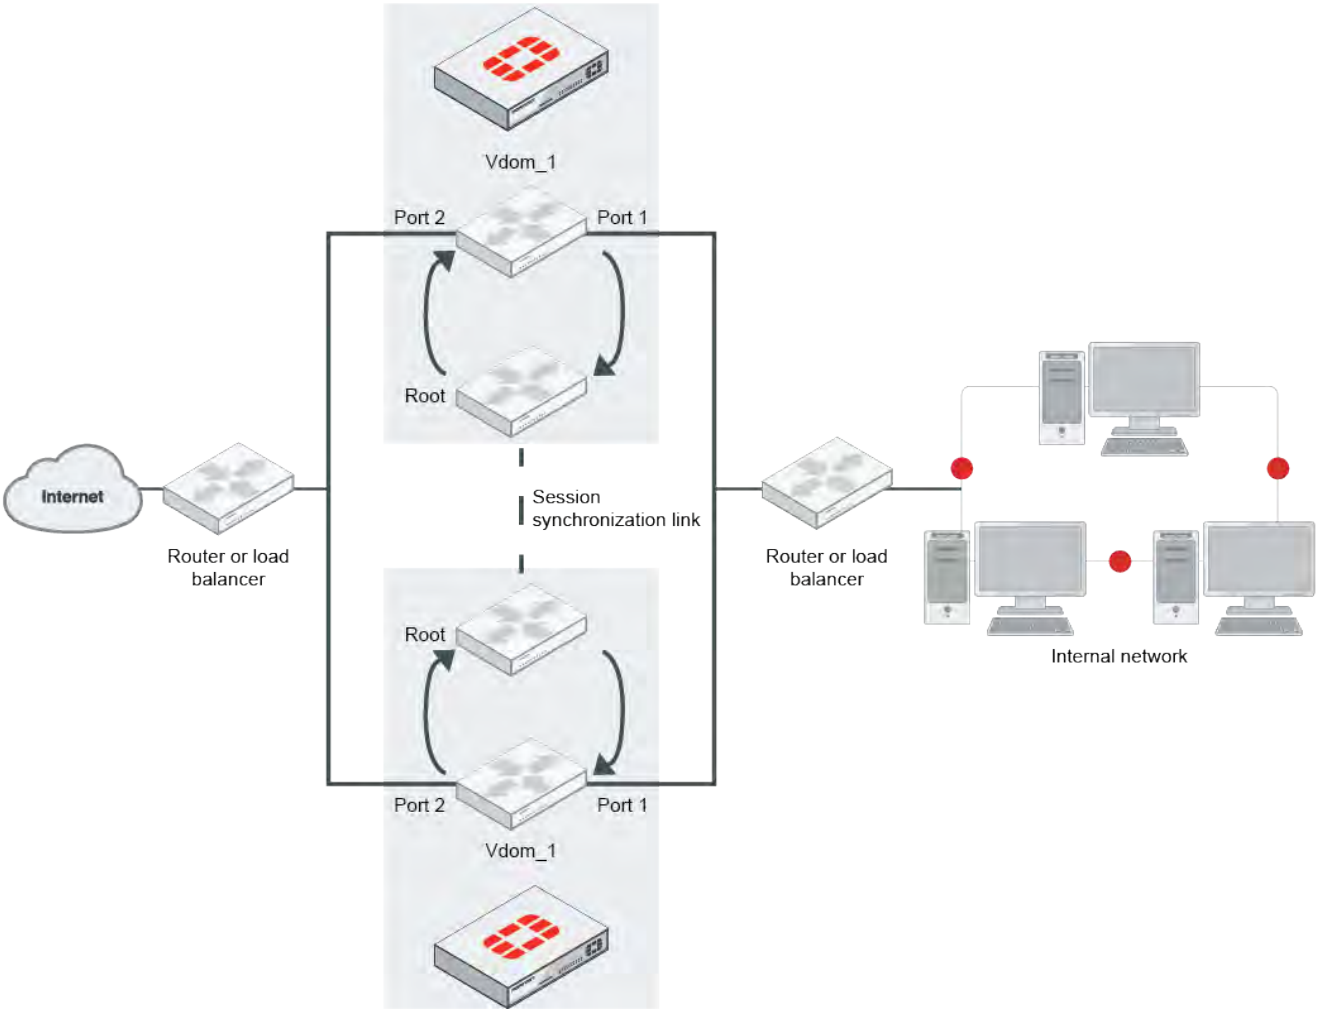
```
config system ha
    set ssd-failover enable
end
```


Netskope Exhibit 1017

FGSP (session synchronization) peer setup

The FortiGate Session Life Support Protocol (FGSP) is a proprietary HA solution for only sharing sessions between two entities and is based on a peer-to-peer structure. The entities could be standalone FortiGates or an FGCP cluster.

Connect all necessary interfaces as per the topology diagram below. Interfaces may be changed depending on the models in use. Interface names in the topology diagram are for example purposes only.



To setup an FGSP peer through the CLI:

These instructions assume that the device has been connected to the console, the CLI is accessible, and that all FortiGates have been factory reset.

1. Connect all necessary interfaces as per the topology diagram.
2. Enter the following command to change the FortiGate unit host name:

```
config system global
    set hostname Example1_host(Example2_host, etc)
end
```

3. On each FGSP peer device, enter the following command:

```
config system cluster-sync
   set peerip xx.xx.xx.xx    --->> peer's interface IP for session info to be
passed.
end
```

4. Set up identical firewall policies.
   FGSP peers share the same session information which goes from the same incoming interface (example: port1) to the outgoing interface (example: port2). Firewall policies should be identical as well, and can be copied from one device to its peer.

To test the setup:

1. Initiate TCP traffic (like HTTP access) to go through *FortiGateA*.
2. Check the session information.
   For example:
   ```
   diagnose sys session filter src xxx.xxx.xxx.xxx (your PCs IP)
   diagnose sys session list
   ```
3. Use the same command on *FortiGateB* to determine if the same session information appeared.

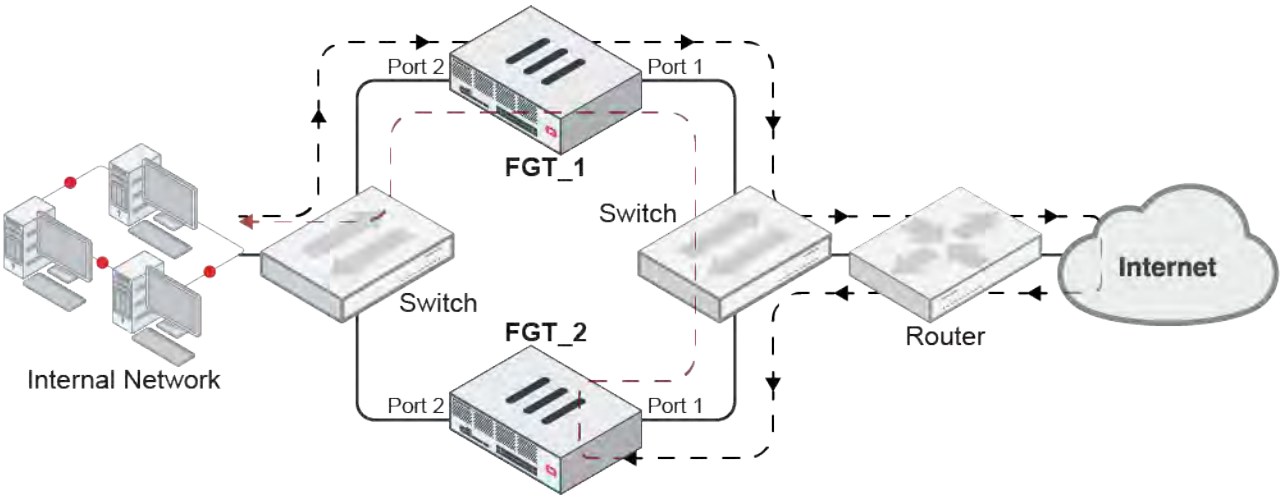
UTM inspection on asymmetric traffic in FGSP

When traffic passes asymmetrically through FGSP peers, UTM inspection can be supported by always forwarding traffic back to the session owner for processing. The session owner is the FortiGate that receives the first packet of the session.

In this example, traffic from the internal network first hits FGT_1, but the return traffic is routed to FGT_2. Consequently, traffic bounces from FGT_2 port1 to FGT_1 port1 using FGT_1's MAC address. Traffic is then inspected by FGT_1.

This example requires the following settings:

- The internal and outgoing interfaces of both FortiGates in the FGSP pair are in the same subnet.
- Both peers have layer 2 access with each other.



To configure FTG_1:

1. Configure the cluster, setting the peer IP to the IP address of FGT_2:

```
config system cluster-sync
    edit 1
        set peerip 10.2.2.2
    next
end
```

2. Configure FGSP cluster attributes:

```
config system standalone-cluster
    set standalone-group-id 1
    set group-member-id 0
    set layer2-connection available
    unset session-sync-dev
end
```

3. Configure the firewall policy:

```
config firewall policy
    edit 1
        set srcintf "port2"
        set dstintf "port1"
        set srcaddr "all"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ALL"
        set utm-status enable
        set av-profile "default"
        set logtraffic all
        set nat enable
    next
end
```

To configure FTG_2:

1. Configure the cluster, setting the peer IP to the IP address of FGT_1:

```
config system cluster-sync
    edit 1
        set peerip 10.2.2.1
    next
end
```

2. Configure FGSP cluster attributes:

```
config system standalone-cluster
    set standalone-group-id 1
    set group-member-id 1
    set layer2-connection available
    unset session-sync-dev
end
```


Netskope Exhibit 1017

3. Configure the firewall policy:

```
config firewall policy
    edit 1
        set srcintf "port2"
        set dstintf "port1"
        set srcaddr "all"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ALL"
        set utm-status enable
        set av-profile "default"
        set logtraffic all
        set nat enable
    next
end
```

Results

Capture packets on FGT_2 to see that traffic bounced from FGT_2 to FGT_1 over the traffic interface.

```
FGT_2 # diagnose sniffer packet any 'host 10.1.100.15 and host 172.6.200.55' 4
interfaces=[any]
filters=[host 10.1.100.15 and host 172.16.200.55]
91.803816 port1 in 172.16.200.55.80 -> 10.1.100.15.40008: syn 2572073713 ack 261949279
92.800480 port1 in 172.16.200.55.80 -> 10.1.100.15.40008: syn 2572073713 ack 261949279
92.800486 port1 out 172.16.200.55.80 -> 10.1.100.15.40008: syn 2572073713 ack 261949279
92.800816 port1 in 172.16.200.55.80 -> 10.1.100.15.40008: syn 2572073713 ack 261949279
92.800818 port1 out 172.16.200.55.80 -> 10.1.100.15.40008: syn 2572073713 ack 261949279
```

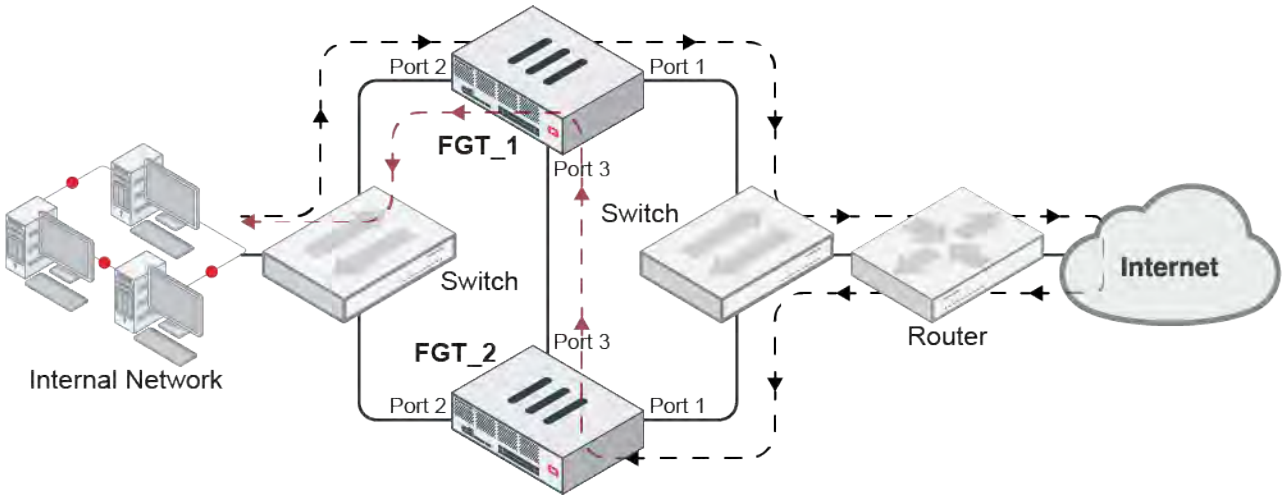
UTM inspection on asymmetric traffic on L3

When traffic passes asymmetrically through FGSP peers, UTM inspection can be supported by always forwarding traffic back to the session owner for processing. The session owner is the FortiGate that receives the first packet of the session.

For networks where L2 connectivity is not available, such as cloud environments, traffic bound for the session owner are forwarded through the peer interface using a UDP connection.

In this example, traffic from the internal network first hits FGT_1, but the return traffic is routed to FGT_2. Consequently, return traffic is packed and sent from FGT_2 to FGT_1 using UDP encapsulation between two peer interfaces (port 3). Traffic is then inspected by FGT_1.

Netskope Exhibit 1017



To configure FTG_1:

1. Configure the cluster, setting the peer IP to the IP address of FGT_2:

```
config system cluster-sync
    edit 1
        set peerip 10.2.2.2
    next
end
```

2. Configure FGSP cluster attributes:

```
config system standalone-cluster
    set standalone-group-id 1
    set group-member-id 0
    set layer2-connection unavailable
    unset session-sync-dev
end
```

3. Configure the firewall policy:

```
config firewall policy
    edit 1
        set srcintf "port2"
        set dstintf "port1"
        set srcaddr "all"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ALL"
        set utm-status enable
        set av-profile "default"
        set logtraffic all
        set nat enable
    next
end
```

To configure FTG_2:

1. Configure the cluster, setting the peer IP to the IP address of FGT_1:

```
config system cluster-sync
    edit 1
        set peerip 10.2.2.1
    next
end
```

2. Configure FGSP cluster attributes:

```
config system standalone-cluster
    set standalone-group-id 1
    set group-member-id 1
    set layer2-connection unavailable
    unset session-sync-dev
end
```

3. Configure the firewall policy:

```
config firewall policy
    edit 1
        set srcintf "port2"
        set dstintf "port1"
        set srcaddr "all"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ALL"
        set utm-status enable
        set av-profile "default"
        set logtraffic all
        set nat enable
    next
end
```

Encryption for L3 on asymmetric traffic in FGSP

In scenarios where asymmetric routing between FGSP members occurs, the return traffic can be encrypted and routed back to the session owner on Layer 3 (L3).

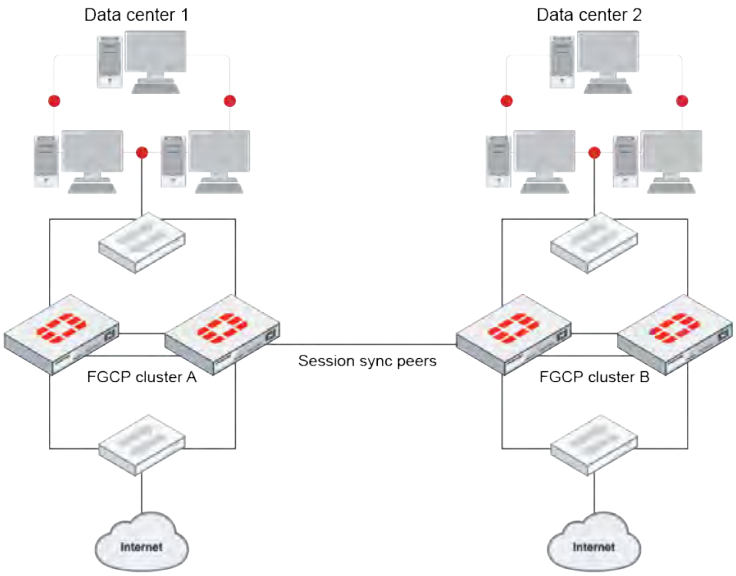
To encrypt L3 traffic in FGSP:

1. Run the following on both FortiGates:

```
config system standalone-cluster
    set encryption enable
    set psksecret xxxxxxxxx
end
```

Synchronizing sessions between FGCP clusters

Synchronizing sessions between FGCP clusters is useful when data centers in different locations are used for load-balancing, and traffic must be shared and flow freely based on demand.



There are some limitations when synchronizing sessions between FGCP clusters:

- All FortiGates must have the same model and generation, hardware configuration, and FortiOS version.
- Currently, a total of four clusters can share sessions.

To configure session synchronization between two clusters:

1. Configure the two clusters (see HA active-passive cluster setup on page 646 or HA active-active cluster setup on page 648).
2. On cluster A, configure the peer IP for the interface:

```
config system interface
    edit "port5"
        set vdom "root"
        set ip 10.10.10.1 255.255.255.0
        set allowaccess ping https ssh snmp http telnet
    next
end
```

   In this example, cluster A uses port5 and its IP address, 10.10.10.1, is reachable from another cluster.
3. On cluster A, configure cluster and session synchronization:

```
config system cluster-sync
    edit 1
        set peerip 10.10.10.2
    next
end
```

4. On cluster A, configure additional FGSP attributes as needed:

```
config system standalone-cluster
    set standalone-group-id 1
    set group-member-id 0
    set session-sync-dev <interface>
end
```

   The `standalone-group-id` must match between FGSP members. The `group-member-id` is unique for each FGCP cluster. `session-sync-dev` is an optional command to specify the interfaces to sync sessions.

5. On cluster B, configure the peer IP for the interface:

```
config system interface
    edit "port5"
        set vdom "root"
        set ip 10.10.10.2 255.255.255.0
        set allowaccess ping https ssh snmp http telnet
    next
end
```

In this example, cluster B uses port5 and its IP address, 10.10.10.2, is reachable from another cluster.

6. On cluster B, configure cluster and session synchronization:

```
config system cluster-sync
    edit 1
        set peerip 10.10.10.1
    next
end
```

7. On cluster B, configure additional FGSP attributes as needed:

```
config system standalone-cluster
    set standalone-group-id 1
    set group-member-id 1
    set session-sync-dev <interface>
end
```

Using standalone configuration synchronization

You can configure synchronization from one standalone FortiGate to another standalone FortiGate (`standalone-config-sync`). With the exception of some configurations that do not sync (settings that identify the FortiGate to the network), the rest of the configurations are synced, such as firewall policies, firewall addresses, and UTM profiles.

This option is useful in situations when you need to set up FGSP peers, or when you want to quickly deploy several FortiGates with the same configurations. You can set up `standalone-config-sync` for multiple members.



`standalone-config-sync` is an independent feature and should be used with caution as there are some limitations. We recommend disabling it once the configurations have been synced over.

Limitations

When standalone configuration synchronization is enabled, there are some limitations, including but not limited to the following:

- Network interruptions occur during firmware upgrades: when upgrading the firmware, all members in the `standalone-config-sync` group are upgraded simultaneously. This creates downtime if the FortiGates are the only outgoing gateway in the network. We recommend disabling the option before upgrading firmware.
- Some unwanted configurations might be synced: the current design and implementation of `standalone-config-sync` is based on requirements from specific customers. Thus, some users may find that unwanted parts of the configurations are synced. Should this occur, we recommend disabling the option and modifying those configurations manually.
- The wrong primary device might be picked accidentally: `standalone-config-sync` is derived from the HA primary unit selection mechanism. All members in the group will join the selection process in the same way as a the

Netskope Exhibit 1017

HA cluster selection process. It is important to select the correct device as the primary, otherwise the wrong device could be selected and existing configurations could be overwritten.

- Layer 2 heartbeat connections must be present: similar to HA heartbeat requirements, one or more layer 2 heartbeat connections are needed to sync configurations between the primary and secondary devices.

Setting up standalone configuration synchronization

Two or more standalone FortiGates should be connected to each other with one or more heartbeat interfaces, either back-to-back or via a switch. In the following example, the device supplying the configurations is called "conf-primary," and the devices receiving the configurations are called "conf-seconds."



To set up standalone configuration synchronization:

1. Configure the conf-primary device for the group:

```
config system ha
    set hbdev ha1 50 ha2 100
    set priority 255
    set override enable
    set standalone-config-sync enable
end
```

2. Configure the conf-primary device as needed to be functional.
3. Configure the other group members as conf-seconds:

```
config system ha
    set standalone-config-sync enable
end
```

4. Wait 10–15 minutes for the configurations to sync over.
5. Verify the synchronization status:

```
get sys ha status
...
Configuration Status:
  FG201E4Q17900771(updated 3 seconds ago): out-of-sync
  FG201ETK19900991(updated 1 seconds ago): in-sync
...
```

If all members are `in-sync`, this means all members share the same configurations, except those that should not be synced. If any members are `out-of-sync`, this means the member failed to sync with the primary device.



Debugging is similar when a cluster is out of sync.

Troubleshoot an HA formation

The following are requirements for setting up an HA cluster or FGSP peers.

Cluster members must have:

- The same model.
- The same hardware configuration.
- The same connections.
- The same generation.



The requirement to have the same generation is done as a best practice as it avoids issues that can occur later on. If you are unsure if the FortiGates are from the same generation, please contact customer service.

Troubleshooting common HA formation errors

One member keeps shutting down during HA setup (hard drive failure):

If one member has a hard drive failure but the other does not, the one with the hard drive failure will be shut down during HA setup. In this case, RMA the member to resolve the issue.

All members are primary devices and members can't see other members:

Typically, this is a heartbeat issue. It is recommended that for a two-member cluster, you use a back-to-back connection for heartbeat communication. If there are more than three members in the cluster, a separate switch should be used to connect all heartbeat interfaces.

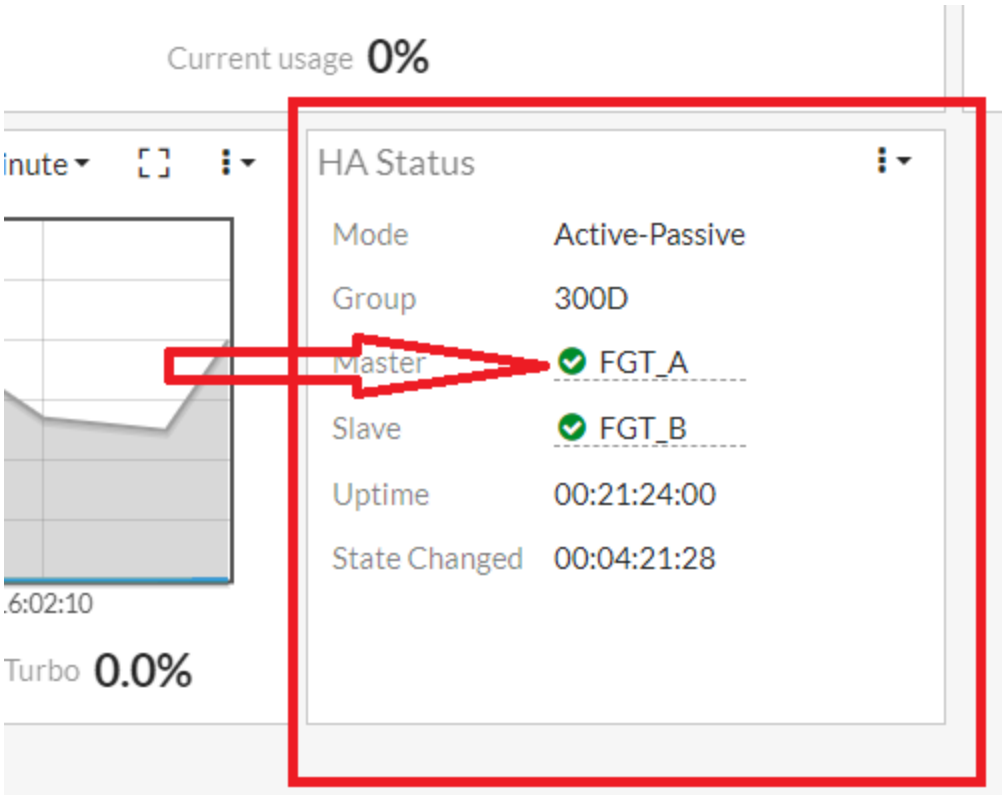
Check HA sync status

The HA sync status can be viewed in the GUI through either a widget on the *Dashboard* or on the *System* > *HA* page. It can also be confirmed through the CLI. When a cluster is out of sync, administrators should correct the issue as soon as possible as it affects the configuration integrity and can cause issues to occur.

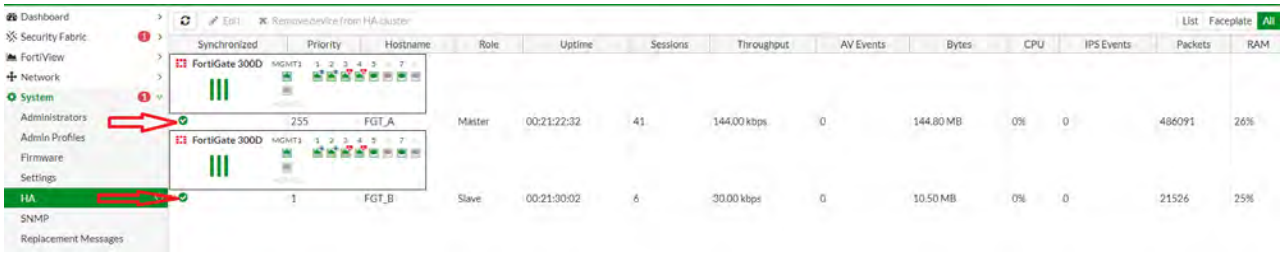
HA sync status in the GUI

- Dashboard widget:Following HA setup, the *HA Status* widget can be added to the *Dashboard*. The widget shows the HA sync status by displaying a green checkmark next to each member in sync. A red mark indicates the member is out of sync.

Netskope Exhibit 1017



- *System* > *HA* page:The same set of icons will be displayed on the *System* > *HA* page to indicate if the member is in sync.



HA sync status in the CLI

In the CLI, run the `get sys ha status` command to see if the cluster is in sync. The sync status is reported under *Configuration Status*.

Disabling stateful SCTP inspection

There is an option in FortiOS to disable stateful SCTP inspection. This option is useful when FortiGates are deployed in a high availability (HA) cluster that uses the FortiGate Clustering Protocol (FGCP) and virtual clustering in a multihoming topology. In this configuration, the primary stream control transmission protocol (SCTP) path traverses the primary FortiGate node by using its active VDOM (for example, VDOM1), and the backup SCTP path traverses the other passive FortiGate node by using its active VDOM (for example, VDOM2).

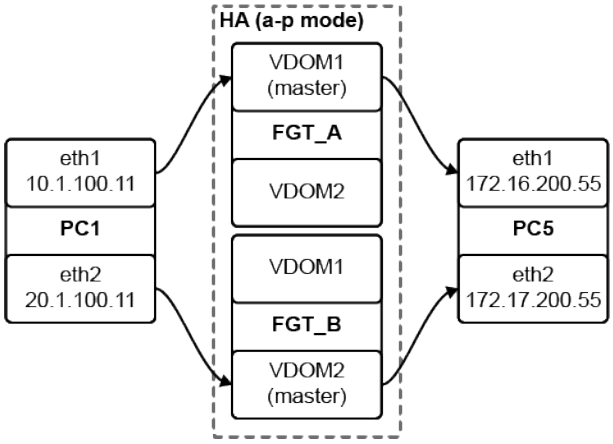
When stateful SCTP inspection is enabled, SCTP heartbeat traffic fails by means of the backup path because the primary path goes through a different platform and VDOM. Since there is no state sharing between VDOMs, the passive

FortiGate is unaware of the original SCTP session and drops the heartbeats because of no associated sessions. When stateful SCTP inspection is disabled, the passive node permits the SCTP heartbeats to pass.

When set to `enable`, SCTP session creation without SCTP INIT is enabled. When set to `disable`, SCTP session creation without SCTP INIT is disabled (this is the default setting):

```
config system settings
  set sctp-session-without-init {enable | disable}
end
```

The following is an example topology and scenario:



In this example, FGT_A and FGT_B are in HA a-p mode with two virtual clusters. Two primaries exist on different FortiGate units. PC1 eth1 can access PC5 eth1 through Vdom1, and PC1 eth2 can access PC5 eth2 through Vdom2.

On PC5, to listen for an SCTP connection:

```
sctp_darn -H 172.16.200.55 -B 172.17.200.55 -P 2500 -l
```

On PC1, to start an SCTP connection:

```
sctp_darn -H 10.1.100.11 -B 20.1.100.11 -P 2600 -c 172.16.200.55 -c 172.17.200.55 -p 2500
     -s
```

An SCTP four-way handshake is on one VDOM, and a session is created on that VDOM. With the default configuration, there is no session on any other VDOM, and the heartbeat on another path (another VDOM) is dropped. After enabling `sctp-session-without-init`, the other VDOM creates the session when it receives the heartbeat, and the heartbeat is forwarded:

```
config system settings
  set sctp-session-without-init enable
end
```

Upgrading FortiGates in an HA cluster

You can upgrade the firmware on an HA cluster in the same way as on a standalone FortiGate. During a firmware upgrade, the cluster upgrades the primary unit and all of the subordinate units to the new firmware image.



Before upgrading a cluster, back up your configuration (Configuration backups on page 108, schedule a maintenance window, and make sure that you are using a supported upgrade path (https://docs.fortinet.com/upgrade-tool).

Netskope Exhibit 1017

Uninterrupted upgrade

An uninterrupted upgrade occurs without interrupting communication in the cluster.

To upgrade the cluster firmware without interrupting communication, the following steps are followed. These steps are transparent to the user and the network, and might result in the cluster selecting a new primary unit.

1. The administrator uploads a new firmware image using the GUI or CLI. See Firmware on page 593 for details.
2. The firmware is upgraded on all of the subordinate units.
3. A new primary unit is selected from the upgraded subordinates.
4. The firmware is upgraded on the former primary unit.
5. Primary unit selection occurs, according to the standard primary unit selection process.

If all of the subordinate units crash or otherwise stop responding during the upgrade process, the primary unit will continue to operate normally, and will not be upgraded until at least one subordinate rejoins the cluster.

Interrupted upgrade

An interrupted upgrade upgrades all cluster members at the same time. This takes less time than an uninterrupted upgrade, but it interrupts communication in the cluster. Interrupted upgrade is disabled by default.

To enable interrupted upgrade:

```
config system ha
    set uninterruptible-upgrade disable
end
```

HA cluster setup examples

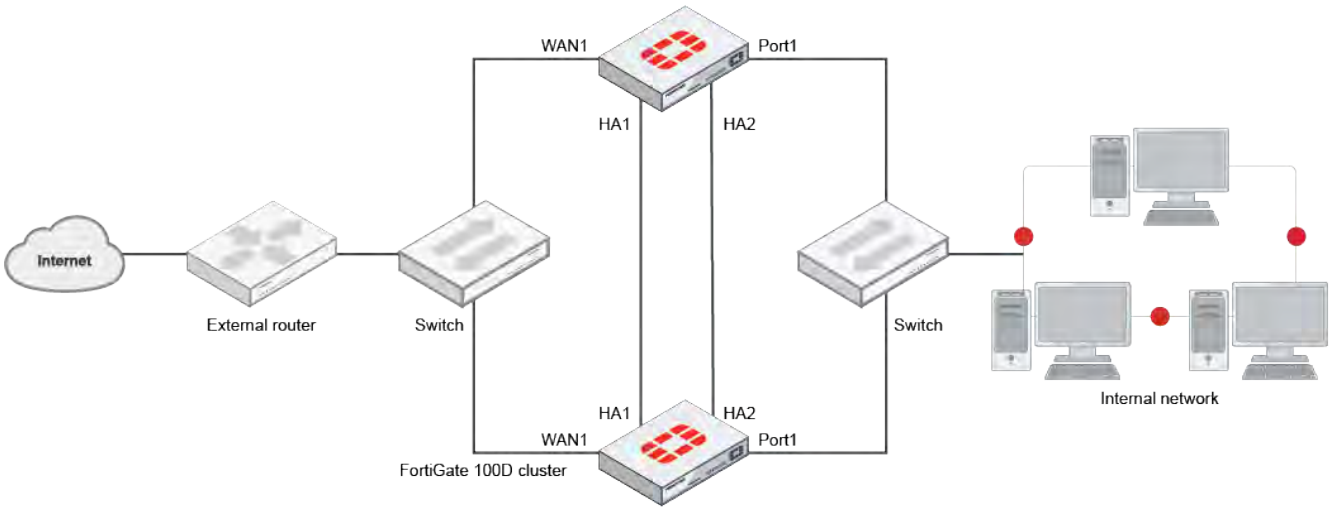
The following examples provide instructions on HA cluster setup:

- HA active-passive cluster setup on page 646
- HA active-active cluster setup on page 648
- HA virtual cluster setup on page 650
- HA using a hardware switch to replace a physical switch on page 652

HA active-passive cluster setup

An HA Active-Passive (A-P) cluster can be set up using the GUI or CLI.

This example uses the following network topology:

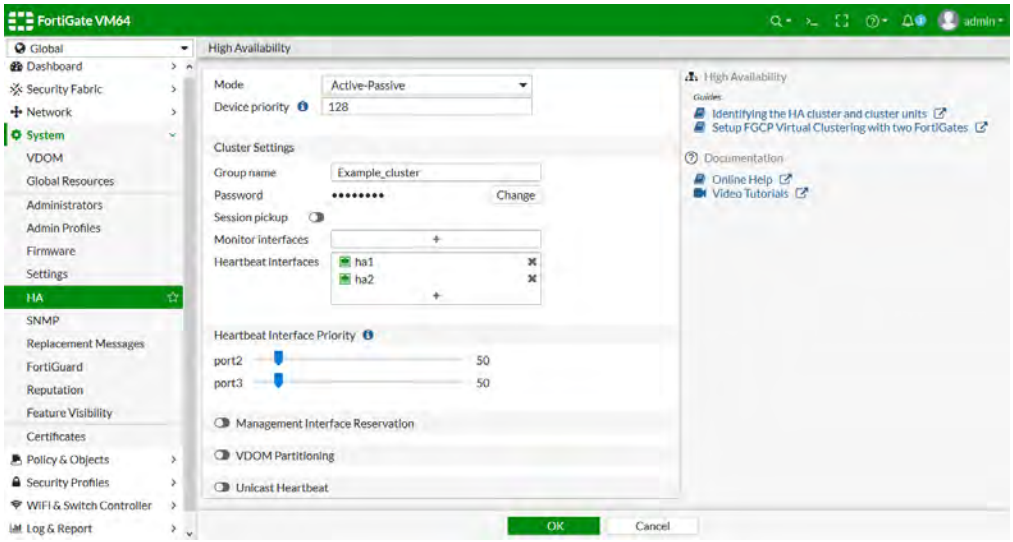


To set up an HA A-P cluster using the GUI:

1. Make all the necessary connections as shown in the topology diagram.
2. Log into one of the FortiGates.
3. Go to *System* > *HA* and set the following options:

Mode	Active-Passive
Device priority	128 or higher
Group name	Example_cluster
Heartbeat interfaces	ha1 and ha2

Except for the device priority, these settings must be the same on all FortiGates in the cluster.



4. Leave the remaining settings as their default values. They can be changed after the cluster is in operation.
5. Click *OK*.
   The FortiGate negotiates to establish an HA cluster. Connectivity with the FortiGate may be temporarily lost as the HA cluster negotiates and the FGCP changes the MAC addresses of the FortiGate's interfaces.

6. Factory reset the other FortiGate that will be in the cluster, configure GUI access, then repeat steps 1 to 5, omitting setting the device priority, to join the cluster.

To set up an HA A-P cluster using the CLI:

1. Make all the necessary connections as shown in the topology diagram.
2. Log into one of the FortiGates.
3. Change the hostname of the FortiGate:

```
config system global
    set hostname Example1_host
end
```

   Changing the host name makes it easier to identify individual cluster units in the cluster operations.
4. Enable HA:

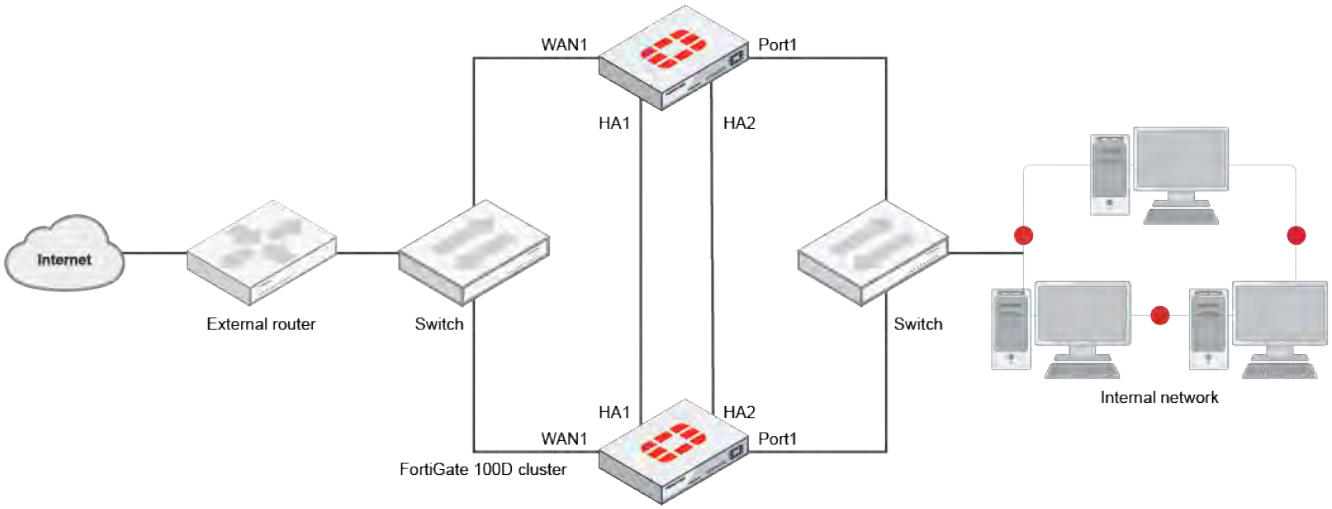
```
config system ha
    set mode a-p
    set group-name Example_cluster
    set hbdev ha1 10 ha2 20
end
```

5. Leave the remaining settings as their default values. They can be changed after the cluster is in operation.
6. Repeat steps 1 to 5 on the other FortiGate devices to join the cluster.

HA active-active cluster setup

An HA Active-Active (A-A) cluster can be set up using the GUI or CLI.

This example uses the following network topology:



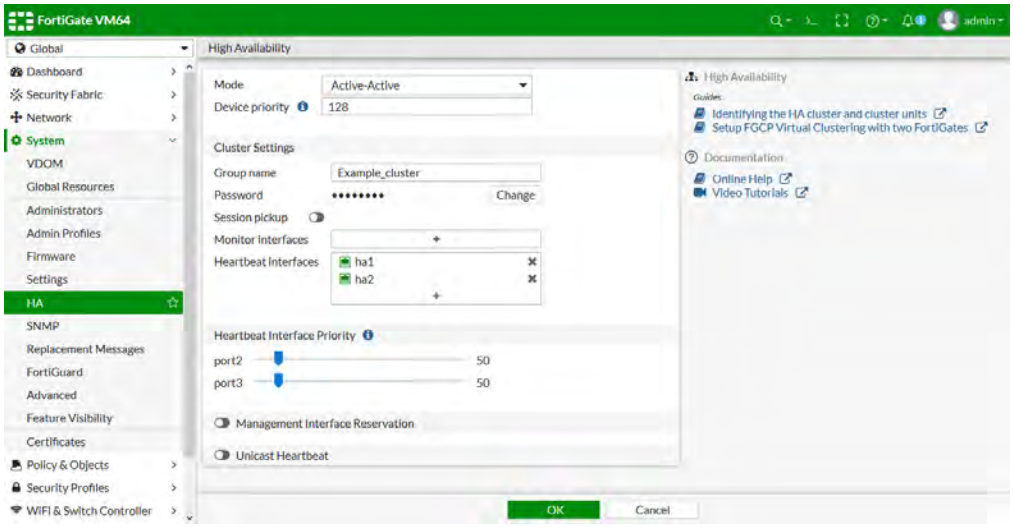
To set up an HA A-A cluster using the GUI:

1. Make all the necessary connections as shown in the topology diagram.
2. Log into one of the FortiGates.

3. Go to *System* > *HA* and set the following options:

Mode	Active-Active
Device priority	128 or higher
Group name	Example_cluster
Heartbeat interfaces	ha1 and ha2

Except for the device priority, these settings must be the same on all FortiGates in the cluster.



4. Leave the remaining settings as their default values. They can be changed after the cluster is in operation.
5. Click *OK*.
   The FortiGate negotiates to establish an HA cluster. Connectivity with the FortiGate may be temporarily lost as the HA cluster negotiates and the FGCP changes the MAC addresses of the FortiGate's interfaces.
6. Factory reset the other FortiGate that will be in the cluster, configure GUI access, then repeat steps 1 to 5, omitting setting the device priority, to join the cluster.

To set up an HA A-A cluster using the CLI:

1. Make all the necessary connections as shown in the topology diagram.
2. Log into one of the FortiGates.
3. Change the hostname of the FortiGate:

```
config system global
    set hostname Example1_host
end
```

   Changing the host name makes it easier to identify individual cluster units in the cluster operations.
4. Enable HA:

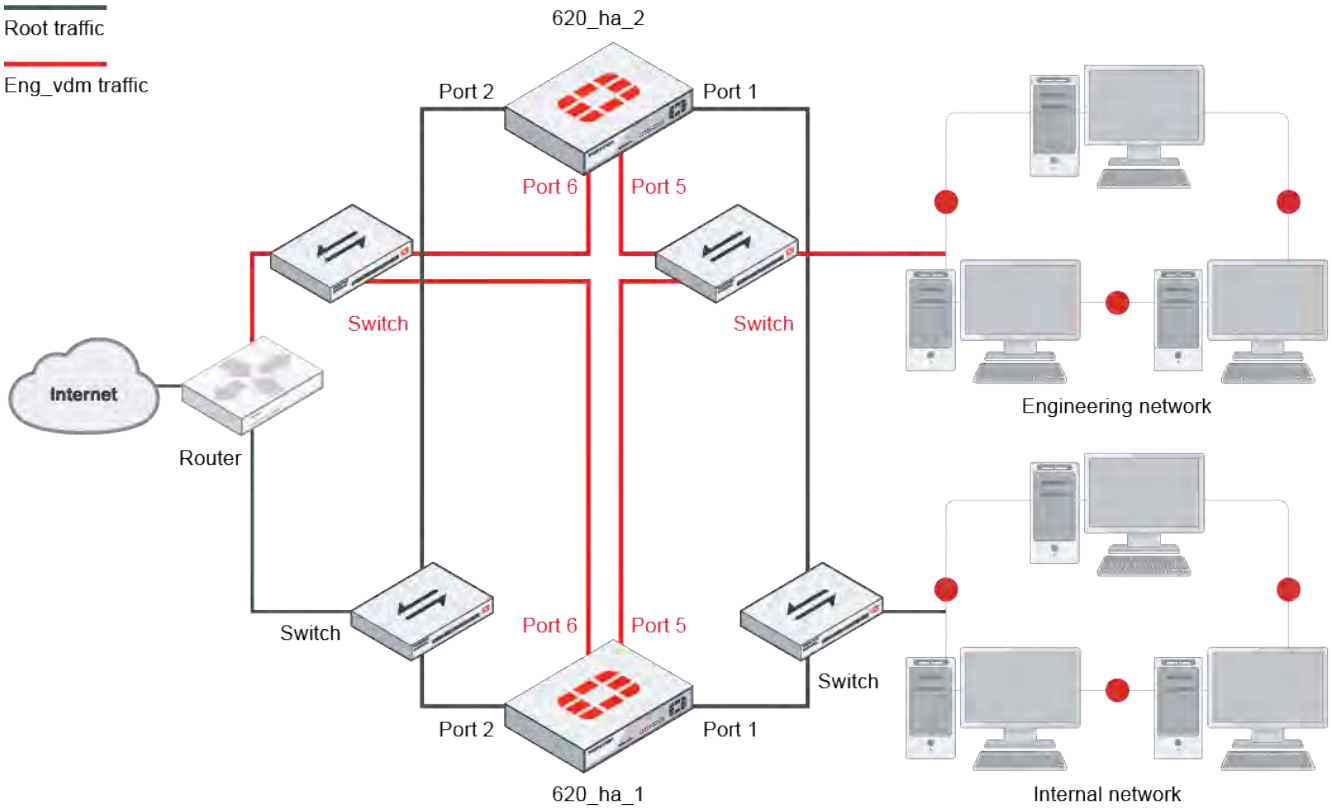
```
config system ha
    set mode a-a
    set group-name Example_cluster
    set hbdev ha1 10 ha2 20
end
```

5. Leave the remaining settings as their default values. They can be changed after the cluster is in operation.
6. Repeat steps 1 to 5 on the other FortiGate devices to join the cluster.

HA virtual cluster setup

An HA virtual cluster can be set up using the GUI or CLI.

This example uses the following network topology:



HA virtual clusters are based on VDOMs and are more complicated than regular clusters.

The root VDOM can only be associated with virtual cluster 1.

The VDOM that is assigned as the management VDOM can also only be associated with virtual cluster 1.

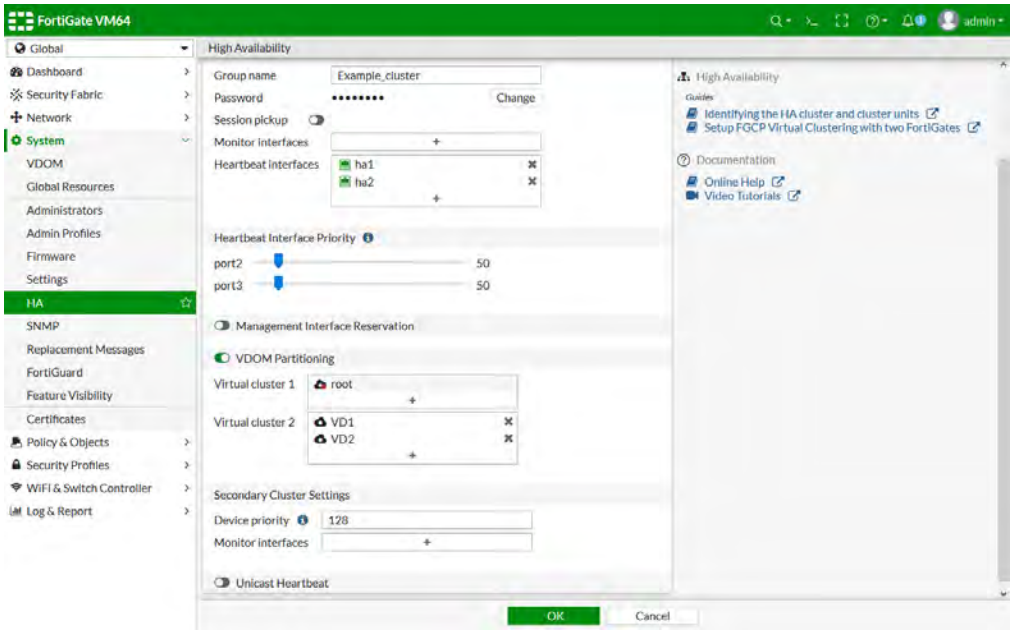
To set up an HA virtual cluster using the GUI:

1. Make all the necessary connections as shown in the topology diagram.
2. Log into one of the FortiGates.

3. Go to *System* > *HA* and set the following options:

Mode	Active-Passive
Device priority	128 or higher
Group name	Example_cluster
Heartbeat interfaces	ha1 and ha2

   Except for the device priority, these settings must be the same on all FortiGates in the cluster.
4. Leave the remaining settings as their default values. They can be changed after the cluster is in operation.
5. Click *OK*.
   The FortiGate negotiates to establish an HA cluster. Connectivity with the FortiGate may be temporarily lost as the HA cluster negotiates and the FGCP changes the MAC addresses of the FortiGate's interfaces.
6. Factory reset the other FortiGate that will be in the cluster, configure GUI access, then repeat steps 1 to 5, omitting setting the device priority, to join the cluster.
7. Go to *System* > *Settings* and enable *Virtual Domains*.
8. Click *Apply*. You will be logged out of the FortiGate.
9. Log back into the FortiGate, ensure that you are in the global VDOM, and go to *System* > *VDOM*.
10. Create two new VDOMs, such as VD1 and VD2:
    a. Click *Create New*. The *New Virtual Domain* page opens.
    b. Enter a name for the VDOM in the *Virtual Domain* field, then click *OK* to create the VDOM.
    c. Repeat these steps to create a second new VDOM.
11. Implement a virtual cluster by moving the new VDOMs to *Virtual cluster 2*:
    a. Go to *System* > *HA*.
    b. Enable *VDOM Partitioning*.
    c. Click on the *Virtual cluster 2* field and select the new VDOMs.



    d. Click *OK*.

To set up an HA virtual cluster using the CLI:

1. Make all the necessary connections as shown in the topology diagram.
2. Set up a regular A-P cluster. See HA active-passive cluster setup on page 646.
3. Enable VDOMs:

```
config system global
    set vdom-mode multi-vdom
end
```

   You will be logged out of the FortiGate.
4. Create two VDOMs:

```
config vdom
    edit VD1
    next
    edit VD2
    next
end
```

5. Reconfigure the HA settings to be a virtual cluster:

```
config global
    config system ha
        set vcluster2 enable
        config secondary-vcluster
            set vdom "VD1" "VD2"
        end
    end
end
```

HA using a hardware switch to replace a physical switch

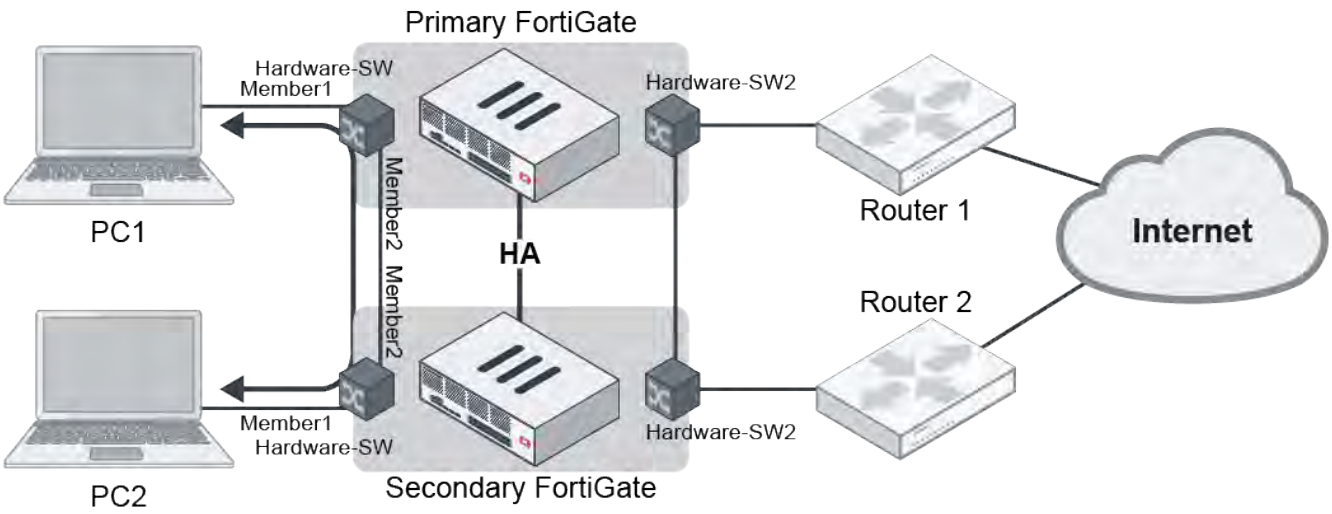
Using a hardware switch to replace a physical switch is not recommended, as it offers no redundancy or interface monitoring.

- If one FortiGate loses power, all of the clients connected to that FortiGate device cannot go to another device until that FortiGate recovers.
- A hardware switch cannot be used as a monitor interface in HA. Any incoming or outgoing link failures on hardware member interfaces will not trigger failover; this can affect traffic.

Netskope Exhibit 1017

Examples

The examples use the following topology:



Traffic between hardware switches

When using Hardware switch in HA environment, a client device connected to the hardware switch on the primary FortiGate can communicate with client devices connected to the hardware switch on secondary FortiGates as long as there is a direct connection between the two switches.

No configuration is required after setting up the hardware switches. If a client connected to both of the hardware switches needs to reach destinations outside of the cluster, the firewall must be configured for it.

To configure the FortiGate devices:

1. Connect the devices as shown in the topology diagram.
2. On each FortiGate, configure HA:

```
config system ha
    set mode a-a
    set group-name Example_cluster
    set hbdev ha1 10 ha2 20
end
```

3. On the primary FortiGate, configure the hardware switch:

```
config system virtual-switch
    edit Hardware-SW
        set physical-switch sw0
        config port
            edit port3
            next
            edit port5
            next
        end
    next
end
```

4. On each FortiGate, configure the IP addresses on the hardware switches:

```
config system interface
    edit Hardware-SW
        set ip 6.6.6.1 255.255.255.0
        set allowaccess ping ssh http https
    next
end
```

After configuring the hardware switches, PC1 and PC2 can now communicate with each other.

Traffic passes through FortiGate

If client device needs to send traffic through the FortiGate, additional firewall configuration on the FortiGate is required.

All traffic from the hardware switches on either the primary or secondary FortiGate reaches the primary FortiGate first. The traffic is then directed according to the HA mode and firewall configuration.

To configure the FortiGate devices:

1. Connect the devices as shown in the topology diagram.
2. On each FortiGate, configure HA:

```
config system ha
    set mode a-a
    set group-name Example_cluster
    set hbdev ha1 10 ha2 20
end
```

3. On the primary FortiGate, configure the hardware switch:

```
config system virtual-switch
    edit Hardware-SW
        set physical-switch sw0
        config port
            edit port3
            next
            edit port5
            next
        end
    next
    edit Hardware-SW2
        set physical-switch sw0
        config port
            edit port1
            next
        end
    next
end
```

4. On each FortiGate, configure the IP addresses on the hardware switch:

```
config system interface
    edit Hardware-SW
        set ip 6.6.6.1 255.255.255.0
        set allowaccess ping ssh http https
    next
    edit Hardware-SW2
```


Netskope Exhibit 1017

```
        set ip 172.16.200.1 255.255.255.0
        set allowaccess ping ssh http https
    next
end
```

5. On each FortiGate, configure a firewall policy:

```
config firewall policy
    edit 1
        set srcintf Hardware-SW
        set dstintf Hardware-SW2
        set srcaddr all
        set dstaddr all
        set service ALL
        set action accept
        set schedule always
        set nat enable
    next
end
```

6. On each FortiGate, configure a static route:

```
config router static
    edit 1
        set device Hardware-SW2
        set gateway 172.16.200.254
    next
end
```

Traffic from PC1 and PC2 can now reach destinations outside of the FortiGate cluster.

Routing data over the HA management interface

In an HA environment, the `ha-direct` option allows data from services such as syslog, FortiAnalyzer, FortiManager, SNMP, and NetFlow to be routed over the outgoing interface.

The following example shows how NetFlow data can be routed over the HA management interface mgmt1.

To route NetFlow data over the HA management interface:

1. On the primary unit (FortiGate A), configure the HA and mgmt1 interface settings:

```
(global) # config system ha
    set group-name "test-ha"
    set mode a-p
    set password ENC
    set hbdev "port6" 50
    set hb-interval 4
    set hb-lost-threshold 10
    set session-pickup enable
    set ha-mgmt-status enable
    config ha-mgmt-interfaces
        edit 1
            set interface "mgmt1"
        next
    end
```


Netskope Exhibit 1017

```
    set override enable
    set priority 200
    set ha-direct enable
end
```

```
(global) # config system interface
    edit "mgmt1"
        set ip 10.6.30.111 255.255.255.0
        set allowaccess ping https ssh http telnet fgfm
        set type physical
        set dedicated-to management
        set role lan
        set snmp-index 1
    next
end
```

2. On the secondary unit (FortiGate B), configure the HA and mgmt1 interface settings:

```
(global) # config system ha
    set group-name "test-ha"
    set mode a-p
    set password ENC
    set hbdev "port6" 50
    set hb-interval 4
    set hb-lost-threshold 10
    set session-pickup enable
    set ha-mgmt-status enable
    config ha-mgmt-interfaces
        edit 1
            set interface "mgmt1"
        next
    end
    set override enable
    set priority 100
    set ha-direct enable
end
```

```
(global) # config system interface
    edit "mgmt1"
        set ip 10.6.30.112 255.255.255.0
        set allowaccess ping https ssh http telnet fgfm
        set type physical
        set dedicated-to management
        set role lan
        set snmp-index 1
    next
end
```

3. On the primary unit (FortiGate A), configure the NetFlow setting:

```
(global) # config system netflow
    set collector-ip 10.6.30.59
end
```

4. Verify that NetFlow uses the mgmt1 IP:

```
(global) # diagnose test application sflowd 3
```


Netskope Exhibit 1017

5. Verify that the NetFlow packets are being sent by the mgmt1 IP:

```
(vdom1) # diagnose sniffer packet any 'udp and port 2055' 4
interfaces=[any]
filters=[udp and port 2055]
8.397265 mgmt1 out 10.6.30.111.1992 -> 10.6.30.59.2055: udp 60
23.392175 mgmt1 out 10.6.30.111.1992 -> 10.6.30.59.2055: udp 188
23.392189 mgmt1 out 10.6.30.111.1992 -> 10.6.30.59.2055: udp 60
^C
3 packets received by filter
0 packets dropped by kernel
```

6. On the secondary device (FortiGate B), change the priority so that it becomes the primary:

```
(global) # config system ha
    set priority 250
end
```

7. Verify the NetFlow status on FortiGate A, which is using the new primary's mgmt1 IP:

```
(global) # diagnose test application sflowd 3
```

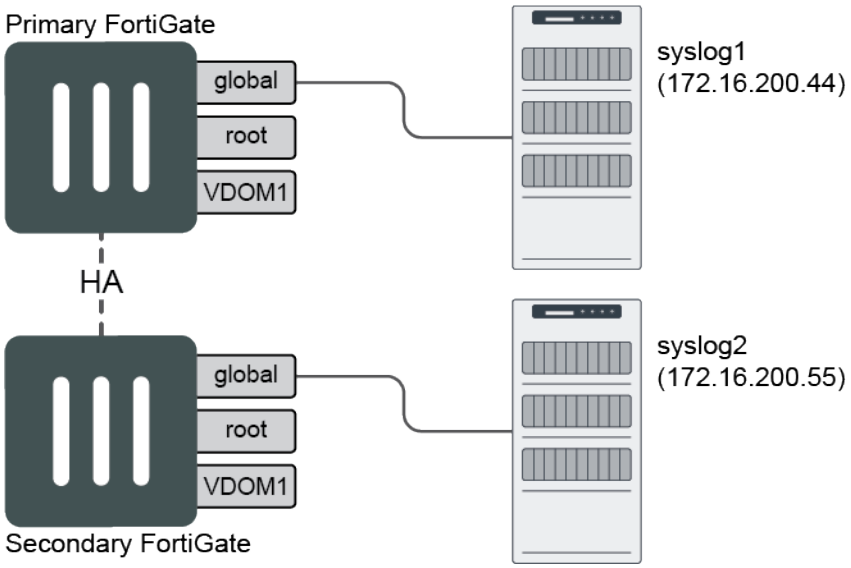
8. Verify that the NetFlow packets use the new source IP on FortiGate B:

```
(vdom1) # diagnose sniffer packet any 'udp and port 2055' 4
interfaces=[any]
filters=[udp and port 2055]
7.579574 mgmt1 out 10.6.30.112.3579 -> 10.6.30.59.2055: udp 60
22.581830 mgmt1 out 10.6.30.112.3579 -> 10.6.30.59.2055: udp 60
29.038336 mgmt1 out 10.6.30.112.3579 -> 10.6.30.59.2055: udp 1140
^C
3 packets received by filter
0 packets dropped by kernel
```

Override FortiAnalyzer and syslog server settings

In an HA cluster, secondary devices can be configured to use different FortiAnalyzer devices and syslog servers than the primary device. VDOMs can also override global syslog server settings.

Configure a different syslog server on a secondary HA device



To configure the primary HA device:

1. Configure a global syslog server:

```
config global
    config log syslog setting
        set status enable
        set server 172.16.200.44
        set facility local6
        set format default
    end
end
```

2. Set up a VDOM exception to enable setting the global syslog server on the secondary HA device:

```
config global
    config system vdom-exception
        edit 1
            set object log.syslogd.setting
        next
    end
end
```

To configure the secondary HA device:

1. Configure a global syslog server:

```
config global
    config log syslogd setting
        set status enable
        set server 172.16.200.55
        set facility local5
    end
end
```

2. After the primary and secondary device synchronize, generate logs on the secondary device.

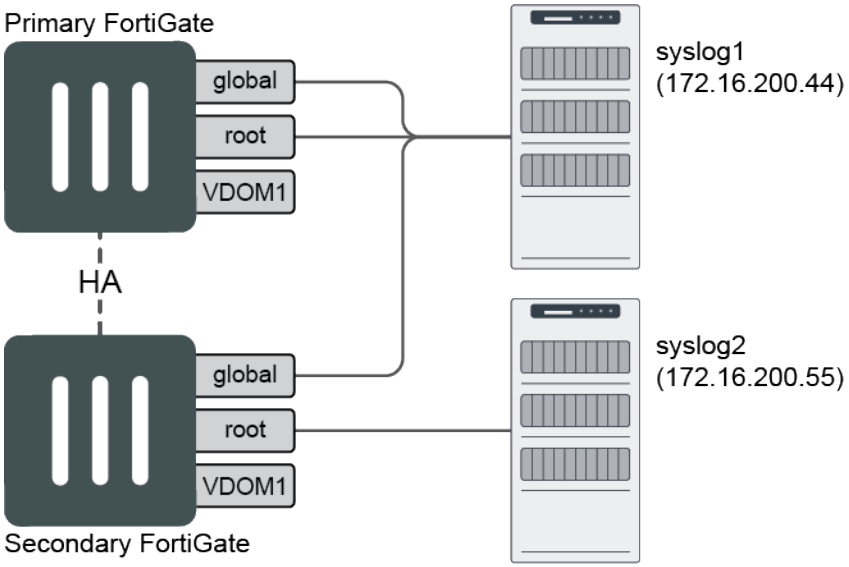
To confirm that logs are been sent to the syslog server configured on the secondary device:

1. On the primary device, retrieve the following packet capture from the secondary device's syslog server:

```
# diagnose sniffer packet any "host 172.16.200.55" 6
interfaces=[any]
filters=[host 172.16.200.55]

266.859494 port2 out 172.16.200.2.7434 -> 172.16.200.55.514: udp 278
0x0000   0000 0000 0000 0009 0f09 0004 0800 4500        ..............E.
0x0010   0132 f3c7 0000 4011 9d98 ac10 c802 ac10        .2....@.........
0x0020   c837 1d0a 0202 011e 4b05 3c31 3734 3e64        .7......K.<174>d
0x0030   6174 653d 3230 3230 2d30 332d 3134 2074        ate=2020-03-14.t
0x0040   696d 653d 3132 3a30 303a 3035 2064 6576        ime=12:00:05.dev
0x0050   6e61 6d65 3d22 466f 7274 6947 6174 652d        name="FGT-81E-Sl
0x0060   3831 455f 4122 2064 6576 6964 3d22 4647        ave-A".devid="FG
0x0070   5438 3145 3451 3136 3030 3030 3438 2220        T81E4Q16000048".
0x0080   6c6f 6769 643d 2230 3130 3030 3230 3032        logid="010002002
0x0090   3722 2074 7970 653d 2265 7665 6e74 2220        7".type="event".
0x00a0   7375 6274 7970 653d 2273 7973 7465 6d22        subtype="system"
0x00b0   206c 6576 656c 3d22 696e 666f 726d 6174        .level="informat
0x00c0   696f 6e22 2076 643d 2276 646f 6d31 2220        ion".vd="vdom1".
0x00d0   6576 656e 7474 696d 653d 3135 3834 3231        eventtime=158421
0x00e0   3234 3035 3835 3938 3335 3639 3120 747a        2405859835691.tz
0x00f0   3d22 2d30 3730 3022 206c 6f67 6465 7363        ="-0700".logdesc
0x0100   3d22 4f75 7464 6174 6564 2072 6570 6f72        ="Outdated.repor
0x0110   7420 6669 6c65 7320 6465 6c65 7465 6422        t.files.deleted"
0x0120   206d 7367 3d22 4465 6c65 7465 2031 206f        .msg="Delete.1.o
0x0130   6c64 2072 6570 6f72 7420 6669 6c65 7322        ld.report.files"
```

Configure a different syslog server in the root VDOM on a secondary HA device



To configure the primary HA device:

1. Configure a global syslog server:

```
config global
    config log syslog setting
        set status enable
        set server 172.16.200.44
        set facility local6
        set format default
    end
end
```

2. Set up a VDOM exception to enable `syslog-override` in the secondary HA device root VDOM:

```
config global
    config system vdom-exception
        edit 1
            set object log.syslogd.override-setting
            set scope inclusive
            set vdom root
        next
    end
end
```

3. In the VDOM, enable `syslog-override` in the log settings, and set up the override syslog server:

```
config root
    config log setting
        set syslog-override enable
    end
    config log syslog override-setting
        set status enable
        set server 172.16.200.44
        set facility local6
        set format default
    end
end
```

After `syslog-override` is enabled, an override syslog server must be configured, as logs will not be sent to the global syslog server.

To configure the secondary HA device:

1. Configure an override syslog server in the root VDOM:

```
config root
    config log syslogd override-setting
        set status enable
        set server 172.16.200.55
        set facility local5
        set format default
    end
end
```

2. After the primary and secondary device synchronize, generate logs in the root VDOM on the secondary device.

Netskope Exhibit 1017

To confirm that logs are been sent to the syslog server configured for the root VDOM on the secondary device:

1. On the primary device, retrieve the following packet capture from the syslog server configured in the root VDOM on the secondary device:

```
# diagnose sniffer packet any "host 172.16.200.55" 6
interfaces=[any]
filters=[host 172.16.200.55]

156.759696 port2 out 172.16.200.2.1165 -> 172.16.200.55.514: udp 277
0x0000   0000 0000 0000 0009 0f09 0004 0800 4500        ..............E.
0x0010   0131 f398 0000 4011 9dc8 ac10 c802 ac10        .1....@.........
0x0020   c837 048d 0202 011d af5f 3c31 3734 3e64        .7......._<174>d
0x0030   6174 653d 3230 3230 2d30 332d 3134 2074        ate=2020-03-14.t
0x0040   696d 653d 3131 3a33 353a 3035 2064 6576        ime=11:35:05.dev
0x0050   6e61 6d65 3d22 466f 7274 6947 6174 652d        name="FGT-81E-Sl
0x0060   3831 455f 4122 2064 6576 6964 3d22 4647        ave-A".devid="FG
0x0070   5438 3145 3451 3136 3030 3030 3438 2220        T81E4Q16000048".
0x0080   6c6f 6769 643d 2230 3130 3030 3230 3032        logid="010002002
0x0090   3722 2074 7970 653d 2265 7665 6e74 2220        7".type="event".
0x00a0   7375 6274 7970 653d 2273 7973 7465 6d22        subtype="system"
0x00b0   206c 6576 656c 3d22 696e 666f 726d 6174        .level="informat
0x00c0   696f 6e22 2076 643d 2272 6f6f 7422 2065        ion".vd="root".e
0x00d0   7665 6e74 7469 6d65 3d31 3538 3432 3130        venttime=1584210
0x00e0   3930 3537 3539 3334 3132 3632 2074 7a3d        905759341262.tz=
0x00f0   222d 3037 3030 2220 6c6f 6764 6573 633d        "-0700".logdesc=
0x0100   224f 7574 6461 7465 6420 7265 706f 7274        "Outdated.report
0x0110   2066 696c 6573 2064 656c 6574 6564 2220        .files.deleted".
0x0120   6d73 673d 2244 656c 6574 6520 3220 6f6c        msg="Delete.2.ol
0x0130   6420 7265 706f 7274 2066 696c 6573 22          d.report.files"
```

Force HA failover for testing and demonstrations



This command should only be used for testing, troubleshooting, maintenance, and demonstrations.

Do not use it in a live production environment outside of an active maintenance window.

HA failover can be forced on an HA primary device. The device will stay in a failover state regardless of the conditions. The only way to remove the failover status is by manually turning it off.

Syntax

```
execute ha failover set <cluster_id>
execute ha failover unset <cluster_id>
```

Variable	Description
<cluster_id>	The cluster ID is 1 for any cluster that is not in virtual cluster mode, and can be 1 or 2 if virtual cluster mode is enabled.

Example

To manually force an HA failover:

```
# execute ha failover set 1
Caution: This command will trigger an HA failover.
It is intended for testing purposes.
Do you want to continue? (y/n)y
```

To view the failover status:

```
# execute ha failover status
failover status: set
```

To view the system status of a device in forced HA failover:

```
# get system ha status
HA Health Status: OK
Model: FortiGate-300D
Mode: HA A-P
Group: 240
Debug: 0
Cluster Uptime: 0 days 2:11:46
Cluster state change time: 2020-03-12 17:38:04
Master selected using:
    <2020/03/12 17:38:04> FGT3HD3914800153 is selected as the master because it has EXE_
FAIL_OVER flag set.
    <2020/03/12 15:27:26> FGT3HD3914800069 is selected as the master because it has the
largest value of override priority.
ses_pickup: disable
override: enable
Configuration Status:
    FGT3HD3914800069(updated 4 seconds ago): in-sync
    FGT3HD3914800153(updated 3 seconds ago): in-sync
System Usage stats:
    FGT3HD3914800069(updated 4 seconds ago):
        sessions=5, average-cpu-user/nice/system/idle=0%/0%/0%/100%, memory=30%
    FGT3HD3914800153(updated 3 seconds ago):
        sessions=41, average-cpu-user/nice/system/idle=0%/0%/0%/99%, memory=30%
HBDEV stats:
    FGT3HD3914800069(updated 4 seconds ago):
        port3: physical/1000auto, up, rx-bytes/packets/dropped/errors=15914162/42929/0/0,
tx=15681840/39505/0/0
        port5: physical/1000auto, up, rx-bytes/packets/dropped/errors=17670346/52854/0/0,
tx=20198409/54692/0/0
    FGT3HD3914800153(updated 3 seconds ago):
        port3: physical/1000auto, up, rx-bytes/packets/dropped/errors=16636700/45544/0/0,
tx=15529791/39512/0/0
        port5: physical/1000auto, up, rx-bytes/packets/dropped/errors=20199928/54699/0/0,
tx=17672146/52862/0/0
Slave : FortiGate-300D  , FGT3HD3914800069, HA cluster index = 1
Master: FortiGate-300D  , FGT3HD3914800153, HA cluster index = 0
number of vcluster: 1
vcluster 1: standby 169.254.0.1
```


Netskope Exhibit 1017

```
Slave : FGT3HD3914800069, HA operating index = 1
Master: FGT3HD3914800153, HA operating index = 0
```

To stop the failover status:

```
# execute ha failover unset 1
Caution: This command may trigger an HA failover.
It is intended for testing purposes.
Do you want to continue? (y/n)y
```

To view the system status of a device after forced HA failover is disabled:

```
# get system ha status
HA Health Status: OK
Model: FortiGate-300D
Mode: HA A-P
Group: 240
Debug: 0
Cluster Uptime: 0 days 2:14:55
Cluster state change time: 2020-03-12 17:42:17
Master selected using:
    <2020/03/12 17:42:17> FGT3HD3914800069 is selected as the master because it has the
largest value of override priority.
    <2020/03/12 17:38:04> FGT3HD3914800153 is selected as the master because it has EXE_
FAIL_OVER flag set.
    <2020/03/12 15:27:26> FGT3HD3914800069 is selected as the master because it has the
largest value of override priority.
ses_pickup: disable
override: enable
Configuration Status:
    FGT3HD3914800069(updated 3 seconds ago): in-sync
    FGT3HD3914800153(updated 2 seconds ago): in-sync
System Usage stats:
    FGT3HD3914800069(updated 3 seconds ago):
        sessions=0, average-cpu-user/nice/system/idle=0%/0%/0%/100%, memory=30%
    FGT3HD3914800153(updated 2 seconds ago):
        sessions=38, average-cpu-user/nice/system/idle=0%/0%/0%/100%, memory=30%
HBDEV stats:
    FGT3HD3914800069(updated 3 seconds ago):
        port3: physical/1000auto, up, rx-bytes/packets/dropped/errors=16302442/43964/0/0,
tx=16053848/40454/0/0
        port5: physical/1000auto, up, rx-bytes/packets/dropped/errors=18161941/54088/0/0,
tx=20615650/55877/0/0
    FGT3HD3914800153(updated 2 seconds ago):
        port3: physical/1000auto, up, rx-bytes/packets/dropped/errors=17033009/46641/0/0,
tx=15907891/40462/0/0
        port5: physical/1000auto, up, rx-bytes/packets/dropped/errors=20617180/55881/0/0,
tx=18163135/54091/0/0
Master: FortiGate-300D  , FGT3HD3914800069, HA cluster index = 1
Slave : FortiGate-300D  , FGT3HD3914800153, HA cluster index = 0
number of vcluster: 1
vcluster 1: work 169.254.0.2
Master: FGT3HD3914800069, HA operating index = 0
Slave : FGT3HD3914800153, HA operating index = 1
```


Netskope Exhibit 1017

Querying autoscale clusters for FortiGate VM

When a FortiGate VM secondary device is added to a cluster, the new secondary member can query the cluster about its autoscale environment. FortiManager can then run this query on the new secondary member to update its autoscale record.

To view cluster information from a secondary member:

```
# diagnose sys ha checksum autoscale-cluster
```

Cluster information sample

Sample cloud topology:

```
FGT_BYOL; primary; 10.0.0.6; FGVM04TM00000066
FGT_BYOL; secondary; 10.0.0.7; FGVM00000000056
FGT_PAYG; secondary; 10.0.0.4; FGTAZ000000000CD
FGT_PAYG; secondary; 10.0.0.5; FGTAZ00000000003D
```

From the secondary device, you can see cluster checksums and the primary device.

```
# diagnose sys ha checksum autoscale-cluster
================== FGTAZ000000000CD ==================
is_autoscale_master()=0
debugzone
global: 56 49 b3 02 f2 b7 5b 82 ec 2d c2 1a ff 80 8c 79
root: bf 18 cf 83 1e 04 c3 04 4c e4 66 bc 38 fe 3a dc
all: 77 06 d0 89 6e 06 c0 86 17 98 53 72 33 85 ae ff
checksum
global: 56 49 b3 02 f2 b7 5b 82 ec 2d c2 1a ff 80 8c 79
root: bf 18 cf 83 1e 04 c3 04 4c e4 66 bc 38 fe 3a dc
all: 77 06 d0 89 6e 06 c0 86 17 98 53 72 33 85 ae ff
================== FGVM04TM00000066 ==================
is_autoscale_master()=1
debugzone
global: 56 49 b3 02 f2 b7 5b 82 ec 2d c2 1a ff 80 8c 79
root: bf 18 cf 83 1e 04 c3 04 4c e4 66 bc 38 fe 3a dc
all: 77 06 d0 89 6e 06 c0 86 17 98 53 72 33 85 ae ff
checksum
global: 56 49 b3 02 f2 b7 5b 82 ec 2d c2 1a ff 80 8c 79
root: bf 18 cf 83 1e 04 c3 04 4c e4 66 bc 38 fe 3a dc
all: 77 06 d0 89 6e 06 c0 86 17 98 53 72 33 85 ae ff
================== FGVM00000000056 ==================
is_autoscale_master()=0
debugzone
global: 56 49 b3 02 f2 b7 5b 82 ec 2d c2 1a ff 80 8c 79
root: bf 18 cf 83 1e 04 c3 04 4c e4 66 bc 38 fe 3a dc
all: 77 06 d0 89 6e 06 c0 86 17 98 53 72 33 85 ae ff
checksum
global: 56 49 b3 02 f2 b7 5b 82 ec 2d c2 1a ff 80 8c 79
root: bf 18 cf 83 1e 04 c3 04 4c e4 66 bc 38 fe 3a dc
all: 77 06 d0 89 6e 06 c0 86 17 98 53 72 33 85 ae ff
================== FGTAZ00000000003D ==================
is_autoscale_master()=0
debugzone
global: 56 49 b3 02 f2 b7 5b 82 ec 2d c2 1a ff 80 8c 79
```


Netskope Exhibit 1017

```
root: bf 18 cf 83 1e 04 c3 04 4c e4 66 bc 38 fe 3a dc
all: 77 06 d0 89 6e 06 c0 86 17 98 53 72 33 85 ae ff
checksum
global: 56 49 b3 02 f2 b7 5b 82 ec 2d c2 1a ff 80 8c 79
root: bf 18 cf 83 1e 04 c3 04 4c e4 66 bc 38 fe 3a dc
all: 77 06 d0 89 6e 06 c0 86 17 98 53 72 33 85 ae ff
```

To get ha sync information from the secondary device:

```
# get test hasync 50
autoscale_count=69. current_jiffies=41235125
    10.0.0.6, timeo=31430, serial_no=FGVM04TM19001766
    10.0.0.7, timeo=31430, serial_no=FGVM04TM19008156
    10.0.0.5, timeo=31430, serial_no=FGTAZR7UZRKKNR3D
```

SNMP

SNMP enables you to monitor hardware on your network. You can configure the hardware, such as the FortiGate SNMP agent, to report system information and send traps (alarms or event messages) to SNMP managers. SNMP traps alert you to events that happen, such as when a log disk is full or a virus is detected.

The FortiGate SNMP implementation is read-only. SNMP v1/v2c, and v3 compliant SNMP managers have read-only access to FortiGate system information through queries, and can receive trap messages from the FortiGate unit.

- Interface access on page 665
- MIB files on page 666
- SNMP agent on page 666
- SNMP v1/v2c communities on page 667
- SNMP v3 users on page 669
- Important SNMP traps on page 670
- SNMP traps and query for monitoring DHCP pool on page 672

Interface access

Before a remote SNMP manager can connect to the FortiGate SNMP agent, you must configure one or more FortiGate interfaces to accept SNMP connections.

To configure a FortiGate interface to accept SNMP connections in the GUI:

1. Go to *Network* > *Interfaces*.
2. Edit the interface.
3. In the *Administrative Access* options, enable *SNMP*.
4. Click *OK*.

To configure a FortiGate interface to accept SNMP connections in the CLI:

```
config system interface
    edit <interface>
```

```
        append allowaccess snmp
        set snmp-index <integer>
        config ipv6
            append ip6-allowaccess snmp
        end
    next
end
```

MIB files

The FortiGate SNMP agent supports Fortinet proprietary MIBs, as well as the parts of RFC 2665 and RFC 1213 that apply to FortiGate unit configuration.

Your SNMP manager may already include standard and private MIBs in a compiled database that is ready to use. You must add the Fortinet proprietary MIBs to this database to have access to Fortinet specific information.

MIB file or RFC	Description
FORTINET-CORE-MIB.mib	The Fortinet core MIB includes all system configuration and trap information that is common to all Fortinet products. Your SNMP manager requires this information to monitor Fortinet device settings and receive traps from the FortiGate SNMP agent.
FORTINET-FORTIGATE-MIB.mib	The FortiGate MIB includes all system configuration information and trap information that is specific to FortiGate units. Your SNMP manager requires this information to monitor FortiGate settings and receive traps from the FortiGate SNMP agent.
RFC-1213 (MIB II)	The FortiGate SNMP agent supports MIB II groups with the following exceptions: • No support for the EGP group from MIB II (RFC 1213, section 3.11 and 6.10). • Protocol statistics returned for MIB II groups (IP/ICMP/TCP/UDP/etc.) do not accurately capture all Fortinet traffic activity. More accurate information can be obtained from the information reported by the Fortinet MIB.
RFC-2665 (Ethernet-like MIB)	The FortiGate SNMP agent supports Ethernet-like MIB information. FortiGate SNMP does not support for the dot3Tests and dot3Errors groups.

To download the MIB files:

1. Go to *System* > *SNMP*.
2. Click *Download FortiGate MIB File* and save the file to the management computer.
3. Click *Download Fortinet Core MIB File* and save the file to the management computer.

SNMP agent

The SNMP agent sends SNMP traps originating on the FortiGate to an external monitoring SNMP manager defined in a SNMP community. The SNMP manager can monitor the FortiGate system to determine if it is operating properly, or if any critical events occurring.

The description, location, and contact information for this FortiGate system will be part of the information that the SNMP manager receives. This information is useful if the SNMP manager is monitoring many devices, and enables faster responses when the FortiGate system requires attention.

To configure the SNMP agent in the GUI:

1. Go to *System* > *SNMP*.
2. Enable *SNMP Agent*.
3. Enter a description of the agent.
4. Enter the location of the FortiGate unit.
5. Enter a contact or administrator for the SNMP Agent or FortiGate unit.
6. Click *Apply*.

To configure the SNMP agent in the CLI:

```
config system snmp sysinfo
    set status enable
    set description <string>
    set contact-info <string>
    set location <string>
end
```

SNMP v1/v2c communities

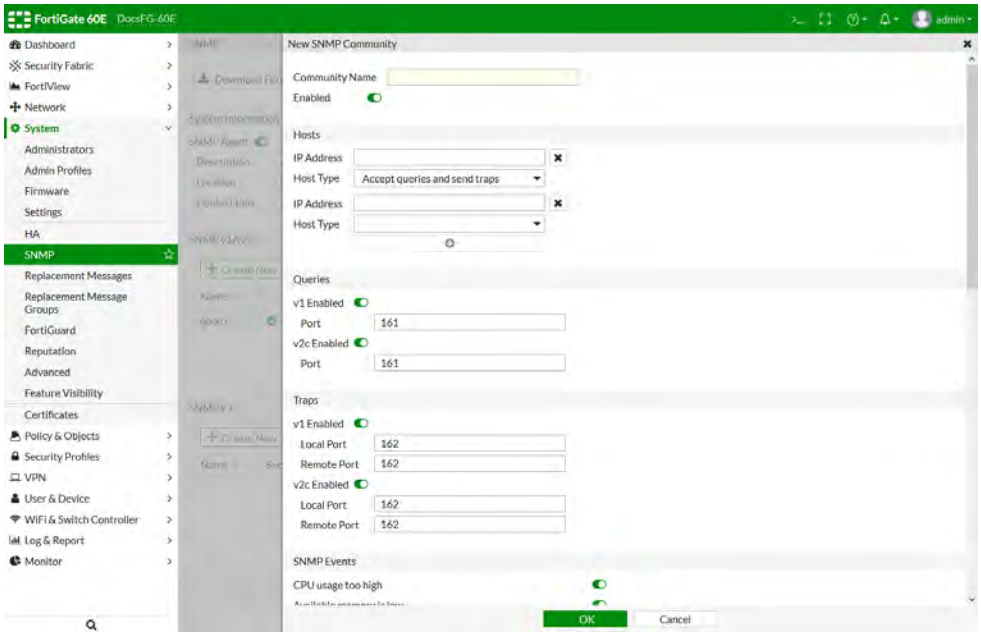
An SNMP community is a grouping of equipment for network administration purposes. A single device can belong to multiple communities.

You must add an SNMP community to the FortiGate so that the SNMP manager can receive traps and system information. Up to three communities can be added.

Netskope Exhibit 1017

To create a n SNMP v1/v2c community in the GUI:

1. Go to *System* > *SNMP*.
2. In the *SNMP v1/v2c* table, click *Create New*.



3. Enter a *Community Name* and enable the community.
4. In the *Hosts* section, enter the *IP Address* and select the *Host Type* for each SNMP manager.
5. In the *Queries* section, enable or disable v1 and v2c queries, then enter the port numbers that the SNMP managers in this community use for them.
6. In the *Traps* section, enable or disable v1 and v2c traps, then enter the local and remote port numbers that the SNMP managers in this community use for them.
7. In the *SNMP Events* section, enable or disable the events that activate traps in this community.
8. Click *OK*.

To create a n SNMP v1/v2c community in the CLI:

```
config system snmp community
    edit 2
        set name <string>
        set status {enable | disable}
        config hosts
            edit <host_id>
                set ip <ip/mask>
                set source-ip <class_ip>
                set ha-direct {enable | disable}
                set host-type {any | query | trap}
            next
        end
        set query-v1-port <port_number>
        set query-v1-status {enable | disable}
        set query-v2c-port <port_number>
        set query-v2c-status {enable | disable}
        set trap-v1-lport <port_number>
        set trap-v1-rport <port_number>
```

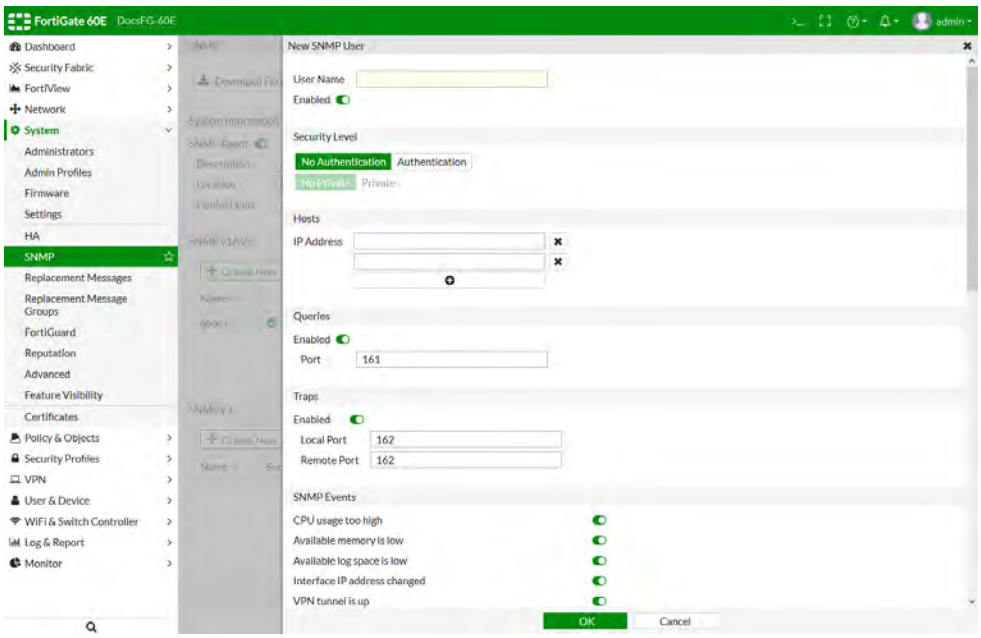
```
        set trap-v1-status {enable | disable}
        set trap-v2c-lport <port_number>
        set trap-v2c-rport <port_number>
        set trap-v2c-status {enable | disable}
        set events <events>
    next
end
```

SNMP v3 users

Authentication is used to ensure the identity of users. Privacy allows for encryption of SNMP v3 messages to ensure confidentiality of data. These protocols provide a higher level of security than is available in SNMP v1 and v2c, which use community strings for security. Both authentication and privacy are optional.

To create a n SNMP v3 user in the GUI:

1. Go to *System* > *SNMP*.
2. In the *SNMP v3* table, click *Create New*.



3. Enter a *User Name* and enable the user.
4. In the *Security Level* section, configure the security level:
   - *No Authentication*: No authentication or encryption.
   - *Authentication*: Select the authentication algorithm and password.
   - *Authentication and Private*: Select both the authentication and encryption algorithms and password.
5. In the *Hosts* section, enter the *IP Address* for each SNMP manager.
6. In the *Queries* section, enable or disable queries, then enter the port number that the SNMP managers use for them.
7. In the *Traps* section, enable or disable traps, then enter the local and remote port numbers that the SNMP managers use for them.
8. In the *SNMP Events* section, enable or disable the events that activate traps.
9. Click *OK*.

To create an SNMP v3 user in the CLI:

```
config system snmp user
    edit <user>
        set status {enable | disable}
        set trap-status {enable | disable}
        set trap-lport <port_number>
        set trap-rport <port_number>
        set queries {enable | disable}
        set query-port <port_number>
        set notify-hosts <class_ip> ... <class_ip>
        set source-ip <class_ip>
        set ha-direct {enable | disable}
        set events <events>
        set security-level {no-auth-no-priv | auth-no-priv | auth-priv}
        set auth-proto {md5 | sha | sha224 | sha256 | sha384 | sha512}
        set auth-pwd <password>
        set priv-proto {aes | des | aes256 | aes256cisco}
        set priv-pwd <password>
    next
end
```

Important SNMP traps

Link Down and Link Up traps

This trap is sent when a FortiGate port either goes down or is brought up.

For example, the following traps are generated when the state of port34 is set to down using `set status down`, and then brought up using `set status up`:

```
NET-SNMP version 5.7.3 2019-01-31 14:11:48 10.1.100.1(via UDP: [10.1.100.1]:162->
[10.1.100.11]:162) TRAP, SNMP v1, community REGR-SYS SNMPv2-MIB::snmpTraps Link Down Trap
(0) Uptime: 0:14:44.95 IF-MIB::ifIndex.42 = INTEGER: 42 IF-MIB::ifAdminStatus.42 = INTEGER:
down(2) IF-MIB::ifOperStatus.42 = INTEGER: down(2) FORTINET-CORE-MIB::fnSysSerial.0 =
STRING: FG140P3G15800330 SNMPv2-MIB::sysName.0 = STRING: FortiGate-140D-POE

2019-01-31 14:11:48 <UNKNOWN> [UDP: [10.1.100.1]:162->[10.1.100.11]:162]: DISMAN-EVENT-
MIB::sysUpTimeInstance = Timeticks: (88495) 0:14:44.95 SNMPv2-MIB::snmpTrapOID.0 = OID: IF-
MIB::linkDown IF-MIB::ifIndex.42 = INTEGER: 42 IF-MIB::ifAdminStatus.42 = INTEGER: down(2)
IF-MIB::ifOperStatus.42 = INTEGER: down(2) FORTINET-CORE-MIB::fnSysSerial.0 = STRING:
FG140P3G15800330 SNMPv2-MIB::sysName.0 = STRING: FortiGate-140D-POE 2019-01-31 14:12:01
10.1.100.1(via UDP: [10.1.100.1]:162->[10.1.100.11]:162) TRAP, SNMP v1, community REGR-SYS
SNMPv2-MIB::snmpTraps Link Up Trap (0) Uptime: 0:14:57.98 IF-MIB::ifIndex.42 = INTEGER: 42
IF-MIB::ifAdminStatus.42 = INTEGER: up(1) IF-MIB::ifOperStatus.42 = INTEGER: up(1) FORTINET-
CORE-MIB::fnSysSerial.0 = STRING: FG140P3G15800330 SNMPv2-MIB::sysName.0 = STRING:
FortiGate-140D-POE

2019-01-31 14:12:01 <UNKNOWN> [UDP: [10.1.100.1]:162->[10.1.100.11]:162]: DISMAN-EVENT-
MIB::sysUpTimeInstance = Timeticks: (89798) 0:14:57.98 SNMPv2-MIB::snmpTrapOID.0 = OID: IF-
MIB::linkUp IF-MIB::ifIndex.42 = INTEGER: 42 IF-MIB::ifAdminStatus.42 = INTEGER: up(1) IF-
MIB::ifOperStatus.42 = INTEGER: up(1) FORTINET-CORE-MIB::fnSysSerial.0 = STRING:
FG140P3G15800330 SNMPv2-MIB::sysName.0 = STRING: FortiGate-140D-POE
```


Netskope Exhibit 1017

fgFmTrapIfChange trap

This trap is sent when any changes are detected on the interface. The change can be very simple, such as giving an IPV4 address.

For example, the user has given the IP address of 1.2.3.4/24 to port 1 and the EMS Manager has detected the following trap:

```
DISMAN-EXPRESSION-MIB::sysUpTimeInstance = Timeticks: (7975058) 22:09:10.58 SNMPv2-
MIB::snmpTrapOID.0 = OID: FORTINET-FORTIGATE-MIB::fgFmTrapIfChange FORTINET-CORE-
MIB::fnSysSerial.0 = STRING: FG140P3G15800330 IF-MIB::ifName.45 = STRING: port1 FORTINET-
FORTIGATE-MIB::fgManIfIp.0 = IpAddress: 1.2.3.4 FORTINET-FORTIGATE-MIB::fgManIfMask.0 =
IpAddress: 255.255.255.0 FORTINET-FORTIGATE-MIB::fgManIfIp6.0 = STRING: 0:0:0:0:0:0:0:0
```

entConfigChange trap

The change to the interface in the previous example has also triggered the *ConfChange Trap* which is sent along with the *fgFmTrapIfChange* trap:

```
2018-11-15 09:30:23 FGT_A [UDP: [172.16.200.1]:162->[172.16.200.55]:162]: DISMAN-EXPRESSION-
MIB::sysUpTimeInstance = Timeticks: (8035097) 22:19:10.97 SNMPv2-MIB::snmpTrapOID.0 = OID:
ENTITY-MIB::entConfigChange
```

fgTrapDeviceNew trap

This trap is triggered when a new device, like a FortiSwitch, is connected to the FortiGate.

For example, the following scenario has given the device a new trap for adding FortiAP on a PoE interface a FortiGate 140D-POE. The trap has important information about the device name, device MAC address, and when it was last seen.

```
2018-11-15 11:17:43 UDP/IPv6: [2000:172:16:200::1]:162 [UDP/IPv6: [2000:172:16:200::1]:162]:
DISMAN-EXPRESSION-MIB::sysUpTimeInstance = Timeticks: (520817) 1:26:48.17 SNMPv2-
MIB::snmpTrapOID.0 = OID: FORTINET-FORTIGATE-MIB::fgTrapDeviceNew FORTINET-CORE-
MIB::fnSysSerial.0 = STRING: FG140P3G15800330 SNMPv2-MIB::sysName.0 = STRING: FGT_A IF-
MIB::ifIndex.0 = INTEGER: 0 FORTINET-FORTIGATE-MIB::fgVdEntIndex.0 = INTEGER: 0 FORTINET-
FORTIGATE-MIB::fgDeviceCreated.0 = Gauge32: 5 FORTINET-FORTIGATE-MIB::fgDeviceLastSeen.0 =
Gauge32: 5 FORTINET-FORTIGATE-MIB::fgDeviceMacAddress.0 = STRING: 90:6c:ac:f9:97:a0

2018-11-15 11:17:43 FGT_A [UDP: [172.16.200.1]:162->[172.16.200.55]:162]: DISMAN-EXPRESSION-
MIB::sysUpTimeInstance = Timeticks: (520817) 1:26:48.17 SNMPv2-MIB::snmpTrapOID.0 = OID:
FORTINET-FORTIGATE-MIB::fgTrapDeviceNew FORTINET-CORE-MIB::fnSysSerial.0 = STRING:
FG140P3G15800330 SNMPv2-MIB::sysName.0 = STRING: FGT_A IF-MIB::ifIndex.0 = INTEGER: 0
FORTINET-FORTIGATE-MIB::fgVdEntIndex.0 = INTEGER: 0 FORTINET-FORTIGATE-
MIB::fgDeviceCreated.0 = Gauge32: 5 FORTINET-FORTIGATE-MIB::fgDeviceLastSeen.0 = Gauge32: 5
FORTINET-FORTIGATE-MIB::fgDeviceMacAddress.0 = STRING: 90:6c:ac:f9:97:a0
```

fgTrapAvOversize trap

The `fgTrapAvOversize` trap is generated when the antivirus scanner detects an oversized file:

```
019-01-31 13:22:04 10.1.100.1(via UDP: [10.1.100.1]:162->[10.1.100.11]:162) TRAP, SNMP v1,
community REGR-SYS FORTINET-FORTIGATE-MIB::fgt140P Enterprise Specific Trap (602) Uptime: 1
day, 3:41:10.31 FORTINET-CORE-MIB::fnSysSerial.0 = STRING: FG140P3G15800330 SNMPv2-
MIB::sysName.0 = STRING: FortiGate-140D-POE 2019-01-31 13:22:29 <UNKNOWN> [UDP:
```


Netskope Exhibit 1017

```
[10.1.100.1]:162->[10.1.100.11]:162]: DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks:
(9967031) 1 day, 3:41:10.31 SNMPv2-MIB::snmpTrapOID.0 = OID: FORTINET-FORTIGATE-
MIB::fgTrapAvOversize FORTINET-CORE-MIB::fnSysSerial.0 = STRING: FG140P3G15800330 SNMPv2-
MIB::sysName.0 = STRING: FortiGate-140D-POE
```

SNMP traps and query for monitoring DHCP pool

The SNMP DHCP event contains three traps and one query.

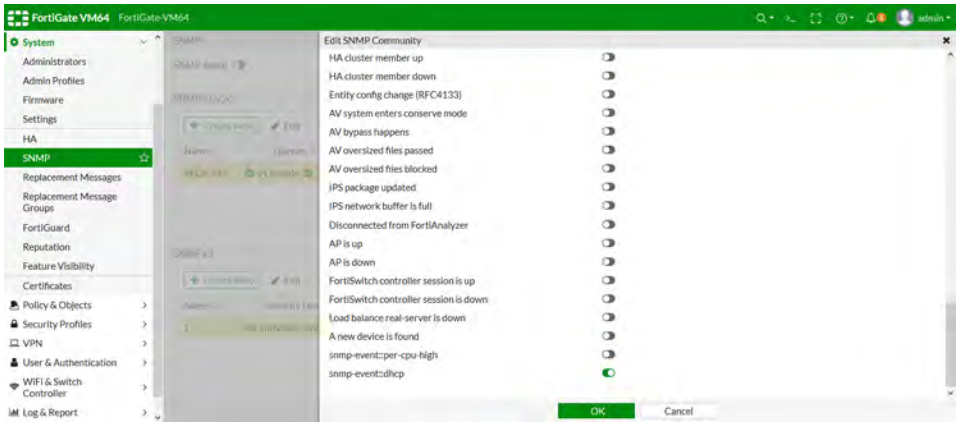
Traps are sent when:

- DHCP server IP pool usage reaches 90%
- DHCP server detect an IP address that is already in use
- DHCP client receives DHCP NAK

SNMP queries are accepted for DHCP lease usage information (OID = 1.3.6.1.4.1.12356.101.23). The query result is based on the leased out percentage.

To enable the SNMP DHCP event in the GUI:

1. Go to *System* > *SNMP*.
2. Click *Create New* in either the *SNMP v1/v2c* table or *SNMP v3* table, or edit an existing community or user.
3. Configure the settings as required.
4. In the *SNMP Events* list, enable *snmp-event::dhcp*.



5. Click *OK*.

To enable the SNMP DHCP event in the CLI:

```
config system snmp community
    edit 1
        set name "REGR-SYS"
        config hosts
            edit 1
                set ip 10.1.100.11 255.255.255.255
            next
            edit 2
                set ip 172.16.200.55 255.255.255.255
            next
        end
```

```
        set events dhcp
    next
end

config system snmp user
    edit "1"
        set notify-hosts 172.10.1.0 172.20.1.0
        set events dhcp
        set security-level auth-priv
        set auth-proto sha384
        set auth-pwd ********************
        set priv-proto aes256
        set priv-pwd *********************
    next
end
```

Replacement messages

The replacement message list in *System* > *Replacement Messages* enables you to view and customize replacement messages. Highlight the replacement messages you want to edit and customize the message content to your requirements. Hit *Save* when done. If you do not see the message you want to edit, select the *Extended View* option in the upper right-hand corner of the screen.

If you make a mistake, select *Restore Default* to return to the original message and code base.

Replacement message images

You can add images to replacement messages on:

- Disclaimer pages
- Login pages
- Declined disclaimer pages
- Login failed pages
- Login challenge pages
- Keepalive pages



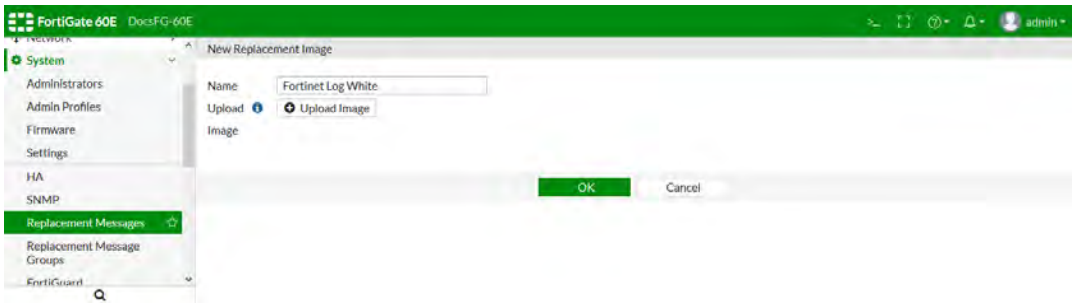
Supported image formats are GIF, JPEG, TIFF, and PNG. The maximum file size supported is 24KB.

Adding images to replacement messages

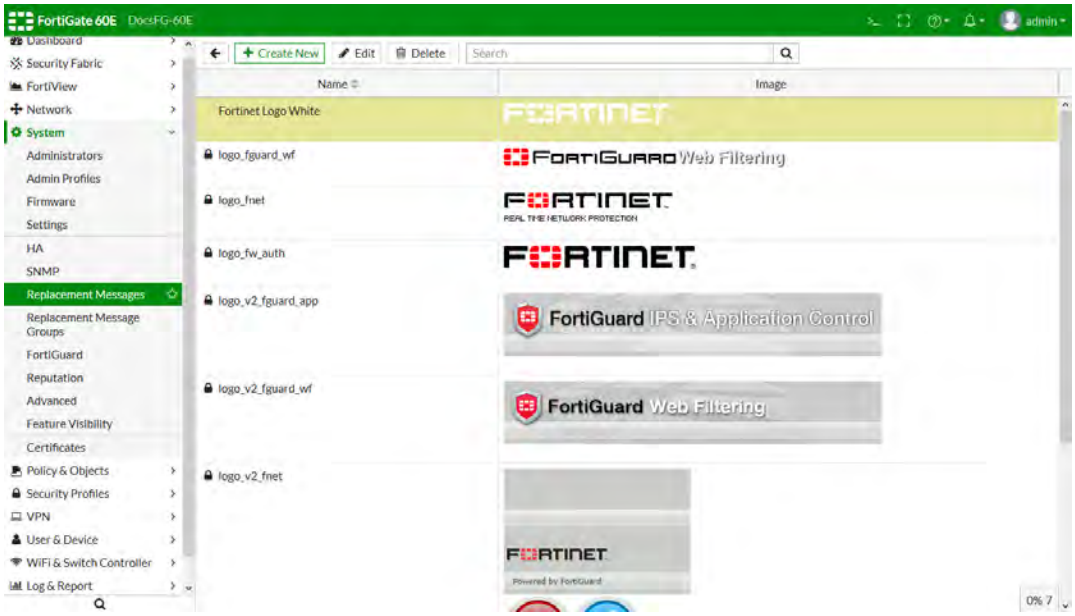
To add images to replacement messages in the GUI:

1. Go to *System* > *Replacement Messages*.
2. Click *Manage Images* at the top of the page.
3. Click *Create New*.
4. Enter a name for the image.

5. Click *Upload Image* and locate the file.



6. Click *OK*.



To add images to replacement messages in the CLI:

```
config system replacemsg-image
    edit <image_name>
        set image-type {gif | jpg | tiff | png}
        set image-base64 <string>
    next
end
```

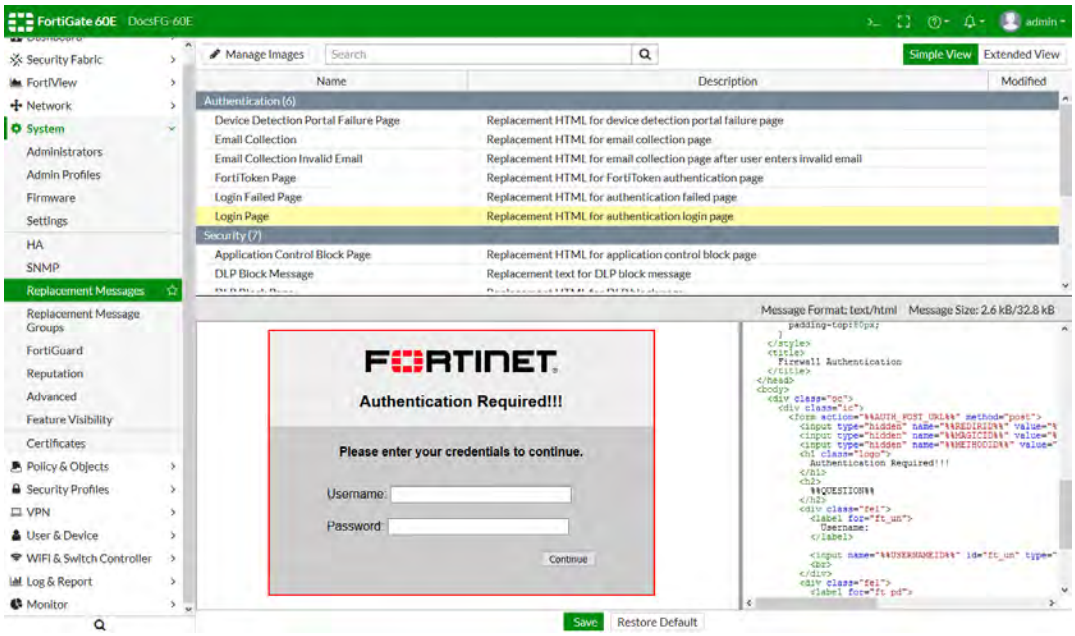
Modifying replacement messages

Replacement messages can be modified to include an HTML message or content that suits your organization. A list of common replacement messages appear in the main window. Select *Extended View* to see the entire list and all categories for replacement messages.

To modify a replacement message in the GUI:

1. Go to *System* > *Replacement Messages*.
2. Select the replacement message you want to edit.

3. In the bottom pane of the GUI the message will be displayed on the left alongside the HTML code on the right. Edit the HTML code as needed.
   The message view changes in real-time as you edit the content.



4. Click *Save*.

To modify a replacement message in the CLI:

```
config system replacemsg <message_category> <message_type>
    set buffer <string>
    set header {none | http | 8bit}
    set format {none | text | html}
end
```

For example, to modify the *Traffic Quota Limit Exceeded Page* message:

```
config system replacemsg traffic-quota "per-ip-shaper-block"
    set buffer "<html>
  <head>
    <title>
      Traffic Quota Control
    </title>
  </head>
  <body>
    <font size=2>
      <table width=\"100%\">
        <tr>
          <td bgcolor=#3300cc align=\"center\" colspan=2>
            <font color=#ffffff>
              <b>
                Traffic blocked because exceeded session quota
              </b>
            </font>
          </td>
        </tr>
      </table>
```

```
      <br>
      <br>
      Traffic blocked because it exceeded the per IP shaper session quota. Please contact
the system administrator.
      <br>
      %%QUOTA_INFO%%
      <br>
      <br>
      <hr>
    </font>
  </body>
</html>"
    set header http
    set format html
end
```

Replacement message groups

Replacement message groups allow you to customize replacement messages for individual policies and profiles.

There are two types of replacement message groups:

- `utm`: Used with UTM settings in firewall policies. Messages in the following categories can be customized: mail, http, webproxy, ftp, nntp, fortiguard-wf, spam, alertmail, admin, sslvpn, nac-quar, traffic-quota, utm, custom-message, and icap.
- `auth`: Used with authentication pages in firewall policies. Messages in the following categories can be customized: webproxy and auth.

The messages added to a group do not need to be customized. The body content, header type, and format of a message will use the default values if not customized.

To create or edit a replacement message group in the CLI:

```
config system replacemsg-group
    edit <group>
        set group-type {auth | utm}
        config <message_category>
            edit <message_type>
                set buffer <message>
                set header {none | http | 8bit}
                set format {none | text | html}
            next
        end
    next
end
```

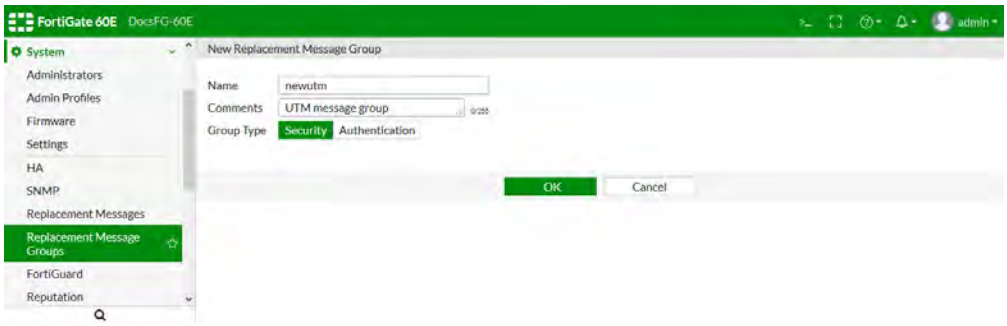
To create a replacement message group in the GUI:

1. Make replacement message groups visible in the GUI with the following CLI command:

   ```
   config system settings
       set gui-replacement-message-groups enable
   end
   ```


Netskope Exhibit 1017

2. Go to *System* > *Replacement Message Groups*.
3. Click *Create New*.
4. Enter a name for the new group.
5. Optionally, enter a comment describing the group.
6. Select the *Group Type*, either *Security* or *Authentication*.



7. Click *OK*.

Example

In this example, two replacement message groups are created. The UTM type message group includes custom mail related messages, changes the formats of some spam related message, and is assigned to an email filter profile. The authentication type message group has a custom authentication success message that is applied to a proxy-based firewall policy that with the email filter profile assigned.

To create the replacement message groups and use them in a profile and a policy in the CLI:

1. Create the replacement message groups:

```
config system replacemsg-group
    edit "newutm"
        set group-type utm
        config mail
            edit "partial"
                set buffer "Fragmented emails are blocked, sorry."
            next
            edit "email-av-fail"
                set buffer "The email has been blocked for reasons."
            next
        end
        config spam
            edit "submit"
                set header http
                set format html
            next
            edit "reversedns"
                set header http
                set format html
            next
        end
    next
    edit "newauth"
        set group-type auth
```

```
        config auth
            edit "auth-success-msg"
                set buffer "Welcome to the firewall. Your authentication has been
accepted, please reconnect."
                set header none
                set format text
            next
        end
    next
end
```

2. Apply the message groups:

```
config emailfilter profile
    edit "newmsgs"
        set replacemsg-group "newutm"
    next
end

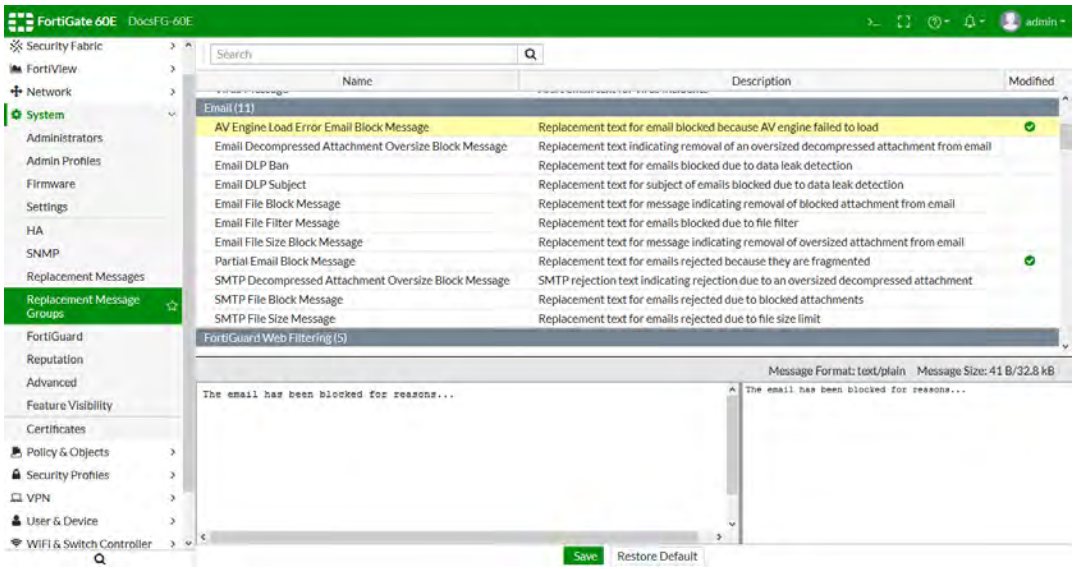
config firewall policy
    edit 1
        ...
        set replacemsg-override-group "newauth"
        set inspection-mode proxy
        set emailfilter-profile "newmsgs"
        ...
    next
end
```

To create the replacement message groups and use them in a profile and a policy in the GUI:

1. Create the *Security* replacement message groups:
   a. Go to *System* > *Replacement Message Groups*.
   b. Click *Create New*.
   c. Enter *newutm* in the *Name* field.
   d. Enter *UTM message group* in the *Comments* field.
   e. Select *Security* as the *Group Type*.
   f. Click *OK*.
2. Customize the replacement messages in the *newutm* group:
   a. Go to *System* > *Replacement Message Groups*.
   b. Edit the *newutm* group.

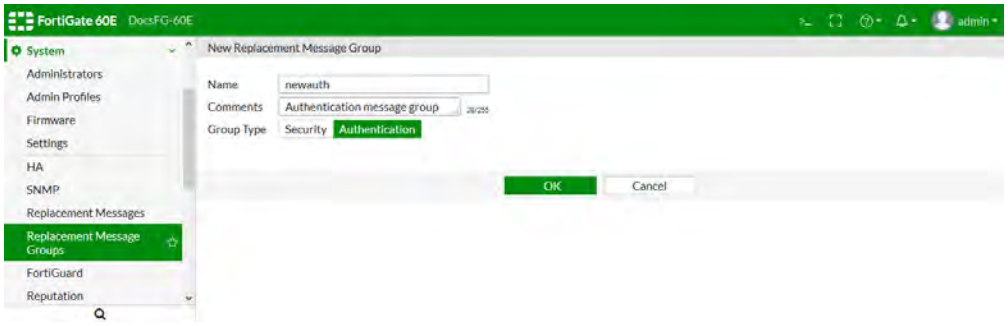
Netskope Exhibit 1017

c. Locate the *AV Engine Load Error Email Block Message*, edit the message, then click *Save*.



d. Locate the *Partial Email Block Message*, edit the message, then click *Save*.

3. Create the *Authentication* replacement message group:
   a. Go to *System Replacement Message Groups*.
   b. Click *Create New*.
   c. Enter *newauth* in the *Name* field.
   d. Enter *Authentication message group* in the *Comments* field.
   e. Select *Authentication* as the *Group Type*.



   f. Click *OK*.
4. Apply the *newutm* replacement message group to an email filter profile using the CLI.
5. Apply the *newauth* replacement message group and the email filter profile to a firewall policy using the CLI.

FortiGuard

FortiGuard services can be purchased and registered to your FortiGate unit. The FortiGate must be connected to the Internet in order to automatically connect to the FortiGuard Distribution Network (FDN) to validate the license and download FDN updates.

The FortiGuard subscription update services include:

- Antivirus (AV)
- Intrusion Protection Service (IPS)
- Application Control
- Antispam
- Web Filtering
- Web Application Firewall (WAF)

To view FDN support contract information, go to *System* > *FortiGuard*. The *License Information* table shows the status of your FortiGate's support contract.

- IPv6 FortiGuard connections on page 680
- Configuring antivirus and IPS options on page 681
- Manual updates on page 681
- Automatic updates on page 682
- Sending malware statistics to FortiGuard on page 684
- Update server location on page 685
- Filtering on page 685
- Override FortiGuard servers on page 686
- Online security tools on page 687
- FortiGuard anycast and third-party SSL validation on page 687
- Using FortiManager as a local FortiGuard server on page 690
- Cloud service communication statistics on page 691
- IoT detection service on page 692

IPv6 FortiGuard connections

The Fortinet DNS can resolve FortiGuard related servers to both IPv4 and IPv6 addresses. FortiOS daemons (update, forticldd, url) connect using either IPv4 or IPv6 addresses. The first available connection will be used for updates or the rating service.

To configure an interface and route for IPv6:

```
config system interface
    edit "wan1"
        set vdom "root"
        config ipv6
            set ip6-address 2000:172:16:200::1/64
        end
    next
end

config router static6
    edit 1
        set gateway 2000:172:16:200::254
        set device "wan1"
    next
end
```


Netskope Exhibit 1017

To configure push updates:

```
config system autoupdate push-update
    set status enable
    set override enable
    set address "2620:101:9005:3860::94"
end
```

Configuring antivirus and IPS options

To configure antivirus and IPS options:

1. Go to *System* > *FortiGuard*
2. Scroll down to the *AntiVirus & IPS Updates* section.
3. Configure the antivirus and IPS options for connecting and downloading definition files:

Accept push updates	Enable to allow updates to be sent automatically to your FortiGate. New definitions will be added as soon as they are released by FortiGuard. See Push updates on page 683.
Use override push	Only available if *Accept push updates* is enabled. See Override push on page 684.
Scheduled Updates	Enable to schedule updates to be sent to the FortiGate at the specified time. See Scheduled updates on page 682.
Improve IPS quality	Enable to send information to the FortiGuard servers when an attack occurs. This can help keep the FortiGuard database current as attacks evolve, and improve IPS signatures.
Use extended IPS signature package	Enable to use the extended IPS database, that includes protection from legacy attacks, along with the regular IPS database that protects against the latest common and in-the-wild attacks.
Update AV & IPS Definitions	Click to manually initiate an FDN update.

4. Click *Apply*.

Manual updates

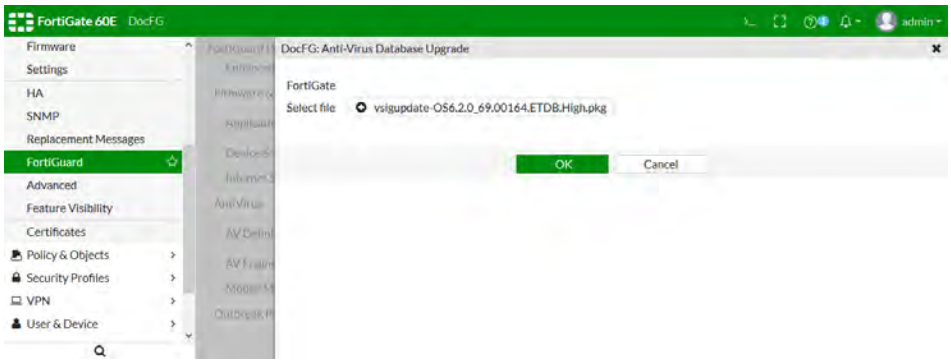
When needed, FortiGuard Distribution Network (FDN) updates can be manually uploaded.

To manually update the signature definitions files:

1. Log in to the Fortinet Support website.
2. Go to *Download* > *FortiGuard Service Updates*.
3. Select your *OS Version* from the dropdown list.
4. Locate your device in the table, and download the signature definitions files.
5. On the FortiGate, go to *System FortiGuard*.
6. In the *License Information* table, locate the row of the definitions that you are updating, and click *Upgrade Database* in the rightmost column.

Netskope Exhibit 1017

7. In the pane that opens, click *Upload*, locate the downloaded definitions file on your computer, then click *Open*. The download may take a few minutes to complete.



8. Click *OK*.

Automatic updates

The FortiGate can be configured to request updates from FDN on a schedule, or via push notification.

Scheduled updates

Scheduling updates ensures that the virus and IPS definitions are downloaded to your FortiGate on a regular basis.

Updating definitions can cause a brief disruption in traffic that is currently being scanned while the FortiGate unit applies the new signature database. Updates should be scheduled during off-peak hours when network usage is at a minimum to ensure that network activity will not be affected by downloading the definitions files.

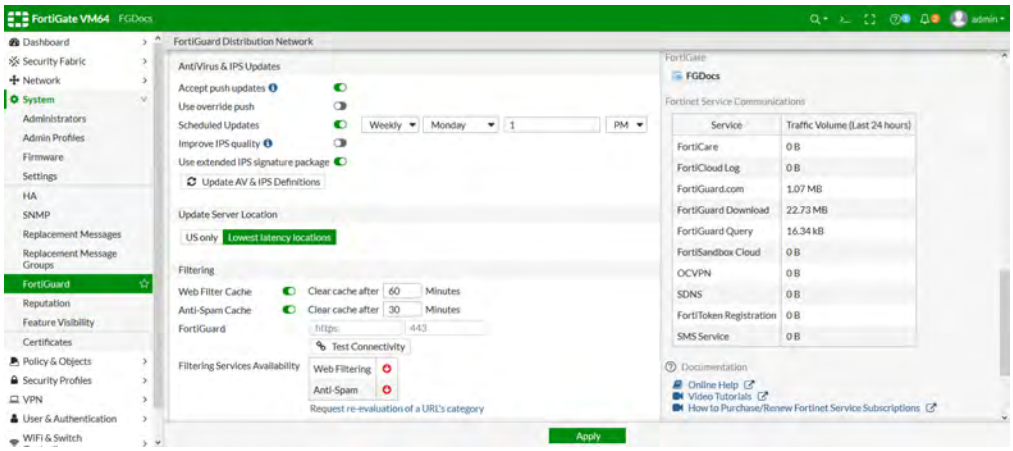


A schedule of once a week means any urgent updates will not be pushed until the scheduled time. If an urgent update is required, click the *Update AV & IPS Definitions* button to manually update the definitions.

To configure scheduled updates in the GUI:

1. Go to *System* > *FortiGuard*
2. Scroll down to the *AntiVirus & IPS Updates* section.
3. Enable *Scheduled Updates*.

4. Configure the update schedule.



5. Click *Apply*.

To configure scheduled updates in the CLI:

```
config system autoupdate schedule
    set status enable
    set frequency {every | daily | weekly}
    set time <hh:mm>
    set day <day_of_week>
end
```

Push updates

Push updates enable you to get immediate updates when new viruses or intrusions are discovered and new signatures are created. This ensures that the latest signature are sent to the FortiGate as soon as possible.

When a push notification occurs, the FortiGuard server sends a notice to the FortiGate that a new signature definition file available. The FortiGate then initiates a download of the definition file. For maximum security, both scheduled and push updates should be enabled.

To enable push updates - GUI:

1. Go to *System* > *FortiGuard*
2. Scroll down to the *AntiVirus & IPS Updates* section.
3. Enable *Accept push updates*.
4. Click *Apply*.

To enable push updates in the CLI:

```
config system autoupdate push-update
    set status enable
    set override {enable | disable}
    set address <vip_address>
end
```

Override push

If the FortiGate is behind a NAT device (or another FortiGate), or if your organization provides updates using their own FortiGuard server, an override server must be used to ensure that the FortiGate receives push update notifications. The FDS will connect to the NAT device when attempting to reach the FortiGate, and the NAT device must be configured to forward FDS traffic to the FortiGate on UDP port 9443.

Push updates must be enabled to configure a push update override.

For example, if the NAT device is another FortiGate:

1. On the FortiGate NAT device, add a port forwarding virtual IP address in *Policy & Objects* > *Virtual IPs*. See for details.
2. On the FortiGate NAT device, add a security policy that connects to the internet and includes the port forwarding VIP.
3. On the internal FortiGate device, configure *Push update override*.

To configure push update override in the GUI:

1. Go to *System* > *FortiGuard*
2. Scroll down to the *AntiVirus & IPS Updates* section.
3. Enable *Accept push updates*.
4. Enable *Use override push*.
5. Enter the IP address and port number configured on the NAT device.
6. Click *Apply*.

To configure push update override in the CLI:

```
config system autoupdate push-update
    set status enable
    set override {enable | disable}
    set address <vip_address>
end
```

Sending malware statistics to FortiGuard

FortiGate devices periodically send encrypted antivirus, IPS, botnet IP list, and application control statistics to FortiGuard. Included with these data is the IP address and serial number of the FortiGate, and the country that it is in. This information is never shared with external parties, Fortinet Privacy Policy.

The malware statistics are used to improve various aspects of FortiGate malware protection. For example, antivirus data allow FortiGuard to determine what viruses are currently active. Signatures for those viruses are kept in the Active AV Signature Database that is used by multiple Fortinet products.Inactive virus signatures are moved to the Extended AV Signature Database (see Configuring antivirus and IPS options on page 681). When events for inactive viruses start appearing in the malware data, the signatures are moved back into the AV Signature Database.

The FortiGate and FortiGuard servers go through a 2-way SSL/TLS 1.2 authentication before any data is transmitted. The certificates used in this process must be trusted by each other and signed by the Fortinet CA server.

The FortiGate only accepts data from authorized FortiGuard severs. Fortinet products use DNS to find FortiGuard servers and periodically update their FortiGate server list. All other servers are provided by a list that is updated through the encrypted channel.

Netskope Exhibit 1017

Malware statistics are accumulated and sent every 60 minutes by default.

To configure sharing this information, use the following CLI command:

```
config system global
    set fds-statistics {enable | disable}
    set fds-statistics-period <minutes>
end
```



The submission of malware data is in accordance with the Fortinet Privacy Policy.

There is no sensitive or personal information included in these submissions. Only malware statistics are sent.

Fortinet uses the malware statistics collected in this manner to improve the performance of the FortiGate services and to display statistics on the Fortinet Support website for customers registered FortiGate devices.

Fortinet may also publish or share statistics or results derived from this malware data with various audiences. The malware statistics shared in this way do not include any customer data.

Update server location

The location of the FortiGuard update server that the FortiGate connects to can be set to either only servers in the USA only, or to the servers with the lowest latency.

On hardware FortiGate devices, the default is *Lowest latency locations*. On VM devices, the default is *US only*.

To configure the update server location in the GUI:

1. Go to *System* > *FortiGuard*
2. Scroll down to the *Update Server Location* section.
3. Select *US only* or *Lowest latency locations*.
4. Click *Apply*.

To configure the update server location in the CLI:

```
config system fortiguard
    set update-server-location {usa | any}
end
```

Filtering

Web filtering is used to block access to harmful, inappropriate, and dangerous web sites (see FortiGuard filter on page 901).

Email filtering is used to detect and block spam messages (see FortiGuard-based filters on page 985).

To configure filtering in the GUI:

1. Go to *System* > *FortiGuard*
2. Scroll down to the *Filtering* section.

3. Configure the settings as needed:

Web Filter Cache	Enable/disable web filter cache, and set the amount of time that the FortiGate will store a blocked IP address or URL locally. After the time expires, the FortiGate contacts the FDN to verify the address.
Anti-Spam Cache	Enable/disable email filter cache, and set the amount of time that the FortiGate will store an email address locally.
FortiGuard	The protocol and port used to contact the FortiGuard servers. These options can be changed in the CLI.
Filtering Service Availability	The status of the filtering service. Click *Test Connectivity* if the filtering service is not available.
Request re-evaluation of a URL's category	Click to re-evaluate a URL category rating on the FortiGuard web filter service.

4. Click *Apply*.

To configure filtering in the CLI:

```
config system fortiguard
    set protocol {https | udp}
    set port {443 | 53 | 8888}
    set antispam-force-off {enable | disable}
    set antispam-cache {enable | disable}
    set antispam-cache-ttl <integer>
    set antispam-cache-mpercent <percent>
    set antispam-timeout <integer>
    set webfilter-force-off {enable | disable}
    set webfilter-cache {enable | disable}
    set webfilter-cache-ttl <integer>
    set webfilter-timeout <integer>
end
```



When anycast is enabled (by default) the protocol is HTTPS and the port is 443.

Override FortiGuard servers

By default, FortiOS will update signature packages and query rating servers using public FortiGuard servers. This list can be overridden by adding servers to the override server list. Communication with public FortiGuard servers can also be disabled.

To add an override FortiGuard server in the GUI:

1. Go to *System* > *FortiGuard*
2. Scroll down to the *Override FortiGuard Servers* section.
3. In the table, click *Create New*. The *Create New Override FortiGuard Server* pane opens.
4. Select the server address type: *IPv4*, *IPv6*, or *FQDN*.

Netskope Exhibit 1017

5. Enter the server address of the selected type in the *Address* field.
6. Select the type of server: *AntiVirus & IPS Updates*, *Filtering*, or *Both*.
7. Click *Apply*.

To add an override FortiGuard server in the CLI:

```
config system central-management
    set type fortiguard
    config server-list
        edit <integer>
            set server-type {update rating}
            set server-address <ip_address>
        next
    end
end
```

Online security tools

FortiGuard Labs provides a number of online security tools, including but not limited to:

- **URL lookup**
  Enter a website address to see if it has been rated and what category and classification it is filed as. If you find a site that has been wrongly categorized, use this page to request that the site be re-evaluated: https://www.fortiguard.com/webfilter
- **Threat Encyclopedia**
  Browse FortiGuard Labs extensive encyclopedia of threats. Search for viruses, botnet C&C, IPS, endpoint vulnerabilities, and mobile malware: https://www.fortiguard.com/encyclopedia
- **Application Control**
  Browse FortiGuard Labs extensive encyclopedia of applications: https://www.fortiguard.com/appcontrol

FortiGuard anycast and third-party SSL validation

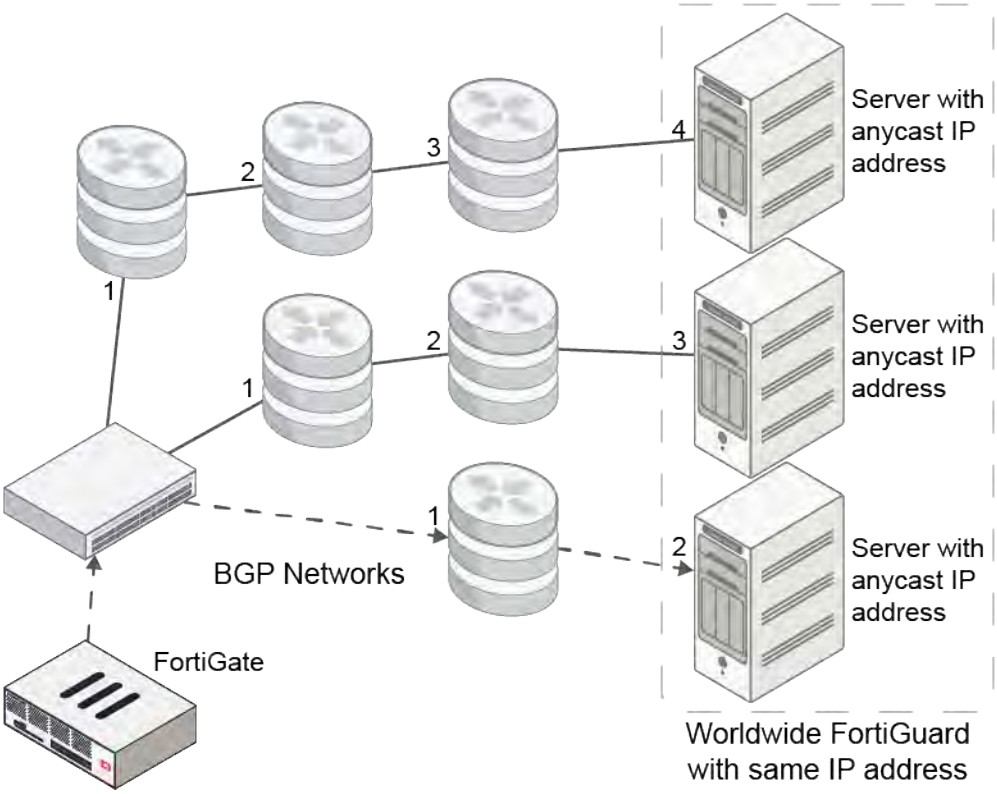
Anycast optimizes routing performance to FortiGuard servers. It is the default FortiGuard access mode.

Using Fortinet DNS servers, the FortiGate receives a single IP address for the domain name of each FortiGuard service. BGP routing optimization is transparent to the FortiGate. The domain name of each FortiGuard service is the common name in that service's certificate, which is signed by a third-party intermediate CA. The FortiGuard server uses third-party certificate verification and the Online Certificate Status Protocol (OCSP) stapling check, so that the FortiGate can always validate the FortiGuard server certificate efficiently.

FortiGate will only complete the TLS handshake with an anycast server that has a good OCSP status for its certificate. Any other status will result in a failed SSL connection. OCSP stapling is reflected on the signature interval so that *good* means that the certificate is not revoked at that timestamp. The FortiGuard servers query the CA's OCSP responder every four hours and update its OCSP status. If the FortiGuard is unable to reach the OCSP responder, it will keep the last known OCSP status for up to seven days. This cached OCSP status will be sent out immediately when a client connection request is made, optimizing the response time.

FortiGuard represents all cloud based servers; see Anycast and unicast services for details.

The anycast server has one IP address to match its domain name. The FortiGate connects with a single server address, using HTTPS and port 443, regardless of where the FortiGate is located.

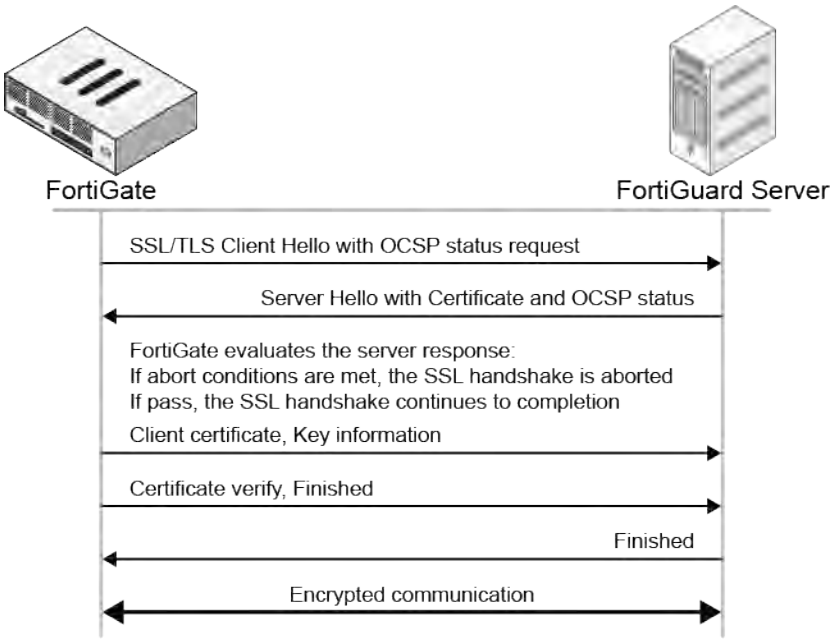


To configure the anycast FortiGuard access mode:

```
config system fortiguard
    set fortiguard-anycast {enable | disable}
    set fortiguard-anycast-source {fortinet | aws}
end
```

Connection process

The following process is used to connect to an anycast server:



1. The FortiGate embeds the CA_bundle certificate, which includes the root CA with CRL list and third-party intermediate CA, in the root CA level.
2. The FortiGate finds the FortiGuard IP address from its domain name from DNS.
3. The FortiGate starts a TLS handshake with the FortiGuard IP address. The client hello includes an extension of the *status request*.
4. The FortiGuard servers provide a certificate with its OCSP status: *good*, *revoked*, or *unknown*.
5. The FortiGate verifies the CA chain against the root CA in the CA_bundle.
6. The FortiGate verifies the intermediate CA's revoke status against the root CA's CRL.
7. The FortiGate verifies the FortiGuard certificate's OCSP status:

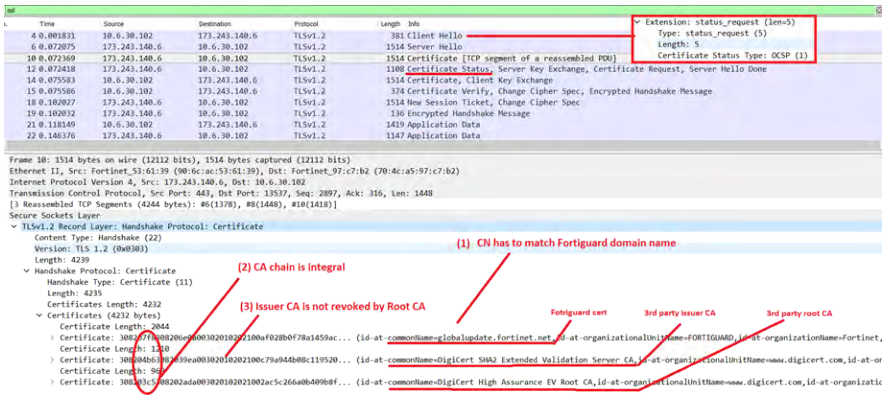
```
OCSP Response Data:
    OCSP Response Status: successful (0x0)
    Response Type: Basic OCSP Response
    Version: 1 (0x0)
    Responder Id: 3DD350A5D6A0ADEEF34A600A65D321D4F8F8D60F
    Produced At: Aug 20 07:50:58 2019 GMT
    Responses:
    Certificate ID:
      Hash Algorithm: sha1
      Issuer Name Hash: 49F4BD8A18BF760698C5DE402D683B716AE4E686
      Issuer Key Hash: 3DD350A5D6A0ADEEF34A600A65D321D4F8F8D60F
      Serial Number: 02555C9F3901B799DF1873402FA9392D
    Cert Status: good
    This Update: Aug 20 07:50:58 2019 GMT
    Next Update: Aug 27 07:05:58 2019 GMT
```

   Abort conditions include:

   - The CN in the server's certificate does not match the domain name resolved from the DNS.
   - The OCSP status is not good.
   - The issuer-CA is revoked by the root-CA.
8. Once the SSL handshake is established, the FortiGate can engage the server.

Netskope Exhibit 1017

Example Wireshark PCAP:



Using FortiManager as a local FortiGuard server

FortiManager can provide a local FortiGuard server with port 443 access.

Anycast FortiGuard settings force the rating process to use port 443, even with an override server. Using a unique address in the same subnet as the FortiManager access IP address, the FortiManager can provide local FortiGuard updates and rating access with a dedicated IP address and port 443.

To use a FortiManager as a local FortiGuard server:

```
config system central-management
    set type fortimanager
    set fmg "172.18.37.148"
    config server-list
        edit 1
            set server-type update
            set server-address 172.18.37.150
        next
        edit 2
            set server-type rating
            set server-address 172.18.37.149
        next
    end
    set fmg-update-port 443
    set include-default-servers enable
end
```

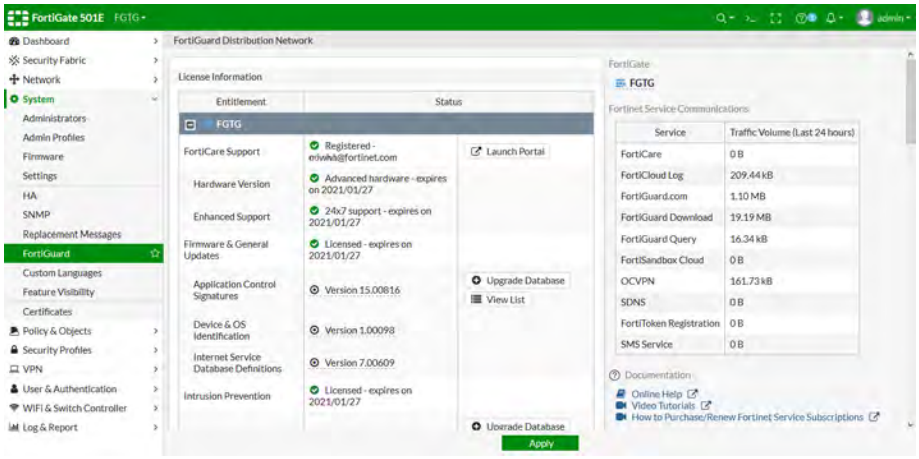
When `fmg-update-port` is set to `443`, the update process will use port 443 to connect to the override update server, which is the local FortiGuard server in the FortiManager. If this is not set, the update process will use port 8890, and the server address setting has to be the FortiManager access IP address. Override FortiGuard services come from the server list that is the local FortiGuard server in the FortiManager, and use the traditional, non-OCSP TLS handshake. If override servers in the FortiManager are not available, the default FortiGuard servers are connected, and the anycast OCSP TLS handshake is used.

Cloud service communication statistics

Fortinet service communications statistics are displayed on the *FortiGuard* page. The statistics correspond with the output from `diagnose sys service-communication`. The traffic volume values in the GUI are the sums of data from the last 24 hours.

To view Fortinet service communications statistics:

1. Go to *System* > *FortiGuard*.
   The *Fortinet Service Communications* statistics are displayed on the right-side of the screen:



These statistics correspond to the following:

```
# diagnose sys service-communication
FortiCare:
The last 1 hour(in bytes):  0 0 0 0 0 0 0 0 0 0 0 0
The last 24 hours(in bytes):  0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
The last 7 days(in bytes):  0 0 0 0 0 0 0
FortiGuard Download:
The last 1 hour(in bytes):  0 0 0 0 0 0 0 0 0 0 0 0
The last 24 hours(in bytes):  0 0 195752 0 21051904 36342800 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0 0
The last 7 days(in bytes):  57590456 0 0 0 0 0 0
FortiGuard Query:
The last 1 hour(in bytes):  0 0 0 0 0 0 0 0 0 0 0 0
The last 24 hours(in bytes):  0 0 2805 0 1298 1709 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
The last 7 days(in bytes):  5812 0 0 0 0 0 0
FortiCloud Log:
The last 1 hour(in bytes):  0 0 0 0 0 0 0 0 0 0 0 0
The last 24 hours(in bytes):  0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
The last 7 days(in bytes):  0 0 0 0 0 0 0
FortiSandbox Cloud:
The last 1 hour(in bytes):  0 0 0 0 0 0 0 0 0 0 0 0
The last 24 hours(in bytes):  0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
The last 7 days(in bytes):  0 0 0 0 0 0 0
FortiGuard.com:
The last 1 hour(in bytes):  2014 0 1329 0 1329 0 1329 0 1329 0 2020 103930
The last 24 hours(in bytes):  3343 112595 9032861 18584815 17757745 16054191 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0
The last 7 days(in bytes):  61545550 0 0 0 0 0 0
```


Netskope Exhibit 1017

```
FortiToken Registration:
The last 1 hour(in bytes):  0 0 0 0 0 0 0 0 0 0 0 0
The last 24 hours(in bytes):  0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
The last 7 days(in bytes):  0 0 0 0 0 0 0
SMS Service:
The last 1 hour(in bytes):  0 0 0 0 0 0 0 0 0 0 0 0
The last 24 hours(in bytes):  0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
The last 7 days(in bytes):  0 0 0 0 0 0 0
```

IoT detection service

Internet of Things (IoT) detection is a subscription service that allows FortiGate to detect unknown devices in FortiGuard that are not detected by the local Device Database (CIDB). When the service is activated, FortiGate can send device information to the FortiGuard collection server. When a new device is detected, FortiGate queries the results from the FortiGuard query for more information about the device.

This feature requires an IoT Detection Service license.

FortiGate device requirements:

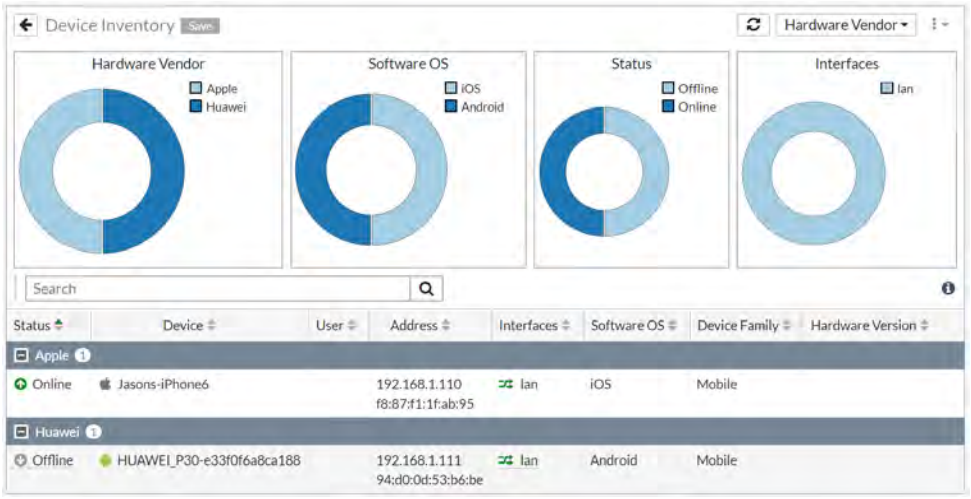
The FortiGate device must be:

- Registered with FortiCare
- Connected to an anycast FortiGuard server

How the service works:

1. Enable Device Detection on an interface..
2. FortiGate uses the interface to detect device traffic flow.
3. Upon detecting traffic from an unknown device, FortiGate sends the device data to the FortiGuard collection server.
4. The collection server returns data about the new device to the FortiGuard query server.
5. If the device signature does not appear in the local Device Database (CIDB) or some fields are not complete, FortiGate queries FortiGuard for more information about the device.

To view the latest device information in the GUI, go to *Dashboard* > *Users & Devices* and expand the *Device Inventory* widget.



To debug the daemon in the CLI:

1. Disable the local device database in order to force all queries to go to FortiGuard.
   ```
   diagnose src-vis local-sig disable
   ```
2. Enable iotd debugs.
   ```
   diagnose debug application iotd -1
   diagnose debug enable
   ```
   FortiGate sends the device data to the FortiGuard collection server.
   ```
   FortiWiFi-60E # [iotd] recv request from caller size:61
   [iotd] service:collect hostname: ip: fd:-1 request tlv_len:41
   [iotd] txt(.....y...w.....Jasons-iPhone6....579=23..)
   [iotd] hex
        (02010007017903060f77fc0203000e4a61736f6e732d6950686f6e65360204000835373 93d32330cf
        f)
   [iotd] service:collect hostname:qadevcollect.fortinet.net ip: fd:-1 got server hostname
   [iotd] service:collect hostname:qadevcollect.fortinet.net ip:192.168.100.133 fd:-1 got
        server ip
   [iotd] service:collect hostname:qadevcollect.fortinet.net ip:192.168.100.133 fd:13
        socket created
   [iotd] service:collect hostname:qadevcollect.fortinet.net ip:192.168.100.133 fd:13
        connecting
   [iotd] fd:13 monitor event:pollout
   [iotd] service:collect hostname:qadevcollect.fortinet.net ip:192.168.100.133 fd:13 build
        req packet
   [iotd] service:collect hostname:qadevcollect.fortinet.net ip:192.168.100.133 fd:13
        collect resp:1(pending)
   ```
   The FortiGuard collection server returns new device data to the FortiGuard query server.
   ```
   [iotd] service:query hostname:qadevquery.fortinet.net ip:192.168.100.248 fd:17 got query
        resp
   [iotd] service:query hostname:qadevquery.fortinet.net ip:192.168.100.248 fd:17 id:0
        total_len:48 header_len:16 tlv_len:32 confidence:100 mac:f8:87:f1:1f:ab:95
   [iotd] service:query hostname:qadevquery.fortinet.net ip:192.168.100.248 fd:17
        remaining_len:32 type:1 len:6
   [iotd] service:query hostname:qadevquery.fortinet.net ip:192.168.100.248 fd:17 got tlv
        category:'Mobile'
   [iotd] service:query hostname:qadevquery.fortinet.net ip:192.168.100.248 fd:17
        remaining_len:24 type:2 len:6
   [iotd] service:query hostname:qadevquery.fortinet.net ip:192.168.100.248 fd:17 got tlv
        sub_category:'Mobile'
   [iotd] service:query hostname:qadevquery.fortinet.net ip:192.168.100.248 fd:17
        remaining_len:16 type:3 len:5
   [iotd] service:query hostname:qadevquery.fortinet.net ip:192.168.100.248 fd:17 got tlv
        vendor:'Apple'
   [iotd] service:query hostname:qadevquery.fortinet.net ip:192.168.100.248 fd:17
        remaining_len:9 type:4 len:0
   [iotd] service:query hostname:qadevquery.fortinet.net ip:192.168.100.248 fd:17
        remaining_len:7 type:5 len:3
   [iotd] service:query hostname:qadevquery.fortinet.net ip:192.168.100.248 fd:17 got tlv
        os:'iOS'
   [iotd] service:query hostname:qadevquery.fortinet.net ip:192.168.100.248 fd:17
        remaining_len:2 type:6 len:0
   [iotd] service:query hostname:qadevquery.fortinet.net ip:192.168.100.248 fd:17 send
        query response to caller size:48
   [iotd] txt(............d0 ...Mobile..Mobile..Apple....iOS..)
   [iotd] hex
        (f887f11fab950000000000006430200001064d6f62696c6502064d6f62696c6503054170706c650400
        0503694f530600)
   ```


Netskope Exhibit 1017

```
[iotd] service:query hostname:qadevquery.fortinet.net ip:192.168.100.248 fd:17 read
     resp:0(good)
```

3. The query returns the device information including the information source (`src fortiguard`).

```
diagnose user device list
   vd root/0  f8:87:f1:1f:ab:95  gen 26  req OUA/34
      created 503s  gen 23  seen 102s  lan  gen 7
      ip 192.168.1.110  src arp
      hardware vendor 'Apple'  src fortiguard  id 0  weight 100
      type 'Mobile'  src fortiguard  id 0  weight 100
      family 'Mobile'  src fortiguard  id 0  weight 100
      os 'iOS'  src fortiguard  id 0  weight 100
      host 'Jasons-iPhone6'  src dhcp
```

Feature visibility

Feature visibility is used to control which features are visible in the GUI. This allows features that are not in use to be hidden. Some features are also invisible by default and must be made visible before they can be configure in the GUI.

The visibility of a feature does not affect its functionality or configuration. Invisible features can still be configured using the CLI.

To change the visibility of features:

1. Go to *System* > *Feature Visibility*.
2. Change the visibility of the features as required. To simplify setting security features, a feature set can be selected from the dropdown list.
   For information about what settings each options affects, click on the + icon to the right of the feature name.
   Changes are listed on the right of the screen.
3. Click *Apply*.

Security feature presets

Six system presets are available:

- *NGFW*: for networks that require application control and protection from external attacks.
- *ATP*: for networks that require protection from viruses and other external threats.
- *WF*: for networks that require web filtering.
- *NGFW + ATP*: for networks that require protection from external threats and attacks.
- *UTM*: for networks that require protection from external threats and wish to use security features that control network usage. This is the default setting.
- *Full UTM*: for networks that require the normal UTM features, as well as antivirus, application control, endpoint control, and web filtering.
- *Custom* should be chosen for networks that require customization of available features (including the ability to select all features).

Netskope Exhibit 1017

Certificates

The following topics provide instructions about certificates:

- Microsoft CA deep packet inspection on page 695
- Purchase and import a signed SSL certificate on page 700

Microsoft CA deep packet inspection

In most production environments, you want to use a certificate issued be your own PKI for deep packet inspection (DPI).

An existing Microsoft root CA can be used to issue a subordinate CA (sub CA) certificate that is installed as a DPI certificate on the FortiGate.

Complete the following steps to create your own sub CA certificate and use it for DPI:

1. Create a Microsoft sub CA certificate
2. Export the certificate and private key
3. Import the certificate and private key into the FortiGate
4. Configure a firewall policy for DPI
5. Verify that the sub CA certificate is being used for DPI

The FortiGate firewall uses information in the original web server certificate, then issues a new certificate signed by the Microsoft DPI certificate. The FortiGate then sends this certificate with the issuing DPI certificate to the client's web browser when the SSL session is being established.

The browser verifies that the certificate was issued by a valid CA, then looks for the issuing CA of the Microsoft DPI certificate in its loca trusted root CA store to complete the path to trusted root CA.

The Microsoft CA root certificate is normally deployed to all client PCs in the Windows domain, so the client can complete the certificate path up to a trusted root CA. The FortiGate now controlsand can inspect the two HTTPS sessions: one with the external web server, and one with the client PC.

Create a Microsoft sub CA certificate

A Microsoft sub CA certificate can be created on a Microsoft CA server, or remotely using a web browser.

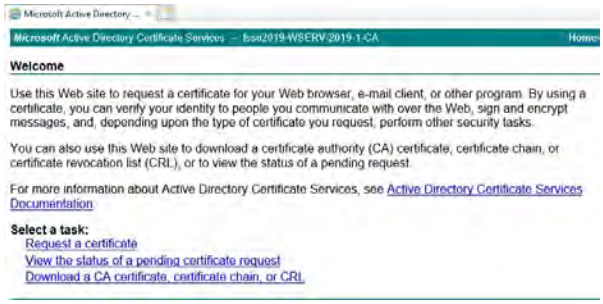
Creating a certificate remotely requires that the web enrollment option is configured on the Microsoft CA server. Remote certificate requests require HTTPS; requests are not allowed with HTTP.

To create a Microsoft sub CA certificate remotely:

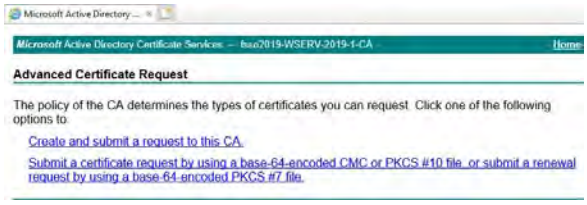
1. Open a web browser and go to one of the following URLs:
   - *https://<FQDN-CA-server>/CertSrv*
   - *https://<IP-CA-server>/CertSrv*.

Netskope Exhibit 1017

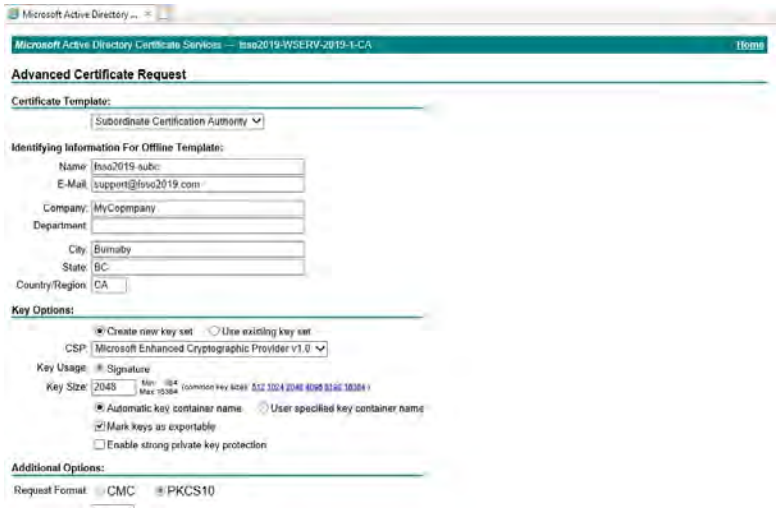
2. Log in to a domain administrator account that has web enrollment rights.



3. Click *Request a certificate*.
4. Click *advanced certificate request*.



5. Click *Create and submit a request to this CA*, then click *Yes* in the *Web Access Confirmation* warning.
6. For the *Certificate Template*, select *Subordinate Certification Authority*.
7. Enable *Mark keys as exportable*.
8. Fill out the remaining information according to your security policy.

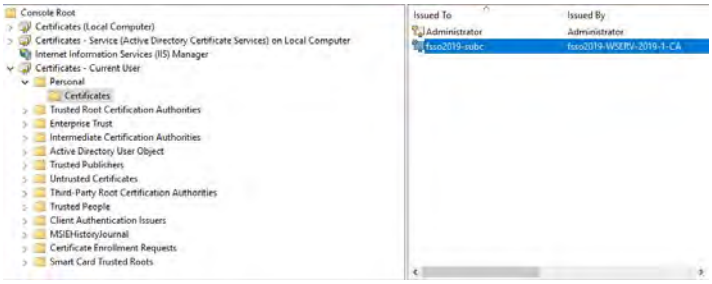


9. Submit the request.
10. Click *Yes* in the *Web Access Confirmation* warning.
11. Click *Install this certificate*.
    The certificate and private key are located in the current user's certificate store.

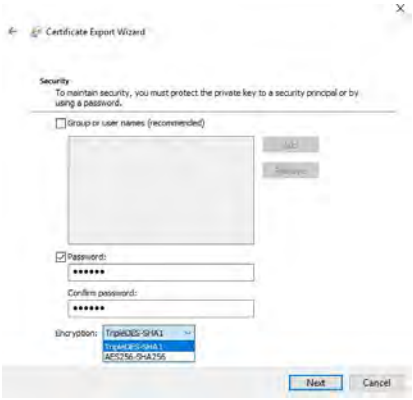
Export the certificate and private key

To export the certificate and private key:

1. Open the Microsoft Management Console (MMC) and add the *Certificate Snap-in*.
2. Go to the user's certificate store to locate the sub CA certificate that you just installed.



3. Right-click on the certificate and select *All Tasks* > *Export*.
4. Click *Next* to start the *Microsoft Certificate Export Wizard*.
5. Follow the steps in the wizard:
   - When asked, select *Yes, export the private key*.
   - Only the PKCS #12 (.PFX) format is available, and it requires a password.
   - When selecting the encryption type, select *TripleDES-SHA1* if you are using an older version of FortiOS (5.6.9 and earlier). Otherwise, select *AES256-SHA256*.



6. Complete the wizard, and save the DPI certificate to a local folder.

Import the certificate and private key into the FortiGate

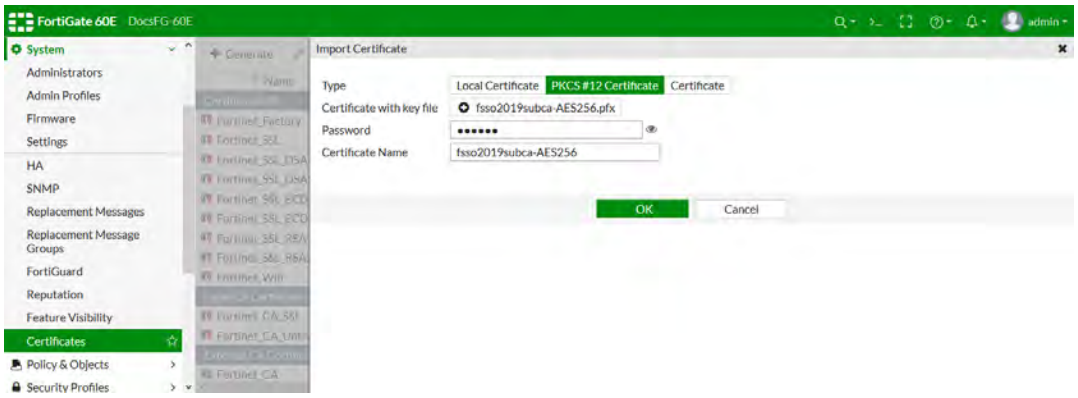
The certificate can be imported from the local computer using the GUI, or from a TFTP server using the CLI.

After importing the certificate, you can view it in the GUI to verify that it was successfully imported.

To import the certificate and private key into the FortiGate in the GUI:

1. Go to *System* > *Certificates*.
2. Select *Import* > *Local Certificate*.
3. Set *Type* to *PKCS #12 Certificate*.
4. Click *Upload* and locate the certificate file.
5. Enter the *Password*.

6. Optionally, modify the *Certificate Name*.



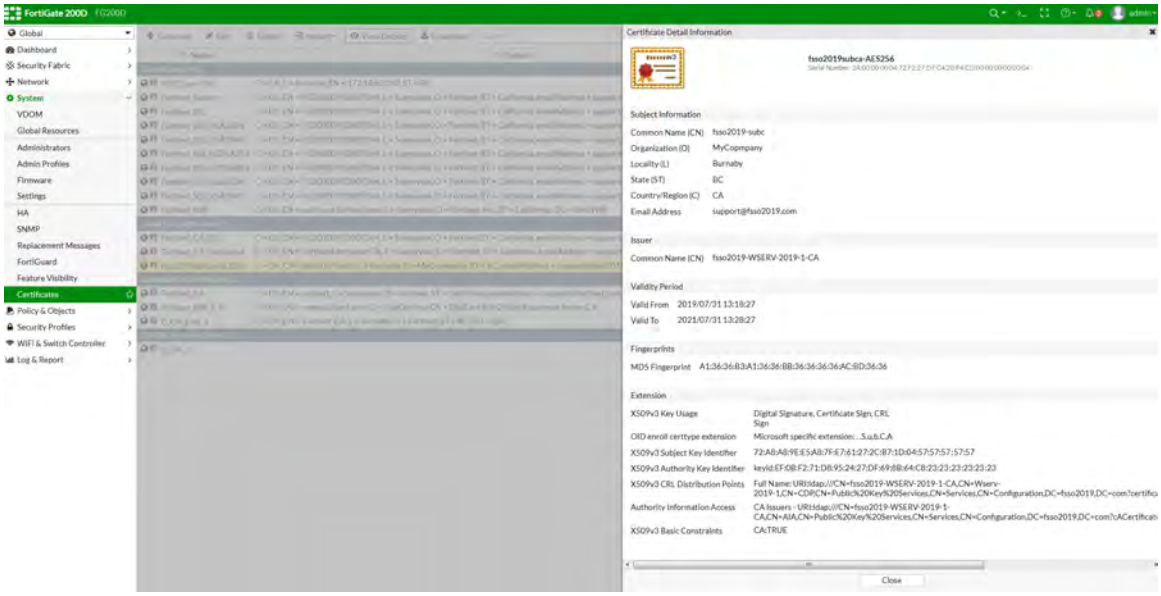
7. Click *OK*.

To import the certificate and private key into the FortiGate in the CLI:

```
execute vpn certificate local import <certificate file name> <tftp ip address> <password>
```

To verify that the certificate was imported:

1. Go to *System* > *Certificates*.
2. Locate the newly imported certificate in the table.
3. Select the certificate and click *View Details* to view the certificate details.

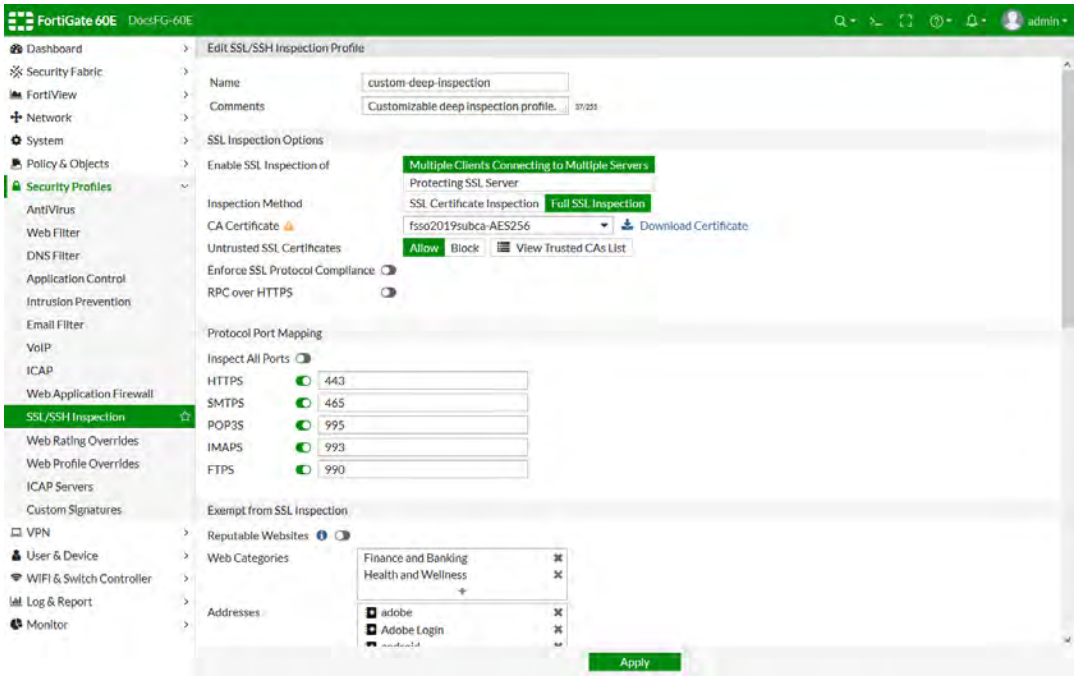


Configure a firewall policy for DPI

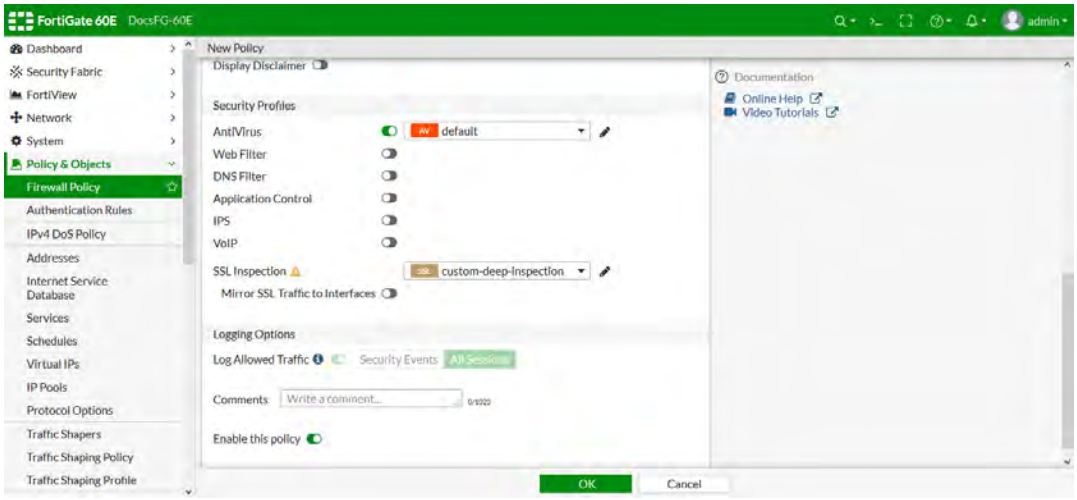
The certificate is used in an SSL/SSH inspection profile that is then used in a firewall policy.

To configure a firewall policy for DPI:

1. Go to *Security Profiles* > *SSL/SSH Inspection*.
2. Click *Create New*.
3. Configure the inspection profile, selecting the new certificate



4. Click *Apply*.
5. Go to *Policy & Objects* > *Firewall Policy*.
6. Create a new policy, or edit an existing policy.
7. In the *SSL Inspection* field, select the new SSL inspection profile.



8. Configure the remaining settings as needed.
9. Click *OK*.

Netskope Exhibit 1017

Verify that the sub CA certificate is being used for DPI

You can verify that the certificate is being used for resigning web server certificates when a user connects to an external HTTPS website.

To verify that the certificate is being used:

1. On a client PC that is behind the FortiGate, go to an external HTTPS website.
   When connecting to the website, no certificate warning should be shown.
2. In your web browser, view the certificate and certificate path.
   The methods for doing this vary depending on the browser. See your browsers documentation for information.

Purchase and import a signed SSL certificate

A signed SSL certificate can be used when configuring SSL VPN, for administrator GUI access, and for other functions that require a certificate.



Before creating a certificate, you must have a registered domain.

Follow these instructions to purchase, import, and use a signed SSL certificate:

- Purchase, setup, and download an SSL certificate package from a certificate authority
- Generate a CSR
- Import the signed certificate into your FortiGate
- Configure your FortiGate device to use the signed certificate

Purchase, setup, and download an SSL certificate package from a certificate authority

SSL certificate packages can be purchased from any Certificate Authority (CA), such as DigiCert, GoDaddy, or GlobalSign.



A third party CA might not sign a certificate with an intranet name or IP address. For details, see Can I request a certificate for an intranet name or IP address?

The process for purchasing, setting up, and downloading a certificate will vary depending on the CA that is used, and if a CSR must be generated on the FortiGate.

To purchase a certificate package:

1. Create an account with your chosen vendor, or use the account that you used to purchase your domain.
2. Locate the SSL Certificates page.
3. Purchase a basic SSL certificate for domain validation only. If required, a more secure SSL certificate can be purchased.

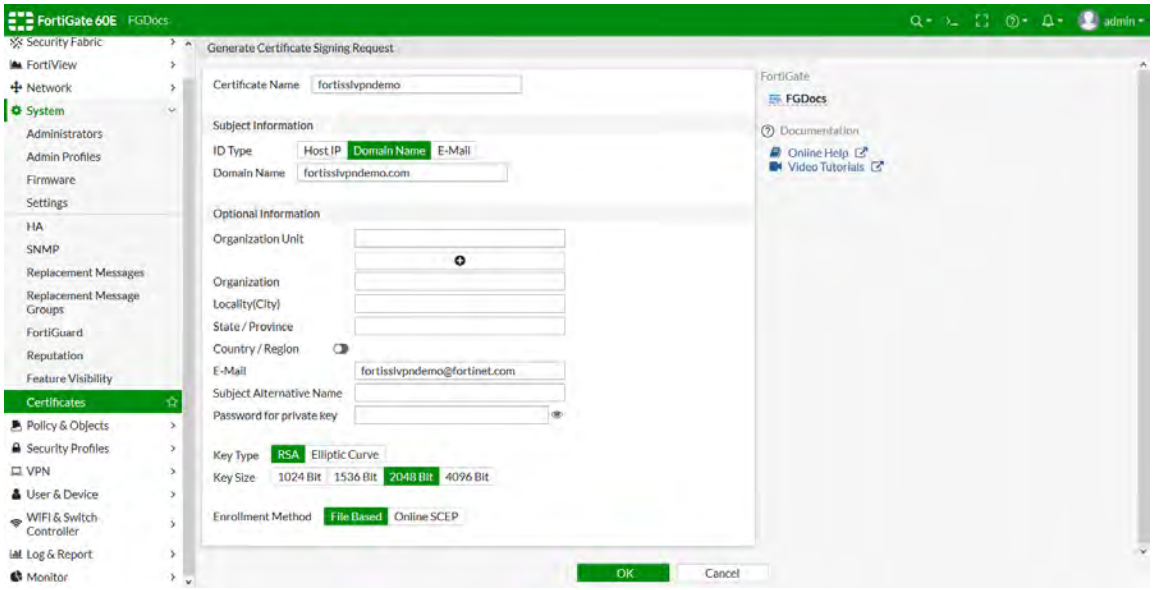
4. If required, load the CSR, either by uploaded the text file or copying and pasting the contents into the requisite text box. See Generate a CSR on page 701 for information on generating the CSR on the FortiGate.
5. If required, set the server type to *Other*.
6. Verify the certificate per the requirements of the CA.
7. Download the signed certificate to your computer.
8. Import the signed certificate into your FortiGate; see Import the signed certificate into your FortiGate on page 702.

Generate a CSR

Some CAs can auto-generate the CSR during the signing process, or provide tools for creating CSRs. If necessary, a CSR can be created in your FortiGate device's GUI.

To generate a CSR on your FortiGate:

1. Go to *System* > *Certificates*. By default, the *Certificate* option is not visible, see Feature visibility on page 694 for information.
2. Click *Generate*. The *Generate Certificate Signing Request* page opens.



3. Configure the CSR request:
   - Ensure that the certificate has a unique name.
   - Set the *ID Type* to *Domain Name* and enter a *Domain Name*.
   - An email address is required.
   - Ensure that the *Key Size* is set to *2048 Bit*.
   - Set the *Enrollment Method* to *File Based*.
4. Click *OK*.
   The CSR will be added to the certificate list with a status of *PENDING*.
5. In the certificate list, select the new CSR then click *Download* to save the CSR to your computer.
   The CSR file can be opened in any text editor, and will resemble the following:

```
-----BEGIN CERTIFICATE REQUEST-----
MIICuTCCAaECAQAwSzEcMBoGA1UEAxMTZm9ydGlzc2x2cG5kZW1vLmNvbTErMCkG
CSqGSIb3DQEJARYcZm9ydGlzc2x2cG5kZW1vQGZvcnRpbmV0LmNvbTCCASIwDQYJ
```

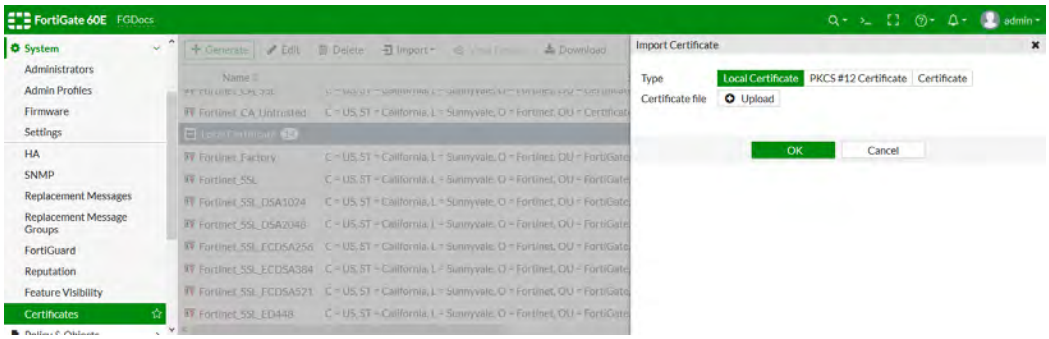

Netskope Exhibit 1017

```
KoZIhvcNAQEBBQADggEPADCCAQoCggEBAMtnpNoR20NH2+UEX/NsyCmZhQqc4af3
Be1u9iOoNbo9Fk42gw47r71moAN+1jTL/Tcp3hRhXtpgoI7Zh3vjZnBbD2wwU8Ow
U7d1h5MULyMehR9r4T6OAJl4KbKPt5u90r5SpIb6mM1OIKvzMncuRS66rW1St0KP
mp/f6QjpjMrthnyJkCejgyTA1YwWNuT9BcO6PTkxBqVMLaRP6TUH6He9uhOx1Cj/
5tzvSdAozZIr2moMieQy0lNd6oQcgpdzaB9QN41+cZOlUXRCMPoH7E4KUe3/Gnis
+NMdQ8rIBijvWCXrKj20wb6sUEjAGJkcXlqVHWYCKWXl6Owejmc4ipkCAwEAAaAp
MCcGCSqGSIb3DQEJDjEaMBgwCQYDVR0TBAIwADALBgNVHQ8EBAMCBaAwDQYJKoZI
hvcNAQELBQADggEBAJKhtz2BPIKeHH9HcJKnfBKL+a6vu1l+1sW+YqnyD+3oR9ec
0eCmLnPxyyxsVel/tRsUg4DTfmooLNDhOjgfMsWxAGUQgrDH2k87cw6kiDAPCqv1
b+hFPNKZQSd09+HXAvOpXrMlrw5YdSaoRnau6Q02yUIYennKTIzFIscgh1mk4FSe
mb12DhPF+QydDCGDgtqnQbfxlDC0WmDcmxwa/0ZktoQhhhEbYgJ2O7l4TMqOxs/q
AZgwJlSNGBALLA2AxkIRUMKUteDdXz0QE8xNrvZpLTbWCNIpYJdRRqSd5C1w2VF4
CFgugTjFaJ13kYmBimeMRQsFtjLV5AxN+bUUsnQ=
-----END CERTIFICATE REQUEST-----
```

Import the signed certificate into your FortiGate

To import the signed certificate into your FortiGate:

1. Unzip the file downloaded from the CA.
   There should be two CRT files: a CA certificate with *bundle* in the file name, and a local certificate.
2. Log in to your FortiGate unit and go to *System* > *Certificates*.
3. Click *Import* > *Local Certificate*.



4. Upload the local certificate file, then click *OK*.
5. The status of the certificate will change from *PENDING* to *OK*.
6. Click *Import* > *CA Certificate*.
7. Set the *Type* to *File*, upload the CA certificate file, then click *OK*.
   The CA certificate will be listed in the *CA Certificates* section of the certificates list.

Configure your FortiGate device to use the signed certificate

After the signed certificates have been imported, you can use it when configuring SSL VPN, for administrator GUI access, and for other functions that require a certificate.

To configure your FortiGate to use the signed certificate for SSL VPN:

1. Go to *VPN* > *SSL-VPN Settings*.
2. Set *Server Certificate* to the new certificate.

3. Configure other settings as needed.
4. Click *Apply*.

For more information on configuring SSL VPN, see SSL VPN on page 1282 and the Setup SSL VPN video in the Fortinet Video Library.

To configure using the certificate for administrator GUI access in the CLI:

```
config system global
    set admin-server-cert fortisslvpndemo
end
```

To change the certificate that is used for administrator GUI access in the GUI:

1. Go to *System* > *Settings*.
2. In the *Administration Settings* section, change *HTTPS server certificate* as needed.
3. Click *Apply*. You will be logged out of FortiOS.

Configuration scripts

Configuration scripts are text files that contain CLI command sequences. They can be created using a text editor or copied from a CLI console, either manually or using the *Record CLI Script* function.

Scripts can be used to run the same task on multiple devices. For example, if your devices use the same security policies, you can enter or record the commands to create those policies in a script, and then run the script on each device. You could also create the policies in the GUI, and then copy and paste the CLI commands from the *CLI Console* using the *show* command.

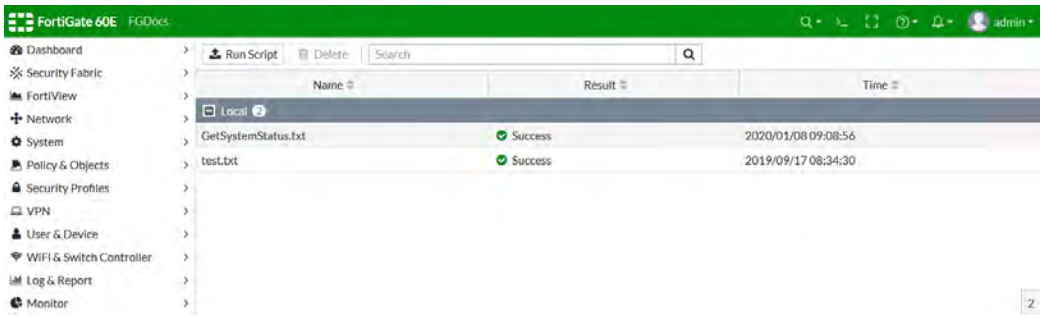
If the FortiGate is managed by FortiManager, scripts can be uploaded to FortiManager and then run on any other FortiGates that are managed by that FortiManager. See Scripts in the FortiManager Administration Guide.



A comment line in a script starts with the number sign (#). Comments are not executed.

To run a script using the GUI:

1. Click on your username and select *Configuration* > *Scripts*.
2. Click *Run Script*.
3. Select the text file containing the script on your management computer, then click *OK*.
   The script runs immediately, and the *Script Execution History* table is updated, showing if the script ran successfully.



Workspace mode

Workspace mode allows administrators to make a batch of changes that are not implemented until the transaction is committed. Prior to committing, the changes can be reverted or edited as needed without impacting current operations.

When an object is edited in workspace mode it is locked, preventing other administrators from editing that object. A warning message will be shown to let the administrator know that the object is currently being configured in another transaction.

All administrators can use workspace mode; their permissions in workspace mode are the same as defined in their account profile.

A workspace mode transaction times out after five minutes if there is no activity. When a transaction times out, all changes are discarded. A warning message will be shown to let the administrator know that a timeout is imminent, or has already happened:

```
config transaction id=1 will expire in 30 seconds
config transaction id=1 will expire in 20 seconds
config transaction id=1 will expire in 10 seconds
config transaction id=1 has expired
```

The following commands are not changeable in a workspace transaction:

```
config system console
config system resource-limits
config system elbc
config system global
   set split-port
   set vdom-admin
   set management-vdom
   set wireless-mode
   set internal-switch-mode
end
config system settings
   set opmode
end
config system npu
config system np6
config system wireless
   set mode
end
config system vdom-property
config system storage
```

The `execute batch` command cannot be used in or to start workspace mode.

To use workspace mode:

1. Start workspace mode:
   ```
   execute config-transaction
   ```
   Once in workspace mode, the administrator can make configuration changes, all of which are made in a local CLI process that is not viewable by other processes.
2. Commit configuration changes:
   ```
   execute config-transaction commit
   ```
   After performing the commit, the changes are available for all other processes, and are also made in the kernel.
3. Abort configuration changes:
   ```
   execute config-transaction abort
   ```
   If changes are aborted, no changes are made to the current configuration or the kernel.

Diagnose commands

```
diagnose sys config-transaction show txn-meta
```

Show config transaction meta information. For example:

```
# diagnose sys config-transaction show txn-meta
txn_next_id=8, txn_nr=2
```

```
diagnose sys config-transaction show txn-info
```

Show config transaction information. For example:

```
# diagnose sys config-transaction show txn-info
current_jiffies=680372

txn_id=6, expire_jiffies=706104, clicmd_fpath='/dev/cmdb/txn/6_EiLl9G.conf'
txn_id=7, expire_jiffies=707427, clicmd_fpath='/dev/cmdb/txn/7_UXK6wY.conf'
```

```
diagnose sys config-transaction show txn-entity
```

Show config transaction entity. For example:

```
# diagnose sys config-transaction show txn-entity
vd='global', cli-node-oid=37(system.vdom), txn_id=7. location: fileid=0, storeid=0,
     pgnr=0, pgidx=0
vd='global', cli-node-oid=46(system.interface), txn_id=7. location: fileid=3,
     storeid=0, pgnr=0, pgidx=0
```

```
diagnose sys config-transaction show txn-lock
```

Show transaction lock status. For example:

```
# diagnose sys config-transaction show txn-lock
type=-1, refcnt=0, value=256, pid=128
```

```
diagnose sys config-transaction status
```

Show the transaction status in the current CLI.

Netskope Exhibit 1017

Policy and Objects

This section contains topics on configuring policies and traffic shaping:

- Policies on page 706
- Objects on page 790
- Traffic shaping on page 814

Policies

The firewall policy is the axis around which most features of the FortiGate revolve. Many firewall settings end up relating to or being associated with the firewall policies and the traffic they govern. Any traffic going through a FortiGate has to be associated with a policy. These policies are essentially discrete compartmentalized sets of instructions that control the traffic flow going through the firewall. These instructions control where the traffic goes, how it is processed, if it is processed, and whether or not it is allowed to pass through the FortiGate.

When the firewall receives a connection packet, it analyzes the source address, destination address, and service (by port number). It also registers the incoming interface, the outgoing interface it needs to use, and the time of day. Using this information, the FortiGate firewall attempts to locate a security policy that matches the packet. If a policy matches the parameters, then the FortiGate takes the required action for that policy. If it is *Accept*, the traffic is allowed to proceed to the next step. If the action is *Deny* or a match cannot be found, the traffic is not allowed to proceed.

The two basic actions at the initial connection are either *Accept* or *Deny*:

- If the action is *Accept*, the policy permits communication sessions. There may be other packet processing instructions, such as requiring authentication to use the policy or restrictions on the source and destination of the traffic.
- If the action is *Deny*, the policy blocks communication sessions, and you can optionally log the denied traffic. If no security policy matches the traffic, the packets are dropped. A *Deny* security policy is needed when it is required to log the denied traffic, also called *violation traffic*.

One other action can be associated with the policy:

- *IPsec*: this is an *Accept* action that is specifically for IPsec VPNs.

The following topics provide instructions on configuring policies:

- Firewall policy parameters on page 707
- Profile-based NGFW vs policy-based NGFW on page 708
- NGFW policy mode application default service on page 712
- Policy views and policy lookup on page 714
- Policy with source NAT on page 715
- Policy with destination NAT on page 728
- Policy with Internet Service on page 741
- NAT64 policy and DNS64 (DNS proxy) on page 758
- NAT46 policy on page 762
- Multicast processing and basic Multicast policy on page 765

Netskope Exhibit 1017

- Local-in policies on page 766
- IPv4/IPv6 access control lists on page 768
- Mirroring SSL traffic in policies on page 769
- Inspection mode per policy on page 772
- OSPFv3 neighbor authentication on page 774
- Firewall anti-replay option per policy on page 776
- Enabling advanced policy options in the GUI on page 777
- Recognize anycast addresses in geo-IP blocking on page 778
- Matching GeoIP by registered and physical location on page 779
- Authentication policy extensions on page 780
- HTTP to HTTPS redirect for load balancing on page 781
- GTPv2 in policies on page 783
- Use active directory objects directly in policies on page 784
- FortiGate Cloud / FDN communication through an explicit proxy on page 788
- No session timeout on page 789

Firewall policy parameters

For traffic to flow through the FortiGate firewall, there must be a policy that matches its parameters:

- Incoming interface(s)
- Outgoing interface(s)
- Source address(es)
- User(s) identity
- Destination address(es)
- Internet service(s)
- Schedule
- Service

Without all six (possibly eight) of these things matching, the traffic is declined.

Traffic flow initiated from each direction requires a policy, that is, if sessions can be initiated from both directions, each direction requires a policy.

Just because packets can go from point A to point B on port X does not mean that the traffic can flow from point B to point A on port X. A policy must be configured for each direction.

When designing a policy, there is often reference to the traffic flow, but most communication is two-way so trying to determine the direction of the flow might be confusing. If traffic is HTTP web traffic, the user sends a request to the website, but most of the traffic flow will be coming from the website to the user or in both directions? For the purposes of determining the direction for a policy, the important factor is the direction of the initiating communication. The user is sending a request to the website, so this is the initial communication; the website is responding so the traffic is from the user's network to the Internet.



FortiOS does not perform a reverse-path check on reply traffic that matches an allowed session based on the IP tuple. The request traffic can be sent on one interface and the reply traffic could return on another interface.

Netskope Exhibit 1017

Profile-based NGFW vs policy-based NGFW

Profile-based next-generation firewall (NGFW) mode is the traditional mode where you create a profile (antivirus, web filter, and so on) and then apply the profile to a policy.

In policy-based NGFW mode, you allow applications and URL categories to be used directly in security policies, without requiring web filter or application control profiles.

In policy-based mode:

- Central NAT is always enabled. If no Central SNAT policy exists, you must create one. See Central SNAT on page 722 for more information.
- Pre-match rules are defined separately from security policies, and define broader rules, such as SSL inspection and user authentication.

If your FortiGate operates in NAT mode, rather than enabling source NAT in individual NGFW policies, go to *Policy & Objects* > *Central SNAT* and add source NAT policies that apply to all matching traffic. In many cases, you may only need one SNAT policy for each interface pair.

The NGFW mode is set per VDOM, and it is only available when the VDOM inspection mode is flow-based. You can operate your entire FortiGate or individual VDOMs in NGFW policy mode.



Switching from profile-based to policy-based mode converts your policies to policy-based. To avoid issues, you could create a new VDOM for the policy-based mode. We recommend backing up your configuration before switching modes. SeeConfiguration backups on page 108 for information.

Enabling policy-based NGFW mode

To enable policy-based NGFW mode without VDOMs in the GUI:

1. Go to *System* > *Settings*.
2. In *NGFW Mode*, select *Policy-based*.
3. Click *Apply*.

To enable policy-based NGFW mode with VDOMs in the GUI:

1. Go to *System* > *VDOM* .
2. Double-click a VDOM to edit the settings.
3. In *NGFW Mode*, select *Policy-based*.
4. Click *OK*.

To enable policy-based NGFW mode without VDOMs in the CLI:

```
config system settings
    set ngfw-mode policy-based
end
```

To enable policy-based NGFW mode with VDOMs in the CLI:

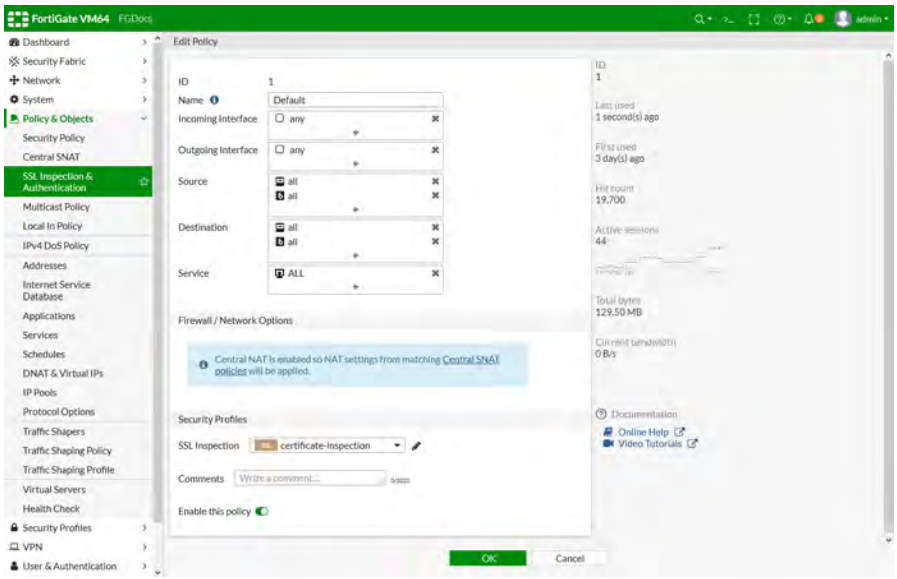
```
config vdom
    edit <vdom>
        config system settings
            set ngfw-mode policy-based
        end
    next
end
```

Security and SSL Inspection & Authentication policies

Security policies work with SSL Inspection & Authentication policies to inspect traffic. To allow traffic from a specific user or user group, both Security and SSL Inspection & Authentication policies must be configured. A default SSL Inspection & Authentication policy with the certificate-inspection SSL Inspection profile is preconfigured. Traffic will match the SSL Inspection & Authentication policy first. If the traffic is allowed, packets are sent to the IPS engine for application, URL category, user, and user group match, and then, if enabled, UTM inspection (antivirus, IPS, DLP, and email filter) is performed.

SSL Inspection & Authentication policies are used to pre-match traffic before sending the packets to the IPS engine:

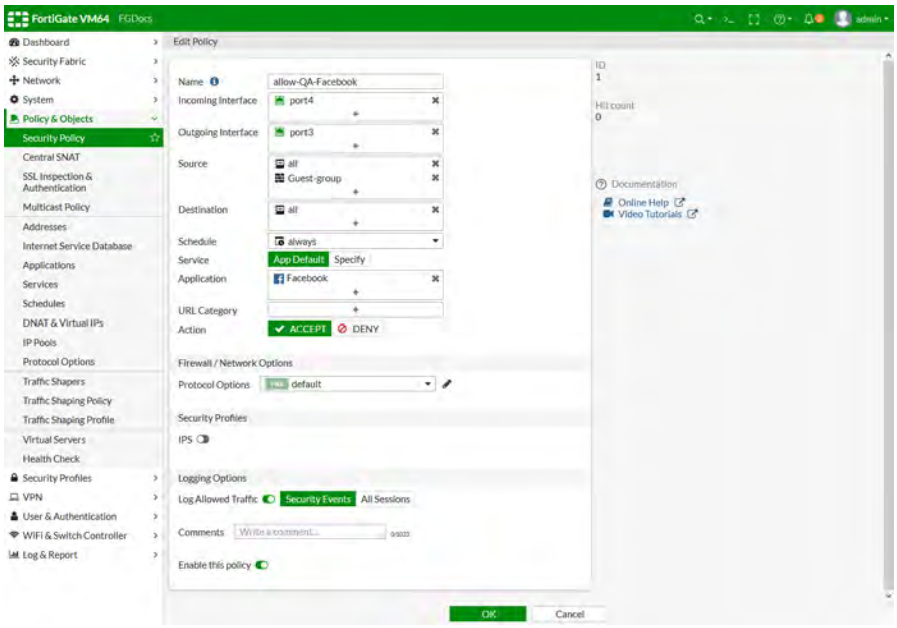
- There are no schedule or action options; traffic matching the policy is always redirected to the IPS engine.
- SSL inspection, formerly configured in the VDOM settings, is configured in an SSL Inspection & Authentication policy.
- Users and user groups that require authentication must be configured in an SSL Inspection & Authentication policy.



Security policies work with SSL Inspection & Authentication policies to inspect traffic:

- Applications and URL categories can be configured directly in the policy.
- Users and user groups that require authentication must also be configured in a security policy.
- The available actions are *Accept* or *Deny*.
- The *Service* option can be used to enforce the standard port for the selected applications. See NGFW policy mode application default service on page 712 for details.
- UTM inspection is configured in a security policy.

Netskope Exhibit 1017



To configure policies for Facebook and Gmail access in the CLI:

1. Configure an SSL Inspection & Authentication policy:

```
config firewall policy
    edit 1
        set name "Policy-1"
        set srcintf "port18"
        set dstintf "port17"
        set srcaddr "all"
        set dstaddr "all"
        set service "ALL"
        set ssl-ssh-profile "new-deep-inspection"
        set groups "Dev" "HR" "QA" "SYS"
    next
end
```

2. Configure security policies:

```
config firewall security-policy
    edit 2
        set name "allow-QA-Facebook"
        set srcintf "port18"
        set dstintf "port17"
        set srcaddr "all"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set application 15832
        set groups "Dev" "QA"
    next
    edit 4
        set name "allow-QA-Email"
        set srcintf "port18"
        set dstintf "port17"
```

```
            set srcaddr "all"
            set dstaddr "all"
            set action accept
            set schedule "always"
            set url-category 23
            set groups "QA"
        next
    end
```

Logs

In the application control and web filter logs, `securityid` maps to the security policy ID.

Application control log:

```
date=2019-06-17 time=16:35:47 logid="1059028704" type="utm" subtype="app-ctrl"
eventtype="signature" level="information" vd="vd1" eventtime=1560814547702405829 tz="-0700"
appid=15832 user="Jack" group="QA" srcip=10.1.100.102 dstip=157.240.3.29 srcport=56572
dstport=443 srcintf="port18" srcintfrole="undefined" dstintf="port17"
dstintfrole="undefined" proto=6 service="P2P" direction="incoming" policyid=1
sessionid=42445 appcat="Social.Media" app="Facebook" action="pass" hostname="external-sea1-
1.xx.fbcdn.net" incidentserialno=1419629662 url="/" securityid=2 msg="Social.Media:
Facebook," apprisk="medium" scertcname="*.facebook.com" scertissuer="DigiCert SHA2 High
Assurance Server CA"
```

Web filter log:

```
date=2019-06-17 time=16:42:41 logid="0317013312" type="utm" subtype="webfilter"
eventtype="ftgd_allow" level="notice" vd="vd1" eventtime=1560814961418114836 tz="-0700"
policyid=4 sessionid=43201 user="Jack" group="QA" srcip=10.1.100.102 srcport=56668
srcintf="port18" srcintfrole="undefined" dstip=172.217.3.165 dstport=443 dstintf="port17"
dstintfrole="undefined" proto=6 service="HTTPS" hostname="mail.google.com"
action="passthrough" reqtype="direct" url="/" sentbyte=709 rcvdbyte=0 direction="outgoing"
msg="URL belongs to an allowed category in policy" method="domain" cat=23 catdesc="Web-based
Email" securityid=4
```

Traffic logs:

```
date=2019-06-17 time=16:35:53 logid="0000000013" type="traffic" subtype="forward"
level="notice" vd="vd1" eventtime=1560814553778525154 tz="-0700" srcip=10.1.100.102
srcport=56572 srcintf="port18" srcintfrole="undefined" dstip=157.240.3.29 dstport=443
dstintf="port17" dstintfrole="undefined" poluuid="b740d418-8ed3-51e9-5a7b-114e99ab6370"
sessionid=42445 proto=6 action="server-rst" user="Jack" group="QA" policyid=1
policytype="consolidated" centralnatid=1 service="HTTPS" dstcountry="United States"
srccountry="Reserved" trandisp="snat" transip=172.16.200.2 transport=56572 duration=6
sentbyte=276 rcvdbyte=745 sentpkt=5 rcvdpkt=11 appid=15832 app="Facebook"
appcat="Social.Media" apprisk="medium" utmaction="allow" countapp=1 utmref=65531-294

2: date=2019-06-17 time=16:47:45 logid="0000000013" type="traffic" subtype="forward"
level="notice" vd="vd1" eventtime=1560815265058557636 tz="-0700" srcip=10.1.100.102
srcport=56668 srcintf="port18" srcintfrole="undefined" dstip=172.217.3.165 dstport=443
dstintf="port17" dstintfrole="undefined" poluuid="b740d418-8ed3-51e9-5a7b-114e99ab6370"
sessionid=43201 proto=6 action="timeout" user="Jack" group="QA" policyid=1
policytype="consolidated" centralnatid=1 service="HTTPS" dstcountry="United States"
srccountry="Reserved" trandisp="snat" transip=172.16.200.2 transport=56668 duration=303
sentbyte=406 rcvdbyte=384 sentpkt=4 rcvdpkt=4 appcat="unscanned" utmaction="allow"
countweb=1 utmref=65531-3486
```


Netskope Exhibit 1017

Other NGFW policy-based mode options

You can combine *Application Control* and *Web Filter* in the same NGFW mode policy. If the policy accepts applications or URL categories, you can also apply *AntiVirus*, *DNS Filter*, *IPS* profiles, and logging options.

NGFW policy mode application default service

In NGFW policy-based mode, the application default service enforces applications running only on their default service port. The applications specified in the policy are monitored, and if traffic is detected from a nonstandard port, it is blocked, and a log entry is recorded with a *port-violation* event type.

If you are not using the default ports, and need to pick specific services, select *Specify* to select the required services.

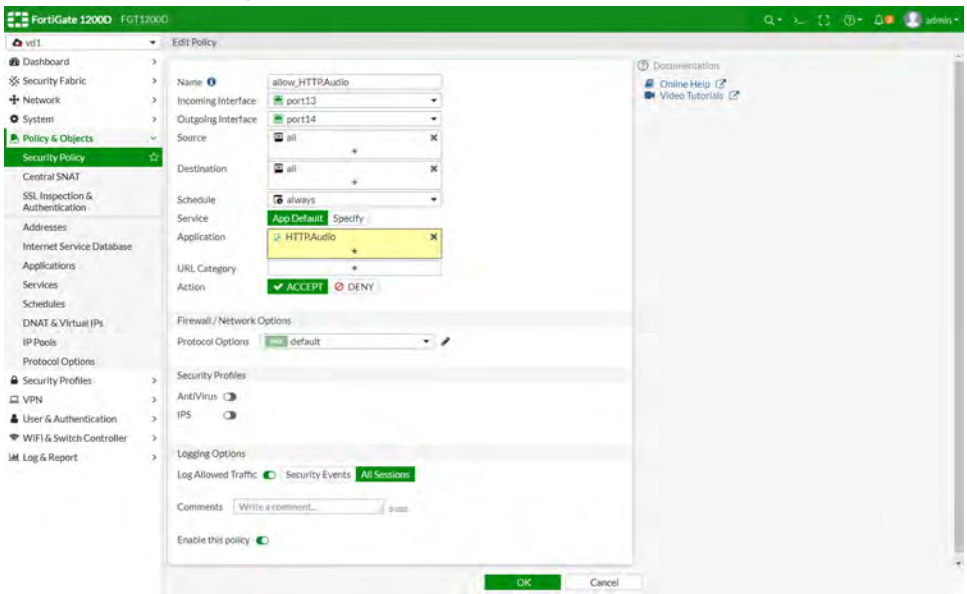
Example

In this example, the standard port is enforced for HTTPS traffic using the HTTP.Audio application.

First, an SSL Inspection & Authentication policy is created do to traffic pre-match, and then a security policy is created to allow the HTTP.Audio application when using the default port. Fetching an MP3 file from an HTTP server using port 443 is allowed, but is blocked when using a nonstandard port, such as 8443.

To enforce the HTTP.Audio application using the default port in the GUI:

1. Create a new SSL Inspection & Authentication policy, or use the default policy.
2. Go to *Policy & Objects* > *Security Policy*, and click *Create New*.
3. Enter a name for the policy, such as *allow_HTTP.Audio*.
4. Configure the ports as needed.
5. Set *Service* to *App Default*.
6. In the *Application* field, select *HTTP.Audio*.
7. Set the *Action* to *Accept*.



8. Click *OK*.

To enforce the HTTP.Audio application using the default port in the CLI:

1. Create a firewall policy:

```
config firewall policy
    edit 1
        set name "consolidated_all"
        set srcintf "port13"
        set dstintf "port14"
        set srcaddr "all"
        set dstaddr "all"
        set service "ALL"
        set ssl-ssh-profile "new-deep-inspection"
    next
end
```

2. Create a security policy:

```
config firewall security-policy
    edit 1
        set name "allow_HTTP.Audio"
        set srcintf "port13"
        set dstintf "port14"
        set srcaddr "all"
        set enforce-default-app-port enable
        set action accept
        set schedule "always"
        set logtraffic all
        set application 15879
    next
end
```

Logs

The application logs show logs with an event type of `port-violation` for traffic on port 8443 that is blocked, and an event type of `signature` for traffic on port 443 that is allowed.

Blocked:

```
2: date=2019-06-18 time=16:15:40 logid="1060028736" type="utm" subtype="app-ctrl"
eventtype="port-violation" level="warning" vd="vd1" eventtime=1560899740218875746 tz="-0700"
appid=15879 srcip=10.1.100.22 dstip=172.16.200.216 srcport=52680 dstport=8443
srcintf="port13" srcintfrole="undefined" dstintf="port14" dstintfrole="undefined" proto=6
service="HTTPS" direction="incoming" policyid=1 sessionid=5041 appcat="Video/Audio"
app="HTTP.Audio" action="block" hostname="172.16.200.216" incidentserialno=1906780850
url="/app_data/story.mp3" securityid=2 msg="Video/Audio: HTTP.Audio," apprisk="elevated"
```

Allowed:

```
1: date=2019-06-18 time=16:15:49 logid="1059028704" type="utm" subtype="app-ctrl"
eventtype="signature" level="information" vd="vd1" eventtime=1560899749258579372 tz="-0700"
appid=15879 srcip=10.1.100.22 dstip=172.16.200.216 srcport=54527 dstport=443
srcintf="port13" srcintfrole="undefined" dstintf="port14" dstintfrole="undefined" proto=6
service="HTTPS" direction="incoming" policyid=1 sessionid=5064 appcat="Video/Audio"
app="HTTP.Audio" action="pass" hostname="172.16.200.216" incidentserialno=1139663486
url="/app_data/story.mp3" securityid=2 msg="Video/Audio: HTTP.Audio," apprisk="elevated"
```


Netskope Exhibit 1017

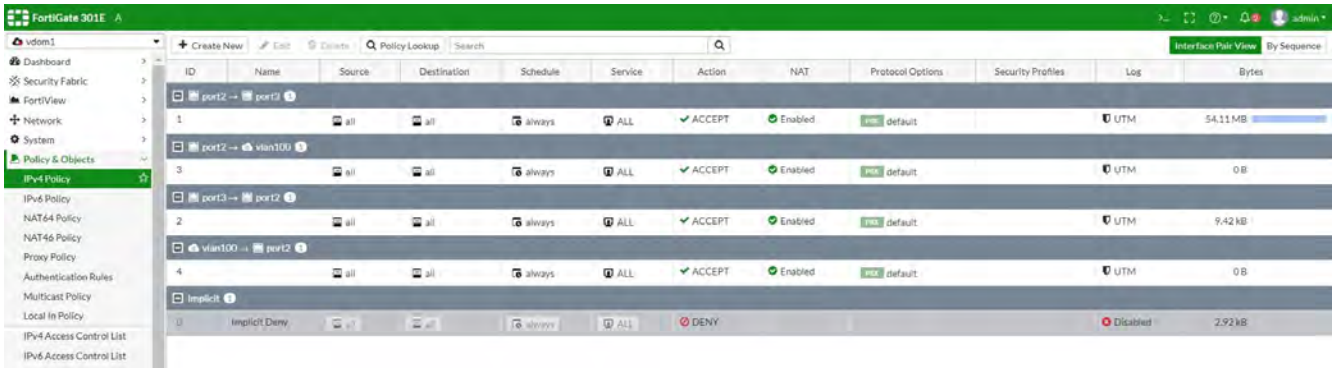
Policy views and policy lookup

This topic provides a sample of firewall policy views and firewall policy lookup.

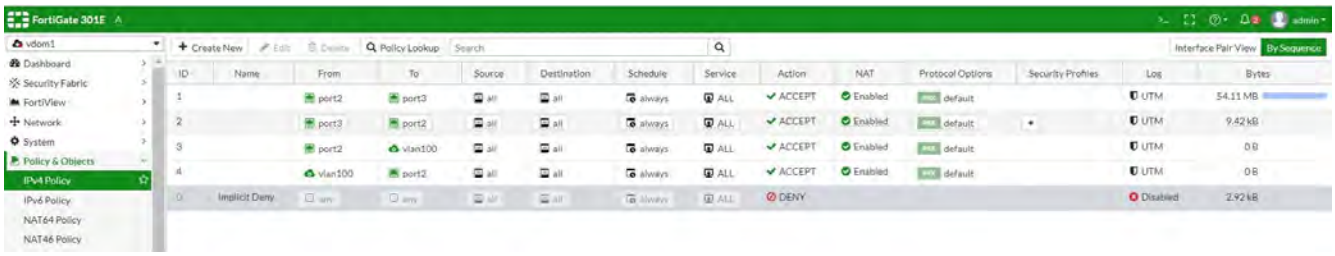
Policy views

In *Policy & Objects* policy list page, there are two policy views: *Interface Pair View* and *By Sequence* view.

Interface Pair View displays the policies in the order that they are checked for matching traffic, grouped by the pairs of Incoming and Outgoing interfaces. For example, all policies referencing traffic from WAN1 to DMZ are in one section. The policies referencing traffic from DMZ to WAN1 are in another section. The sections are collapsible so that you only need to look at the sections you want.



By Sequence displays policies in the order that they are checked for matching traffic without any grouping.



The default display is *Interface Pair View*. You can switch between the two views except if *any* or multiple-interfaces are applied in the policy.

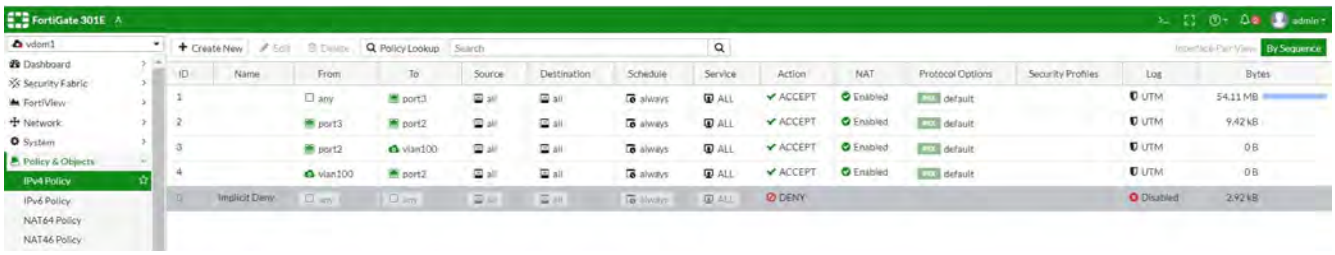
How *Any* or multiple-interfaces policy can change the *Interface Pair View*

The FortiGate unit automatically changes the view on the policy list page to *By Sequence* whenever there is a policy containing *any* or multiple-interfaces as the *Source* or *Destination* interface. If the *Interface Pair View* is grayed out, it is likely that one or more policies have used the *any* or multiple-interfaces.

When you use the *any* or multiple-interfaces, the policy goes into multiple sections because it might be any one of a number of interface pairings. Policies are divided into sectioned using the interface pairings, for example, port1 to port2.

Each section has its own policy order. The order in which a policy is checked for matching criteria to a packet's information is based solely on the position of the policy within its section or within the entire list of policies. If the policy is in multiple sections, FortiGate cannot place the policy in order in multiple sections. Therefore the view can only be *By Sequence*.

Netskope Exhibit 1017



Policy lookup

Firewall policy lookup is based on the `Source_interfaces/Protocol/Source_Address/Destination_Address` that matches the `source-port` and `dst-port` of the protocol. Use this tool to find out which policy matches specific traffic from a number of policies. After completing the lookup, the matching firewall policy is highlighted on the policy list page.

The Policy Lookup tool has the following requirements:

- Transparent mode does not support Policy lookup function.
- When executing the policy lookup, you need to confirm whether the relevant route required for the policy work already exists.

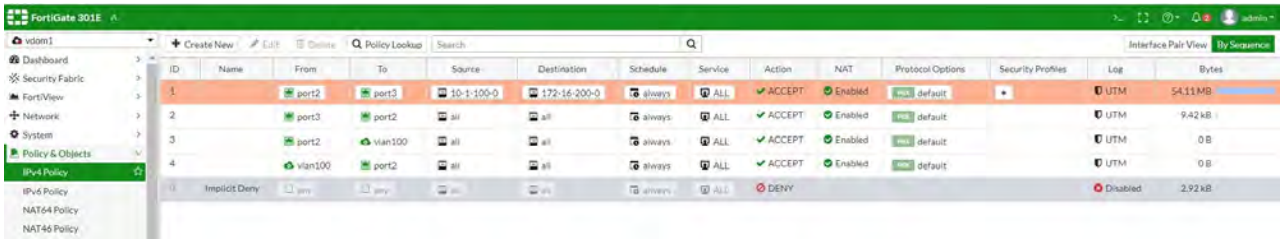
Sample configuration

This example uses the TCP protocol to show how policy lookup works:

1. In *Policy & Objects* policy list page, click *Policy Lookup* and enter the traffic parameters.



2. Click *Search* to display the policy lookup results.



Policy with source NAT

The following topics provide instructions on configuring policies with source NAT:

- Static SNAT on page 716
- Dynamic SNAT on page 717
- Central SNAT on page 722
- Configuring an IPv6 SNAT policy on page 725
- SNAT policies with virtual wire pairs on page 727

Static SNAT

Network Address Translation (NAT) is the process that enables a single device such as a router or firewall to act as an agent between the Internet or Public Network and a local or private network. This agent acts in real time to translate the source or destination IP address of a client or server on the network interface. For the source IP translation, this enables a single public address to represent a significantly larger number of private addresses. For the destination IP translation, the firewall can translate a public destination address to a private address. So we don't have to configure a real public IP address for the server deployed in a private network.

We can subdivide NAT into two types: source NAT (SNAT) and destination NAT (DNAT). This topic is about SNAT, We support three NAT working modes: static SNAT, dynamic SNAT, and central SNAT.

In static SNAT all internal IP addresses are always mapped to the same public IP address. This is a port address translation, Since we have 60416 available port numbers, this one public IP address can handle the conversion of 60,416 internal IP addresses. See example below.

Internal Source IP	Source Port	Translated Source IP	Translated Source Port
10.1.100.1	11110	172.16.200.1	5117
10.1.100.1	11111	172.16.200.1	5118
10.1.100.2	11112	172.16.200.1	5119
.....	.....	172.16.200.1	.....
.....	.....	172.16.200.1	65533

FortiGate firewall configurations commonly use the Outgoing Interface address.

Sample configuration

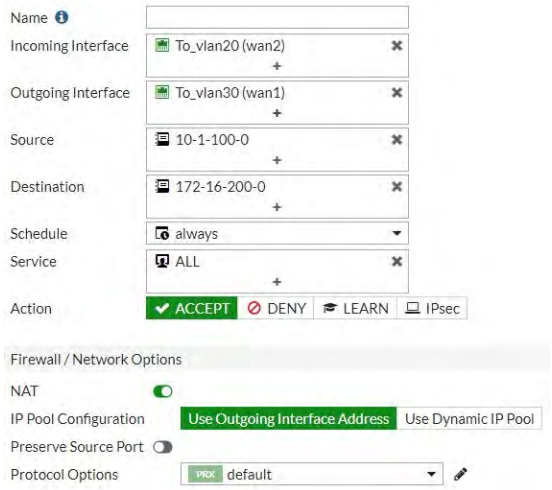
The following example of static SNAT uses an internal network with subnet 10.1.100.0/24 (vlan20) and an external/ISP network with subnet 172.16.200.0/24 (vlan30).

When the clients in internal network need to access the servers in external network, We need to translate IP addresses from 10.1.100.0/24 to an IP address 172.16.200.0/24, In this example, we implement static SNAT by creating a firewall policy.

To configure static NAT:

1. In *Policy & Objects* > *Firewall Policy*, click *Create New*.
2. Enter the required policy parameters.
3. Enable *NAT* and select *Use Outgoing Interface Address*.
4. If needed, enable *Preserve Source Port*.
   Enable *Preserve Source Port* to keep the same source port for services that expect traffic to come from a specific source port.
   Disable *Preserve Source Port* to allow more than one connection through the firewall for that service.

Netskope Exhibit 1017



For packets that match this policy, its source IP address is translated to the IP address of the outgoing interface.

Dynamic SNAT

Dynamic SNAT maps the private IP addresses to the first available public address from a pool of addresses. In the FortiGate firewall, this can be done by using IP pools. IP pools is a mechanism that allows sessions leaving the FortiGate firewall to use NAT. An IP pool defines a single IP address or a range of IP addresses to be used as the source address for the duration of the session. These assigned addresses are used instead of the IP address assigned to that FortiGate interface.

IP pool types

FortiGate uses four types of IPv4 IP pools. This topic focuses on some of the differences between them.

Overload

This type of IP pool is similar to static SNAT mode. We need to define an external IP range that contains one or more IP addresses. When there is only one IP address it is almost the same as static SNAT, the outgoing interface address is used. When it contains multiple IP addresses, it is equivalent to an extended mode of static SNAT.

For instance, if we define an overload type IP pool with two external IP addresses (172.16.200.1—172.16.200.2), since there are 60,416 available port numbers per IP, this IP pool can handle 60,416*2 internal IP addresses.

Original Source IP	Original Source Port	Translated Source IP	Translated Source Port
10.1.100.1	11110	172.16.200.1	5117
10.1.100.2	11111	172.16.200.1	5118
.....	.....	172.16.200.1	.....
.....	.....	172.16.200.1	65533
.....	.....	172.16.200.2	5117
.....	.....	.....	.....
.....	.....	172.16.200.2	65533

The mapped IP address can be calculated from the source IP address. The index number of the address in the pool is the remainder of the source IP address, in decimal, divided by the number addresses in the pool.



To calculate the decimal value of the source IP address, either use an online calculator, or use the following equation:

$$a.b.c.d = a * (256)^3 + b * (256)^2 + c * (256) + d$$

For example:

$$192.168.0.1 = 192 * (256)^3 + 168 * (256)^2 + 0 * (256) + 1 = 3232235521$$

If there is one IP pool, where:

- P_1 = the first address in the IP pool
- R_1 = the number of IP addresses in the IP pool
- X = the source IP address as a decimal number
- Y = the mapped IP address

Then the equation to determine the mapped address is:

$$Y = P_1 + X \bmod R_1$$

For example:

IP pool	Source IP address
172.26.73.20 to 172.26.73.90	192.168.1.200

1. Convert the source IP address to a decimal number:
$$192 * (256)^3 + 168 * (256)^2 + 1 * (256) + 200 = 3232235976$$
2. Determine the number of IP addresses in the pool:
$$172.26.73.90 - 172.26.73.20 = 71$$
3. Find the remainder of the source IP address divided by the number of addresses in the pool:
$$3232235976 \bmod 71 = 26$$
4. Add the remainder to the first IP address in the pool:
$$172.26.73.20 + 26 = 172.26.73.46$$
So, the mapped IP address is **172.26.73.46**.

If there are multiple IP pools, the calculation is similar to when there is only one pool.

If there are two IP pools, where:

- P_1 = the first address in the first IP pool
- P_2 = the first address in the second IP pool
- R_1 = the number of IP addresses in the first IP pool
- R_2 = the number of IP addresses in the second IP pool
- X = the source IP address as a decimal number
- Y = the mapped IP address

Then the equations to determine the mapped address are:

$$\text{If } X \bmod (R_1 + R_2) >= R_1, \text{ then } Y = P_2 + X \bmod R_2$$

$$\text{If } X \bmod (R_1 + R_2) < R_1, \text{ then } Y = P_1 + X \bmod R_1$$

For example:

IP pools	Source IP address
pool01: 172.26.73.20 to 172.26.73.90	192.168.1.200
pool02: 172.26.75.50 to 172.26.75.150	

1. Convert the source IP address to a decimal number:
$$192 * (256)^3 + 168 * (256)^2 + 1 * (256) + 200 = 3232235976$$
2. Determine the total number of IP addresses in the pools:
$$(172.26.73.90 - 172.26.73.20) + (172.26.75.50 - 172.26.75.150) = 71 + 101 = 172$$
3. Find the remainder of the source IP address divided by the number of addresses in the pools:
$$3232235976 \bmod 172 = 108$$
4. The remainder is greater than the number of addresses in pool01, so the address is selected from pool02 and the remainder is recalculated based only on pool02:
$$3232235976 \bmod 101 = 40$$
5. Add the new remainder to the first IP address in pool02:
$$172.26.75.50 + 40 = 172.26.75.90$$
So, the mapped IP address is **172.26.75.90**.

One-to-one

This type of IP pool means that the internal IP address and the external (translated) IP address match one-to-one. The port address translation (PAT) is disabled when using this type of IP pool. For example, if we define a one-to-one type IP pool with two external IP addresses (172.16.200.1 - 172.16.200.2), this IP pool only can handle two internal IP addresses.

Fixed port range

For the overload and one-to-one IP pool types, we do not need to define the internal IP range. For the fixed port range type of IP pool, we can define both internal IP range and external IP range. Since each external IP address and the number of available port numbers is a specific number, if the number of internal IP addresses is also determined, we can calculate the port range for each address translation combination. So we call this type fixed port range. This type of IP pool is a type of port address translation (PAT).

For instance, if we define one external IP address (172.16.200.1) and ten internal IP addresses (10.1.100.1-10.1.100.10), we have translation IP+Port combination like following table:

Original Source IP	Original Source Port	Translated Source IP	Translated Source Port Range
10.1.100.1	.....	172.16.200.1	5117~11157
10.1.100.2	.....	172.16.200.1	11158~17198
10.1.100.3	.....	172.16.200.1	.....
10.1.100.4	.....	172.16.200.1	.....
10.1.100.5	.....	172.16.200.1	.....
10.1.100.6	.....	172.16.200.1	.....
10.1.100.7	.....	172.16.200.1	.....
10.1.100.8	.....	172.16.200.1	.....
10.1.100.9	.....	172.16.200.1	53445~59485
10.1.100.10	.....	172.16.200.1	59486~65526

Netskope Exhibit 1017

Port block allocation

This type of IP pool is also a type of port address translation (PAT). It gives users a more flexible way to control the way external IPs and ports are allocated. Users need to define *Block Size/Block Per User* and external IP range. *Block Size* means how many ports each Block contains. *Block per User* means how many blocks each user (internal IP) can use.

The following is a simple example:

- **External IP Range**: 172.16.200.1—172.16.200.1
- **Block Size**: 128
- **Block Per User**: 8

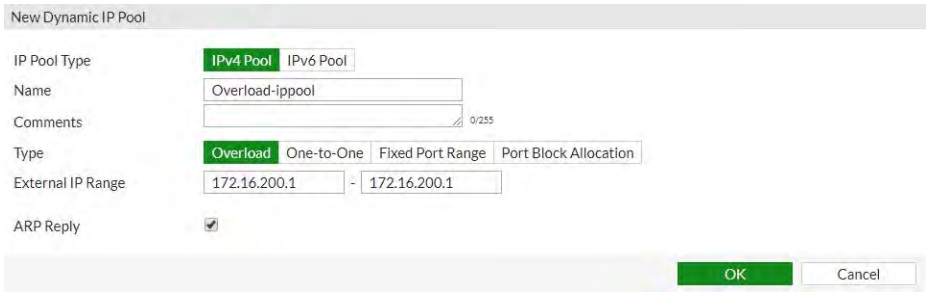
Result:

- **Total-PBAs**: 472 (60416/128)
- **Maximum ports can be used per User (Internal IP Address)**: 1024 (128*8)
- **How many Internal IP can be handled**: 59 (60416/1024 or 472/8)

Sample configuration

To configure overload IP pool in the GUI:

1. In *Policy & Objects* > *IP Pools*, click *Create New*.
2. Select *IPv4 Pool* and then select *Overload*.

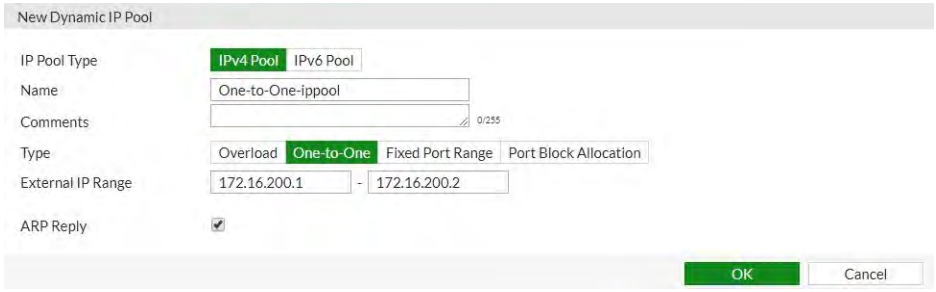


To configure overload IP pool in the CLI:

```
config firewall ippool
    edit "Overload-ippool"
        set startip 172.16.200.1
        set endip 172.16.200.1
    next
end
```

To configure one-to-one IP pool using the GUI:

1. In *Policy & Objects* > *IP Pools*, click *Create New*.
2. Select *IPv4 Pool* and then select *One-to-One*.

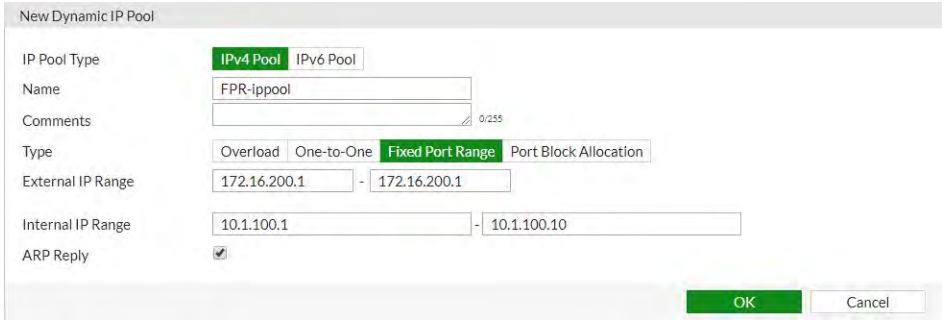


To configure one-to-one IP pool in the CLI:

```
config firewall ippool
    edit "One-to-One-ippool"
        set type one-to-one
        set startip 172.16.200.1
        set endip 172.16.200.2
    next
end
```

To configure fixed port range IP pool in the GUI:

1. In *Policy & Objects* > *IP Pools*, click *Create New*.
2. Select *IPv4 Pool* and then select *Fixed Port Range*.

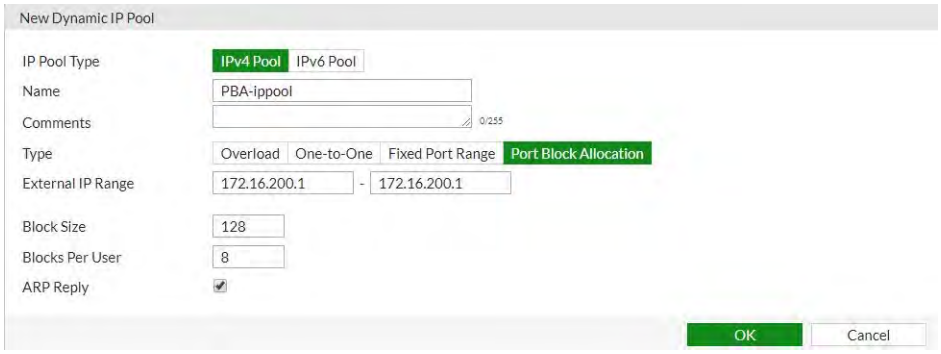


To configure fixed port range IP pool in the CLI:

```
config firewall ippool
    edit "FPR-ippool"
        set type fixed-port-range
        set startip 172.16.200.1
        set endip 172.16.200.1
        set source-startip 10.1.100.1
        set source-endip 10.1.100.10
   next
end
```

To configure port block allocation IP pool in the GUI:

1. In *Policy & Objects* > *IP Pools*, click *Create New*.
2. Select *IPv4 Pool* and then select *Port Block Allocation*.



To configure port block allocation IP pool in the CLI:

```
config firewall ippool
    edit PBA-ippool
        set type port-block-allocation
        set startip 172.16.200.1
        set endip 172.16.200.1
        set block-size 128
        set num-blocks-per-user 8
    next
end
```

Central SNAT

The central SNAT table enables you to define and control (with more granularity) the address translation performed by FortiGate. With the NAT table, you can define the rules for the source address or address group, and which IP pool the destination address uses.

While similar in functionality to IP pools where a single address is translated to an alternate address from a range of IP addresses, with IP pools there is no control over the translated port. When using the IP pool for source NAT, you can define a fixed port to ensure the source port number is unchanged. If no fixed port is defined, the port translation is randomly chosen by FortiGate. With the central NAT table, you have full control over both the IP address and port translation.

FortiGate reads the NAT rules from the top down until it hits a matching rule for the incoming address. This enables you to create multiple NAT policies that dictate which IP pool is used based on the source address. NAT policies can be rearranged within the policy list. NAT policies are applied to network traffic after a security policy.

The central SNAT table allows you to create, edit, delete, and clone central SNAT entries.

Central SNAT notes

- The central NAT feature in not enabled by default.
- If central NAT is enabled, the NAT option under IPv4 policies is skipped and SNAT must be done via `central-snat-map`. The firewall policy list and dialog boxes have messages and redirection links to show this information.
- If NGFW mode is policy-based, then it is assumed that central NAT (specifically SNAT) is enabled implicitly.

Netskope Exhibit 1017

- The option to toggle NAT in `central-snat-map` policies has been added. Previously it was only shown in NGFW policy-based mode.
- In the central SNAT policy dialog box, the port mapping fields for the original port have been updated to accept ranges.
- If per VDOM NAT is enabled, NAT is skipped in firewall policy.
- The central SNAT window contains a table of all the central SNAT policies.

Sample configuration

To enable or disable central SNAT using the CLI:

```
config system settings
   set central-nat [enable | disable]
end
```

When central NAT is enabled, *Policy & Objects* displays the Central SNAT section.

To create central SNAT using the GUI:

1. In *Policy & Objects* > *Central SNAT*.
   The right pane displays a table of Central SNAT entries.
2. To create a new entry, click *Create New* in the right pane.
   To edit an entry, double-click the policy you want to edit.
3. To set the *Incoming Interface*, click + in that field.
4. In the pane on the right, select an interface to add it.
   You can select multiple interfaces.
5. To set the *Outgoing Interface*, click click + in that field.
6. In the pane on the right, select an interface to add it.
   You can select multiple interfaces.
7. To set the *Source Address*, click click + in that field.
8. In the pane on the right, select an address to add it.
   You can select multiple addresses.
9. To set the *Destination Address*, click click + in that field.
10. In the pane on the right, select an address to add it.
    You can select multiple addresses.
11. In *NAT* > *IP Pool Configuration*, select either *Use Outgoing Interface Address* or *Use Dynamic IP Pool*.
    If you select *Use Dynamic IP Pool*, click + and select which IP pool to use.
12. Select one of the following *Protocol* parameters.
    - *ANY*. Use any protocol traffic.
    - *TCP*. Use TCP traffic only. Protocol number is set to *6*.
    - *UDP*. Use UDP traffic only. Protocol number is set to *17*.
    - *SCTP*. Use SCTP traffic only. Protocol number is set to *132*.
    - *Specify*. You can specify the traffic filter protocol by setting the protocol number.
13. If you use the *Overload* type of IP pool, you can enable *Explicit Port Mapping*.
    a. If you enable *Explicit Port Mapping*, set the *Original Source Port* to the start number of the source port range.
    b. Set the *Translated Port* to the start number of the translated port range.
14. Click *OK*.

Netskope Exhibit 1017

To configure central SNAT using the CLI:

```
config firewall central-snat-map
edit <policyID number>set status [enable|disable]
  set orig-addr <valid address object preconfigured on the FortiGate>
  set srcintf <name of interface on the FortiGate>
  set dst-addr <valid address object preconfigured on the FortiGate>
  set dstintf <name of interface on the FortiGate>
  set protocol <integer for protocol number>
  set orig-port <integer for original port number>
  set nat-port <integer for translated port number>
  set comments <string>
end
```

To set NAT to be not available regardless of NGFW mode:

```
config firewall central-snat-map
    edit 1
        set orig-addr "192-86-1-86"
        set srcintf "port23"
        set dst-addr "192-96-1-96"
        set dstintf "port22"
        set nat-ippool "pool1"
        set protocol 17
        set orig-port 2896-2897
        set nat enable
    next
end
```

To hide NAT port if NAT IP pool is not set or if NAT is disabled:

```
config firewall central-snat-map
    edit 1
        set orig-addr "192-86-1-86"
        set srcintf "port23"
        set dst-addr "192-96-1-96"
        set dstintf "port22"
        set nat-ippool "pool1"
        set protocol 17
        set orig-port 2896-2897
        set nat disable
    next
end
```

To change original port to accept range:

```
config firewall central-snat-map
    edit 1
        set orig-addr "192-86-1-86"
        set srcintf "port23"
        set dst-addr "192-96-1-96"
        set dstintf "port22"
        set nat-ippool "pool1"
        set protocol 17
        set orig-port 2896-2897 (help text changed to: Original port or port range).
```


Netskope Exhibit 1017

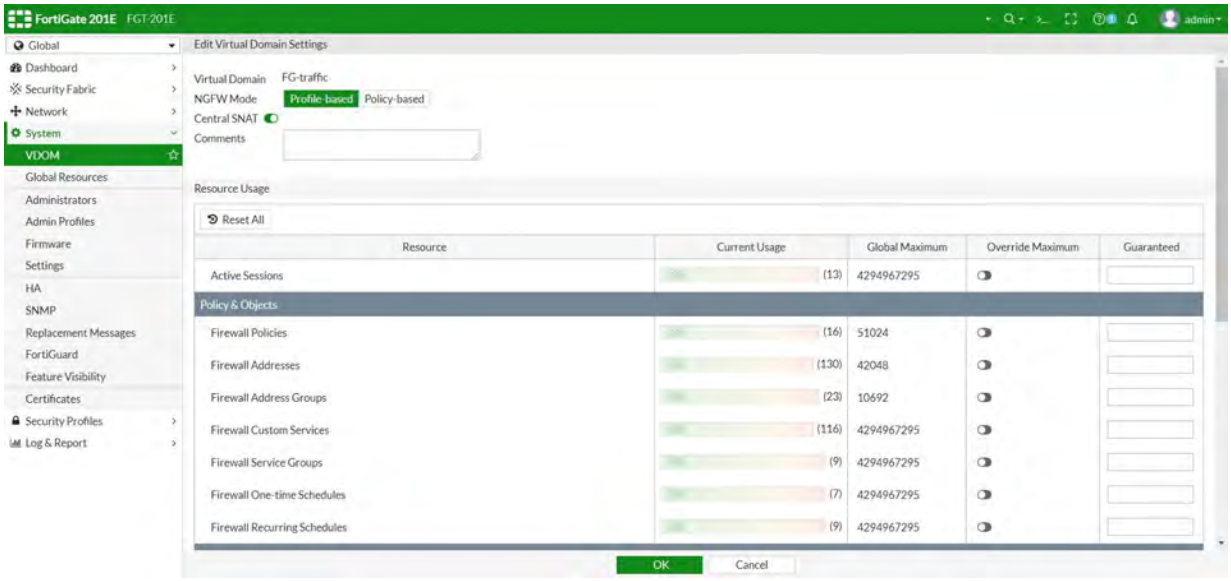
```
        set nat-port 35804-35805
    next
end
```

Configuring an IPv6 SNAT policy

IPv4 and IPv6 central SNAT maps are displayed in the same table.

To configure an IPv6 policy with central SNAT in the GUI:

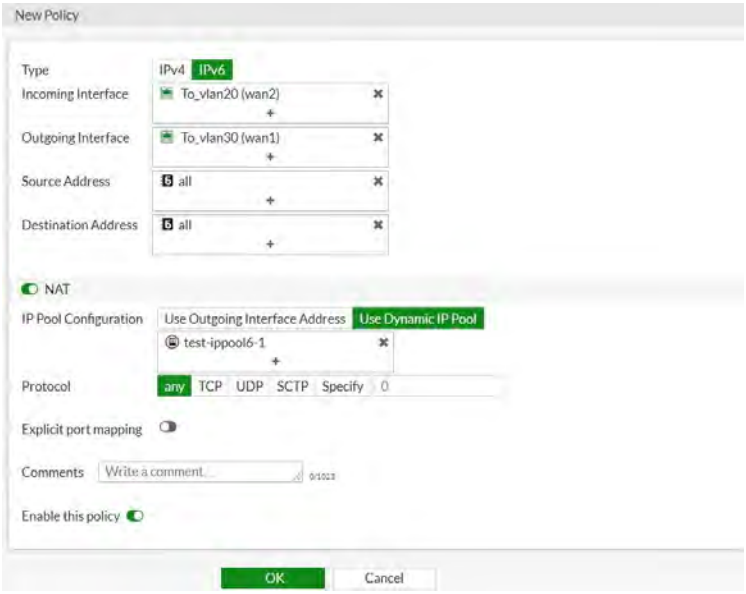
1. Enable central SNAT:
   a. In the Global VDOM, go to *System* > *VDOM*.
   b. Select a VDOM and click *Edit*. The *Edit Virtual Domain Settings* pane opens.
   c. Enable *Central SNAT*.
   d. Click *OK*.



2. Go in to the VDOM with central SNAT enabled (FG-traffic in this example).
3. Go *Policy & Objects* > *Central SNAT* and click *Create New*.
4. Configure the policy settings:
   a. For *Type*, select *IPv6*.
   b. Enter the interface, address, and IP pool information.
   c. Configure the other settings as needed.

Netskope Exhibit 1017

d. Click *OK*.



The matching SNAT traffic will be handled by the IPv6 central SNAT map.

To configure an IPv6 policy with central SNAT in the CLI:

1. Enable central SNAT:

```
config vdom
    edit FG-traffic
        config system settings
            set central-nat enable
        end
    next
end
```

2. Create an IPv6 central SNAT policy:

```
config vdom
    edit FG-traffic
        config firewall central-snat-map
            edit 2
                set type ipv6
                set srcintf "wan2"
                set dstintf "wan1"
                set orig-addr6 "all"
                set dst-addr6 "all"
                set nat-ippool6 "test-ippool6-1"
            next
        end
    next
end
```

3. Verify the SNAT traffic:

```
(FG-traffic) # diagnose sniffer packet any icmp6 4
interfaces=[any]
filters=[icmp6]
```

```
3.602891 wan2 in 2000:10:1:100::41 -> 2000:172:16:200::55: icmp6: echo request seq 0
3.602942 wan1 out 2000:172:16:200::199 -> 2000:172:16:200::55: icmp6: echo request seq 0
3.603236 wan1 in 2000:172:16:200::55 -> 2000:172:16:200::199: icmp6: echo reply seq 0
3.603249 wan2 out 2000:172:16:200::55 -> 2000:10:1:100::41: icmp6: echo reply seq 0
4.602559 wan2 in 2000:10:1:100::41 -> 2000:172:16:200::55: icmp6: echo request seq 1
4.602575 wan1 out 2000:172:16:200::199 -> 2000:172:16:200::55: icmp6: echo request seq 1
4.602956 wan1 in 2000:172:16:200::55 -> 2000:172:16:200::199: icmp6: echo reply seq 1
4.602964 wan2 out 2000:172:16:200::55 -> 2000:10:1:100::41: icmp6: echo reply seq 1
^C
8 packets received by filter
0 packets dropped by kernel
```

SNAT policies with virtual wire pairs

Source NAT (SNAT) can be configured in IPv4 and IPv6 policies with virtual wire pair (VWP) interfaces.

To configure a policy using SNAT and a virtual wire pair:

1. Create the virtual wire pair interface:

```
config system virtual-wire-pair
    edit "test-vw-1"
        set member "port1" "port4"
    next
end
```

2. Create the IP pool. The IP pool must have a different subnet than the VWP peers:

```
config firewall ippool
    edit "vwp-pool-1"
        set startip 172.16.222.99
        set endip 172.16.222.100
    next
end
```

3. Configure the policy:

```
config firewall policy
    edit 88
        set srcintf "port4"
        set dstintf "port1"
        set srcaddr "all"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ALL"
        set logtraffic all
        set ippool enable
        set poolname "vwp-pool-1"
        set nat enable
    next
end
```

4. Verify the IP pool functions as expected and traffic passes through:

```
# diagnose sniffer packet any icmp 4
interfaces=[any]
```


Netskope Exhibit 1017

```
filters=[icmp]
23.438095 port4 in 172.16.200.11 -> 172.16.200.156: icmp: echo request
23.438126 port1 out 172.16.222.100 -> 172.16.200.156: icmp: echo request
23.438492 port1 in 172.16.200.156 -> 172.16.222.100: icmp: echo reply
23.438501 port4 out 172.16.200.156 -> 172.16.200.11: icmp: echo reply
24.439305 port4 in 172.16.200.11 -> 172.16.200.156: icmp: echo request
24.439319 port1 out 172.16.222.100 -> 172.16.200.156: icmp: echo request
24.439684 port1 in 172.16.200.156 -> 172.16.222.100: icmp: echo reply
24.439692 port4 out 172.16.200.156 -> 172.16.200.11: icmp: echo reply

8 packets received by filter
0 packets dropped by kernel
```

Policy with destination NAT

The following topics provide instructions on configuring policies with destination NAT:

- Static virtual IPs on page 728
- Virtual IP with services on page 729
- Virtual IPs with port forwarding on page 731
- Virtual server on page 733

Static virtual IPs

Mapping a specific IP address to another specific IP address is usually called Destination NAT (DNAT). When this central NAT table is not used, FortiOS calls this a Virtual IP address (VIP). DNAT, or VIP, is are used to map an external IP address to an IP address or address range. The mapping can include all TCP/UDP ports or, if port forwarding is enabled, it only refers to the specific configured ports. As the central NAT table is disabled by default, the term VIP is usually used.

VIPs are typically used to NAT external or public IP addresses to internal or private IP addresses. Using a VIP between two internal interfaces made up of private IP addresses is possible, but rare, because the two networks can just use the IP addresses of the networks without any address translation. Using a VIP for traffic going from the inside to the internet is supported, but unlikely to be required.

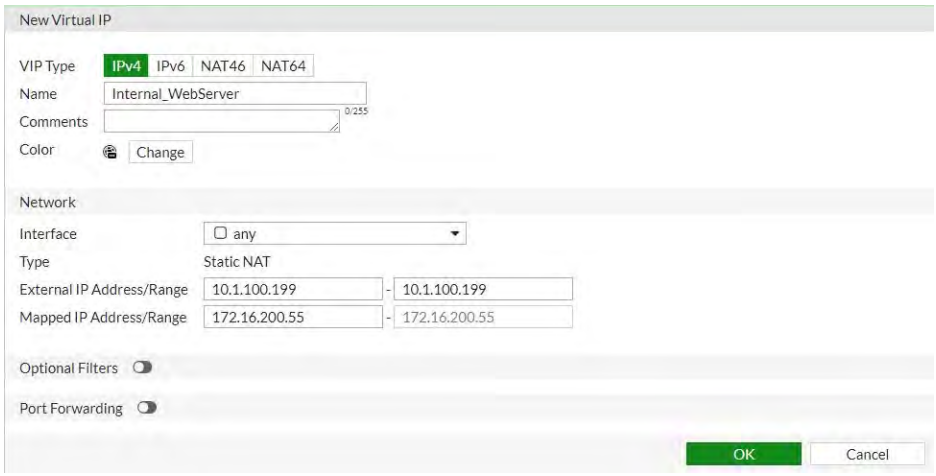
Sample configuration

To create a virtual IP using the GUI:

1. In *Policy & Objects* > *Virtual IPs*.
2. Click *Create New* and select *Virtual IP*.
3. Select a *VIP Type* based on the IP versions used:
   - If IPv4 is on both sides of the FortiGate unit, select *IPv4*.
   - If IPv6 is on both sides of the FortiGate unit, select *IPv6*.
   - If traffic goes from an IPv4 network to an IPv6 network, select *NAT46*.
   - If traffic goes from an IPv6 network to an IPv4 network, select *NAT64*.

Netskope Exhibit 1017

4. Enter a unique name for the virtual IP and fill in the other fields.



To create a virtual IP using the CLI:

```
config firewall vip
    edit "Internal_WebServer"
        set extip 10.1.100.199
        set extintf "any"
        set mappedip "172.16.200.55"
    next
end
```

To apply a virtual IP to policy using the CLI:

```
config firewall policy
    edit 8
        set name "Example_Virtual_IP_in_Policy"
        set srcintf "wan2"
        set dstintf "wan1"
        set srcaddr "all"
        set dstaddr "Internal_WebServer"
        set action accept
        set schedule "always"
        set service "ALL"
        set nat enable
    next
end
```

Virtual IP with services

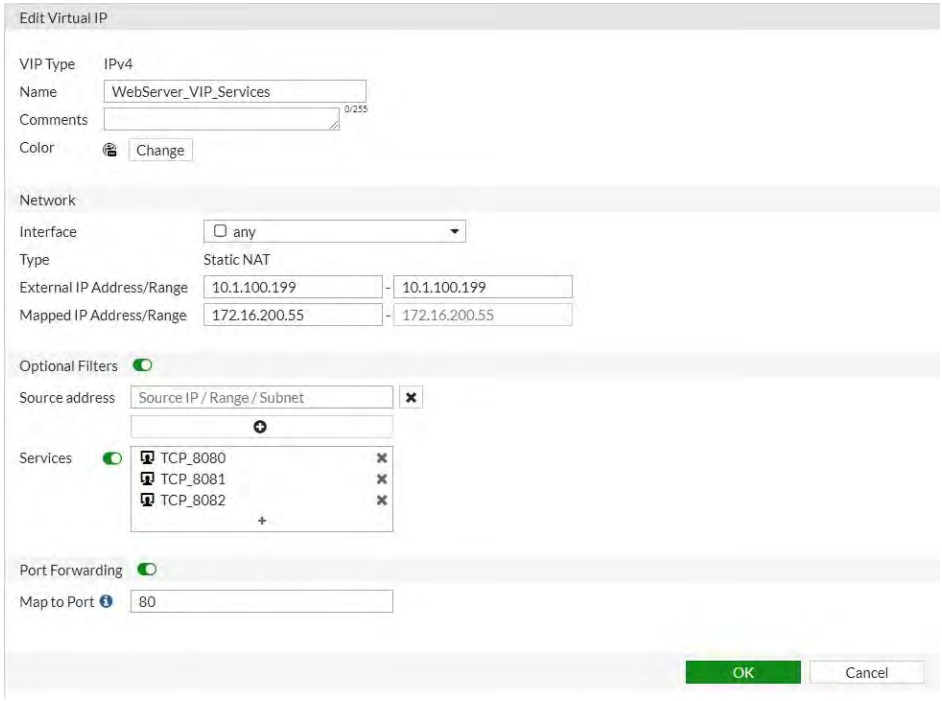
Virtual IP with services is a more flexible virtual IP mode. This mode allows users to define services to a single port number mapping.

This topic shows how to use virtual IP with services enabled. This example has one public external IP address. We map TCP ports 8080, 8081, and 8082 to an internal WebServer TCP port 80. This allows remote connections to communicate with a server behind the firewall.

Sample configuration

To create a virtual IP with services using the GUI:

1. In *Policy & Objects* > *Virtual IPs*.
2. Click *Create New* and select *Virtual IP*.
3. For *VIP Type*, select *IPv4*.
4. Enter a unique name for the virtual IP and fill in the other fields.
5. Configure the fields in the *Network* section. For example:
   - Set *Interface* to *any*.
   - Set *External IP Address/Range* to *10.1.100.199*.
   - Set *Mapped IP Address/Range* to *172.16.200.55*.
6. Enable *Optional Filters* and then enable *Services*.
7. In the *Services* field, click + to display the Services pane.
8. In the *Services* pane, select *TCP_8080*, *TCP_8081*, and *TCP_8082*.
9. Enable *Port Forwarding*.
10. Set *Map to Port* to *80*.



11. Click *OK*.

To see the results:

1. Apply the above virtual IP to the Firewall policy.
2. The results are:
   - Access 10.1.100.199:8080 from external network and FortiGate maps to 172.16.200.55:80 in internal network.
   - Access 10.1.100.199:8081 from external network and FortiGate maps to 172.16.200.55:80 in internal network.
   - Access 10.1.100.199:8082 from external network and FortiGate maps to 172.16.200.55:80 in internal network.

Netskope Exhibit 1017

To create a virtual IP with services using the CLI:

```
config firewall vip
    edit "WebServer_VIP_Services"
        set service "TCP_8080" "TCP_8081" "TCP_8082"
        set extip 10.1.100.199
        set extintf "any"
        set portforward enable
        set mappedip "172.16.200.55"
        set mappedport 80
    next
end
```

Virtual IPs with port forwarding

If you need to hide the internal server port number or need to map several internal servers to the same public IP address, enable port-forwarding for Virtual IP.

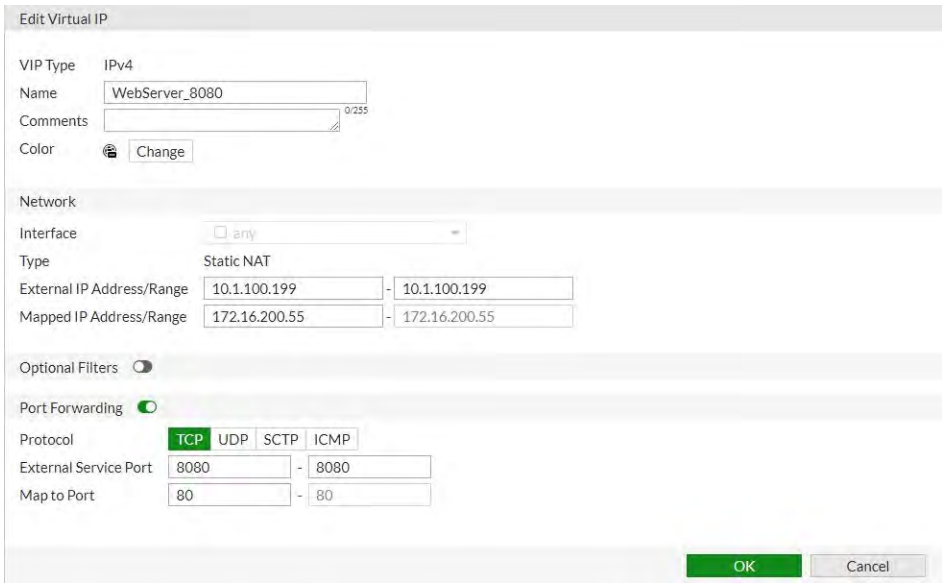
This topic shows how to use virtual IPs to configure port forwarding on a FortiGate unit. This example has one public external IP address. We map TCP ports 8080, 8081, and 8082 to different internal WebServers' TCP port 80. This allows remote connections to communicate with a server behind the firewall.

Sample configuration

To create a virtual IP with port forwarding using the GUI:

1. In *Policy & Objects* > *Virtual IPs*.
2. Click *Create New* and select *Virtual IP*.
3. For *VIP Type*, select *IPv4*.
4. Enter a unique name for the virtual IP and fill in the other fields.
5. Configure the fields in the *Network* section. For example:
   - Set *Interface* to *any*.
   - Set *External IP Address/Range* to *10.1.100.199*.
   - Set *Mapped IP Address/Range* to *172.16.200.55*.
6. Leave *Optional Filters* disabled.
7. Enable *Port Forwarding*.
8. Configure the fields in the *Port Forwarding* section. For example:
   - Set *Protocol* to *TCP*.
   - Set *External Service Port* to *8080 - 8080*.
   - Set *Map to Port* to *80 - 80*.

Netskope Exhibit 1017



9. Click *OK*.
10. Follow the above steps to create two additional virtual IPs.
    a. For one virtual IP:
        - Use a different *Mapped IP Address/Range*, for example, *172.16.200.56*.
        - Set *External Service Port* to *8081 - 8081*.
        - Use the same *Map to Port* numbers: *80 - 80*.
    b. For the other virtual IP:
        - Use a different *Mapped IP Address/Range*, for example, *172.16.200.57*.
        - Set *External Service Port* to *8082 - 8082*.
        - Use the same *Map to Port* numbers: *80 - 80*.
11. Create a *Virtual IP Group* and put the above three virtual IPs into that group.



To see the results:

1. Apply the above virtual IP to the Firewall policy.
2. The results are:
    - Access 10.1.100.199:8080 from external network and FortiGate maps to 172.16.200.55:80 in internal network.
    - Access 10.1.100.199:8081 from external network and FortiGate maps to 172.16.200.56:80 in internal network.
    - Access 10.1.100.199:8082 from external network and FortiGate maps to 172.16.200.57:80 in internal network

Netskope Exhibit 1017

Virtual server

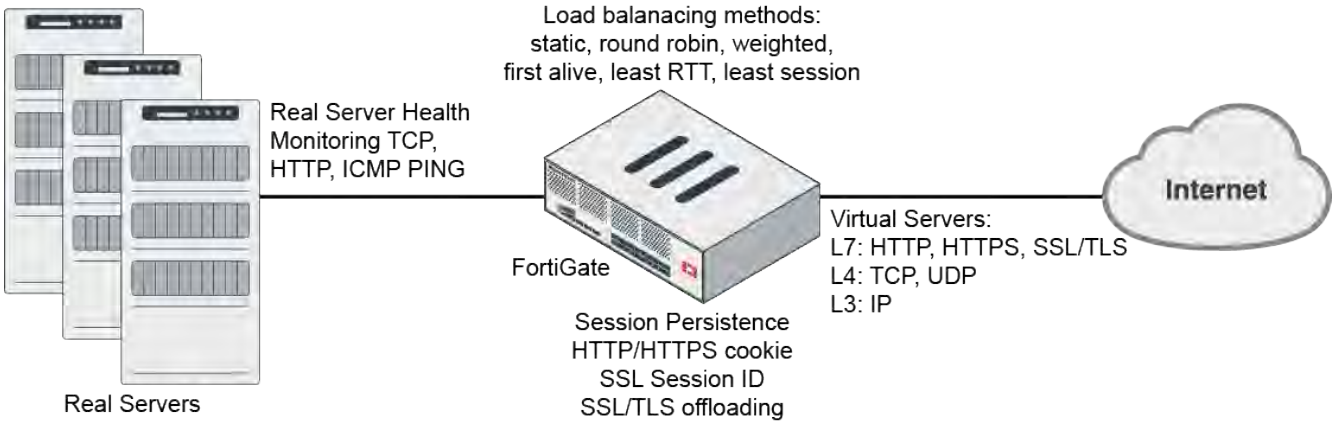
This topic shows a special virtual IP type: virtual server. Use this type of VIP to implement server load balancing.

The FortiOS server load balancing contains all the features of a server load balancing solution. You can balance traffic across multiple backend servers based on multiple load balancing schedules including:

- Static (failover)
- Round robin
- Weighted (to account for different sized servers or based on the health and performance of the server including round trip time and number of connections)

The load balancer supports HTTP, HTTPS, IMAPS, POP3S, SMTPS, SSL/TLS, and generic TCP/UDP and IP protocols. Session persistence is supported based on the SSL session ID based on an injected HTTP cookie, or based on the HTTP or HTTPS host. SSL/TLS load balancing includes protection from protocol downgrade attacks. Server load balancing is supported on most FortiGate devices and includes up to 10,000 virtual servers on high end systems.

Sample topology



SSL/TLS offloading

FortiGate SSL/TLS offloading is designed for the proliferation of SSL/TLS applications. The key exchange and encryption/decryption tasks are offloaded to the FortiGate unit where they are accelerated using FortiASIC technology which provides significantly more performance than a standard server or load balancer. This frees up valuable resources on the server farm to give better response to business operations. Server load balancing offloads most SSL/TLS versions including SSL 3.0, TLS 1.0, and TLS 1.2, and supports full mode or half mode SSL offloading with DH key sizes up to 4096 bits.

FortiGate SSL offloading allows the application payload to be inspected before it reaches your servers. This prevents intrusion attempts, blocks viruses, stops unwanted applications, and prevents data leakage. SSL/TLS content inspection supports TLS versions 1.0, 1.1, and 1.2 and SSL versions 1.0, 1.1, 1.2, and 3.0.

Virtual server requirements

When creating a new virtual server, you must configure the following options:

- Virtual Server Type.
- Load Balancing Methods.

Netskope Exhibit 1017

- Health check monitoring (optional).
- Session persistence (optional).
- Virtual Server IP (External IP Address).
- Virtual Server Port (External Port).
- Real Servers (Mapped IP Address & Port).

Virtual server types

Select the protocol to be load balanced by the virtual server. If you select a general protocol such as IP, TCP, or UDP, the virtual server load balances all IP, TCP, or UDP sessions. If you select specific protocols such as HTTP, HTTPS, or SSL, you can apply additional server load balancing features such as *Persistence* and *HTTP Multiplexing*.

HTTP	Select *HTTP* to load balance only HTTP sessions with the destination port number that matches the *Virtual Server Port* setting. Change *Virtual Server Port* to match the destination port of the sessions to be load balanced (usually port 80 for HTTP sessions). You can enable *HTTP Multiplexing*. You can also set *Persistence* to *HTTP Cookie* to enable cookie-based persistence.
HTTPS	Select *HTTPS* to load balance only HTTPS sessions with the destination port number that matches the *Virtual Server Port* setting. Change *Virtual Server Port* to match the destination port of the sessions to be load balanced (usually port 443 for HTTPS sessions). You can enable *HTTP Multiplexing*. You can also set *Persistence* to *HTTP Cookie* to enable cookie-based persistence, or you can set *Persistence* to *SSL Session ID*.
IMAPS	Select *IMAPS* to load balance only IMAPS sessions with the destination port number that matches the *Virtual Server Port* setting. Change *Virtual Server Port* to match the destination port of the sessions to be load balanced (usually port 993 for IMAPS sessions). You can also set *Persistence* to *SSL Session ID*.
POP3S	Select *POP3S* to load balance only POP3S sessions with the destination port number that matches the *Virtual Server Port* setting. Change *Virtual Server Port* to match the destination port of the sessions to be load balanced (usually port 995 for POP3S sessions). You can also set *Persistence* to *SSL Session ID*.
SMTPS	Select *SMTPS* to load balance only SMTPS sessions with the destination port number that matches the *Virtual Server Port* setting. Change *Virtual Server Port* to match the destination port of the sessions to be load balanced (usually port 465 for SMTPS sessions). You can also set *Persistence* to *SSL Session ID*.
SSL	Select *SSL* to load balance only SSL sessions with the destination port number that matches the *Virtual Server Port* setting. Change *Virtual Server Port* to match the destination port of the sessions to be load balanced. You can also set *Persistence* to *SSL Session ID*.
TCP	Select *TCP* to load balance only TCP sessions with the destination port number that matches the *Virtual Server Port* setting. Change *Virtual Server Port* to match the destination port of the sessions to be load balanced.
UDP	Select *UDP* to load balance only UDP sessions with the destination port number that matches the *Virtual Server Port* setting. Change *Virtual Server Port* to match the destination port of the sessions to be load balanced.
IP	Select *IP* to load balance all sessions accepted by the security policy that contains this virtual server.

Netskope Exhibit 1017

Load balancing methods

The load balancing method defines how sessions are load balanced to real servers.

All load balancing methods do not send traffic to real servers that are down or not responding. FortiGate can only determine if a real server is not responding by using a health check monitor. You should always add at least one health check monitor to a virtual server or to real servers; otherwise load balancing might try to distribute sessions to real servers that are not functioning.

Static	The traffic load is statically spread evenly across all real servers. Sessions are not assigned according to how busy individual real servers are. This load balancing method provides some persistence because all sessions from the same source address always go to the same real server. Because the distribution is stateless, so if a real server is added, removed, or goes up or down, the distribution is changed and persistence might be lost.
Round Robin	Directs new requests to the next real server. This method treats all real servers as equals regardless of response time or the number of connections. This method does not direct requests to real servers that down or non responsive.
Weighted	Real servers with a higher weight value receive a larger percentage of connections. Set the real server weight when adding a real server.
Least Session	Directs requests to the real server that has the least number of current connections. This method works best in environments where the real servers or other equipment you are load balancing all have similar capabilities. This load balancing method uses the FortiGate session table to track the number of sessions being processed by each real server. The FortiGate unit cannot detect the number of sessions actually being processed by a real server.
Least RTT	Directs sessions to the real server with the lowest round trip time. The round trip time is determined by a ping health check monitor. The default is 0 if no ping health check monitors are added to the virtual server.
First Alive	Directs sessions to the first live real server. This load balancing schedule provides real server failover protection by sending all sessions to the first live real server. If a real server fails, all sessions are sent to the next live real server. Sessions are not distributed to all real servers so all sessions are processed by the first real server only.
HTTP Host	Load balances HTTP host connections across multiple real servers using the host's HTTP header to guide the connection to the correct real server.

Health check monitoring

In the FortiGate GUI, you can configure health check monitoring so that the FortiGate unit can verify that real servers are able respond to network connection attempts. If a real server responds to connection attempts, the load balancer continues to send sessions to it. If a real server stops responding to connection attempts, the load balancer assumes that the server is down and does not send sessions to it. The health check monitor configuration determines how the load balancer tests real servers. You can use a single health check monitor for multiple load balancing configurations. You can configure TCP, HTTP, and Ping health check monitors. You usually set the health check monitor to use the same protocol as the traffic being load balanced to it. For example, for an HTTP load balancing configuration, you would normally use an HTTP health check monitor.

Netskope Exhibit 1017

Session persistence

Use persistence to ensure a user is connected to the same real server every time the user makes an HTTP, HTTPS, or SSL request that is part of the same user session. For example, if you are load balancing HTTP and HTTPS sessions to a collection of eCommerce web servers, when users make a purchase, they will be starting multiple sessions as they navigate the eCommerce site. In most cases, all the sessions started by this user during one eCommerce session should be processed by the same real server. Typically, the HTTP protocol keeps track of these related sessions using cookies. HTTP cookie persistence ensure all sessions that are part of the same user session are processed by the same real server.

When you configure persistence, the FortiGate unit load balances a new session to a real server according to the load balance method. If the session has an HTTP cookie or an SSL session ID, the FortiGate unit sends all subsequent sessions with the same HTTP cookie or SSL session ID to the same real server.

Real servers

Add real servers to a load balancing virtual server to provide information the virtual server requires to send sessions to the server. A real server configuration includes the IP address of the real server and port number the real server receives sessions on. The FortiGate unit sends sessions to the real server's IP address using the destination port number in the real server configuration.

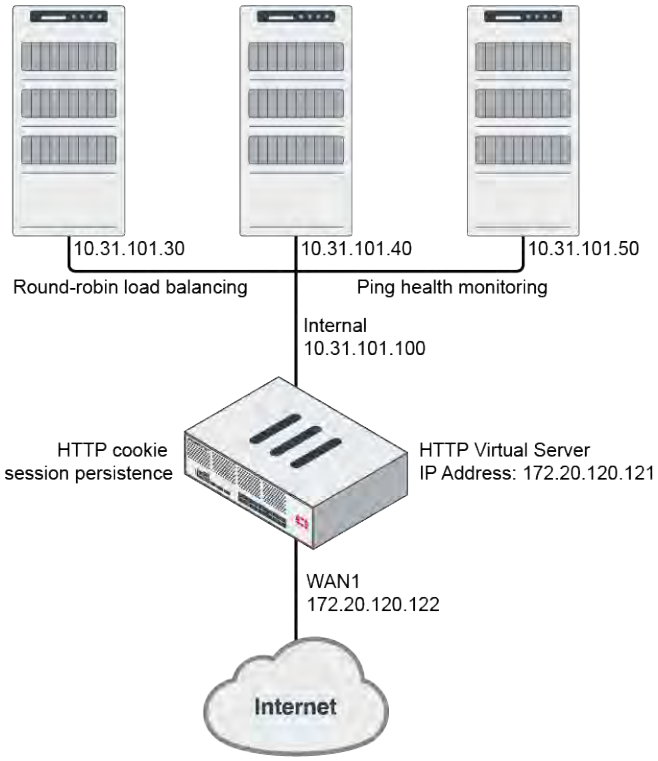
When configuring a real server, you can also specify the weight (if the load balance method is set to *Weighted*) and you can limit the maximum number of open connections between the FortiGate unit and the real server. If the maximum number of connections is reached for the real server, the FortiGate unit automatically switches all further connection requests to other real servers until the connection number drops below the limit. Setting *Maximum Connections* to 0 means that the FortiGate unit does not limit the number of connections to the real server.

Sample of HTTP load balancing to three real web servers

This example describes the steps to configure the load balancing configuration below. In this configuration, a FortiGate unit is load balancing HTTP traffic from the Internet to three HTTP servers on the internal network. HTTP sessions are accepted at the wan1 interface with destination IP address 172.20.120.121 on TCP port 8080, and forwarded from the internal interface to the web servers. When forwarded, the destination address of the session is translated to the IP address of one of the web servers.

This load balancing configuration also includes session persistence using HTTP cookies, round-robin load balancing, and TCP health monitoring for the real servers. Ping health monitoring consists of the FortiGate unit using ICMP ping to ensure the web servers can respond to network traffic.

Netskope Exhibit 1017



General steps:

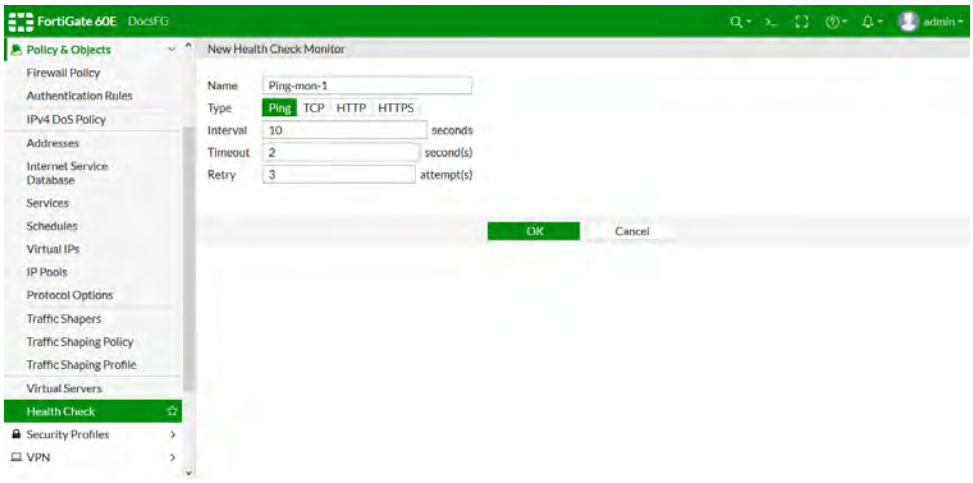
1. Create a health check monitor.
   A ping health check monitor causes the FortiGate to ping the real servers every 10 seconds. If one of the servers does not respond within 2 seconds, the FortiGate unit will retry the ping 3 times before assuming that the HTTP server is not responding.
2. Create a load balance virtual server with three real servers.
3. Add the load balancing virtual server to a policy as the destination address.

To see the virtual servers and health check monitors options in the GUI, *Load Balance* must be selected in *Feature Visibility* > *Additional Features*. See Feature visibility on page 694 on page 1 for details.

Configure a load balancing virtual server in the GUI

To create a health check monitor:

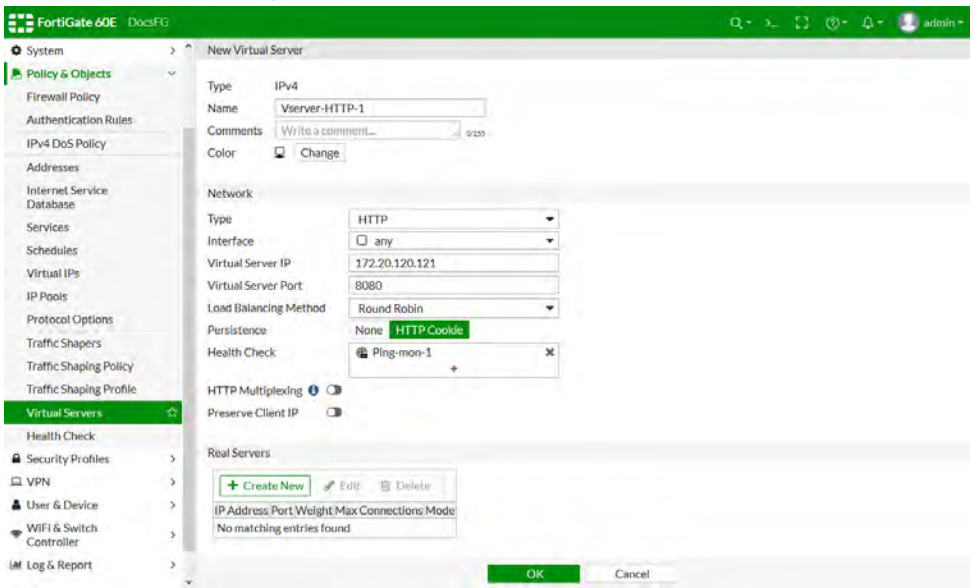
1. Go to *Policy & Objects* > *Health Check*.
2. Click *Create New*.
3. Set the following:
   - *Name* to *Ping-mon-1*
   - *Type* to *Ping*
   - *Interval* to *10* seconds
   - *Timeout* to *2* seconds
   - *Retry* to *3* attempt(s)



4. Click *OK*.

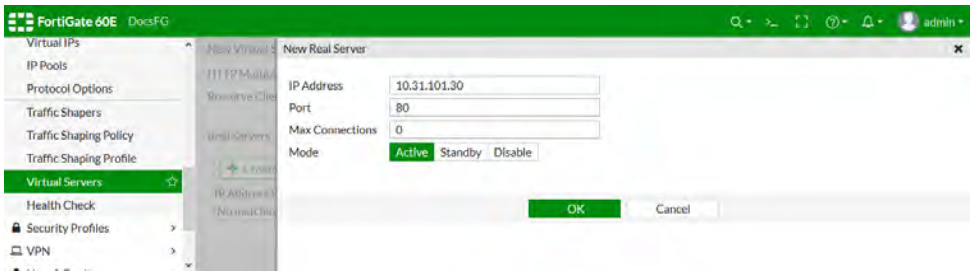
To create a virtual server:

1. Go to *Policy & Objects* > *Virtual Servers*.
2. Click *Create New*.
3. Set the following:
   - *Name* to *Vserver-HTTP-1*
   - *Type* to *HTTP*
   - *Interface* to *wan1*
   - *Virtual Server IP* to *172.20.120.121*
   - *Virtual Server Port* to *8080*
   - *Load Balance Method* to *Round Robin*
   - *Persistence* to *HTTP Cookie*
   - *Health Check* to *Ping-mon-1*



4. In the *Real Servers* table, click *Create New*.

5. Set the following for the first real server:
   - *IP Address* to *10.31.101.30*
   - *Port* to *80*
   - *Max Connections* to *0*
   - *Mode* to *Active*

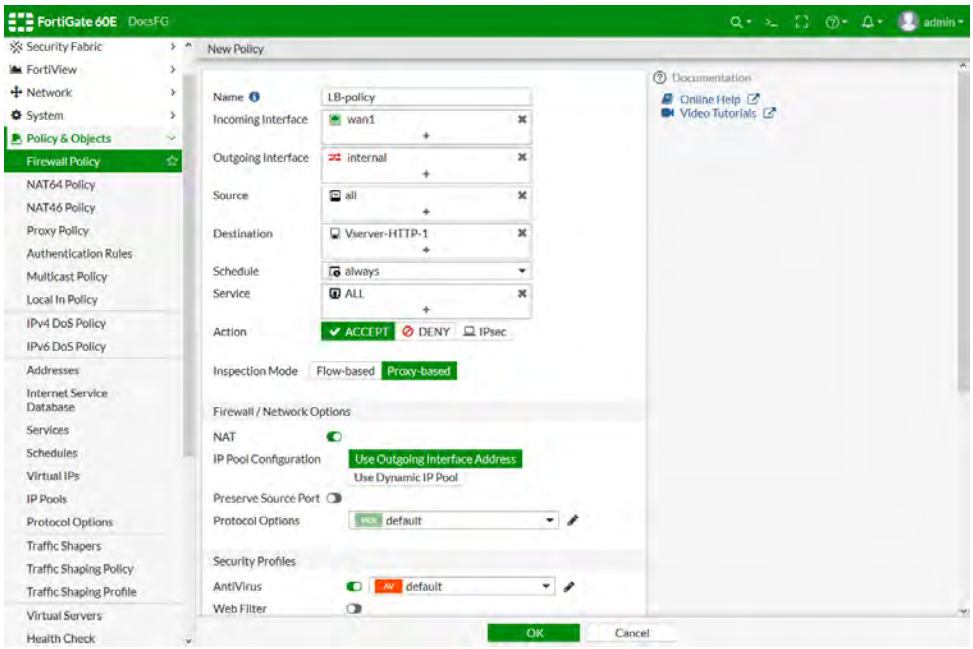


6. Configure two more real servers with IP addresses 10.31.101.40 and 10.31.101.50, and the remaining settings the same as the first real server.
7. Click *OK*.

To create a security policy that includes the load balance virtual server as the destination address:

1. Go to *Policy & Objects* > *Firewall Policy*.
2. Click *Create New*.
3. Set the *Inspection Mode* to *Proxy-based*. The new virtual server will not be available if the inspection mode is *Flow-based*.
4. Set the following:
   - *Name* to *LB-policy*
   - *Incoming Interface* to *wan1*
   - *Outgoing Interface* to *internal*
   - *Source* to *all*
   - *Destination* to *Vserver-HTTP-1*
   - *Schedule* to *always*
   - *Service* to *ALL*
   - *Action* to *ACCEPT*
5. Enable *NAT* and set *IP Pool Configuration* to *Use Outgoing Interface Address*.

6. Enable *AntiVirus* and select an antivirus profile.



7. Click *OK*.

Configure a load balancing virtual server in the CLI

To configure HTTP load balancing to three real web servers in the CLI:

1. Create a health check monitor:

```
config firewall ldb-monitor
    edit "Ping-mon-1"
        set type ping
        set interval 10
        set timeout 2
        set retry 3
    next
end
```

2. Create a virtual server:

```
config firewall vip
    edit "Vserver-HTTP-1"
        set type server-load-balance
        set extip 172.20.120.121
        set extintf "any"
        set server-type http
        set monitor "Ping-mon-1"
        set ldb-method round-robin
        set persistence http-cookie
        set extport 8080
        config realservers
            edit 1
                set ip 10.31.101.30
                set port 80
```

```
            next
            edit 2
                set ip 10.31.101.40
                set port 80
            next
            edit 3
                set ip 10.31.101.50
                set port 80
            next
        end
    next
end
```

3. Add the load balancing virtual server to a policy as the destination address:

```
config firewall policy
    edit 2
        set name "LB-policy"
        set inspection-mode proxy
        set srcintf "wan1"
        set dstintf "internal"
        set srcaddr "all"
        set dstaddr "Vserver-HTTP-1"
        set action accept
        set schedule "always"
        set service "ALL"
        set utm-status enable
        set ssl-ssh-profile "certificate-inspection"
        set av-profile "default"
        set fsso disable
        set nat enable
    next
end
```

Results

Traffic accessing 172.20.120.121:8080 is forwarded in turn to the three real servers.

If the access request has an http-cookie, FortiGate forwards the access to the corresponding real server according to the cookie.

Policy with Internet Service

The following topics provide instructions on configuring policies with Internet Service:

- Using Internet Service in policy on page 742
- Using custom Internet Service in policy on page 744
- Using extension Internet Service in policy on page 746
- Global IP address information database on page 749
- IP reputation filtering on page 751
- Internet service groups in policies on page 752
- Allow creation of ISDB objects with regional information on page 756

Netskope Exhibit 1017

Using Internet Service in policy

This topic shows how to apply a predefined Internet Service entry into a policy.

The Internet Service Database is a comprehensive public IP address database that combines IP address range, IP owner, service port number, and IP security credibility. The data comes from the FortiGuard service system. Information is regularly added to this database, for example, geographic location, IP reputation, popularity & DNS, and so on. All this information helps users define Internet security more effectively. You can use the contents of the database as criteria for inclusion or exclusion in a policy.

From FortiOS version 5.6, Internet Service is included in the firewall policy. It can be applied to a policy only as a destination object. From version 6.0, Internet Service can be applied both as source and destination objects in a policy. You can also apply Internet Services to shaping policy.

There are three types of Internet Services you can apply to a firewall policy:

- Predefined Internet Services
- Custom Internet Services
- Extension Internet Services

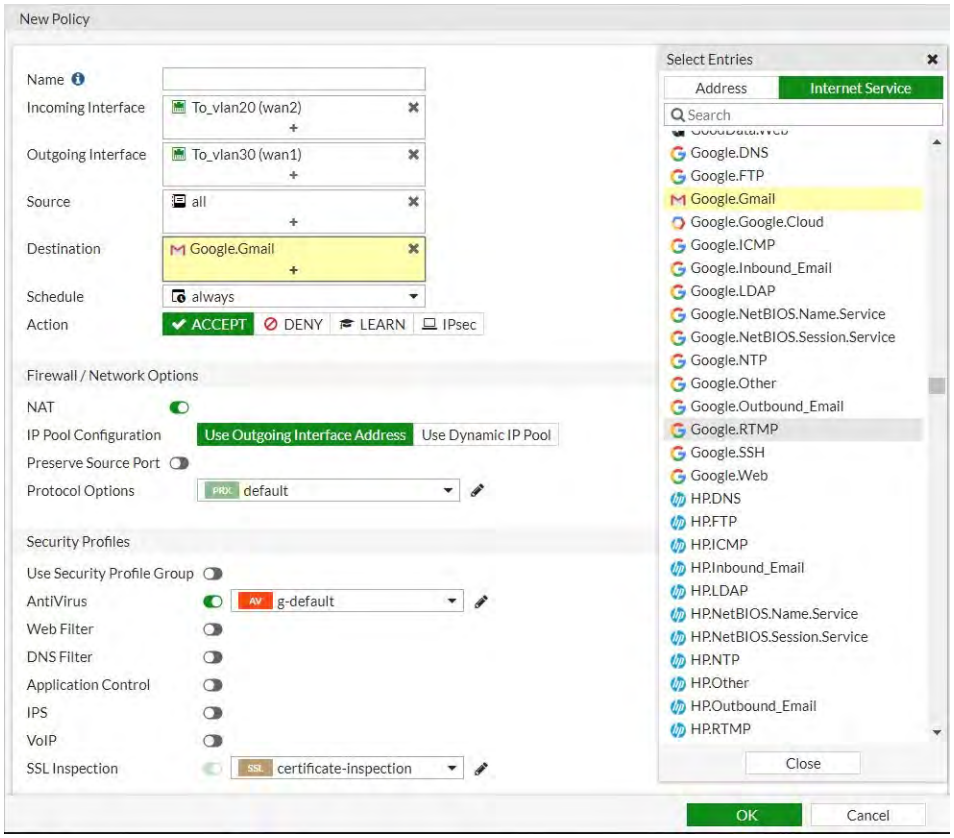
Sample configuration

To apply a predefined Internet Service entry to a policy using the GUI:

1. Go to *Policy & Objects* and create a new policy.
2. In the *Source* or *Destination* field, click +.
3. In the *Select Entries* pane, click *Internet Service*.

Netskope Exhibit 1017

4. Locate and click *Google.Gmail*.



5. Configure the other fields and then click *OK*.

To apply a predefined Internet Service entry to a policy using the CLI:

In the CLI, enable the `internet-service` first and then use its ID to apply the policy.

This example uses Google Gmail and its ID is 65646. Each Internet Service has a unique ID.

```
config firewall policy
    edit 9
        set name "Internet Service in Policy"
        set srcintf "wan2"
        set dstintf "wan1"
        set srcaddr "all"
        set internet-service enable
        set internet-service-id 65646
        set action accept
        set schedule "always"
        set utm-status enable
        set av-profile "g-default"
        set ssl-ssh-profile "certificate-inspection"
        set nat enable
    next
end
```


Netskope Exhibit 1017

To diagnose an Internet Service entry using the CLI:

```
# diagnose internet-service id-summary 65646
Version: 0000600096
Timestamp: 201902111802
Total number of IP ranges: 444727
Number of Groups: 7
Group(0), Singularity(20), Number of IP ranges(142740)
Group(1), Singularity(19), Number of IP ranges(1210)
Group(2), Singularity(16), Number of IP ranges(241)
Group(3), Singularity(15), Number of IP ranges(38723)
Group(4), Singularity(10), Number of IP ranges(142586)
Group(5), Singularity(8), Number of IP ranges(5336)
Group(6), Singularity(6), Number of IP ranges(113891)
Internet Service: 65646(Google.Gmail)
Number of IP range: 60
Number of IP numbers: 322845
Singularity: 15
Reputation: 5(Known and verified safe sites such as Gmail, Amazon, eBay, etc.)
Icon Id: 510
Second Level Domain: 53(gmail.com)
Direction: dst
Data source: isdb
```

Result

Because the IP and services related to Google Gmail on the Internet are included in this Internet Service (65646), all traffic to Google Gmail is forwarded by this policy.

Using custom Internet Service in policy

Custom Internet Services can be created and used in firewall policies.

When creating a custom Internet Service, you must set following elements:

- IP or IP ranges
- Protocol number
- Port or port ranges
- Reputation

You must use CLI to create a custom Internet Service, except for geographic based services (see Allow creation of ISDB objects with regional information on page 756).

CLI syntax

```
config firewall internet-service-custom
    edit <name>
       set comment <comment>
       set reputation {1 | 2 | 3 | 4 | 5}
       config entry
           edit <ID>
               set protocol <protocol #>
               set dst <object_name>
               config port-range
```


Netskope Exhibit 1017

```
                    edit <ID>
                        set start-port <port #>
                        set end-port <port #>
                    next
                end
            next
        end
    end
end
```

Sample configuration

To configure a custom Internet Service:

```
config firewall internet-service-custom
   edit "test-isdb-1"
       set comment "Test Custom Internet Service"
       set reputation 4
       config entry
          edit 1
              set protocol 6
              config port-range
                  edit 1
                      set start-port 80
                      set end-port 443
                  next
              end
              set dst "10-1-100-0"
          next
          edit 2
              set protocol 6
              config port-range
                  edit 1
                      set start-port 80
                      set end-port 80
                  next
              end
              set dst "172-16-200-0"
          next
       end
   next
end
```

To apply a custom Internet Service into a policy:

```
config firewall policy
    edit 1
        set name "Internet Service in Policy"
        set srcintf "wan2"
        set dstintf "wan1"
        set srcaddr "all"
        set internet-service enable
        set internet-service-id 65646
        set internet-service-custom "test-isdb-1"
        set action accept
```


Netskope Exhibit 1017

```
        set schedule "always"
        set utm-status enable
        set av-profile "g-default"
        set ssl-ssh-profile "certificate-inspection"
        set nat enable
    next
end
```

Result

In addition to the IP address, IP address ranges, and services allowed by Google.Gmail, this policy also allows the traffic which access to 10.1.100.0/24 and TCP/80-443 and 172.16.200.0/24 and TCP/80.

Using extension Internet Service in policy

Extension Internet Service lets you add custom or remove existing IP address and port ranges to an existing predefined Internet Service entries. Using an extension type Internet Service is actually editing a predefined type Internet Service entry and adding IP address and port ranges to it.

When creating an extension Internet Service and adding custom ranges, you must set following elements:

- IP or IP ranges
- Protocol number
- Port or port ranges

You must use CLI to add custom IP address and port entries into a predefined Internet Service.

You must use GUI to remove entries from a predefined Internet Service.

Custom extension Internet Service CLI syntax

```
config firewall internet-service-extension
    edit <ID #>
        set comment <comment>
        config entry
            edit <ID #>
                set protocol <number #>
                set dst <object_name>
                config port-range
                    edit <ID #>
                        set start-port <number #>
                        set end-port <number #>
                    next
                end
            next
        end
    end
end
```


Netskope Exhibit 1017

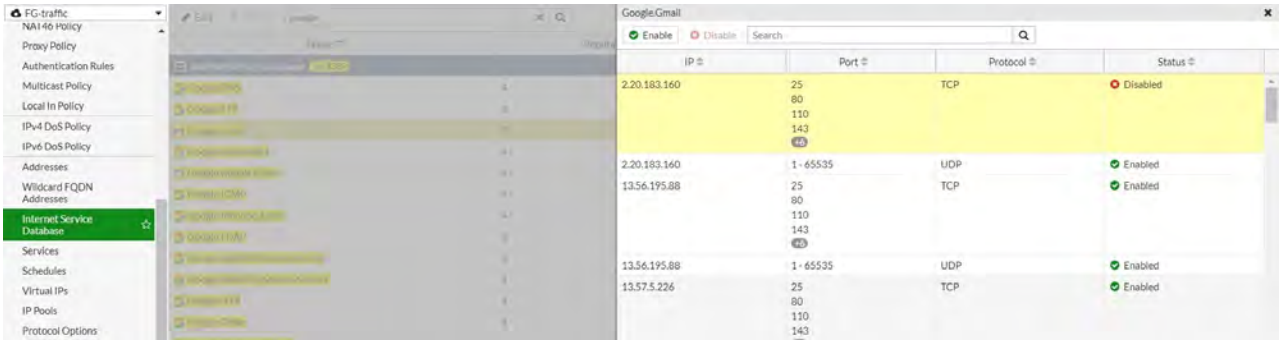
Sample configuration

To configure an extension Internet Service using the CLI:

```
config firewall internet-service-extension
    edit 65646
        set comment "Test Extension Internet Service 65646"
        config entry
            edit 1
                set protocol 6
                config port-range
                    edit 1
                        set start-port 80
                        set end-port 443
                    next
                end
                set dst "172-16-200-0"
            next
            edit 2
                set protocol 17
                config port-range
                    edit 1
                        set start-port 53
                        set end-port 53
                    next
                end
                set dst "10-1-100-0"
            next
        end
    next
end
```

To remove IP address and port entries from an existing Internet Service:

1. Go to *Policy & Objects* > *Internet Service Database*.
2. Search for *Google.Gmail*.
3. Select *Google.Gmail* and click *Edit*.
4. Locate the *IP* entry you want to remove and click *Disable* beside that entry.



5. Click *Return*.
6. When you complete the actions in the GUI, the CLI automatically generates the configuration from your GUI actions:

```
config firewall internet-service-extension
    edit 65646
```

```
set comment "Test Extension Internet Service 65646"
config entry
    edit 1
        set protocol 6
        config port-range
            edit 1
                set start-port 80
                set end-port 443
            next
        end
        set dst "172-16-200-0"
    next
    edit 2
        set protocol 17
        config port-range
            edit 1
                set start-port 53
                set end-port 53
            next
        end
        set dst "10-1-100-0"
    next
end
config disable-entry
    edit 1
        set protocol 6
        config port-range
            edit 1
                set start-port 25
                set end-port 25
            next
            edit 2
                set start-port 80
                set end-port 80
            next
            edit 3
                set start-port 110
                set end-port 110
            next
            edit 4
                set start-port 143
                set end-port 143
            next
            edit 5
                set start-port 443
                set end-port 443
            next
            edit 6
                set start-port 465
                set end-port 465
            next
            edit 7
                set start-port 587
                set end-port 587
            next
            edit 8
```


Netskope Exhibit 1017

```
                        set start-port 993
                        set end-port 993
                    next
                    edit 9
                        set start-port 995
                        set end-port 995
                    next
                    edit 10
                        set start-port 2525
                        set end-port 2525
                    next
                end
                config ip-range
                    edit 1
                        set start-ip 2.20.183.160
                        set end-ip 2.20.183.160
                    next
                end
            next
        end
    next
end
```

To apply an extension Internet Service into policy using the CLI:

```
config firewall policy
    edit 9
        set name "Internet Service in Policy"
        set srcintf "wan2"
        set dstintf "wan1"
        set srcaddr "all"
        set internet-service enable
        set internet-service-id 65646
        set action accept
        set schedule "always"
        set utm-status enable
        set av-profile "g-default"
        set ssl-ssh-profile "certificate-inspection"
        set nat enable
    next
end
```

Result

In addition to the IP addresses, IP address ranges, and services allowed by Google.Gmail, this policy also allows the traffic which accesses 10.1.100.0/24 and UDP/53 and 172.16.200.0/24 and TCP/80-443. At the same time, the traffic that accesses 2.20.183.160 is dropped because this IP address and port is disabled from Google.Gmail.

Global IP address information database

The Internet Service and IP Reputation databases download details about public IP address, including: ownership, known services, geographic location, blocklisting information, and more. The details are available in drilldown information, tooltips, and other mechanisms in the FortiView and other pages.

Netskope Exhibit 1017

The global IP address database is an integrated database containing all public IP addresses, and is implemented in the Internet Service Database.

To view the owner of the IP address:

```
(global) # get firewall internet-service-owner ?
    id    Internet Service owner ID.
    1  Google
    2  Facebook
    3  Apple
    4  Yahoo
    5  Microsoft
    ......
    115  Cybozu
    116  VNC
```

To check for any known service running on an IP address:

```
(global) # diagnose internet-service info FG-traffic 6 80 8.8.8.8
    Internet Service: 65537(Google.Web)
```

To check GeoIP location and blocklist information:

```
(global) # diagnose internet-service id 65537 | grep 8.8.8.8
```

To check a known malicious server:

```
(global) # diagnose internet-service id-summary 3080383
    Version: 0000600096
    Timestamp: 201902111802
    Total number of IP ranges: 444727
    Number of Groups: 7
    Group(0), Singularity(20), Number of IP ranges(142740)
    Group(1), Singularity(19), Number of IP ranges(1210)
    Group(2), Singularity(16), Number of IP ranges(241)
    Group(3), Singularity(15), Number of IP ranges(38723)
    Group(4), Singularity(10), Number of IP ranges(142586)
    Group(5), Singularity(8), Number of IP ranges(5336)
    Group(6), Singularity(6), Number of IP ranges(113891)
    Internet Service: 3080383(Botnet.C&C.Server)
    Number of IP range: 111486
    Number of IP numbers: 111486
    Singularity: 20
    Reputation: 1(Known malicious sites related to botnet servers, phishing sites, etc.)
    Icon Id: 591
    Second Level Domain: 1(other)
    Direction: dst
    Data source: irdb
```

To check questionable usage:

```
(global) # diagnose internet-service id-summary 2818238
    Version: 0000600096
    Timestamp: 201902111802
    Total number of IP ranges: 444727
```


Netskope Exhibit 1017

```
    Number of Groups: 7
    Group(0), Singularity(20), Number of IP ranges(142740)
    Group(1), Singularity(19), Number of IP ranges(1210)
    Group(2), Singularity(16), Number of IP ranges(241)
    Group(3), Singularity(15), Number of IP ranges(38723)
    Group(4), Singularity(10), Number of IP ranges(142586)
    Group(5), Singularity(8), Number of IP ranges(5336)
    Group(6), Singularity(6), Number of IP ranges(113891)
    Internet Service: 2818238(Tor.Relay.Node)
    Number of IP range: 13718
    Number of IP numbers: 13718
    Singularity: 20
    Reputation: 2(Sites providing high risk services such as TOR, proxy, P2P, etc.)
    Icon Id: 43
    Second Level Domain: 1(other)
    Direction: dst
    Data source: irdb

(global) # diagnose internet-service id-summary 2818243
    Version: 0000600096
    Timestamp: 201902111802
    Total number of IP ranges: 444727
    Number of Groups: 7
    Group(0), Singularity(20), Number of IP ranges(142740)
    Group(1), Singularity(19), Number of IP ranges(1210)
    Group(2), Singularity(16), Number of IP ranges(241)
    Group(3), Singularity(15), Number of IP ranges(38723)
    Group(4), Singularity(10), Number of IP ranges(142586)
    Group(5), Singularity(8), Number of IP ranges(5336)
    Group(6), Singularity(6), Number of IP ranges(113891)
    Internet Service: 2818243(Tor.Exit.Node)
    Number of IP range: 1210
    Number of IP numbers: 1210
    Singularity: 19
    Reputation: 2(Sites providing high risk services such as TOR, proxy, P2P, etc.)
    Icon Id: 43
    Second Level Domain: 1(other)
    Direction: src
    Data source: irdb
```

IP reputation filtering

There are currently five reputation levels in the Internet Service Database (ISDB), and custom reputation levels can be defined in a custom internet service. You can configure firewall policies to filter traffic according to the desired reputation level. If the reputation level of either the source or destination IP address is equal to or greater than the level set in the policy, then the packet is forwarded, otherwise, the packet is dropped.

The five default reputation levels are:

1	Known malicious sites, such as phishing sites or sites related to botnet servers
2	High risk services sites, such as TOR, proxy, and P2P
3	Unverified sites

Netskope Exhibit 1017

4	Reputable social media sites, such as Facebook and Twitter
5	Known and verified safe sites, such as Gmail, Amazon, and eBay

The default minimum reputation level in a policy is zero, meaning that the reputation filter is disabled.

For IP addresses that are not included in the ISDB, the default reputation level is three.

The default reputation direction is `destination`.

To set the reputation level and direction in a policy using the CLI:

```
config firewall policy
    edit 1
        set srcintf "wan2"
        set dstintf "wan1"
        set dstaddr "all"
        set reputation-minimum 3
        set reputation-direction source
        set action accept
        set schedule "always"
        set service "ALL"
        set logtraffic all
        set auto-asic-offload disable
        set nat enable
    next
end
```

Packets from the source IP address with reputation levels three, four, or five will be forwarded by this policy.



In a policy, if `reputation-minimum` is set, and the `reputation-direction` is `destination`, then the `dstaddr`, `service`, and `internet-service` options are removed from the policy.

If `reputation-minimum` is set, and the `reputation-direction` is `source`, then the `srcaddr`, and `internet-service-src` options are removed from the policy.

Internet service groups in policies

This feature provides support for Internet Service Groups in traffic shaping and firewall policies. Service groups can be used as the source and destination of the policy. Internet Service Groups are used as criteria to match traffic; the shaper will be applied when the traffic matches.

To use a group as a destination, `internet-service` must be enabled. To use a group as a source, `internet-service-src` must be enabled.

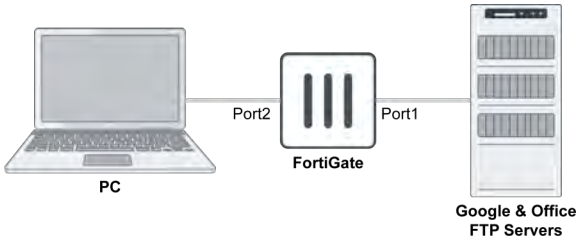
The following CLI variables are available in the `firewall policy` and `firewall shaping-policy` commands:

Variable	Description
internet-service-group <string>	Internet Service group name.
internet-service-custom-group <string>	Custom Internet Service group name.

Variable	Description
internet-service-src-group <string>	Internet Service source group name.
internet-service-src-custom-group <string>	Custom Internet Service source group name.

Examples

The following examples use the below topology.



Example 1

In this example, the PC is allowed to access Google, so all Google services are put into an Internet Service Group.

To configure access to Google services using an Internet Service Group using the CLI:

1. Create a Service Group:

```
config firewall internet-service-group
    edit "Google_Group"
        set direction destination
        set member 65537 65538 65539 65540 65542 65543 65544 65545 65550 65536 65646
    next
end
```

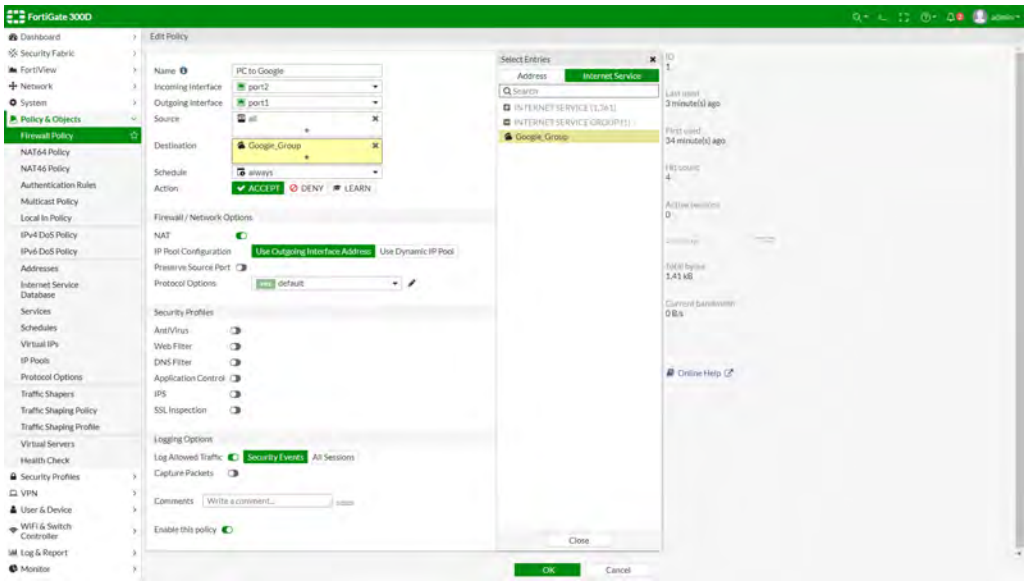
2. Create a firewall policy to allow access to all Google Services from the PC:

```
config firewall policy
    edit 1
        set name "PC to Google"
        set srcintf "port2"
        set dstintf "port1"
        set srcaddr "PC"
        set internet-service enable
        set internet-service-group "Google_Group"
        set action accept
        set schedule "always"
        set fsso disable
        set nat enable
    next
end
```

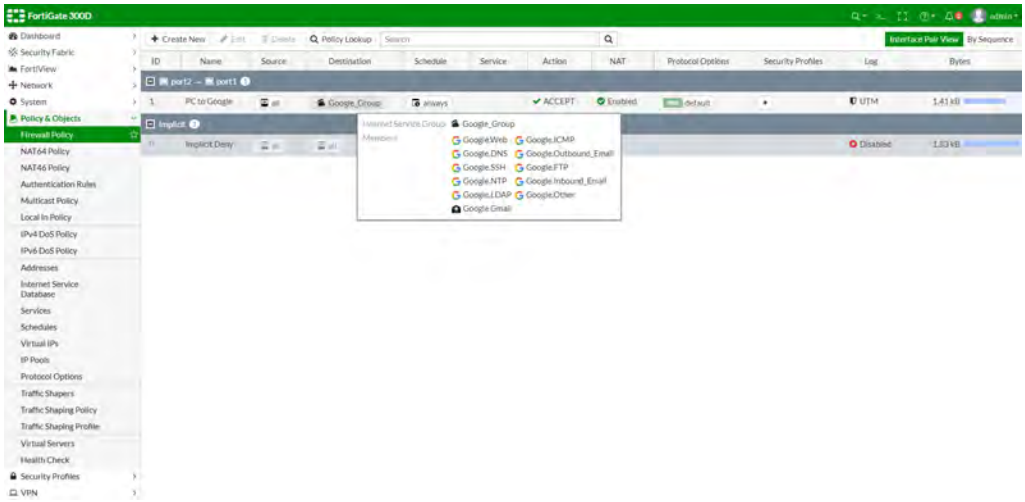
To configure access to Google services using an Internet Service Group in the GUI:

1. On the FortiGate, create a Service Group using the CLI.
2. Go to *Policy & Objects* > *Firewall Policy*, and create a new policy.

3. Set the *Destination* as the just created Internet Service Group.



4. Configure the remaining options as shown, then click *OK*.
   On the policy page, hover over the group to view a list of its members.



Example 2

In this example, two office FTP servers are put into an Internet Custom Service Group, and the PC connection to the FTP servers is limited to 1Mbps.

To put two FTP servers into a custom service group and limit the PC connection speed to them using the CLI:

1. Create custom internet services for the internal FTP servers:

```
config firewall internet-service-custom
    edit "FTP_PM"
        config entry
```

```
            edit 1
                config port-range
                    edit 1
                        set start-port 21
                        set end-port 21
                    next
                end
                set dst "PM_Server"
            next
        end
    next
    edit "FTP_QA"
        config entry
            edit 1
                config port-range
                    edit 1
                        set start-port 21
                        set end-port 21
                    next
                end
                set dst "QA_Server"
            next
        end
    next
end
```

2. Create a custom internet server group and add the just created custom internet services to it:

```
config firewall internet-service-custom-group
    edit "Internal_FTP"
        set member "FTP_QA" "FTP_PM"
    next
end
```

3. Create a traffic shaper to limit the maximum bandwidth:

```
config firewall shaper traffic-shaper
    edit "Internal_FTP_Limit_1Mbps"
        set guaranteed-bandwidth 500
        set maximum-bandwidth 1000
        set priority medium
    next
end
```

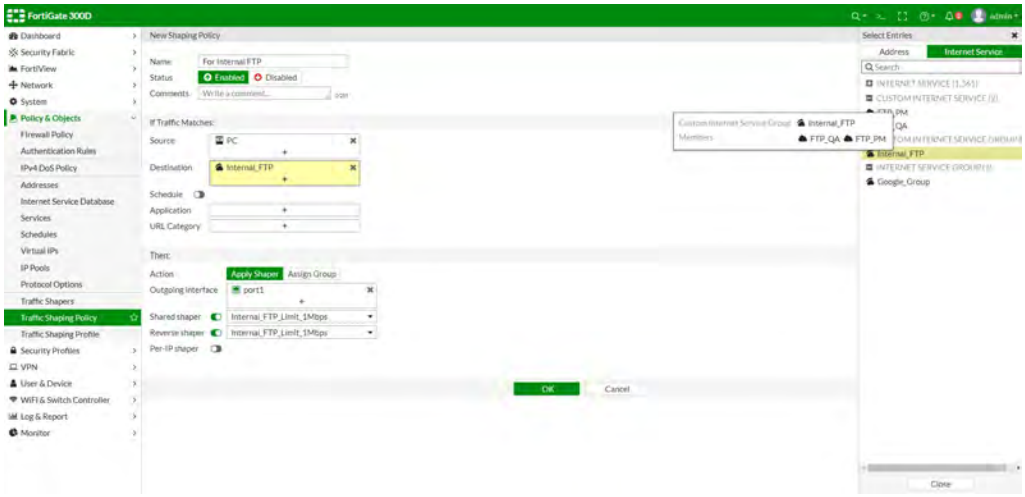
4. Create a firewall shaping policy to limit the speed from the PC to the internal FTP servers:

```
config firewall shaping-policy
    edit 1
        set name "For Internal FTP"
        set internet-service enable
        set internet-service-custom-group "Internal_FTP"
        set dstintf "port1"
        set traffic-shaper "Internal_FTP_Limit_1Mbps"
        set traffic-shaper-reverse "Internal_FTP_Limit_1Mbps"
        set srcaddr "PC"
    next
end
```


Netskope Exhibit 1017

To put two FTP servers into a custom service group and limit the PC connection speed to the using the GUI:

1. Create custom internet services for the internal FTP servers using the CLI.
2. Create a custom internet server group and add the just created custom internet services to it using the CLI.
3. Create a traffic shaper to limit the maximum bandwidth:
   a. Go to *Policy & Objects* > *Traffic Shapers*, and click *Create New*.
   b. Enter a *Name* for the shaper, such as *Internal_FTP_Limit_1Mbps*.
   c. Set the *Traffic Priority* to *Medium*.
   d. Enable *Max Bandwidth* and set it to *1000*.
   e. Enable *Guaranteed Bandwidth* and set it to *500*.
   f. Click *OK*.
4. Create a firewall shaping policy to limit the speed from the PC to the internal FTP servers:
   a. Go to *Policy & Objects* > *Traffic Shaping Policy*, and click *Create New*.
   b. Set the *Destination* as the just created Custom Internet Service Group, and apply the just create traffic shaper.



   c. Configure the remaining options as shown, then click *OK*.

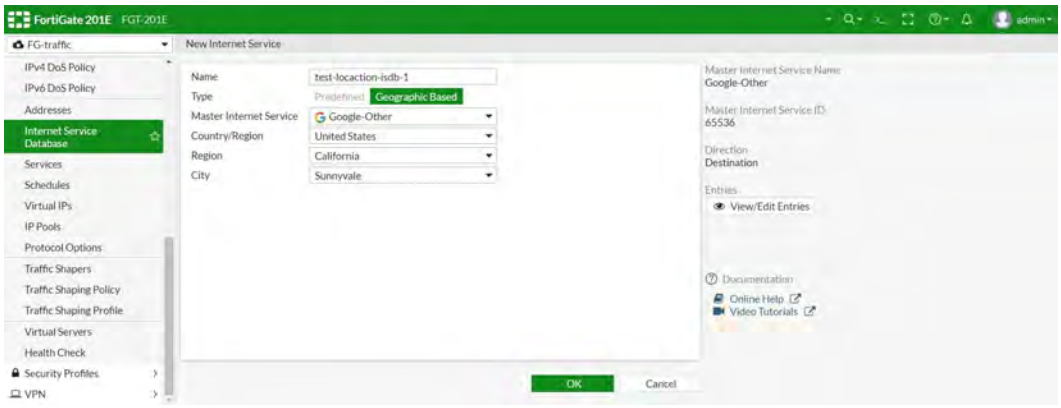
Allow creation of ISDB objects with regional information

Geographic-based Internet Service Database (ISDB) objects allow users to define a country, region, and city. These objects can be used in firewall policies for more granular control over the location of the parent ISDB object. ISDB objects are now referenced in policies by name instead of ID.

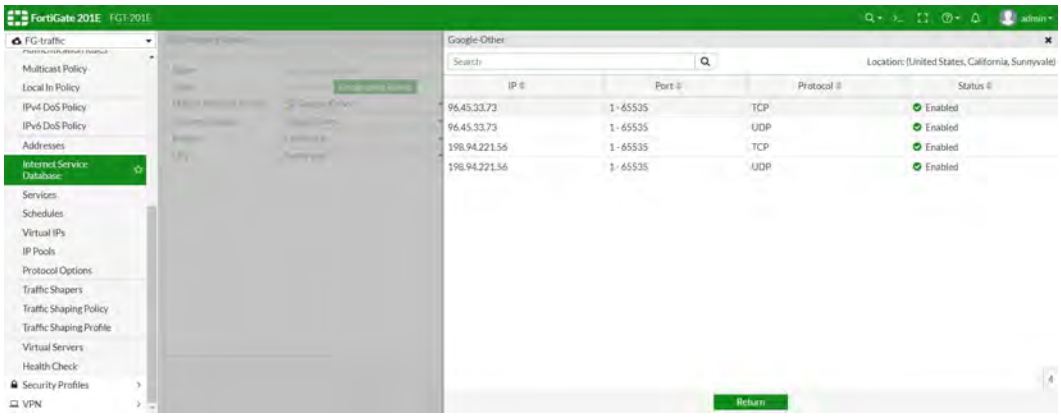
To apply a location-based ISDB object to a policy in the GUI:

1. Create the ISDB object:
   a. Go to *Policy & Objects* > *Internet Service Database* > *Create New*.
   b. For *Type*, select *Geographic Based*, and configure the other settings as needed.

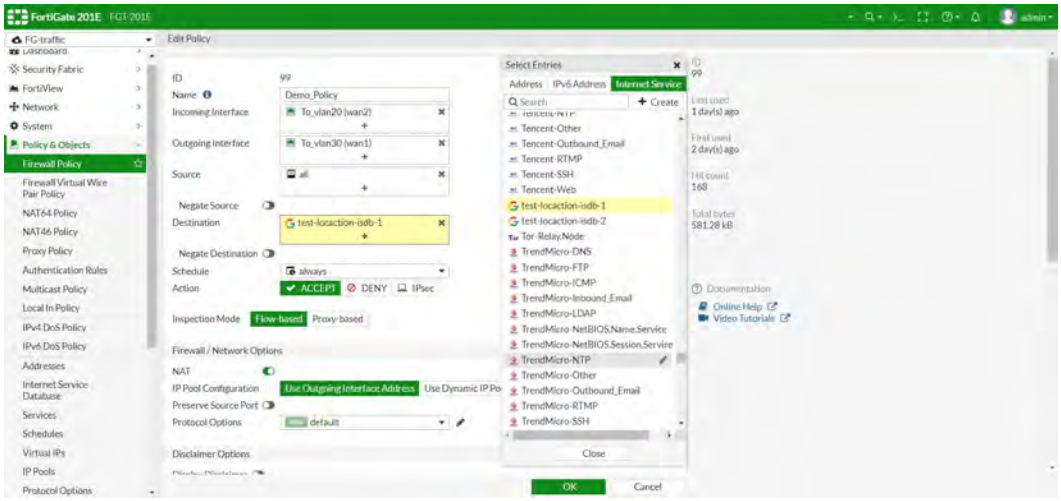
c. Click *OK*.



2. View the IP ranges in the location-based internet service:
   a. Go to *Policy & Objects* > *Internet Service Database* .
   b. In the table, hover over the object created in step 1 and click *View/Edit Entries*. The list of IPs is displayed:



   c. Click *Return*.
3. Add the ISDB object to a policy:
   a. Go to *Policy & Objects* > *Firewall Policy*. Create a new policy or edit an existing policy.
   b. For *Destination*, click *Internet Service* and select the ISDB object created in step 1.
   c. Configure the other settings as needed.



   d. Click *OK*.

To apply a location-based ISDB object to a policy in the CLI:

1. Create the ISDB object:

```
config firewall internet-service-name
    edit "test-locaction-isdb-1"
        set type location
        set internet-service-id 65536
        set country-id 840
        set region-id 283
        set city-id 23352
    next
end
```

2. View the IP ranges in the location-based internet service:

```
# diagnose internet-service id 65536 | grep "country(840) region(283) city(23352)"
```

3. Add the ISDB object to a policy:

```
config firewall policy
    edit 99
        set name "Demo_Policy"
        set srcintf "wan2"
        set dstintf "wan1"
        set srcaddr "all"
        set internet-service enable
        set internet-service-name "test-locaction-isdb-1"
        set action accept
        set schedule "always"
        set logtraffic all
        set logtraffic-start enable
        set auto-asic-offload disable
        set comments "1"
        set nat enable
    next
end
```

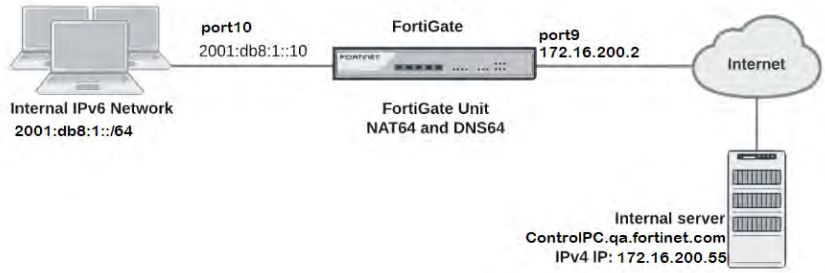
NAT64 policy and DNS64 (DNS proxy)

NAT64 policy translates IPv6 addresses to IPv4 addresses so that a client on an IPv6 network can communicate transparently with a server on an IPv4 network.

NAT64 policy is usually implemented in combination with the DNS proxy called DNS64. DNS64 synthesizes AAAA records from A records and is used to synthesize IPv6 addresses for hosts that only have IPv4 addresses. DNS proxy and DNS64 are interchangeable terms.

Netskope Exhibit 1017

Sample topology



In this example, a host on the internal IPv6 network communicates with `ControlPC.qa.fortinet.com` that only has IPv4 address on the Internet.

1. The host on the internal network does a DNS lookup for `ControlPC.qa.fortinet.com` by sending a DNS query for an AAAA record for `ControlPC.qa.fortinet.com`.
2. The DNS query is intercepted by the FortiGate DNS proxy. The DNS proxy performs an A-record query for `ControlPC.qa.fortinet.com` and gets back an RRSet containing a single A record with the IPv4 address `172.16.200.55`.
3. The DNS proxy then synthesizes an AAAA record. The IPv6 address in the AAAA record begins with the configured NAT64 prefix in the upper 96 bits and the received IPv4 address in the lower 32 bits. By default, the resulting IPv6 address is `64:ff9b::172.16.200.55`.
4. The host on the internal network receives the synthetic AAAA record and sends a packet to the destination address `64:ff9b::172.16.200.55`.
5. The packet is routed to the FortiGate internal interface (port10) where it is accepted by the NAT64 security policy.
6. The FortiGate unit translates the destination address of the packets from IPv6 address `64:ff9b::172.16.200.55` to IPv4 address `172.16.200.55` and translates the source address of the packets to `172.16.200.200` (or another address in the IP pool range) and forwards the packets out the port9 interface to the Internet.

Sample configuration

To enable display for IPv6, NAT46/NAT64, and DNS Database using the GUI:

1. Go to *System* > *Feature Visibility*.
2. In the *Basic Features* section, enable *IPv6*.
3. In the *Additional Features* section, enable the following features:
   - *NAT46 & NAT64*
   - *DNS Database*
4. Click *Apply*.

To enable display for IPv6, NAT46/NAT64, and DNS Database using the CLI:

```
config system global
   set gui-ipv6 enable
end
config system settings
   set gui-nat46-64 enable
   set gui-dns-database enable
end
```

To enable DNS proxy on the IPv6 interface using the GUI:

1. Go to *Network* > *DNS Servers*.
2. In *DNS Service on Interface*, click *Create New*.
3. For *Interface*, select *port10*.
4. Click *OK*.

To enable DNS proxy on the IPv6 interface using the CLI:

```
config system dns-server
   edit "port10"
       set mode forward-only
   next
end
```

To configure IPv6 DHCP server using the CLI:

```
config system dhcp6 server
   edit 1
       set subnet 2001:db8:1::/64
       set interface "port10"
       config ip-range
           edit 1
               set start-ip 2001:db8:1::11
               set end-ip 2001:db8:1::20
           next
       end
       set dns-server1 2001:db8:1::10
   next
end
```

To enable NAT64 and related settings using the CLI:

Enabling NAT64 with the `config system nat64` command means that all IPv6 traffic received by the current VDOM can be subject to NAT64 if the source and destination address matches an NAT64 security policy.

By default, the setting `always-synthesize-aaaa-record` is enabled. If you disable this setting, the DNS proxy (DNS64) will attempt to find an AAAA records for queries to domain names and therefore resolve the host names to IPv6 addresses. If the DNS proxy cannot find an AAAA record, it synthesizes one by adding the NAT64 prefix to the A record.

`nat64-prefix` setting is the `nat64` prefix. By default, it is `64:ff9b::/96`.

```
config system nat64
   set status enable
end
```

To create NAT64 policy using the GUI:

1. Add an IPv4 firewall address for the external network.
   a. Go to *Policy & Object* > *Addresses*.
   b. Click *Create New*.
   c. For *Name*, enter *external-net4*.
   d. For *IP/Network*, enter *17216.200.0/24*.

Netskope Exhibit 1017

   e. For *Interface*, select *port9*.
   f. Click *OK*.
2. Add an IPv6 firewall address for the internal network.
   a. Go to *Policy & Object* > *Addresses*.
   b. Click *Create New*.
   c. Change *Category* to *IPv6 Address*.
   d. For *Name*, enter *internal-net6*.
   e. For *IPv6 Address*, enter *2001:db8:1::/48*.
   f. Click *OK*.
3. Add an IP pool containing the IPv4 address that is used as the source address of the packets exiting port9.
   a. Go to *Policy & Object* > *IP Pools*.
   b. Click *Create New*.
   c. For *Name*, enter *exit-pool4*.
   d. For *External IP Range*, enter *172.16.200.200-172.16.200.210*.
   e. Click *OK*.
4. Add a NAT64 policy that allows connections from the internal IPv6 network to the external IPv4 network.
   a. Go to *Policy & Object* > *NAT64 Policy*.
   b. Click *Create New*.
   c. For *Incoming Interface*, select *port10*.
   d. For *Outgoing Interface*, select *port9*.
   e. For *Source Address*, select *internal-net6*.
   f. For *Destination Address*, select *external-net4*.
   g. Set *IP Pool Configuration* to *Use Dynamic IP Pool* and select the IP pool *exit-pool4*.
   h. Click *OK*.

To create NAT64 policy using the CLI:

```
config firewall address
    edit "external-net4"
        set associated-interface "port9"
        set subnet 172.16.200.0 255.255.255.0
    next
end
config firewall address6
    edit "internal-net6"
        set ip6 2001:db8:1::/48
    next
end
config firewall ippool
    edit "exit-pool4"
        set startip 172.16.200.200
        set endip 172.16.200.210
    next
end
config firewall policy64
    edit 1
        set srcintf "port10"
        set dstintf "port9"
        set srcaddr "internal-net6"
        set dstaddr "external-net4"
```

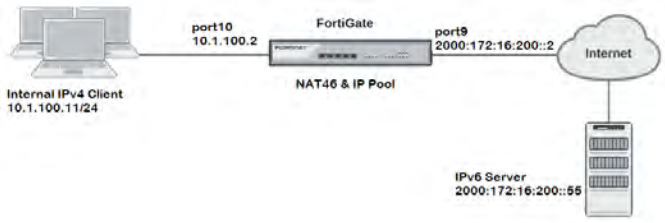

Netskope Exhibit 1017

```
        set action accept
        set schedule "always"
        set service "ALL"
        set ippool enable
        set poolname "exit-pool4"
    next
end
```

NAT46 policy

NAT46 refers to the mechanism that allows IPv4 addressed hosts to communicate with IPv6 hosts. Without such a mechanism, IPv4 environments cannot connect to IPv6 networks.

Sample topology



In this example, an IPv4 client tries to connect to an IPv6 server. A VIP is configured on FortiGate to map the server IPv6 IP address `2000:172:16:200:55` to an IPv4 address `10.1.100.55`. On the other side, an IPv6 IP pool is configured and the source address of packets from client are changed to the defined IPv6 address. In this setup, the client PC can access the server by using IP address `10.1.100.55`.

Sample configuration

To enable display for IPv6 and NAT46/NAT64 using the GUI:

1. Go to *System* > *Feature Visibility*.
2. In the *Basic Features* section, enable *IPv6*.
3. In the *Additional Features* section, enable *NAT46 & NAT64*.
4. Click *Apply*.

To enable display for IPv6 and NAT46/NAT64 using the CLI:

```
config system global
    set gui-ipv6 enable
end
config system settings
    set gui-nat46-64 enable
end
```

To configure VIP46 using the GUI:

1. Go to *Policy & Object* > *Virtual IPs*.
2. Click *Create New*.

3. For *Name*, enter *vip46_server*.
4. For *External IP Address/Range*, enter *10.1.100.55- 10.1.100.55*.
5. For *Mapped IP Address/Range*, enter *2000:172:16:200::55*.
6. Click *OK*.

To configure VIP46 using the CLI:

```
config firewall vip46
    edit "vip46_server"
        set extip 10.1.100.55
        set mappedip 2000:172:16:200::55
    next
end
```

To configure IPv6 IP pool using the GUI:

1. Go to *Policy & Object* > *IP Pools*.
2. Click *Create New*.
3. For *Name*, enter *client_expternal*.
4. For *External IP Range*, enter *2000:172:16:201::11- 2000:172:16:201::20*.
5. Click *OK*.

To configure IPv6 IP pool using the CLI:

```
config firewall ippool6
    edit "client_external"
        set startip 2000:172:16:201::11
        set endip 2000:172:16:201::20
    next
end
```

To enable NAT64 and configure address prefix using the CLI:

```
config system nat64
    set status enable
    set secondary-prefix-status enable
    config secondary-prefix
        edit "1"
            set nat64-prefix 2000:172:16:201::/96
        next
    end
end
```

To create NAT46 policy using the GUI:

1. Go to *Policy & Object* > *NAT46 Policy*.
2. Click *Create New*.
3. For *Incoming Interface*, select *port10*.
4. For *Outgoing Interface*, select *port9*.
5. For *Source Address*, select *all*.
6. For *Destination Address*, select *vip46_server*.

Netskope Exhibit 1017

7. Set *IP Pool Configuration* to *Use Dynamic IP Pool* and select the IP pool *client_expernal*.
8. Click *OK*.

To create NAT46 policy using the CLI:

```
config firewall policy46
    edit 1
        set srcintf "port10"
        set dstintf "port9"
        set srcaddr "all"
        set dstaddr "vip46_server"
        set action accept
        set schedule "always"
        set service "ALL"
        set ippool enable
        set poolname "client_external"
    next
end
```

Sample troubleshooting

Example to trace flow to see the whole process.

```
# diagnose debug flow filter saddr 10.1.100.11
# diagnose debug flow show function-name enable
show function name
# diagnose debug flow show iprope enable
show trace messages about iprope
# diagnose debug flow trace start 5

id=20085 trace_id=1 func=print_pkt_detail line=5401 msg="vd-root:0 received a packet
(proto=1, 10.1.100.11:27592->10.1.100.55:2048) from port10. type=8, code=0, id=27592,
seq=1."
id=20085 trace_id=1 func=init_ip_session_common line=5561 msg="allocate a new session-
000003b9"
id=20085 trace_id=1 func=iprope_dnat_check line=4948 msg="in-[port10], out-[]"
id=20085 trace_id=1 func=iprope_dnat_tree_check line=822 msg="len=1"
id=20085 trace_id=1 func=__iprope_check_one_dnat_policy line=4822 msg="checking gnum-100000
policy-1"
id=20085 trace_id=1 func=get_vip46_addr line=998 msg="find DNAT46: IP-2000:172:16:200::55,
port-27592"
id=20085 trace_id=1 func=__iprope_check_one_dnat_policy line=4904 msg="matched policy-1,
act=accept, vip=1, flag=100, sflag=2000000"
id=20085 trace_id=1 func=iprope_dnat_check line=4961 msg="result: skb_flags-02000000, vid-1,
ret-matched, act-accept, flag-00000100"
id=20085 trace_id=1 func=fw_pre_route_handler line=183 msg="VIP-10.1.100.55:27592, outdev-
unkown"
id=20085 trace_id=1 func=__ip_session_run_tuple line=3220 msg="DNAT 10.1.100.55:8-
>10.1.100.55:27592"
id=20085 trace_id=1 func=vf_ip_route_input_common line=2594 msg="find a route: flag=80000000
gw-10.1.100.55 via root"
id=20085 trace_id=1 func=ip4_nat_af_input line=601 msg="nat64 ipv4 received a packet
proto=1"
id=20085 trace_id=1 func=__iprope_check line=2112 msg="gnum-100012, check-ffffffffa0024ebe"
id=20085 trace_id=1 func=__iprope_check_one_policy line=1873 msg="checked gnum-100012
```


Netskope Exhibit 1017

```
policy-1, ret-matched, act-accept"
id=20085 trace_id=1 func=__iprope_user_identity_check line=1677 msg="ret-matched"
id=20085 trace_id=1 func=get_new_addr46 line=1047 msg="find SNAT46: IP-2000:172:16:201::13
(from IPPOOL), port-27592"
id=20085 trace_id=1 func=__iprope_check_one_policy line=2083 msg="policy-1 is matched, act-
accept"
id=20085 trace_id=1 func=__iprope_check line=2131 msg="gnum-100012 check result: ret-
matched, act-accept, flag-08050500, flag2-00200000"
id=20085 trace_id=1 func=iprope_policy_group_check line=4358 msg="after check: ret-matched,
act-accept, flag-08050500, flag2-00200000"
id=20085 trace_id=1 func=resolve_ip6_tuple line=4389 msg="allocate a new session-00000081"
```

Multicast processing and basic Multicast policy

You need to add firewall policies to allow packets to pass from one interface to another. Multicast packets require multicast security policies. Similar to firewall policies, in a multicast policy, the administrator specifies the source interface, destination interfaces, the allowed source address ranges, and destination addresses of the multicast traffic. You can also use multicast policies to configure source NAT and destination NAT for multicast packets.

Multicast forwarding in NAT mode

When `multicast-forward` is enabled, the FortiGate forwards any multicast IP packets in which the TTL is 2 or higher to all interfaces and VLAN interfaces except the receiving interface. The TTL in the IP header is reduced by 1. Even though the multicast packets are forwarded to all interfaces, you must add multicast policies to allow multicast packets through the FortiGate.

If `multicast-forward` is disabled, then FortiGate unit drops packets that have multicast source or destination addresses.

In NAT mode, there is a per-VDOM configuration to disable forwarding any multicast traffic. This command is only available in NAT mode.

```
config system settings
    set multicast-forward <disable|enable(default)>
end
```

You can also use the `multicast-ttl-notchange` option so that FortiGate doesn't increase the TTL value for forwarded multicast packets. Use this option only if packets are expiring before reaching the multicast router.

```
config system settings
    set multicast-ttl-notchange enable
end
```

Multicast processing in transparent mode

When `multicast-skip-policy` is enabled, no check is performed based on multicast policy. A multicast packet received on an interface is flooded unconditionally to all interfaces (except the incoming interface) belonging to the same forwarding domain. Multicast packets are forwarded even when there is no multicast policy or the multicast policy is set to deny. To forward multicast traffic based on multicast policy, `multicast-skip-policy` must be disabled.

In transparent mode, there is a per-VDOM configuration to skip policy check and forward all multicast traffic. This command is only available in transparent mode, and is disabled by default.

Netskope Exhibit 1017

```
config system settings
    set multicast-skip-policy {disable | enable}
end
```

Sample configuration

To allow RIP2 packets from port1 to port2 using the GUI:

1. Go to *Policy & Object* > *Multicast Policy*.
2. Click *Create New*.
3. For *Incoming Interface*, select *port1*.
4. For *Outgoing Interface*, select *port2*.
5. For *Source Address*, select *10.10.0.10/32*.
6. For *Destination Address*, select *RIPv2*.
7. Click *OK*.

To allow RIP2 packets from port1 to port2 using the CLI:

```
config firewall address
    edit "10.10.0.10/32"
        set subnet 10.10.0.10 255.255.255.255
    next
end
config firewall multicast-address
    edit "RIPv2"
        set start-ip 224.0.0.9
        set end-ip 224.0.0.9
    next
end
config firewall multicast-policy
    edit 2
        set srcintf "port1"
        set dstintf "port2"
        set srcaddr "10.10.0.10/32"
        set dstaddr "RIPv2"
    next
end
```

Local-in policies

While security profiles control traffic flowing through the FortiGate, local-in policies control inbound traffic that is going to a FortiGate interface.

Administrative access traffic (HTTPS, PING, SSH, and others) can be controlled by allowing or denying the service in the interface settings. Trusted hosts can be configured under an administrator to restrict the hosts that can access the administrative service.

Local-in policies allow administrators to granularly define the source and destination addresses, interface, and services. Traffic destined for the FortiGate interface specified in the policy that meets the other criteria is subject to the policies action.

Netskope Exhibit 1017

Local-in policies can be used to restrict administrative access or other services, such as VPN, that can be specified as services. You can define source addresses or address groups to restrict access from. For example, by using a geographic type address you can restrict a certain geographic set of IP addresses from accessing the FortiGate.

By default, no local-in policies are defined, so there are no restrictions on local-in traffic.



Local-in policies can only be created or edited in the CLI. You can view the existing local-in policies in the GUI by enabling it in *System > Feature Visibility* under the *Additional Features* section. This page does not list the custom local-in policies.

To configure a local-in policy using the CLI:

```
config firewall {local-in-policy | local-in-policy6}
    edit <policy_number>
        set intf <interface>
        set srcaddr <source_address> [source_address] ...
        set dstaddr <destination_address> [destination_address] ...
        set action {accept | deny}
        set service <service_name> [service_name] ...
        set schedule <schedule_name>
        set comments <string>
    next
end
```

For example, to prevent the source subnet 10.10.10.0/24 from pinging port1, but allow administrative access for PING on port1:

```
config firewall address
    edit "10.10.10.0"
        set subnet 10.10.10.0 255.255.255.0
    next
end
config firewall local-in-policy
    edit 1
        set intf "port1"
        set srcaddr "10.10.10.0"
        set dstaddr "all"
        set service "PING"
        set schedule "always"
    next
end
```

To test the configuration:

1. From the PC at 10.10.10.12, start a continuous ping to port1:

   ```
   ping 192.168.2.5 -t
   ```

2. On the FortiGate, enable debug flow:

   ```
   # diagnose debug flow filter addr 10.10.10.12
   # diagnose debug flow filter proto 1
   # diagnose debug enable
   # diagnose debug flow trace start 10
   ```

3. The output of the debug flow shows that traffic is dropped by local-in policy 1:

Netskope Exhibit 1017

```
# id=20085 trace_id=1 func=print_pkt_detail line=5746 msg="vd-root:0 received a packet
(proto=1, 10.10.10.12:1->192.168.2.5:2048) from port1. type=8, code=0, id=1, seq=128."
id=20085 trace_id=1 func=init_ip_session_common line=5918 msg="allocate a new session-
0017c5ad"
id=20085 trace_id=1 func=vf_ip_route_input_common line=2615 msg="find a route:
flag=80000000 gw-192.168.2.5 via root"
id=20085 trace_id=1 func=fw_local_in_handler line=474 msg="iprope_in_check() check
failed on policy 1, drop"
```

Additional options

To disable or re-enable the local-in policy, use the `set status {enable | disable}` command.

To dedicate the interface as an HA management interface, use the `set ha-mgmt-intf-only enable` command.

IPv4/IPv6 access control lists

Access control lists (ACL) in the FortiOS firmware is a granular or more specifically targeted blocklist. ACL drop IPv4 and IPv6 packets at the physical network interface before the packets are analyzed by the CPU. On a busy appliance, this can really improve performance.

ACL is available on FortiGates with NP6-accelerated interfaces. ACL checking is one of the first things that happens to the packet and checking is done by the NP6 processor. The result is very efficient protection that does not use CPU or memory resources.

The following platforms support ACL:

- FGT_100D, FGT_100E, FGT_100EF, FGT_101E.
- FGT_140D, FGT_140D_POE, FGT_140E, FGT_140E_POE.
- FGT_301E, FGT_500E, FGT_501E.
- FGT_1200D, FGT_1500D, FGT_1500DT.
- FGT_2000E, FGT_2500E.
- FGT_3000D, FGT_3100D, FGT_3200D, FGT_3700D.
- FGT_3800D, FGT_3810D, FGT_3815D.
- FGT_3960E, FGT_3980E.

Limitation

The configuration of ACL allows you to specify which interface the ACL is applied to. You should be aware of a hardware limitation. The ACL is a Layer 2 function and is offloaded to the ISF hardware. Therefore, no CPU resources are used in the processing of the ACL. It is handled by the inside switch chip which can do hardware acceleration, which increases the performance of the FortiGate. The drawback is that the ACL function is only supported on switch fabric driven interfaces. It also cannot be applied to hardware switch interfaces or their members. Ports such as WAN1 or WAN2 on some models that use network cards that connect to the CPU through a PCIe bus do not support ACL.

Netskope Exhibit 1017

Sample configuration

To block all IPv4 and IPv6 Telnet traffic from port2 to Company_Servers using the CLI:

```
config firewall acl
    edit 1
       set interface "port2"
       set srcaddr "all"
       set dstaddr "Company_Servers"
       set service "TELNET"
    next
end
config firewall acl6
    edit 1
        set interface "port2"
        set srcaddr "all"
        set dstaddr "Company_Servers_v6"
        set service "TELNET"
    next
end
```

Sample troubleshooting

To check the number of packets drop by an ACL:

```
# diagnose firewall acl counter
ACL id 1 dropped 0 packets
```

To clear the packet drop counter:

```
# diagnose firewall acl clearcounter
```

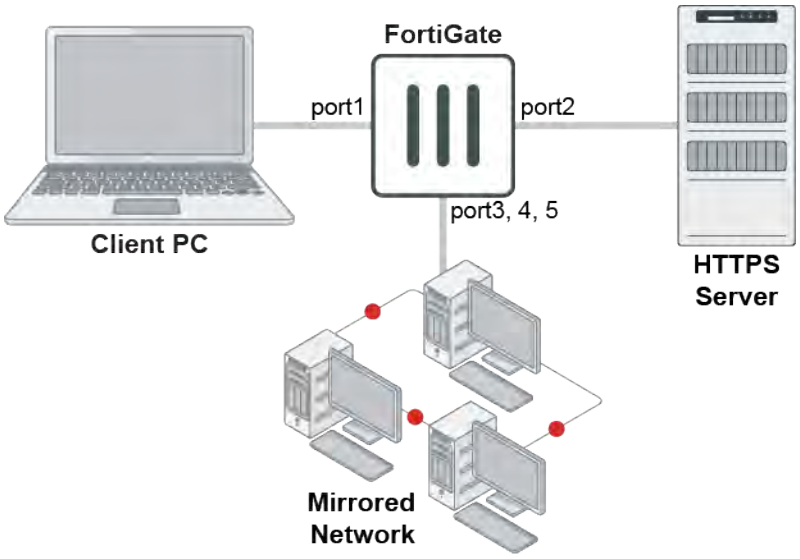
Use the same commands for IPv6 ACL.

```
# diagnose firewall acl
counter          Show number of packets dropped by ACL.
counter6         Show number of packets dropped by ACL6.
clearcounter     Clear ACL packet counter.
clearcounter6    Clear ACL6 packet counter.
```

Mirroring SSL traffic in policies

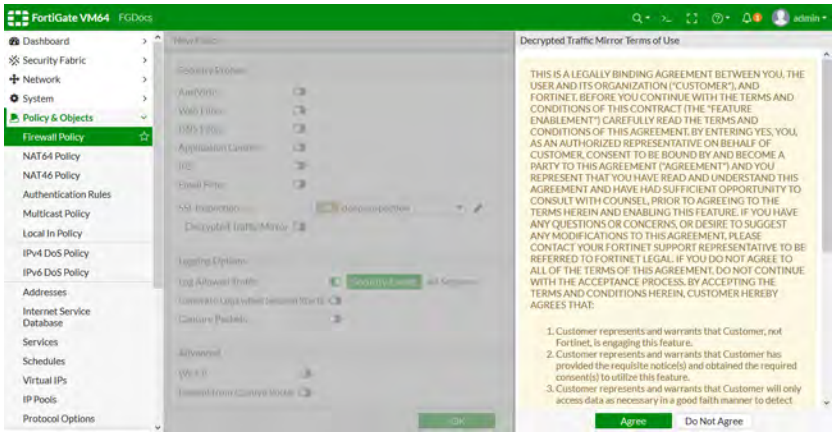
SSL mirroring allows the FortiGate to decrypt and mirror traffic to a designated port. A new decrypted traffic mirror profile can be applied to IPv4, IPv6, and explicit proxy firewall policies in both flow and proxy mode. Full SSL inspection must be used in the policy for the traffic mirroring to occur.

SSL inspection is automatically enabled when you enable a security profile on the policy configuration page.



To configure SSL mirroring in a policy in the GUI:

1. Go to *Policy & Objects* > *Firewall Policy*.
2. Create a new policy, or edit an existing one.
3. Configure the interfaces, sources, and other required information.
4. In the *Security Profiles* section, for *SSL Inspection*, select *deep-inspection*, or another profile that uses *Full SSL Inspection*.
5. Enable *Decrypted Traffic Mirror* . The terms of use agreement opens.



6. Click *Agree* to accept the terms.
7. In the drop-down list, select a decrypted traffic mirror, or click *Create* to create a new one.
   In this example, a new decrypted traffic mirror is created using the port3 interface.



8. Click *OK* to save the policy.

To configure SSL mirroring in proxy mode in the CLI:

1. Create the decrypted traffic mirror profile:

```
config firewall decrypted-traffic-mirror
    edit SSL-to-port3
        set dstmac ff:ff:ff:ff:ff:ff
        set traffic-type ssl
        set traffic-source client
        set interface port3
    next
end
```

2. Configure the policy to enable SSL traffic mirroring:

```
config firewall policy
    edit 1
        set name "mirror-policy"
        set srcintf "port1"
        set dstintf "port2"
        set srcaddr "all"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ALL"
        set nat enable
        set ssl-ssh-profile "deep-inspection"
        set decrypted-traffic-mirror "SSL-to-port3"

THIS IS A LEGALLY BINDING AGREEMENT BETWEEN YOU, THE USER AND ITS ORGANIZATION
("CUSTOMER"), AND FORTINET. BEFORE YOU CONTINUE WITH THE TERMS AND CONDITIONS OF THIS
CONTRACT (THE "FEATURE ENABLEMENT") CAREFULLY READ THE TERMS AND CONDITIONS OF THIS
AGREEMENT. BY ENTERING YES, YOU, AS AN AUTHORIZED REPRESENTATIVE ON BEHALF OF CUSTOMER,
CONSENT TO BE BOUND BY AND BECOME A PARTY TO THIS AGREEMENT ("AGREEMENT") AND YOU
REPRESENT THAT YOU HAVE READ AND UNDERSTAND THIS AGREEMENT AND HAVE HAD SUFFICIENT
OPPORTUNITY TO CONSULT WITH COUNSEL, PRIOR TO AGREEING TO THE TERMS HEREIN AND ENABLING
THIS FEATURE. IF YOU HAVE ANY QUESTIONS OR CONCERNS, OR DESIRE TO SUGGEST ANY
MODIFICATIONS TO THIS AGREEMENT, PLEASE CONTACT YOUR FORTINET SUPPORT REPRESENTATIVE TO
BE REFERRED TO FORTINET LEGAL. IF YOU DO NOT AGREE TO ALL OF THE TERMS OF THIS
AGREEMENT, DO NOT CONTINUE WITH THE ACCEPTANCE PROCESS. BY ACCEPTING THE TERMS AND
CONDITIONS HEREIN, CUSTOMER HEREBY AGREES THAT:

1. Customer represents and warrants that Customer, not Fortinet, is engaging this
feature.

2. Customer represents and warrants that Customer has provided the requisite notice(s)
and obtained the required consent(s) to utilize this feature.

3. Customer represents and warrants that Customer will only access data as necessary in
a good faith manner to detect malicious traffic and will put in place processes and
controls to ensure this occurs.

4. Customer represents and warrants that Customer has the right to enable and utilize
this feature, and Customer is fully in compliance with all applicable laws in so doing.

5. Customer shall indemnify Fortinet in full for any of the above certifications being
untrue.
```


Netskope Exhibit 1017

```
6. Customer shall promptly notify Fortinet Legal in writing of any breach of these Terms
and Conditions and shall indemnify Fortinet in full for any failure by Customer or any
of its employees or representatives to abide in full by the Terms and Conditions above.

7. Customer agrees that these Terms and Conditions shall be governed by the laws of the
State of California, without regards to the choice of laws provisions thereof and
Customer hereby agrees that any dispute related to these Terms and Conditions shall be
resolved in Santa Clara County, California, USA, and Customer hereby consents to
personal jurisdiction in Santa Clara County, California, USA.

Do you want to continue? (y/n) y
    next
end
```

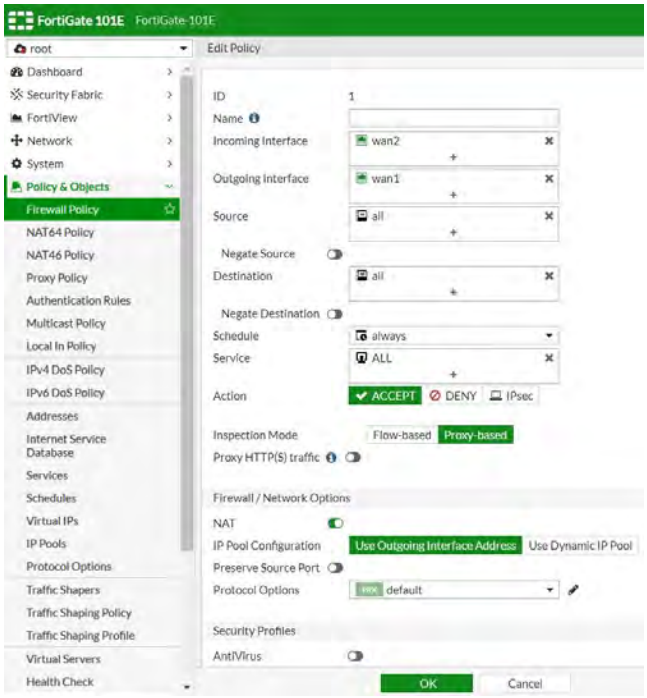
Inspection mode per policy

Inspection mode is configured on a per-policy basis in NGFW mode. This gives you more flexibility when setting up different policies.

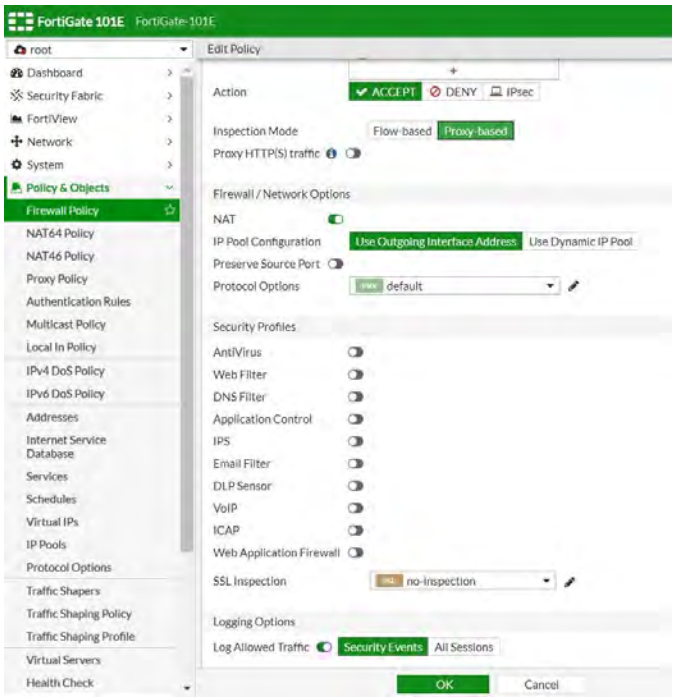
When configuring a firewall policy, you can select a *Flow-based* or *Proxy-basedInspection Mode*. The default setting is *Flow-based*.

To configure inspection mode in a policy:

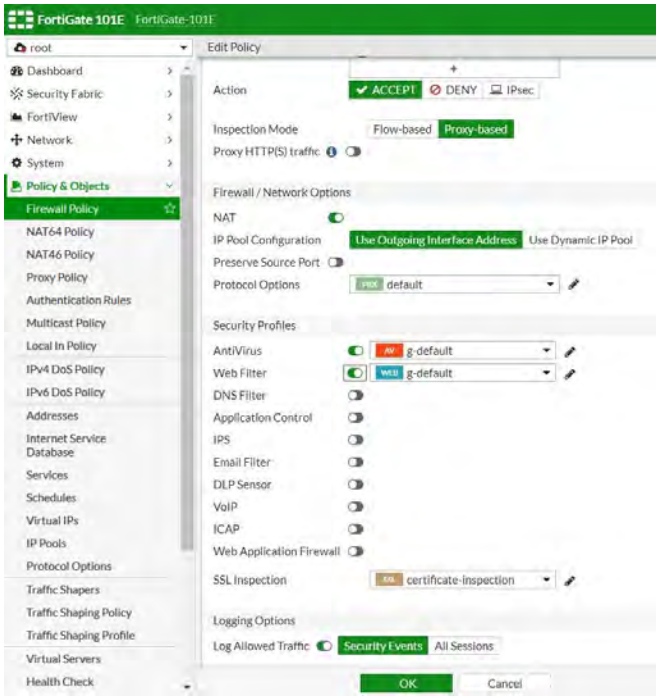
1. Go to *Policy & Objects* > *Firewall Policy*.
2. Create a new policy, or edit an existing policy.
3. Configure the policy as needed.
   a. If you change the *Inspection Mode* to *Proxy-based*, the *Proxy HTTP(S) traffic* option displays.



b. In the *Security Profiles* section, if no security profiles are enabled, the default *SSL Inspection* is *no-inspection*.



c. In the *Security Profiles* section, if you enable any security profile, the *SSL Inspection* changes to *certificate-inspection*.



To see the inspection mode changes using the CLI:

```
config firewall policy
    edit 1
```

```
        set srcintf "wan2"
        set dstintf "wan1"
        set srcaddr "all"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ALL"
        set inspection-mode proxy
        set nat enable
    next
end
```

To see the HTTP and SSH policy redirect settings when inspection mode is set to proxy using the CLI:

```
config firewall policy
    edit 1
        set srcintf "wan2"
        set dstintf "wan1"
        set srcaddr "all"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ALL"
        set inspection-mode proxy
        set http-policy-redirect enable
        set ssh-policy-redirect enable
        set nat enable
    next
end
```

To see the default SSL-SSH policy set to no inspection using the CLI:

```
config firewall policy
    edit 1
       show fu | grep ssl-ssh-profile
        set ssl-ssh-profile "no-inspection"
    next
end
```

OSPFv3 neighbor authentication

OSPFv3 neighbor authentication is available for enhanced IPv6 security.

To configure an OSPF6 interface:

```
config router ospf6
    config ospf6-interface
        edit <name>
            set authentication {none | ah | esp | area}
            set key-rollover-interval <integer>
            set ipsec-auth-alg {md5 | sha1 | sha256 | sha384 | sha512}
            set ipsec-enc-alg {null | des | 3des | aes128 | aes192 | aes256}
            config ipsec-keys
                edit <spi>
```

```
                    set auth-key <string>
                    set enc-key <string>
                next
            end
        next
    end
end
```

To configure an OSPF6 virtual link:

```
config router ospf6
    config area
        edit <id>
            config virtual-link
                edit <name>
                    set authentication {none | ah | esp | area}
                    set key-rollover-interval <integer>
                    set ipsec-auth-alg {md5 | sha1 | sha256 | sha384 | sha512}
                    set ipsec-enc-alg {null | des | 3des | aes128 | aes192 | aes256}
                    config ipsec-keys
                        edit <spi>
                            set auth-key <string>
                            set enc-key <string>
                        next
                    end
                next
            end
        next
    end
end
```

To configure an OSPF6 area:

```
config router ospf6
    config area
        edit <id>
            set authentication {none | ah | esp}
            set key-rollover-interval <integer>
            set ipsec-auth-alg {md5 | sha1 | sha256 | sha384 | sha512}
            set ipsec-enc-alg {null | des | 3des | aes128 | aes192 | aes256}
            config ipsec-keys
                edit <spi>
                    set auth-key <string>
                    set enc-key <string>
                next
            end
        next
    end
end
```


Netskope Exhibit 1017

CLI command descriptions

Command	Description
<id>	Area entry IP address.
authentication {none | ah | esp | area}	Authentication mode: • `none`: Disable authentication • `ah`: Authentication Header • `esp`: Encapsulating Security Payload • `area`: Use the routing area authentication configuration
key-rollover-interval <integer>	Enter an integer value (300 - 216000, default = 300).
ipsec-auth-alg {md5 | sha1 | sha256 | sha384 | sha512}	Authentication algorithm.
ipsec-enc-alg {null | des | 3des | aes128 | aes192 | aes256}	Encryption algorithm.
<spi>	Security Parameters Index.
auth-key <string>	Authentication key should be hexadecimal numbers. Key length for each algorithm: • MD5: 16 bytes • SHA1: 20 bytes • SHA256: 32 bytes • SHA384:48 bytes • SHA512:84 bytes If the key is shorter than the required length, it will be padded with zeroes.
enc-key <string>	Encryption key should be hexadecimal numbers. Key length for each algorithm: • DES: 8 bytes • 3DES: 24 bytes • AES128: 16 bytes • AES192: 24 bytes • AES256: 32 bytes If the key is shorter than the required length, it will be padded with zeroes.

Firewall anti-replay option per policy

When the global anti-replay option is disabled, the FortiGate does not check TCP flags in packets. The per policy anti-replay option overrides the global setting. This allows you to control whether or not TCP flags are checked per policy.

To enable the anti-replay option so TCP flags are checked using the CLI:

```
config firewall policy
   edit 1
      set name "policyid-1"
      set srcintf "wan2"
      set dstintf "wan1"
```

```
        set srcaddr "all"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ALL"
        set anti-replay enable
        set logtraffic all
        set nat enable
    next
end
```

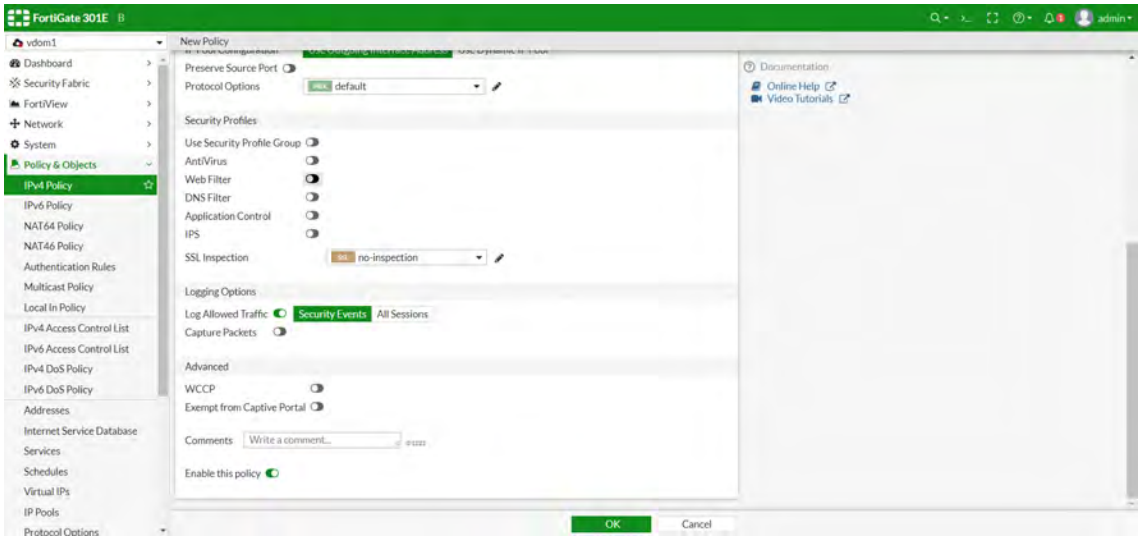
Enabling advanced policy options in the GUI

Advanced policy options can be enabled so you can configure the options in the GUI.

To enable advanced policy options:

```
config system settings
    set gui-advanced-policy enable
end
```

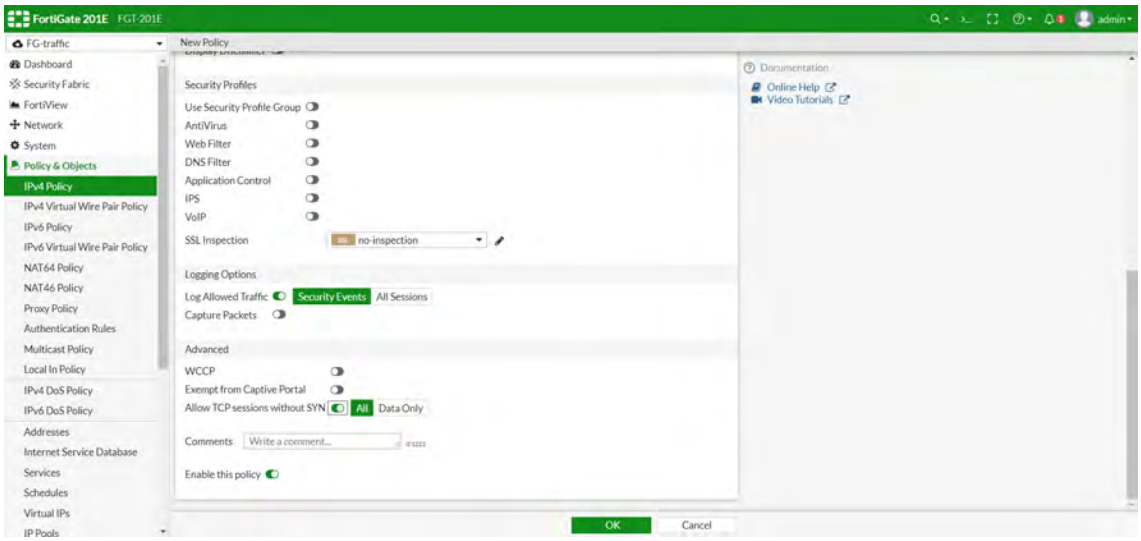
Advanced policy options are now available when creating or editing a policy in the GUI:



To enable configuring TCP sessions without SYN:

```
config system settings
    set tcp-session-without-syn enable
end
```

TCP sessions without SYN can now be configured when creating or editing a policy in the GUI:



Recognize anycast addresses in geo-IP blocking

An anycast IP can be advertised from multiple locations and the router selects a path based on latency, distance, cost, number of hops, and so on. This technique is widely used by providers to route users to the closest server. Since the IP is hosted in multiple geographic locations, there is no way to specify one single location to that IP.

In FortiOS 6.4.0, there is an option to bypass anycast IP ranges in geo-IP blocking. The ISDB contains a list of confirmed anycast IP ranges that can be used for this purpose.

When the source or destination is set to `geoip`, you can enable the `geoip-anycast` option. Once enabled, IPs where the anycast option is set to `1` in `geoip_db` are bypassed in country matching and blocking.



You can only use the CLI to configure this feature.

To enable the geoip-anycast option using the CLI:

```
config firewall policy
   edit 1
      set name "policyid-1"
      set srcintf "wan2"
      set dstintf "wan1"
      set srcaddr "all"
      set dstaddr "test-geoip-CA_1"
      set action accept
      set schedule "always"
      set service "ALL"
      set geoip-anycast enable
      set logtraffic all
      set nat enable
   next
end
```

To check the geoip-anycast option for an IP address using the CLI:

```
diagnose geoip ip2country 1.0.0.1
   1.0.0.1 - Australia, is anycast ip
```

The anycast IP is `1.0.0.1`.

Matching GeoIP by registered and physical location

IP addresses have both a physical and registered location in the geography IP database. Sometimes these two locations are different. The `geoip-match` command allows users to match an IPv4 address in an firewall policy to its physical or registered location when a GeoIP is used as a source or destination address. IPv6 policies currently support geography address objects but do not support `geoip-match`.

In the following example, the physical location of 220.243.219.10 is CA (Canada), the registered location is CN (China), and it is not an anycast IP.

To configure GeoIP matching based on registered location:

1. Create a firewall policy to match the IP:

```
config firewall policy
    edit 1
        set name "policy_id_1"
        set srcintf "wan2"
        set dstintf "wan1"
        set srcaddr "all"
        set dstaddr "test-geoip-CA"
        set action accept
        set schedule "always"
        set service "ALL"
        set geoip-match registered-location
        set logtraffic all
        set auto-asic-offload disable
        set nat enable
    next
end
```

   Since CA is applied as a destination address and registered location IP matching is enabled, if the destination IP of the traffic is 220.243.219.10, then the traffic will be blocked because the registered location is CN.

2. Verify that the policy is blocking traffic from the IP address:

```
# diagnose sniffer packet any icmp 4
interfaces=[any]
filters=[icmp]
5.383798 wan2 in 10.1.100.41 -> 220.243.219.10: icmp: echo request
6.381982 wan2 in 10.1.100.41 -> 220.243.219.10: icmp: echo request
7.382608 wan2 in 10.1.100.41 -> 220.243.219.10: icmp: echo request
^C
3 packets received by filter
0 packets dropped by kernel
```

To configure GeoIP matching based on physical location:

1. Create a firewall policy to match the IP:

```
config firewall policy
    edit 1
        set name "policy_id_1"
        set srcintf "wan2"
        set dstintf "wan1"
        set srcaddr "all"
        set dstaddr "test-geoip-CA"
        set action accept
        set schedule "always"
        set service "ALL"
        set geoip-match physical-location
        set logtraffic all
        set auto-asic-offload disable
        set nat enable
    next
end
```

   Since CA is applied as a destination address and physical location IP matching is enabled, if the destination IP of the traffic is 220.243.219.10, then the traffic will pass through.

2. Verify that the policy is allowing traffic from the IP address:

```
# diagnose sniffer packet any icmp 4
interfaces=[any]
filters=[icmp]
5.273985 wan2 in 10.1.100.41 -> 220.243.219.10: icmp: echo request
5.274176 wan1 out 172.16.200.10 -> 220.243.219.10: icmp: echo request
6.274426 wan2 in 10.1.100.41 -> 220.243.219.10: icmp: echo request
6.274438 wan1 out 172.16.200.10 -> 220.243.219.10: icmp: echo request
7.273978 wan2 in 10.1.100.41 -> 220.243.219.10: icmp: echo request
7.273987 wan1 out 172.16.200.10 -> 220.243.219.10: icmp: echo request
^C
6 packets received by filter
0 packets dropped by kernel
```

Authentication policy extensions

By default, unauthenticated traffic is permitted to fall to the next policy. This means that unauthenticated users are only forced to authenticate against a policy when there are no other matching policies. To avoid this, you can force authentication to always take place.

To set that authentication requirement:

```
config user setting
    set auth-on-demand {always | implicitly}
end
```

Where:

`always`	Always trigger firewall authentication on demand.

`implicitly` (default)	Implicitly trigger firewall authentication on demand. This is the default setting (and the behavior in FortiOS 6.0 and earlier).

In the following example, authentication is required; traffic that would otherwise be allowed by the second policy is instead blocked by the first policy.

To use forced authentication:

```
config user setting
    set auth-on-demand always
end
config firewall policy
    edit 1
        set name "QA to Database"
        set srcintf "port10"
        set dstintf "port9"
        set srcaddr "QA_subnet"
        set dstaddr "Database"
        set action accept
        set schedule "always"
        set service "ALL"
        set fsso disable
        set groups "qa_group"
        set nat enable
    next
    edit 2
        set name "QA to Internet"
        set srcintf "port10"
        set dstintf "port9"
        set srcaddr "QA_subnet"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ALL"
        set fsso disable
        set nat enable
    next
end
```

HTTP to HTTPS redirect for load balancing

You can configure a virtual server with HTTP to HTTPS redirect enabled. When enabled, a virtual server can convert a client's HTTP requests to HTTPS requests. Through this mandatory conversion, HTTP traffic is converted to HTTPS traffic. This conversion improves the security of the user network.

You can only enable this feature by using the CLI. After you enable this feature, traffic flows as follows:

- When FortiGate receives an HTTP request for an external IP, such as 10.1.100.201 in the following example, FortiGate sends an HTTP 303 response back to the original client and redirects HTTP to HTTPS, instead of forwarding the HTTP request to the real backend servers.
- The client browser restarts the TCP session to HTTPS.
- The HTTPS session comes to the FortiGate where a matching firewall policy allows the HTTPS traffic and establishes a secure SSL connection, and then forwards the request to the real backend servers.

Netskope Exhibit 1017

To configure virtual server with HTTPS redirect enabled:

1. Create a virtual server with `server-type` set to `http`:
```
config firewall vip
   edit "virtual-server-http"
      set type server-load-balance
      set extip 10.1.100.201
      set extintf "wan2"
      set server-type http
      set ldb-method round-robin
      set extport 80
      config realservers
         edit 1
            set ip 172.16.200.44
            set port 80
         next
         edit 2
            set ip 172.16.200.55
            set port 80
         next
      end
   next
end
```
2. Create a virtual server with `server-type` set to `https` and with the same external IP address:
```
config firewall vip
   edit "virtual-server-https"
      set type server-load-balance
      set extip 10.1.100.201
      set extintf "wan2"
      set server-type https
      set ldb-method round-robin
      set extport 443
         config realservers
            edit 1 set ip 172.16.200.44
            set port 443
         next
         edit 2
            set ip 172.16.200.55
            set port 443
         next
      end
      set ssl-certificate "Fortinet_CA_SSL"
   next
end
```
3. Enable the `http-redirect` option for the virtual server with `server-type` set to `http`:
```
config firewall vip
   edit "virtual-server-http"
      set http-redirect enable
   next
end
```
4. Add the two virtual servers to a policy:
```
config firewall policy
   edit 9
      set srcintf "wan2"
      set dstintf "wan1"
      set srcaddr "all"
```


Netskope Exhibit 1017

```
        set dstaddr "virtual-server-http" "virtual-server-https"
        set action accept
        set schedule "always"
        set service "ALL"
        set inspection-mode proxy set logtraffic all
        set auto-asic-offload disable
        set nat enable
      next
    end
```

GTPv2 in policies

You can use GTPv2 in the policy section of a GTP profile.

GTPv2 is only available for the following advanced fields:

- `apnmember`
- `apn-sel-mode`
- `messages`
- `max-apn-restriction`
- `imsi-prefix`
- `msisdn-prefix`
- `rat-type`
- `mei`
- `uli`

GTPv2 support includes the following changes for overall GTP support:

- `rai` is no longer supported in any GTP version.
- `uli` can coexist with CGI, SAI, RAI, TAI, ECGI, or LAI, each of which has the pattern MCC.MNC.ID or MCC.MNC.ID.ID2.
- `mei` can take IMEI (15 digits) or IMEISV (16 digits).

To configure a new GTPv2 policy:

```
config firewall gtp
    edit "gtpv2"
        config policy-v2
            edit 1
                set messages create-ses-req
            next
        end
    next
end
```

To configure the new ULI format:

This example matches packets with TAI 510-519.01-09.d02a and ECGI 505.02.1409900-14099ff.

```
config firewall gtp
    edit "gtpv2"
        config policy-v2
            edit 1
```

```
                set messages create-ses-req
                set uli "0" "0" "0" "51*.01-09.d02a" "505.02.14099*"
            next
        end
    next
end
```

To configure the GTPv2 message type:

```
config firewall gtp
    edit "gtpv2"
        config policy-v2
            edit 1
                set messages create-ses-req create-ses-res modify-bearer-req modify-bearer-
res
            next
        end
    next
end
```

To configure the RAT type:

This example includes virtual and NB-IoT types.

```
config firewall gtp
    edit "gtpv2"
        config policy-v2
            edit 1
                set messages create-ses-req
                set rat-type virtual nbiot
            next
        end
    next
end
```

Use active directory objects directly in policies

Active Directory (AD) groups can be used directly in identity-based firewall policies. You do not need to add remote AD groups to local FSSO groups before using them in policies.

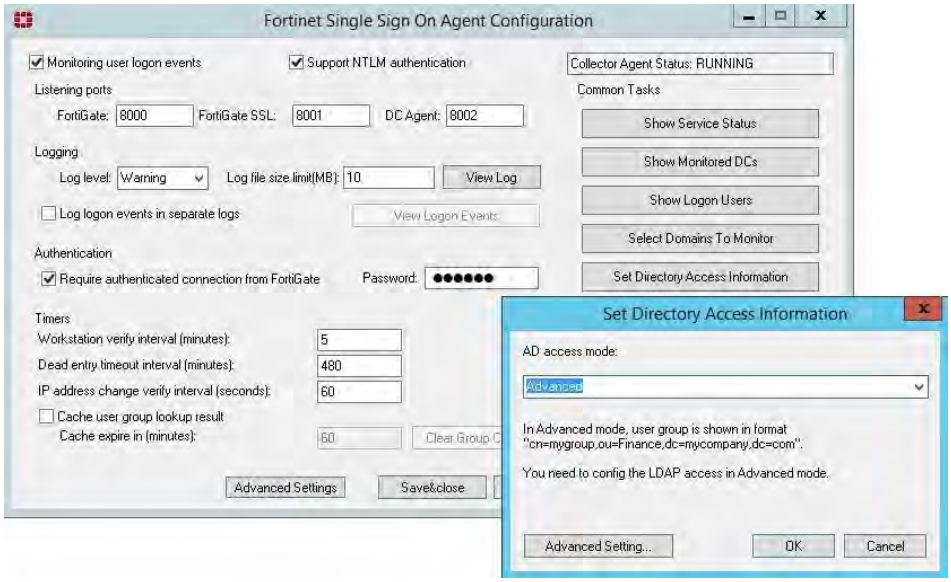
FortiGate administrators can define how often group information is updated from AD LDAP servers.

To retrieve and use AD user groups in policies:

1. Set the FSSO Collector Agent AD access mode on page 785
2. Add an LDAP server on page 785
3. Create the FSSO collector that updates the AD user groups list on page 786
4. Use the AD user groups in a policy on page 787

Set the FSSO Collector Agent AD access mode

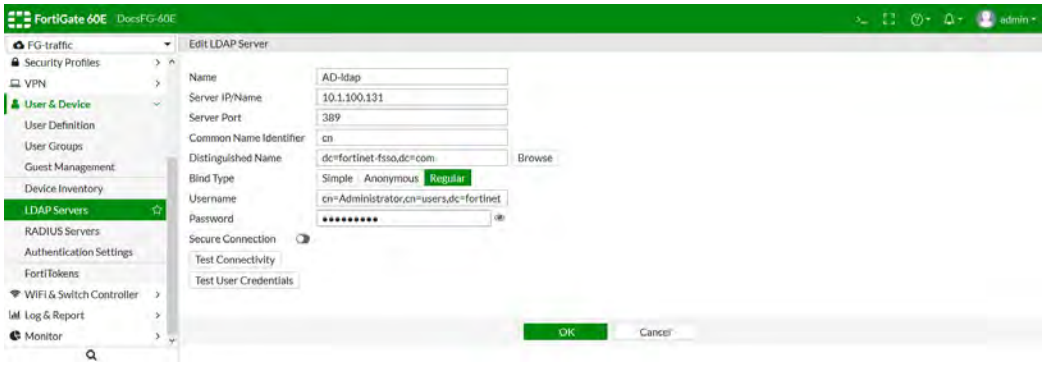
To use this feature, you must set FSSO Collector Agent to *Advanced* AD access mode. If the FSSO Collector Agent is running in the default mode, FortiGate cannot correctly match user group memberships.



Add an LDAP server

To add an LDAP server in the GUI:

1. Go to *User & Authentication* > *LDAP Servers*.
2. Click *Create New*.
3. Configure the settings as needed.



4. If secure communication over TLS is supported by the remote AD LDAP server:
   a. Enable *Secure Connection* .
   b. Select the protocol.
   c. Select the certificate from the CA that issued the AD LDAP server certificate.
   If the protocol is LDAPS, the port will automatically change to 636.
5. Click *OK*.

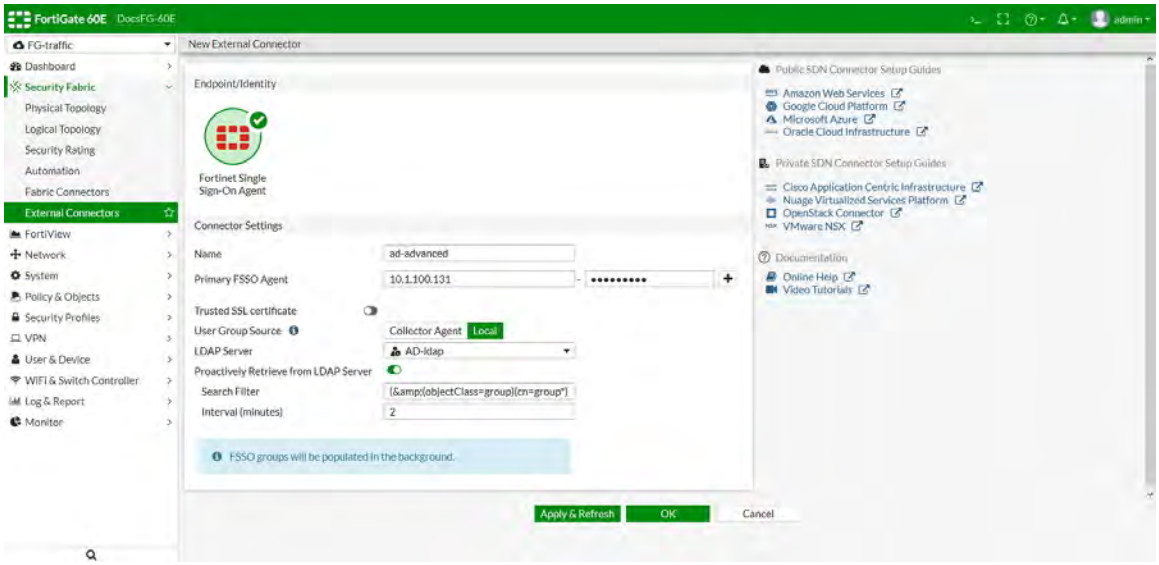
To add an LDAP server in the CLI:

```
config user ldap
    edit "AD-ldap"
        set server "10.1.100.131"
        set cnid "cn"
        set dn "dc=fortinet-fsso,dc=com"
        set type regular
        set username "cn=Administrator,cn=users,dc=fortinet-fsso,dc=com"
        set password XXXXXXXXXXXXXXXXXXXXXXXX
    next
end
```

Create the FSSO collector that updates the AD user groups list

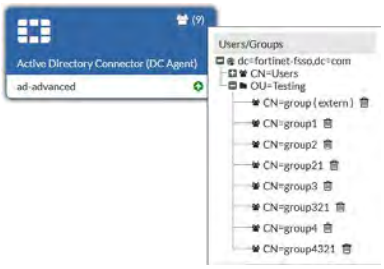
To create an FSSO agent connector in the GUI:

1. Go to *Security Fabric* > *External Connectors*.
2. Click *Create New*.
3. In the *Endpoint/Identity* section, click *Fortinet Single Sign-On Agent*.
4. Fill in the *Name*
5. Set the *Primary FSSO Agent* to the IP address of the FSSO Collector Agent, and enter its password.
6. Set the *User Group Source* to *Local*.
7. Set the *LDAP Server* to the just created *AD-ldap* server.
8. Enable *Proactively Retrieve from LDAP Server*.
9. Set the *Search Filter* to *(&(objectClass=group)(cn=group*))*.
   The default search filter retrieves all groups, including Microsoft system groups. In this example, the filter is configured to retrieve *group1*, *group2*, etc, and not groups like *grp199*.
   The filter syntax is not automatically checked; if it is incorrect, the FortiGate might not retrieve any groups.
10. Set the *Interval (minutes)* to configure how often the FortiGate contacts the remote AD LDAP server to update the group information.



11. Click *OK*.

12. To view the AD user groups that are retrieved by the FSSO agent, hover the cursor over the group icon on the fabric connector listing.



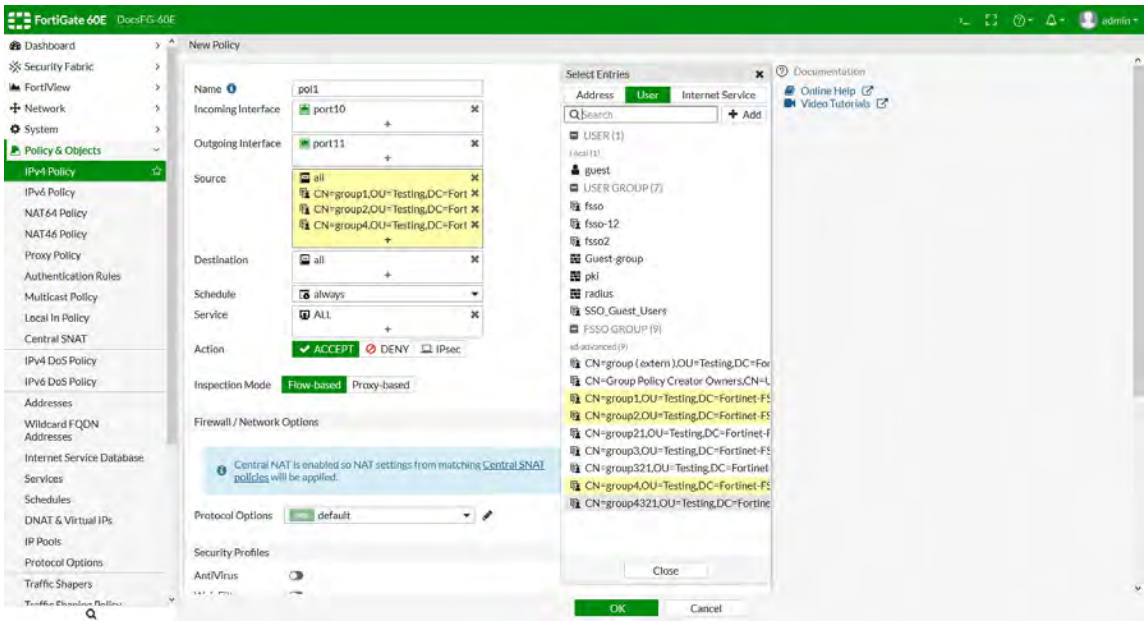
To create an FSSO agent connector in the CLI:

```
config user fsso
    edit "ad-advanced"
        set server "10.1.100.131"
        set password XXXXXXXXXXXXXX
        set ldap-server "AD-ldap"
        set ldap-poll enable
        set ldap-poll-interval 2
        set ldap-poll-filter "(&amp;(objectClass=group)(cn=group*))"
    next
end
```

You view the retrieved AD user groups with the `show user adgrp` command.

Use the AD user groups in a policy

The AD user groups retrieved by the FortiGate can be used directly in firewall policies.

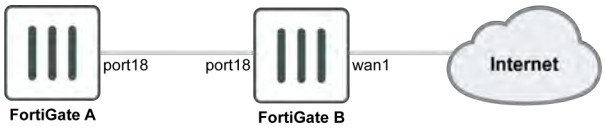


Netskope Exhibit 1017

FortiGate Cloud / FDN communication through an explicit proxy

Explicit proxy communication to FortiGate Cloud and FortiGuard servers from FortiGate is enabled. A proxy server can be configured in the FortiGuard settings so that all FortiGuard connections under the `forticldd` process can be established through the proxy server.

Not all FortiGuard services are supported by these proxy settings. For example, web filter service traffic to FortiGuard will not be directed to the configured proxy.



To configure a proxy server and communicate with FortiGate Cloud though it:

1. Configure FortiGate B as a proxy server:

```
config firewall proxy-policy
    edit 1
        set proxy explicit-web
        set dstintf "wan1"
        set srcaddr "all"
        set dstaddr "all"
        set service "webproxy"
        set action accept
        set schedule "always"
        set logtraffic all
        set users "guest1"
    next
end
config user local
    edit "guest1"
        set type password
        set passwd 123456
    next
end
config authentication scheme
    edit "local-basic"
        set method basic
        set user-database "local-user-db"
    next
end
config authentication rule
    edit "local-basic-rule"
        set srcaddr "all"
        set ip-based disable
        set active-auth-method "local-basic"
    next
end
```

2. Configure a firewall policy on FortiGate B to allow FortiGate A to get DNS resolution:

```
config firewall policy
    edit 1
        set name "dns"
        set srcintf "port18"
        set dstintf "wan1"
        set srcaddr "all"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "DNS"
        set fsso disable
        set nat enable
    next
end
```

3. Configure the FortiGuard proxy settings on FortiGate A:

```
config system fortiguard
    set proxy-server-ip 10.2.2.2
    set proxy-server-port 8080
    set proxy-username "guest1"
    set proxy-password 123456
end
```

4. On FortiGate A, log in to FortiGate Cloud to activate the logging service:

```
execute fortiguard-log login <username> <password>
```

5. On FortiGate A, view the `forticldd` debug message to see the connection to the log controller through the proxy server:

```
# diagnose test application forticldd 1
```

No session timeout

To allow clients to permanently connect with legacy medical applications and systems that do not have keepalive or auto-reconnect features, the session timeout can be set to never for firewall services, policies, and VDOMs.

The options to disable session timeout are hidden in the CLI.

To set the session TTL value of a custom service to never:

```
config firewall service custom
    edit "tcp_23"
        set tcp-portrange 23
        set session-ttl never
    next
end
```

To set the session TTL value of a policy to never:

```
config firewall policy
    edit 201
        set srcintf "wan1"
        set dstintf "wan2"
```

```
        set srcaddr "all"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "TCP_8080"
        set logtraffic disable
        set session-ttl never
        set nat enable
    next
end
```

To set the session TTL value of a VDOM to never:

```
config system session-ttl
    set default never
    config port
        edit 1
            set protocol 6
            set timeout never
            set start-port 8080
            set end-port 8080
        next
    end
end
```

To view a session list with the timeout set to never:

```
# diagnose sys session list

session info: proto=6 proto_state=01 duration=9 expire=never timeout=never flags=00000000
sockflag=00000000 sockport=0 av_idx=0 use=3
origin-shaper=
reply-shaper=
per_ip_shaper=
class_id=0 ha_id=0 policy_dir=0 tunnel=/ vlan_cos=0/255
state=log may_dirty f00
statistic(bytes/packets/allow_err): org=2290/42/1 reply=2895/34/1 tuples=2
tx speed(Bps/kbps): 238/1 rx speed(Bps/kbps): 301/2
orgin->sink: org pre->post, reply pre->post dev=18->17/17->18 gwy=172.16.200.55/10.1.100.41
hook=post dir=org act=snat 10.1.100.41:34256->172.16.200.55:23(172.16.200.10:34256)
hook=pre dir=reply act=dnat 172.16.200.55:23->172.16.200.10:34256(10.1.100.41:34256)
pos/(before,after) 0/(0,0), 0/(0,0)
misc=0 policy_id=9 auth_info=0 chk_client_info=0 vd=1
serial=00000b27 tos=ff/ff app_list=0 app=0 url_cat=0
rpdb_link_id = 00000000 ngfwid=n/a
dd_type=0 dd_mode=0
npu_state=0x000001 no_offload
no_ofld_reason:  disabled-by-policy
total session 1
```

Objects

The following topics provide information about objects:

- Address group exclusions on page 791
- MAC addressed-based policies on page 792
- ISDB well-known MAC address list on page 795
- Dynamic policy — fabric devices on page 796
- FSSO dynamic address subtype on page 798
- ClearPass integration for dynamic address objects on page 802
- Group address objects synchronized from FortiManager on page 806
- Using wildcard FQDN addresses in firewall policies on page 808
- IPv6 geography-based addresses on page 810
- Array structure for address objects on page 812

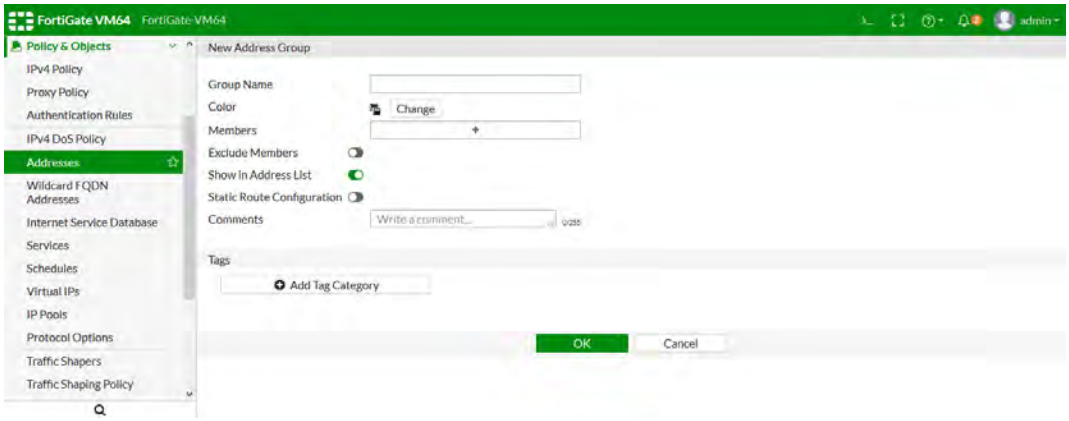
Address group exclusions

Specific IP addresses or ranges can be subtracted from the address group with the *Exclude Members* setting in IPv4 address groups.

This feature is only supported for IPv4 address groups, and only for addresses with a *Type* of *IP Range* or *Subnet*.

To exclude addresses from an address group using the GUI:

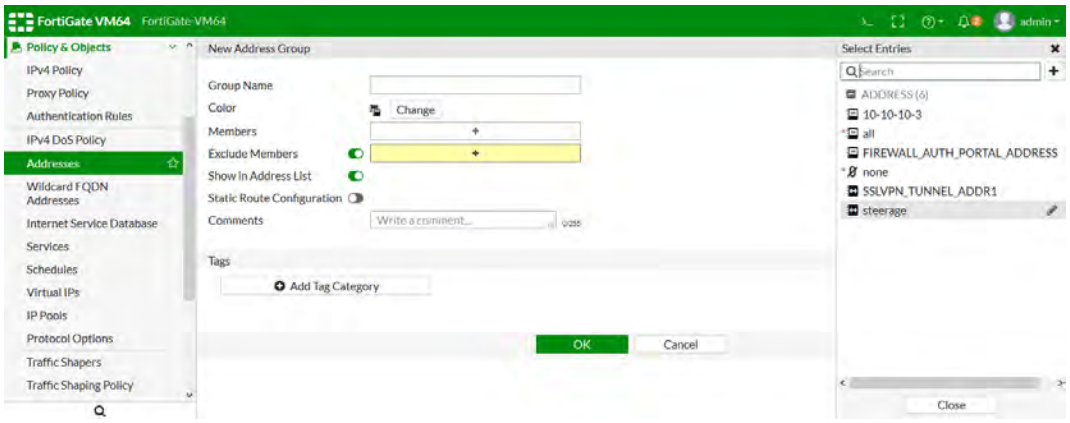
1. Go to *Policy & Objects* > *Addresses*.
2. Create a new address group, or edit an existing address group.



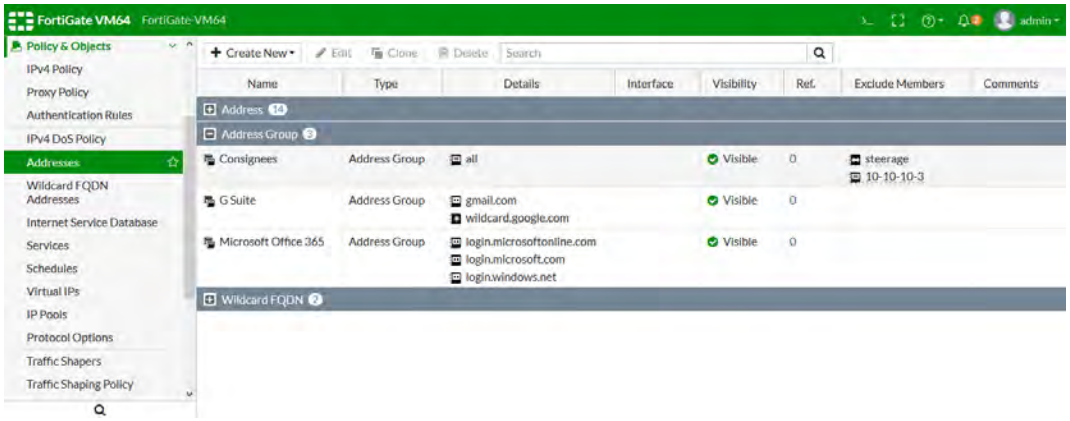
3. Enable *Exclude Members*. The *Select Entries* pane opens.
4. Select the addresses you want to exclude from the group.

Netskope Exhibit 1017

5. Click *OK*.



The excluded members are listed in the *Exclude Members* column.



To exclude addresses from an address group using the CLI:

```
config firewall addrgrp
   edit <address group>
      set exclude enable
      set exclude-member <address> <address> ... <address>
   next
end
```

MAC addressed-based policies

MAC address ranges can be added to the following IPv4 policies:

- Firewall
- Virtual wire pair
- ACL
- Central SNAT
- DoS

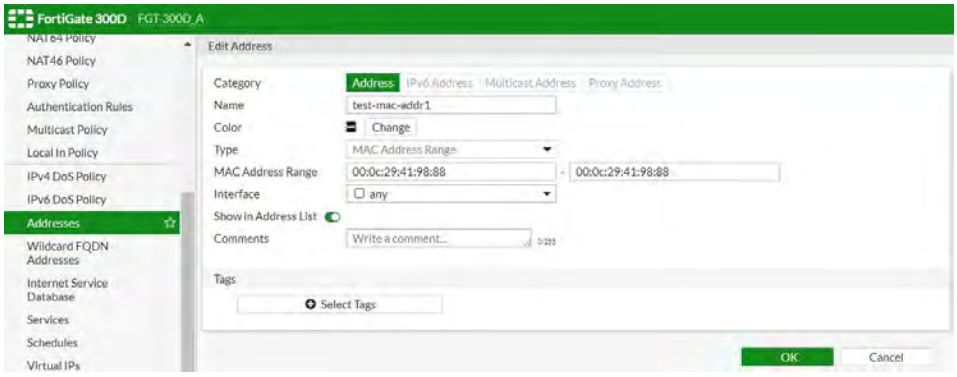
A MAC address is a link layer-based address type and it cannot be forwarded across different IP segments.

FortiOS only supports the MAC address type as source address for policies in NAT mode VDOM. When you use the MAC address type in a policy as source address in NAT mode VDOM, IP address translation (NAT) is still performed according to the rules defined in the policy. The MAC address type only works for source address matching. It does not have any association with NAT actions.

For policies in transparent mode or the virtual wire pair interface, you can use the MAC address type as source or destination address.

To configure a MAC address range using the GUI:

1. Go to *Policy & Objects* > *Addresses* to create or edit an address:
   a. For *Category*, select *Address*.
   b. For *Type*, select *MAC Address Range*.
   c. Enter the address range in the empty fields.
   d. Configure the other fields as needed.
   e. Click *OK*.

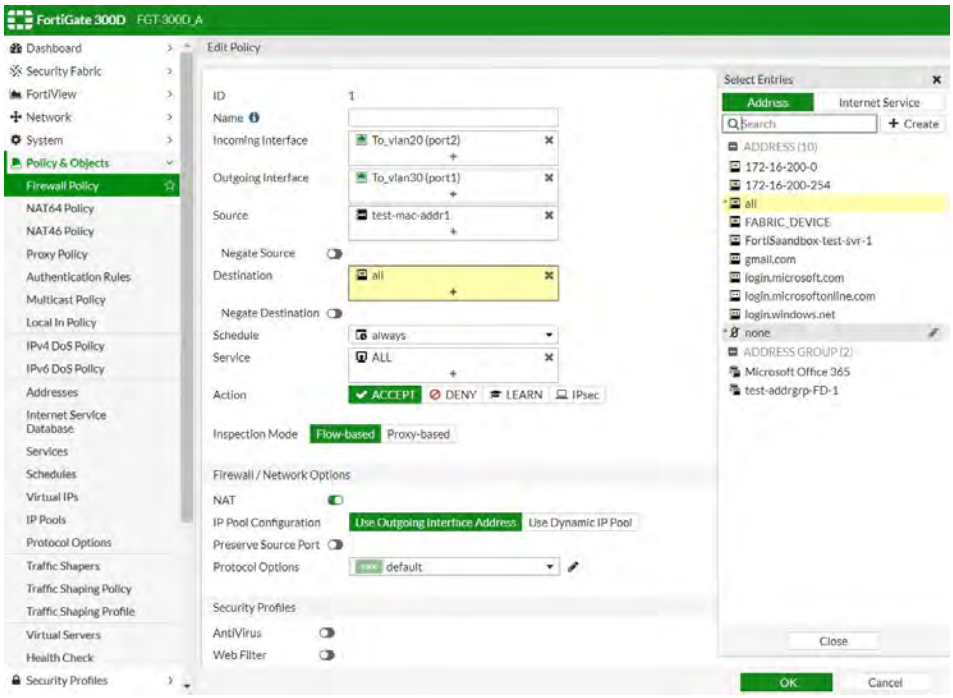


2. Go to *Policy & Objects* > *Firewall Policy* to apply the address type to a policy in NAT mode VDOM:
   a. For *Source*, select the MAC address you just configured.
   b. For *Destination*, select an address.



In NAT mode VDOM, this address type cannot be used as destination address.

c. Click *OK*.



To configure a MAC address range using the CLI:

1. Create a new MAC address range type:

```
config firewall address
    edit <object_name>
        set type mac
        set start-mac <mac_address_start #>
        set end-mac <mac_address_end #>
    next
end
```

2. Apply the address type to a policy. In transparent mode or the virtual wire pair interface, this address type can be mixed with other address types in the policy:

```
config firewall address
    edit "test-mac-addr1"
        set type mac
        set start-mac 00:0c:29:41:98:88
        set end-mac 00:0c:29:41:98:88
    next
end
config firewall policy
    edit 1
        set srcintf "port2"
        set dstintf "port1"
        set srcaddr "test-mac-addr1" "10-1-100-42"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ALL"
```


Netskope Exhibit 1017

```
            set logtraffic all
            set nat enable
        next
    end
```

ISDB well-known MAC address list

The Internet Service Database (ISDB) includes well-known vendor MAC address range lists. The lists can only be used for source MAC addresses in IPv4 policies, and include the vendor name and the MAC address ranges that the vendor belongs to.

To view the vendor list:

```
# diagnose vendor-mac id
Please input Vendor MAC ID.
ID: 1 name: "Asus"
ID: 2 name: "Acer"
ID: 3 name: "Amazon"
ID: 4 name: "Apple"
ID: 5 name: "Xiaomi"
ID: 6 name: "BlackBerry"
ID: 7 name: "Canon"
ID: 8 name: "Cisco"
ID: 9 name: "Linksys"
ID: 10 name: "D-Link"
ID: 11 name: "Dell"
ID: 12 name: "Ericsson"
ID: 13 name: "LG"
ID: 14 name: "Fujitsu"
ID: 15 name: "Fitbit"
ID: 16 name: "Fortinet"
ID: 17 name: "OPPO"
ID: 18 name: "Hitachi"
ID: 19 name: "HTC"
ID: 20 name: "Huawei"
ID: 21 name: "HP"
ID: 22 name: "IBM"
ID: 23 name: "Juniper"
ID: 24 name: "Lenovo"
ID: 25 name: "Microsoft"
ID: 26 name: "Motorola"
ID: 27 name: "Netgear"
ID: 28 name: "Nokia"
ID: 29 name: "Nintendo"
ID: 30 name: "PaloAltoNetworks"
ID: 31 name: "Polycom"
ID: 32 name: "Samsung"
ID: 33 name: "Sharp"
ID: 34 name: "Sony"
ID: 35 name: "Toshiba"
ID: 36 name: "VMware"
ID: 37 name: "Vivo"
ID: 38 name: "Zyxel"
ID: 39 name: "ZTE"
```

To view the MAC address ranges for a vendor:

```
# diagnose vendor-mac id 16
Vendor MAC: 16(Fortinet)
Version: 0000700021
Timestamp: 201908081432
Number of MAC ranges: 6
00:09:0f:00:00:00 - 00:09:0f:ff:ff:ff
04:d5:90:00:00:00 - 04:d5:90:ff:ff:ff
08:5b:0e:00:00:00 - 08:5b:0e:ff:ff:ff
70:4c:a5:00:00:00 - 70:4c:a5:ff:ff:ff
90:6c:ac:00:00:00 - 90:6c:ac:ff:ff:ff
e8:1c:ba:00:00:00 - e8:1c:ba:ff:ff:ff
```

To query the vendor of a specific MAC address or range:

```
# diagnose vendor-mac match 00:09:0f:ff:ff:ff 48
Vendor MAC: 16(Fortinet), matched num: 1
```

To use the vendor ID in a firewall policy:

```
config firewall policy
    edit 9
        set name "policy_id_9"
        set uuid 6150cf30-308d-51e9-a7a3-bcbd05d61f93
        set srcintf "wan2"
        set dstintf "wan1"
        set srcaddr "all"
        set dstaddr "all"
        set vendor-mac 36 16
        set action accept
        set schedule "always"
        set service "ALL"
        set logtraffic all
        set auto-asic-offload disable
        set nat enable
    next
end
```

Only packets whose source MAC address belong to Fortinet or VMware are passed by the policy.

Dynamic policy — fabric devices

The dynamic address group represents the configured IP addresses of all Fortinet devices connected to the Security Fabric. It currently includes FortiManager, FortiAnalyzer, FortiClient EMS, FortiMail, FortiAP(s), and FortiSwitch(es). Like other dynamic address groups for fabric connectors, it can be used as an IPv4 address in firewall policies and objects.

The list of firewall addresses includes a default address object called `FABRIC_DEVICE`. You can apply the `FABRIC_DEVICE` object to the following types of policies:

- Firewall policy (including virtual wire pairs)
- IPv4 shaping policy

Netskope Exhibit 1017

- IPv4 ACL policy
- `policy64` and `policy46` (IPv4 only)

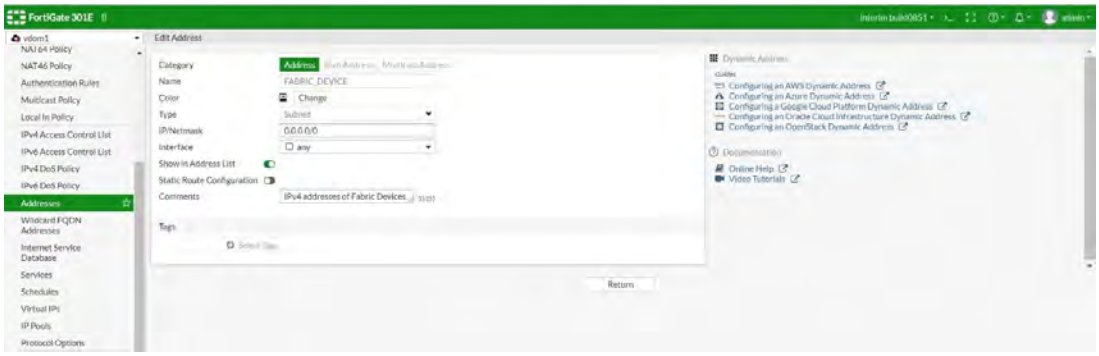
You cannot apply the `FABRIC_DEVICE` object to the following types of policies:

- IPv4 explicit proxy policy

You also cannot use the `FABRIC_DEVICE` object with the following settings:

- Custom extension on `internet-service`
- Exclusion of `addrgrp`

Initially the `FABRIC_DEVICE` object does not have an address value. The address value is populated dynamically as things change. As a result, you cannot edit the `FABRIC_DEVICE` object, add any addresses to the object, or remove any addresses from the object. The *Edit Address* pane in the GUI only has a *Return* button because the object is read-only:

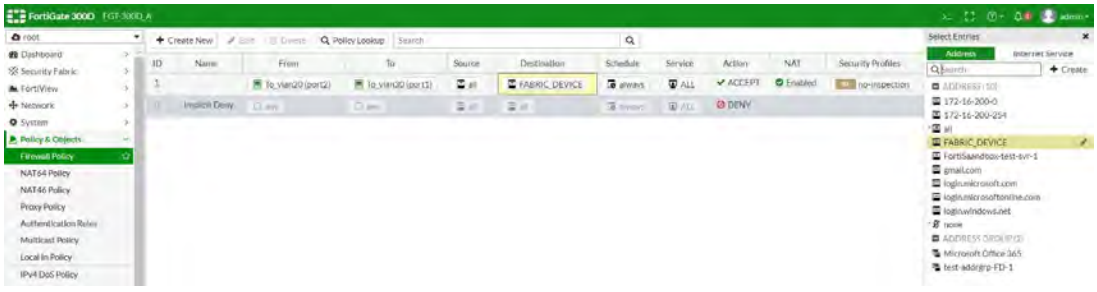


The `FABRIC_DEVICE` object address values are populated based on:

- FortiAnalyzer IP (from the *Fabric Settings* pane)
- FortiManager IP (from the *Fabric Settings* pane)
- FortiMail IP (from the *Fabric Settings* pane)
- FortiClient EMS IP (from the *Fabric Settings* pane)
- FortiAP IPs (from the *FortiAP Setup* pane or DHCP)
- FortiSwitch IPs (from the *FortiSwitch Setup* page or DHCP)

To apply the FABRIC_DEVICE object to a firewall policy using the GUI:

1. Go to *Policy & Objects* > *Firewall Policy*.
2. Create a new policy or edit an existing policy.
3. For the *Destination* field, select *FABRIC_DEVICE* from the list of address entries.



4. Configure the rest of the policy as needed.
5. Click *OK*.

To apply the FABRIC_DEVICE object to an IPv4 policy using the CLI:

```
(root) # show full-configuration firewall address FABRIC_DEVICE
config firewall address
    edit "FABRIC_DEVICE"
        set uuid ffde44a6-9eab-51ea-62f2-d5facde7af77
        set type ipmask
        set comment "IPv4 addresses of Fabric Devices."
        set associated-interface ''
        set color 0
        set allow-routing disable
        set subnet 0.0.0.0 0.0.0.0
    next
end
```

```
config firewall policy
    edit 1
        set srcintf "port2"
        set dstintf "port1"
        set srcaddr "all"
        set dstaddr "FABRIC_DEVICE"
        set action accept
        set schedule "always"
        set service "ALL"
        set utm-status enable
        set fsso disable
        set nat enable
    next
end
```

Diagnose command

You can use the diagnose command to list IP addresses of Fortinet devices that are configured in the Security Fabric.

To run the diagnose command using the CLI:

```
(root) # diagnose firewall sf-addresses list

FabricDevices: 172.18.64.48
FortiAnalyzer: 172.18.60.25
FortiSandbox: 172.18.52.154
FortiManager: 172.18.28.31
FortiClientEMS: 172.18.62.6
FortiAP:
FortiSwitch:
FortiAP/SW-DHCP:
```

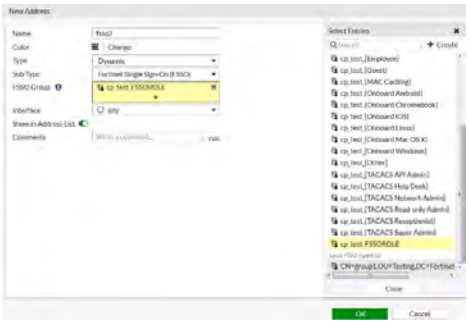
FSSO dynamic address subtype

The Fortinet Single Sign-ON (FSSO) dynamic firewall address subtype can be used in policies that support dynamic address types. The FortiGate will update the dynamic address used in firewall policies based on the source IP information for the authenticated FSSO users.

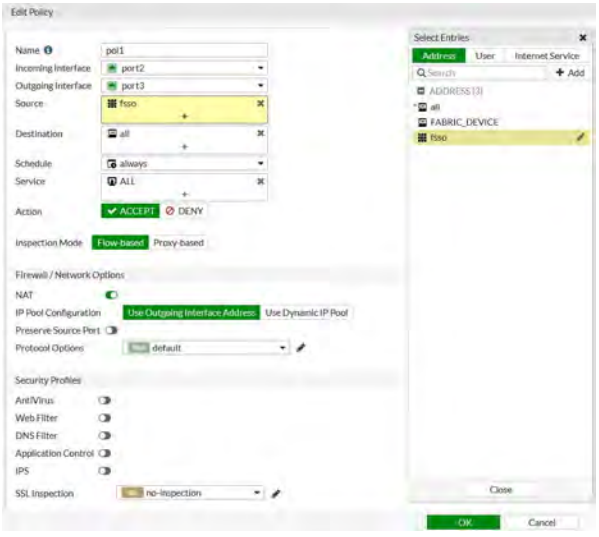
It can also be used with FSSO group information that is forwarded by ClearPass Policy Manager (CPPM) via FortiManager, and other FSSO groups provided by the FSSO collector agent or FortiNAC.

To configure FSSO dynamic addresses with CPPM and FortiManager in the GUI:

1. Create the dynamic address object:
   a. Go to *Policy & Objects* > *Addresses*, and click *Create New* > *Address*.
   b. For *Type*, select *Dynamic*.
   c. For *Sub Type*, select *Fortinet Single Sign-On (FSSO)*. The *Select Entries* pane opens and displays all available FSSO groups.
   d. Select one or more groups.
   e. Click *OK* to save the configuration.



   In the address table, there will be an error message for the address you just created (*Unresolved dynamic address: fsso*). This is expected because there are currently no authenticated FSSO users (based on source IP) in the local FSSO user list.
2. Add the dynamic address object to a firewall policy:
   a. Go to *Policy & Objects* > *Firewall Policy*.
   b. Create a new policy or edit an existing policy.
   c. For *Source*, add the dynamic FSSO address object you just created.
   d. Configure the rest of the policy as needed.
   e. Click *OK* to save your changes.



3. Test the authentication to add a source IP address to the FSSO user list:
   a. Log in as user and use CPPM for user authentication to connect to an external web server. After successful authentication, CPPM forwards the user name, source IP address, and group membership to the FortiGate via FortiManager.
   b. Go to *Monitor* > *Firewall User Monitor* to view the user name (*fsso1*) and IP address.

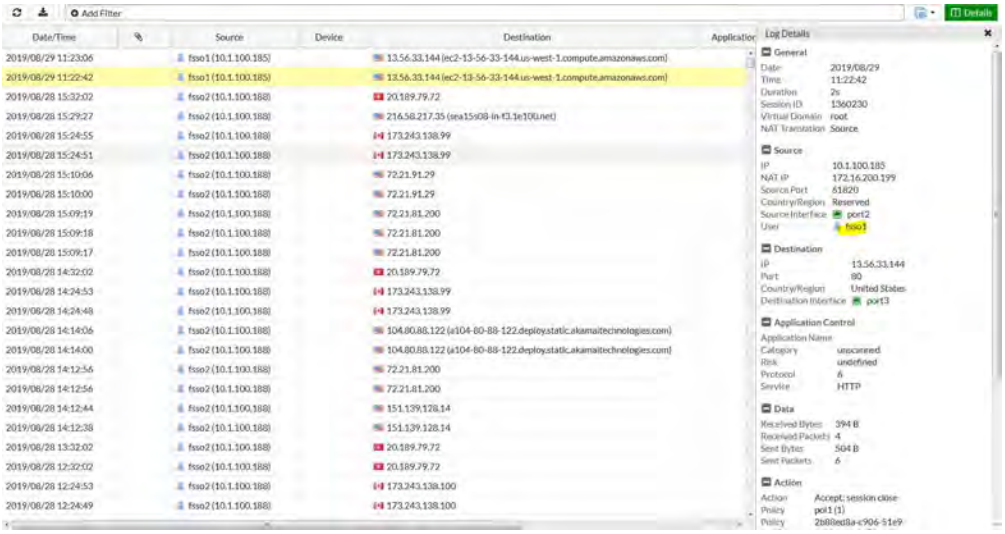


   c. Go to *Policy & Objects* > *Addresses* to view the updated address table. The error message no longer appears.
   d. Hover over the dynamic FSSO address to view the IP address (*fsso resolves to: 10.1.100.185*).



To verify user traffic in the GUI:

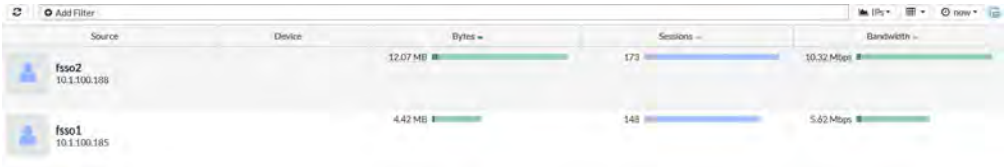
1. Go to *Log & Report* > *Forward Traffic*.
   Details for the user *fsso1* are visible in the traffic log:



   - If another user is authenticated by CPPM, then the dynamic address *fsso* entry in the address table will be updated. The IP address for user *fsso2* (*10.1.100.188*) is now visible:



2. Go to *FortiView* > *Sources* to verify that the users were able to successfully pass the firewall policy.



Netskope Exhibit 1017



If a user logs off and CPPM receives log off confirmation, then CPPS updates the FortiGate FSSO user list via FortiManager. The user IP address is deleted from the dynamic FSSO address, and the user is no longer be able to pass the firewall policy.

To configure FSSO dynamic addresses with CPPM and FortiManager in the CLI:

1. Create the dynamic address object:

```
config firewall address
    edit "fsso"
        set type dynamic
        set sub-type fsso
        set fsso-group "cp_test_FSSOROLE"
    next
end
```

2. Add the dynamic address object to a policy:

```
config firewall policy
    edit 1
        set name "pol1"
        set srcintf "port2"
        set dstintf "port3"
        set srcaddr "fsso"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ALL"
        set logtraffic all
        set fsso disable
        set nat enable
    next
end
```

To verify user traffic in the CLI:

1. Check the FSSO user list:

```
diagnose debug authd fsso list
----FSSO logons----
IP: 10.1.100.185  User: fsso1  Groups: cp_test_FSSOROLE  Workstation:  MemberOf: FSSO-
CPPM cp_test_FSSOROLE
Total number of logons listed: 1, filtered: 0
----end of FSSO logons----
```

2. Check the authenticated firewall users list:

```
diagnose firewall auth list
10.1.100.185, fsso1
type: fsso, id: 0, duration: 2928, idled: 2928
server: FortiManager
packets: in 0 out 0, bytes: in 0 out 0
group_id: 2 33554433
group_name: FSSO-CPPM cp_test_FSSOROLE
----- 1 listed, 0 filtered ------
```

After user traffic passes through the firewall, the nu

Netskope Exhibit 1017

```
diagnose firewall auth list
10.1.100.185, fsso1
type: fsso, id: 0, duration: 3802, idled: 143
server: FortiManager
packets: in 1629 out 1817, bytes: in 2203319 out 133312
group_id: 2 33554433
group_name: FSSO-CPPM cp_test_FSSOROLE
----- 1 listed, 0 filtered ------
```

ClearPass integration for dynamic address objects

ClearPass Policy Manager (CPPM) can gather information about the statuses of network hosts, for example, the latest patches or virus infections. Based on this information, CPPM send the IP addresses and current states, such as Healthy or Infected, to the FortiGate.

On the FortiGate, the IP addresses received from CPPM are added to a dynamic firewall address with the *clearpass-spt* subtype. This address can be used in any policy that supports dynamic addresses, such as Firewall or SSL-VPN policies.

In this example, you create two dynamic IP addresses that are used in two firewall policies (deny and allow). One policy allows traffic (host state = Healthy), and the other denies traffic (host state = Infected). When CPPM sends the information, the IP addresses are assigned according to their host state: Healthy or Infected.

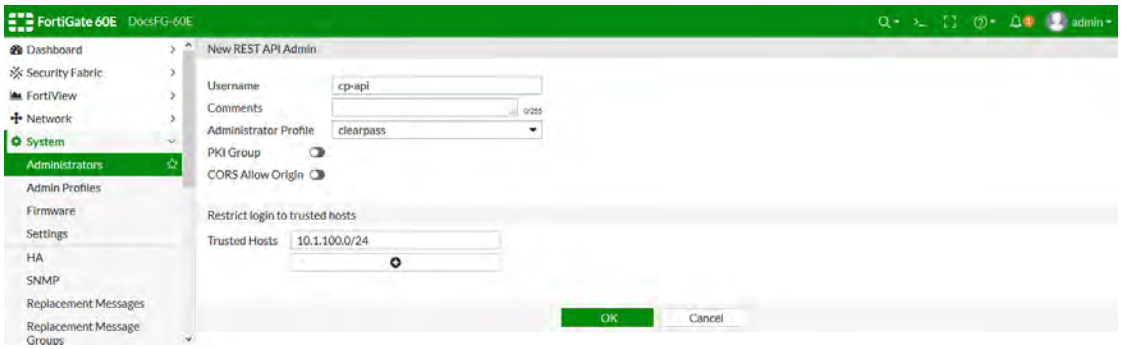
You can then verify that traffic from the Infected host is denied access by the deny policy, and traffic from the Healthy host is allowed access by the allow policy.

Create a REST API administrator

A RESET API administrator is required to generate an authorization token for REST API messages, and to limit hosts that can send REST API messages to the FortiGate.

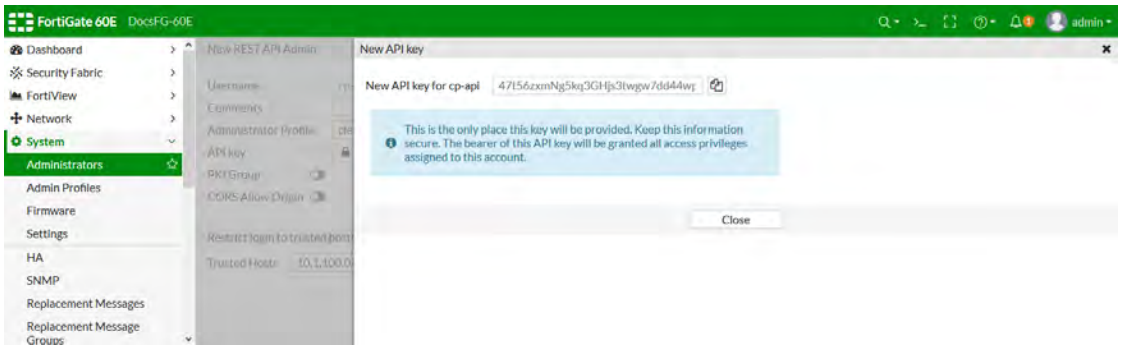
To create a REST API administrator in the GUI:

1. Go to *System* > *Administrators*.
2. Click *Create New* > *REST API Admin*.
3. Configure the *Username* and other information as needed.
4. Disable *PKI Group*.
5. In the *Trusted Hosts* field, enter *10.1.100.0/24*.



For this example, an administrator profile called *clearpass* was created with full read/write access. See Administrator profiles on page 587 for details.

6. Click *OK*.
   The *New API key* pane opens.



   The API key is the REST API authorization token that is used in REST API messages sent by CPPM to the FortiGate.
7. Copy the API key to a secure location. A new key can be generated if this one is lost or compromised.
8. Click *Close*.

To create a REST API administrator in the CLI:

```
config system api-user
    edit "cpi-back"
        set accprofile "clearpass"
        config trusthost
            edit 1
                set ipv4-trusthost 10.1.100.0 255.255.255.0
            next
        end
    next
end
```

```
execute api-user generate-key cp-api
    New API key: 0f1HxGHh9r9p74k7qgfHNH40p51bjs
    NOTE: The bearer of this API key will be granted all access privileges assigned to the
api-user cp-api.
```

Create dynamic IP addresses with the clearpass subtype

Two dynamic IP addresses are required, one for the allow policy, and the other for the deny policy.

To create the dynamic IP addresses:

```
config firewall address
    edit "cppm"
        set type dynamic
        set sub-type clearpass-spt
        set clearpass-spt healthy
        set comment ''
        set visibility enable
        set associated-interface ''
        set color 0
    next
    edit "cppm-deny"
```

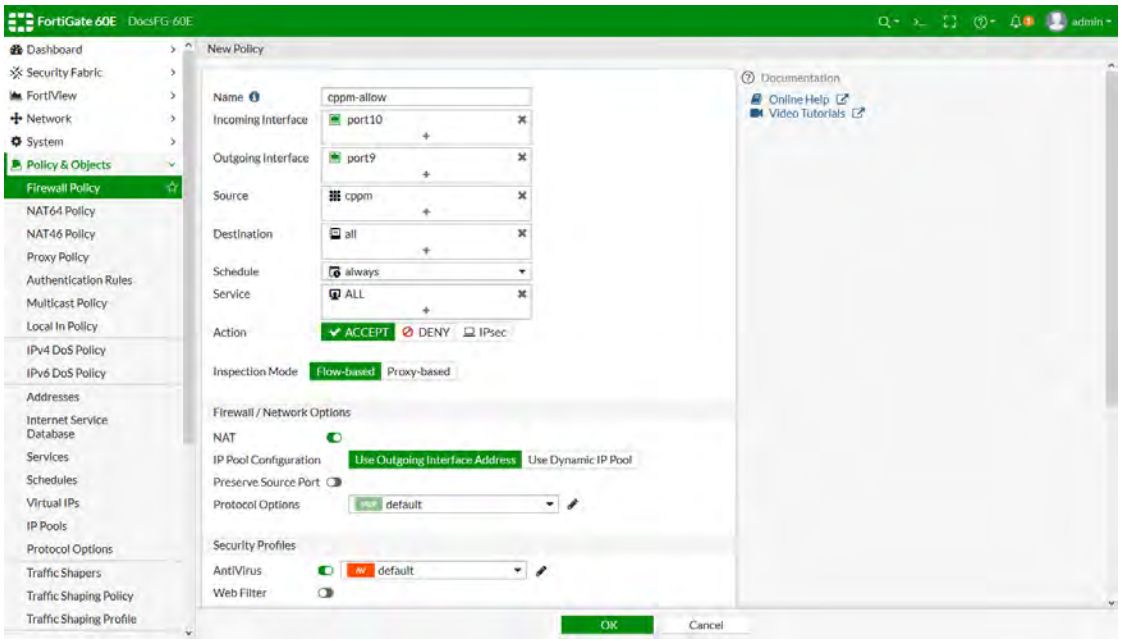
```
        set type dynamic
        set sub-type clearpass-spt
        set clearpass-spt infected
        set comment ''
        set visibility enable
        set associated-interface ''
        set color 0
    next
end
```

Create firewall policies

Two firewall policies are required, one to accept traffic (*cppm-allow*), and the other to deny traffic (*cppm-deny*).

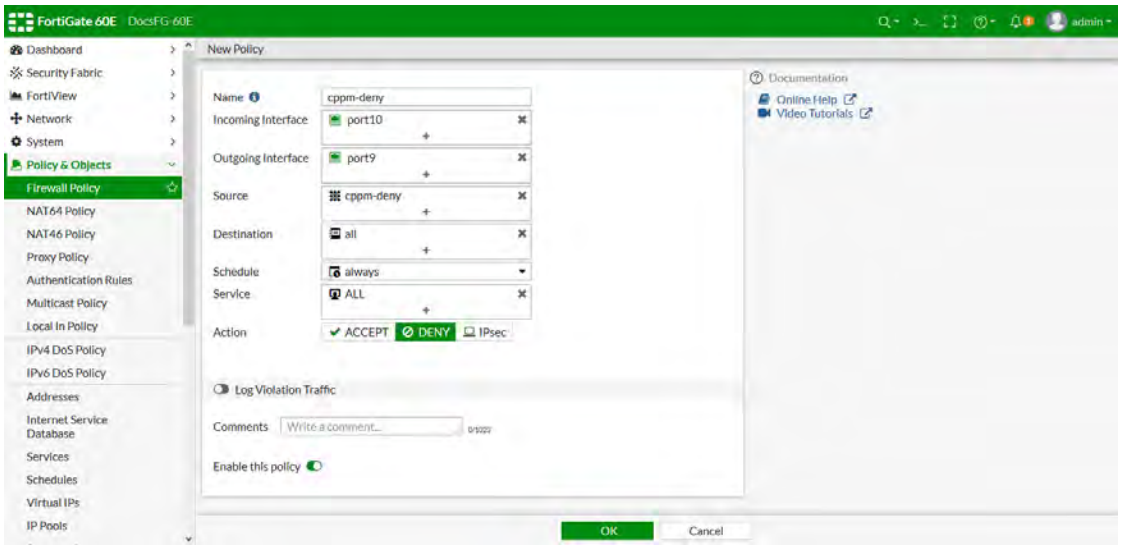
To create the firewall policies in the GUI:

1. Go to *Policy & Objects* > *Firewall Policy*.
2. Configure the allow policy:
   a. Click *Create New*.
   b. Enter a name for the policy.
   c. Set *Source* set to *cppm*.
   d. Set *Action* to *ACCEPT*.
   e. Configure the remaining settings as needed.



   f. Click *OK*.
3. Configure the deny policy:
   a. Click *Create New*.
   b. Enter a name for the policy.
   c. Set *Source* set to *cppm-deny*.
   d. Set *Action* to *DENY*.

e. Configure the remaining settings as needed.



f. Click *OK*.

To create the firewall policies in the CLI:

```
config firewall address
    edit "cppm"
        set type dynamic
        set sub-type clearpass-spt
        set clearpass-spt healthy
        set comment ''
        set visibility enable
        set associated-interface ''
        set color 0
    next
    edit "cppm-deny"
        set type dynamic
        set sub-type clearpass-spt
        set clearpass-spt infected
        set comment ''
        set visibility enable
        set associated-interface ''
        set color 0
    next
end
```

Verification

Go to *Log & Report* > *Forward Traffic* to review traffic logs and ensure that traffic is allowed or denied as expected.

To verify that FortiGate addresses are assigned correctly, enter the following CLI command:

```
diagnose firewall dynamic list
List all dynamic addresses:
cppm-deny: ID(141)
        ADDR(10.1.100.188)
```

```
cppm: ID(176)
        ADDR(10.1.100.185)
        ADDR(10.1.100.186)
```

Group address objects synchronized from FortiManager

Address objects from external connectors that are learned by FortiManager are synchronized to FortiGate. These objects can be grouped together with the FortiGate CLI to simplify selecting connector objects in the FortiGate GUI. Multiple groups can be created.

This option is only available for objects that are synchronized from FortiManager.

To add an object to a connector group:

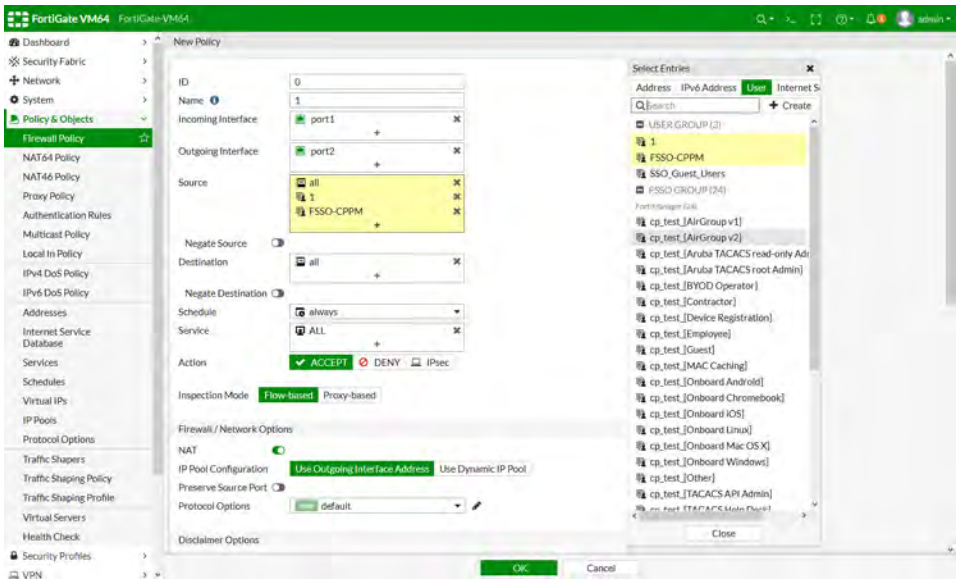
```
config user adgrp
    edit <object_name>
        set server-name "FortiManager"
        set connector-source <group_name>
    next
end
```

Example

In this example, objects learned by the FortiManager from an Aruba ClearPass device are synchronized to the FortiGate. Some of the objects are then added to a group called *ClearPass* to make them easier to find in the object list when creating a firewall policy.



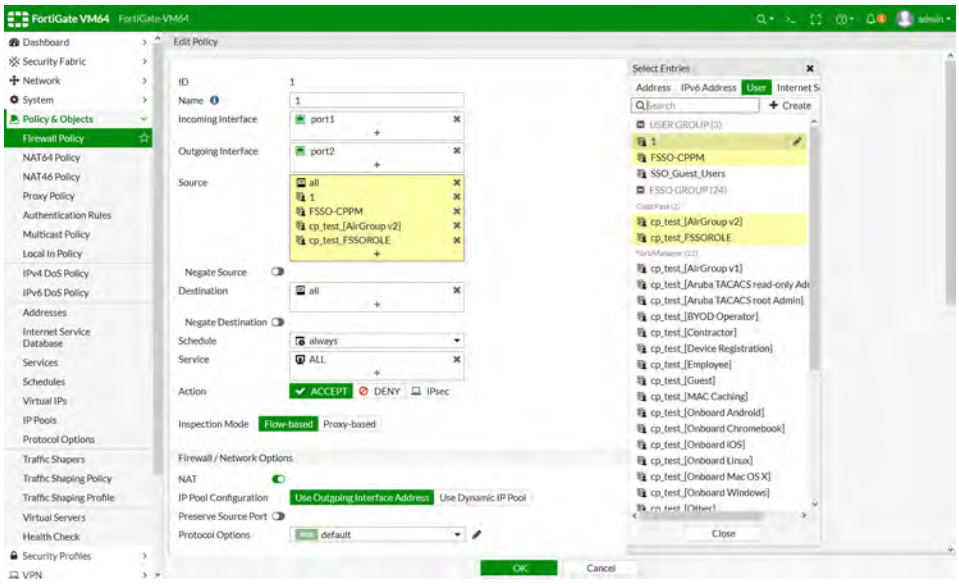
Prior to being grouped, the synchronized objects are listed under the FortiManager heading in the object lists.



To add some of the objects to a group:

```
config user adgrp
    edit "cp_test_FSSOROLE"
        set server-name "FortiManager"
        set connector-source "ClearPass"
    next
    edit "cp_test_[AirGroup v2]"
        set server-name "FortiManager"
        set connector-source "ClearPass"
    next
end
```

The objects are now listed under the *ClearPass* heading.



Using wildcard FQDN addresses in firewall policies

You can use wildcard FQDN addresses in firewall policies.

The firewall policy types that support wildcard FQDN addresses include IPv4, IPv6, ACL, local, shaping, NAT64, NAT46, and NGFW.

For wildcard FQDN addresses to work, the FortiGate should allow DNS traffic to pass through. Clients behind the FortiGate should use the same DNS server(s) as the FortiGate to ensure the FortiGate and the clients are resolving to the same addresses.

Initially, the wildcard FQDN object is empty and contains no addresses. When the client tries to resolve a FQDN address, the FortiGate will analyze the DNS response. The IP address(es) contained in the answer section of the DNS response will be added to the corresponding wildcard FQDN object.

Since FortiGate must analyze the DNS response, it does not work with DNS over HTTPS.

When the wildcard FQDN gets the resolved IP addresses, FortiOS loads the addresses into the firewall policy for traffic matching.

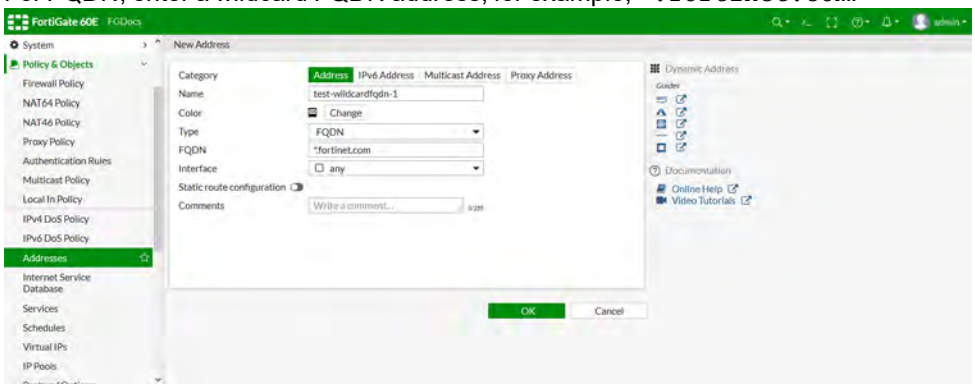
The FortiGate will keep the IP addresses in the FQDN object table as long as the DNS entry itself has not expired. Once it expires, the IP address is removed from the wildcard FQDN object until another query is made. At any given time, a single wildcard FQDN object may have up to 1000 IP addresses.

The DNS expiry TTL value is set by the authoritative name server for that DNS record. If the TTL for a specific DNS record is very short and you would like to cache the IP address longer, then you can extend it with the CLI. See To extend the TTL for a DNS record in the CLI: on page 810

For more information, see FQDN address firewall object type.

To create a wildcard FQDN using the GUI:

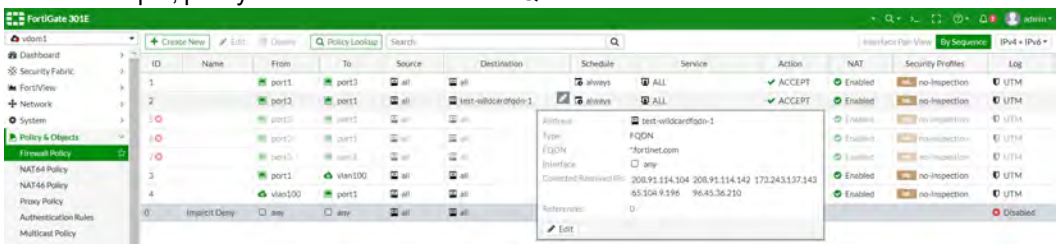
1. Go to *Policy & Objects* > *Addresses* and click *Create New* > *Address*.
2. Specify a *Name*.
3. For *Type*, select *FQDN*.
4. For *FQDN*, enter a wildcard FQDN address, for example, `*.fortinet.com`.



5. Click *OK*.

To use a wildcard FQDN in a firewall policy using the GUI:

1. Go to *Policy & Objects* > *Firewall Policy* and click *Create New*..
2. For *Destination*, select the wildcard FQDN.
3. Configure the rest of the policy as needed.
4. Click *OK*.
   In this example, policy ID 2 uses the wildcard FQDN:



To create a wildcard FQDN using the CLI:

```
config firewall address
    edit "test-wildcardfqdn-1"
        set type fqdn
        set fqdn "*.fortinet.com"
    next
end
```

To use wildcard FQDN in a firewall policy using the CLI:

```
config firewall policy
    edit 2
        set srcintf "port3"
        set dstintf "port1"
        set srcaddr "all"
        set dstaddr "test-wildcardfqdn-1"
        set action accept
        set schedule "always"
        set service "ALL"
        set auto-asic-offload disable
        set nat enable
    next
end
```

To use the diagnose command to list resolved IP addresses of wildcard FQDN objects:

```
# diagnose firewall fqdn list

List all FQDN:

*.fortinet.com: ID(48) ADDR(96.45.36.159) ADDR(192.168.100.161) ADDR(65.39.139.161)
```

Alternatively:

```
# diagnose test application dnsproxy 6

worker idx: 0

vfid=0 name=*.fortinet.com ver=IPv4  min_ttl=3266:0, cache_ttl=0 , slot=-1, num=3,
wildcard=1
```


Netskope Exhibit 1017

```
      96.45.36.159 (ttl=68862:68311:68311) 192.168.100.161 (ttl=3600:3146:3146)
65.39.139.161

(ttl=3600:3481:3481)
```

To use the diagnose command for firewall policies which use wildcard FQDN:

```
# diagnose firewall iprope list 100004
...
destination fqdn or dynamic address (1):*.fortinet.com ID(48) uuid_idx=57 ADDR
(208.91.114.104) ADDR(208.91.114.142) ADDR(173.243.137.143) ADDR(65.104.9.196) ADDR
(96.45.36.210)
...
```

To extend the TTL for a DNS record in the CLI:

In this the example the `set cache-ttl` value has been extended to 3600 seconds.

```
config firewall address
   edit "fortinet.com"
      set type fqdn
      set fqdn "www.fortinet.com”
      set cache-ttl 3600
   next
end
```

IPv6 geography-based addresses

Geography-based IPv6 addresses can be created and applied to IPv6 firewall policies.



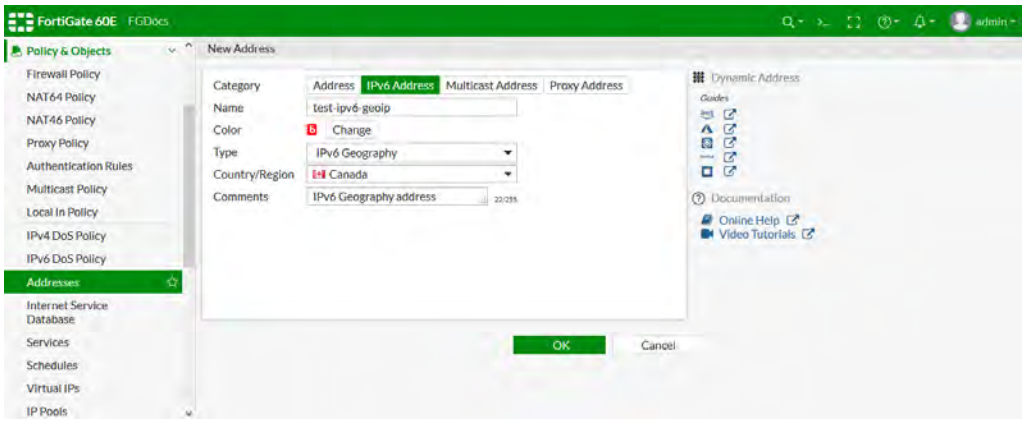
IPv6 geography-based addresses do not support `geoip-override` or `geoip-anycast`.

To create an IPv6 geography-based address in the GUI:

1. Go to *Policy and Objects* > *Addresses*.
2. Click *Create New* > *Address*.
3. Set *Category* to *IPv6 Address*.
4. Enter a name for the address.
5. Set *Type* to *IPv6 Geography*.
6. Select the *Country/Region* from the list.

Netskope Exhibit 1017

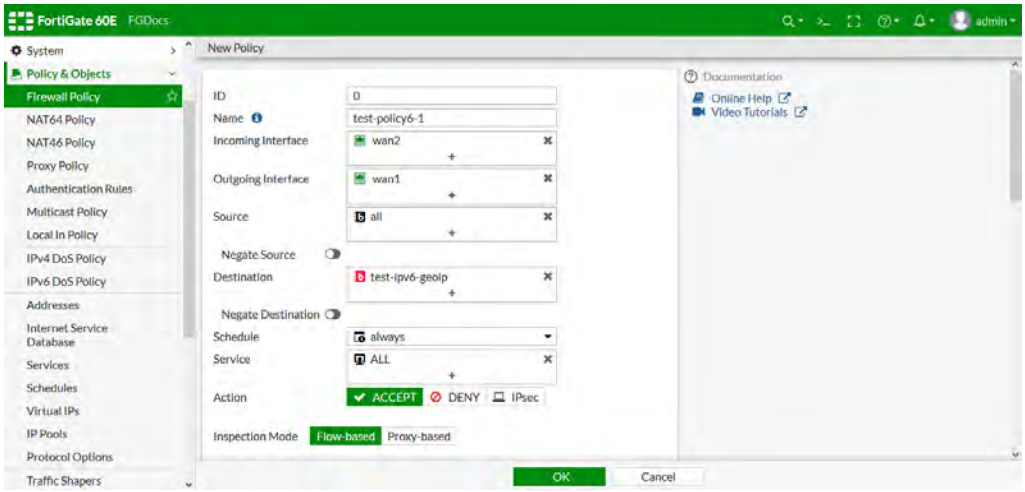
7. Optionally, enter comments.



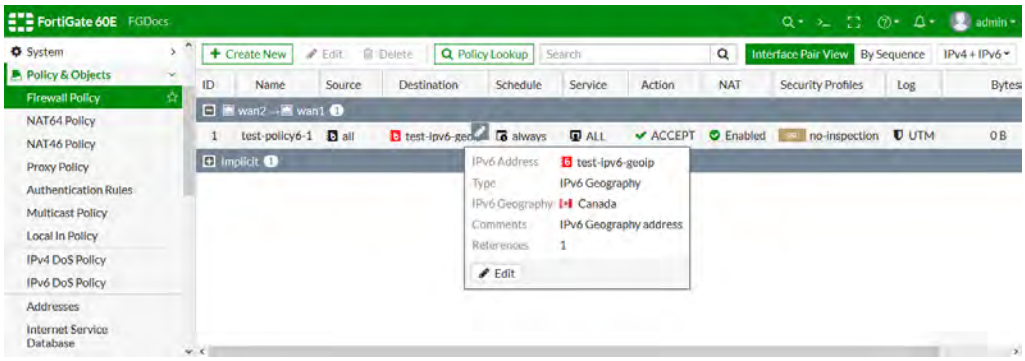
8. Click *OK*.

To use the IPv6 geography address in a policy:

1. Go to *Policy & Objects* > *Firewall Policy*.
2. Edit an existing policy, or create a new one, using the IPv6 geography address as the *Source* or *Destination Address*.



3. In the policy list, hover over the address to view details.



To configure an IPv6 geography-based address in the CLI:

1. Create an IPv6 geography-based address:

```
config firewall address6
    edit "test-ipv6-geoip"
        set type geography
        set color 6
        set comment "IPv6 Geography address"
        set country "CA"
    next
end
```

2. Use the IPv6 geography-based address in a policy:

```
config firewall policy
    edit 1
        set name "test-policy6-1"
        set srcintf "wan2"
        set dstintf "wan1"
        set srcaddr6 "all"
        set dstaddr6 "test-ipv6-geoip"
        set action accept
        set schedule "always"
        set service "ALL"
        set nat enable
    next
end
```

Array structure for address objects

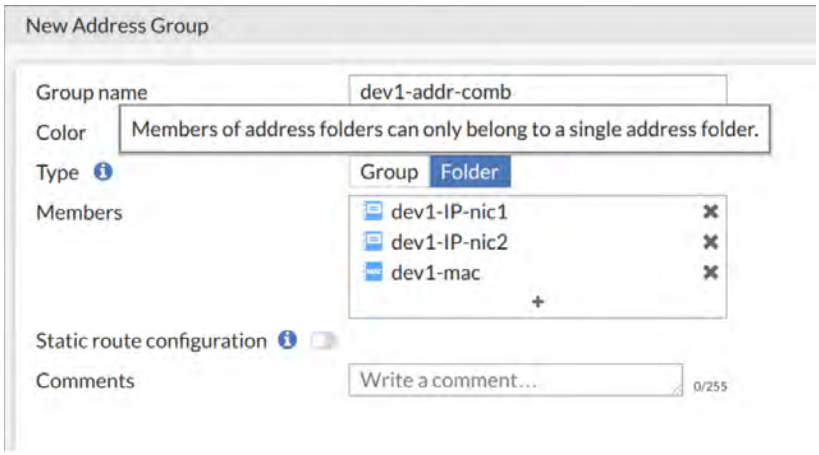
Some address objects logically belong to the same device, such as two IPs from the same computer. These address objects can be grouped into an address folder, which is an exclusive list of address objects that do not appear in other address groups or folders.

In the CLI, the folder type can be set after the member list is already populated. If the member list contains an incompatible entry, then the setting will be discarded when the `next`/`end` command is issued. If the folder type is set before the member list is populated, then the possible member entry list will be filtered according to the selected type.

To create an address folder in the GUI:

1. Go to *Policy & Objects* > *Addresses*.
2. Click *Create New* > *Address Group* and enter a name.
3. For *Type*, select *Folder*.
4. For *Members*, click the + to add the addresses. Address folders and groups are exclusive, so the *Select Entries* window filters out address objects that are a member of an existing group or folder.

Netskope Exhibit 1017



5. Click *OK*.
6. In the address table, expand the *Address Group* section to view the folder (*dev1-addr-comb*). The expandable folder view shows the address folder's child objects:



To configure an address folder in the CLI:

notes

```
config firewall addrgrp
    edit "safe-network1-devices"
        set type folder
        set member "dev1-addr-comb" "dev2-addr-comb"
        set comment ''
        set exclude disable
        set color 13
    next
end
```

```
config firewall addrgrp
    edit "dev1-addr-comb"
        set type folder
        set member "dev1-IP-nic1" "dev1-IP-nic2" "dev1-mac"
        set comment ''
        set exclude disable
        set color 18
    next
end
```

```
config firewall addrgrp
    edit "dev2-addr-comb"
        set type folder
        set member "dev2-IP-nic1" "dev2-IP-nic2" "dev2-IP-nic3" "dev2-mac"
        set comment ''
        set exclude disable
        set color 5
    next
end
```

Traffic shaping

QoS (quality of service) is the capability to adjust quality aspects of your overall network traffic, including techniques such as priority-based queuing and traffic policing. Because bandwidth is finite and some types of traffic are slow, jitter or packet loss sensitive, bandwidth intensive, or critical for operations, QoS is a useful tool to optimize the performance of various applications in your network. QoS is especially important for managing voice and streaming multimedia traffic because these types of traffic can rapidly consume bandwidth and are sensitive to latency. You can implement QoS on FortiGate devices using the following techniques:

Technique	Description
Traffic policing	The FortiGate drops packets that do not conform to the configured bandwidth limitations. Note that excessive traffic policing can degrade network performance rather than improve it.
Traffic shaping	The FortiGate ensures that traffic consumes bandwidth at least at the guaranteed rate by assigning a greater priority queue to the traffic if the guaranteed rate is not being met. The FortiGate ensures that traffic does not consume more than the maximum configured bandwidth. Traffic that exceeds the maximum rate is subject to traffic policing.
Queuing	The FortiGate transmits packets in the order of their assigned priority queue for that physical interface. All traffic in a higher priority traffic queue must be completely transmitted before traffic in lower priority queues is transmitted.

When determining how to configure QoS, it is helpful to know when a FortiGate uses each technique in the overall traffic processing flow and the considerations for each technique. After the FortiGate accepts packets, it classifies the traffic and may apply traffic policing at additional points during traffic processing. The FortiGate may also apply QoS techniques, such as prioritization and traffic shaping. Traffic shaping consists of both traffic policing to enforce bandwidth limits and adjusting priority queues to help packets achieve the guaranteed rate.

Traffic shaping accuracy is optimal for security policies without a protection profile where no FortiGate content inspection is processed.



You can enable traffic shaping in *System* > *Feature Visibility* under the *Additional Features* section.

The following topics provide information about configuring traffic shaping policies:

- Determining your QoS requirements on page 815
- Packet rates on page 816
- Changing traffic shaper bandwidth unit of measurement on page 818
- Shared traffic shaper on page 818
- Per-IP traffic shaper on page 822
- Type of Service-based prioritization and policy-based traffic shaping on page 825
- Interface-based traffic shaping profile on page 828
- Interface-based traffic shaping with NP acceleration on page 837
- Classifying traffic by source interface on page 838
- Configuring traffic class IDs on page 839
- Traffic shaping schedules on page 842
- QoS assignment and rate limiting for quarantined VLANs on page 843
- Weighted random early detection queuing on page 844

Determining your QoS requirements

Before implementing QoS, you should identify the types of traffic that:

- Are important to your organization
- Use high amounts of bandwidth
- Are sensitive to latency or packet loss

Discovering the needs and relative importance of each traffic type on your network will help you design an appropriate overall approach, including how you configure each available QoS component technique. Some organizations discover they only need to configure bandwidth limits for some services. Other organizations determine they need to fully configure interface and security policy bandwidth limits for all services, and prioritize the queuing of critical services relative to traffic rate.

For example, your organization wants to guarantee sufficient bandwidth for revenue-producing e-commerce traffic. You need to ensure that customers complete transactions and do not experience service delays. At the same time, you need to ensure low latency for voice over IP (VoIP) traffic that sales and customer support teams use, while traffic latency and bursts may be less critical to the success of other network applications, such as long term, resumable file transfers.

Best practices

The following list includes recommendations and considerations when configuring QoS in your network:

- Ensure maximum bandwidth limits at the source interface and security policy are not too low. This can cause the FortiGate to discard an excessive number of packets.
- Consider the ratios of how packets are distributed between the available queues, and which queues are used by which types of services. Assigning most packets to the same priority queue can reduce the effects of configuring prioritization. Assigning a lot of high bandwidth services to high priority queues may take too much bandwidth away from lower priority queues and cause increased or indefinite latency. For example, you may want to prioritize a latency-sensitive service, such as SIP, over a bandwidth-intensive service, such as FTP. Also consider that bandwidth guarantees can affect queue distribution, and assign packets to queue 0 instead of their regular queue in high-volume situations.
- Decide whether or not to guarantee bandwidth because it causes the FortiGate to assign packets to queue 0 if the guaranteed packet rate is not being met. When you compare queuing behavior for low and high bandwidth

Netskope Exhibit 1017

situations, this means the effect of prioritization only becomes visible as traffic volumes rise and exceed their guarantees. Because of this, you might want only some services to use bandwidth guarantees. This way, you can avoid the possibility that all traffic uses the same queue in high-volume situations, which negates the effects of configuring prioritization.

- Configure prioritization for all through traffic by either ToS (type of service)-based priority or security policy priority, not both, to simplify analysis and troubleshooting. Traffic subject to both ToS-based and security policy priorities use a combined priority from both parts of the configuration. Traffic subject to only one of the prioritization methods will use only that priority. If you configure both methods, or if you configure either method for only a subset of traffic, packets that apply to the combined configuration may receive a lower priority queue than packets that apply to only one of the priority methods, as well as packets that do not apply to the configured prioritization. For example, if both the ToS-based priority and security policy priority dictate that a packet should receive a medium priority, in the absence of bandwidth guarantees, a packet will use queue 3. If only ToS-based priority is configured, the packet will use queue 1. If only security policy priority is configured, the packet will use queue 2. If no prioritization is configured, the packet will use queue 0.
  - Because you can configure QoS using a combination of security policies and ToS-based priorities, and to distribute traffic over the six possible queues for each physical interface, the results of those configurations can be more difficult to analyze because of their complexity. In those cases, prioritization behavior can vary by several factors, including: traffic volume, ToS or differentiated services (DiffServ) markings, and correlation of session to a security policy.



The FortiGate does not prioritize traffic based on the differentiated services code point (DSCP) marking configured in the security policy. However, ToS-based prioritization can be used for ingress traffic.

- Use the UDP protocol to obtain more accurate testing results. Packets that are discarded by traffic shapers impact flow-control mechanisms, such as TCP.
- Do not oversubscribe outbandwidth throughput. For example, sum [guaranteed bandwidth] < outbandwidth. For accurate bandwidth calculations, you must set the outbandwidth parameter on interfaces.

Packet rates

The formula for packet rates specified for maximum bandwidth or guaranteed bandwidth is:

rate = amount / time

where rate is in Kbps

Burst size cannot exceed the configured maximum bandwidth. The FortiGate drops packets that exceed the configured maximum bandwidth. Packets deduct from the amount of bandwidth available to subsequent packets, and available bandwidth regenerates at a fixed rate. As a result, the available bandwidth for a packet may be less than the configured rate, down to a minimum of 0 Kbps.

Alternatively, rate calculation and behavior can be described using the token bucket metaphor. A traffic flow has an associated bucket, which represents burst size bounds and is the size of the configured bandwidth limit. The bucket receives tokens, which represent available bandwidth at the fixed configured rate. As time passes, tokens are added to the bucket up to capacity, and excess tokens are discarded. When a packet arrives at the FortiGate, the packet must deduct bandwidth tokens from the bucket equal to its size in order to leave the FortiGate. If there are not enough tokens, the packet cannot leave the FortiGate and is dropped.

Bursts are not redistributed over a longer interval, so bursts are propagated rather than smoothed. However, peak size is limited. The maximum burst size is the capacity of the bucket, which is the configured bandwidth limit. The actual size varies depending on the current number of tokens in the bucket, which may be less than the capacity of the bucket due to

Netskope Exhibit 1017

deductions made by previous packets and the fixed rate at which tokens accumulate. A depleted bucket refills at the rate of the configured bandwidth limit. Bursts cannot borrow tokens from other time intervals.

By limiting traffic peaks and token regeneration, the available bandwidth may be less than the capacity of the bucket, but the limit of the total amount per time interval is ensured. Total bandwidth use during each interval of one second is, at most, the integral of the configured rate.

Rate discrepancy

You may observe that external clients, such as FTP or BitTorrent, initially report rates between the maximum bandwidth and twice the amount of the maximum bandwidth depending on the size of their initial burst. For example, when a connection is initiated following a period of no network activity. The apparent discrepancy in rates is caused by a difference in perspective when delimiting time intervals. A burst from the client may initially consume all tokens in the bucket, and before the end of one second as the bucket regenerates, is allowed to consume almost another bucket worth of bandwidth. From the perspective of the client, this equals one time interval. However, from the perspective of the FortiGate, the bucket cannot accumulate tokens when it is full. Therefore, the time interval for token regeneration begins after the initial burst and does not contain the burst. These different points of reference result in an initial discrepancy equal to the size of the burst. The client's rate contains it, but the FortiGate's rate does not. However, if the connection is sustained to its limit and time progresses over an increasing number of intervals, this discrepancy decreases in importance relative to the bandwidth total. The client reported rate will eventually approach the configured rate limit for the FortiGate.

Example

The maximum bandwidth is 50 Kbps, there has been no network activity for one or more seconds, and the bucket is full. A burst from an FTP client immediately consumes 50 kilobits. Because the bucket completely regenerates over one second, by the time another second elapses from the initial burst, traffic can consume another 49.999 kilobits, for a total of 99.999 kilobits between the two points in time. From the vantage point of an external FTP client regulated by this bandwidth limit, it initially appears that the bandwidth limit is 99.999 Kbps. This is almost twice the configured limit of 50 Kbps. However, bucket capacity only regenerates at the configured rate of 50 Kbps, and the connection can only consume a maximum of 50 kilobits during each subsequent second. The result is that as bandwidth consumption is averaged over an increasing number of time intervals, each of which are limited to 50 Kbps, the effect of the first interval's doubled bandwidth size diminishes proportionately, and the client's reported rate eventually approaches the configured rate limit. The following table shows the effects of a 50 Kbps limit on client reported rates:

Total size transferred (kilobits)	Time (seconds)	Rate reported by client (Kbps)
99.999 (50 + 49.999)	1	99.999
149.999	2	74.999
199.999	3	66.666
249.999	4	62.499
299.999	5	59.998
349.999	6	58.333

Guaranteed bandwidth can also be described using a token bucket metaphor. However, because this feature attempts to achieve or exceed a rate rather than limit it, the FortiGate does not discard non-conforming packets, as it does for

Netskope Exhibit 1017

maximum bandwidth. Instead, when the flow does not achieve the rate, the FortiGate increases the packet priority queue, in an effort to increase the rate.

Guaranteed and maximum bandwidth rates apply to the bidirectional total for all sessions controlled by the security policy. For example, an FTP connection may entail two separate connections for the data and control portion of the session. Some packets may be reply traffic rather than initiating traffic. All packets for both connections are counted when calculating the packet rate for comparison with the guaranteed and maximum bandwidth rate.

Changing traffic shaper bandwidth unit of measurement

Bandwidth speeds are measured in kilobits per second (Kbps), and bytes that are sent and received are measured in megabytes (MB). In some cases, this can cause confusion depending on whether your ISP uses kilobits per second (Kbps), kilobytes per second (KBps), megabits per second (Mbps), or gigabits per second (Gbps).

You can change the unit of measurement for traffic shapers in the CLI.

To change the bandwidth unit of measurement for a shared traffic shaper:

```
config firewall shaper traffic-shaper
    edit <traffic_shaper_name>
        set bandwidth-unit {kbps | mbps | gbps}
    next
end
```

To change the bandwidth unit of measurement for a per-IP traffic shaper:

```
config firewall shaper per-ip-shaper
    edit <traffic_shaper_name>
        set bandwidth-unit {kbps | mbps | gbps}
    next
end
```

Shared traffic shaper

Shared traffic shaper is used in a firewall shaping policy to indicate the priority and guaranteed and maximum bandwidth for a specified type of traffic use.

The maximum bandwidth indicates the largest amount of traffic allowed when using the policy. You can set the maximum bandwidth to a value between 1 and 16776000 Kbps. The GUI displays an error if any value outside this range is used. If you want to allow unlimited bandwidth, use the CLI to enter a value of 0.

The guaranteed bandwidth ensures that there is a consistent reserved bandwidth available. When setting the guaranteed bandwidth, ensure that the value is significantly less than the interface's bandwidth capacity. Otherwise, the interface will allow very little or no other traffic to pass through, potentially causing unwanted latency.

In a shared traffic shaper, the administrator can prioritize certain traffic as high, medium, or low. FortiOS provides bandwidth to low priority connections only when high priority connections do not need the bandwidth. For example, you should assign a high traffic priority to a policy for connecting a secure web server that needs to support e-commerce traffic. You should assign less important services a low priority.

When you configure a shared traffic shaper, you can apply bandwidth shaping per policy or for all policies. By default, a shared traffic shaper applies traffic shaping evenly to all policies that use the shared traffic shaper.

Netskope Exhibit 1017

When configuring a per-policy traffic shaper, FortiOS applies the traffic shaping rules defined for each security policy individually. For example, if a per-policy traffic shaper is configured with a maximum bandwidth of 1000 Kbps, any security policies that have that traffic shaper enabled get 1000 Kbps of bandwidth each.

If a traffic shaper for all policies is configured with a maximum bandwidth of 1000 Kbps, all policies share the 1000 Kbps on a first-come, first-served basis.

The configuration is as follows:

```
config firewall shaper traffic-shaper
   edit "traffic_shaper_name"
      set per-policy enable
   next
end
```

The shared traffic shaper selected in the traffic shaping policy affects traffic in the direction defined in the policy. For example, if the source port is LAN and the destination is WAN1, the traffic shaping affects the flow in this direction only, affecting the outbound traffic's upload speed. You can define the traffic shaper for the policy in the opposite direction (reverse shaper) to affect the inbound traffic's download speed. In this example, that would be from WAN1 to LAN.

The following example shows how to apply different speeds to different types of service. The example configures two shared traffic shapers to use in two firewall shaping policies. One policy guarantees a speed of 10 Mbps for VoIP traffic. The other policy guarantees a speed of 1 Mbps for other traffic. In the example, FortiOS communicates with a PC using port10 and the Internet using port9.

To configure shared traffic shapers in the FortiOS GUI:

1. Create a firewall policy:
   a. Go to *Policy & Objects* > *Firewall Policy*.
   b. Click *Create New*.
   c. In the *Name* field, enter Internet Access.
   d. From the *Incoming Interface* dropdown list, select *port10*.
   e. From the *Outgoing Interface* dropdown list, select *port9*.
   f. For the *Source* and *Destination* fields, select *all*.
   g. From the *Schedule* dropdown list, select *always*.
   h. For the *Service* field, select *ALL*.
   i. Click *OK*.
2. Create the shared traffic shapers:
   a. Go to *Policy & Objects* > *Traffic Shapers*.
   b. Click *Create New*.
   c. In the *Name* field, enter 10Mbps. This shaper is for VoIP traffic.
   d. From the *Traffic Priority* dropdown list, select *High*.
   e. Enable *Max Bandwidth* and enter 20000. This equates to 20 Mbps.
   f. Enable *Guaranteed Bandwidth* and enter 10000. This equates to 10 Mbps.
   g. Click *OK*.
   h. Repeat the process above to create another traffic shaper named 1Mbps. Set the *Traffic Priority* to *Low*, the *Max Bandwidth* to 10000, and the *Guaranteed Bandwidth* to 1000.
3. Create a firewall shaping policy:
   a. Go to *Policy & Objects* > *Traffic Shaping Policy*.
   b. Click *Create New*.
   c. In the *Name* field, enter VoIP_10Mbps_High. This policy is for VoIP traffic.

Netskope Exhibit 1017

d. For the *Source* and *Destination* fields, select *all*.
e. For the *Service* field, select all VoIP services.
f. For the *Outgoing Interface* field, select *port9*.
g. Enable *Shared shaper*. Select 10Mbps from the dropdown list.
h. Enable *Reverse shaper*. Select 10Mbps from the dropdown list.
i. Click *OK*.
j. Repeat the process above to create a firewall shaping policy named Other_1Mbps_Low for other traffic. Set the *Source* and *Destination* to *all*, *Service* to *ALL*, *Outgoing Interface* to *port9*, and *Shared shaper* and *Reverse shaper* to 1Mbps.

To configure shared traffic shapers using the FortiOS CLI:

1. Create a firewall policy:

```
config firewall policy
   edit 1
      set name "Internet Access"
      set srcintf "port10"
      set dstintf "port9"
      set srcaddr "all"
      set dstaddr "all"
      set action accept
      set schedule "always"
      set service "ALL"
      set fsso disable
      set nat enable
   next
end
```

2. Create the shared traffic shapers:

```
config firewall shaper traffic-shaper
   edit "10Mbps"
      set guaranteed-bandwidth 10000
      set maximum-bandwidth 20000
   next
   edit "1Mbps"
      set guaranteed-bandwidth 1000
      set maximum-bandwidth 10000
      set priority low
   next
end
```

3. Create a firewall shaping policy:

```
config firewall shaping-policy
   edit 1
      set name "VOIP_10Mbps_High"
      set service "H323" "IRC" "MS-SQL" "MYSQL" "RTSP" "SCCP" "SIP" "SIP-MSNmessenger"
      set dstintf "port9"
      set traffic-shaper "10Mbps"
      set traffic-shaper-reverse "10Mbps"
      set srcaddr "all"
      set dstaddr "all"
   next
   edit 2
      set name "Other_1Mbps_Low"
      set service "ALL"
      set dstintf "port9"
```


Netskope Exhibit 1017

```
      set traffic-shaper "1Mbps"
      set traffic-shaper-reverse "1Mbps"
      set srcaddr "all"
      set dstaddr "all"
   next
end
```

To troubleshoot shared traffic shapers:

1. To check if specific traffic is attached to the correct traffic shaper, run the `diagnose firewall iprope list 100015` command. The example output shows the traffic attached to the 10Mbps and 1Mbps shapers:

```
# diagnose firewall iprope list 100015

policy index=1 uuid_idx=0 action=accept
flag (0):
shapers: orig=10Mbps(2/1280000/2560000)
cos_fwd=0 cos_rev=0
group=00100015 av=00000000 au=00000000 split=00000000
host=4 chk_client_info=0x0 app_list=0 ips_view=0
misc=0 dd_type=0 dd_mode=0
zone(1): 0 -> zone(1): 38
source(1): 0.0.0.0-255.255.255.255, uuid_idx=0,
dest(1): 0.0.0.0-255.255.255.255, uuid_idx=0,
service(15):
   [6:0x0:0/(1,65535)->(1720,1720)] helper:auto
   [6:0x0:0/(1,65535)->(1503,1503)] helper:auto
   [17:0x0:0/(1,65535)->(1719,1719)] helper:auto
   [6:0x0:0/(1,65535)->(6660,6669)] helper:auto
   [6:0x0:0/(1,65535)->(1433,1433)] helper:auto
   [6:0x0:0/(1,65535)->(1434,1434)] helper:auto
   [6:0x0:0/(1,65535)->(3306,3306)] helper:auto
   [6:0x0:0/(1,65535)->(554,554)] helper:auto
   [6:0x0:0/(1,65535)->(7070,7070)] helper:auto
   [6:0x0:0/(1,65535)->(8554,8554)] helper:auto
   [17:0x0:0/(1,65535)->(554,554)] helper:auto
   [6:0x0:0/(1,65535)->(2000,2000)] helper:auto
   [6:0x0:0/(1,65535)->(5060,5060)] helper:auto
   [17:0x0:0/(1,65535)->(5060,5060)] helper:auto
   [6:0x0:0/(1,65535)->(1863,1863)] helper:auto

policy index=2 uuid_idx=0 action=accept
flag (0):
shapers: orig=1Mbps(4/128000/1280000)
cos_fwd=0 cos_rev=0
group=00100015 av=00000000 au=00000000 split=00000000
host=4 chk_client_info=0x0 app_list=0 ips_view=0
misc=0 dd_type=0 dd_mode=0
zone(1): 0 -> zone(1): 38
source(1): 0.0.0.0-255.255.255.255, uuid_idx=0,
dest(1): 0.0.0.0-255.255.255.255, uuid_idx=0,
service(1):
   [0:0x0:0/(0,0)->(0,0)] helper:auto
```

2. To check if the correct traffic shaper is applied to the session, run the `diagnose sys session list` command. The example output shows that the 1Mbps shaper is applied to the session:

```
# diagnose sys session list
```


Netskope Exhibit 1017

```
session info: proto=6 proto_state=01 duration=11 expire=3599 timeout=3600 flags=00000000
     sockflag=00000000 sockport=0 av_idx=0 use=5
origin-shaper=1Mbps prio=4 guarantee 128000Bps max 1280000Bps traffic 1050Bps drops 0B
reply-shaper=
per_ip_shaper=
class_id=0 shaping_policy_id=2 ha_id=0 policy_dir=0 tunnel=/ helper=ftp vlan_cos=0/255
state=may_dirty npu npd os mif route_preserve
statistic(bytes/packets/allow_err): org=868/15/1 reply=752/10/1 tuples=2
tx speed(Bps/kbps): 76/0 rx speed(Bps/kbps): 66/0
orgin->sink: org pre->post, reply pre->post dev=39->38/38->39 gwy=172.16.200.55/0.0.0.0
hook=post dir=org act=snat 10.1.100.11:58241->172.16.200.55:21(172.16.200.1:58241)
hook=pre dir=reply act=dnat 172.16.200.55:21->172.16.200.1:58241(10.1.100.11:58241)
pos/(before,after) 0/(0,0), 0/(0,0)
misc=0 policy_id=1 auth_info=0 chk_client_info=0 vd=4
serial=0003255f tos=ff/ff app_list=0 app=0 url_cat=0
rpdb_link_id = 00000000
dd_type=0 dd_mode=0
npu_state=0x100000
npu info: flag=0x00/0x00, offload=0/0, ips_offload=0/0, epid=0/0, ipid=0/0,
     vlan=0x0000/0x0000
vlifid=0/0, vtag_in=0x0000/0x0000 in_npu=0/0, out_npu=0/0, fwd_en=0/0, qid=0/0
no_ofld_reason: offload-denied helper
total session 1
```

3. To check statuses of shared traffic shapers, run the `diagnose firewall shaper traffic-shaper list` command. The output should resemble the following:

```
# diagnose firewall shaper traffic-shaper list

name 10Mbps
maximum-bandwidth 2500 KB/sec
guaranteed-bandwidth 1250 KB/sec
current-bandwidth 0 B/sec
priority 2
tos ff
packets dropped 0
bytes dropped 0

name 1Mbps
maximum-bandwidth 1250 KB/sec
guaranteed-bandwidth 125 KB/sec
current-bandwidth 0 B/sec
priority 4
tos ff
packets dropped 0
bytes dropped 0
```

Per-IP traffic shaper

With per-IP traffic shaping, you can limit each IP address's behavior to avoid a situation where one user uses all of the available bandwidth. In addition to controlling the maximum bandwidth used per IP address, you can also define the maximum number of concurrent sessions for an IP address. For example, if you apply a per-IP shaper of 1 Mbps to your entire network, FortiOS allocates each user/IP address 1 Mbps of bandwidth. Even if the network consists of a single user, FortiOS allocates them 1 Mbps. If there are ten users, each user gets 1 Mbps of bandwidth, totaling 10 Mbps of outgoing traffic.

Netskope Exhibit 1017

For shared shapers, all users share the set guaranteed and maximum bandwidths. For example, if you set a shared shaper for all PCs using an FTP service to 10 Mbps, all users uploading to the FTP server share the 10 Mbps.

Shared shapers affect upload speed. If you want to limit the download speed from the FTP server in the example, you must configure the shared shaper as a reverse shaper. Per-IP shapers apply the speed limit on both upload and download operations.

The following example shows how to apply a per-IP shaper to a traffic shaping policy. This shaper assigns each user a maximum bandwidth of 1 Mbps and allows each user to have a maximum of ten concurrent connections to the FTP server. In the example, FortiOS communicates with users using port10 and the FTP server using port9.

To configure a per-IP shaper in the FortiOS GUI:

1. Create a firewall policy:
   a. Go to *Policy & Objects* > *Firewall Policy*.
   b. Click *Create New*.
   c. In the *Name* field, enter FTP Access.
   d. From the *Incoming Interface* dropdown list, select *port10*.
   e. From the *Outgoing Interface* dropdown list, select *port9*.
   f. For the *Source* and *Destination* fields, select *all* and *FTP_Server*, respectively.
   g. From the *Schedule* dropdown list, select *always*.
   h. For the *Service* field, select *ALL*.
   i. Click *OK*.
2. Create the per-IP traffic shaper:
   a. Go to *Policy & Objects* > *Traffic Shapers*.
   b. Click *Create New*.
   c. For *Type*, select *Per-IP*.
   d. In the *Name* field, enter FTP_Max_1M. This shaper is for VoIP traffic.
   e. Enable *Max Bandwidth* and enter 1000. This equates to 1 Mbps.
   f. Enable *Max Concurrent Connections* and enter 10. This means that each user can have up to ten concurrent connections to the FTP server.
   g. Click *OK*.
3. Create a firewall shaping policy:
   a. Go to *Policy & Objects* > *Traffic Shaping Policy*.
   b. Click *Create New*.
   c. In the *Name* field, enter FTP speed 1M.
   d. For the *Source* fields, select the users that need to access the FTP server.
   e. For the *Destination* field, select *FTP_Server*.
   f. For the *Service* field, select *ALL*.
   g. For the *Outgoing Interface* field, select *port9*.
   h. Enable *Per-IP shaper*. Select FTP_Max_1M from the dropdown list.
   i. Click *OK*.

To configure a per-IP traffic shaper using the FortiOS CLI:

1. Create a firewall policy:

```
config firewall policy
   edit 1
      set name "FTP Access"
```


Netskope Exhibit 1017

```
      set srcintf "port10"
      set dstintf "port9"
      set srcaddr "all"
      set dstaddr "FTP_Server"
      set action accept
      set schedule "always"
      set service "ALL"
      set fsso disable
      set nat enable
   next
end
```

2. Create the per-IP traffic shaper:

```
config firewall shaper per-ip-shaper
   edit "FTP_Max_1M"
      set max-bandwidth 1000
      set max-concurrent-session 10
   next
end
```

3. Create a firewall shaping policy:

```
config firewall shaping-policy
   edit 1
      set name "FTP speed 1M"
      set service "ALL"
      set dstintf "port9"
      set per-ip-shaper "FTP_Max_1M"
      set srcaddr "PC1" "WinPC" "PC2"
      set dstaddr "FTP_Server"
   next
end
```

To troubleshoot per-IP traffic shapers:

1. To check if specific traffic is attached to the correct traffic shaper, run the `diagnose firewall iprope list 100015` command. The example output shows the traffic attached to the FTP_Max_1M shaper:

```
# diagnose firewall iprope list 100015

policy index=3 uuid_idx=0 action=accept
flag (0):
shapers: per-ip=FTP_Max_1M
cos_fwd=0 cos_rev=0
group=00100015 av=00000000 au=00000000 split=00000000
host=2 chk_client_info=0x0 app_list=0 ips_view=0
misc=0 dd_type=0 dd_mode=0
zone(1): 0 -> zone(1): 38
source(3): 10.1.100.11-10.1.100.11, uuid_idx=30, 10.1.100.143-10.1.100.143, uuid_idx=32,
     10.1.100.22-10.1.100.22, uuid_idx=31,
dest(1): 172.16.200.55-172.16.200.55, uuid_idx=89,
service(1):
   [0:0x0:0/(0,65535)->(0,65535)] helper:auto
```

2. To check if the correct traffic shaper is applied to the session, run the `diagnose sys session list` command. The example output shows that the FTP_Max_1M shaper is applied to the session:

```
# diagnose sys session list
session info: proto=6 proto_state=01 duration=36 expire=3567 timeout=3600 flags=00000000
     sockflag=00000000 sockport=0 av_idx=0 use=4
origin-shaper=
reply-shaper=
```


Netskope Exhibit 1017

```
per_ip_shaper=FTP_Max_1M
class_id=0 shaping_policy_id=3 ha_id=0 policy_dir=0 tunnel=/ helper=ftp vlan_cos=0/255
state=may_dirty per_ip npu npd mif route_preserve
statistic(bytes/packets/allow_err): org=506/9/1 reply=416/6/1 tuples=2
tx speed(Bps/kbps): 0/0 rx speed(Bps/kbps): 0/0
orgin->sink: org pre->post, reply pre->post dev=39->38/38->39 gwy=172.16.200.55/0.0.0.0
hook=post dir=org act=snat 10.1.100.11:58275->172.16.200.55:21(172.16.200.1:58275)
hook=pre dir=reply act=dnat 172.16.200.55:21->172.16.200.1:58275(10.1.100.11:58275)
pos/(before,after) 0/(0,0), 0/(0,0)
misc=0 policy_id=1 auth_info=0 chk_client_info=0 vd=2
serial=0000211a tos=ff/ff app_list=0 app=0 url_cat=0
rpdb_link_id = 00000000
dd_type=0 dd_mode=0
npu_state=0x100000
npu info: flag=0x00/0x00, offload=0/0, ips_offload=0/0, epid=0/0, ipid=0/0,
     vlan=0x0000/0x0000
vlifid=0/0, vtag_in=0x0000/0x0000 in_npu=0/0, out_npu=0/0, fwd_en=0/0, qid=0/0
no_ofld_reason: offload-denied helper
```

3. To check statuses of per-IP traffic shapers, run the `diagnose firewall shaper per-ip-shaper list` command. The output should resemble the following:

```
# diagnose firewall shaper per-ip-shaper list

name FTP_Max_1M
maximum-bandwidth 125 KB/sec
maximum-concurrent-session 10
tos ff/ff
packets dropped 0
bytes dropped 0
addr=10.1.100.11 status: bps=0 ses=3
```

Type of Service-based prioritization and policy-based traffic shaping

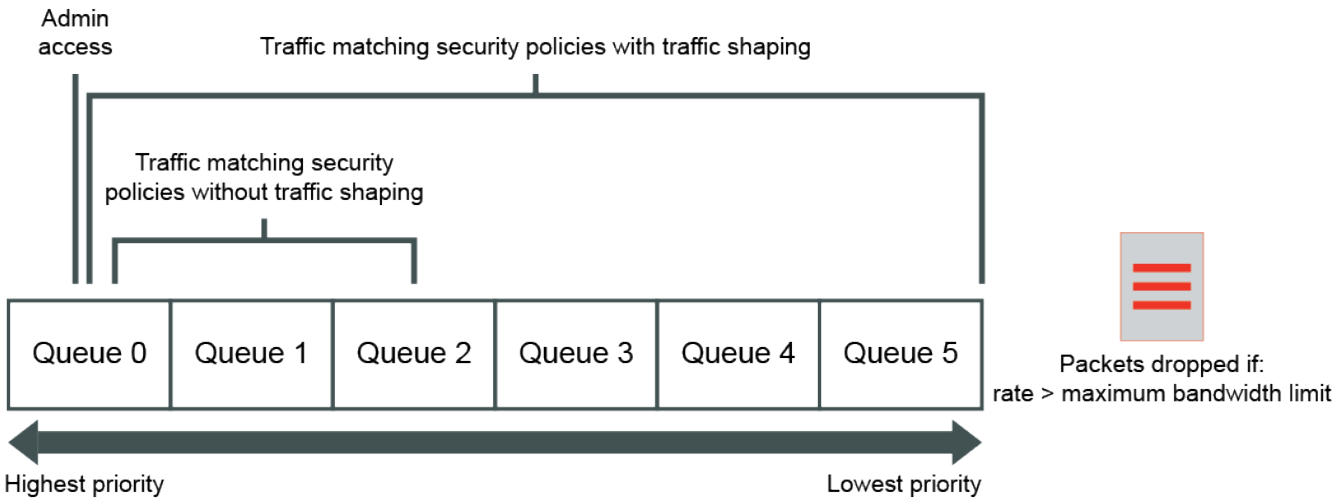
Priority queues

After packet acceptance, FortiOS classifies traffic and may apply Quality of Service (QoS) techniques, such as prioritization and traffic shaping. Traffic shaping consists of a mixture of traffic policing to enforce bandwidth limits and priority queue adjustment to assist packets in achieving the guaranteed rate.

If you have configured prioritization, FortiOS prioritizes egressing packets by distributing them among first in first out (FIFO) queues associated with each possible priority number. Each physical interface has six priority queues. Virtual interfaces use the priority queues of the physical interface that they are bound to.

The physical interface's six queues are queue 0 to 5, where queue 0 is the highest priority queue. You might observe that your traffic uses only a subset of those six queues. For example, some traffic may always use a certain queue number. Queuing may also vary by the packet rate or mixture of services. Some queue numbers may only be used by through traffic for which you have configured traffic shaping in the security policy that applies to that traffic session.

Netskope Exhibit 1017



- Administrative access traffic always uses queue 0.
- Traffic matching firewall policies without traffic shaping may use queue 0, 1, or 2. The queue is selected based on the priority value you have configured for packets with that ToS bit value, if you have configured ToS-based priorities.
- Traffic matching firewall shaping policies with traffic shaping enabled can use any queue. The queue is selected based on whether the packet rate is currently below the guaranteed bandwidth (queue 0), or above the guaranteed bandwidth. Packets at rates greater than the maximum bandwidth limit are dropped.

Priority types

Packets can be assigned a priority in one of three types:

- On entering ingress – for packets flowing through the firewall.
- Upon generation – for packets generated by the firewall (including packets generated due to AV proxying).
- On passing through a firewall policy – for packets passing through a firewall policy (firewall shaping policy) that has a traffic shaper defined.

ToS priority

The first and second types, ingress priority and priority for generated packets, are controlled by two different CLI settings:

```
config system global
    set traffic-priority-level {high | medium | low}
end
config system tos-based-priority
    edit 1
        set tos [0-15]      <---- type of service bit in the IP datagram header with a value
between 0 and 15
        set priority (high | medium | low)      <---- priority of this type of service
    next
end
```

Each priority level is mapped to a value as follows:

ToS priority	Value
High	0
Medium	1
Low	2



ToS-based traffic prioritization cannot be used to apply bandwidth limits and guarantees, but can be used to prioritize traffic at per-packet levels.

Example

In the following example configuration, packets with ToS bit values of 10 are prioritized as medium and packets with ToS bit values of 20 are prioritized as high. All the other traffic is prioritized as low.

```
config system global
    set traffic-priority-level low
end
config system tos-based-priority
    edit 1
        set tos 10
        set priority medium
    next
    edit 2
        set tos 20
        set priority high
    next
end
```

Firewall shaping policy priority

You can enable traffic shaping in a firewall shaping policy. In the shared traffic shaper, you can set the firewall priority to high, medium, or low:

```
config firewall shaper traffic-shaper
    edit 1
        set priority {high | medium | low}
    next
end
```

As the priority in a traffic shaper is set to high by default, you must set some traffic at a lower priority to see results. Each priority level is mapped to a value as follows:

Firewall policy priority	Value
High (default)	1
Medium	2
Low	3

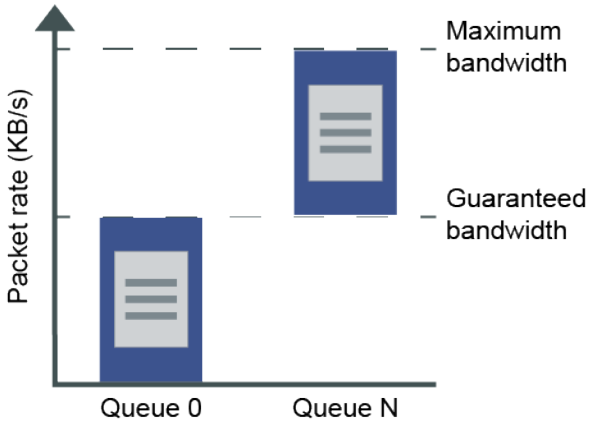
Combination of two priority types

To combine the two priority types, the global or ingress ToS-based priority value is combined with the firewall policy priority value:

ToS priority (0, 1, 2) + policy priority (1, 2, 3) = total priority (queue number)

Consider the following scenarios:

- If the current packet rate is less than the guaranteed bandwidth, packets use priority queue 0. Packet priority is 0.
- If the current packet rate exceeds the maximum bandwidth, excess packets are dropped.
- If the current packet rate is greater than the guaranteed bandwidth but less than the maximum bandwidth, FortiOS assigns a priority queue by adding the ToS-based priority and the firewall priority.
  For example, if you have enabled traffic shaping in the security policy and the security policy's traffic priority is low (value 3), and the priority normally applied to packets with that ToS bit is medium (value 1), the packets have a total packet priority of 4, and use priority queue 4.



Interface-based traffic shaping profile

A traffic shaping policy can be used for interface-based traffic shaping by organizing traffic into 30 class IDs. The shaping profile defines the percentage of the interface bandwidth that is allocated to each class. Each traffic class ID is shaped to the assigned speed according to the outgoing bandwidth limit configured to the interface.

Traffic classification

A shaping policy classifies traffic and organizes it into different class IDs, based on matching criteria. For traffic matching a criteria, you can choose to put it into 30 different shaping classes, identified by class ID 2 to 31.

You must select an outgoing interface for the traffic. The shaping policy is only applied when the traffic goes to one of the selected outgoing interfaces.

Criterion	Description
Source	• Address: match the source address of the traffic to the selected address or address group. • User: use the user credentials of the traffic to match the selected user or user group. At least one address, address group, or internet service must also be

Criterion	Description
	selected. • Internet service: match the traffic to the selected internet service. Internet services cannot be used if addresses or address or groups are used.
Destination	• Address: match the destination address of the traffic to the selected address or address group. • Internet service: match the traffic to the selected internet service. Internet services cannot be used if addresses or address or groups are used.
Schedule	Match the current date and time to the selected schedule. You can select a one-time schedule, recurring schedule, or schedule group. This setting is optional.
Service	Match the service of the traffic to the selected service or service group.
Application	Match the application of the traffic to the selected application, application category, or application group. Application control must be enabled in the related firewall policy to know the application of the traffic. See Application control on page 963 for more information.
URL category	Match the URL of the traffic to the selected URL category. Web filter must be enabled in the related firewall policy to know the URL of the traffic. See Web filter on page 894 for more information.



When multiple items are selected in one criterion, it is considered a match when traffic matches any one of them.

Traffic prioritization

Shaping profiles define how different shaping classes of traffic are prioritized. For each class, you can define three prioritization strategies: guaranteed bandwidth, maximum bandwidth, and priority.

For each shaping profile, a default shaping class must be defined. Traffic is prioritized based on the default shaping group in the following two circumstances:

- All traffic to the outgoing interface that does not match to any shaping policy
- Traffic with a shaping group that is not defined in a shaping profile

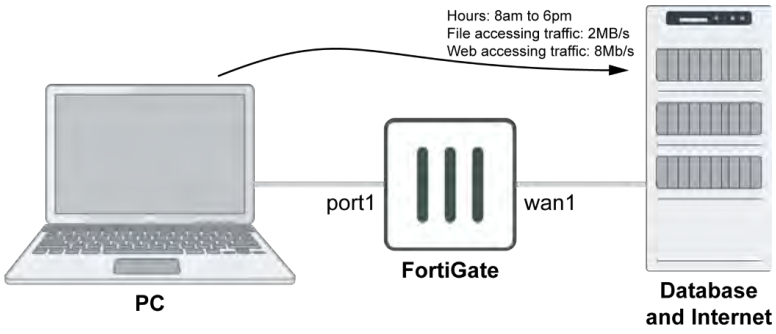
Prioritization strategy	Description
Guaranteed bandwidth	The percentage of the link speed that is reserved for the shaping group. The total guaranteed bandwidth for all shaping groups cannot exceed 100%.
Maximum bandwidth	The maximum percentage of the link speed that the shaping group can use.
Priority	The shaping class priority: top, critical, high, medium, or low. When groups are competing for bandwidth on the interface, the group with the higher priority wins.

Applying a shaping profile to an interface

Traffic shaping is accomplished by configuring the outgoing bandwidth and outgoing shaping profile on an interface. The shaping profile uses the outgoing bandwidth of the interface as the maximum link speed, and it only works when the outgoing bandwidth is configured.

This example shows how to apply interface-based traffic shaping to web and file accessing traffic according to a schedule:

- The link speed of the wan1 interface is 10 Mb/s.
- File access can use up to 2 Mb/s from 8:00 AM to 6:00 PM.
- Web access can use 8 Mb/s from 8:00 AM to 6:00 PM.



Putting the traffic into shaping classes

To create a recurring schedule in the GUI:

1. Go to *Policy & Objects* > *Schedules*.
2. Click *Create New* > *Schedule*.
3. Configure a recurring schedule called *Day_Hours* for everyday from 8:00 AM to 6:00 PM.
4. Click *OK*.

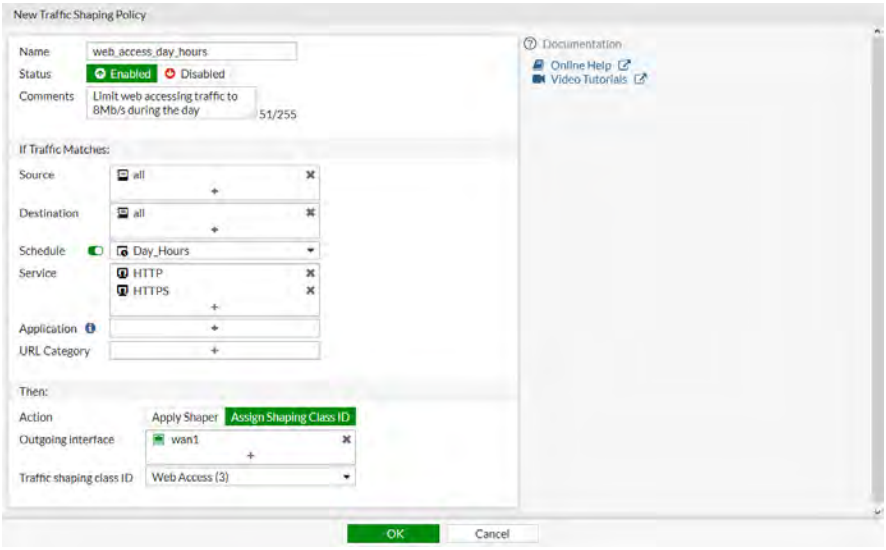
To create a traffic shaping policy and class ID for the web accessing traffic in the GUI:

1. Go to *Policy & Objects* > *Traffic Shaping Policy* and click *Create New*.
2. Enter a name for the policy, such as *web_access_day_hours*.
3. Enable *Schedule* and select the schedule you just created.
4. Set *Service* to web accessing services, such as *HTTP* and *HTTPS*.
5. Set *Action* to *Assign Shaping Class ID*, and *Outgoing interface* to *wan1*
6. Click the *Traffic shapping class ID* drop down then click *Create*.
7. Enter a value for the *ID* (integer) and a description for the *Name*, such as *Web Access*.



8. Click *OK*.

9. Set *Traffic shaping class ID* to the just created class ID.



10. Configure the remaining settings as required.
11. Click *OK*.

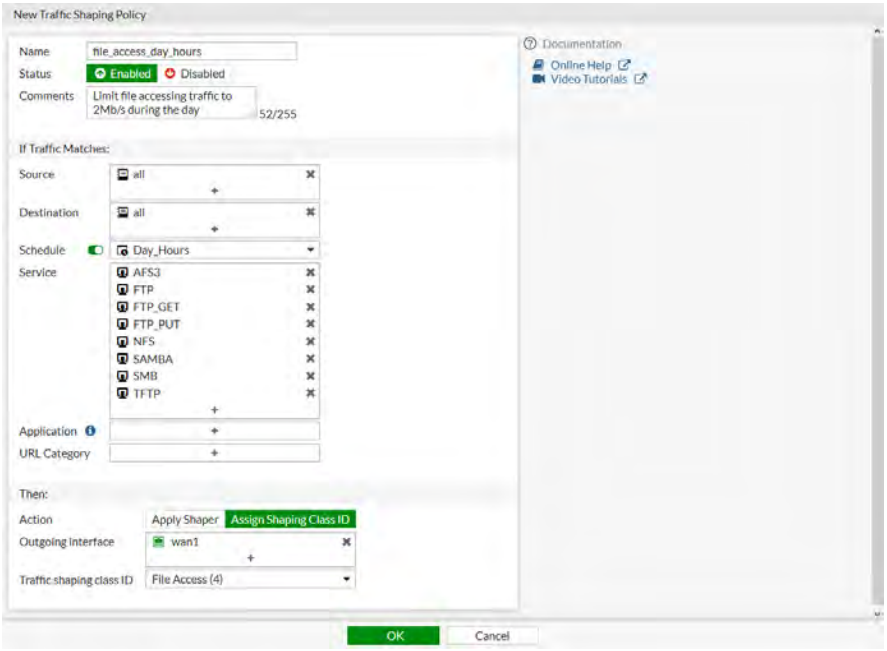
To create a traffic shaping policy and class ID for the file accessing traffic in the GUI:

1. Go to *Policy & Objects* > *Traffic Shaping Policy* and click *Create New*.
2. Enter a name for the policy, such as *file_access_day_hours*.
3. Enable *Schedule* and select the schedule you just created.
4. Set *Service* to file accessing services, such as *ASF3*, *FTP* and *SMB*.
5. Set *Action* to *Assign Shaping Class ID*, and *Outgoing interface* to *wan1*
6. Click the *Traffic shapping class ID* drop down then click *Create*.
7. Enter a value for the *ID* (integer) and a description for the *Name*, such as *File Access*.



8. Click *OK*.

9. Set *Traffic shaping class ID* to the just created class ID.



10. Configure the remaining settings as required.
11. Click *OK*.

To put the traffic into shaping classes in the CLI:

1. Create a recurring schedule:

```
config firewall schedule recurring
    edit "Day_Hours"
        set start 08:00
        set end 18:00
        set day sunday monday tuesday wednesday thursday friday saturday
    next
end
```

2. Create the traffic class IDs:

```
config firewall traffic-class
    edit 3
        set class-name "Web Access"
    next
    edit 4
        set class-name "File Access"
    next
end
```

3. Create the web and file accessing traffic shaping policies:

```
config firewall shaping-policy
    edit 2
        set name "web_access_day_hours"
        set comment "Limit web accessing traffic to 8Mb/s in day time"
        set service "HTTP" "HTTPS"
        set schedule "Day_Hours"
```

```
            set dstintf "wan1"
            set class-id 3
            set srcaddr "all"
            set dstaddr "all"
        next
        edit 3
            set name "file_access_day_hours"
            set comment "Limit file accessing traffic to 2Mb/s during the day"
            set service "AFS3" "FTP" "FTP_GET" "FTP_PUT" "NFS" "SAMBA" "SMB" "TFTP"
            set schedule "Day_Hours"
            set dstintf "wan1"
            set class-id 4
            set srcaddr "all"
            set dstaddr "all"
        next
    end
```

Allocating bandwidth to the shaping classes

A traffic shaping profile defines the guaranteed and maximum bandwidths each class receives. In this example, file access can use up to 2 Mb/s and web access can use 8 Mb/s from 8:00 AM to 6:00 PM.

To create a traffic shaping profile using the GUI:

1. Go to *Policy & Objects* > *Traffic Shaping Profile* and click *Create New*.
2. Enter a name for the profile, such as *Day_Hours_Profile*.
3. Configure a default traffic shaping class:
   This class has a high priority, meaning that when the other classes have reached their guaranteed bandwidths, this default class will use the rest of the available bandwidth.
   a. In the *Traffic Shaping Classes* table click *Create New*.
   b. Click the *Traffic shapping class ID* drop down then click *Create*.
   c. Enter a name for the class, such as *Default Access*.



   d. Click *OK*.
   e. Set *Traffic shaping class ID* to the just created class ID.

f. Configure the following settings, then click *OK*:

Guaranteed bandwidth	30
Maximum bandwidth	100
Priority	High



4. Configure a web accessing traffic shaping class:
When other types of traffic are competing for bandwidth, this class is guaranteed to 6 Mb/s, or 60% of the bandwidth.
   a. In the *Traffic Shaping Classes* table click *Create New*.
   b. Configure the following settings, then click *OK*:

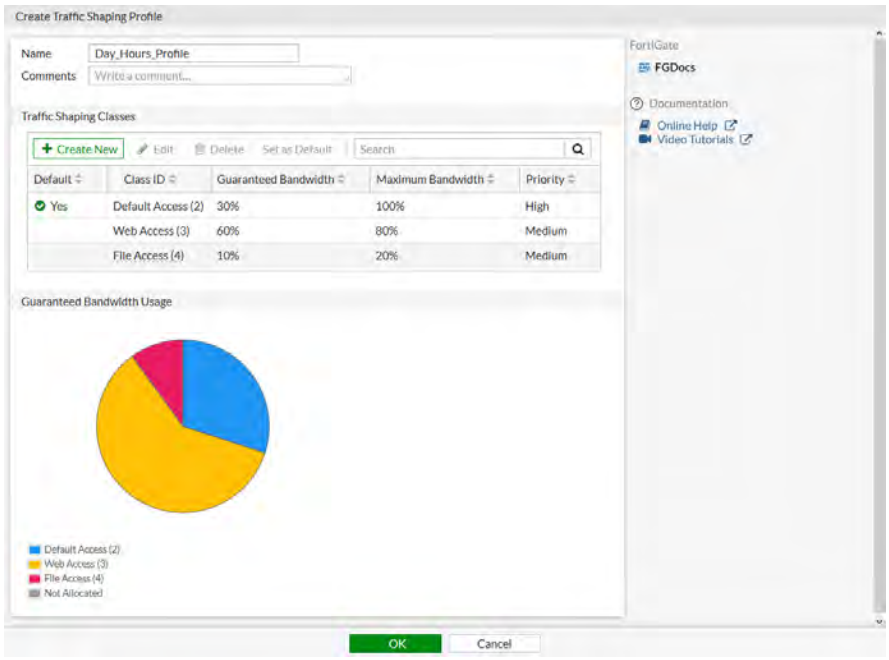
Traffic shaping class ID	Web Access
Guaranteed bandwidth	60
Maximum bandwidth	80
Priority	High



5. Configure a file accessing traffic shaping class:
When other types of traffic are competing for bandwidth, this group is guaranteed to 1 Mb/s, or 10% of the bandwidth.
   a. In the *Traffic Shaping Classes* table click *Create New*.
   b. Configure the following settings, then click *OK*:

Traffic shaping class ID	File Access
Guaranteed bandwidth	10
Maximum bandwidth	20
Priority	High





6. Click *OK*.

To create a traffic shaping profile using the CLI:

```
config firewall shaping-profile
    edit "Day_Hours_Profile"
        set default-class-id 2
        config shaping-entries
            edit 1
                set class-id 2
                set guaranteed-bandwidth-percentage 30
                set maximum-bandwidth-percentage 100
            next
            edit 2
                set class-id 3
                set priority medium
                set guaranteed-bandwidth-percentage 60
                set maximum-bandwidth-percentage 80
            next
            edit 3
                set class-id 4
                set priority medium
                set guaranteed-bandwidth-percentage 10
                set maximum-bandwidth-percentage 20
            next
        end
```

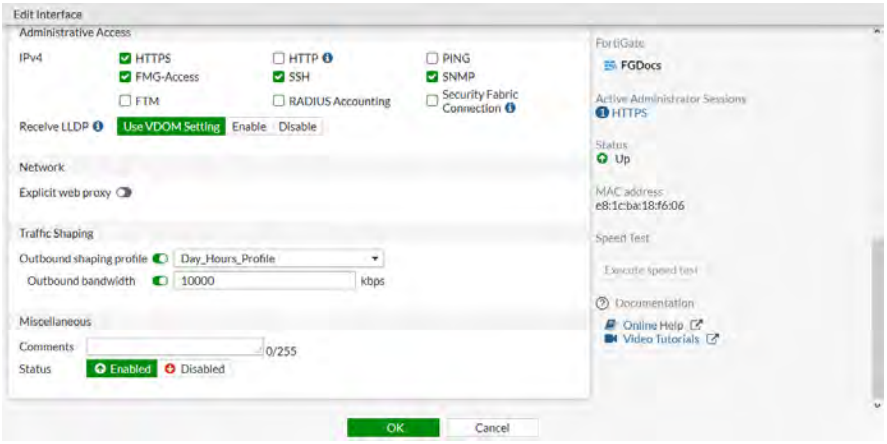
```
    next
end
```

Defining the available bandwidth on an interface

In this example, the link speed of the wan1 interface is 10 Mb/s.

To set the bandwidth of the wan1 interface in the GUI:

1. Go to *Network* > *Interfaces*.
2. Edit the wan1 interface.
3. Under Traffic Shaping, enable *Outbound shaping profile* and select the profile that you just created, *Day_Hours_Profile*.
4. Enable *Outbound Bandwidth* and set it to *10000* Kbps.



5. Click *OK*.

To set the bandwidth of the wan1 interface in the CLI:

```
config system interface
    edit "wan1"
        ....
        set egress-shaping-profile "Day_Hours_Profile"
        set outbandwidth 10000
        ....
    next
end
```

Diagnose commands

To check that the specific traffic is put into the correct shaping group or class ID:

```
# diagnose firewall iprope list 100015
```

To check the speed limit for each class ID on an interface:

```
# diagnose netlink interface list wan1
```

Interface-based traffic shaping with NP acceleration

Interface-based traffic shaping with NP acceleration is supported on some devices.

An administrator configures the WAN interface's maximum outbound bandwidth and, based on that, creates a traffic shaping profile with a percentage based shaper. This allows for proper QoS and traffic shaping. VLAN interfaces are not supported.



This feature is supported on FortiGate 600E, 500E, and 300E models.

To configure interface-based traffic shaping:

1. Enable NPU offloading when doing interface-based traffic shaping according to the egress-shaping-profile:

```
config system npu
    set intf-shaping-offload enable
end
```

2. Configure shaping profiles:

```
config firewall shaping-profile
    edit "sdwan"
        set default-class-id 4
        config shaping-entries
            edit 1
                set class-id 4
                set guaranteed-bandwidth-percentage 3
                set maximum-bandwidth-percentage 5
            next
            edit 2
                set class-id 3
                set priority medium
                set guaranteed-bandwidth-percentage 50
                set maximum-bandwidth-percentage 100
            next
            edit 3
                set class-id 2
                set priority low
                set guaranteed-bandwidth-percentage 1
                set maximum-bandwidth-percentage 5
            next
        end
    next
end
```

   The class number is limited to 16.

3. Configure a traffic shaper and shaping policy:

```
config firewall shaper traffic-shaper
    edit "Transactional"
        set priority medium
    next
end
```

```
config firewall shaping-policy
    edit 1
        set service "ALL"
        set dstintf "any"
        set traffic-shaper "Transactional"
        set class-id 3
        set srcaddr "all"
        set dstaddr "all"
    next
end
```

4. Apply the egress shaping profile on the interface:

```
config system interface
    edit "port2"
        set vdom "root"
        set ip 10.1.100.23 255.255.255.0
        set allowaccess ping
        set type physical
        set outbandwidth 500
        set egress-shaping-profile "sdwan"
        set snmp-index 4
    next
end
```

5. Configure a firewall policy:

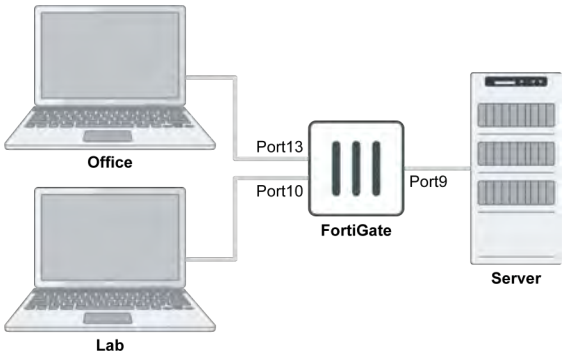
```
config firewall policy
    edit 3
        set srcintf "port2"
        set dstintf "port1"
        set srcaddr "all"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ALL"
        set logtraffic all
        set nat enable
    next
end
```

Classifying traffic by source interface

In firewall shaping policies, you can classify traffic by source interface with the following command:

```
config firewall shaping-policy
  edit 1
    set srcintf <interface_name>
    ......
  next
end
```

Sample configuration



For this example, there are two shaping policies:

- Policy 1 is for traffic from the Office to the Server, with the speed limited to 5 MB/s.
- Policy 2 is for traffic from the Lab to the Server, with the speed limited to 1 MB/s.

To configure the traffic shaping policy:

```
config firewall shaping-policy
    edit 1
        set name "Office_Speed_5MB"
        set service "ALL"
        set srcintf "port13"
        set dstintf "port9"
        set traffic-shaper "5MB/s"
        set traffic-shaper-reverse "5MB/s"
        set srcaddr "all"
        set dstaddr "all"
    next
    edit 2
        set name "Lab_Speed_1MB"
        set service "ALL"
        set srcintf "port10"
        set dstintf "port9"
        set traffic-shaper "1MB/s"
        set traffic-shaper-reverse "1MB/s"
        set srcaddr "all"
        set dstaddr "all"
    next
end
```

Configuring traffic class IDs

As of FortiOS 6.2.2, you can configure traffic class IDs with a descriptive name in the GUI or CLI. Class IDs can help you correlate traffic shaping policy and profile entries.

GUI configurations

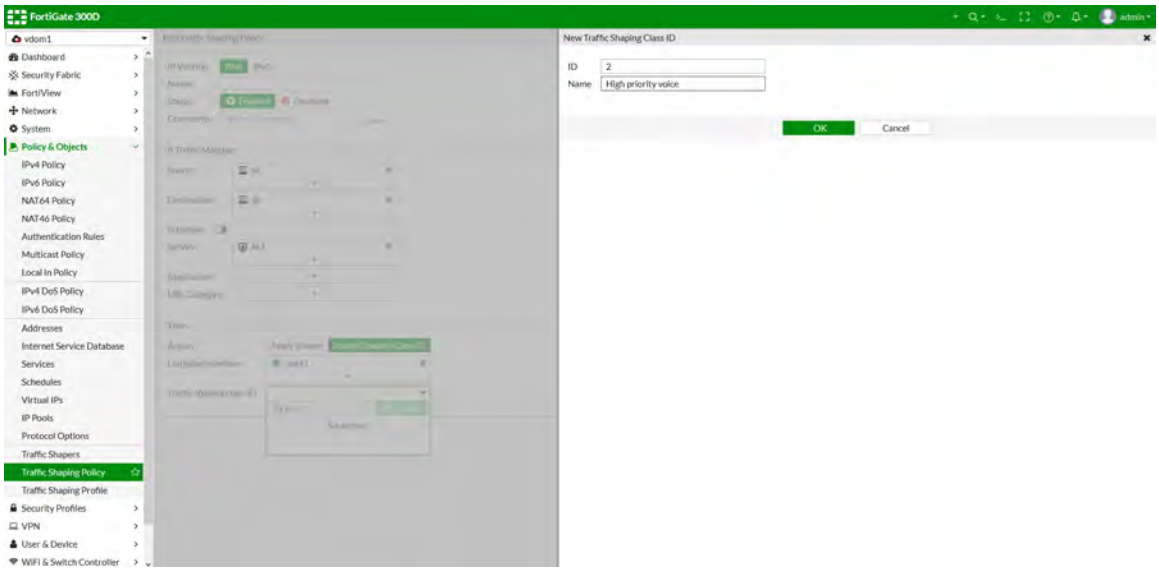
Within the GUI, there are three locations to configure the traffic class ID:

- Traffic shaping policy
- Traffic shaping profile
- Interface

Assign Shaping Class ID replaces the *Assign Group* functionality in earlier versions of FortiOS.

To configure the traffic class ID in a traffic shaping policy:

1. Go to *Policy & Objects* > *Traffic Shaping Policy*.
2. Edit an existing policy, or create a new one.
3. In the *Then: Action* section, click *Assign Shaping Class ID*.
4. In *Traffic shaping class ID*, click *Create*.
5. Enter a value for the *ID* (integer) and a description for the *Name*.
6. Click *OK* to save the class ID.

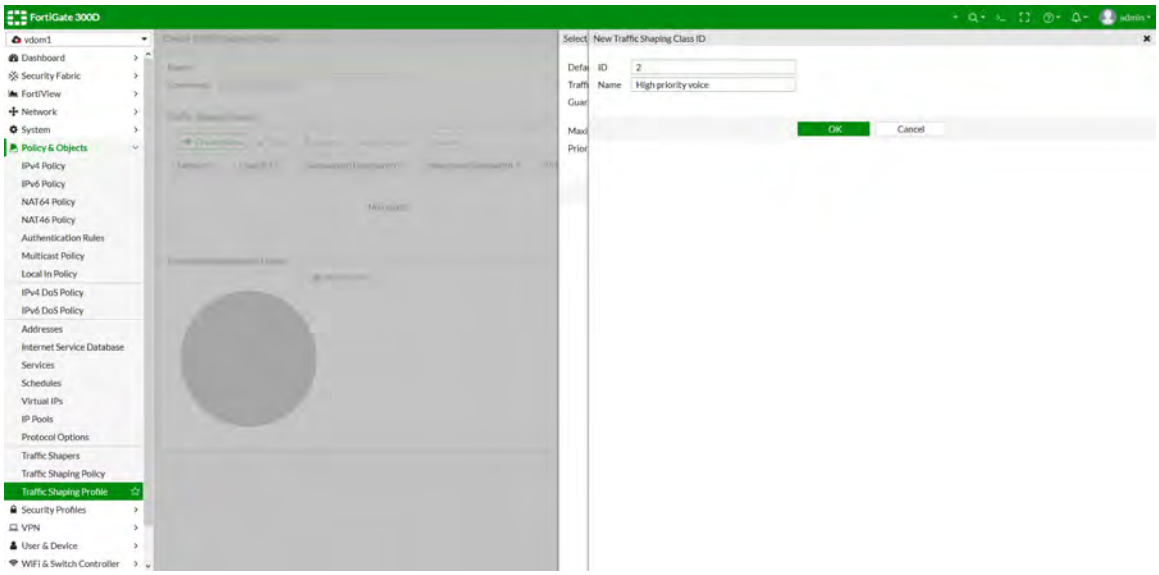


7. Configure the rest of the policy as needed.
8. Click *OK* to save the policy.

To configure the traffic class ID in a traffic shaping profile:

1. Go to *Policy & Objects* > *Traffic Shaping Profile*.
2. Edit an existing profile, or create a new one.
3. In the *Traffic Shaping Classes* section, click *Create New*. The *Select Traffic Shaping Class ID* window opens.
4. Click *Create*. The *New Traffic Shaping Class ID* window opens.
5. Enter a value for the *ID* (integer) and a description for the *Name*.

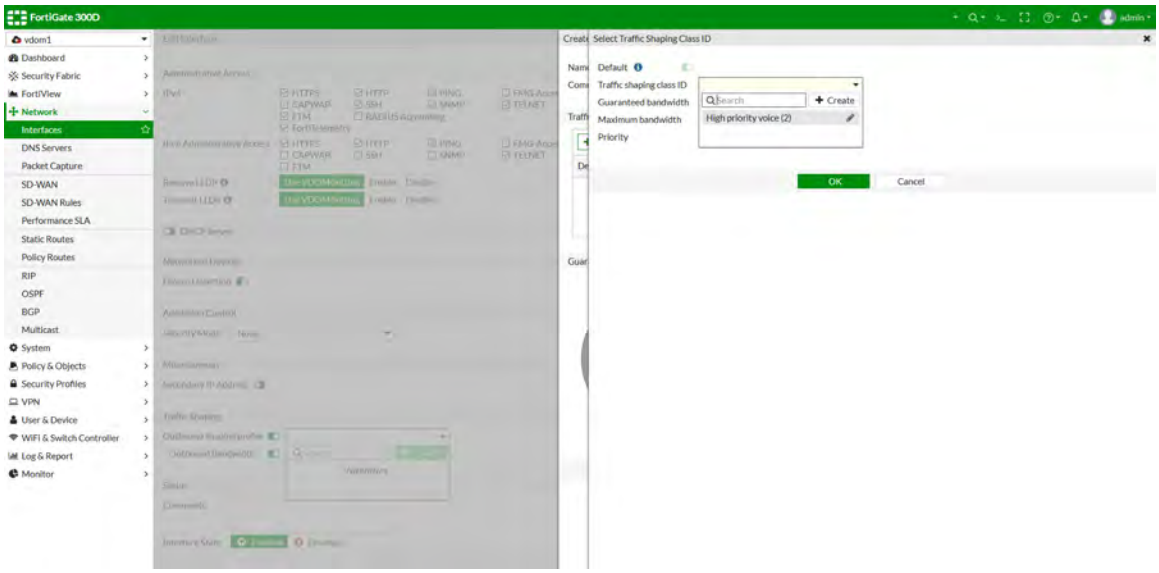
6. Click *OK* to save the class ID.



7. Click *OK* to add the class ID.
8. Configure the rest of the profile as needed.
9. Click *OK* to save the profile.

To configure the traffic class ID in an interface:

1. Go to *Network* > *Interfaces*.
2. Edit an existing interface, or create a new one.
3. In the *Traffic Shaping* section, enable *Outbound shaping profile* and *Outbound Bandwidth*.
4. Click *Create*. The *Create Traffic Shaping Profile* window opens.
5. Click *Create New*. The *Select Traffic Shaping Class ID* window opens.
6. Select an existing class ID, or create a new one.
7. Click *OK* to save the class ID.



8. Click *OK* to add the class ID .

Netskope Exhibit 1017

9. Configure the rest of the interface as needed.
10. Click *OK* to save the interface.

CLI configuration

To configure the traffic class ID in the CLI:

```
config firewall traffic-class
    edit 2
        set class-name "High priority voice."
    next
    ...
end
```

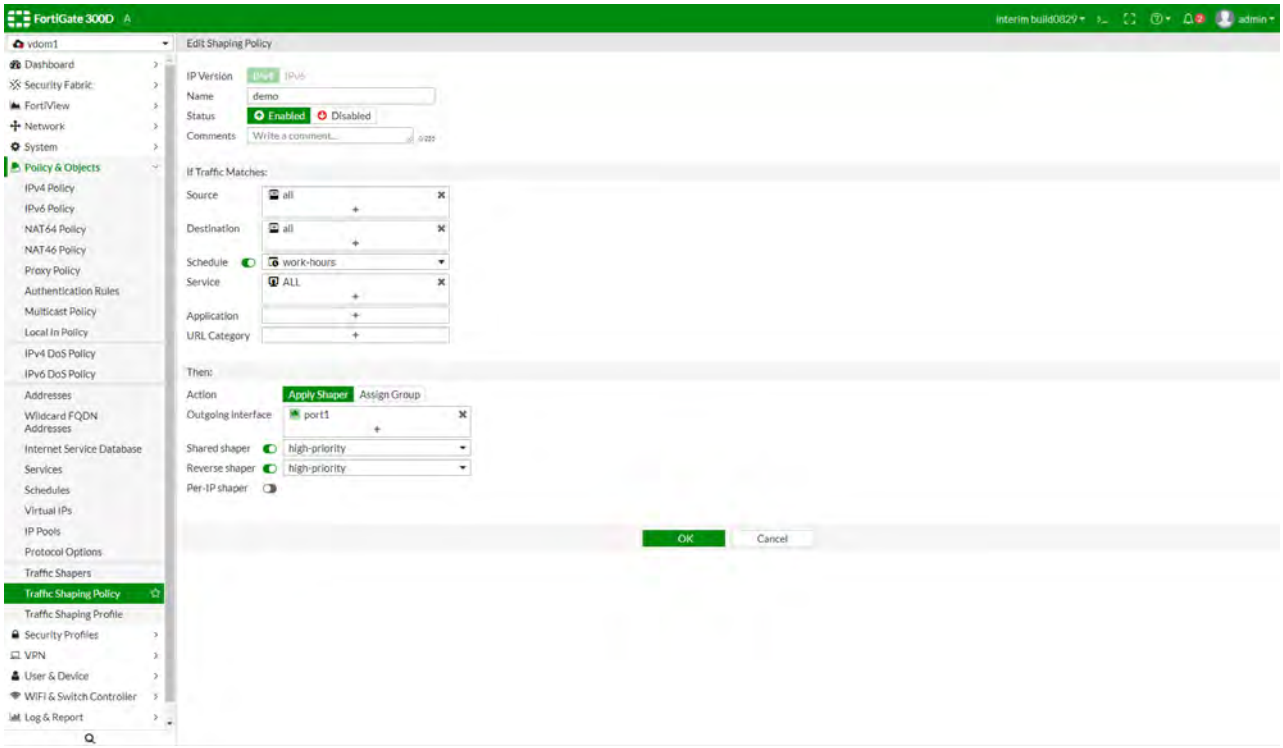
Traffic shaping schedules

In a shaping policy, there are many matching criteria available for administrators to match a specific traffic and apply a traffic shaper or shaping group to the traffic, including using schedules. This feature gives shaping policy the ability to apply different shaping profiles at different times. Administrators can select a one-time schedule, recurring schedule, or schedule group.

Schedule is not a mandatory setting. If it is not set, then the current date and time are not used to match the traffic.

To configure a traffic shaping policy in the GUI:

1. Navigate to *Policy & Objects* > *Traffic Shaping Policy*.
2. Create or edit a *Traffic Shaping Policy*.
3. Enable *Schedule* and select a schedule option.



4. Configure other options and click *OK*.

To configure a traffic shaping policy in the CLI:

```
config firewall schedule recurring
    edit "work-hours"
        set start 07:00
        set end 20:00
        set day monday tuesday wednesday thursday friday
    next
end

config firewall shaping-policy
    edit 1
        set name "demo"
        set service "ALL"
        set schedule "work-hours"  <<< Can select schedule from one-time schedule, recurring
schedule or schedule group
        set dstintf "port1"
        set traffic-shaper "high-priority"
        set traffic-shaper-reverse "high-priority"
        set srcaddr "all"
        set dstaddr "all"
    next
end
```

To troubleshoot a traffic shaping policy in the CLI:

The selected schedule is listed in the `iprope`.

```
diagnose firewall iprope list 100015

policy index=1 uuid_idx=0 action=accept
flag (0):
schedule(work-hours)
shapers: orig=high-priority(2/0/134217728) reply=high-priority(2/0/134217728)
cos_fwd=0  cos_rev=0
group=00100015 av=00000000 au=00000000 split=00000000
host=1 chk_client_info=0x0 app_list=0 ips_view=0
misc=0 dd_type=0 dd_mode=0
zone(1): 0 -> zone(1): 9
source(1): 0.0.0.0-255.255.255.255, uuid_idx=28,
dest(1): 0.0.0.0-255.255.255.255, uuid_idx=28,
service(1):
    [0:0x0:0/(0,65535)->(0,65535)] helper:auto
```

QoS assignment and rate limiting for quarantined VLANs

When devices are quarantined, they are isolated from the rest of the network. However, they can still impact the network if not controlled beyond isolation. A quarantined host, which offers heavy traffic, could congest the network and create a DOS-style reduction in service to authorized hosts.

Within the quarantined VLAN, two restrictions are available within the network:

Netskope Exhibit 1017

- Traffic policing (also known as rate limiting)
- QoS (Quality of Service) assignment (also known as priority assignment)

Each quarantined host's traffic can be subject to rate limiting and priority adjustment. This reduces the impact that any quarantined host can have on authorized traffic on the network.

To configure QoS assignment and rate limiting for quarantined VLANs:

1. Configure a traffic policy, or use the default "quarantine" policy:

```
config switch-controller traffic-policy
    edit "quarantine"
        set description "Rate control for quarantined traffic"
        set guaranteed-bandwidth 163840
        set guaranteed-burst 8192
        set maximum-burst 163840
        set cos-queue 0
    next
end
```

2. Configure an interface:

```
config system interface
    edit "qtn.aggr1"
        set vdom "root"
        set ip 10.254.254.254 255.255.255.0
        set description "Quarantine VLAN"
        set security-mode captive-portal
        set replacemsg-override-group "auth-intf-qtn.aggr1"
        set device-identification enable
        set snmp-index 30
        set switch-controller-access-vlan enable
        set switch-controller-traffic-policy "quarantine"
        set color 6
        set interface "aggr1"
        set vlanid 4093
    next
end
```

By default, `switch-controller-traffic-policy` is empty. You need to apply the necessary traffic policy (not only limited to "quarantine").

Weighted random early detection queuing

You can use the weighted random early detection (WRED) queuing function within traffic shaping.

This topic includes three parts:

- Traffic shaping with queuing
- Burst control in queuing mode
- Multi-stage DSCP marking and class ID in traffic shapers

You cannot configure or view WRED in the GUI; you must use the CLI.



WRED is not supported when traffic is offloaded to an NPU.

Traffic shaping with queuing

Traffic shaping has a queuing option. Use this option to fine-tune the queue by setting the profile queue size or performing random early drop (RED) according to queue usage.

This example shows setting the profile queue size limit to 5 so that the queue can contain a maximum of five packets and more packets are dropped.

To set the profile queue size limit:

```
config firewall shaping-profile
    edit "profile"
        set type queuing
        set default-class-id 31
        config shaping-entries
            edit 31
                set class-id 31
                set guaranteed-bandwidth-percentage 5
                set maximum-bandwidth-percentage 10
                set limit 5  <range from 5 to 10000; default: 1000>
            next
        end
    next
end
```

This example shows performing RED according to queue usage by setting `red-probability`, `min`, and `max`. Setting `red-probability` to `10` means start to drop packets when queue usage reaches the `min` setting. When queue usage reaches the `max` setting, drop 10% of the packets.

- Level 1: when queue is less than `min` packets, drop 0% of packets.
- Level 2: when queue reaches `min` packets, start to drop packets.
- Level 3: when queue usage is between `min` and `max` packets, drop 0–10% of packets by proportion.
- Level 4: when queue (average queue size) is more than `max` packets, drop 100% of packets.

To set RED according to queue usage:

```
config firewall shaping-profile
    edit "profile"
        set type queuing
        set default-class-id 31
        config shaping-entries
            edit 31
                set class-id 31
                set guaranteed-bandwidth-percentage 5
                set maximum-bandwidth-percentage 10
                set red-probability 10 <range from 0 to 20; default: 0 no drop>
                set min 100 <range from 3 to 3000>
                set max 300 <range from 3 to 3000>
            next
        end
```

```
    next
end
```

To troubleshoot this function, use the following diagnose commands:

```
diagnose netlink intf-class list <intf>
diagnose netlink intf-qdisc list <intf>
```

Burst control in queuing mode

In a hierarchical token bucket (HTB) algorithm, each traffic class has buckets to allow a burst of traffic. The maximum burst is determined by the bucket size `burst` (for guaranteed bandwidth) and `cburst` (for maximum bandwidth). The shaping profile has `burst-in-msec` and `cburst-in-msec` parameters for each shaping entry (`class id`) to control the bucket size.

This example uses the outbandwidth of the interface as 1 Mbps and the maximum bandwidth of class is 50%.

burst = burst-in-msec * guaranteed bandwidth = 100 ms × 1 Mbps x 50% = 50000 b = 6250 B

cburst = cburst-in-msec * maximum bandwidth = 200 ms × 1 Mbps x 50% = 100000 b = 12500 B

The following example sets `burst-in-msec` to 100 and `cburst-in-msec` to 200.

To set burst control in queuing mode:

```
config firewall shaping-profile
    edit "profile"
        set type queuing
        set default-class-id 31
        config shaping-entries
            edit 31
                set class-id 31
                set guaranteed-bandwidth-percentage 5
                set maximum-bandwidth-percentage 50
                set burst-in-msec 100 <range from 0 to 2000>
                set cburst-in-msec 200 <range from 0 to 2000>
            next
        end
    next
end
```

Multi-stage DSCP marking and class ID in traffic shapers

Traffic shapers have a multi-stage method so that packets are marked with a different differentiated services code point (DSCP) and `class id` at different traffic speeds. Marking packets with a different DSCP code is for the next hop to classify the packets. The FortiGate benefits by marking packets with a different `class id`. Combined with the egress interface shaping profile, the FortiGate can handle the traffic differently according to its `class id`.

Rule	DSCP code	Class ID
speed < guarantee bandwidth	`diffservcode`	`class id` in shaping policy
guarantee bandwidth < speed < exceed bandwidth	`exceed-dscp`	`exceed-class-id`
exceed bandwidth < speed	`maximum-dscp`	`exceed-class-id`

Netskope Exhibit 1017

This example sets the following parameters:

- When the current bandwidth is less than 50 Kbps, mark packets with `diffservcode` 100000 and set `class id` to 10.
- When the current bandwidth is between 50 Kbps and 100 Kbps, mark packets with `exceed-dscp` 111000 and set `exceed-class-id` to 20.
- When the current bandwidth is more than 100 Kbps, mark packets with `maximum-dscp` 111111 and set `exceed-class-id` to 20.

To set multi-stage DSCP marking and class ID in a traffic shaper:

```
config firewall shaper traffic-shaper
    edit "50k-100k-150k"
        set guaranteed-bandwidth 50
        set maximum-bandwidth 150
        set diffserv enable
        set dscp-marking-method multi-stage
        set exceed-bandwidth 100
        set exceed-dscp 111000
        set exceed-class-id 20
        set maximum-dscp 111111
        set diffservcode 100000
    next
end
```

```
config firewall shaping-policy
     edit 1
         set service "ALL"
         set dstintf PORT2
         set srcaddr "all"
         set dstaddr "all"
         set class-id 10
     next
end
```

Traffic shapers also have an `overhead` option that defines the per-packet size overhead used in rate computation.

To set the traffic shaper overhead option:

```
config firewall shaper traffic-shaper
    edit "testing"
        set guaranteed-bandwidth 50
        set maximum-bandwidth 150
        set overhead 14 <range from 0 to 100>
    next
end
```


Netskope Exhibit 1017

Examples

Enabling RED for FTP traffic from QA

This first example shows how to enable RED for FTP traffic from QA. This example sets a maximum of 10% of the packets to be dropped when queue usage reaches the maximum value.

To configure the firewall address:

```
config firewall address
    edit QA_team
        set subnet 10.1.100.0/24
    next
end
```

To set the shaping policy to classify traffic into different class IDs:

```
config firewall shaping-policy
    edit 1
        set service HTTPS HTTP
        set dstintf port1
        set srcaddr QA_team
        set dstaddr all
        set class-id 10
    next
    edit 2
        set service FTP
        set dstintf port1
        set srcaddr QA_team
        set dstaddr all
        set class-id 20
    next
end
```

To set the shaping policy to define the speed of each class ID:

```
config firewall shaping-profile
    edit QA_team_profile
       set type queuing
       set default-class-id 30
            config shaping-entries
                edit 1
                    set class-id 10
                    set guaranteed-bandwidth-percentage 50
                    set maximum-bandwidth-percentage 100
                next
                edit 2
                    set class-id 20
                    set guaranteed-bandwidth-percentage 30
                    set maximum-bandwidth-percentage 60
                    set red-probability 10
                next
                edit 3
                    set class-id 30
                    set guaranteed-bandwidth-percentage 20
```


Netskope Exhibit 1017

```
                    set maximum-bandwidth-percentage 50
                next
            end
    next
end
```

To apply the shaping policy to the interface:

```
config sys interface
    edit port1
        set outbandwidth 10000
        set egress-shaping-profile QA_team_profile
    next
end
```

To use diagnose commands to troubleshoot:

```
# diagnose netlink intf-class list port1
class htb 1:1 root rate 1250000Bps ceil 1250000Bps burst 1600B/8 mpu 0B overhead 0B cburst
1600B/8 mpu 0B overhead 0B level 7 buffer [00004e20] cbuffer [00004e20]
 Sent 11709 bytes 69 pkt (dropped 0, overlimits 0 requeues 0)
 rate 226Bps 2pps backlog 0B 0p
 lended: 3 borrowed: 0 giants: 0
 tokens: 18500 ctokens: 18500
class htb 1:10 parent 1:1 leaf 10: prio 1 quantum 62500 rate 625000Bps ceil 1250000Bps burst
1600B/8 mpu 0B overhead 0B cburst 1600B/8 mpu 0B overhead 0B level 0 buffer [00009c40]
cbuffer [00004e20]
 Sent 0 bytes 0 pkt (dropped 0, overlimits 0 requeues 0)
 rate 0Bps 0pps backlog 0B 0p
 lended: 0 borrowed: 0 giants: 0
 tokens: 40000 ctokens: 20000
class htb 1:20 parent 1:1 leaf 20: prio 1 quantum 37500 rate 375000Bps ceil 750000Bps burst
1599B/8 mpu 0B overhead 0B cburst 1599B/8 mpu 0B overhead 0B level 0 buffer [0001046a]
cbuffer [00008235]
 Sent 0 bytes 0 pkt (dropped 0, overlimits 0 requeues 0)
 rate 0Bps 0pps backlog 0B 0p
 lended: 0 borrowed: 0 giants: 0
 tokens: 66666 ctokens: 33333
class htb 1:30 parent 1:1 leaf 30: prio 1 quantum 25000 rate 250000Bps ceil 625000Bps burst
1600B/8 mpu 0B overhead 0B cburst 1600B/8 mpu 0B overhead 0B level 0 buffer [000186a0]
cbuffer [00009c40]
 Sent 11709 bytes 69 pkt (dropped 0, overlimits 0 requeues 0)
 rate 226Bps 2pps backlog 0B 0p
 lended: 66 borrowed: 3 giants: 0
 tokens: 92500 ctokens: 37000
class red 20:1 parent 20:0
```

```
 # diagnose netlink intf-qdisc list port1
qdisc htb 1: root refcnt 5 r2q 10 default 30 direct_packets_stat 0 ver 3.17
 Sent 18874 bytes 109 pkt (dropped 0, overlimits 5 requeues 0)
 backlog 0B 0p
qdisc pfifo 10: parent 1:10 refcnt 1 limit 1000p
 Sent 0 bytes 0 pkt (dropped 0, overlimits 0 requeues 0)
 backlog 0B 0p
qdisc red 20: parent 1:20 refcnt 1 limit 4000000B min 300000B max 1000000B ewma 9 Plog 23
Scell_log 20 flags 0
```


Netskope Exhibit 1017

```
 Sent 0 bytes 0 pkt (dropped 0, overlimits 0 requeues 0)
 backlog 0B 0p
  marked 0 early 0 pdrop 0 other 0
qdisc pfifo 30: parent 1:30 refcnt 1 limit 1000p
 Sent 18874 bytes 109 pkt (dropped 0, overlimits 0 requeues 0)
 backlog 0B 0p
```

Marking QA traffic with a different DSCP

This second example shows how to mark QA traffic with a different DSCP according to real-time traffic speed.

To configure the firewall address:

```
config firewall address
    edit QA_team
        set subnet 10.1.100.0/24
    next
end
```

To configure the firewall shaper traffic shaper:

```
config firewall shaper traffic-shaper
    edit "500k-1000k-1500k"
        set guaranteed-bandwidth 500
        set maximum-bandwidth 1500
        set diffserv enable
        set dscp-marking-method multi-stage
        set exceed-bandwidth 1000
        set exceed-dscp 111000
        set maximum-dscp 111111
        set diffservcode 100000
    next
end
```

```
config firewall shaping-policy
    edit QA_team
        set service "ALL"
        set dstintf port1
        set traffic-shaper "500k-1000k-1500k"
        set traffic-shaper-reverse "500k-1000k-1500k"
        set srcaddr "QA_team"
        set dstaddr "all"
    next
end
```


Netskope Exhibit 1017

Security Profiles

This section contains information about configuring FortiGate security features, including:

- Inspection modes on page 851
- Antivirus on page 856
- Web filter on page 894
- DNS filter on page 942
- Application control on page 963
- Intrusion prevention on page 974
- Email filter on page 979
- Data leak prevention on page 993
- VoIP solutions on page 1004
- ICAP on page 1015
- Web application firewall on page 1021
- SSL & SSH Inspection on page 1024
- Custom signatures on page 1035
- Overrides on page 1038



If you are unable to view a security profile feature, go to *System* > *Feature Visibility* to enable it.

Inspection modes

FortiOS supports flow-based and proxy-based inspection in firewall policies. You can select the inspection mode when configuring a policy.

Flow-based inspection takes a snapshot of content packets and uses pattern matching to identify security threats in the content.

Proxy-based inspection reconstructs content that passes through the FortiGate and inspects the content for security threats.

Certain security profiles allows users to display flow-based or proxy-based feature sets.

This following topics provide information about inspection modes for various security profile features:

- Flow mode inspection (default mode) on page 852
- Proxy mode inspection on page 852
- Inspection mode feature comparison on page 854

Netskope Exhibit 1017

Flow mode inspection (default mode)

When a firewall policy's inspection mode is set to *flow*, traffic flowing through the policy will not be buffered by the FortiGate. Unlike proxy mode, the content payload passing through the policy will be inspected on a packet by packet basis with the very last packet held by the FortiGate until the scan returns a verdict. If a violation is detected in the traffic, a reset packet is issued to the receiver, which terminates the connection, and prevents the payload from being sent successfully.

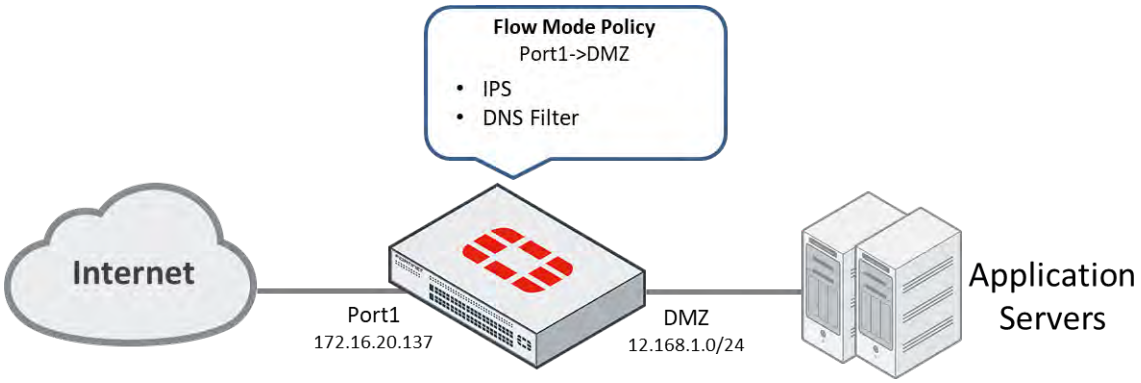
Flow-based inspection identifies and blocks security threats in real time as they are identified. All applicable flow-based security modules are applied simultaneously in one single pass, using Direct Filter Approach (DFA) pattern matching to identify possible attacks or threats. Pattern matching is offloaded and accelerated by CP8 or CP9 processors.

Flow-based inspection typically requires lower processing resources than proxy-based inspection and does not change packets, unless a threat is found and packets are blocked.

Use case

It is recommended that flow inspection is applied to policies that prioritize traffic throughput, such as allowing connections to be made towards a streaming or file server.

You have an application server which accepts connections from users for the daily quiz show app, HQ. Each HQ session sees 500,000+ participants, and speed is very important because participants have less than 10 seconds to answer the quiz show questions.



In this scenario, a flow inspection policy is recommended to prioritize throughput. The success of the application depends on providing reliable service for large numbers of concurrent users. We will apply an IPS sensor to this policy to protect the server from external DOS attacks.

Proxy mode inspection

When a firewall policy's inspection mode is set to *proxy*, traffic flowing through the policy will be buffered by the FortiGate for inspection. This means that the packets for a file, email message, or web page will be held by the FortiGate until the entire payload is inspected for violations (virus, spam, or malicious web links). After FortiOS has finished the inspection, the payload is either released to the destination (if traffic is clean) or dropped and replaced with a replacement message (if traffic contains violations).

To optimize inspection, the policy can be configured to block or ignore files or messages that exceed a certain size. To prevent the receiving end user from timing out, client comforting can be applied, which allows small portions of the payload to be sent while it is undergoing inspection.

Proxy mode provides the most thorough inspection of the traffic; however, its thoroughness sacrifices performance, making its throughput slower than that of a flow-mode policy. Under normal traffic circumstances, the throughput difference between a proxy-based and flow-based policy is not significant.

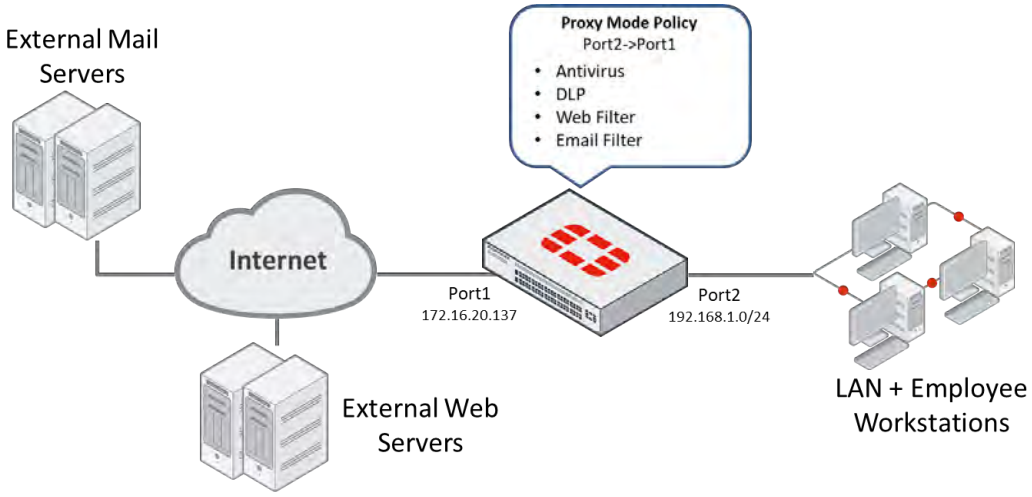
Use case

We want traffic inspection to be as thorough as possible to avoid any data leaks from exiting the LAN and any malicious content from entering it.

The following scenarios demonstrate common use cases for proxy inspection.

Scenario 1

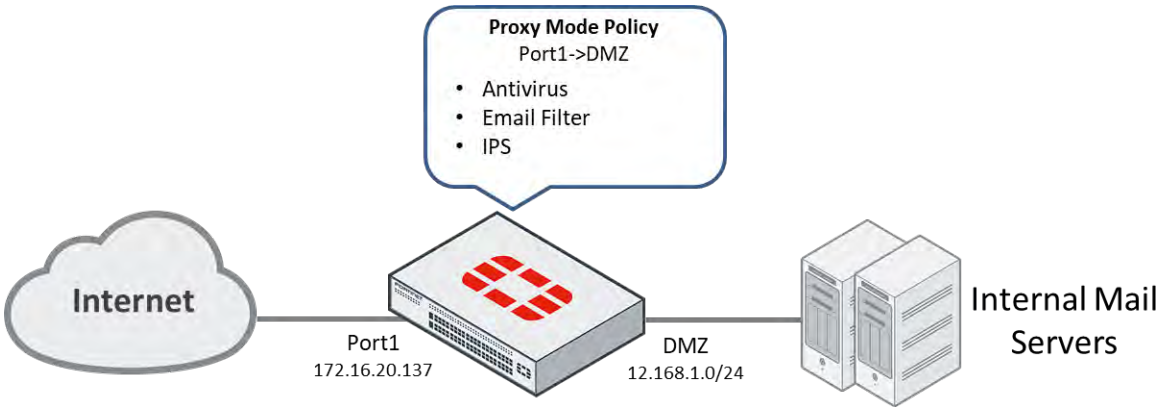
Your organization deals with sensitive data on a regular basis and a data leak would significantly harm your business. At the same time, you wish to protect your employees from malicious content, such as viruses and phishing emails, which could be used to gain access to your network and the sensitive data on your systems.



In this scenario, a proxy inspection policy is recommended to prioritize network security. We want traffic inspection to be as thorough as possible to avoid any data leaks from exiting the LAN and any malicious content from entering it. On this policy, we will apply the virus filter, DLP filter, Web Filter, and email filter all operating in proxy mode.

Scenario 2

You have a corporate mail server in your domain, which is used by your employees for everyday business activities. You want to protect your employees from phishing emails and viruses. At the same time, you want to also protect your web servers from external attacks.



In this scenario, a proxy inspection policy is recommended to prioritize the safety of employee emails. Applying the antivirus and email filter in this mode allows us to most reliably filter out any malware and spam emails received by the mail servers via SMTP or MAPI. The IPS sensor can be used to prevent DOS attacks on the mail servers.

Inspection mode feature comparison

The following table shows which UTM profile can be configured on a flow mode or proxy mode inspection policy.

Some UTM profiles are hidden in the GUI and can only be configured using the CLI. To configure profiles in a firewall policy in CLI, enable the `utm-status` setting.

Some profiles might have feature differences between flow-based and proxy-based Inspection. From the GUI and CLI, you can set the *Feature set* option to be *Flow-based* or *Proxy-based* to display only the settings for that mode.

	Flow Mode Inspection Policy		Proxy Mode Inspection Policy		Feature set option
UTM Profile	GUI	CLI	GUI	CLI	
AntiVirus	Yes	Yes	Yes	Yes	GUI/CLI
Web Filter	Yes	Yes	Yes	Yes	GUI/CLI
DNS Filter	Yes	Yes	Yes	Yes	N/A
Application Control	Yes	Yes	Yes	Yes	N/A
Intrusion Prevention System	Yes	Yes	Yes	Yes	N/A
Email Filter	Yes	Yes	Yes	Yes	GUI/CLI
Data Leak Prevention	No	Yes	No	Yes	CLI
VoIP	Yes	Yes	Yes	Yes	N/A
ICAP	No	No	Yes	Yes	N/A

Netskope Exhibit 1017

UTM Profile	Flow Mode Inspection Policy		Proxy Mode Inspection Policy		Feature set option
	GUI	CLI	GUI	CLI	
Web Application Firewall	No	No	Yes	Yes	N/A
SSL/SSH Inspection	Yes	Yes	Yes	Yes	CLI

The following sections outline differences between flow-based and proxy-based inspection for a security profile.

Feature comparison between AntiVirus inspection modes

The following table indicates which AntiVirus features are supported by their designated scan modes.

Part1	Replacement Message	Content Disarm	Mobile Malware	Virus Outbreak	Sandbox Inspection	NAC Quarantine
Proxy	Yes	Yes	Yes	Yes	Yes	Yes
Flow (hybrid scan)	Yes*	No	Yes	Limited	Yes	Yes

*IPS Engine caches the URL and a replacement message is presented after the second attempt.

Part 2	Archive Blocking	Emulator	Client Comforting	Infection Quarantine	Heuristics	Treat EXE as Virus
Proxy	Yes	Yes	Yes	Yes (1)	Yes	Yes (2)
Flow (hybrid scan)	Yes	Yes	No	Limited	Yes	Yes (2)

1. Only available on FortiGate models with HDD or when FortiAnalyzer or FortiGate Cloud is connected and enabled.
2. Only applies to inspection on IMAP, POP3, SMTP, and MAPI protocols.

Feature comparison between Web Filter inspection modes

The following table indicates which Web Filter features are supported by their designated inspection modes.

	FortiGuard Category-Based Filter	Category Usage Quota	Override Blocked Categories	File Filter	Search Engines	Static URL Filter	Rating Option	Proxy Option	Web Profile Override
Proxy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Flow	Yes (1)	No	Yes (2)	No	No	Yes	Yes	Limited (3)	No

1. Local Category and Remote Category filters do not support the warning and authenticate actions.
2. Local Category and Remote Category filters cannot be overridden.

Netskope Exhibit 1017

3. Only HTTP POST Action is supported.

Feature comparison between Email Filter inspection modes

The following tables indicate which Email Filters are supported by the specified inspection modes for local filtering and FortiGuard-assisted filtering.

Local Filtering	Banned Word Check	Block/Allow List	HELO/ EHLO DNS Check	Return Address DNS Check	DNSBL/ ORBL Check	MIME Header Check	File Filter
Proxy	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Flow	Yes	Yes	No	No	No	Yes	No

FortiGuard-Assisted Filtering	Phishing URL Check	Anti-Spam Block List Check	Submit Spam to FortiGuard	Spam Email Checksum Check	Spam URL Check
Proxy	Yes	Yes	Yes	Yes	Yes
Flow	No	No	No	No	No

Feature comparison between DLP inspection modes

The following table indicates which DLP filters are supported by their designated inspection modes.

	Credit Card Filter	SSN Filter	Regex Filter	File-Type Filter	File-Pattern Filter	Fingerprint Filter	Watermark Filter	Encrypted Filter	File-Size Filter
Proxy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Flow	Yes	Yes	Yes	Yes	Yes	No	No	Yes	Yes*

*File-size filtering only works if file size is present in the protocol exchange.

Antivirus

FortiOS offers the unique ability to implement both flow-based and proxy-based antivirus concurrently, depending on the traffic type, users, and locations. Flow-based antivirus offers higher throughput performance.

FortiOS includes two preloaded antivirus profiles:

- *default*
- *wifi-default*

You can customize these profiles, or you can create your own to inspect certain protocols, remove viruses, analyze suspicious files with FortiSandbox, and apply botnet protection to network traffic. Once configured, you can add the antivirus profile to a firewall policy.



This functionality requires a subscription to FortiGuard Antivirus.

Protocol comparison between antivirus inspection modes

The following table indicates which protocols can be inspected by the designated antivirus scan modes.

	HTTP	FTP	IMAP	POP3	SMTP	NNTP	MAPI	CIFS	SSH
Proxy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes*	Yes
Flow	Yes	Yes	Yes	Yes	Yes	Yes	No	Yes	No

* Proxy mode antivirus inspection on CIFS protocol has the following limitations:

- Cannot detect infections within some archive files.
- Cannot detect oversized files.

Other antivirus differences between inspection modes

Starting from 6.4.0, the scan mode option is no longer available for flow-based AV.

This means that AV no longer exclusively uses the default or legacy scan modes when handling traffic on flow-based firewall policies. Instead, AV in flow-based policies uses a hybrid of the two scan modes. Flow AV may use a pre-filtering database for malware detection in some circumstances as opposed to the full AV signature database in others. The scan method is determined by the IPS engine algorithm that is based on the type of file being scanned.

In contrast, proxy mode maintains the scan mode option, which can be toggled between default or legacy mode. In default mode, the WAD daemon uses a stream-based approach, while legacy mode disables this stream-based approach. Proxy default scan-mode uses pre-scanning and stream-based scanning for HTTP traffic.

Stream-based scanning provides the following AV improvements:

- Archive files (ZIP, GZIP, BZIP2, TAR, ISO) that exceed the oversize limit are uncompressed and scanned for infections.
- The contents of large archive files are scanned without having to buffer the entire file.
- Small files are scanned locally by the WAD daemon if only AV scanning is needed in the policy.
- File filtering on HTTP/HTTPS is handled locally by the WAD daemon.

This means that the overall memory usage is optimized when an archive file is scanned, and better security is achieved by scanning archives that would otherwise be bypassed.

However, stream-based scanning has limitations on the more complex features that it can scan. For the following features, traffic will be automatically handed off to the scanunit daemon for scanning (as in the case of legacy mode):

- Heuristic AV scan
- DLP
- Quarantine
- FortiGuard outbreak prevention and external block list
- Content disarm

Netskope Exhibit 1017

To configure the scan mode:

```
config antivirus profile
    edit <name>
        set feature-set proxy
        ...
        set scan-mode {default | legacy}
    next
end
```

The following topics provide information about antivirus profiles:

- Content disarm and reconstruction for antivirus on page 859
- FortiGuard outbreak prevention on page 862
- External malware block list for antivirus on page 864
- Checking flow antivirus statistics on page 867
- CIFS support on page 870
- Databases on page 858

The following topics provide information about sandbox inspection with antivirus:

- Using FortiSandbox appliance with antivirus on page 876
- Using FortiSandbox Cloud with antivirus on page 886

Databases

The AntiVirus scanning engine uses a virus signatures database to record the unique attributes of each infection. The AntiVirus scan searches for these signatures and when one is discovered, the FortiGate unit determines if the file is infected and takes action.

All FortiGate units have the normal AntiVirus signature database. Some models have additional databases that you can use. The database you use depends on your network and security needs, and on your FortiGate model.

The extended virus definitions database is the default setting and provides comprehensive antivirus protection. Low end FortiGate models cannot support the extreme database and the `use-extreme-db` setting is only available on models that support the extreme database.

Extended	This is the default setting. This database includes currently spreading viruses, as determined by the FortiGuard Global Security Research Team, plus recent viruses that are no longer active. These viruses may have been spreading within the last year but have since nearly or completely disappeared.
Extreme	Includes the extended database, plus a large collection of zoo viruses. These are viruses that have not spread in a long time and are largely dormant. Some zoo viruses might rely on operating systems and hardware that are no longer widely used.

To change the antivirus database:

```
config antivirus settings
    set use-extreme-db {enable | disable}
end
```

Content disarm and reconstruction for antivirus

Content disarm and reconstruction (CDR) allows the FortiGate to sanitize Microsoft documents and PDF (disarm) by removing active content such as hyperlinks, embedded media, javascript, macros, and so on from the Office document files without affecting the integrity of its textual content (reconstruction).

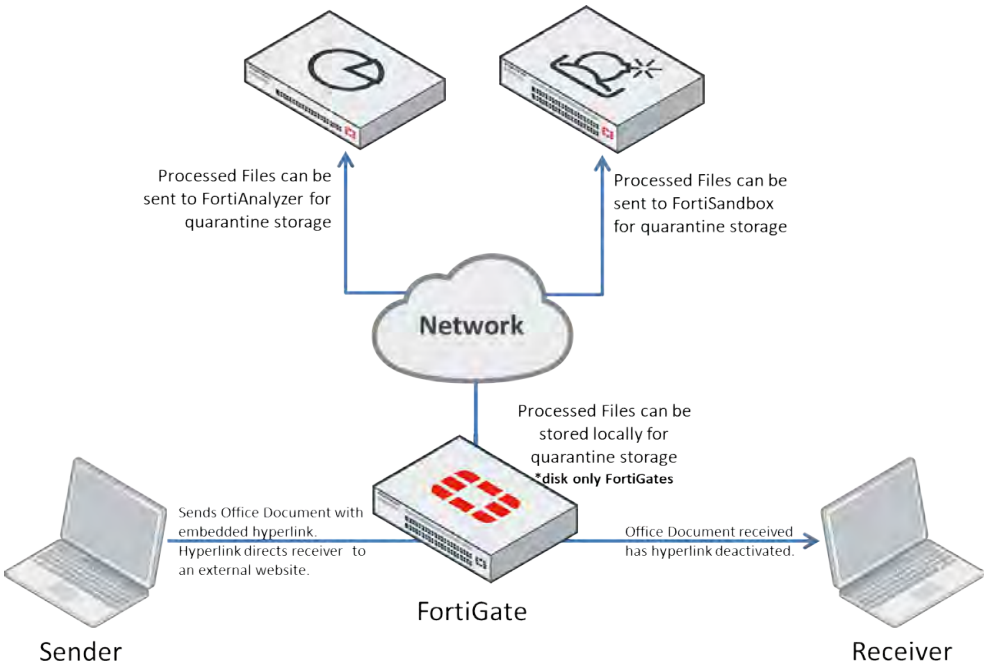
This feature allows network admins to protect their users from malicious Office document files.

Files processed by CDR can have the original copy quarantined on the FortiGate, allowing admins to observe them. These original copies can also be obtained in the event of a false positive.

Support and limitations

- CDR can only be performed on Microsoft Office document and PDF files.
- Local disk CDR quarantine is only possible on FortiGate models that contain a hard disk.
- CDR is only supported on HTTP, SMTP, POP3, IMAP.
  - SMTP splice and client-comfort mode is not supported.
- CDR does not work on flow-based inspection modes.
- CDR can only work on files in .ZIP type archives.

Network topology example

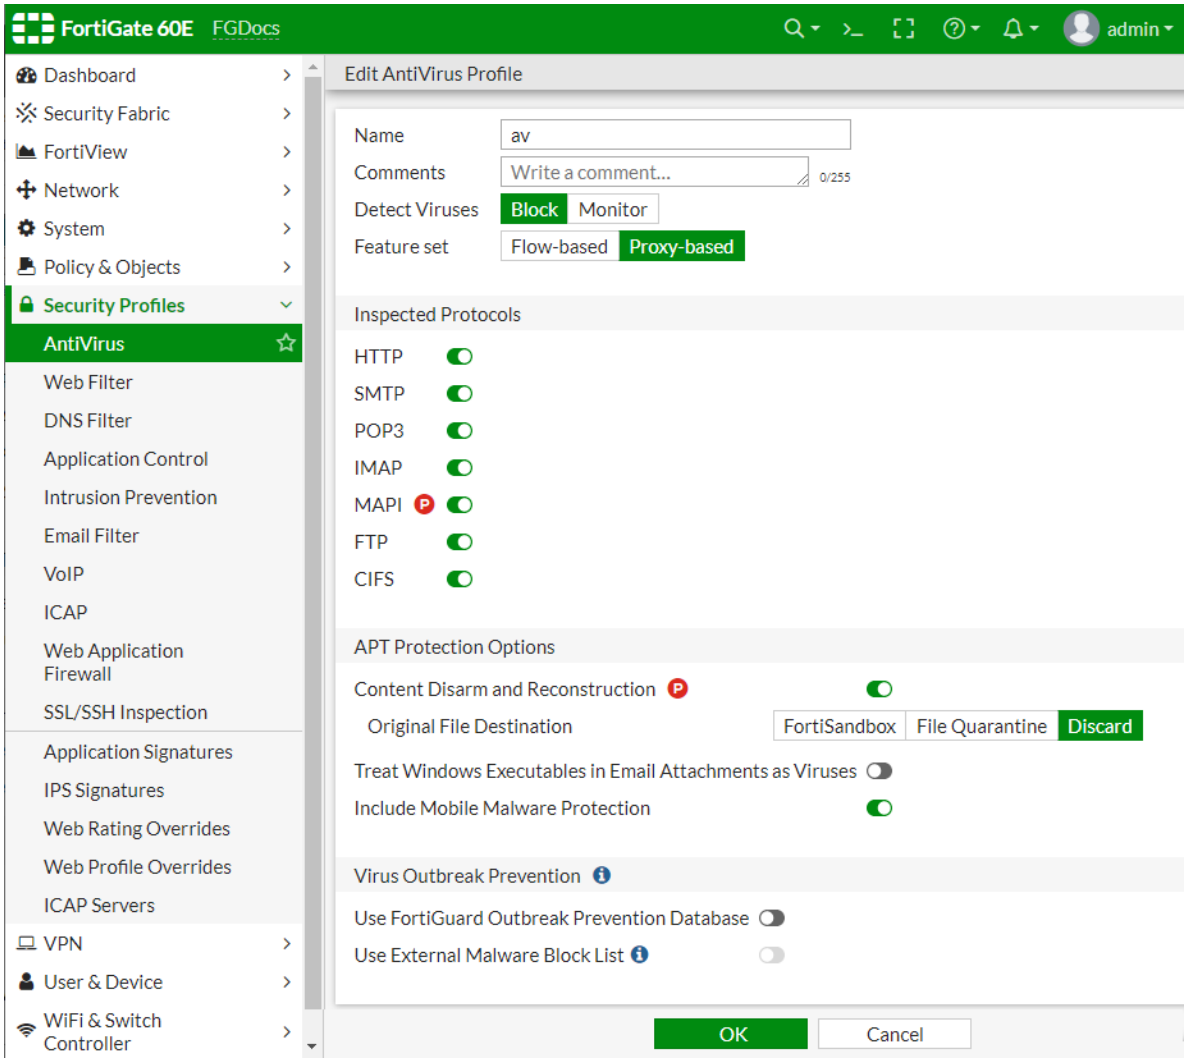


Configuring the feature

In order to configure antivirus to work with CDR, you must enable CDR on your antivirus profile, set the quarantine location, and then fine-tune the CDR detection parameters.

To enable CDR on your antivirus profile:

1. Go to *Security Profiles* > *AntiVirus*.
2. Edit an antivirus profile, or create a new one.
3. Enable the toggle for *Content Disarm and Reconstruction* under *APT Protection Options*.



To set a quarantine location:

1. Go to *Security Profiles* > *AntiVirus*.
2. Edit an antivirus profile, or create a new one.

3. Select a quarantine location from the available options: *Discard*, *File Quarantine*, or *FortiSandbox*.

Discard	The default setting, which discards the original document file.
File Quarantine	Saves the original document file to disk (if possible) or a connected FortiAnalyzer based on the FortiGate's log settings, visible through *Config Global > Config Log FortiAnalyzer Setting*.
FortiSandbox	Saves the original document file to a connected FortiSandbox.

4. Click *Apply*.

To fine-tune CDR detection parameters in the CLI:

- Select which active content to detect/process:
  By default, all active office and PDF content types are enabled. To fine-tune CDR to ignore certain content, you must disable that particular content parameter. The example below configures the CDR to ignore Microsoft Office macros.

```
config antivirus profile
    edit av
        config content-disarm
            set ?
                original-file-destination    Destination to send original file if active
content is removed.
                office-macro             Enable/disable stripping of macros in Microsoft
Office documents.
                office-hylink            Enable/disable stripping of hyperlinks in
Microsoft Office documents.
                office-linked            Enable/disable stripping of linked objects in
Microsoft Office documents.
                office-embed             Enable/disable stripping of embedded objects in
Microsoft Office documents.
                office-dde               Enable/disable stripping of Dynamic Data
Exchange events in Microsoft Office documents.
                office-action            Enable/disable stripping of PowerPoint action
events in Microsoft Office documents.
                pdf-javacode             Enable/disable stripping of JavaScript code in
PDF documents.
                pdf-embedfile            Enable/disable stripping of embedded files in
PDF documents.
                pdf-hyperlink            Enable/disable stripping of hyperlinks from PDF
documents.
                pdf-act-gotor            Enable/disable stripping of PDF document
actions that access other PDF documents.
                pdf-act-launch           Enable/disable stripping of PDF document
actions that launch other applications.
                pdf-act-sound            Enable/disable stripping of PDF document
actions that play a sound.
                pdf-act-movie            Enable/disable stripping of PDF document
actions that play a movie.
                pdf-act-java             Enable/disable stripping of PDF document
actions that execute JavaScript code.
                pdf-act-form             Enable/disable stripping of PDF document
actions that submit data to other targets.
                cover-page               Enable/disable inserting a cover page into the
```


Netskope Exhibit 1017

```
disarmed document.
                detect-only             Enable/disable only detect disarmable files, do
not alter content.
            set office-macro disable
        end
    next
end
```

- Detect but do not modify active content:
  By default, CDR will disarm any detected documents containing active content. To prevent CDR from disarming documents, you can set it to operate in detect-only mode. To do this, the option *detect-only* must be enabled.

```
config antivirus profile
    edit av
        config content-disarm
            set detect-only enable
        end
    next
end
```

- Enable/disable the CDR cover page:
  By default, a cover page will be attached to the file's content when the file has been processed by CDR. To disable the cover page, the `cover-page` parameter needs to be disabled.

```
config antivirus profile
    edit av
        config content-disarm
            set cover-page disable
        end
    next
end
```

FortiGuard outbreak prevention

FortiGuard Virus Outbreak Protection Service (VOS) allows the FortiGate antivirus database to be subsidized with third-party malware hash signatures curated by FortiGuard. The hash signatures are obtained from FortiGuard's Global Threat Intelligence database. The antivirus database queries FortiGuard with the hash of a scanned file. If FortiGuard returns a match, the scanned file is deemed to be malicious.

FortiGuard VOS can be used in both proxy-based and flow-based policy inspections across all supported protocols.

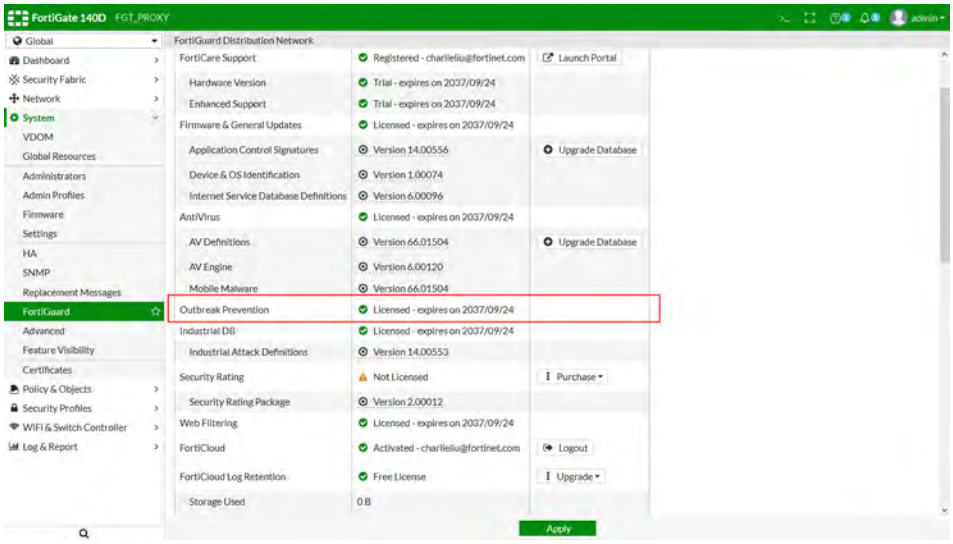


The FortiGate must be registered with a valid FortiGuard outbreak prevention license.

Netskope Exhibit 1017

To verify FortiGuard antivirus license information:

1. Go to *System* > *FortiGuard* and locate the *Outbreak Prevention* section in the table.



2. See the instructions in the video, How to Purchase or Renew FortiGuard Services, if required.

To enable FortiGuard outbreak prevention:

1. Go to *Security Profiles* > *AntiVirus*.
2. Edit an antivirus profile, or create a new one.
3. Under *Virus Outbreak Protection*, enable *Use FortiGuard Outbreak Prevention Database*.
4. Click *OK*.

To verify FortiGuard antivirus license information:

```
# diagnose debug rating
Locale       : english

Service      : Web-filter
Status       : Enable
License      : Contract

Service      : Antispam
Status       : Disable

Service      : Virus Outbreak Prevention
Status       : Enable
License      : Contract

-=- Server List (Tue Feb 19 16:36:15 2019) -=-

IP                   Weight    RTT Flags  TZ    Packets  Curr Lost Total Lost
   Updated Time
192.168.100.185        -218      2 DI     -8        113          0          0 Tue Feb
19 16:35:55 2019
```

To enable all scanunit debug categories:

```
# diagnose sys scanunit debug all
Set meta-category: all(0xffffffff)
Enabled categories(0xffffffff): daemon job quarantine analytics outbreak-prevention dlp
antispam file-filter
```

```
# diagnose debug enable
# su 4739 open
su 4739 req vfid 1 id 1 ep 0 new request, size 313, policy id 1, policy type 0
su 4739 req vfid 1 id 1 ep 0 received; ack 1, data type: 0
su 4739 job 1 request info:
su 4739 job 1   client 10.1.100.11:39412 server 172.16.200.44:80
su 4739 job 1   object_name 'zhvo_test.com'
su 4739 file-typing NOT WANTED options 0x0 file_filter no
su 4739 enable databases 0b (core mmdb extended)
su 4739 job 1 begin http scan
su 4739 scan file 'zhvo_test.com' bytes 68
su 4739 job 1 outbreak-prevention scan, level 0, filename 'zhvo_test.com'
su 4739 scan result 0
su 4739 job 1 end http scan
su 4739 job 1 inc pending tasks (1)
su 4739 not wanted for analytics: analytics submission is disabled (m 0 r 0)
su 4739 job 1 suspend
su 4739 outbreak-prevention recv error
su 4739 ftgd avquery id 0 status 1
su 4739 job 1 outbreak-prevention infected entryid=0
su 4739 report AVQUERY infection priority 1
su 4739 insert infection AVQUERY SUCCEEDED loc (nil) off 0 sz 0 at index 0 total infections
1 error 0
su 4739 job 1 dec pending tasks 0
su 4739 job 1 send result
su 4739 job 1 close
su 4739 outbreak-prevention recv error
```

External malware block list for antivirus

The external malware block list is a new feature introduced in FortiOS 6.2.0, which falls under the umbrella of outbreak prevention.

This feature provides another means of supporting the AV Database by allowing users to add their own malware signatures in the form of MD5, SHA1, and SHA256 hashes.

This feature provides a mechanism for antivirus to retrieve an external malware hash list from a remote server and polls the hash list every *n* minutes for updates.

Support and limitations

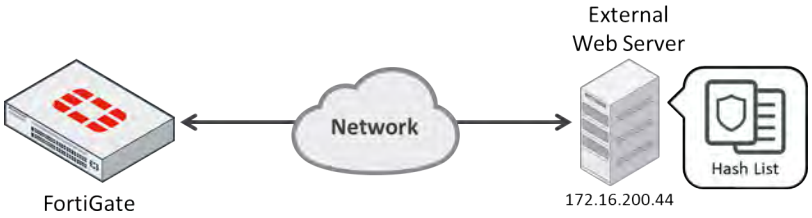
Malware detection using the external malware block list can be used in both proxy-based and flow-based policy inspections.

Just like FortiGuard outbreak prevention, external dynamic block list is not supported in AV quick scan mode.

Using different types of hashes simultaneously may slow down the performance of malware scanning. For this reason, users are recommended to only using one type of hash (either MD5, SHA1, or SHA256), not all three simultaneously.

Netskope Exhibit 1017

Network topology example



Configuring the feature

To configure antivirus to work with external block list:

1. Create the malware hash list
   The malware hash list follows a strict format in order for its contents to be valid. Malware hash signature entries must be separated into each line. A valid signature needs to follow the format below:

```
# MD5 Entry with hash description
aa67243f746e5d76f68ec809355ec234  md5_sample1

# SHA1 Entry with hash description
a57983cb39e25ab80d7d3dc05695dd0ee0e49766  sha1_sample2

# SHA256 Entry with hash description
ae9bc0b4c5639d977d720e4271da06b50f7c60d1e2070e9c75cc59ab30e49379  sha256_sample1

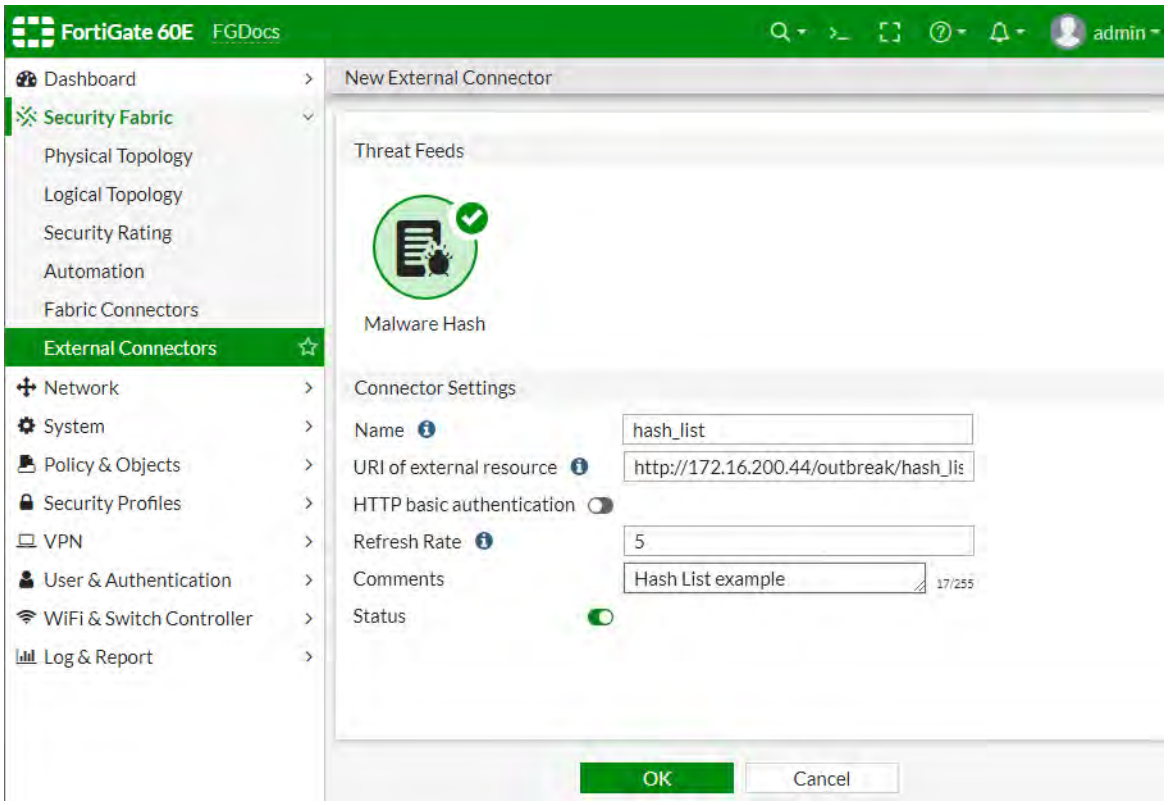
# Entry without hash description
0289b0d967cb7b1fb1451339c7b9818a621903090e0020366ab415c549212521

# Invalid entries
7688499dc71b932feb126347289c0b8a_md5_sample2
7614e98badca10b5e2d08f8664c519b7a906fbd5180ea5d04a82fce9796a4b87sha256_sample3
```

2. Configure the external malware block list source:
   a. Go to *Global* > *Security Fabric* > *External Connectors* and click *Create New*.
   b. Select *Malware Hash*.

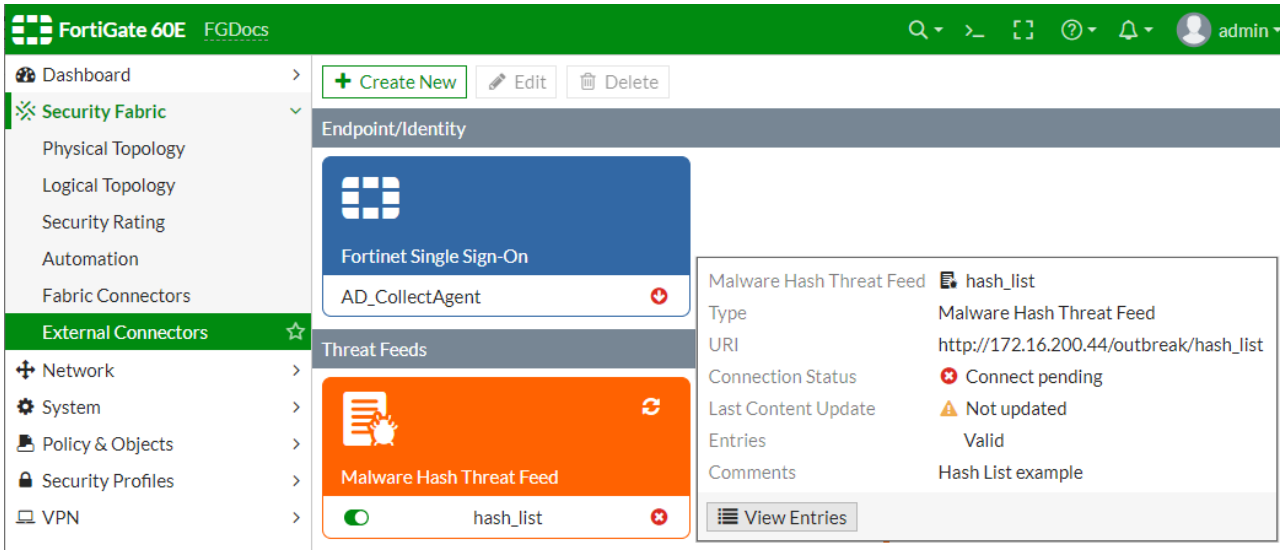
Netskope Exhibit 1017

c. Fill out the fields as shown. The URI must point to the malware hash list on the remote server.

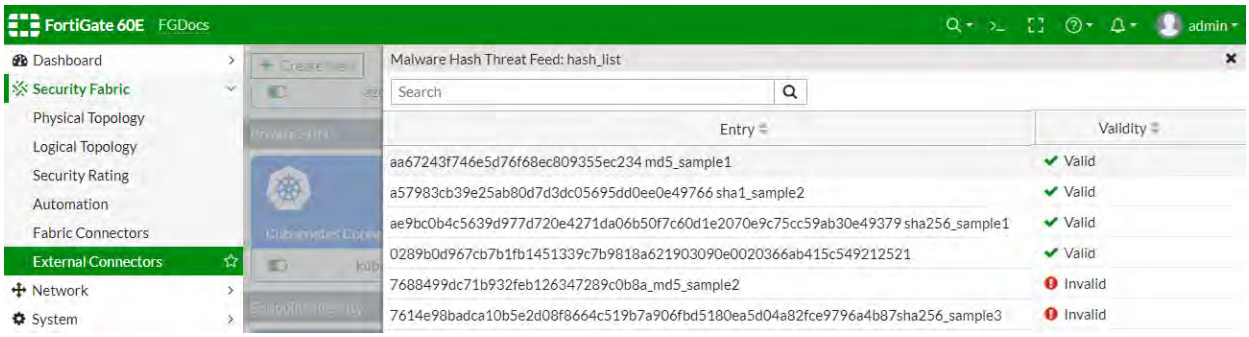


d. Click *OK*.
The malware hash source object is now created.

3. To view entries inside the malware block list, hover the pointer over the malware hash card and click the *View Entries* button.



The malware hash threat feed displays (*hash_list*):



4. Enable the external malware block list in the antivirus profile:

```
config antivirus profile
    edit <av_profile>
        config outbreak-prevention
            set external-blocklist enable
        end
    next
end
```

Antivirus is now ready to use the external malware block list.

Diagnostics and debugging

Check if the `scanunit` daemon has updated itself with the external hashes:

```
FGT_PROXY # config global
FGT_PROXY (global) # diagnose sys scanunit malware-list list
md5 'aa67243f746e5d76f68ec809355ec234' profile 'hash_list' description 'md5_sample1'
sha1 'a57983cb39e25ab80d7d3dc05695dd0ee0e49766' profile 'hash_list' description 'sha1_
sample2'
sha256 '0289b0d967cb7b1fb1451339c7b9818a621903090e0020366ab415c549212521' profile 'hash_
list' description ''
sha256 'ae9bc0b4c5639d977d720e4271da06b50f7c60d1e2070e9c75cc59ab30e49379' profile 'hash_
list' description 'sha256_sample1'
```

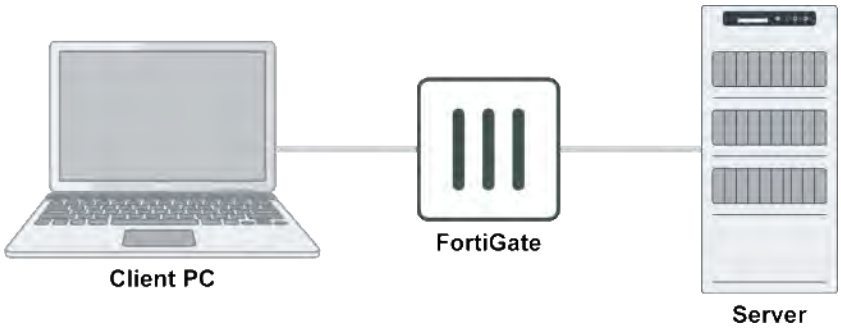
Checking flow antivirus statistics

This feature provides a flow antivirus statistics check, and an API for SNMP to get AV statistics.

Two CLI commands are used to show and clear the antivirus statistics:

```
diagnose ips av stats show
diagnose ips av stats clear
```

This example uses the following topology:



To check flow antivirus statistics:

1. Create an antivirus profile:

```
config antivirus profile
    edit "av-test"
        config http
            set options scan avmonitor
        end
        config ftp
            set options scan quarantine
        end
    next
end
```

2. Enable the profile on a firewall policy:

```
config firewall policy
    edit 1
        set name "policy1"
        set srcintf "port2"
        set dstintf "port1"
        set srcaddr "all"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ALL"
        set utm-status enable
        set fsso disable
        set av-profile "av-test"
        set ssl-ssh-profile "custom-deep-inspection"
        set nat enable
    next
end
```

3. On the client PC, download the EICAR Standard Anti-Virus Test File via HTTP.
4. Check the antivirus statistics on the FortiGate. As the action is set to monitor for HTTP, `HTTP virus detected` is increased by 1:

```
# diagnose ips av stats show
AV stats:
HTTP virus detected: 1
HTTP virus blocked: 0
SMTP virus detected: 0
SMTP virus blocked: 0
```


Netskope Exhibit 1017

```
POP3 virus detected: 0
POP3 virus blocked: 0
IMAP virus detected: 0
IMAP virus blocked: 0
NNTP virus detected: 0
NNTP virus blocked: 0
FTP virus detected: 0
FTP virus blocked: 0
SMB virus detected: 0
SMB virus blocked: 0
```

5. On the client PC, download the EICAR file via FTP.
6. Check the antivirus statistics on the FortiGate. As the action is set to quarantine for FTP, `FTP virus detected` and `FTP virus blocked` are both increased by 1:

```
# diagnose ips av stats show
AV stats:
HTTP virus detected: 1
HTTP virus blocked: 0
SMTP virus detected: 0
SMTP virus blocked: 0
POP3 virus detected: 0
POP3 virus blocked: 0
IMAP virus detected: 0
IMAP virus blocked: 0
NNTP virus detected: 0
NNTP virus blocked: 0
FTP virus detected: 1
FTP virus blocked: 1
SMB virus detected: 0
SMB virus blocked: 0
```

7. Check the antivirus statistics using snmpwalk:

```
root:~# snmpwalk -c public -v 1 10.1.100.6 1.3.6.1.4.1.12356.101.8.2.1.1
iso.3.6.1.4.1.12356.101.8.2.1.1.1.1 = Counter32: 2  (fgAvVirusDetected)
iso.3.6.1.4.1.12356.101.8.2.1.1.2.1 = Counter32: 1  (fgAvVirusBlocked)
iso.3.6.1.4.1.12356.101.8.2.1.1.3.1 = Counter32: 1  (fgAvHTTPVirusDetected)
iso.3.6.1.4.1.12356.101.8.2.1.1.4.1 = Counter32: 0
iso.3.6.1.4.1.12356.101.8.2.1.1.5.1 = Counter32: 0
iso.3.6.1.4.1.12356.101.8.2.1.1.6.1 = Counter32: 0
iso.3.6.1.4.1.12356.101.8.2.1.1.7.1 = Counter32: 0
iso.3.6.1.4.1.12356.101.8.2.1.1.8.1 = Counter32: 0
iso.3.6.1.4.1.12356.101.8.2.1.1.9.1 = Counter32: 0
iso.3.6.1.4.1.12356.101.8.2.1.1.10.1 = Counter32: 0
iso.3.6.1.4.1.12356.101.8.2.1.1.11.1 = Counter32: 1  (fgAvFTPVirusDetected)
iso.3.6.1.4.1.12356.101.8.2.1.1.12.1 = Counter32: 1  (fgAvFTPVirusBlocked)
iso.3.6.1.4.1.12356.101.8.2.1.1.13.1 = Counter32: 0
iso.3.6.1.4.1.12356.101.8.2.1.1.14.1 = Counter32: 0
iso.3.6.1.4.1.12356.101.8.2.1.1.15.1 = Counter32: 0
iso.3.6.1.4.1.12356.101.8.2.1.1.16.1 = Counter32: 0
iso.3.6.1.4.1.12356.101.8.2.1.1.17.1 = Counter32: 0
iso.3.6.1.4.1.12356.101.8.2.1.1.18.1 = Counter32: 0
iso.3.6.1.4.1.12356.101.8.2.1.1.19.1 = Counter32: 0
iso.3.6.1.4.1.12356.101.8.2.1.1.20.1 = Counter32: 0
iso.3.6.1.4.1.12356.101.8.2.1.1.21.1 = Counter32: 0
iso.3.6.1.4.1.12356.101.8.2.1.1.22.1 = Counter32: 0
```


Netskope Exhibit 1017

8. Optionally, reset the antivirus statistics to zero:

```
diagnose ips av stats clear
```

CIFS support

File filtering and antivirus scanning for proxy-based inspection on Common Internet File System (CIFS) traffic is supported.

File filtering for CIFS is performed by inspecting the first 4 KB of the file to identify the file's magic number. If a match occurs, CIFS file filtering prevents the CIFS command that contains that file from running.

The CIFS security profile handles the configuration of file filtering on CIFS. The antivirus profile handles the antivirus configuration for CIFS scanning.

For a CIFS profile to be available for assignment in a policy, the policy must use proxy inspection mode. See Proxy mode inspection on page 852 for details.

The following are not supported by CIFS scanning in proxy inspection mode:

- File types and infections within archive files cannot be detected.
- Oversized files cannot be detected.
- Special condition archive files (encrypted, corrupted, mailbomb, and so on) marked by the antivirus engine are blocked automatically.

Supported file types

File filter supports the following file types:

File Type Name	Description
7z	Match 7-zip files
arj	Match arj compressed files
cab	Match Windows cab files
lzh	Match lzh compressed files
rar	Match rar archives
tar	Match tar files
zip	Match zip files
bzip	Match bzip files
gzip	Match gzip files
bzip2	Match bzip2 files
xz	Match xz files
bat	Match Windows batch files
msc	Match msc files

File Type Name	Description
uue	Match uue files
mime	Match mime files
base64	Match base64 files
binhex	Match binhex files
bin	Match bin files
elf	Match elf files
exe	Match Windows executable files
hta	Match hta files
html	Match html files
jad	Match jad files
class	Match class files
cod	Match cod files
javascript	Match javascript files
msoffice	Match MS-Office files. For example, doc, xls, ppt, and so on.
msofficex	Match MS-Office XML files. For example, docx, xlsx, pptx, and so on.
fsg	Match fsg files
upx	Match upx files
petite	Match petite files
aspack	Match aspack files
prc	Match prc files
sis	Match sis files
hlp	Match Windows help files
activemime	Match activemime files
jpeg	Match jpeg files
gif	Match gif files
tiff	Match tiff files
png	Match png files
bmp	Match bmp files
unknown	Match unknown files
mpeg	Match mpeg files

File Type Name	Description
mov	Match mov files
mp3	Match mp3 files
wma	Match wma files
wav	Match wav files
pdf	Match pdf files
avi	Match avi files
rm	Match rm files
torrent	Match torrent files
msi	Match Windows Installer msi bzip files
mach-o	Match Mach object files
dmg	Match Apple disk image files
.net	Match .NET files
xar	Match xar archive files
chm	Match Windows compiled HTML help files
iso	Match ISO archive files
crx	Match Chrome extension files

Configure file-type filtering and antivirus scanning on CIFS traffic

To configure file-type filtering and antivirus scanning on CIFS traffic:

1. Configure a CIFS domain controller on page 872
2. Configure a CIFS profile on page 873
3. Configure an antivirus profile on page 874

Configure a CIFS domain controller

The domain controller must be configured when CIFS traffic is encrypted. The configuration tells the FortiGate the network location of the domain controller and the superuser credentials.

To configure the CIFS domain controller:

```
config credential-store domain-controller
    edit "SERVER_NAME"
        set domain-name "EXAMPLE.COM"
        set username "admin-super"
        set password **********
        set ip 172.16.201.40
    next
end
```

Configure a CIFS profile

To create a CIFS profile, configure the server credential type and add file filter entries.

Set the CIFS server credential type

The CIFS server credential type can be `none`, `credential-replication`, or `credential-keytab`.

none

The CIFS profile assumes the CIFS traffic is unencrypted. This is the default value.

```
config cifs profile
    edit "cifs"
        set server-credential-type none
    next
end
```

credential-replication

To decrypt CIFS traffic, FortiOS obtains the session key from the domain controller by logging in to the superuser account. The domain controller must be configured.

```
config cifs profile
    edit "cifs"
        set server-credential-type credential-replication
        set credential-store domain-controller "SERVER_NAME"
    next
end
```

Variable	Description
domain-controller <string>	The previously configured domain to decrypt CIFS traffic for.

credential-keytab

To decrypt CIFS traffic, FortiOS uses a series of keytab values. This method is used when the SMB connection is authenticated by Kerberos. Keytab entries must be configured, and are stored in FortiOS in plaintext.

```
config cifs profile
    edit "cifs"
        set server-credential-type credential-keytab
        config server-keytab
            edit "keytab1"
                set keytab
"BQIAAABFAAEAC0VYQU1QTEUuQ09NAAdleGFtcGxlAAAAAVUmAlwBABIAILdV5P6NXT8RrTvapcMJQxDYCjRQiD0Bzxh
wS9h0VgyM"
            next
        end
    next
end
```

Variable	Description
keytab <keytab>	Base64 encoded keytab file containing the credentials of the server.

Configure CIFS profile file filtering

Multiple file filter entries can be added to a profile.

To configure a file filter entry in a CIFS profile:

```
config cifs profile
    edit "cifs"
        config file-filter
            set status {enable | disable}
            set log {enable | disable}
            config entries
                edit <filter>
                    set comment <string>
                    set action {log | block}
                    set direction {incoming | outgoing | any}
                    set file-type <file_type>
                next
            end
        end
    next
end
```

Variable	Description
status {enable | disable}	Enable/disable file filter (default = enable).
log {enable | disable}	Enable/disable file filter logging (default = enable).
comment <string>	A brief comment describing the entry.
action {log | block}	The action to take for matched files: • `log`: Allow the content and write a log message (default). • `block`: Block the content and write a log message.
direction {incoming | outgoing | any}	Match files transmitted in the session's originating (`incoming`) and/or reply (`outgoing`) direction (default = any).
file-type <file_type>	The file types to be matched (default = none). See Supported file types on page 870 for details.

Configure an antivirus profile

The antivirus profile handles the antivirus configuration for CIFS scanning.

To configure an antivirus profile:

```
config antivirus profile
    edit "av"
        ...
        config cifs
            set options {scan avmonitor quarantine}
            set archive-block {encrypted corrupted partiallycorrupted multipart nested
mailbomb fileslimit timeout unhandled}
            set archive-log {encrypted corrupted partiallycorrupted multipart nested
```

```
mailbomb fileslimit timeout unhandled}
            set emulator {enable | disable}
            set outbreak-prevention {disabled | files | full-archive}
        end
    next
end
```

Variable	Description
options {scan avmonitor quarantine}	Enable/disable CIFS antivirus scanning, monitoring, and quarantine.
archive-block {encrypted corrupted partiallycorrupted multipart nested mailbomb fileslimit timeout unhandled}	Select the archive types to block: • `encrypted`: Block encrypted archives. • `corrupted`: Block corrupted archives. • `partiallycorrupted`: Block partially corrupted archives. • `multipart`: Block multipart archives. • `nested`: Block nested archives. • `mailbomb`: Block mail bomb archives. • `fileslimit`: Block exceeded archive files limit. • `timeout`: Block scan timeout. • `unhandled`: Block archives that FortiOS cannot open.
archive-log {encrypted corrupted partiallycorrupted multipart nested mailbomb fileslimit timeout unhandled}	Select the archive types to log: • `encrypted`: Log encrypted archives. • `corrupted`: Log corrupted archives. • `partiallycorrupted`: Log partially corrupted archives. • `multipart`: Log multipart archives. • `nested`: Log nested archives. • `mailbomb`: Log mail bomb archives. • `fileslimit`: Log exceeded archive files limit. • `timeout`: Log scan timeout. • `unhandled`: Log archives that FortiOS cannot open.
emulator {enable | disable}	Enable/disable the virus emulator (default = enable).
outbreak-prevention {disabled | files | full-archive}	Enable the virus outbreak prevention service: • `disabled`: Disabled (default). • `files`: Analyze files as sent, not the content of archives. • `full-archive`: Analyze files, including the content of archives.

Log examples

File-type detection events generated by CIFS profiles are logged in the `utm-cifs` log category. Antivirus detection over the CIFS protocol generates logs in the `utm-virus` category. See the FortiOS Log Message Reference for more information.

Logs generated by CIFS profile file filter:

```
date=2019-03-28 time=10:39:19 logid="1800063001" type="utm" subtype="cifs" eventtype="cifs-
filefilter" level="notice" vd="vdom1" eventtime=1553794757 msg="File was detected by file
```


Netskope Exhibit 1017

```
filter." direction="incoming" action="passthrough" service="CIFS" srcip=10.1.100.11
dstip=172.16.200.44 srcport=33372 dstport=445 srcintf="wan2" srcintfrole="wan"
dstintf="wan1" dstintfrole="wan" policyid=1 proto=16 profile="cifs" filesize="1154"
filename="virus\\test.png" filtername="2" filetype="png"

date=2019-03-28 time=10:39:12 logid="1800063001" type="utm" subtype="cifs" eventtype="cifs-
filefilter" level="notice" vd="vdom1" eventtime=1553794751 msg="File was detected by file
filter." direction="incoming" action="passthrough" service="CIFS" srcip=10.1.100.11
dstip=172.16.200.44 srcport=33370 dstport=445 srcintf="wan2" srcintfrole="wan"
dstintf="wan1" dstintfrole="wan" policyid=1 proto=16 profile="cifs" filesize="81975"
filename="virus\\screen.png" filtername="2" filetype="png"

date=2019-03-28 time=10:33:55 logid="1800063000" type="utm" subtype="cifs" eventtype="cifs-
filefilter" level="warning" vd="vdom1" eventtime=1553794434 msg="File was blocked by file
filter." direction="incoming" action="blocked" service="CIFS" srcip=10.1.100.11
dstip=172.16.200.44 srcport=33352 dstport=445 srcintf="wan2" srcintfrole="wan"
dstintf="wan1" dstintfrole="wan" policyid=1 proto=16 profile="cifs" filesize="28432"
filename="filetypes\\mpnotify.exe" filtername="3" filetype="exe"

date=2019-03-28 time=10:33:45 logid="1800063000" type="utm" subtype="cifs" eventtype="cifs-
filefilter" level="warning" vd="vdom1" eventtime=1553794424 msg="File was blocked by file
filter." direction="incoming" action="blocked" service="CIFS" srcip=10.1.100.11
dstip=172.16.200.44 srcport=33348 dstport=445 srcintf="wan2" srcintfrole="wan"
dstintf="wan1" dstintfrole="wan" policyid=1 proto=16 profile="cifs" filesize="96528"
filename="filetypes\\winmine.exe" filtername="3" filetype="exe"
```

Logs generated by AV profile for infections detected over CIFS:

```
date=2019-04-09 time=15:19:02 logid="0204008202" type="utm" subtype="virus"
eventtype="outbreak-prevention" level="warning" vd="vdom1" eventtime=1554848342519005401
msg="Blocked by Virus Outbreak Prevention service." action="blocked" service="SMB"
sessionid=177 srcip=10.1.100.11 dstip=172.16.200.44 srcport=37444 dstport=445 srcintf="wan2"
srcintfrole="wan" dstintf="wan1" dstintfrole="wan" policyid=1 proto=6 direction="incoming"
filename="outbreak\\zhvo_test.com" quarskip="File-was-not-quarantined."
virus="503e99fe40ee120c45bc9a30835e7256fff3e46a" dtype="File Hash"
filehash="503e99fe40ee120c45bc9a30835e7256fff3e46a" filehashsrc="fortiguard" profile="av"
analyticssubmit="false" crscore=50 craction=2 crlevel="critical"

date=2019-04-09 time=15:18:59 logid="0211008192" type="utm" subtype="virus"
eventtype="infected" level="warning" vd="vdom1" eventtime=1554848339909808987 msg="File is
infected." action="blocked" service="SMB" sessionid=174 srcip=10.1.100.11
dstip=172.16.200.44 srcport=37442 dstport=445 srcintf="wan2" srcintfrole="wan"
dstintf="wan1" dstintfrole="wan" policyid=1 proto=6 direction="incoming"
filename="sample\\eicar.com" quarskip="File-was-not-quarantined." virus="EICAR_TEST_FILE"
dtype="Virus" ref="http://www.fortinet.com/ve?vn=EICAR_TEST_FILE" virusid=2172 profile="av"
analyticscksum="275a021bbfb6489e54d471899f7db9d1663fc695ec2fe2a2c4538aabf651fd0f"
analyticssubmit="false" crscore=50 craction=2 crlevel="critical"
```

Using FortiSandbox appliance with antivirus

Antivirus can use FortiSandbox to supplement its detection capabilities. In real-world situations, networks are always under the threat of zero-day attacks.

Antivirus can submit potential zero-day viruses to FortiSandbox for inspection. Based on FortiSandbox's analysis, the FortiGate can supplement its own antivirus database with FortiSandbox's database to detect files determined as

Netskope Exhibit 1017

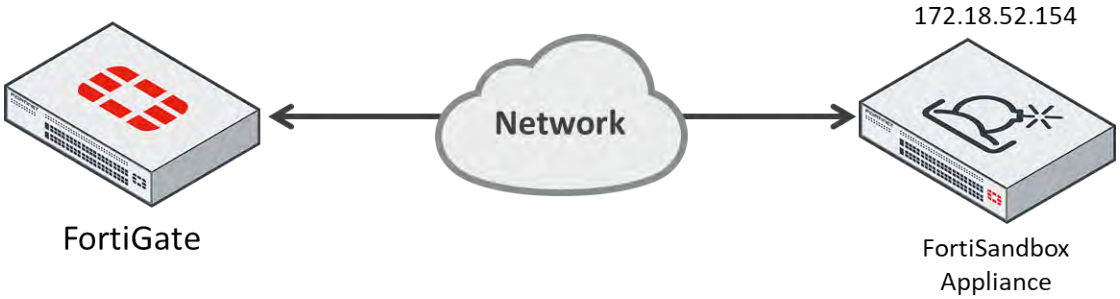
malicious or risky by FortiSandbox. This helps FortiGate antivirus detect zero-day viruses and malware whose signatures are not found in the FortiGate antivirus database.

For more information, see FortiSandbox on page 130.

Support and limitations

- FortiSandbox can be used with antivirus in both proxy-based and flow-based inspection modes.
- With FortiSandbox enabled, full scan mode antivirus can do the following:
  - Submit only suspicious files to FortiSandbox for inspection.
  - Submit every file to FortiSandbox for inspection.
  - Do not submit anything.
- Quick scan mode antivirus cannot submit suspicious files to FortiSandbox. It can only do the following:
  - Submit every file to FortiSandbox for inspection.
  - Do not submit anything.

Network topology example



Configuring the feature

To configure antivirus to work with an external block list, the following steps are required:

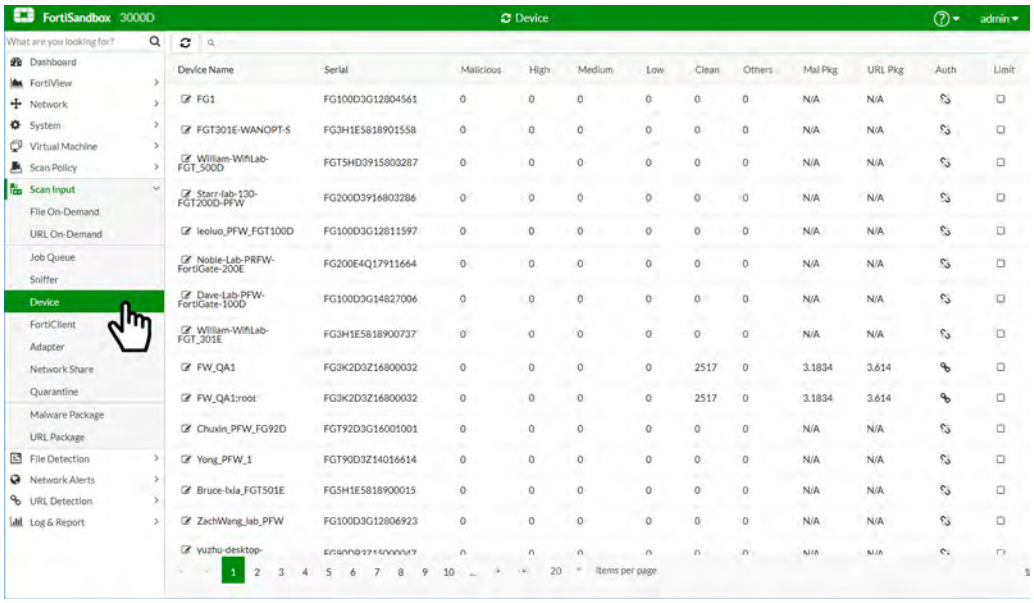
1. Enable FortiSandbox on the FortiGate.
2. Authorize FortiGate on the FortiSandbox.
3. Enable FortiSandbox inspection.
4. Enable use of the FortiSandbox database.

To enable FortiSandbox on the FortiGate:

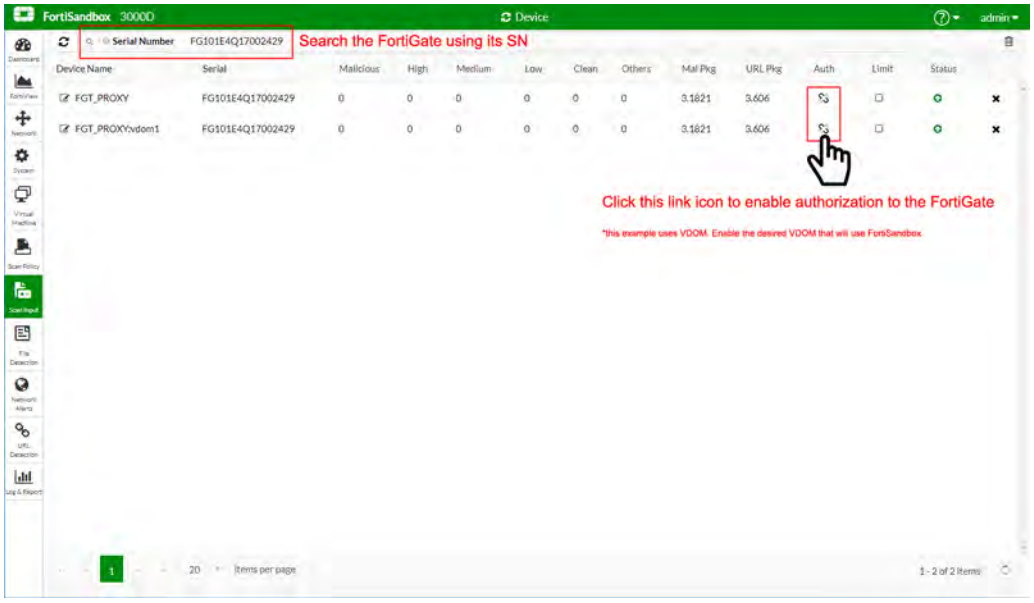
1. Go to *Global* > *Security Fabric* > *Fabric Connectors* and double-click the *FortiSandbox* card.
2. For status, click *Enable*.
3. For *Type*, click *On-Premise*.
4. Enter the IP address of the FortiSandbox, and enter an optional *Notifier email*.
5. At this point, clicking *Test connectivity* returns an unreachable status. This is expected, because the FortiGate is not yet authorized by FortiSandbox.
6. Click *OK* to save the settings.

To authorize FortiGate on the FortiSandbox:

1. In the FortiSandbox Appliance GUI, go to *Scan Input* > *Device*.

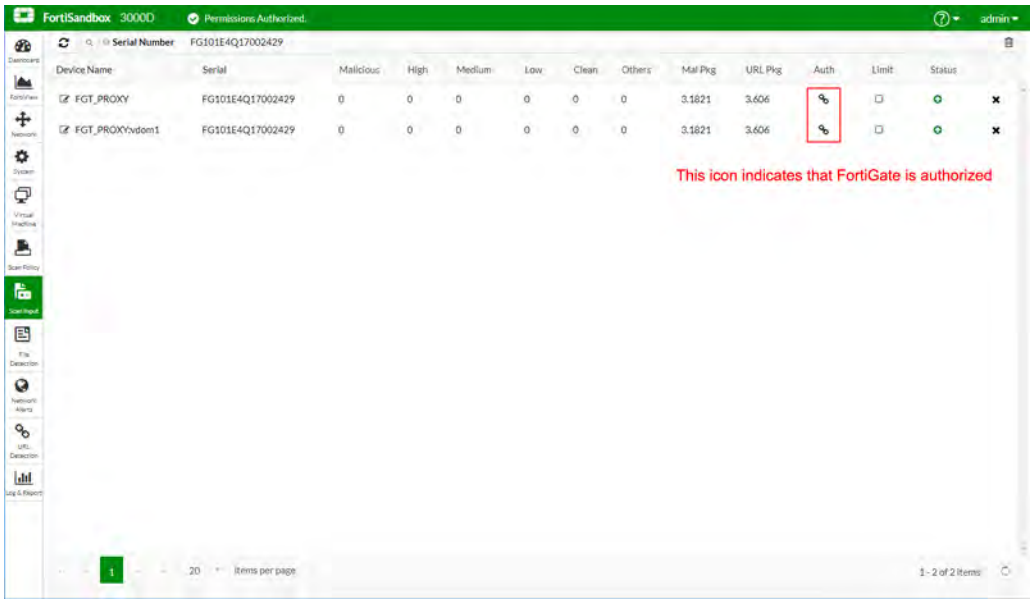


2. Use the FortiGate serial number to quickly locate the desired FortiGate and select the *link* icon to authorize the FortiGate.
3. Enable the desired VDOM in the same manner.

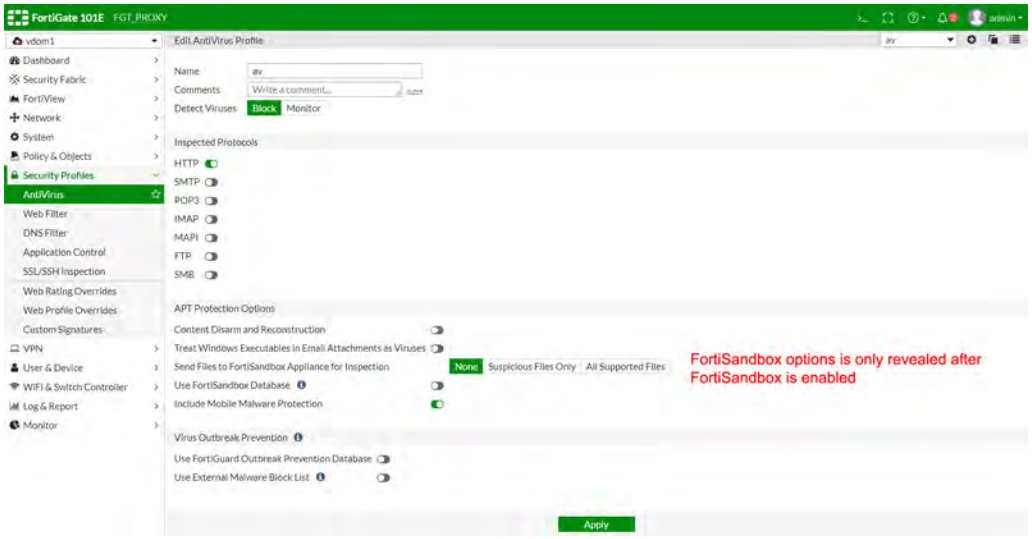


4. The *link* icon changes from an open to closed link. This indicates that the FortiSandbox has authorized this FortiGate.

Netskope Exhibit 1017

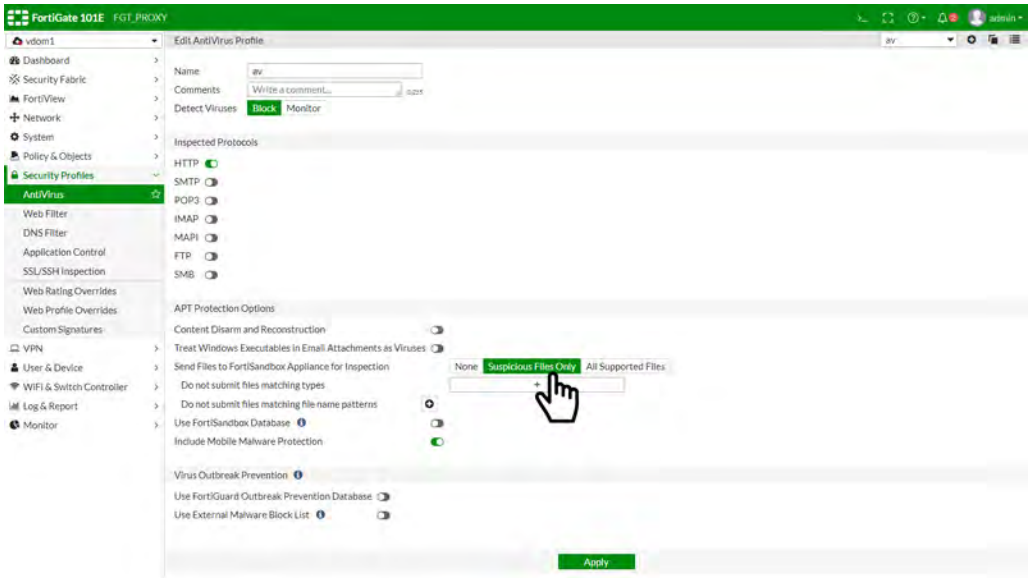


5. In the FortiGate GUI, go to *Global* > *Security Fabric* > *Fabric Connectors*.
6. Click *Test connectivity*. The FortiGate is now authorized and the status displays as *Connected*.
7. FortiSandbox options are now displayed in the *AV Profile* page.

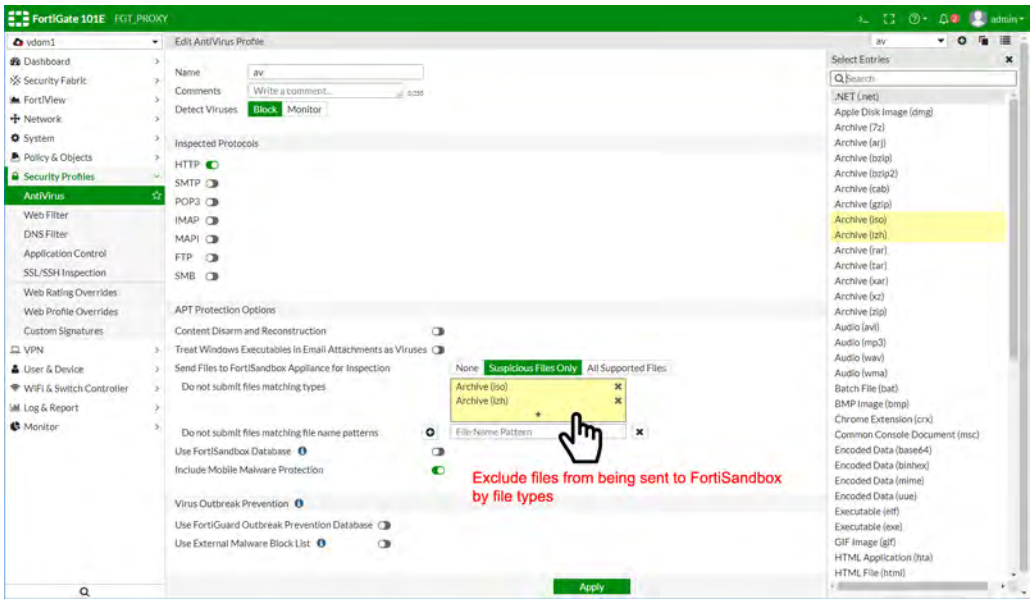


To enable FortiSandbox inspection:

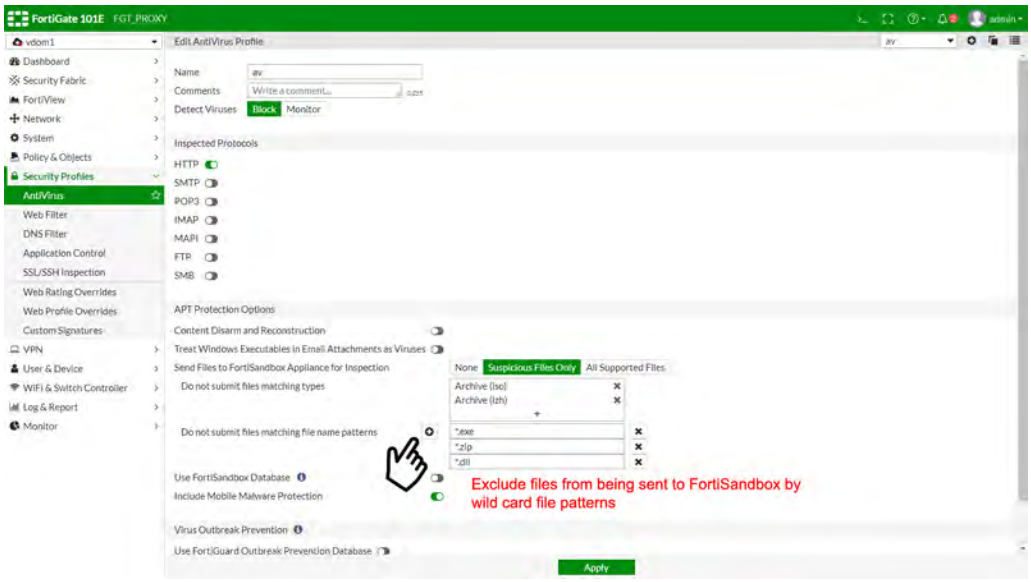
1. Go to *Security Profiles* > *AntiVirus*.
2. Enable FortiSandbox inspection by selecting either *Suspicious Files Only* or *All Supported Files*.



3. Files can be excluded from being sent to FortiSandbox based on their file types by choosing from a list of supported file types.



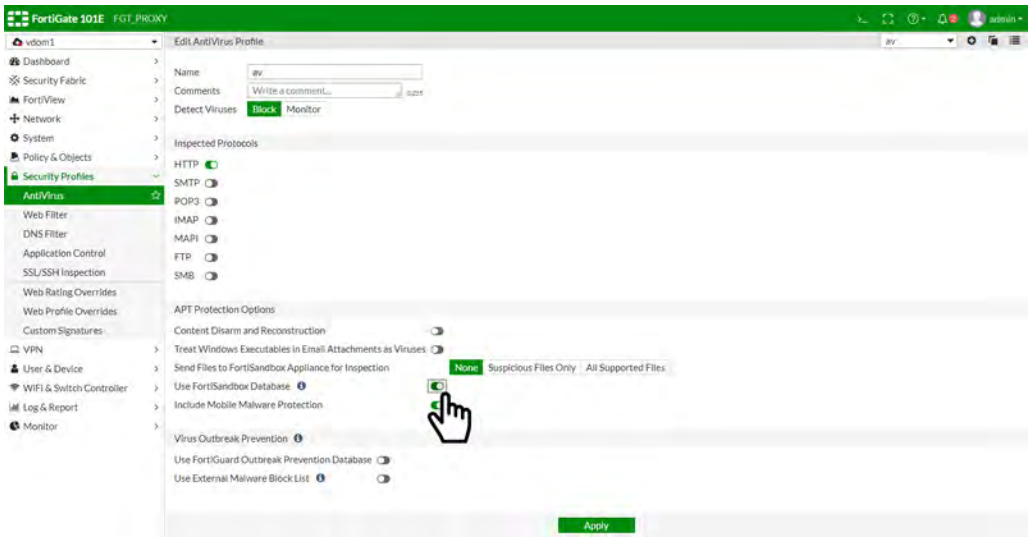
4. Files can also be excluded from being sent to FortiSandbox by using wildcard patterns.



5. Click *Apply*.

To enable use of the FortiSandbox database:

1. Go to *Security Profiles* > *AntiVirus*
2. Enable *Use FortiSandbox Database*.



3. Click *Apply*.

Netskope Exhibit 1017

Diagnostics and Debugging

Debug on the FortiGate side

- Update daemon:

```
FGT_PROXY (global) # diagnose debug application quarantined -1
FGT_PROXY (global) # diagnose debug enable

quar_req_fsa_file()-890: fsa ext list new_version (1547781904)
quar_fsb_handle_quar()-1439: added a req-6 to fortisandbox-fsb5, vfid=1, oftp-name=[].
__quar_start_connection()-908: start server fortisandbox-fsb5-172.18.52.154 in vdom-1
[103] __ssl_cert_ctx_load: Added cert /etc/cert/factory/root_Fortinet_Factory.cer, root
ca Fortinet_CA, idx 0 (default)
[551] ssl_ctx_create_new_ex: SSL CTX is created
[578] ssl_new: SSL object is created
upd_cfg_extract_av_db_version[378]-version=06002000AVDB00201-00066.01026-1901301530
upd_cfg_extract_ids_db_version[437]-version=06002000NIDS02403-00014.00537-1901300043
upd_cfg_extract_ids_db_version[437]-version=06002000APDB00103-00006.00741-1512010230
upd_cfg_extract_ids_db_version[437]-version=06002000ISDB00103-00014.00537-1901300043
upd_cfg_extract_ibdb_botnet_db_version[523]-version=06002000IBDB00101-00004.00401-
1901281000
upd_cfg_extract_av_db_version[378]-version=06002000AVDB00201-00066.01026-1901301530
upd_cfg_extract_ids_db_version[437]-version=06002000NIDS02403-00014.00537-1901300043
upd_cfg_extract_ids_db_version[437]-version=06002000APDB00103-00006.00741-1512010230
upd_cfg_extract_ids_db_version[437]-version=06002000ISDB00103-00014.00537-1901300043
upd_cfg_extract_ibdb_botnet_db_version[523]-version=06002000IBDB00101-00004.00401-
1901281000
upd_cfg_extract_av_db_version[378]-version=06002000AVDB00201-00066.01026-1901301530
upd_cfg_extract_ids_db_version[437]-version=06002000NIDS02403-00014.00537-1901300043
upd_cfg_extract_ids_db_version[437]-version=06002000APDB00103-00006.00741-1512010230
upd_cfg_extract_ids_db_version[437]-version=06002000ISDB00103-00014.00537-1901300043
upd_cfg_extract_ibdb_botnet_db_version[523]-version=06002000IBDB00101-00004.00401-
1901281000
upd_cfg_extract_av_db_version[378]-version=06002000AVDB00201-00066.01026-1901301530
upd_cfg_extract_ids_db_version[437]-version=06002000NIDS02403-00014.00537-1901300043
upd_cfg_extract_ids_db_version[437]-version=06002000APDB00103-00006.00741-1512010230
upd_cfg_extract_ids_db_version[437]-version=06002000ISDB00103-00014.00537-1901300043
upd_cfg_extract_ibdb_botnet_db_version[523]-version=06002000IBDB00101-00004.00401-
1901281000
upd_cfg_extract_av_db_version[378]-version=06002000AVDB00201-00066.01026-1901301530
upd_cfg_extract_ids_db_version[437]-version=06002000NIDS02403-00014.00537-1901300043
upd_cfg_extract_ids_db_version[437]-version=06002000APDB00103-00006.00741-1512010230
upd_cfg_extract_ids_db_version[437]-version=06002000ISDB00103-00014.00537-1901300043
upd_cfg_extract_ibdb_botnet_db_version[523]-version=06002000IBDB00101-00004.00401-
1901281000
quar_remote_recv_send()-731: dev=fortisandbox-fsb2 xfer-status=0
__quar_build_pkt()-408: build req(id=337, type=4) for vdom-vdom1, len=99, oftp_name=
__quar_send()-470: dev buffer -- pos=0, len=99
quar_remote_send()-520: req(id=337, type=4) read response, dev=fortisandbox-fsb2, xfer_
status=1, buflen=12
quar_remote_recv_send()-770: dev-fortisandbox-fsb2, oevent=4, nevent=1, xfer-status=1
quar_remote_recv_send()-731: dev=fortisandbox-fsb3 xfer-status=0
__quar_build_pkt()-408: build req(id=338, type=6) for vdom-vdom1, len=93, oftp_name=
__quar_send()-470: dev buffer -- pos=0, len=93
quar_remote_send()-520: req(id=338, type=6) read response, dev=fortisandbox-fsb3, xfer_
```


Netskope Exhibit 1017

```
status=1, buflen=12
quar_remote_recv_send()-770: dev-fortisandbox-fsb3, oevent=4, nevent=1, xfer-status=1
quar_remote_recv_send()-731: dev=fortisandbox-fsb5 xfer-status=0
__quar_build_pkt()-408: build req(id=340, type=6) for vdom-vdom1, len=93, oftp_name=
__quar_send()-470: dev buffer -- pos=0, len=93
quar_remote_send()-520: req(id=340, type=6) read response, dev=fortisandbox-fsb5, xfer_
status=1, buflen=12
quar_remote_recv_send()-770: dev-fortisandbox-fsb5, oevent=4, nevent=1, xfer-status=1
quar_remote_recv_send()-731: dev=fortisandbox-fsb2 xfer-status=1
quar_remote_recv()-662: dev(fortisandbox-fsb2) received a packet: len=69, type=1
quar_remote_recv()-718: file-[337] is accepted by server(fortisandbox-fsb2).
quar_put_job_req()-332: Job 337 deleted
quar_remote_recv_send()-731: dev=fortisandbox-fsb4 xfer-status=0
__quar_build_pkt()-408: build req(id=339, type=6) for vdom-vdom1, len=93, oftp_name=
__quar_send()-470: dev buffer -- pos=0, len=93
quar_remote_send()-520: req(id=339, type=6) read response, dev=fortisandbox-fsb4, xfer_
status=1, buflen=12
quar_remote_recv_send()-770: dev-fortisandbox-fsb4, oevent=4, nevent=1, xfer-status=1
quar_remote_recv_send()-731: dev=fortisandbox-fsb1 xfer-status=0
__quar_build_pkt()-408: build req(id=336, type=4) for vdom-root, len=98, oftp_name=
__quar_send()-470: dev buffer -- pos=0, len=98
...
__get_analytics_stats()-19: Received an ANALYTICS_STATS request, vfid: 0
__quar_req_handler()-127: Request 0 was handled successfully
__get_analytics_stats()-19: Received an ANALYTICS_STATS request, vfid: 0
__quar_req_handler()-127: Request 0 was handled successfully
__get_analytics_stats()-19: Received an ANALYTICS_STATS request, vfid: 0
__quar_req_handler()-127: Request 0 was handled successfully
__get_analytics_stats()-19: Received an ANALYTICS_STATS request, vfid: 0
__quar_req_handler()-127: Request 0 was handled successfully
quar_fsb_handle_quar()-1439: added a req-6 to fortisandbox-fsb1, vfid=1, oftp-name=[].
__quar_start_connection()-908: start server fortisandbox-fsb1-172.18.52.154 in vdom-1
[103] __ssl_cert_ctx_load: Added cert /etc/cert/factory/root_Fortinet_Factory.cer, root
ca Fortinet_CA, idx 0 (default)
[551] ssl_ctx_create_new_ex: SSL CTX is created
[578] ssl_new: SSL object is created
upd_cfg_extract_av_db_version[378]-version=06002000AVDB00201-00066.01026-1901301530
upd_cfg_extract_ids_db_version[437]-version=06002000NIDS02403-00014.00537-1901300043
upd_cfg_extract_ids_db_version[437]-version=06002000APDB00103-00006.00741-1512010230
upd_cfg_extract_ids_db_version[437]-version=06002000ISDB00103-00014.00537-1901300043
upd_cfg_extract_ibdb_botnet_db_version[523]-version=06002000IBDB00101-00004.00401-
1901281000
quar_remote_recv_send()-731: dev=fortisandbox-fsb1 xfer-status=0
__quar_build_pkt()-408: build req(id=2, type=6) for vdom-vdom1, len=93, oftp_name=
__quar_send()-470: dev buffer -- pos=0, len=93
quar_remote_send()-520: req(id=2, type=6) read response, dev=fortisandbox-fsb1, xfer_
status=1, buflen=12
quar_remote_recv_send()-770: dev-fortisandbox-fsb1, oevent=4, nevent=1, xfer-status=1
quar_remote_recv_send()-731: dev=fortisandbox-fsb1 xfer-status=1
quar_remote_recv()-662: dev(fortisandbox-fsb1) received a packet: len=767, type=1
quar_store_analytics_report()-590: Analytics-report return
file=/tmp/fsb/83bb2d9928b03a68b123730399b6b9365b5cc9a5a77f8aa007a6f1a499a13b18.json.gz,
buf_sz=735
quar_store_analytics_report()-597: The request
'83bb2d9928b03a68b123730399b6b9365b5cc9a5a77f8aa007a6f1a499a13b18' score is 1
quar_remote_recv()-718: file-[2] is accepted by server(fortisandbox-fsb1).
```


Netskope Exhibit 1017

```
quar_put_job_req()-332: Job 2 deleted
quar_monitor_connection_func()-978: monitoring dev fortisandbox-fsb1
quar_monitor_connection_func()-978: monitoring dev fortisandbox-fsb1
__get_analytics_stats()-19: Received an ANALYTICS_STATS request, vfid: 0
__quar_req_handler()-127: Request 0 was handled successfully
__get_analytics_stats()-19: Received an ANALYTICS_STATS request, vfid: 0
__quar_req_handler()-127: Request 0 was handled successfully
quar_monitor_connection_func()-978: monitoring dev fortisandbox-fsb1
quar_stop_connection()-1006: close connection to server(fortisandbox-fsb1)
[193] __ssl_data_ctx_free: Done
[805] ssl_free: Done
[185] __ssl_cert_ctx_free: Done
[815] ssl_ctx_free: Done
[796] ssl_disconnect: Shutdown
```

- Appliance FortiSandbox diagnostics:

```
FGT_PROXY # config global
FGT_PROXY (global) # diagnose test application quarantined 1
Total remote&local devices: 8, any task full? 0
System have disk, vdom is enabled, mgmt=1, ha=2
xfer-fas is enabled: ips-archive dlp-archive, realtime=yes, taskfull=no
    addr=0.0.0.0/514, source-ip=0.0.0.0, keep-alive=no.
    ssl_opt=0, hmac_alg=0
    License=0, content_archive=0, arch_pause=0.

global-fas is disabled.
forticloud-fsb is disabled.
fortisandbox-fsb1 is enabled: analytics, realtime=yes, taskfull=no
    addr=172.18.52.154/514, source-ip=0.0.0.0, keep-alive=no.
    ssl_opt=3, hmac_alg=0
fortisandbox-fsb2 is enabled: analytics, realtime=yes, taskfull=no
    addr=172.18.52.154/514, source-ip=0.0.0.0, keep-alive=no.
    ssl_opt=3, hmac_alg=0
fortisandbox-fsb3 is enabled: analytics, realtime=yes, taskfull=no
    addr=172.18.52.154/514, source-ip=0.0.0.0, keep-alive=no.
    ssl_opt=3, hmac_alg=0
fortisandbox-fsb4 is enabled: analytics, realtime=yes, taskfull=no
    addr=172.18.52.154/514, source-ip=0.0.0.0, keep-alive=no.
    ssl_opt=3, hmac_alg=0
fortisandbox-fsb5 is enabled: analytics, realtime=yes, taskfull=no
    addr=172.18.52.154/514, source-ip=0.0.0.0, keep-alive=no.
    ssl_opt=3, hmac_alg=0
fortisandbox-fsb6 is enabled: analytics, realtime=yes, taskfull=no
    addr=172.18.52.154/514, source-ip=0.0.0.0, keep-alive=no.
    ssl_opt=3, hmac_alg=0
global-faz is disabled.
global-faz2 is disabled.
global-faz3 is disabled.
```

- Checking FortiSandbox analysis statistics:

```
FGT_PROXY (global) # diagnose test application quarantine 7
Total: 0

Statistics:
        vfid: 0, detected: 0, clean: 0, risk_low: 0, risk_med: 0, risk_high: 0, limit_
reached:0
```


Netskope Exhibit 1017

```
        vfid: 3, detected: 0, clean: 0, risk_low: 0, risk_med: 0, risk_high: 0, limit_
reached:0
        vfid: 4, detected: 0, clean: 0, risk_low: 0, risk_med: 0, risk_high: 0, limit_
reached:0

FGT_PROXY (global) #
```

Debug on the FortiSandbox side

- Appliance FortiSandbox OFTP debug:

```
# diagnose-debug device FG101E4Q17002429

[2019/01/31 00:48:21] LOGIN->SUCCEED: Serial(FG101E4Q17002429), HOSTNAME(FGT_PROXY)
[2019/01/31 00:48:21] FG101E4Q17002429 VDOM: vdom1
[2019/01/31 00:48:21] FG101E4Q17002429 suspicious stats START_TIME: 1548290749
[2019/01/31 00:48:21] FG101E4Q17002429 suspicious stats END_TIME: 1548895549
[2019/01/31 00:48:21] FG101E4Q17002429 opd_data_len=37 clean=2 detected=2 risk_low=0
risk_med=0 risk_high=0 sus_limit=0
[2019/01/31 00:48:21] FG101E4Q17002429 ENTERING->HANDLE_SEND_FILE.
[2019/01/31 00:48:21] FG101E4Q17002429 ENTERING->HANDLE_SEND_FILE.
[2019/01/31 00:48:21] FG101E4Q17002429 INCOMING->FGT->VDOM: vdom1
[2019/01/31 00:48:21] FG101E4Q17002429 INCOMING->FGT->VDOM: vdom1
[2019/01/31 00:48:21] FG101E4Q17002429 INCOMING->IMG_VERSION: 6.2.0.0818
[2019/01/31 00:48:21] FG101E4Q17002429 INCOMING->IMG_VERSION: 6.2.0.0818
[2019/01/31 00:48:21] INCOMING->FGT: FG101E4Q17002429, VDOM: vdom1
[2019/01/31 00:48:21] INCOMING->FGT: FG101E4Q17002429, VDOM: vdom1
[2019/01/31 00:48:21] FG101E4Q17002429 INCOMING->TYPE: 0
[2019/01/31 00:48:21] FG101E4Q17002429 INCOMING->TYPE: 1
[2019/01/31 00:48:21] FG101E4Q17002429 INCOMING->VERSION: 3 . 1795
[2019/01/31 00:48:21] FG101E4Q17002429 INCOMING->VERSION: 3 . 595
[2019/01/31 00:48:21] FG101E4Q17002429 VDOM: root
[2019/01/31 00:48:21] FG101E4Q17002429 ENTERING->HANDLE_SEND_FILE.
[2019/01/31 00:48:21] FG101E4Q17002429 suspicious stats START_TIME: 1548290749
[2019/01/31 00:48:21] FG101E4Q17002429 INCOMING->FGT->VDOM: vdom1
[2019/01/31 00:48:21] FG101E4Q17002429 suspicious stats END_TIME: 1548895549
[2019/01/31 00:48:21] FG101E4Q17002429 INCOMING->IMG_VERSION: 6.2.0.0818
[2019/01/31 00:48:21] INCOMING->FGT: FG101E4Q17002429, VDOM: vdom1
[2019/01/31 00:48:21] FG101E4Q17002429 INCOMING->TYPE: 4
[2019/01/31 00:48:21] FG101E4Q17002429 opd_data_len=37 clean=0 detected=0 risk_low=0
risk_med=0 risk_high=0 sus_limit=0
[2019/01/31 00:48:22] FG101E4Q17002429 RETRIEVE->PKG: TYPE: av, ENTRY_VERSION: 1795,
PACKAGE_PATH: /Storage/malpkg/pkg/avsig/avsigrel_1795.pkg
[2019/01/31 00:48:22] FG101E4Q17002429 RETRIEVE->PKG: TYPE: url, ENTRY_VERSION: 595,
PACKAGE_PATH: /Storage/malpkg/pkg/url/urlrel_595.pkg.gz
[2019/01/31 00:48:29] LOGIN->SUCCEED: Serial(FG101E4Q17002429), HOSTNAME(FGT_PROXY)
[2019/01/31 00:48:32] LOGIN->SUCCEED: Serial(FG101E4Q17002429), HOSTNAME(FGT_PROXY)
[2019/01/31 00:48:59] LOGIN->SUCCEED: Serial(FG101E4Q17002429), HOSTNAME(FGT_PROXY)
[2019/01/31 00:49:03] LOGIN->SUCCEED: Serial(FG101E4Q17002429), HOSTNAME(FGT_PROXY)
```


Netskope Exhibit 1017

Using FortiSandbox Cloud with antivirus

Feature overview

FortiSandbox Cloud allows users to take advantage of FortiSandbox features without having to purchase, operate, and maintain a physical appliance. It works the same way as the physical FortiSandbox appliance.

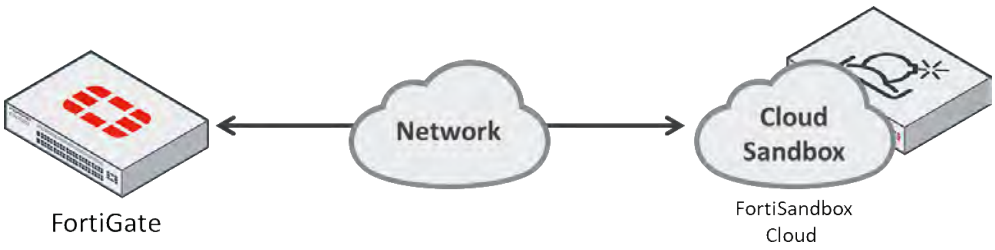
FortiSandbox Cloud allows you to control the region where your traffic is sent to for analysis. This allows you to meet your country's compliances regarding data storage locations.

For more information, see FortiSandbox on page 130.

Support and limitations

- In FortiOS 6.2 and later, users do not require a FortiGate Cloud account to use FortiSandbox Cloud.
- Without a valid AVDB license, FortiGate devices are limited to 100 FortiGate Cloud submissions per day.
- Unlimited FortiGate Cloud submissions are allowed if the FortiGate has a valid AVDB license.
  - There is a limit on how many submissions are sent per minute.
  - The per-minute submission rate is based on the FortiGate model.
- FortiSandbox can be used with antivirus in both proxy-based and flow-based policy inspection modes.
- With FortiSandbox enabled, full scan mode antivirus can do the following:
  - Submit only suspicious files to FortiSandbox for inspection.
  - Submit every file to FortiSandbox for inspection.
  - Do not submit anything.
- Quick scan mode antivirus cannot submit suspicious files to FortiSandbox. It can only do the following:
  - Submit every file to FortiSandbox for inspection.
  - Do not submit anything.

Network topology example



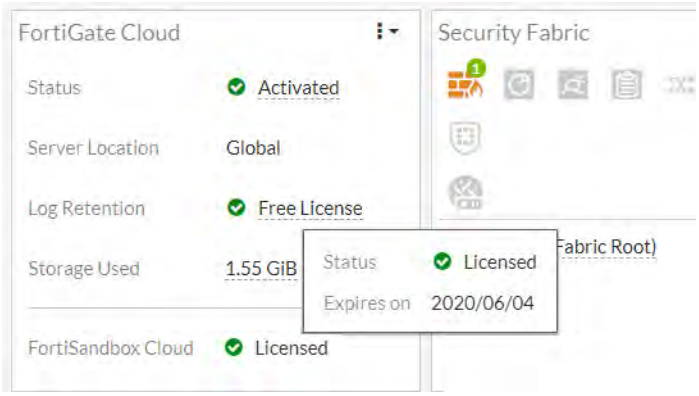
Configuring the feature

To configure antivirus to work with an external block list, the following steps are required:

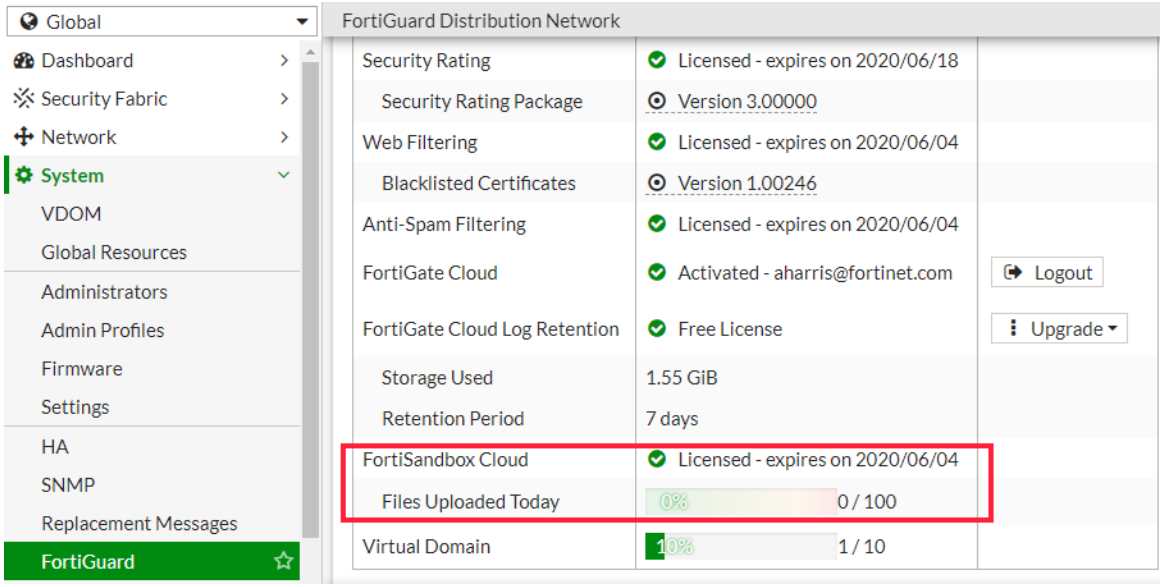
1. Through FortiCare, register the FortiGate device and purchase a FortiGuard antivirus license
2. Enable FortiSandbox Cloud on the FortiGate
3. Enable FortiSandbox inspection
4. Enable the use of the FortiSandbox database

To obtain or renew an AVDB license:

1. Please see the video How to Purchase or Renew FortiGuard Services for FortiGuard antivirus license purchase instructions.
2. Once a FortiGuard license has been purchased or activated, users will be provided with a paid FortiSandbox Cloud license.
   a. Go to *Global* > *Dashboard* > *Status* to view the FortiSandbox Cloud license indicator.



   b. Users can also view this indicator at *Global* > *System* > *FortiGuard*.

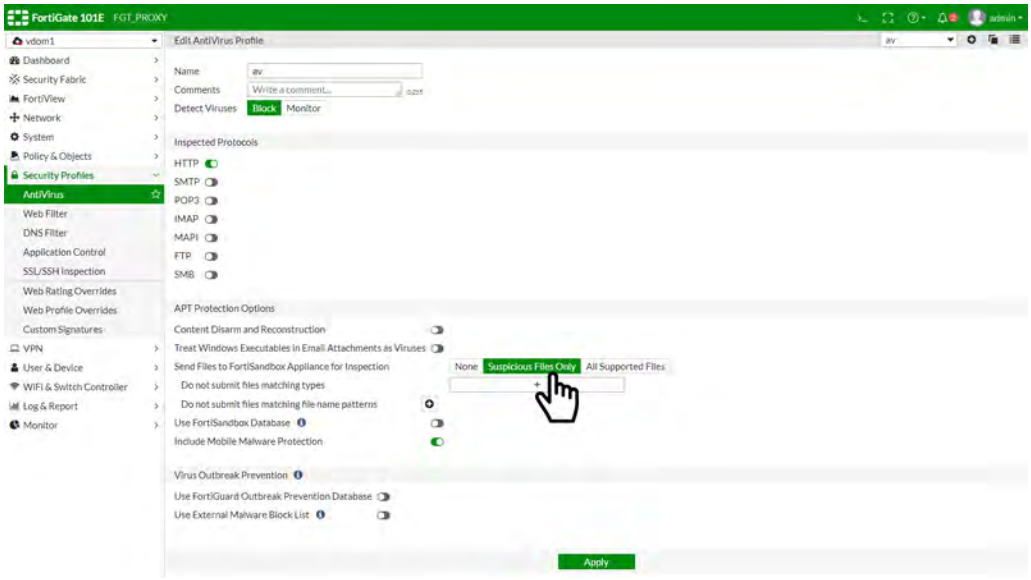


To enable FortiSandbox Cloud on the FortiGate:

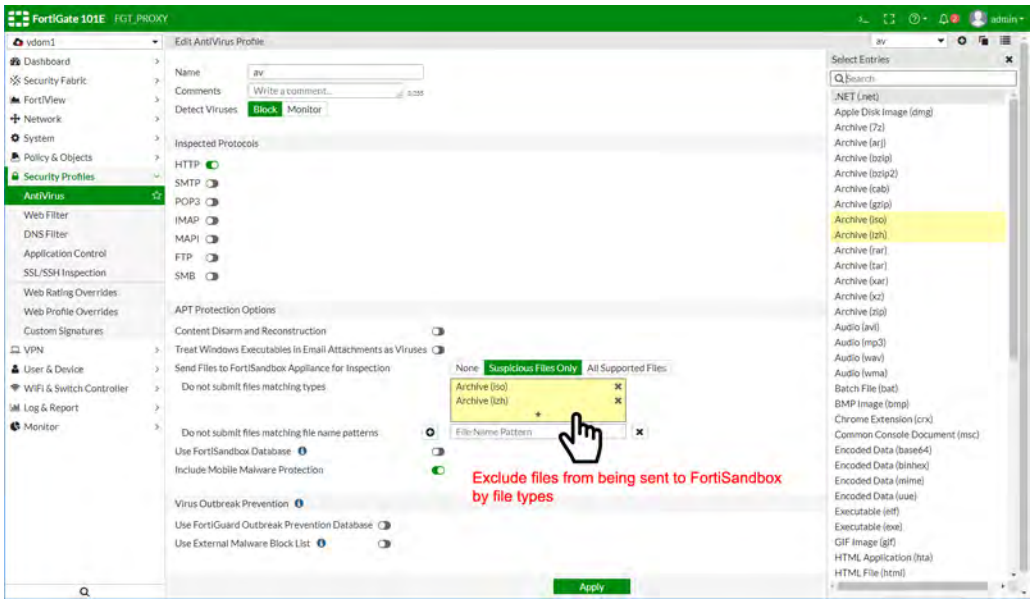
1. Go to *Global* > *Security Fabric* > *Fabric Connectors* and double-click the *FortiSandbox* card.
2. For status, click *Enable*.
3. For *Type*, click *FortiSandbox Cloud* and choose a region from the dropdown list.
4. Click *OK* to save the settings.
5. When the FortiGate is connected to the FortiSandbox Cloud, FortiSandbox's current database version is displayed.

To enable FortiSandbox inspection:

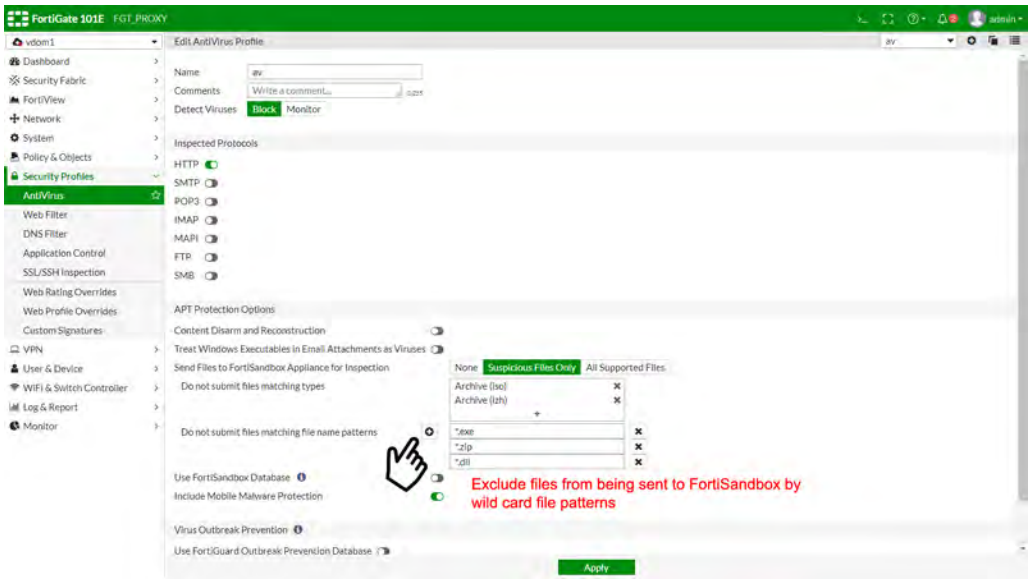
1. Go to *Security Profiles* > *AntiVirus*.
2. Enable FortiSandbox inspection by selecting either *Suspicious Files Only* or *All Supported Files*.



3. Files can be excluded from being sent to FortiSandbox based on their file types by choosing from a list of supported file types.



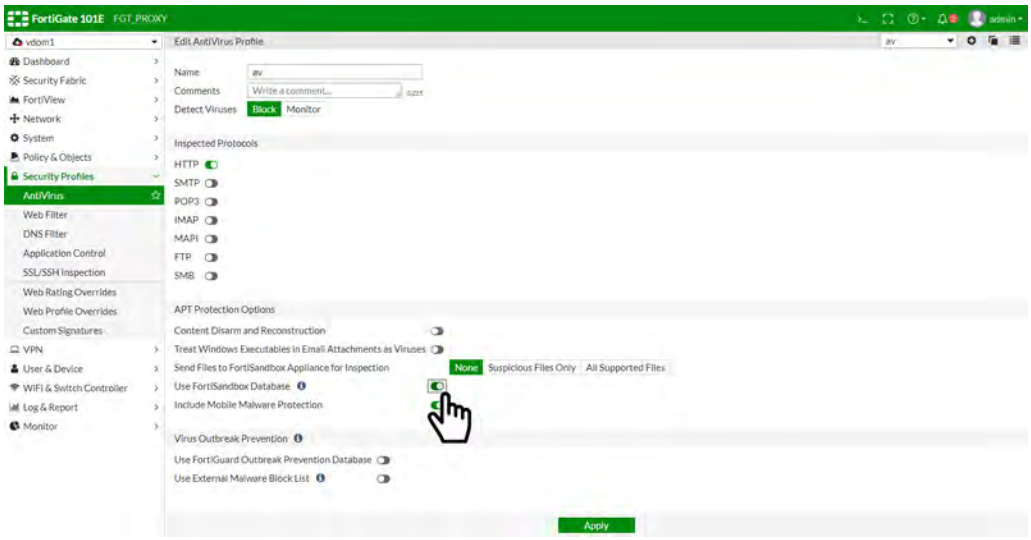
4. Files can also be excluded from being sent to FortiSandbox by using wildcard patterns.



5. Click *Apply*.

To enable the use of the FortiSandbox database:

1. Go to *Security Profiles* > *AntiVirus*.
2. Enable *Use FortiSandbox Database*.



3. Click *Apply*.

Diagnostics and debugging

- Checking FortiGate Cloud controller status:

```
FGT_FL_FULL (global) # diagnose test application forticldd 2
Server: log-controller, task=0/10, watchdog is off
Domain name: logctrl1.fortinet.com
```

```
Address of log-controller: 1
        172.16.95.168:443
        Statistics: total=3, discarded=1, sent=2, last_updated=12163 secs ago
http connection: is not in progress
        Current address: 172.16.95.168:443
        Calls: connect=9, rxtx=12
Current tasks number: 0
Account: name=empty, status=0, type=basic
Current volume: 0B
Current tasks number: 0
Update timer fires in 74240 secs
```

- Checking Cloud APT server status:

```
FGT_FL_FULL (global) # diagnose test application forticldd 3
Debug zone info:
    Domain:
    Home log server: 0.0.0.0:0
    Alt log server: 0.0.0.0:0
    Active Server IP:      0.0.0.0
    Active Server status:  down
    Log quota:      0MB
    Log used:       0MB
    Daily volume:   0MB
    fams archive pause: 0
    APTContract : 1                          <====
    APT server: 172.16.102.51:514            <====
    APT Altserver: 172.16.102.52:514         <====
    Active APTServer IP:      172.16.102.51  <====
    Active APTServer status:  up             <====
```

- FortiSandbox Cloud diagnostics:

```
FGT_FL_FULL (global) # diagnose test application quarantine 1
Total remote&local devices: 4, any task full? 0
System have disk, vdom is enabled, mgmt=3, ha=1
xfer-fas is enabled: ips-archive dlp-archive, realtime=yes, taskfull=no
    addr=0.0.0.0/514, source-ip=0.0.0.0, keep-alive=no.
    ssl_opt=1, hmac_alg=0
    License=0, content_archive=0, arch_pause=0.

global-fas is disabled.
forticloud-fsb is enabled: analytics, realtime=yes, taskfull=no
    addr=172.16.102.51/514, source-ip=0.0.0.0, keep-alive=no.
    ssl_opt=1, hmac_alg=0
fortisandbox-fsb1 is disabled.
fortisandbox-fsb2 is disabled.
fortisandbox-fsb3 is disabled.
fortisandbox-fsb4 is disabled.
fortisandbox-fsb5 is disabled.
fortisandbox-fsb6 is disabled.
global-faz is disabled.
global-faz2 is disabled.
global-faz3 is disabled.
```


Netskope Exhibit 1017

- Checking FortiSandbox Cloud submission statistics:

```
FGT_FL_FULL (global) # diagnose test application quarantined 2
Quarantine daemon state:
QUAR mem: mem_used=0, mem_limit=97269, threshold=72951
dropped(0 by quard, 0 by callers)
pending-jobs=0, tot-mem=0, last_ipc_run=12353, check_new_req=1
alloc_job_failed=0, job_wrong_type=0, job_wrong_req_len=0, job_invalid_qfd=0
tgz_create_failed=0, tgz_attach_failed=0, qfd_mmap_failed=0, buf_attached=0
xfer-fas:
    ips: total=0, handled=0, accepted=0
    quar: total=0, handled=0, accepted=0
    archive: total=0, handled=0, accepted=0
    analytics: total=0, handled=0, accepted=0, local_dups=0
    analytics stats: total=0, handled=0, accepted=0
    last_rx=0, last_tx=0, error_rx=0, error_tx=0
    max_num_tasks=10000, num_tasks=0, mem_used=0, ttl_drops=0, xfer_status=0
forticloud-fsb:
    ips: total=0, handled=0, accepted=0
    quar: total=0, handled=0, accepted=0
    archive: total=0, handled=0, accepted=0
    analytics: total=0, handled=0, accepted=0, local_dups=0
num_buffer=0(per-minute:10) last_min_count=0 last_vol_count=0 next_vol_reset_tm='Sun Feb
17 00:00:00 2019
'
    analytics stats: total=24, handled=24, accepted=24
    last_rx=1224329, last_tx=1224329, error_rx=2, error_tx=0
    max_num_tasks=200, num_tasks=0, mem_used=0, ttl_drops=0, xfer_status=0
```

- Checking FortiSandbox analysis statistics:

```
FGT_FL_FULL (global) # diagnose test application quarantine 7
Total: 0

Statistics:
        vfid: 0, detected: 0, clean: 0, risk_low: 0, risk_med: 0, risk_high: 0, limit_
reached:0
        vfid: 3, detected: 0, clean: 0, risk_low: 0, risk_med: 0, risk_high: 0, limit_
reached:0
        vfid: 4, detected: 0, clean: 0, risk_low: 0, risk_med: 0, risk_high: 0, limit_
reached:0

FGT_FL_FULL (global) #
```

- Update Daemon debug:

```
FGT_FL_FULL (global) # diagnose debug application quarantined -1
FGT_FL_FULL (global) # diagnose debug enable

quar_req_fsa_file()-890: fsa ext list new_version (1547781904)
quar_fsb_handle_quar()-1439: added a req-6 to fortisandbox-fsb5, vfid=1, oftp-name=[].
__quar_start_connection()-908: start server fortisandbox-fsb5-172.18.52.154 in vdom-1
[103] __ssl_cert_ctx_load: Added cert /etc/cert/factory/root_Fortinet_Factory.cer, root
ca Fortinet_CA, idx 0 (default)
[551] ssl_ctx_create_new_ex: SSL CTX is created
[578] ssl_new: SSL object is created
upd_cfg_extract_av_db_version[378]-version=06002000AVDB00201-00066.01026-1901301530
upd_cfg_extract_ids_db_version[437]-version=06002000NIDS02403-00014.00537-1901300043
```


Netskope Exhibit 1017

```
upd_cfg_extract_ids_db_version[437]-version=06002000APDB00103-00006.00741-1512010230
upd_cfg_extract_ids_db_version[437]-version=06002000ISDB00103-00014.00537-1901300043
upd_cfg_extract_ibdb_botnet_db_version[523]-version=06002000IBDB00101-00004.00401-
1901281000
upd_cfg_extract_av_db_version[378]-version=06002000AVDB00201-00066.01026-1901301530
upd_cfg_extract_ids_db_version[437]-version=06002000NIDS02403-00014.00537-1901300043
upd_cfg_extract_ids_db_version[437]-version=06002000APDB00103-00006.00741-1512010230
upd_cfg_extract_ids_db_version[437]-version=06002000ISDB00103-00014.00537-1901300043
upd_cfg_extract_ibdb_botnet_db_version[523]-version=06002000IBDB00101-00004.00401-
1901281000
upd_cfg_extract_av_db_version[378]-version=06002000AVDB00201-00066.01026-1901301530
upd_cfg_extract_ids_db_version[437]-version=06002000NIDS02403-00014.00537-1901300043
upd_cfg_extract_ids_db_version[437]-version=06002000APDB00103-00006.00741-1512010230
upd_cfg_extract_ids_db_version[437]-version=06002000ISDB00103-00014.00537-1901300043
upd_cfg_extract_ibdb_botnet_db_version[523]-version=06002000IBDB00101-00004.00401-
1901281000
upd_cfg_extract_av_db_version[378]-version=06002000AVDB00201-00066.01026-1901301530
upd_cfg_extract_ids_db_version[437]-version=06002000NIDS02403-00014.00537-1901300043
upd_cfg_extract_ids_db_version[437]-version=06002000APDB00103-00006.00741-1512010230
upd_cfg_extract_ids_db_version[437]-version=06002000ISDB00103-00014.00537-1901300043
upd_cfg_extract_ibdb_botnet_db_version[523]-version=06002000IBDB00101-00004.00401-
1901281000
upd_cfg_extract_av_db_version[378]-version=06002000AVDB00201-00066.01026-1901301530
upd_cfg_extract_ids_db_version[437]-version=06002000NIDS02403-00014.00537-1901300043
upd_cfg_extract_ids_db_version[437]-version=06002000APDB00103-00006.00741-1512010230
upd_cfg_extract_ids_db_version[437]-version=06002000ISDB00103-00014.00537-1901300043
upd_cfg_extract_ibdb_botnet_db_version[523]-version=06002000IBDB00101-00004.00401-
1901281000
quar_remote_recv_send()-731: dev=fortisandbox-fsb2 xfer-status=0
__quar_build_pkt()-408: build req(id=337, type=4) for vdom-vdom1, len=99, oftp_name=
__quar_send()-470: dev buffer -- pos=0, len=99
quar_remote_send()-520: req(id=337, type=4) read response, dev=fortisandbox-fsb2, xfer_
status=1, buflen=12
quar_remote_recv_send()-770: dev-fortisandbox-fsb2, oevent=4, nevent=1, xfer-status=1
quar_remote_recv_send()-731: dev=fortisandbox-fsb3 xfer-status=0
__quar_build_pkt()-408: build req(id=338, type=6) for vdom-vdom1, len=93, oftp_name=
__quar_send()-470: dev buffer -- pos=0, len=93
quar_remote_send()-520: req(id=338, type=6) read response, dev=fortisandbox-fsb3, xfer_
status=1, buflen=12
quar_remote_recv_send()-770: dev-fortisandbox-fsb3, oevent=4, nevent=1, xfer-status=1
quar_remote_recv_send()-731: dev=fortisandbox-fsb5 xfer-status=0
__quar_build_pkt()-408: build req(id=340, type=6) for vdom-vdom1, len=93, oftp_name=
__quar_send()-470: dev buffer -- pos=0, len=93
quar_remote_send()-520: req(id=340, type=6) read response, dev=fortisandbox-fsb5, xfer_
status=1, buflen=12
quar_remote_recv_send()-770: dev-fortisandbox-fsb5, oevent=4, nevent=1, xfer-status=1
quar_remote_recv_send()-731: dev=fortisandbox-fsb2 xfer-status=1
quar_remote_recv()-662: dev(fortisandbox-fsb2) received a packet: len=69, type=1
quar_remote_recv()-718: file-[337] is accepted by server(fortisandbox-fsb2).
quar_put_job_req()-332: Job 337 deleted
quar_remote_recv_send()-731: dev=fortisandbox-fsb4 xfer-status=0
__quar_build_pkt()-408: build req(id=339, type=6) for vdom-vdom1, len=93, oftp_name=
__quar_send()-470: dev buffer -- pos=0, len=93
quar_remote_send()-520: req(id=339, type=6) read response, dev=fortisandbox-fsb4, xfer_
status=1, buflen=12
quar_remote_recv_send()-770: dev-fortisandbox-fsb4, oevent=4, nevent=1, xfer-status=1
```


Netskope Exhibit 1017

```
quar_remote_recv_send()-731: dev=fortisandbox-fsb1 xfer-status=0
__quar_build_pkt()-408: build req(id=336, type=4) for vdom-root, len=98, oftp_name=
__quar_send()-470: dev buffer -- pos=0, len=98
...
__get_analytics_stats()-19: Received an ANALYTICS_STATS request, vfid: 0
__quar_req_handler()-127: Request 0 was handled successfully
__get_analytics_stats()-19: Received an ANALYTICS_STATS request, vfid: 0
__quar_req_handler()-127: Request 0 was handled successfully
__get_analytics_stats()-19: Received an ANALYTICS_STATS request, vfid: 0
__quar_req_handler()-127: Request 0 was handled successfully
__get_analytics_stats()-19: Received an ANALYTICS_STATS request, vfid: 0
__quar_req_handler()-127: Request 0 was handled successfully
quar_fsb_handle_quar()-1439: added a req-6 to fortisandbox-fsb1, vfid=1, oftp-name=[].
__quar_start_connection()-908: start server fortisandbox-fsb1-172.18.52.154 in vdom-1
[103] __ssl_cert_ctx_load: Added cert /etc/cert/factory/root_Fortinet_Factory.cer, root
ca Fortinet_CA, idx 0 (default)
[551] ssl_ctx_create_new_ex: SSL CTX is created
[578] ssl_new: SSL object is created
upd_cfg_extract_av_db_version[378]-version=06002000AVDB00201-00066.01026-1901301530
upd_cfg_extract_ids_db_version[437]-version=06002000NIDS02403-00014.00537-1901300043
upd_cfg_extract_ids_db_version[437]-version=06002000APDB00103-00006.00741-1512010230
upd_cfg_extract_ids_db_version[437]-version=06002000ISDB00103-00014.00537-1901300043
upd_cfg_extract_ibdb_botnet_db_version[523]-version=06002000IBDB00101-00004.00401-
1901281000
quar_remote_recv_send()-731: dev=fortisandbox-fsb1 xfer-status=0
__quar_build_pkt()-408: build req(id=2, type=6) for vdom-vdom1, len=93, oftp_name=
__quar_send()-470: dev buffer -- pos=0, len=93
quar_remote_send()-520: req(id=2, type=6) read response, dev=fortisandbox-fsb1, xfer_
status=1, buflen=12
quar_remote_recv_send()-770: dev-fortisandbox-fsb1, oevent=4, nevent=1, xfer-status=1
quar_remote_recv_send()-731: dev=fortisandbox-fsb1 xfer-status=1
quar_remote_recv()-662: dev(fortisandbox-fsb1) received a packet: len=767, type=1
quar_store_analytics_report()-590: Analytics-report return
file=/tmp/fsb/83bb2d9928b03a68b123730399b6b9365b5cc9a5a77f8aa007a6f1a499a13b18.json.gz,
buf_sz=735
quar_store_analytics_report()-597: The request
'83bb2d9928b03a68b123730399b6b9365b5cc9a5a77f8aa007a6f1a499a13b18' score is 1
quar_remote_recv()-718: file-[2] is accepted by server(fortisandbox-fsb1).
quar_put_job_req()-332: Job 2 deleted
quar_monitor_connection_func()-978: monitoring dev fortisandbox-fsb1
quar_monitor_connection_func()-978: monitoring dev fortisandbox-fsb1
__get_analytics_stats()-19: Received an ANALYTICS_STATS request, vfid: 0
__quar_req_handler()-127: Request 0 was handled successfully
__get_analytics_stats()-19: Received an ANALYTICS_STATS request, vfid: 0
__quar_req_handler()-127: Request 0 was handled successfully
quar_monitor_connection_func()-978: monitoring dev fortisandbox-fsb1
quar_stop_connection()-1006: close connection to server(fortisandbox-fsb1)
[193] __ssl_data_ctx_free: Done
[805] ssl_free: Done
[185] __ssl_cert_ctx_free: Done
[815] ssl_ctx_free: Done
[796] ssl_disconnect: Shutdown
```

- Appliance FortiSandbox diagnostics:

```
FGT_PROXY # config global
FGT_PROXY (global) # diagnose test application quarantined 1
```


Netskope Exhibit 1017

```
Total remote&local devices: 8, any task full? 0
System have disk, vdom is enabled, mgmt=1, ha=2
xfer-fas is enabled: ips-archive dlp-archive, realtime=yes, taskfull=no
    addr=0.0.0.0/514, source-ip=0.0.0.0, keep-alive=no.
    ssl_opt=0, hmac_alg=0
    License=0, content_archive=0, arch_pause=0.

global-fas is disabled.
forticloud-fsb is disabled.
fortisandbox-fsb1 is enabled: analytics, realtime=yes, taskfull=no
    addr=172.18.52.154/514, source-ip=0.0.0.0, keep-alive=no.
    ssl_opt=3, hmac_alg=0
fortisandbox-fsb2 is enabled: analytics, realtime=yes, taskfull=no
    addr=172.18.52.154/514, source-ip=0.0.0.0, keep-alive=no.
    ssl_opt=3, hmac_alg=0
fortisandbox-fsb3 is enabled: analytics, realtime=yes, taskfull=no
    addr=172.18.52.154/514, source-ip=0.0.0.0, keep-alive=no.
    ssl_opt=3, hmac_alg=0
fortisandbox-fsb4 is enabled: analytics, realtime=yes, taskfull=no
    addr=172.18.52.154/514, source-ip=0.0.0.0, keep-alive=no.
    ssl_opt=3, hmac_alg=0
fortisandbox-fsb5 is enabled: analytics, realtime=yes, taskfull=no
    addr=172.18.52.154/514, source-ip=0.0.0.0, keep-alive=no.
    ssl_opt=3, hmac_alg=0
fortisandbox-fsb6 is enabled: analytics, realtime=yes, taskfull=no
    addr=172.18.52.154/514, source-ip=0.0.0.0, keep-alive=no.
    ssl_opt=3, hmac_alg=0
global-faz is disabled.
global-faz2 is disabled.
global-faz3 is disabled.
```

Web filter

Web filtering restricts or controls user access to web resources and can be applied to firewall policies.

In FortiOS, there are three main components of web filtering:

- Web content filter: blocks web pages containing words or patterns that you specify.
- URL filter: uses URLs and URL patterns to block or exempt web pages from specific sources, or block malicious URLs discovered by FortiSandbox.
- FortiGuard Web Filtering service: provides many additional categories you can use to filter web traffic.

These components interact with each other to provide maximum control over what users on your network can view and protect your network from many internet content threats.

Web filters are applied in the following order:

1. URL filter
2. FortiGuard Web Filtering
3. Web content filter
4. Web script filter
5. Antivirus scanning

FortiOS includes three preloaded web filter profiles:

- *default*
- *monitor-all* (monitors and logs all URLs visited, flow-based)
- *wifi-default* (default configuration for offloading WiFi traffic)

You can customize these profiles, or you can create your own to manage network user access.

Some features of this functionality require a subscription to FortiGuard Web Filtering.

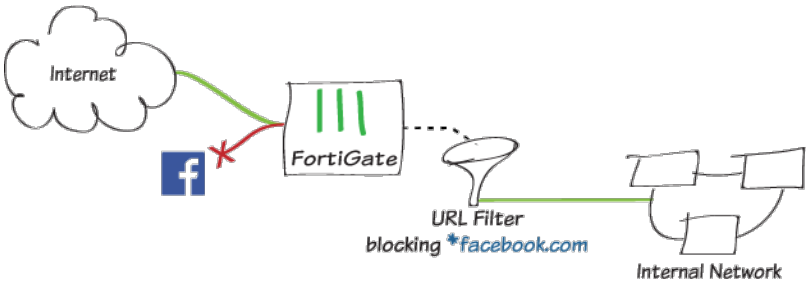
The following topics provide information about web filters:

- URL filter on page 895
- FortiGuard filter on page 901
- Credential phishing prevention on page 908
- Usage quota on page 910
- Web content filter on page 913
- File filter on page 916
- Advanced filters 1 on page 922
- Advanced filters 2 on page 926
- External resources for web filter on page 931
- Reliable web filter statistics on page 937
- Flow-based web filtering on page 940
- URL certificate blocklist on page 941

URL filter

URL filter is also called static URL filter. By adding specific URLs with patterns containing text and regular expressions, FortiGate can allow, block, exempt, and monitor web pages matching any specified URLs or patterns, and can display a replacement message instead.

Sample topology

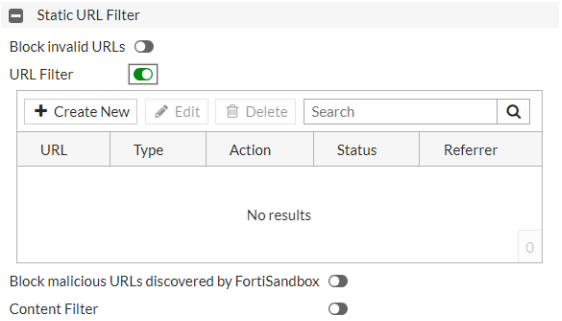


Create URL filter

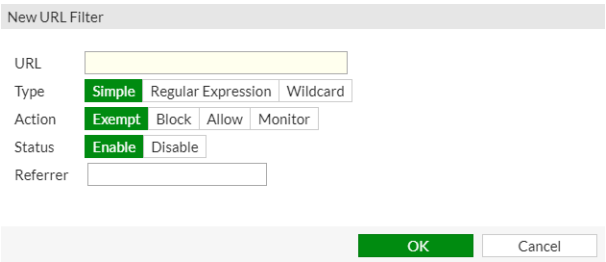
You can create a URL filter using the GUI or CLI. After creating the URL filter, attach it to a web filter profile.

To create URL filter in the GUI:

1. Go to *Security Profiles* > *Web Filter* and click *Create New*, or edit an existing profile.
2. Go the *Static URL Filter* section and enable *URL Filter*.



3. Under *URL Filter*, click *Create New* to display the *New URL Filter* pane.

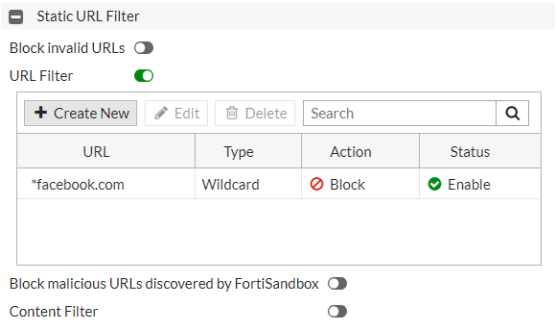


URL Filter Type	Description
Simple	FortiGate tries to strictly match the full context. For example, if you enter *www.facebook.com* in the *URL* field, it only matches traffic with *www.facebook.com*. It won't match *facebook.com* or *message.facebook.com*. When FortiGate finds a match, it performs the selected URL *Action*.
Regular Expression or Wildcard	FortiGate tries to match the pattern based on the rules of regular expressions or wildcards. For example, if you enter **fa** in the *URL* field, it matches all the content that has *fa* such as *www.facebook.com*, *message.facebook.com*, *fast.com*, etc. When FortiGate finds a match, it performs the selected URL *Action*.

For more information, see the URL Filter expressions technical note in https://kb.fortinet.com/kb/documentLink.do?externalID=FD37057.

URL Filter Action	Description
Block	Denies or blocks attempts to access any URL matching the URL pattern. FortiGate displays a replacement message.
Allow	The traffic is passed to the remaining FortiGuard web filters, web content filters, web script filters, antivirus proxy operations, and DLP proxy operations. If the URL does not appear in the URL list, the traffic is permitted.
Monitor	The traffic is processed the same way as the *Allow* action. For the *Monitor* action, a log message is generated each time a matching traffic pattern is established.
Exempt	The traffic is allowed to bypass the remaining FortiGuard web filters, web content filters, web script filters, antivirus scanning, and DLP proxy operations

4. Enter **facebook.com*, select *Wildcard*, and select *Block*.



5. Click *OK*.

After creating the URL filter, attach it to a webfilter profile.

Create URL filter using CLI

To create and enable a URL filter using the CLI, create the URL filter and then attach it to a `webfilter` profile. The CLI commands below show the full configuration of creating a URL filter.

```
config webfilter urlfilter
   edit {id}
   # Configure URL filter lists.
      set name {string}   Name of URL filter list. size[35]
      config entries
         edit {id}
         # URL filter entries.
            set url {string}   URL to be filtered. size[511]
            set type {simple | regex | wildcard}  Filter type (simple, regex, or wildcard).
                    simple    Simple URL string.
                    regex     Regular expression URL string.
                    wildcard  Wildcard URL string.
            set action {exempt | block | allow | monitor}  Action to take for URL filter
matches.
                    exempt   Exempt matches.
                    block    Block matches.
                    allow    Allow matches (no log).
                    monitor  Allow matches (with log).
            set status {enable | disable}   Enable/disable this URL filter.
            set exempt {option}   If action is set to exempt, select the security profile
operations that exempt URLs skip. Separate multiple options with a space.
                    av                   AntiVirus scanning.
                    web-content          Web Filter content matching.
                    activex-java-cookie  ActiveX, Java, and cookie filtering.
                    dlp                  DLP scanning.
                    fortiguard           FortiGuard web filter.
                    range-block          Range block feature.
                    pass                 Pass single connection from all.
                    all                  Exempt from all security profiles.
            set referrer-host {string}   Referrer host name. size[255]
         next
   next
end
```

To create URL filter to filter Facebook using the CLI:

```
config webfilter urlfilter
    edit 1
        set name "webfilter"
        config entries
            edit 1
                set url "*facebook.com"
                set type wildcard
                set action block
            next
        end
    next
end
```

To attach the URL filter to a web filter profile:

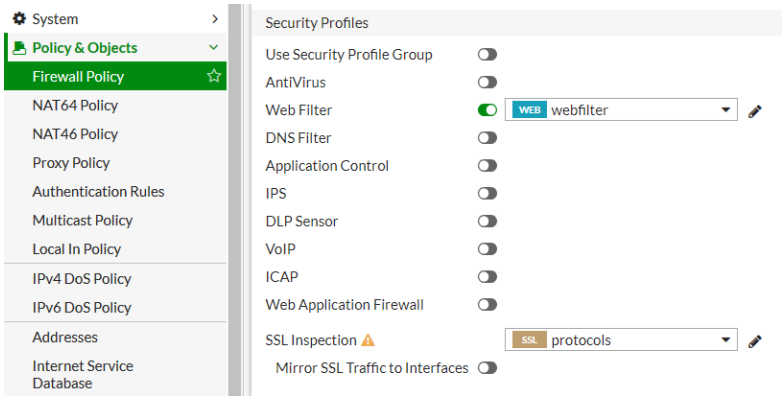
```
config webfilter profile
    edit "webfilter"                <-- the name of the webfilter profile
        config web
            set urlfilter-table 1   <-- the URL filter created with ID number 1
        end
        config ftgd-wf
            unset options
        end
    next
end
```

Attach web filter profile to the firewall policy

After you have created the URL filter and attached it to a web filter profile, you must attach the profile to a firewall policy.

To attach a web filter profile to a firewall policy using the GUI:

1. Go to *Policy & Objects* > *Firewall Policy*.
2. Edit the policy that you want to enable the webfilter.
3. In the *Security Profiles* section, enable *Web Filter* and select the profile you created.

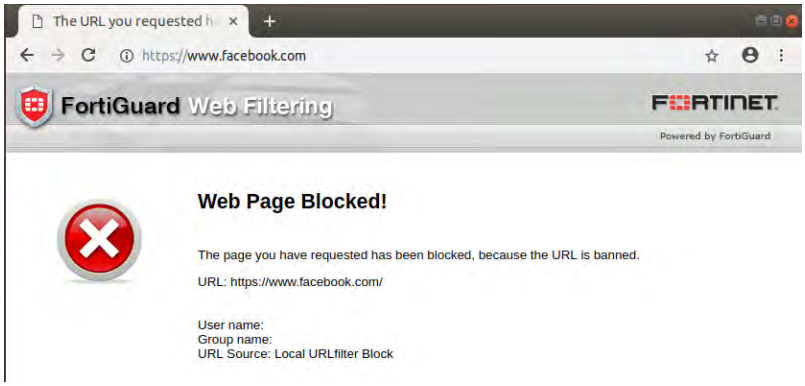


To attach a web filter profile to a firewall policy using the CLI:

```
config firewall policy
    edit 1
        set name "WF"
        set srcintf "wan2"
        set dstintf "wan1"
        set srcaddr "all"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ALL"
        set utm-status enable
        set inspection-mode proxy
        set logtraffic all
        set webfilter-profile "webfilter"    <-- attach the webfilter profile you just
created.
        set profile-protocol-options "protocol"
        set ssl-ssh-profile "protocols"
        set nat enable
    next
end
```

Validate the URL filter results

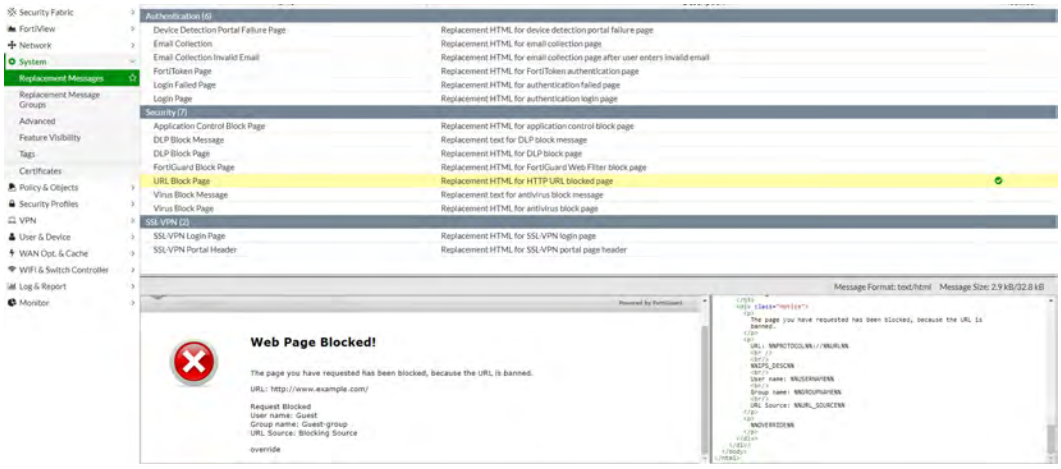
Validate the URL filter results by going to a blocked website. For example, when you go to the Facebook website, you see the replacement message.



To customize the URL web page blocked message:

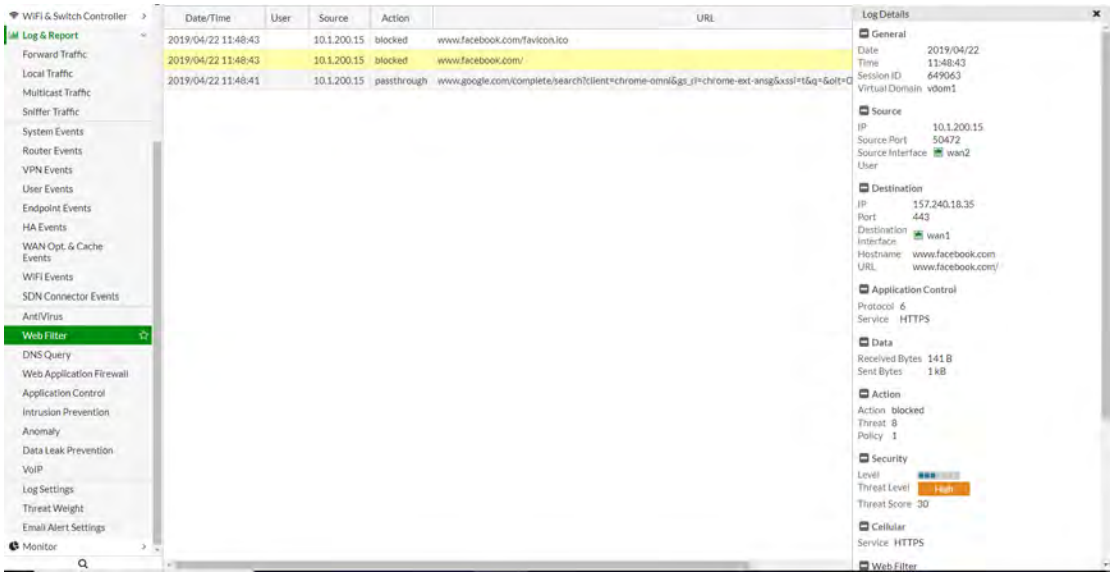
1. Go to *System* > *Replacement Messages*.
2. Go to the *Security* section and select *URL Block Page*.

3. Set up a custom message for blocked pages.



To check web filter logs in the GUI:

1. Go to *Log & Report* > *Web Filter*.



2. If there are too many log entries, click *Add Filter* and select *Event Type* > *urlfilter* to display logs generated by the URL filter.

To check web filter logs in the CLI:

```
# execute log filter category utm-webfilter

# execute log display

1: date=2019-04-22 time=11:48:43 logid="0315012544" type="utm" subtype="webfilter"
eventtype="urlfilter" level="warning" vd="vdom1" eventtime=1555958923322174610
urlfilteridx=0 urlsource="Local URLfilter Block" policyid=1 sessionid=649063
srcip=10.1.200.15 srcport=50472 srcintf="wan2" srcintfrole="wan" dstip=157.240.18.35
dstport=443 dstintf="wan1" dstintfrole="wan" proto=6 service="HTTPS"
hostname="www.facebook.com" profile="webfilter" action="blocked" reqtype="direct" url="/"
```


Netskope Exhibit 1017

```
sentbyte=1171 rcvdbyte=141 direction="outgoing" msg="URL was blocked because it is in the
URL filter list" crscore=30 craction=8 crlevel="high"
```

FortiGuard filter

To use this service, you must have a valid subscription on your FortiGate.

FortiGuard filter enhances the web filter features supplied with your FortiGate unit by sorting billions of web pages into a wide range of categories that users can allow or block.

The FortiGuard Web Filtering service includes over 45 million individual website ratings that apply to more than two billion pages. When the FortiGuard filter is enabled in a web filter and is applied to firewall policies, if a request for a web page appears in traffic controlled by one of the firewall policies, the URL is sent to the nearest FortiGuard server. The URL category or rating is returned. If the category is blocked, the FortiGate shows a replacement message in place of the requested page. If the category is not blocked, the page request is sent to the requested URL as normal.

FortiGuard web filter actions

You can select one of the following FortiGuard web filter actions:

FortiGuard Web Filter Action	Description
Allow	Permit access to the sites in the category.
Block	Prevent access to the sites in the category. Users trying to access a blocked site sees a replacement message indicating the site is blocked.
Monitor	Permits and logs access to sites in the category. You can enable user quotas when you enable this action.
Warning	Displays a message to the user allowing them to continue if they choose.
Authenticate	Requires the user to authenticate with the FortiGate before allowing access to the category or category group.

FortiGuard web filter categories

FortiGuard has many web filter categories including two local categories and a special remote category. For more information on the different categories, see the table below.

FortiGuard Web Filter category	Where to find more information
All URL categories	https://fortiguard.com/webfilter/categories.
Remote category	External resources for web filter on page 931.

The priority of categories is local category > external category > FortiGuard built-in category. If a URL is configured as a local category, it only follows the behavior of local category and not external or FortiGuard built-in category.

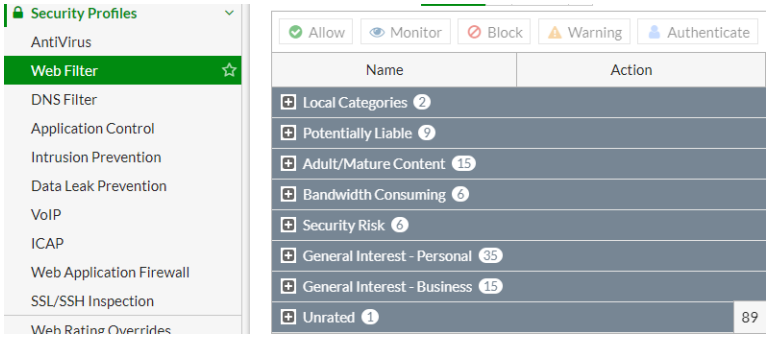
Sample configuration of blocking a web category

This example shows blocking a website based on its category (rating), for example, information technology.

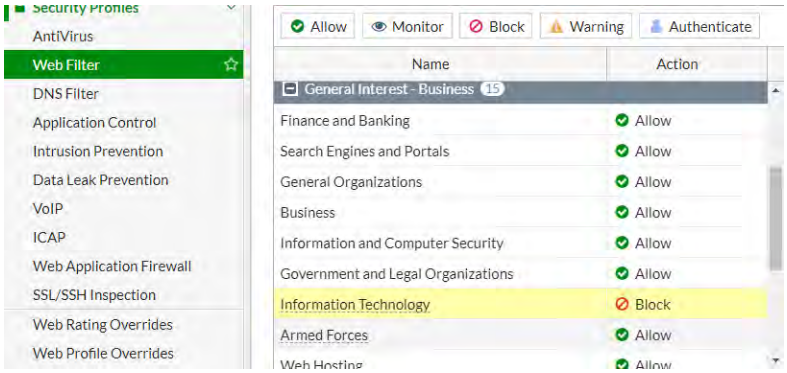
Netskope Exhibit 1017

To block a category in the GUI:

1. Go to *Security Profiles* > *Web Filter* and go to the *FortiGuard category based filter* section.



2. Open the *General Interest - Business* section by clicking the + icon beside it.
3. Select *Information Technology* and then select *Block*.

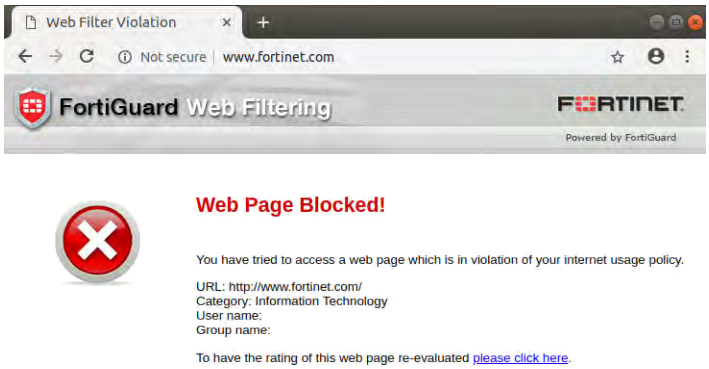


To block a category in the CLI:

```
config webfilter profile
   edit "webfilter"
      config ftgd-wf
         unset options
         config filters
            edit 1
               set category 52    <-- the pre-set id of "information technology" caterogy
               set action block   <-- set action to block
            next
         end
      end
   next
end
```

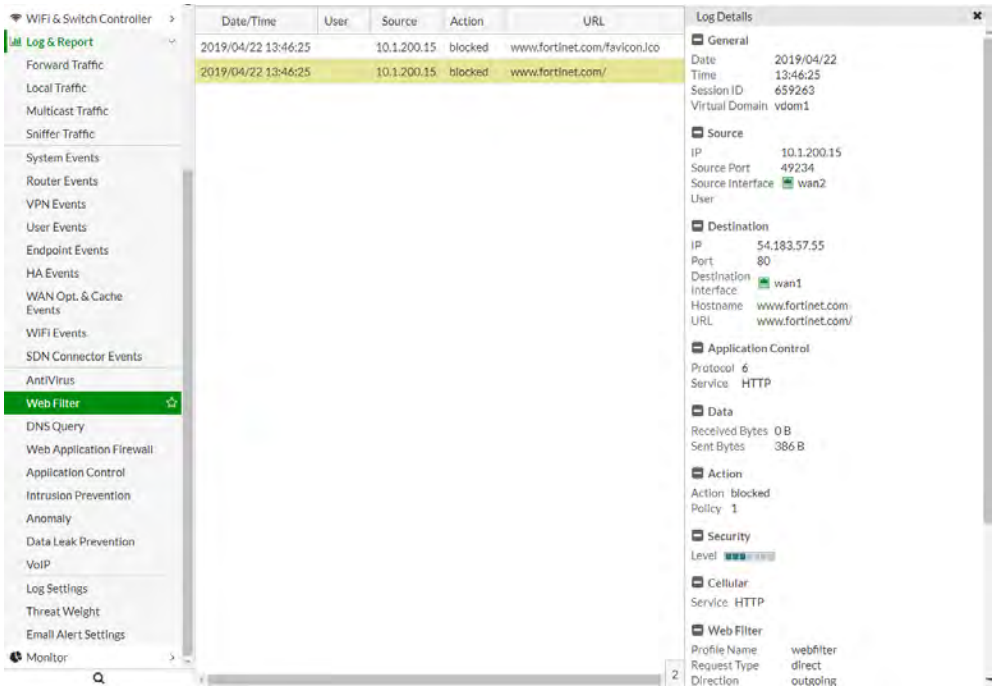
To validate that you have blocked a category:

1. Go to a website belonging to the blocked category, for example, www.fortinet.com, and you see a blocked page and the category that is blocked.



To view the log of a blocked website in the GUI:

1. Go to Log & Report > Web Filter.



To view the log of a blocked website in the CLI:

```
FGT52E-NAT-WF # execute log filter category utm-webfilter

FGT52E-NAT-WF # execute log display

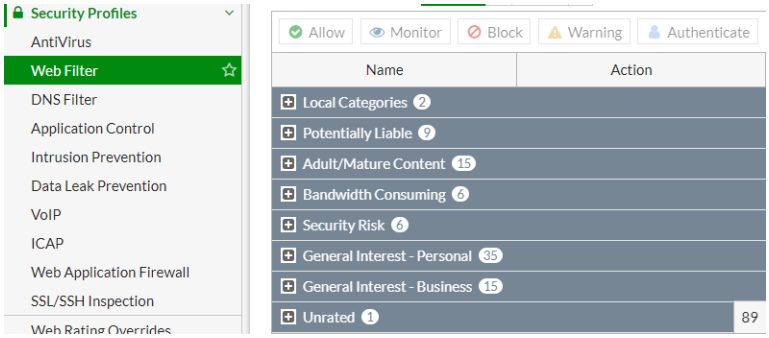
1: date=2019-04-22 time=13:46:25 logid="0316013056" type="utm" subtype="webfilter"
eventtype="ftgd_blk" level="warning" vd="vdom1" eventtime=1555965984972459609 policyid=1
sessionid=659263 srcip=10.1.200.15 srcport=49234 srcintf="wan2" srcintfrole="wan"
dstip=54.183.57.55 dstport=80 dstintf="wan1" dstintfrole="wan" proto=6 service="HTTP"
hostname="www.fortinet.com" profile="webfilter" action="blocked" reqtype="direct" url="/"
sentbyte=386 rcvdbyte=0 direction="outgoing" msg="URL belongs to a denied category in
policy" method="domain" cat=52 catdesc="Information Technology"
```

Sample configuration of issuing a warning

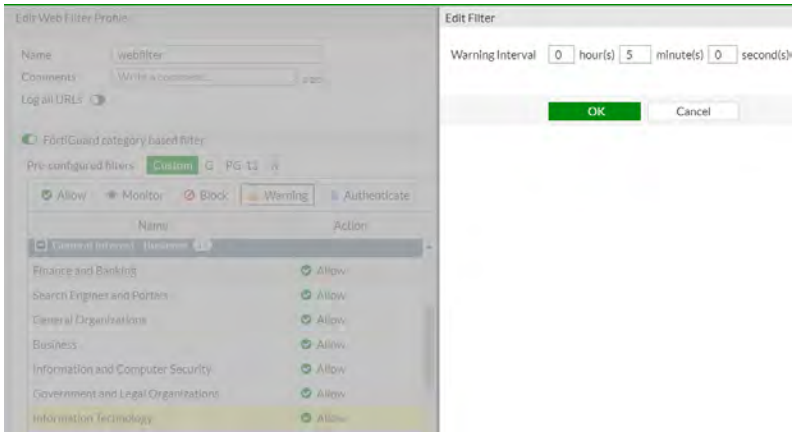
This example shows issuing a warning when a user visits a website based on its category (rating), for example, information technology.

To configure a warning in the GUI:

1. Go to *Security Profiles* > *Web Filter* and go to the *FortiGuard category based filter* section.



2. Open the *General Interest - Business* section by clicking the + icon beside it.
3. Select *Information Technology* and then select *Warning*.



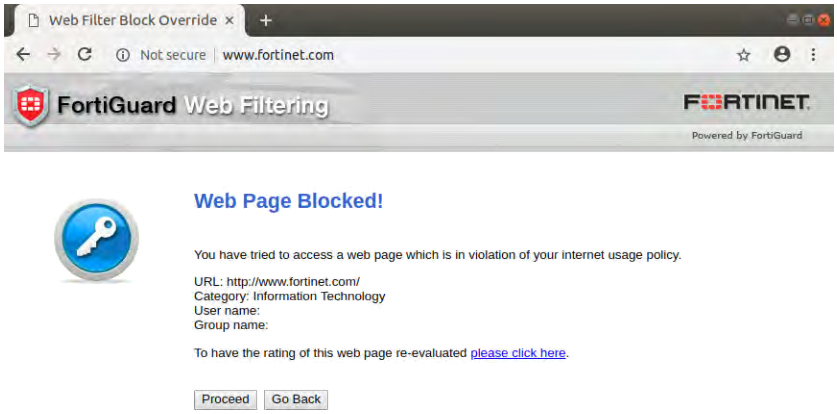
4. Set the *Warning Interval* which is the interval when the warning page appears again after the user chooses to continue.

To configure a warning in the CLI:

```
config webfilter profile
    edit "webfilter"
        config ftgd-wf
            unset options
            config filters
                edit 1
                    set category 52
                    set action warning  <-- set action to warning
                next
            end
        end
    next
end
```

To validate that you have configured the warning:

1. Go to a website belonging to the selected category, for example, www.fortinet.com, and you see a warning page where you can choose to *Proceed* or *Go Back*.

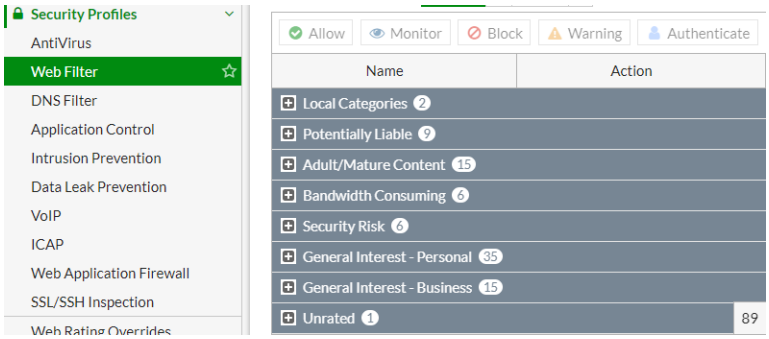


Sample configuration of authenticating a web category

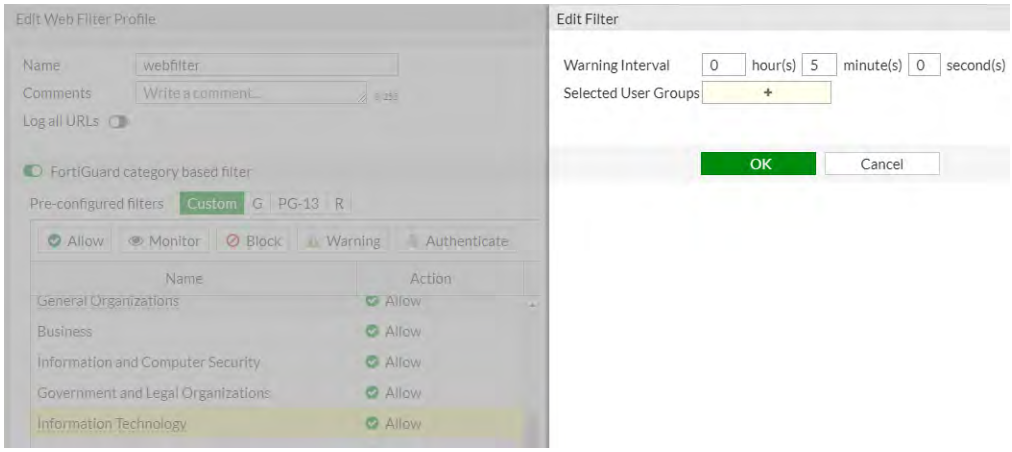
This example shows authenticating a website based on its category (rating), for example, information technology.

To authenticate a category in the GUI:

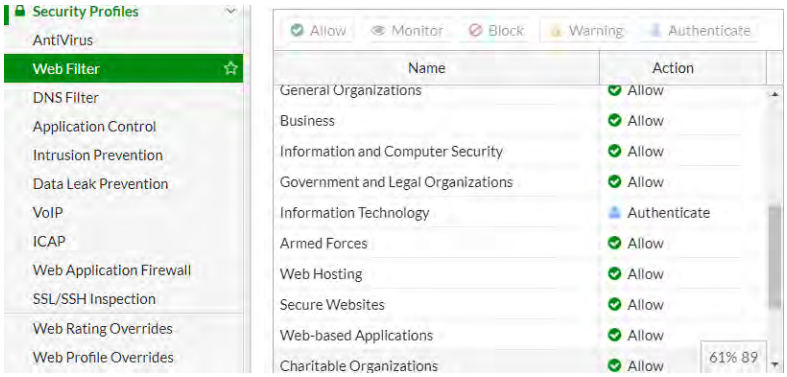
1. Go to *Security Profiles* > *Web Filter* and go to the *FortiGuard category based filter* section.



2. Open the *General Interest - Business* section by clicking the + icon beside it.
3. Select *Information Technology* and then select *Authenticate*.



4. Set the *Warning Interval* which is the interval when the authentication page appears again after authentication.
5. Click the + icon beside Selected User Group and select a user group. You must have a valid user group to use this feature.

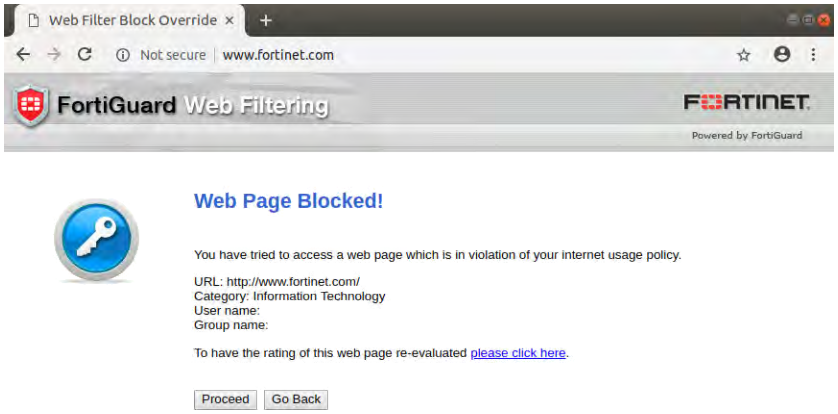


To authenticate a category in the CLI:

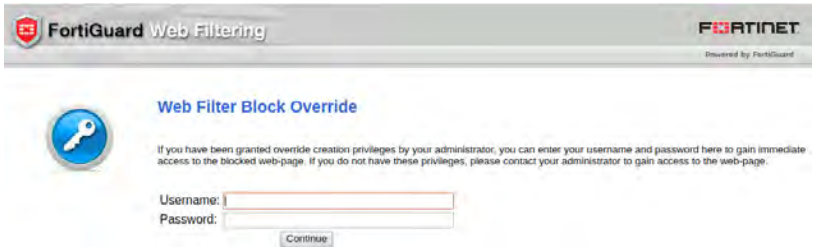
```
config webfilter profile
   edit "webfilter"
     config ftgd-wf
        unset options
        config filters
          edit 1
            set category 52
            set action authenticate          <-- set the action of authenticate
            set auth-usr-grp "local_group"   <-- user to authenticate
          next
        end
     end
   next
end
```

To validate that you have configured authentication:

1. Go to a website belonging to the selected category, for example, www.fortinet.com. First, you see a warning page where you can choose to *Proceed* or *Go Back*.



2. Click *Proceed* to check that the authentication page appears.



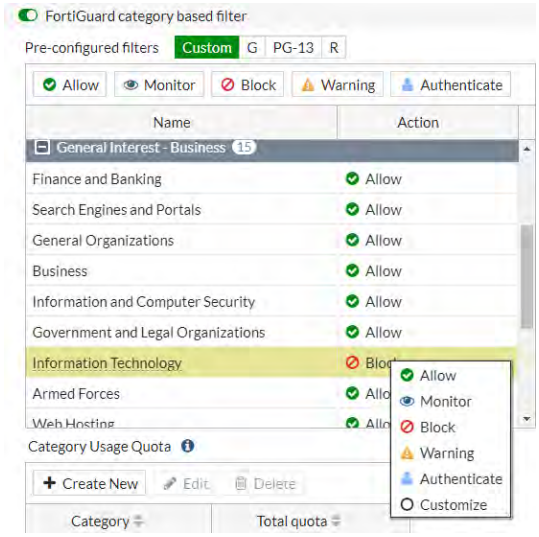
3. Enter the username and password of the user group you selected, and click *Continue*.
   If the credentials are correct, the traffic is allowed through.

Sample customization of the replacement page

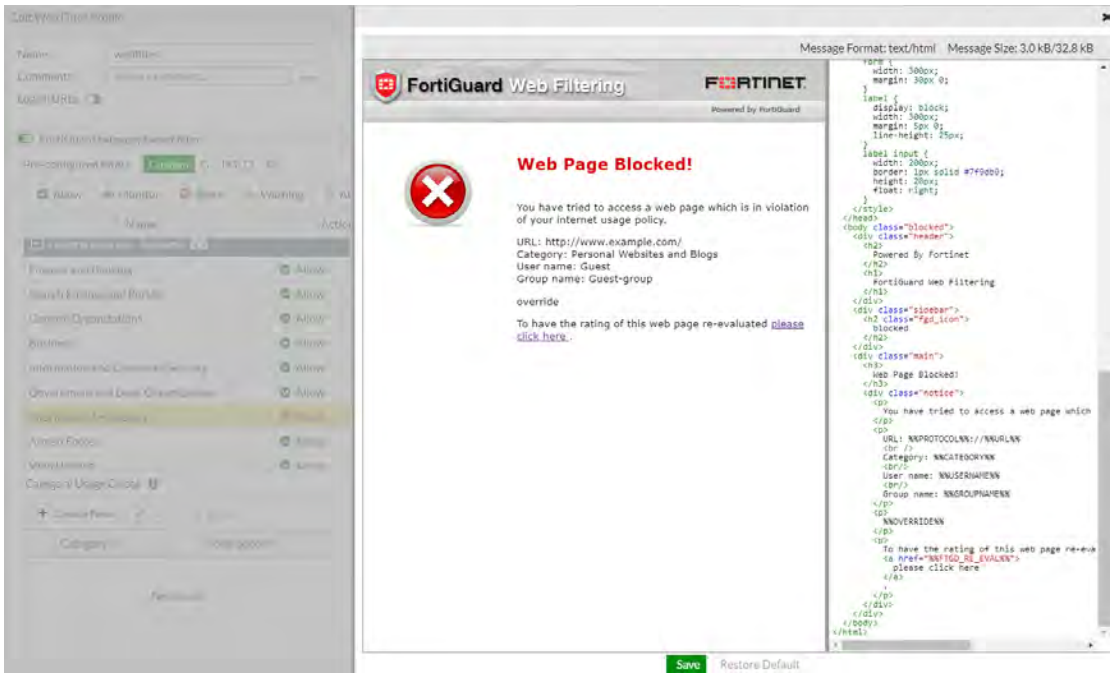
When the FortiGuard Web Filter action is *Block*, *Warning*, or *Authenticate*, there is a *Customize* option for you to customize the replace page.

To customize the replace page:

1. Go to *Security Profiles* > *Web Filter* and go to the *FortiGuard category based filter* section.
2. Right-click the item and select *Customize*.



3. A pane appears for you to customize the page.



Credential phishing prevention

When credential phishing prevention is enabled, the FortiGate scans for corporate credentials submitted to external websites and compares them to sensitive credentials stored in the corporate domain controller. Based on the configured antiphishing rules in proxy mode web filter profiles, the FortiGate will block the URL or alert the user if the credentials match ones that are stored on the corporate domain controller.

- The corporate domain controller must be configured on the credential-store. Credentials are matched based on sAMAccountName. UPN format is not currently supported.
- The antiphishing profile defines the corporate domain controller, antiphishing check option, default action if no rules match, antiphishing status, and so on.
- Inspection entries in the profile define what action occurs when the submission request matches the specified FortiGuard categories.
- The profile scans for pre-defined and custom username and password fields in the HTTP request, such as `username`, `auth`, and `password`. You can evaluate custom fields by configuring custom patterns.
- The URL filter defines individual URLs that the antiphish action (block or log) is applied to when the URL submission request matches.

Web-based URL filter actions and FortiGuard category-based filtering have higher priority than antiphishing URL filter actions and FortiGuard filtering:

- If a request is blocked by the web-based URL filter or FortiGuard filter, there is no further antiphishing scanning. Antiphishing scanning only happens after the web-based URL filtes and FortiGuard filters allow the traffic.
- If a submission matches an entry in the URL filter table that has an antiphishing action, the defined action is taken. No further FortiGuard category-based rules are applied.
- Like firewall rules, the URL filter table and Fortiguard category-based antiphishing rules use a top-down priority. The rule that matches first is the one that is used.

In this example, URLs that match FortiGuard category 37 (social networking) will be blocked and other categories will be logged.

To configure credential phishing prevention:

1. Configure the corporate domain controller:

```
config credential-store domain-controller
    edit "win2016"
        set domain-name "corpserver.local"
        set username "Administrator"
        set password **********
        set ip <server_ip>
    next
end
```



The domain controller entry name must be the hostname of the DC (`win2016` in the example). Both it and the domain name are case sensitive.

2. Configure the antiphishing profile, which includes the FortiGuard category rule:

```
config webfilter profile
    edit "<profile-name>"
        set feature-set proxy
        ...
        config web
            ...
        end
        config antiphish
            set status enable
            set domain-controller "win2016"
            set default-action block
            set check-uri enable
            set check-basic-auth enable
            set max-body-len 65536
            config inspection-entries
                edit "inspect-37"
                    set fortiguard-category 37
                    set action block
                next
                edit "inspect-others"
                    set fortiguard-category all
                    set action log
                next
            end
            config custom-patterns
                edit "customer-name"
                    set category username
                next
                edit "customer-passwd"
                   set category password
                next
            end
        end
```


Netskope Exhibit 1017

```
        ...
        set web-antiphishing-log enable
    next
end
```

- `check-uri` enables support for scanning HTTP GET URI parameters.
- `check-basic-auth` enables support for scanning the HTTP Basic Auth field.

3. Configure the URL filter to scan specific URLs.
   The antiphish action is added to the URL filter table entry, and the URL filter is applied to the webfilter profile.

```
config webfilter urlfilter
    edit 1
        set name "antiphish-table"
        config entries
            edit 1
                set url "www.example.com"
                set type simple
                set antiphish-action block
                set status enable
                set referrer-host ''
            next
        end
    next
end
config webfilter profile
    edit "<profile-name>"
        config web
            set urlfilter-table 1
        end
        ...
    next
end
```

4. Optionally, define custom patterns to scan fields other than the built-in username and password keywords are needed:

```
config webfilter profile
    edit "<profile-name>"
        config custom-patterns
                edit "customer-name"
                    set category username
                next
                edit "customer-passwd"
                   set category password
                next
            end
        end
    next
end
```

Usage quota

In addition to using category and classification blocks and overrides to limit user access to URLs, you can set a daily quota by category, category group, or classification. Quotas allow access for a specified length of time or a specific bandwidth, and is calculated separately for each user. Quotas are reset everyday at midnight.

Netskope Exhibit 1017

Quotas can be set only for the actions of *Monitor*, *Warning*, or *Authenticate*. When the quota is reached, the traffic is blocked and the replacement page displays.

You can only use quotas when inspection mode is *Proxy*.

Sample topology

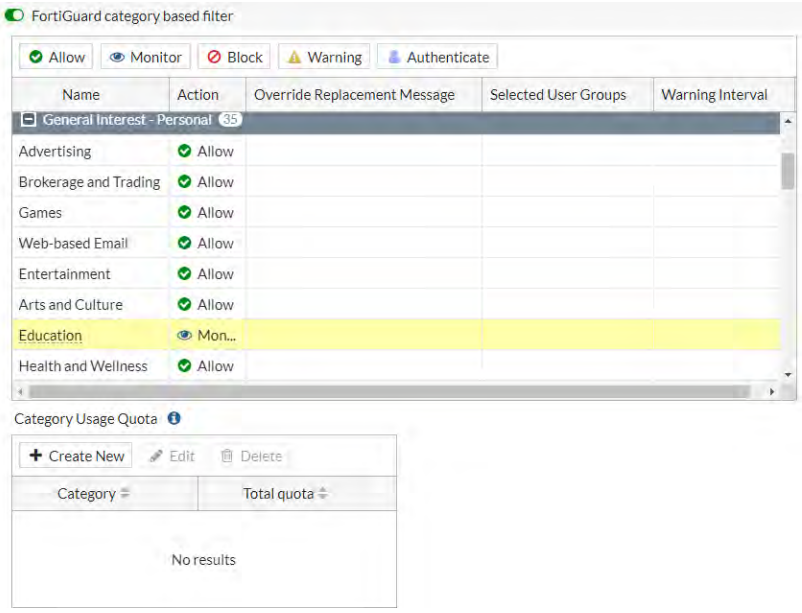


Sample configuration of setting a quota

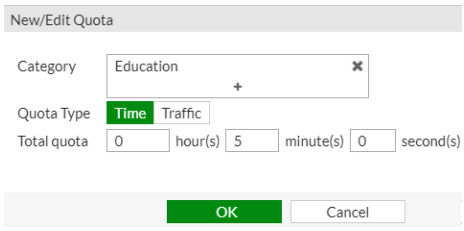
This example shows setting a time quota for a category, for example, the Education category.

To configure a quota in the GUI:

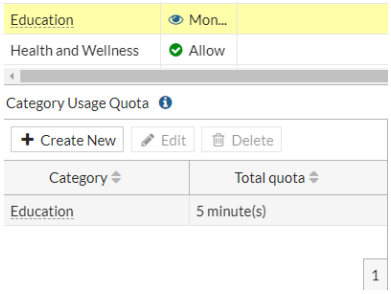
1. Go to *Security Profiles* > *Web Filter* and go to the *FortiGuard category based filter* section.
2. Open the *General Interest - Personal* section by selecting the + icon beside it.
3. Select *Education* and then select *Monitor*.
4. In the *Category Usage Quota* section, select *Create New*.



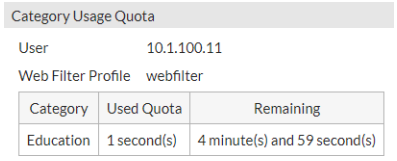
5. In the right pane, select the *Category* field and then select *Education*.
6. For the *Quota Type*, select *Time* and set the *Total quota* to 5 *minute(s)*.



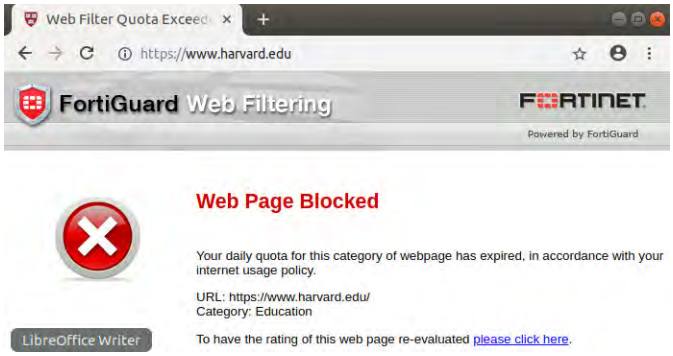
7. Select *OK* and the Category Usage Quota section displays the quota.



8. Validate the configuration by visiting a website in the education category, for example https://www.harvard.edu/. You can view websites in the education category.
9. Check the used and remaining quota in *Monitor* > *FortiGuard Quota*.



10. When the quota reaches its limit, traffic is blocked and the replacement page displays.



To configure a quota in the CLI:

```
config webfilter profile
    edit "webfilter"
        config ftgd-wf
           unset options
           config filters
              edit 1
                 set category 30 <-- the id of education category
              next
           end
```

```
            config quota
               edit 1
                  set category 30
                  set type time
                  set duration 5m
               next
            end
         end
      next
end
```

Web content filter

You can control access to web content by blocking web pages containing specific words or patterns. This helps to prevent access to pages with questionable material. You can specify words, phrases, patterns, wildcards and Perl regular expressions to match content on web pages. You can use multiple web content filter lists and select the best web content filter list for each Web Filter profile.

Pattern type

When you have created the Web Filter content list, you need to add web content patterns to it. There are two types of patterns: wildcard and regular expression.

Wildcard

Use the wildcard setting to block or exempt one word or text strings of up to 80 characters. You can also use wildcard symbols such as **?** or ***** to represent one or more characters. For example, a wildcard expression *forti*.com* matches *fortinet.com* and *forticare.com*. The ***** represents any character appearing any number of times.

Regular expression

Use the regular expression setting to block or exempt patterns of Perl expressions which use some of the same symbols as wildcard expressions but for different purposes. In regular expressions, ***** represents the character before the symbol. For example, *forti*.com* matches *fortiii.com* but not *fortinet.com* or *fortiice.com*. In this case, the symbol ***** represents **i** appearing any number of times.

The maximum number of web content patterns in a list is 5000.

Content evaluation

The web content filter feature scans the content of every web page that is accepted by a security policy. The system administrator can specify banned words and phrases and attach a numerical value, or score, to the importance of those words and phrases. When the web content filter scan detects banned content, it adds the scores of banned words and phrases found on that page. If the sum is higher than a threshold set in the Web Filter profile, FortiGate blocks the page.

The default score for web content filter is 10 and the default threshold is 10. This means that by default, a web page is blocked by a single match.

Banned words or phrases are evaluated according to the following rules:

- The score for each word or phrase is counted only once, even if that word or phrase appears many times in the web page.

Netskope Exhibit 1017

- The score for any word in a phrase without quotation marks is counted.
- The score for a phrase in quotation marks is counted only if it appears exactly as written.

Sample of applying banned pattern rules

The following table is an example of how rules are applied to the contents of a web page. For example, a web page contains only this sentence:

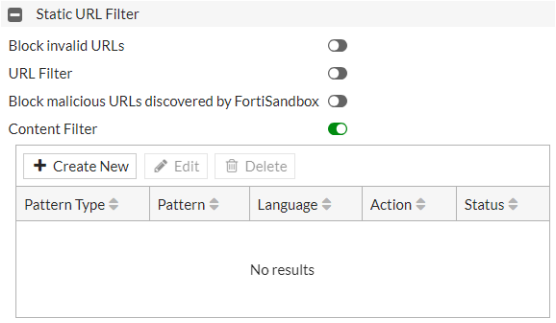
The score for each word or phrase is counted only once, even if that word or phrase appears many times in the web page.

Banned pattern	Assigned score	Score added to the sum for the entire page	Threshold score	Comment
word	20	20	20	Appears twice but only counted once. Web page is blocked.
word phrase	20	40	20	Each word appears twice but only counted once giving a total score of 40. Web page is blocked.
word sentence	20	20	20	“word” appears twice, “sentence” does not appear, but since any word in a phrase without quotation marks is counted, the score for this pattern is 20. Web page is blocked.
"word sentence"	20	0	20	This phrase does not appear exactly as written. Web page is allowed.
"word or phrase"	20	20	20	This phrase appears twice but is counted only once. Web page is blocked.

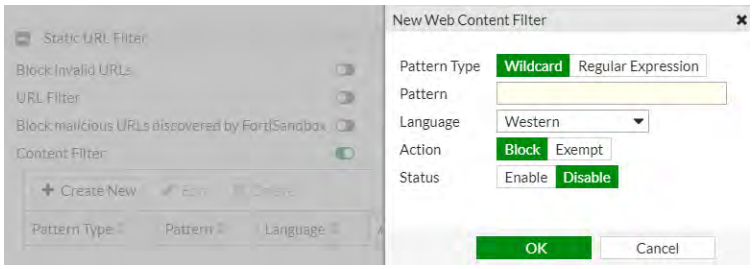
Sample configuration

To configure web content filter in the GUI:

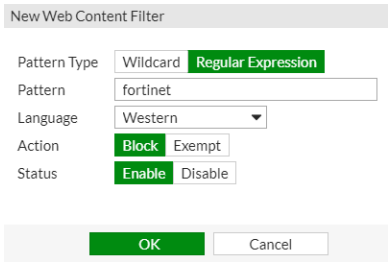
1. Go to *Security Profiles* > *Web Filter* and go to the *Static URL Filter* section.
2. Enable *Content Filter* to display its options.



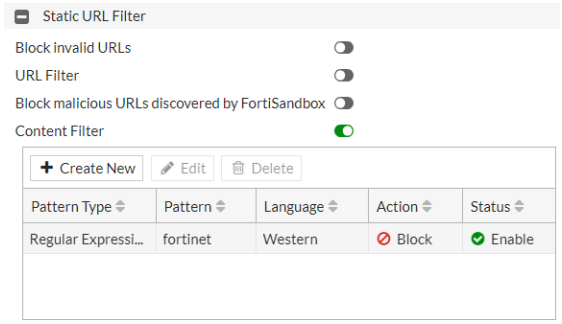
3. Select *Create New* to display the content filter options.



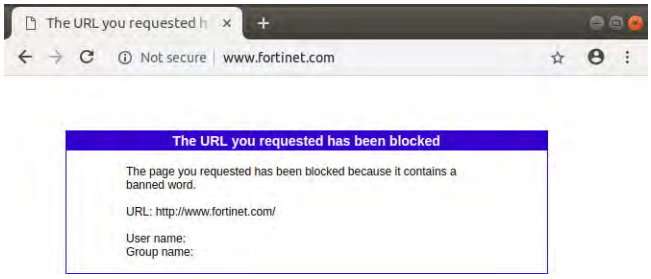
4. For *Pattern Type*, select *Regular Expression* and enter *fortinet* in the *Pattern* field.
   - Leave *Language* as *Western*.
   - Set *Action* to *Block*.
   - Set *Status* to *Enable*.



5. Select *OK* to see the updated *Static URL Filter* section.



6. Validate the configuration by visiting a website with the word *fortinet*, for example, www.fortinet.com. The website is blocked and a replacement page displays.



To configure web content filter in the CLI:

1. Create a content table:

```
config webfilter content
   edit 1                          <-- the id of this content
      set name "webfilter"
```

```
        config entries
           edit "fortinet"            <-- the banned word
             set pattern-type regexp  <-- the type is regular expression
             set status enable
             set lang western
             set score 10             <-- the score for this word is 10
             set action block
           next
        end
     next
end
```

2. Attach the content table to the Web Filter profile:

```
config webfilter profile
   edit "webfilter"
      config web
         set bword-threshold 10  <-- the threshold is 10
         set bword-table 1       <-- the id of content table we created in the previous
step
      end
      config ftgd-wf
         unset options
      end
   next
end
```

File filter

File Filter allows the Web Filter profile to block files passing through a FortiGate based on file type.

HTTP and FTP File Filtering is configurable in Web Filter profile.

File Filtering in Web Filter profile is based on file type (file's meta data) only, and not on file size or file content. You need to configure a DLP sensor to block files based on size or content such as SSN numbers, credit card numbers, or regexp.

File filtering only works on proxy mode policies.

Supported file types

The following file types are supported in File Filter and DLP profiles:

File Type Name	Description
.net	Match .NET files
7z	Match 7-zip files
activemime	Match activemime files
arj	Match arj compressed files
aspack	Match aspack files
avi	Match avi files

File Type Name	Description
base64	Match base64 files
bat	Match Windows batch files
bin	Match bin files
binhex	Match binhex files
bmp	Match bmp files
bzip	Match bzip files
bzip2	Match bzip2 files
cab	Match Windows cab files
chm	Match Windows compiled HTML help files
class	Match class files
cod	Match cod files
crx	Match Chrome extension files
dmg	Match Apple disk image files
elf	Match elf files
exe	Match Windows executable files
flac	Match FLAC files
fsg	Match fsg files
gif	Match gif files
gzip	Match gzip files
hlp	Match Windows help files
hta	Match hta files
html	Match html files
iso	Match ISO archive files
jad	Match jad files
javascript	Match javascript files
jpeg	Match jpeg files
lzh	Match lzh compressed files
mach-o	Match Mach object files
mime	Match mime files
mov	Match mov files

File Type Name	Description
mp3	Match mp3 files
mpeg	Match mpeg files
msc	Match msc files
msi	Match Windows Installer msi bzip files
msoffice	Match MS-Office files. For example, doc, xls, ppt, and so on.
msofficex	Match MS-Office XML files. For example, docx, xlsx, pptx, and so on.
pdf	Match pdf files
petite	Match petite files
png	Match png files
prc	Match prc files
rar	Match rar archives
rm	Match rm files
sis	Match sis files
tar	Match tar files
tiff	Match tiff files
torrent	Match torrent files
unknown*	Match unknown files
upx	Match upx files
uue	Match uue files
wav	Match wav files
wma	Match wma files
xar	Match xar archive files
xz	Match xz files
zip	Match zip files

* This file type is only available in DLP profiles.

Example

In the following example, three file filters are used in the Web Filter profile:

1. Block PDFs from entering our leaving the network (filter1).
2. Log the download of some graphics file-types via HTTP (filter2).
3. Block executable files from leaving to the network over FTP (filter3).

To configure a file-type based web filter in the CLI:

```
config webfilter profile
    edit "webfilter-file-filter"
        config file-filter
            set status enable
            set log enable
            set scan-archive-contents enable
            config entries
                edit "filter1"
                    set comment "Block PDF files"
                    set protocol http ftp
                    set action block
                    set direction any
                    set encryption any
                    set file-type "pdf"
                next
                edit "filter2"
                    set comment "Log graphics files"
                    set protocol http
                    set action log
                    set direction incoming
                    set encryption any
                    set file-type "jpeg" "png" "gif"
                next
                edit "filter3"
                    set comment "Block upload of EXE files"
                    set protocol ftp
                    set action block
                    set direction outgoing
                    set encryption any
                    set file-type "exe"
                next
            end
        end
    next
end
```

After configuring file filters in Web Filter profile, apply it to a firewall policy using the following command:

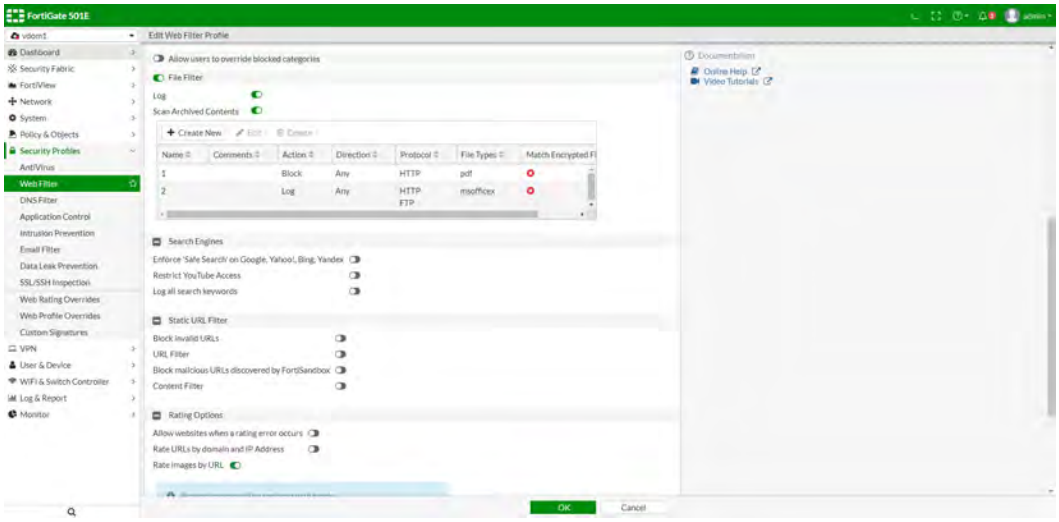
```
config firewall policy
    edit 1
        set name "client-to-internet"
        set srcintf "dmz"
        set dstintf "wan1"
        set srcaddr "all"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ALL"
        set utm-status enable
        set utm-inspection-mode proxy
        set logtraffic all
        set webfilter profile "webfilter-filefilter"
        set profile-protocol-options "protocol"
        set ssl-ssh-profile "protocols"
        set nat enable
```


Netskope Exhibit 1017

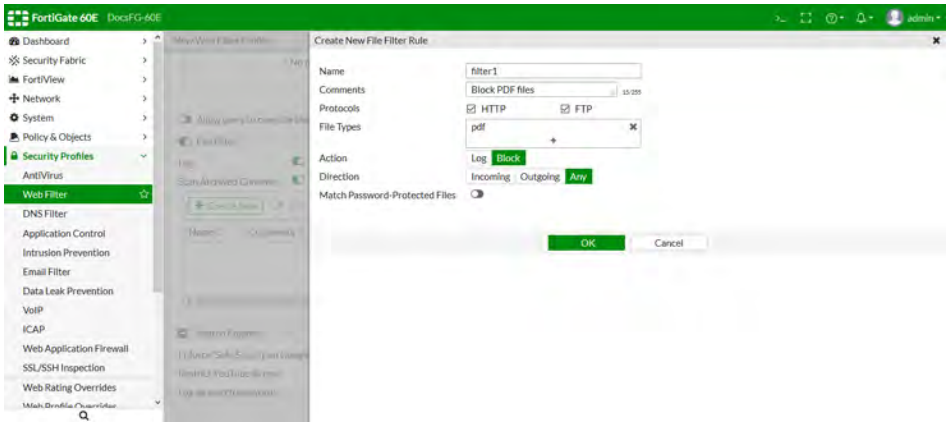
```
    next
end
```

To configure a file-type based web filter in the GUI:

1. Go to *Security Profiles* > *Web Filter*.
2. Click *Create New* or select an existing profile and click *Edit*.

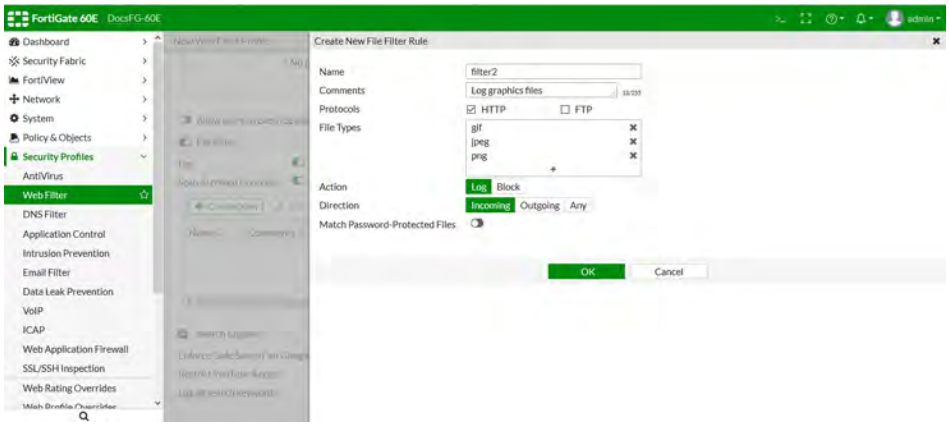


3. Enable *File Filter*.
4. Enable *Log* and *Scan Archived Contents*.
5. In the *File Filter* table, click *Create New*.
6. Configure the filters:
   a. *filter1* blocks PDFs from entering our leaving the network .

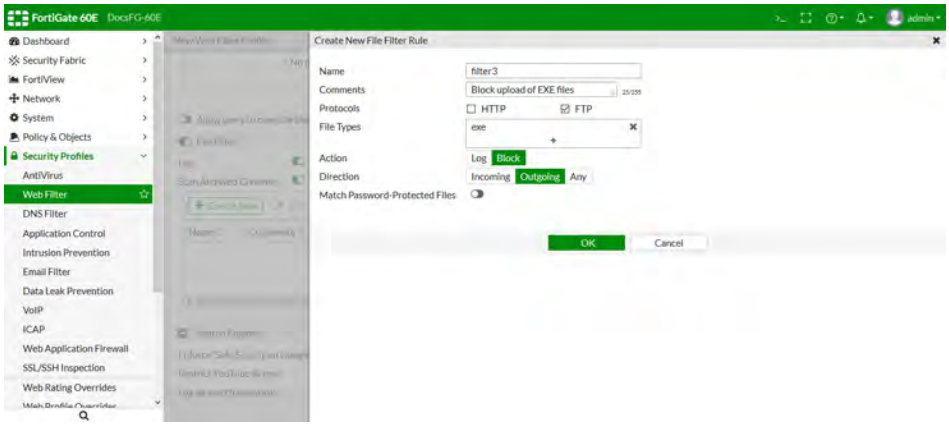


Netskope Exhibit 1017

b. *filter2* logs the download of some graphics file-types via HTTP .



c. *filter3* blocks EXE files from leaving to the network over FTP .



7. Click *OK*.
8. Add the new web filter profile to a firewall policy.

Log examples

File filter block action:

```
1: date=2019-03-19 time=09:42:15 logid="0346012673" type="utm" subtype="webfilter"
eventtype="file_filter" level="warning" vd="vd1" eventtime=1548438135 policyid=1
sessionid=29449 srcip=10.1.100.22 srcport=52816 srcintf="dmz" srcintfrole="undefined"
dstip=172.16.200.55 dstport=80 dstintf="wan1" dstintfrole="undefined" proto=6 service="HTTP"
hostname="172.16.200.55" profile="webfilter-filefilter" action="blocked" reqtype="direct"
url="/app_data/test1.pdf" sentbyte=0 rcvdbyte=0 direction="incoming" filename="test1.pdf"
filtername="filter1" filetype="pdf" msg="File was blocked by file filter."
```

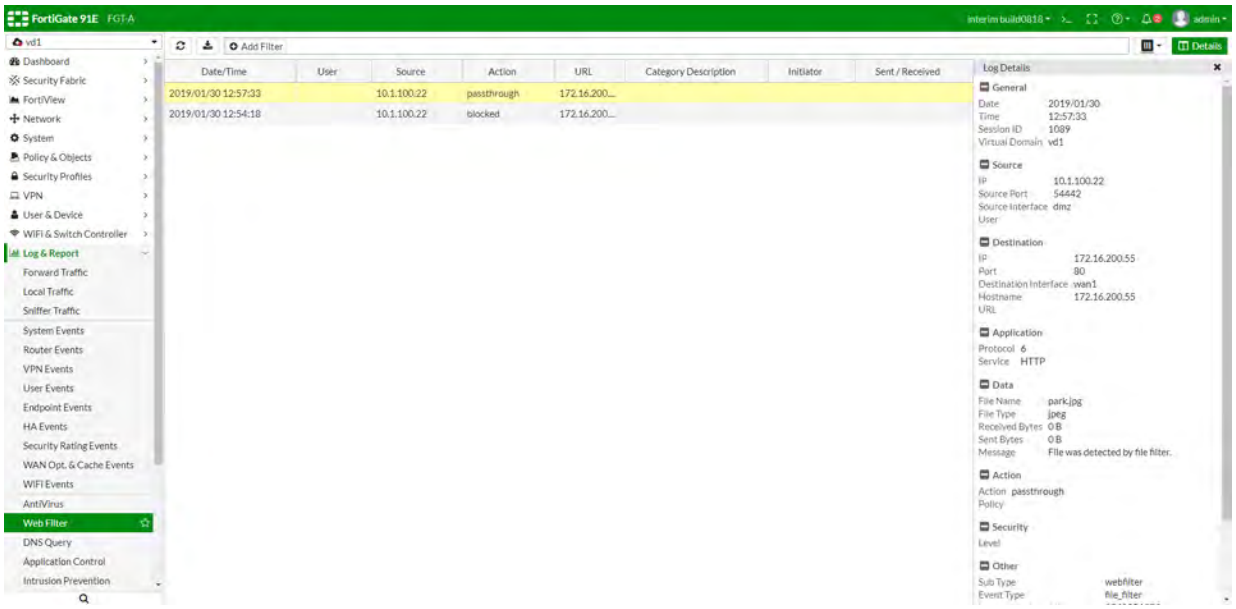
File filter log action:

```
2: date=2019-03-19 time=10:48:23 logid="0346012672" type="utm" subtype="webfilter"
eventtype="file_filter" level="notice" vd="vd1" eventtime=1548442102 policyid=1
sessionid=521 srcip=10.1.100.22 srcport=52894 srcintf="dmz" srcintfrole="undefined"
dstip=172.16.200.55 dstport=80 dstintf="wan1" dstintfrole="undefined" proto=6 service="HTTP"
hostname="172.16.200.55" profile="webfilter-filefilter" action="passthrough"
reqtype="direct" url="/app_data/park.jpg" sentbyte=0 rcvdbyte=0 direction="incoming"
filename="park.jpg" filtername="filter2" filetype="jpeg" msg="File was detected by file
filter."
```


Netskope Exhibit 1017

In the GUI:

Go to *VDOM* > *Log & Report* > *Web Filter*:



Advanced filters 1

This topic gives examples of the following advanced filter features:

- Block malicious URLs discovered by FortiSandbox on page 922
- Allow websites when a rating error occurs on page 923
- Rate URLs by domain and IP address on page 923
- Block invalid URLs on page 924
- Rate images by URL on page 925

Block malicious URLs discovered by FortiSandbox

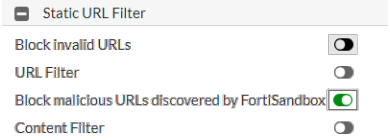
To use this feature, you must be registered to a FortiSandbox and be connected to it.

This feature blocks malicious URLs that FortiSandbox finds.

For information on configuring FortiSandbox, see Using FortiSandbox Cloud with antivirus on page 886.

To enable this feature in the GUI:

1. Go to *Security Profiles* > *Web Filter* and click *Create New*, or edit an existing profile.
2. Go the *Static URL Filter* section and enable *Block malicious URLs discovered by FortiSandbox*.



To enable this feature in the CLI:

```
config webfilter profile
    edit "webfilter"
        config web
            set blacklist enable
        end
    next
end
```

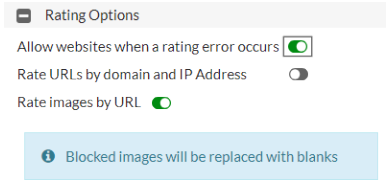
Allow websites when a rating error occurs

If you don't have a FortiGuard license but you have enabled services that need a FortiGuard license, such as FortiGuard filter, then you'll get a rating error message.

Use this setting to allow access to websites that return a rating error from the FortiGuard Web Filter service.

To enable this feature in the GUI:

1. Go to *Security Profiles* > *Web Filter* and go to the *Rating Options* section.
2. Enable *Allow websites when a rating error occurs*.



To enable this feature in the CLI:

```
config webfilter profile
    edit "webfilter"
        config ftgd-wf
            set options error-allow
        end
    next
end
```

Rate URLs by domain and IP address

If you enable this feature, in addition to only sending domain information to FortiGuard for rating, FortiGate always sends both the URL domain name and the TCP/IP packet's IP address (except for private IP addresses) to FortiGuard for the rating.

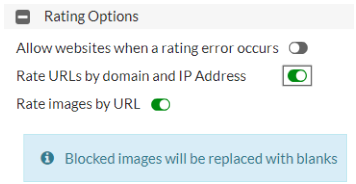
FortiGuard server might return a different category of IP address and URL domain. If they are different, FortiGate uses the rating weight of the IP address or domain name to determine the rating result and decision. This rating weight is hard-coded in FortiGate.

For example, if we use a spoof IP of Google as www.irs.gov, FortiGate will send both the IP address and domain name to FortiGuard to get the rating. In this example, we get two different ratings, one is search engine and portals which belongs to the IP of Google, another is government and legal organizations which belongs to www.irs.gov. As the search engine

and portals has a higher weight than government and legal organizations, this traffic will be rated as search engine and portals and not rated as government and legal organizations.

To enable this feature in the GUI:

1. Go to *Security Profiles* > *Web Filter* and go to the *Rating Options* section.
2. Enable *Rate URLs by domain and IP address*.



To enable this feature in the CLI:

```
config webfilter profile
    edit "webfilter"
        config ftgd-wf
            set options rate-server-ip
        end
    next
end
```

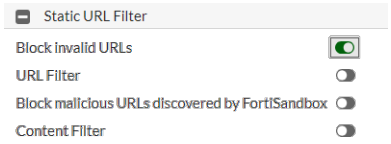
Block invalid URLs

Use this feature to block websites when their SSL certificate CN field does not contain a valid domain name.

For example, this option blocks URLs which contains spaces. If there is a space in the URL, it must be written as: *http://www.example.com/space%20here.html*.

To enable this feature in the GUI:

1. Go to *Security Profiles* > *Web Filter* and go to the *Static URL Filter* section.
2. Enable *Block invalid URLs* .



To enable this feature in the CLI:

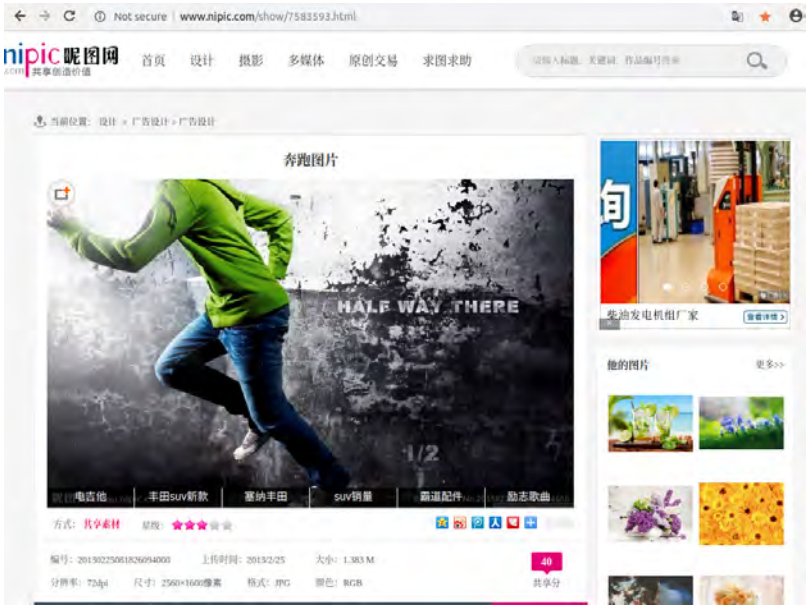
```
config webfilter profile
    edit "webfilter"
        set options block-invalid-url
    next
end
```

Rate images by URL

This feature enable FortiGate to retrieve ratings for individual images in addition to websites. Images in a blocked category are not displayed even if they are part of a site in an allowed category. Blocked images are replaced with blank placeholders. These image file types are rated: GIF, JPEG, PNG, BMP, and TIFF.

This feature requires a valid FortiGuard license, otherwise rating errors will occur. By default, this feature is enabled.

For example, if the Other Adult Materials category is blocked, before enabling *Rate images by URL*, the image is not blocked:

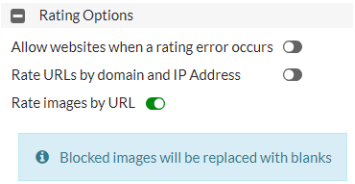


After enabling *Rate images by URL*, images in the Other Adult Materials category are blocked. For example:



To enable this feature in the GUI:

1. Go to *Security Profiles* > *Web Filter* and go to the *Rating Options* section.
2. Enable *Rate images by URL*.



To enable this feature in the CLI:

```
config webfilter profile
    edit "webfilter"
        config ftgd-wf
            unset options
            set rate-image-urls enable
        end
    next
end
```

Advanced filters 2

This topic gives examples of the following advanced filter features:

- Safe search on page 926
- YouTube education filters on page 927
  - Restrict YouTube access on page 927
  - YouTube channel filtering on page 928
- Log all search keywords on page 929
- Restrict Google account usage to specific domains on page 929
- HTTP POST Action on page 930
- Remove Java applets, remove ActiveX, and remove cookies on page 931

These advanced filters are only available when inspection mode is *Proxy*.

Safe search

This feature applies to popular search sites and prevents explicit websites and images from appearing in search results.

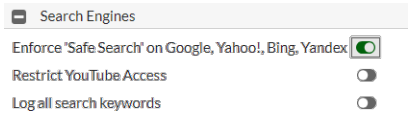
Supported search sites are:

- Google
- Yahoo
- Bing
- Yandex

Netskope Exhibit 1017

To enable this feature in the GUI:

1. Go to *Security Profiles* > *Web Filter* and go to the *Search Engines* section.
2. Enable *Enforce 'Safe Search' on Google, Yahoo!, Bing, Yandex*.



To enable this feature in the CLI:

```
config webfilter profile
    edit "webfilter"
        config web
            set safe-search url
        end
    next
end
```

YouTube education filters

Use these features to limit users' access to YouTube channels, such as in an education environment where you want students and users to be able to access YouTube education videos but not other YouTube videos.

Restrict YouTube access

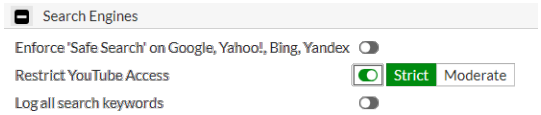
Formerly, YouTube for Schools was a way to access educational videos inside a school network. This YouTube feature lets schools access educational videos on YouTube EDU and to specify the videos accessible within the school network.

When Google stopped supporting YouTube for Schools on July 1, 2016, YouTube safe search also stopped working.

Google provides information on restricting YouTube content such as Restrict YouTube content available to G Suite users. At this time, the options Google offers to restrict inappropriate content includes: DNS, HTTP headers, and Chromebooks..

To enable this feature in the GUI:

1. Go to *Security Profiles* > *Web Filter* and go to the *Search Engines* section.
2. Enable *Restrict YouTube Access* and select *Strict* or *Moderate*.



To enable this feature in the CLI:

```
config webfilter profile
    edit "webfilter"
        config web
            set youtube-restrict strict
        end
    next
end
```

YouTube channel filtering

This Web Filter feature is also called *Restrict YouTube access to specific channels*. Use this feature to block or only allow matching YouTube channels.

The following identifiers are used:

```
given <channel-id>, affect on:
   www.youtube.com/channel/<channel-id>
   www.youtube.com/user/<user-id>
```

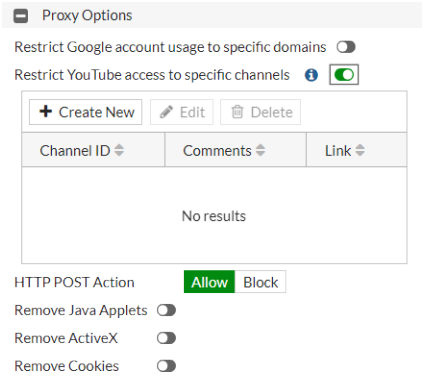
matches channel-id from <meta itemprop="channelId" content="UCGzuiiLdQZu9wxDNJHO_JnA">

```
   www.youtube.com/watch?v=<string>
```

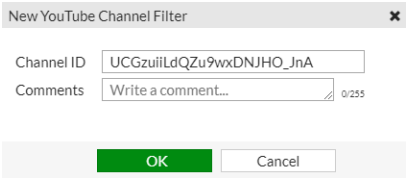
matches channel-id from <meta itemprop="channelId" content="UCGzuiiLdQZu9wxDNJHO_JnA">

To enable this feature in the GUI:

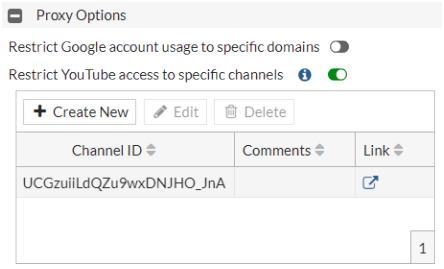
1. Go to *Security Profiles* > *Web Filter* and go to the *Proxy Options* section.
2. Enable *Restrict YouTube access to specific channels*.



3. Select *Create New* and specify the *Channel ID*, for example, `UCGzuiiLdQZu9wxDNJHO_JnA`.



4. Select *OK* and the option shows the *Channel ID* and its *Link*.



To enable this feature in the CLI:

```
config webfilter profile
   edit "webfilter"
```

```
      set youtube-channel-status whitelist  <-- whitelist: only allow the traffic belongs to
this channel id and relative identifiers

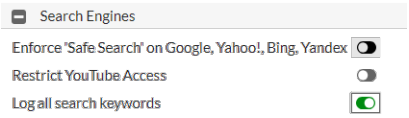
                                                blacklist: only block the traffic belongs to
this channel id and relative identifiers and allow the other traffic pass
      config youtube-channel-filter
         edit 1
            set channel-id "UCGzuiiLdQZu9wxDNJHO_JnA"
         next
      end
   next
end
```

Log all search keywords

Use this feature to log all search phrases.

To enable this feature in the GUI:

1. Go to *Security Profiles* > *Web Filter* and go to the *Search Engines* section.
2. Enable *Log all search keywords*.



To enable this feature in the CLI:

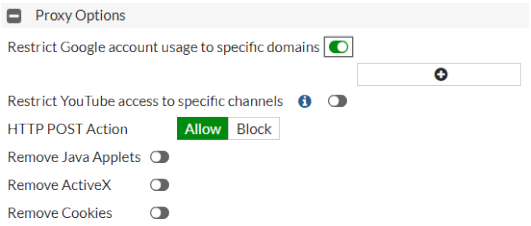
```
config webfilter profile
    edit "webfilter"
        config web
            set log-search enable
        end
    next
end
```

Restrict Google account usage to specific domains

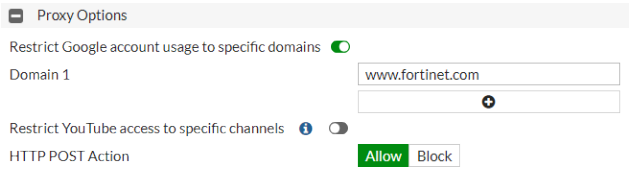
Use this feature to block access to some Google accounts and services while allowing access to accounts in the domains in the exception list.

To enable this feature in the GUI:

1. Go to *Security Profiles* > *Web Filter* and go to the *Proxy Options* section.
2. Enable *Restrict Google account usage to specific domains*.



3. Select the + button and enter the domains that Google can access, for example, www.fortinet.com.



When you try to use Google services like Gmail, only traffic from the domain of www.fortinet.com can go through. Traffic from other domains is blocked.

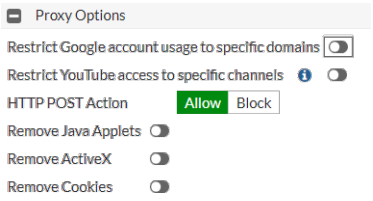
HTTP POST Action

Select the action to take with HTTP POST traffic. HTTP POST is the command used by your browser when you send information, such as a form you have filled-out or a file you are uploading to a web server.

The action options are Allow or Block. The default is Allow.

To enable this feature in the GUI:

1. Go to *Security Profiles* > *Web Filter* and go to the *Proxy Options* section.
2. For *HTTP POST Action*, select *Allow* or *Block*.



To enable this feature in the CLI:

```
config webfilter profile
    edit "webfilter"
        set post-action [normal/block]
        config ftgd-wf
            unset options
        end
    next
end
```

Remove Java applets, remove ActiveX, and remove cookies

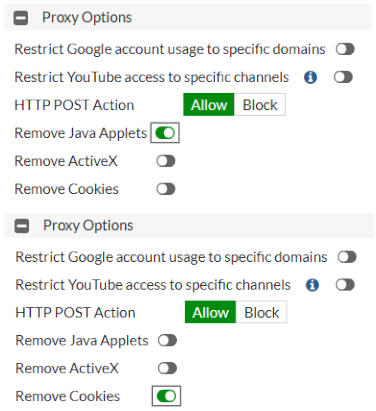
The *Remove Java Applets* feature filters java applets from web traffic. Websites using java applets might not function properly if you enable this filter.

The *Remove ActiveX* feature filters ActiveX scripts from web traffic. Websites using ActiveX might not function properly with if you enable this filter.

The *Remove Cookies* feature filters cookies from web traffic. Websites using cookies might not function properly if you enable this filter.

To enable this feature in the GUI:

1. Go to *Security Profiles* > *Web Filter* and go to the *Proxy Options* section.
2. Select the filters you want to use: *Remove Java Applets*, *Remove ActiveX*, and/or *Remove Cookies*.



To enable this feature in the CLI:

```
config webfilter profile
   edit "webfilter"
      set options activexfilter cookiefilter javafilter <-- enable one or more of
activexfilter cookiefilter javafilter.
      config ftgd-wf
         unset options
      end
   next
end
```

External resources for web filter

External Resources is a new feature introduced in FortiOS 6.0, which provides a capability to import an external blocklist which sits on an HTTP server. This feature helps FortiGate retrieve a dynamic URL/Domain Name/IP Address/Malware hash list from an external HTTP server periodically. FortiGate uses these external resources as Web Filter's remote categories, DNS Filter's remote categories, policy address objects or antivirus profile's malware definitions. If the external resource is updated, FortiGate objects will update dynamically.

External Resource are categorized into 4 types:

- URL list (Type=*category*)
- Domain Name List (Type=*domain*)

- IP Address list (Type=*address*)
- Malware hash list (Type=*malware*)

For Web Filter profile, it can use category type external resources. Category type external resources file is a URL entries list in a plain text file.

When a *category* type external resource is configured in Web Filter profile, it will be treated as a Remote Category. If the URL in a HTTP/HTTPS request matches the entry inside this external resource file, it will be treated as the Remote Category and follow the action configured for this category in Web Filter profile.

External resource type *category* also can be used in *ssl-ssh-profile* configuration for category-based *SSL-Exempt*. When a Remote Category is configured in *ssl-ssh-profile SSL-Exempt*, if a HTTPS request's URL matches in the Remote Category's entry list, HTTPS request with destination for this URL can be exempted from SSL Deep Inspection.

External Resources File Format

External Resources File should follow the following requirements:

- The external resource file is a plain text format file and each URL list/IP Address/Domain Name occupies a single line.
- The file is limited to 10M, line is limited 128K (128 x 1024 entries), and the line length limit is 4K characters.
- The entries limited also follow table size limitation defined by CMDB per model.
- The external resource update period can be set to 1 minute, hourly, daily, weekly, or monthly (43200 min, 30 days).
- The external resource type as category (URL list) and domain (Domain Name list) share the category number range 192-221 (total 30 categories).
- There's no duplicated entry validation for external resources file (entry inside each file or inside different files).

For URL list (Type=category):

Scheme is optional, and will be truncated if found (http://, https:// is not needed).

Wildcard (*) is supported (from 6.2). It supports the '*' at beginning and ending of URL, and not in the middle of URL as follows:

```
+ support *.domain2.com, domain.com.* + not support: domain3.*.com
```

IDN (International Domain Name) and UTF encoding URL is supported (from 6.2).

IPv4,IPv6 format URL is supported. IPv6 in URL list must in [ ] form.

Configure External Resources from CLI

We can use CLI to configure the external resources files that is located on external HTTP Server. Under Global, configure the external resource file location and specify the resource type.

Web Filter will use *category* type external resources as Remote Categories. In the following example, it is configured a file *Ext-Resource-Type-as-Category-1.txt* as type as category, it will be treated in Web Filter as Remote Category, the category name configured as *Ext-Resource-Type-as-Category-1* and category-id as *192*:

```
config system external-resource
   edit "Ext-Resource-Type-as-Category-1"
     set type category  <----
     set category 192   <----
     set resource "http://172.16.200.66/external-resources/Ext-Resource-Type-as-Category-
1.txt"
     set refresh-rate 1
```


Netskope Exhibit 1017

```
   next
end
```

Now in each VDOM, category type external resource can be used in Web Filter as Remote Cateogry. In the example above, URL list in "Ext-Resource-Type-as-Category-1.txt" file will be treated as remote category (category-id 192). Configure the action for this remote category in Web Filter profile and apply it in the policy:

```
config webfilter profile
    edit "webfilter"
        config ftgd-wf
            unset options
            config filters
                edit 1
                    set category 2
                    set action warning
                next
                ......

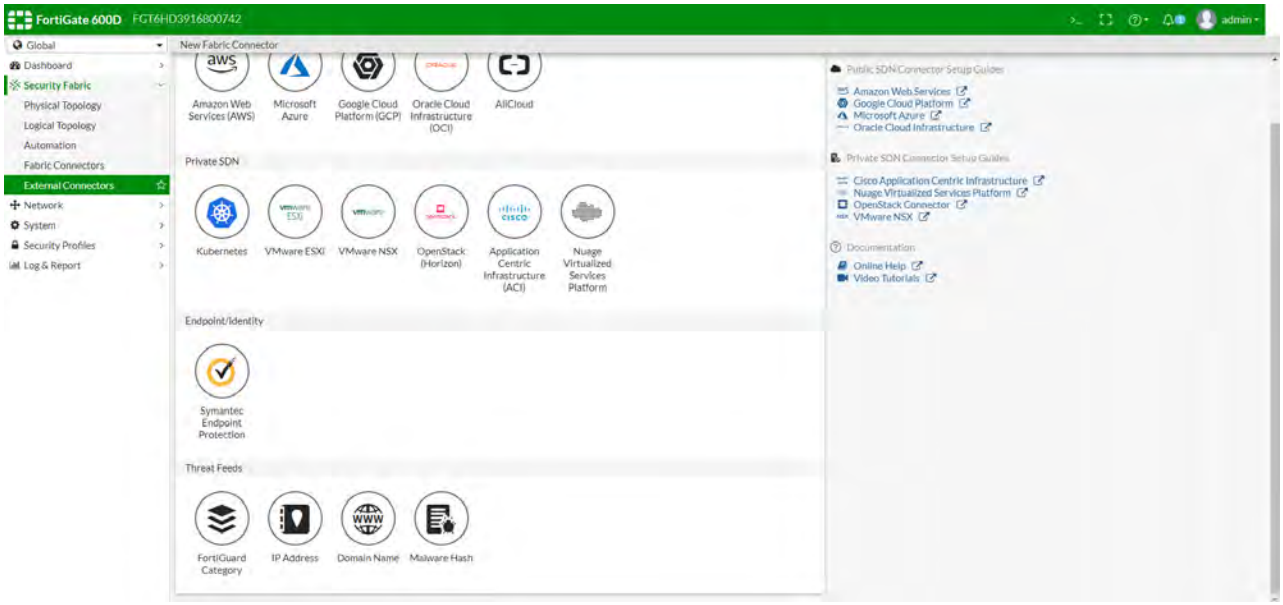
                edit 24
                    set category 192  <----
                    set action block
                next
                edit 25
                    set category 221
                    set action warning
                next
                edit 26
                    set category 193
                next
            end
        end
        set log-all-url enable
    next
end

config firewall policy
    edit 1
        set name "WebFilter"
        set srcintf "port10"
        set dstintf "port9"
        set srcaddr "all"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ALL"
        set utm-status enable
        set logtraffic all
        set webfilter-profile "webfilter"
        set profile-protocol-options "protocol"
        set ssl-ssh-profile "protocols"
        set nat enable
    next
end
```

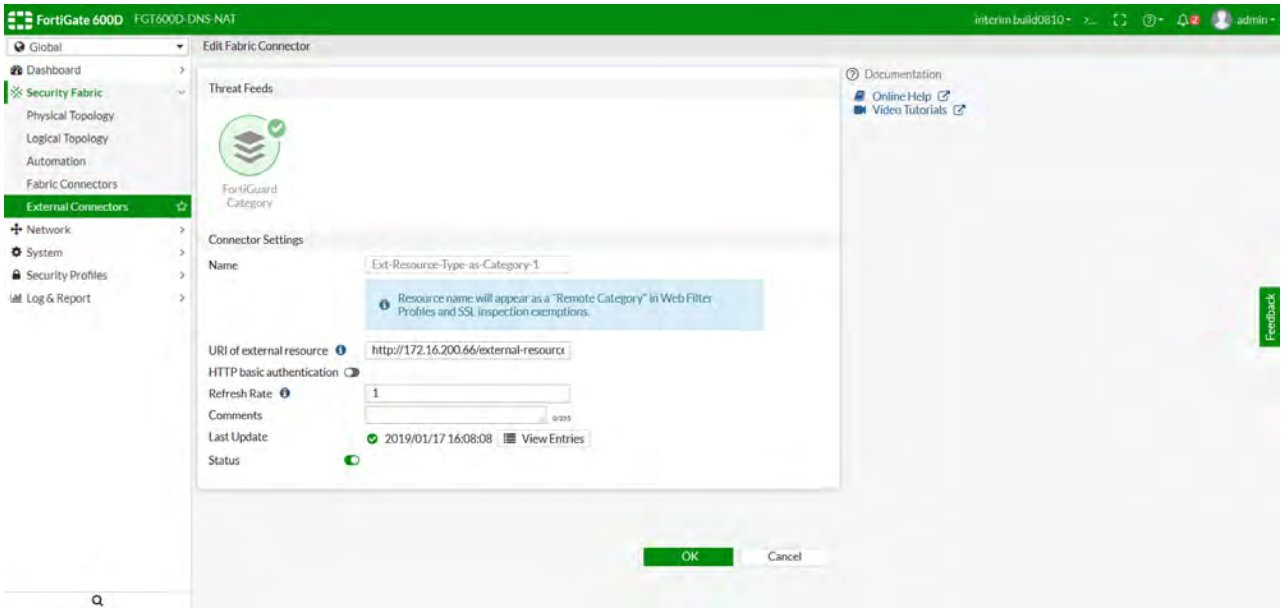
Configure External Resources from GUI

To configure, edit, or view the Entries for external resources from GUI:

1. Go to *Global* > *Security Fabric* > *External Connectors*.
2. Click *Create New*, and in the *Threat Feeds* section, select *FortiGuard Category*.

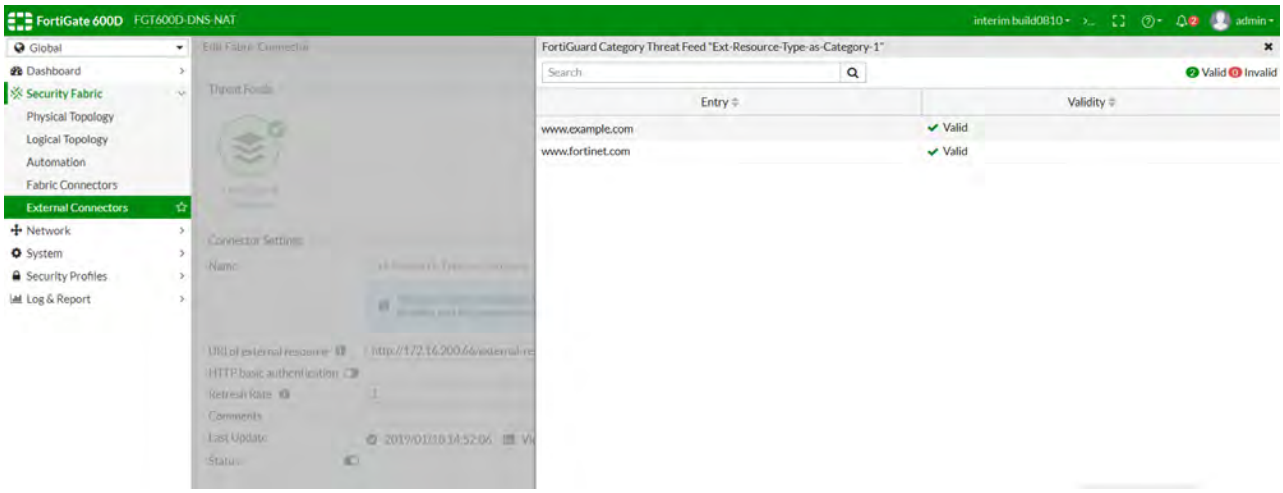


3. Enter the resource name, URI location of the resource file, resource authentication credential, and *Refresh Rate*.

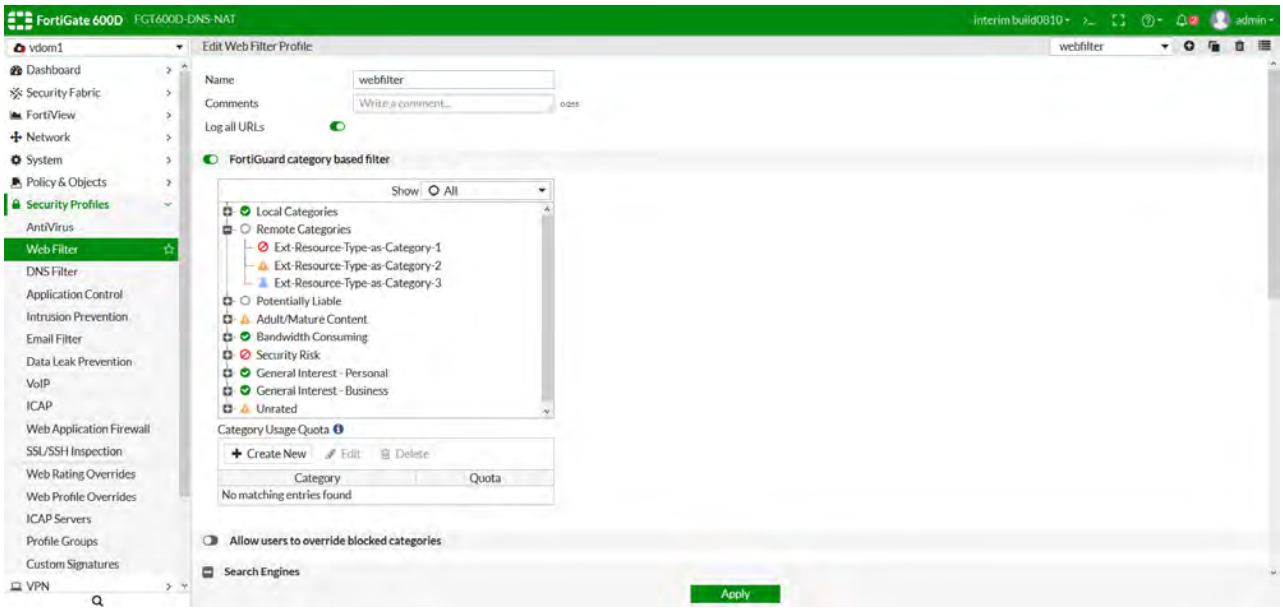


4. Click *OK*.
5. After a few minutes, double-click the *Threat Feeds Object* you just configured. It is shown in the *Edit* page.

6. Click *View Entries* to view the entry list in the external resources file:



7. Go to *VDOM* > *Security Profiles* > *Web Filter*. The configured external resources is shown and configured in each Web Filter Profile:

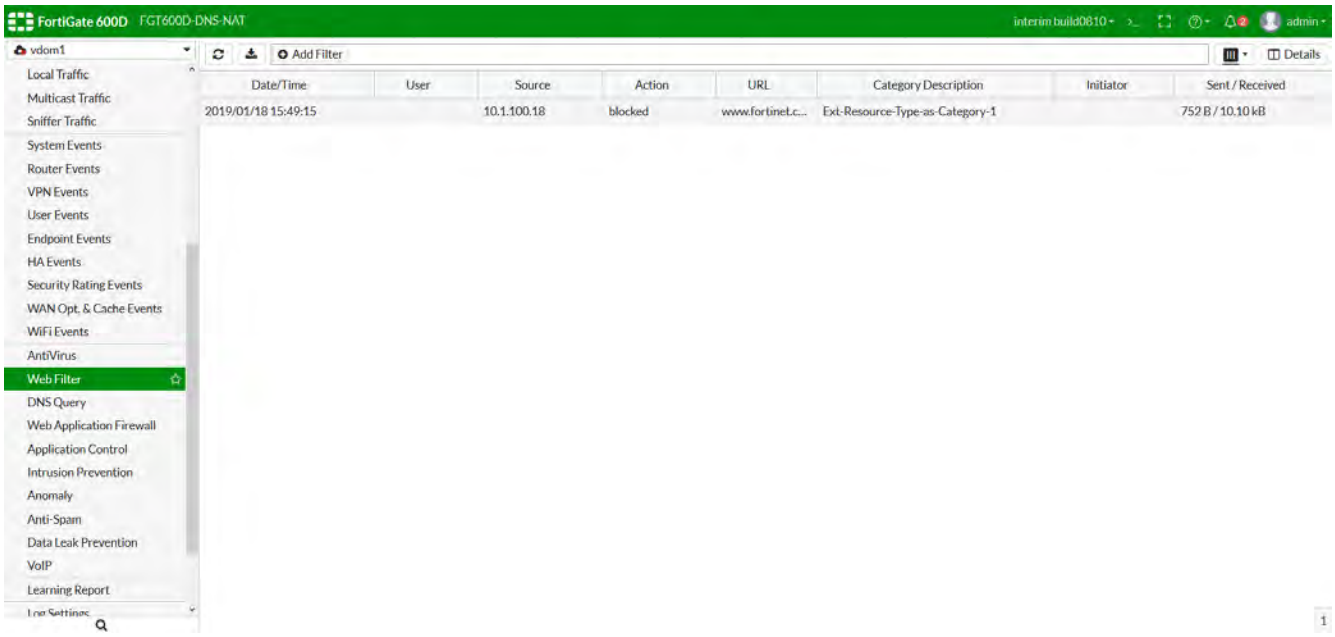


Log Example

If an HTTP/HTTPS request URL is matched in remote category's entry list, it will override its original FortiGuard URL rating and be treated as a remote category.

Go to *VDOM* > *Log & Report* > *Web Filter*:

Netskope Exhibit 1017



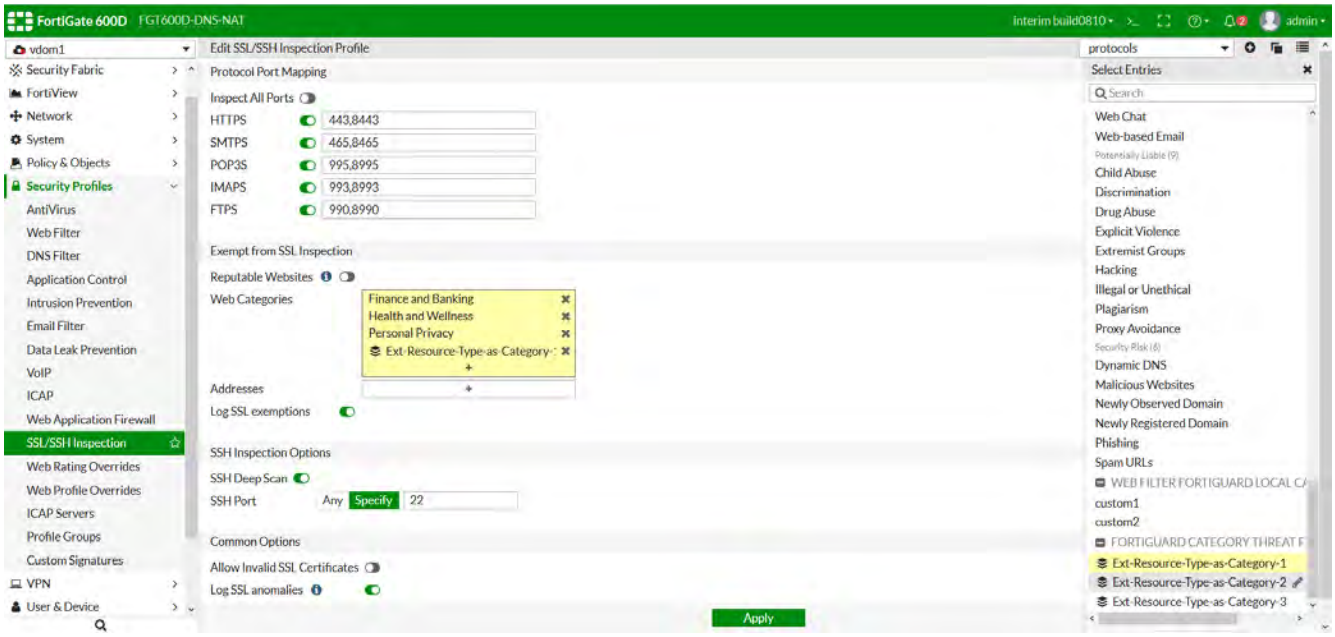
CLI Example:

```
1: date=2019-01-18 time=15:49:15 logid="0316013056" type="utm" subtype="webfilter"
eventtype="ftgd_blk" level="warning" vd="vdom1" eventtime=1547855353 policyid=1
sessionid=88922 srcip=10.1.100.18 srcport=39886 srcintf="port10" srcintfrole="undefined"
dstip=216.58.193.67 dstport=443 dstintf="port9" dstintfrole="undefined" proto=6
service="HTTPS" hostname="www.fortinet.com" profile="webfilter" action="blocked"
reqtype="direct" url="/" sentbyte=752 rcvdbyte=10098 direction="outgoing" msg="URL belongs
to a denied category in policy" method="domain" cat=192 catdesc="Ext-Resource-Type-as-
Category-1"
```

Remote Category in ssl-ssh-profile category-based SSL-Exempt

Remote category can be applied in *ssl-ssh-profile category-based SSL-Exempt*.

Go to *VDOM* > *Security Profiles* > *SSL/SSH Inspection*:



HTTPS request URLs matched in this remote category will be exempted from SSL deep inspection.

Log example:

```
3: date=2019-01-18 time=16:06:21 logid="0345012688" type="utm" subtype="webfilter"
eventtype="ssl-exempt" level="information" vd="vdom1" eventtime=1547856379 policyid=1
sessionid=90080 srcip=10.1.100.18 srcport=39942 srcintf="port10" srcintfrole="undefined"
dstip=216.58.193.67 dstport=443 dstintf="port9" dstintfrole="undefined" proto=6
service="HTTPS" hostname="www.fortinet.com" profile="webfilter" action="passthrough"
reqtype="direct" url="/" sentbyte=517 rcvdbyte=0 direction="outgoing" msg="The SSL session
was exempted." method="domain" cat=192 catdesc="Ext-Resource-Type-as-Category-1"
urlsource="exempt_type_user_cat"
```

Local Category and Remote Category Priority

Web Filter can have both local category and remote category at the same time. There's no duplication check between local category URL override and remote category resource file. For example, a URL like www.example.com may be shown both in remote category entry list and in FortiGate's local category URL override configuration. We recommend avoiding this scenario since FortiGate does not check for duplicates. However, if a URL is duplicated in both local category and remote category, it is rated as local category.

Reliable web filter statistics

FortiOS 6.4 provides command line tools to view the Web Filter statistics report. These command line tools currently fall into either proxy-based or flow-based Web Filter statistics commands.

Proxy-based Web Filter statistics report

- The proxy-based Web Filter statistics command line tools are as follows. These commands are available in both global or per-VDOM command lines.

```
#diagnose wad filter   <----define the interested objects for output
(global) # diagnose wad ?
console-log    Send WAD log messages to the console.
debug          Debug setting.
stats          Show statistics.
filter         Filter for listing sessions or tunnels. <----use filter to filter-out
interested object and output
kxp            SSL KXP diagnostics.
user           User diagnostics.
memory         WAD memory diagnostics.
restore        Restore configuration defaults.
history        Statistics history.
session        Session diagnostics.
tunnel         Tunnel diagnostics.
webcache       Web cache statistics.
worker         Worker diagnostics.
csvc           Cache service diagnostics.

#diagnose wad stat filter list/clear <----list/clear Web Filter/DLP statistics
report
```

- In the example below, there are two VDOMs using proxy-based policies which have Web Filter profiles enabled. The command line can be used to view the proxy-based Web Filter statistics report.

```
(global) # diagnose wad filter ?
list                    Display current filter.
clear                   Erase current filter settings.
src                     Source address range to filter by.
dst                     Destination address range to filter by.
sport                   Source port range to filter by.
dport                   Destination port range to filter by.
vd                      Virtual Domain Name.  <----filter for per-vdom or global
statistics report
explicit-policy         Index of explicit-policy. -1 matches all.
firewall-policy         Index of firewall-policy. -1 matches all.
drop-unknown-session    Enable drop message unknown sessions.
negate                  Negate the specified filter parameter.
protocol                Select protocols to filter by.

FGT_600D-ICAP-NAT (global) # diagnose wad filter vd
<vdom>    Virtual Domain Name.
ALL     all vdoms
root    vdom
vdom1   vdom

FGT_600D-ICAP-NAT (global) # diagnose wad filter vd root  <----filter-out root vdom
statistics
Drop_unknown_session is enabled.

FGT_600D-ICAP-NAT (global) # diagnose wad stats filter list
filtering of vdom root  <----Displayed the WF statistics for root vdom
  dlp          = 0      <----Number of Reuqest that DLP Sensor processed;
```


Netskope Exhibit 1017

```
  content-type = 0      <----Number of Reuqest that matching content-type filter;
  urls:
    examined = 6  <----Number of Request that Proxy Web-Filter(all wad daemons)
examined;
    allowed = 3   <----Number of Request that be allowed in the examined requests;
    blocked = 0   <----Number of Request that be blocked in the examined requests;
    logged = 0    <----Number of Request that be logged in the examined requests;
    overridden = 0 <----Number of Request that be overrided to another Web Filter
profile in the examined requests;

FGT_600D-ICAP-NAT (global) # diagnose wad filter vd vdom1  <----filter-out vdom1
statistics

FGT_600D-ICAP-NAT (global) # diagnose wad stats filter list
filtering of vdom vdom1  <----Displayed the WF statistics for vdom1
  dlp          = 0
  content-type = 0
  urls:
    examined = 13
    allowed = 2
    blocked = 9
    logged = 8
    overridden = 0

FGT_600D-ICAP-NAT (global) # diagnose wad filter vd ALL

FGT_600D-ICAP-NAT (global) # diagnose wad stats filter list
filtering of all accessible vdoms  <----global statistics is sum of two VDOMs
  dlp          = 0
  content-type = 0
  urls:
    examined = 19
    allowed = 5
    blocked = 9
    logged = 8
    overridden = 0
```

Flow-based Web Filter statistics report

- The command to check the flow-based Web Filter statistics is:
  `# diag webfilter stats list <vdom/global>`
  This command is available in both VDOM and global command lines.
- In the example below, the VDOM is using flow-based policies which have Web Filter profiles enabled.

```
FGT # diag webfilter stats list root

Proxy/flow URL filter stats:

request:   9474 <----Number of Requests that Flow Web-Filter(all ips engines) received;

blocked:   8606 <----Number of Request that Flow Web-Filter blocked;

allowed:   868 <----Number of Request that Flow Web-Filter allowed;

overridden:0 <----Number of Request that be overrided to another Web Filter profile in
the examined requests;

logged:    8606 <----Number of Request that be logged in the examined requests;
```


Netskope Exhibit 1017

```
pending:    0 <----Number of Requests that under processing at that moment;
```

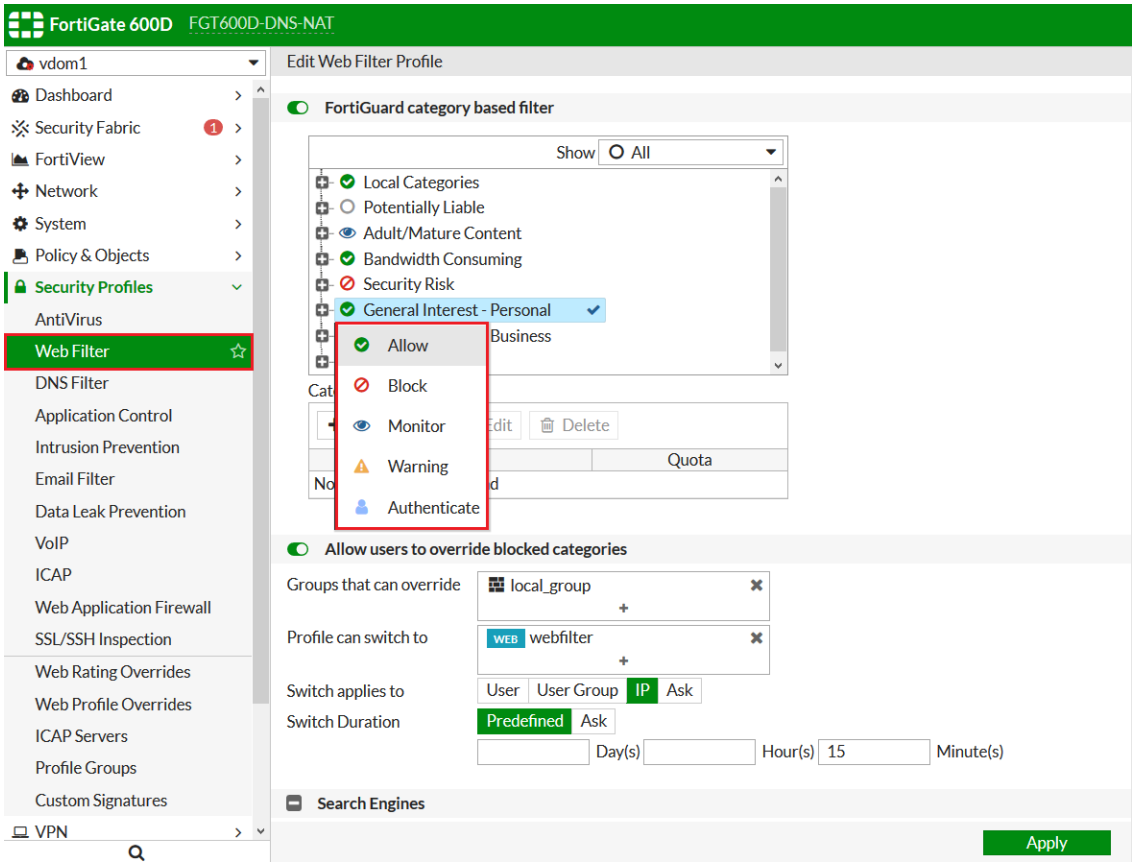
Flow-based web filtering

Flow-based web filtering includes the following options:

- Authenticate: requires user authentication for specific website categories.
- Warn: displays a warning message for specific website categories, but allows users to continue to the website.
- Override: allows users with valid credentials to override their web filter profile.

To enable *Authenticate* and *Warning* web filters:

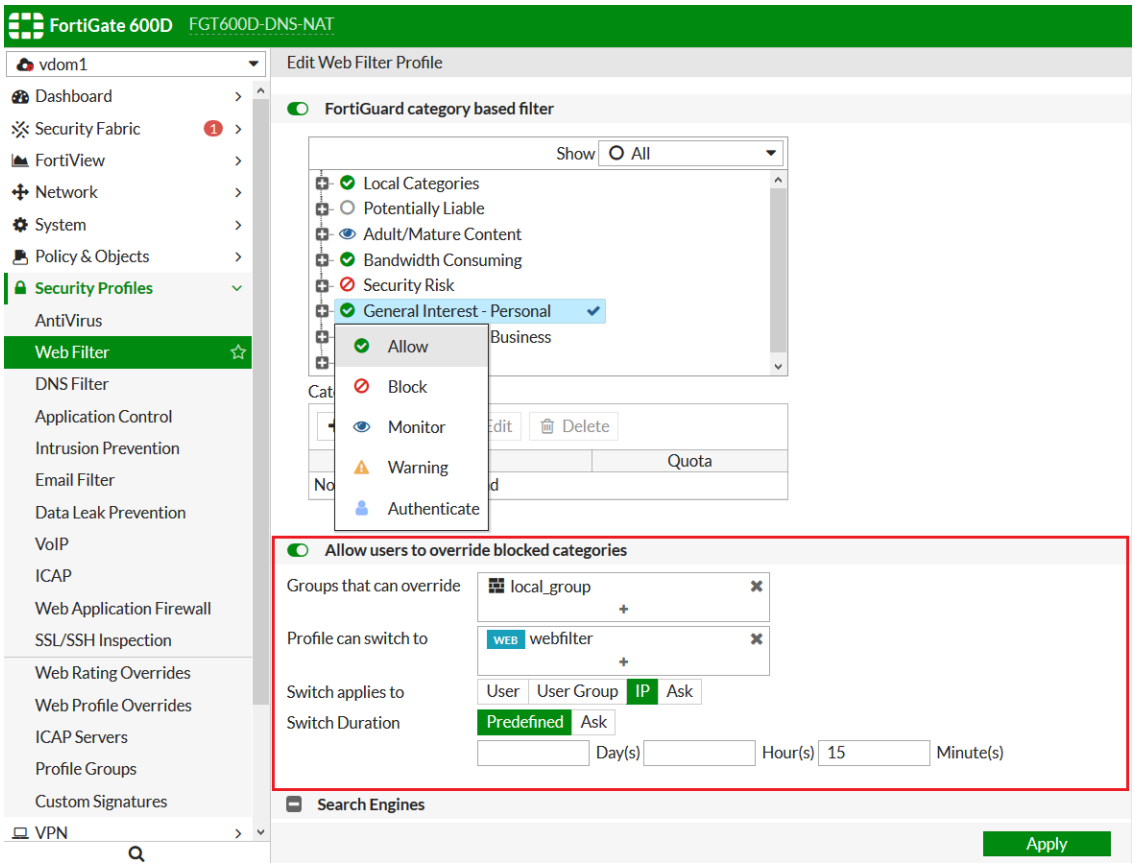
1. Go to *Security Profiles* > *Web Filter*.
2. Edit an existing profile, or create a new one.
3. In the *FortiGuard category based filter* section, right-click on a category.
4. Select the *Authenticate* or *Warning* web filter.
5. Configure other settings as needed.
6. Click *Apply*.



To allow users to override blocked categories:

1. Go to *Security Profiles* > *Web Filter*.
2. Edit an existing profile, or create a new one.

3. Enable *Allow users to override blocked categories*.
4. Enter information in the following fields:
   - *Groups that can override*
   - *Profile can switch to*
   - *Switch applies to*
   - *Switch Duration*
5. Configure other settings as needed.
6. Click *Apply*.



URL certificate blocklist

As increasing numbers of malware have started to use SSL to attempt to bypass IPS, maintaining a fingerprint-based certificate blocklist is useful to block botnet communication that relies on SSL.

This feature adds a dynamic package that is distributed by FortiGuard and is part of the Web Filtering service. It is enabled by default for SSL/SSH profiles, and can be configured using the following CLI commands:

```
config vdom
    edit <vdom>
        config firewall ssl-ssh-profile
            edit "certificate-inspection"
                set block-blacklisted-certificates enable
            next
            edit "deep-inspection"
```


Netskope Exhibit 1017

```
                set block-blacklisted-certificates enable
            next
        end
    next
end
```

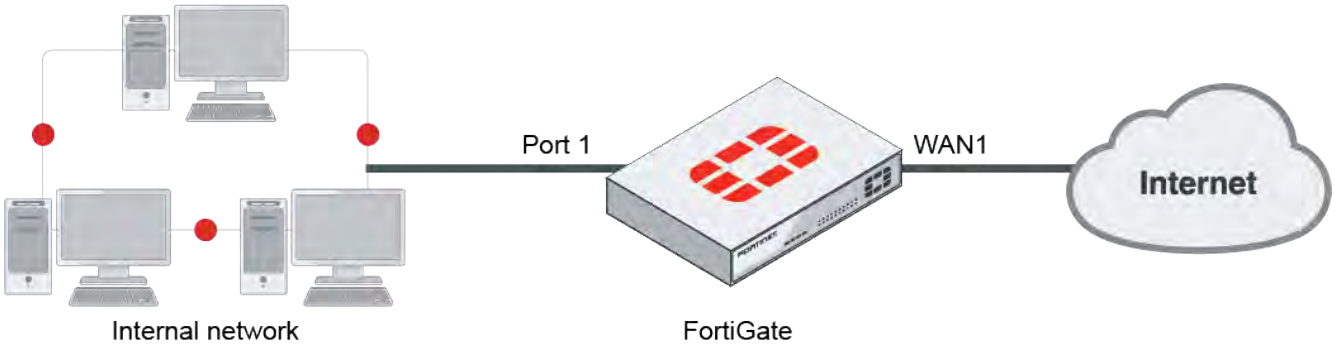
DNS filter

You can apply DNS category filtering to control user access to web resources. You can customize the default profile, or create your own to manage network user access and apply it to a firewall policy, or you can add it to a DNS server on a FortiGate interface.

DNS filtering has the following features:

- FortiGuard Filtering: filters the DNS request based on the FortiGuard domain rating.
- Botnet C&C domain blocking: blocks the DNS request for the known botnet C&C domains.
- External dynamic category domain filtering: allows you to define your own domain category.
- DNS safe search: enforces Google, Bing, and YouTube safe addresses for parental controls.
- Local domain filter: allows you to define your own domain list to block or allow.
- External IP block list: allows you to define an IP block list to block resolved IPs that match this list.
- DNS translation: maps the resolved result to another IP that you define.

The following sample topology is used in the topics of this section. It includes an internal network and a FortiGate that is used as a gateway device that all DNS traffic traverses.



Some features of this functionality require a subscription to FortiGuard Web Filtering.

The following topics provide information about DNS filters:

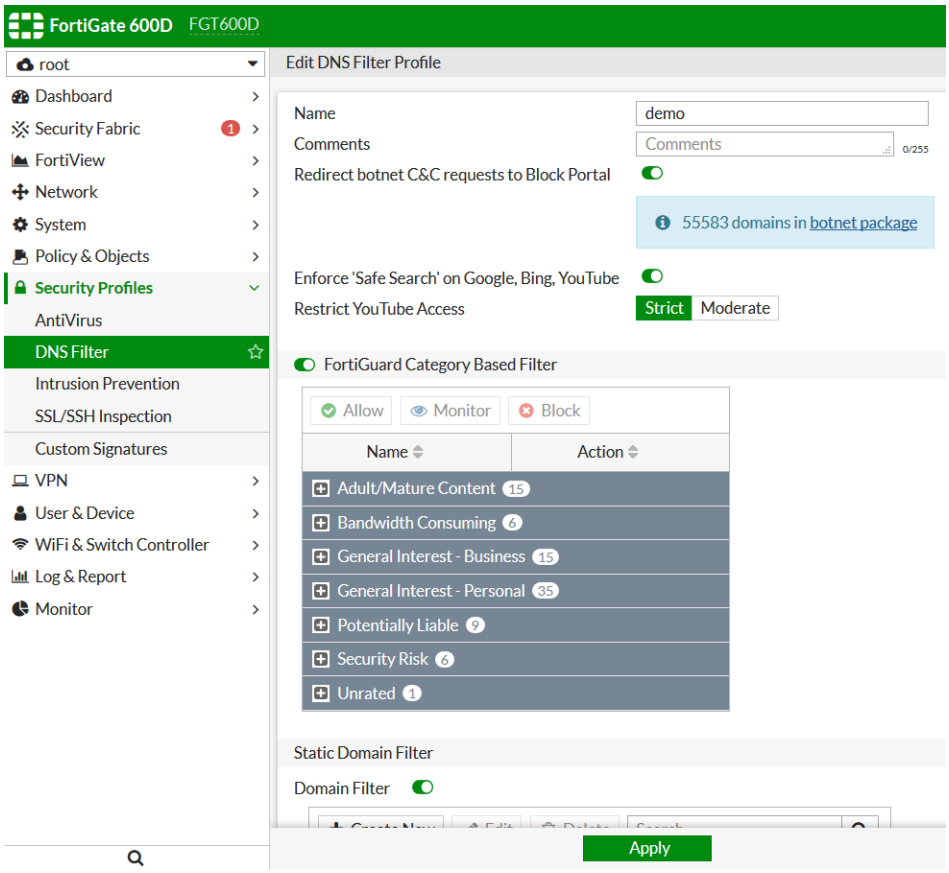
- How to configure and apply a DNS filter profile on page 943
- FortiGuard category-based DNS domain filtering on page 945
- Botnet C&C domain blocking on page 948
- DNS safe search on page 952
- Local domain filter on page 953
- DNS translation on page 956

- Using a FortiGate as a DNS server on page 959
- Troubleshooting for DNS filter on page 961

How to configure and apply a DNS filter profile

To create or configure DNS Filter profile in the GUI:

1. Go to *Security Profiles* > *DNS Filter*.
2. You can modify the default DNS Filter and enable the options you want or you can click + at the top right to create a new DNS Filter.



To create or configure DNS Filter profile in the CLI:

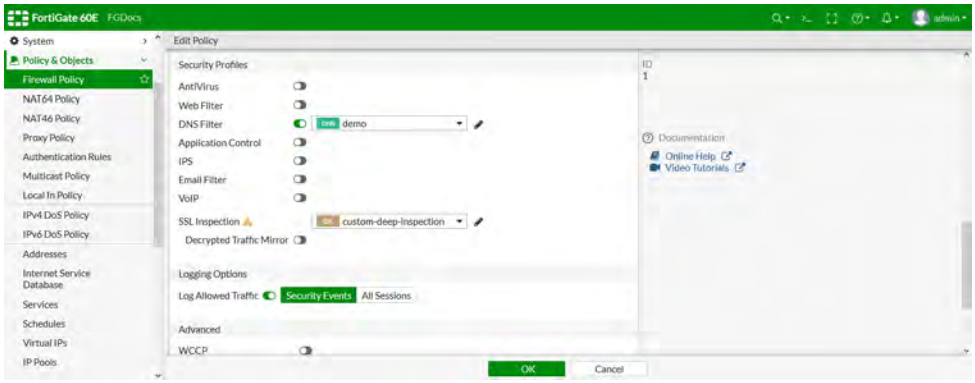
```
config dnsfilter profile
    edit "demo"
        set comment ''
        config domain-filter
            unset domain-filter-table
        end
        config ftgd-dns
            set options error-allow
            config filters
                edit 2
                    set category 2
                    set action monitor
```

```
                next
                edit 7
                    set category 7
                    set action block
                next
                ...
                edit 22
                    set category 0
                    set action monitor
                next
            end
        end
        set log-all-domain enable
        set sdns-ftgd-err-log enable
        set sdns-domain-log enable
        set block-action redirect
        set block-botnet enable
        set safe-search enable
        set redirect-portal 93.184.216.34
        set redirect-portal6 ::
        set youtube-restrict strict
    next
end
```

After you have created the DNS Filter profile, you can apply it to the policy. DNS filters also support IPv6 policies.

To apply DNS Filter profile to the policy in the GUI:

1. Go to *Policy & Objects* > *Firewall Policy*.
2. Edit or create a new policy.
3. In the *Security Profiles* section, enable *DNS Filter* and select the DNS filter.



To apply DNS Filter profile to the policy in the CLI:

```
config firewall policy
    edit 1
        set name "Demo"
        set srcintf "port10"
        set dstintf "port9"
        set srcaddr "all"
        set dstaddr "all"
        set action accept
        set schedule "always"
```

```
        set service "ALL"
        set utm-status enable
        set inspection-mode proxy
        set logtraffic all
        set fsso disable
        set dnsfilter-profile "demo"  <<<====
        set profile-protocol-options "default"
        set ssl-ssh-profile "deep-inspection"
        set nat enable
    next
end
```

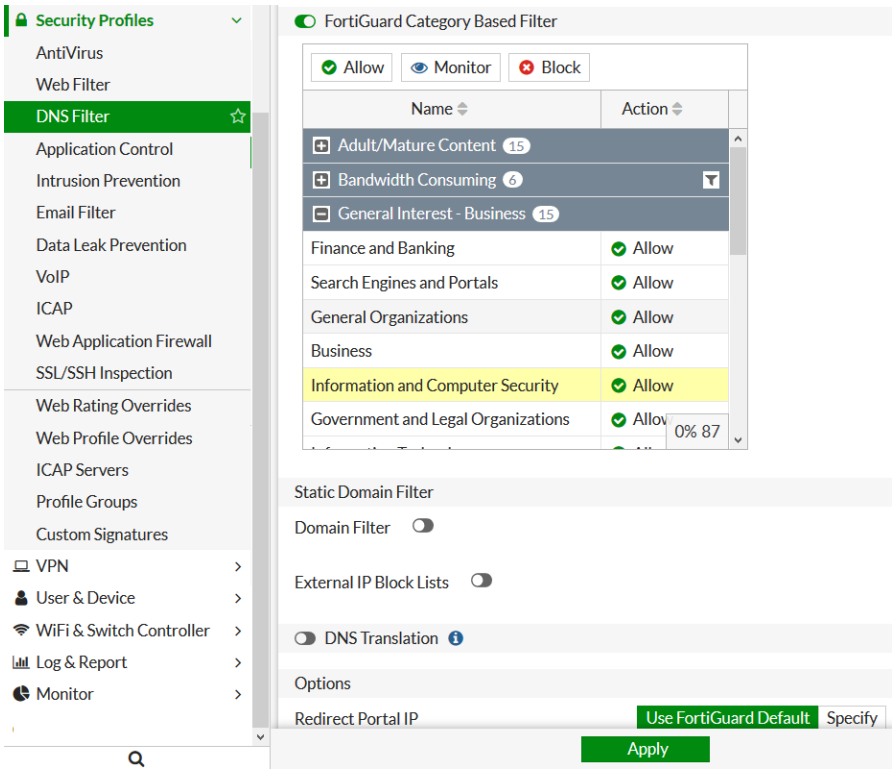
FortiGuard category-based DNS domain filtering

The FortiGate must have a FortiGuard Web Filter license to use FortiGuard Category Based Filter.

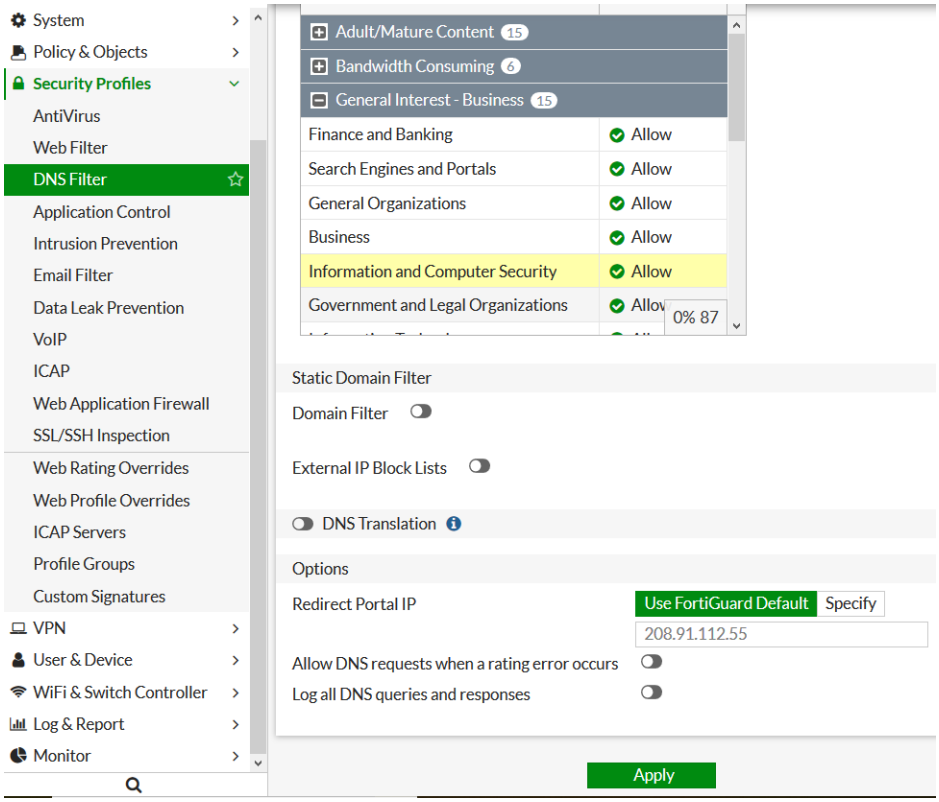
You can use the FortiGuard category-based DNS Domain Filter to inspect DNS traffic. This makes use of FortiGuard's continually updated domain rating database for more reliable protection.

To configure FortiGuard category-based DNS Domain Filter by GUI:

1. Go to *Security Profiles* > *DNS Filter* and edit or create a DNS Filter.
2. Enable *FortiGuard Category Based Filter*.
3. Select the category and then select *Allow*, *Monitor*, or *Block* for that category.



4. If you select *Block*, there are two options:
   - *Redirect Portal IP*. If the DNS query domain will be blocked, FortiGate will use portal IP to replace the resolved IP in DNS response packet. You can use the default portal IP 208.91.112.55 or click *Specify* to enter another portal IP.
   - *Block*. Blocked DNS query has no response return and the DNS query client will timeout.



To configure FortiGuard category-based DNS Domain Filter by CLI:

```
config dnsfilter profile
   edit "demo"
      set comment ''
      config domain-filter
         unset domain-filter-table
      end
      config ftgd-dns
         set options error-allow
         config filters  <<<==== FortiGuard Category Based Filter
             edit 2
                 set category 2
                 set action monitor
             next
             edit 7
                 set category 7
                 set action monitor
             next
            ...
             edit 22
                 set category 0
                 set action monitor
```

```
            next
         end
      end
      set log-all-domain enable
      set sdns-ftgd-err-log enable
      set sdns-domain-log enable
      set block-action redirect/block  <<<==== You can specify Block or Redirect
      set block-botnet enable
      set safe-search enable
      set redirect-portal 93.184.216.34  <<<==== Specify Redirect portal-IP.
      set redirect-portal6 ::
      set youtube-restrict strict
   next
end
```

Sample

To see an example of how this works, from your internal network PC, use a command line tool such as dig or nslookup to do DNS query for some domains, for example:

```
#dig www.example.com
;; ->>HEADER<<- opcode: QUERY; status: NOERROR; id: 61252
;; Flags: qr rd ra; QUERY: 1; ANSWER: 1; AUTHORITY: 13; ADDITIONAL: 11

;; QUESTION SECTION:
;; www.example.com.             IN      A

;; ANSWER SECTION:
www.example.com.        17164   IN      A       93.184.216.34

;; AUTHORITY SECTION:
com.                    20027   IN      NS      h.gtld-servers.net.
com.                    20027   IN      NS      i.gtld-servers.net.
com.                    20027   IN      NS      f.gtld-servers.net.
com.                    20027   IN      NS      d.gtld-servers.net.
com.                    20027   IN      NS      j.gtld-servers.net.
com.                    20027   IN      NS      l.gtld-servers.net.
com.                    20027   IN      NS      e.gtld-servers.net.
com.                    20027   IN      NS      a.gtld-servers.net.
com.                    20027   IN      NS      k.gtld-servers.net.
com.                    20027   IN      NS      g.gtld-servers.net.
com.                    20027   IN      NS      m.gtld-servers.net.
com.                    20027   IN      NS      c.gtld-servers.net.
com.                    20027   IN      NS      b.gtld-servers.net.

;; ADDITIONAL SECTION:
a.gtld-servers.net.     21999   IN      A       192.5.6.30
a.gtld-servers.net.     21999   IN      AAAA    2001:503:a83e::2:30
b.gtld-servers.net.     21997   IN      A       192.33.14.30
b.gtld-servers.net.     21997   IN      AAAA    2001:503:231d::2:30
c.gtld-servers.net.     21987   IN      A       192.26.92.30
c.gtld-servers.net.     20929   IN      AAAA    2001:503:83eb::30
d.gtld-servers.net.     3340    IN      A       192.31.80.30
d.gtld-servers.net.     3340    IN      AAAA    2001:500:856e::30
e.gtld-servers.net.     19334   IN      A       192.12.94.30
e.gtld-servers.net.     19334   IN      AAAA    2001:502:1ca1::30
```


Netskope Exhibit 1017

```
f.gtld-servers.net.     3340    IN      A       192.35.51.30

;; Received 509 B
;; Time 2019-04-05 09:39:33 PDT
;; From 172.16.95.16@53(UDP) in 3.8 ms
```

To check the DNS Filter log in the GUI:

1. Go to *Log & Report* > *DNS Query* to view the DNS traffic that just traverse the FortiGate and the FortiGuard rating for this domain name.



To check the DNS log in the CLI:

```
#execute log filter category utm-dns

# execute log display
2 logs found.
2 logs returned.

1: date=2019-04-05 time=09:39:34 logid="1501054802" type="utm" subtype="dns" eventtype="dns-
response" level="notice" vd="vdom1" eventtime=1554482373 policyid=1 sessionid=50868
srcip=10.1.100.18 srcport=34308 srcintf="port10" srcintfrole="undefined" dstip=172.16.95.16
dstport=53 dstintf="port9" dstintfrole="undefined" proto=17 profile="demo" xid=17647
qname="www.example.com" qtype="A" qtypeval=1 qclass="IN" ipaddr="93.184.216.34" msg="Domain
is monitored" action="pass" cat=52 catdesc="Information Technology"

2: date=2019-04-05 time=09:39:34 logid="1500054000" type="utm" subtype="dns" eventtype="dns-
query" level="information" vd="vdom1" eventtime=1554482373 policyid=1 sessionid=50868
srcip=10.1.100.18 srcport=34308 srcintf="port10" srcintfrole="undefined" dstip=172.16.95.16
dstport=53 dstintf="port9" dstintfrole="undefined" proto=17 profile="demo" xid=17647
qname="www.example.com" qtype="A" qtypeval=1 qclass="IN"
```

Botnet C&C domain blocking

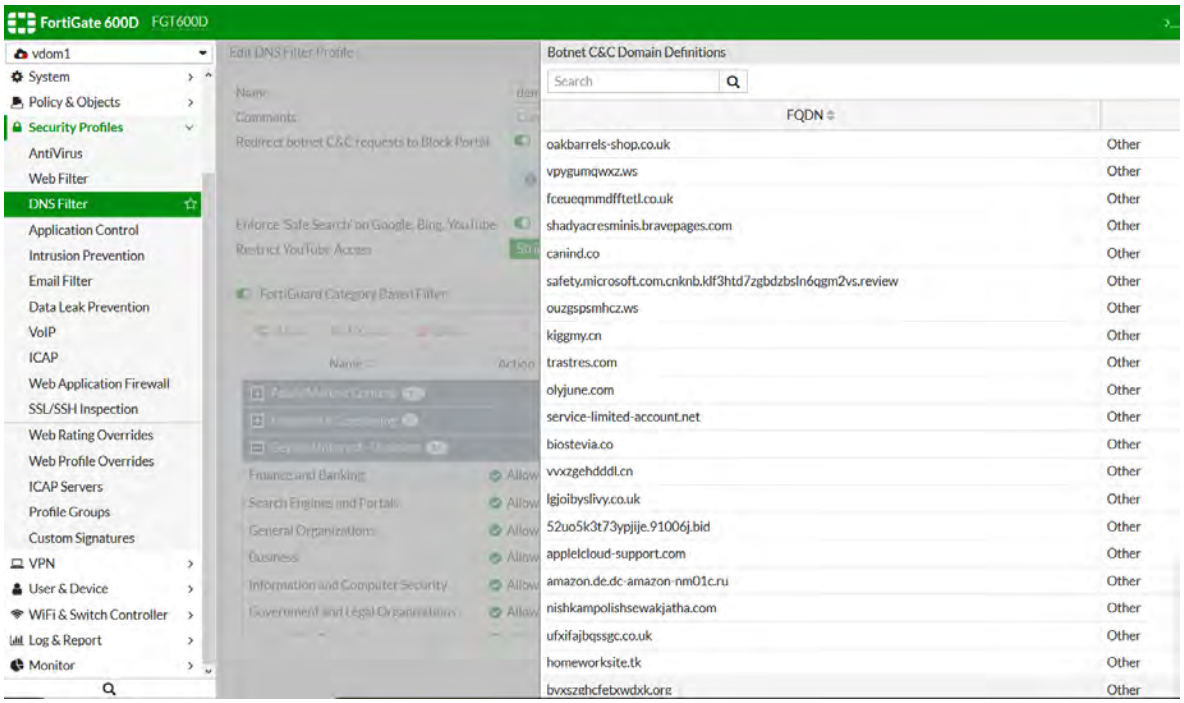
FortiGuard Service continually updates the Botnet C&C domain list (Domain DB). The botnet C&C domain blocking feature can block the botnet website access at the DNS name resolving stage. This provides additional protection for your network.

To configure botnet C&C domain blocking in the GUI:

1. Go to *Security Profiles* > *DNS Filter* and edit or create a DNS Filter.
2. Enable *Redirect botnet C&C requests to Block Portal*.



3. Click the *botnet package* link to see the latest botnet C&C domain list.



Sample

To see an example of how this works, select a botnet domain from that list. Then from your internal network PC, use a command line tool such as dig or nslookup to send a DNS query to traverse the FortiGate to see the query blocked as a botnet domain. For example:

```
#dig canind.co
;; ->>HEADER<<- opcode: QUERY; status: NOERROR; id: 997
;; Flags: qr rd; QUERY: 1; ANSWER: 1; AUTHORITY: 0; ADDITIONAL: 0

;; QUESTION SECTION:
;; canind.co.                   IN      A

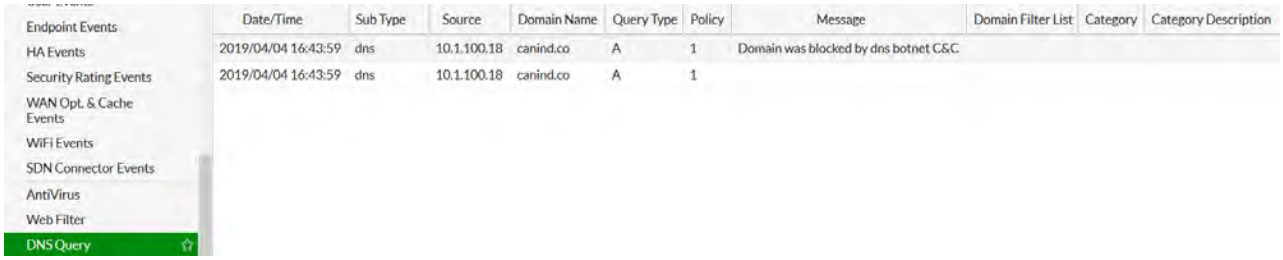
;; ANSWER SECTION:
canind.co.              60      IN      A       208.91.112.55  <<<==== botnet domain query
blocked, redirect with portal-IP.

;; Received 43 B
```

```
;; Time 2019-04-05 09:55:21 PDT
;; From 172.16.95.16@53(UDP) in 0.3 ms
```

To check the DNS Filter log in the GUI:

1. Go to *Log & Report* > *DNS Query* to view the DNS query blocked as a botnet domain.



To check the DNS Filter log in the CLI:

```
(vdom1) # execute log filter category utm-dns

(vdom1) # execute log display
2 logs found.
2 logs returned.

1: date=2019-04-04 time=16:43:59 logid="1501054601" type="utm" subtype="dns" eventtype="dns-
response" level="warning" vd="vdom1" eventtime=1554421439 policyid=1 sessionid=14135
srcip=10.1.100.18 srcport=57447 srcintf="port10" srcintfrole="undefined" dstip=172.16.95.16
dstport=53 dstintf="port9" dstintfrole="undefined" proto=17 profile="demo" xid=24339
qname="canind.co" qtype="A" qtypeval=1 qclass="IN" msg="Domain was blocked by dns botnet
C&C" action="redirect" botnetdomain="canind.co"

2: date=2019-04-04 time=16:43:59 logid="1500054000" type="utm" subtype="dns" eventtype="dns-
query" level="information" vd="vdom1" eventtime=1554421439 policyid=1 sessionid=14135
srcip=10.1.100.18 srcport=57447 srcintf="port10" srcintfrole="undefined" dstip=172.16.95.16
dstport=53 dstintf="port9" dstintfrole="undefined" proto=17 profile="demo" xid=24339
qname="canind.co" qtype="A" qtypeval=1 qclass="IN"
```

Botnet C&C IPDB blocking

FortiGate also maintains a botnet C&C IP address database (botnet IPDB). If a DNS query response IP address (resolved IP address) matches an entry inside the botnet IPDB, this DNS query is also blocked by DNS Filter botnet C&C blocking.

To view the botnet IPDB list in the CLI:

```
(global) # diagnose sys botnet list 9000 10
9000. proto=TCP ip=103.228.28.166, port=80, rule_id=7630075, name_id=3, hits=0
9001. proto=TCP ip=5.9.32.166, port=481, rule_id=4146631, name_id=7, hits=0
9002. proto=TCP ip=91.89.44.166, port=80, rule_id=48, name_id=96, hits=0
9003. proto=TCP ip=46.211.46.166, port=80, rule_id=48, name_id=96, hits=0
9004. proto=TCP ip=77.52.52.166, port=80, rule_id=48, name_id=96, hits=0
9005. proto=TCP ip=98.25.53.166, port=80, rule_id=48, name_id=96, hits=0
9006. proto=TCP ip=70.120.67.166, port=80, rule_id=48, name_id=96, hits=0
9007. proto=TCP ip=85.253.77.166, port=80, rule_id=48, name_id=96, hits=0
```


Netskope Exhibit 1017

```
9008. proto=TCP ip=193.106.81.166, port=80, rule_id=48, name_id=96, hits=0
9009. proto=TCP ip=58.13.84.166, port=80, rule_id=48, name_id=96, hits=0
```

To see an example of how DNS Filter botnet C&C IPDB blocking works, select an IP address from the IPDB list and use Internet reverse lookup service to find its corresponding domain name. Then from your internal network PC, use a command line tool such as dig or nslookup to query this domain and see that it's blocked by DNS Filter botnet C&C blocking. For example:

```
# dig cpe-98-25-53-166.sc.res.rr.com
;; ->>HEADER<<- opcode: QUERY; status: NOERROR; id: 35135
;; Flags: qr rd; QUERY: 1; ANSWER: 1; AUTHORITY: 0; ADDITIONAL: 0

;; QUESTION SECTION:
;; cpe-98-25-53-166.sc.res.rr.com.              IN      A

;; ANSWER SECTION:
cpe-98-25-53-166.sc.res.rr.com. 60      IN      A       208.91.112.55  <<<====  Since
resolved IP address match the botnet IPDB, dns query blocked with redirect portal IP.

;; Received 64 B
;; Time 2019-04-05 11:06:47 PDT
;; From 172.16.95.16@53(UDP) in 0.6 ms
```

To check the DNS Filter log in the GUI:

1. Go to *Log & Report* > *DNS Query* to view the DNS query blocked by botnet C&C IPDB blocking.

Date/Time	Sub Type	Source	Domain Name	Query Type	Policy	Message	Domain Filter List	Category	Category Description
2019/04/05 11:06:48	dns	10.1.100.18	cpe-98-25-53-166.sc.res.rr.com	A	1	Domain was blocked by dns botnet C&C			
2019/04/05 11:06:48	dns	10.1.100.18	cpe-98-25-53-166.sc.res.rr.com	A	1				

To check the DNS Filter log in the CLI:

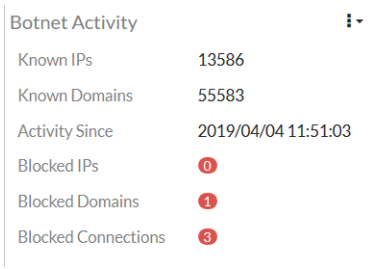
```
1: date=2019-04-05 time=11:06:48 logid="1501054600" type="utm" subtype="dns" eventtype="dns-
response" level="warning" vd="vdom1" eventtime=1554487606 policyid=1 sessionid=55232
srcip=10.1.100.18 srcport=60510 srcintf="port10" srcintfrole="undefined" dstip=172.16.95.16
dstport=53 dstintf="port9" dstintfrole="undefined" proto=17 profile="demo" xid=16265
qname="cpe-98-25-53-166.sc.res.rr.com" qtype="A" qtypeval=1 qclass="IN"
ipaddr="93.184.216.34" msg="Domain was blocked by dns botnet C&C" action="redirect"
botnetip=98.25.53.166

2: date=2019-04-05 time=11:06:48 logid="1500054000" type="utm" subtype="dns" eventtype="dns-
query" level="information" vd="vdom1" eventtime=1554487606 policyid=1 sessionid=55232
srcip=10.1.100.18 srcport=60510 srcintf="port10" srcintfrole="undefined" dstip=172.16.95.16
dstport=53 dstintf="port9" dstintfrole="undefined" proto=17 profile="demo" xid=16265
qname="cpe-98-25-53-166.sc.res.rr.com" qtype="A" qtypeval=1 qclass="IN"
```

To check botnet activity:

1. Go to *Dashboard* > *Status* and see the *Botnet Activity* widget.
   If you cannot find the *Botnet Activity* widget, click the *Settings* button at the bottom right, select *Add Widget*, and add the *Botnet Activity* widget.

Netskope Exhibit 1017



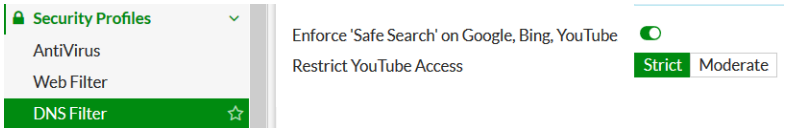
DNS safe search

Enable DNS Filter safe search so that FortiGate responds with the search engine's children and school safe domain or IP address. Users might not be aware of this filter. Explicit contents are filtered by the search engine itself. This feature isn't 100% accurate but it can help you avoid explicit and inappropriate search results.

This feature currently supports Google, Bing, and YouTube.

To configure DNS Filter Safe Search on GUI:

1. Go to *Security Profiles* > *DNS Filter* and edit or create a DNS Filter.
2. Enable *Enforce 'Safe search' on Google, Bing, YouTube*.
3. For *Restrict YouTube Access*, select *Strict* or *Moderate*.



To configure DNS Filter Safe Search on CLI:

```
config dnsfilter profile
    edit "demo"
        config ftgd-dns
            set options error-allow
            config filters
                edit 2
                    set category 2
                next
               ...
            end
        end
        set log-all-domain enable
        set block-botnet enable
        set safe-search enable  <<<==== DNS Filter Safe Search option
    next
end
```

Sample

To see an example of how this works, enable this option. Then from your internal network PC, use a command line tool such as dig or nslookup to do a DNS query on www.bing.com. For example:

```
# dig www.bing.com
;; ->>HEADER<<- opcode: QUERY; status: NOERROR; id: 46568
;; Flags: qr rd ra; QUERY: 1; ANSWER: 2; AUTHORITY: 0; ADDITIONAL: 0

;; QUESTION SECTION:
;; www.bing.com.                IN      A

;; ANSWER SECTION:
www.bing.com.           103     IN      CNAME   strict.bing.com.  <<<====
strict.bing.com.        103     IN      A       204.79.197.220

;; Received 67 B
;; Time 2019-04-05 14:34:52 PDT
;; From 172.16.95.16@53(UDP) in 196.0 ms
```

The DNS query for www.bing.com returns with a CNAME strict.bing.com, and A record for the CNAME. The user's web browser then connects to this address with the same search engine UI but any explicit content search is filtered out. Check the DNS Filter log for the message *DNS Safe Search enforced*.

Date/Time	Sub Type	Source	Domain Name	Query Type	Policy	Message	Domain Filter List	Category	Category Description
2019/04/05 14:34:53	dns	10.1.100.18	www.bing.com	A	1	DNS Safe Search enforced		41	Search Engines and Portals
2019/04/05 14:34:53	dns	10.1.100.18	www.bing.com	A	1				

To check the DNS Filter Safe Search log in the CLI:

```
1: date=2019-04-05 time=14:34:53 logid="1501054804" type="utm" subtype="dns" eventtype="dns-
response" level="notice" vd="vdom1" eventtime=1554500093 policyid=1 sessionid=65955
srcip=10.1.100.18 srcport=36575 srcintf="port10" srcintfrole="undefined" dstip=172.16.95.16
dstport=53 dstintf="port9" dstintfrole="undefined" proto=17 profile="demo" xid=59573
qname="www.bing.com" qtype="A" qtypeval=1 qclass="IN" ipaddr="204.79.197.220" msg="DNS Safe
Search enforced" action="pass" sscname="strict.bing.com" cat=41 catdesc="Search Engines and
Portals"

2: date=2019-04-05 time=14:34:53 logid="1500054000" type="utm" subtype="dns" eventtype="dns-
query" level="information" vd="vdom1" eventtime=1554500092 policyid=1 sessionid=65955
srcip=10.1.100.18 srcport=36575 srcintf="port10" srcintfrole="undefined" dstip=172.16.95.16
dstport=53 dstintf="port9" dstintfrole="undefined" proto=17 profile="demo" xid=59573
qname="www.bing.com" qtype="A" qtypeval=1 qclass="IN"
```

Additional information

For each search engine's safe search specifications, see its specification page:

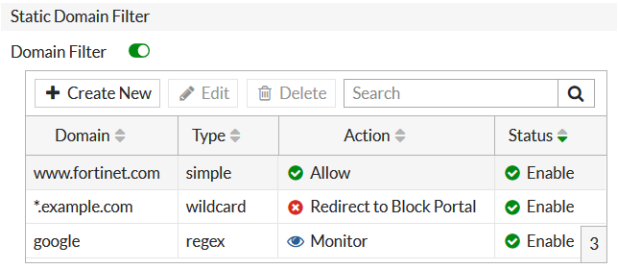
- https://help.bing.microsoft.com/#apex/18/en-US/10003/0
- https://support.google.com/websearch/answer/510?co=GENIE.Platform%3DDesktop&hl=en
- https://support.google.com/youtube/answer/174084?co=GENIE.Platform%3DDesktop&hl=en

Local domain filter

In addition to FortiGuard's category-based domain filter, you can also can define your own local static domain filter to allow or block specific domains.

To configure DNS local domain filter on GUI:

1. Go to *Security Profiles* > *DNS Filter* and edit or create a DNS Filter.
2. In the *Static Domain Filter* section, enable *Domain Filter*.
3. Click *Create New* to create your local domain filter entries.



To configure DNS local domain filter on CLI:

```
config dnsfilter domain-filter
    edit 1
        set name "demo"
        set comment ''
        config entries
            edit 1
                set domain "www.fortinet.com"
                set type simple
                set action allow
                set status enable
            next
            edit 2
                set domain "*.example.com"
                set type wildcard
                set action block
                set status enable
            next
            edit 3
                set domain "google"
                set type regex
                set action monitor
                set status enable
            next
        end
    next
end
```

Wildcard entries are converted to regular expressions by FortiOS. As a result of this conversion, wildcards will match any suffix, as long as there is a word boundary following the search term.

For example:

```
config entries
    edit 1
        set domain "*.host"
        set type wildcard
    next
end
```

will match `wp36.host` and `wp36.host.pressdns.com`, but not `wp36.host123.pressdnds.com`.

To avoid this, use an explicit regular expression search string:

```
config entries
    edit 1
        set domain "^.*\\.host$"
        set type regexp
    next
end
```

To check the DNS local domain filter log in the GUI:

1. Go to *Log & Report* > *DNS Query* to view the DNS query log.

Date/Time	Sub Type	Source	Domain Name	Query Type	Policy	Message	Domain Filter List	Category	Category Description
2019/04/05 15:37:06	dns	10.1.100.18	www.google.com	A	1	Domain belongs to a denied category in policy		41	Search Engines and Por
2019/04/05 15:37:06	dns	10.1.100.18	www.google.com	A	1				
2019/04/05 15:36:59	dns	10.1.100.18	www.example.com	A	1	Domain was blocked because it is in the domain-filter list	demo		
2019/04/05 15:36:59	dns	10.1.100.18	www.example.com	A	1				
2019/04/05 15:36:51	dns	10.1.100.18	www.fortinet.com	A	1	Domain was allowed because it is in the domain-filter list	demo		
2019/04/05 15:36:51	dns	10.1.100.18	www.fortinet.com	A	1				

Since the local domain list "google" action is Monitor, it's blocked by FortiGuard category-based domain filter.

To check the DNS local domain filter log in the CLI:

```
7: date=2019-04-05 time=15:37:06 logid="1501054803" type="utm" subtype="dns" eventtype="dns-response" level="warning" vd="vdom1" eventtime=1554503826 policyid=1 sessionid=69132 srcip=10.1.100.18 srcport=49832 srcintf="port10" srcintfrole="undefined" dstip=172.16.95.16 dstport=53 dstintf="port9" dstintfrole="undefined" proto=17 profile="demo" xid=4612 qname="www.google.com" qtype="A" qtypeval=1 qclass="IN" ipaddr="208.91.112.55" msg="Domain belongs to a denied category in policy" action="redirect" cat=41 catdesc="Search Engines and Portals"

8: date=2019-04-05 time=15:37:06 logid="1500054000" type="utm" subtype="dns" eventtype="dns-query" level="information" vd="vdom1" eventtime=1554503826 policyid=1 sessionid=69132 srcip=10.1.100.18 srcport=49832 srcintf="port10" srcintfrole="undefined" dstip=172.16.95.16 dstport=53 dstintf="port9" dstintfrole="undefined" proto=17 profile="demo" xid=4612 qname="www.google.com" qtype="A" qtypeval=1 qclass="IN"

9: date=2019-04-05 time=15:36:59 logid="1501054400" type="utm" subtype="dns" eventtype="dns-response" level="warning" vd="vdom1" eventtime=1554503818 policyid=1 sessionid=69121
```

```
srcip=10.1.100.18 srcport=40659 srcintf="port10" srcintfrole="undefined" dstip=172.16.95.16
dstport=53 dstintf="port9" dstintfrole="undefined" proto=17 profile="demo" xid=24730
qname="www.example.com" qtype="A" qtypeval=1 qclass="IN" msg="Domain was blocked because it
is in the domain-filter list" action="redirect" domainfilteridx=1 domainfilterlist="demo"

10: date=2019-04-05 time=15:36:59 logid="1500054000" type="utm" subtype="dns"
eventtype="dns-query" level="information" vd="vdom1" eventtime=1554503818 policyid=1
sessionid=69121 srcip=10.1.100.18 srcport=40659 srcintf="port10" srcintfrole="undefined"
dstip=172.16.95.16 dstport=53 dstintf="port9" dstintfrole="undefined" proto=17
profile="demo" xid=24730 qname="www.example.com" qtype="A" qtypeval=1 qclass="IN"

11: date=2019-04-05 time=15:36:51 logid="1501054401" type="utm" subtype="dns"
eventtype="dns-response" level="information" vd="vdom1" eventtime=1554503810 policyid=1
sessionid=69118 srcip=10.1.100.18 srcport=33461 srcintf="port10" srcintfrole="undefined"
dstip=172.16.95.16 dstport=53 dstintf="port9" dstintfrole="undefined" proto=17
profile="demo" xid=53801 qname="www.fortinet.com" qtype="A" qtypeval=1 qclass="IN"
ipaddr="13.56.55.78, 54.183.57.55" msg="Domain was allowed because it is in the domain-
filter list" action="pass" domainfilteridx=1 domainfilterlist="demo"

12: date=2019-04-05 time=15:36:51 logid="1500054000" type="utm" subtype="dns"
eventtype="dns-query" level="information" vd="vdom1" eventtime=1554503810 policyid=1
sessionid=69118 srcip=10.1.100.18 srcport=33461 srcintf="port10" srcintfrole="undefined"
dstip=172.16.95.16 dstport=53 dstintf="port9" dstintfrole="undefined" proto=17
profile="demo" xid=53801 qname="www.fortinet.com" qtype="A" qtypeval=1 qclass="IN"
```

Sequence and priority

In DNS Filter, local domain filter has a higher priority than FortiGuard category-based domain filter.

A DNS query is scanned and matched with local domain filter first. If an entry matches and the local filter entry's action is block, then that DNS query is blocked or redirected.

If local domain filter list has no match, then the FortiGuard category-based domain filter is used. If a DNS query domain name rating belongs to the block category, this query is blocked or redirected. If the FortiGuard category-based filter has no match, then the original resolved IP address is returned to the client DNS resolver.

The local domain filter action can be Block, Allow, or Monitor. If the local domain filter action is Allow and an entry matches, it will skip the FortiGuard category-based domain filter and directly return to client DNS resolver. If the local domain filter action is Monitor and an entry matches, it will go to FortiGuard category-based domain filter scanning and matching.

DNS translation

Using this feature, you can translate a DNS resolved IP address to another IP address you specify on a per-policy basis.

For example, website A has a public address 1.2.3.4. However, when your internal network users visit this website, you want them to connect to an internal host, say, 192.168.3.4. In this case, you can use DNS translation to translate the DNS resolved address 1.2.3.4 to 192.168.3.4. Reverse use of DNS translation is also applicable, for example, if you want public DNS query of your internal server to get a public IP address, then you can translate a DNS resolved private IP to a public IP address.

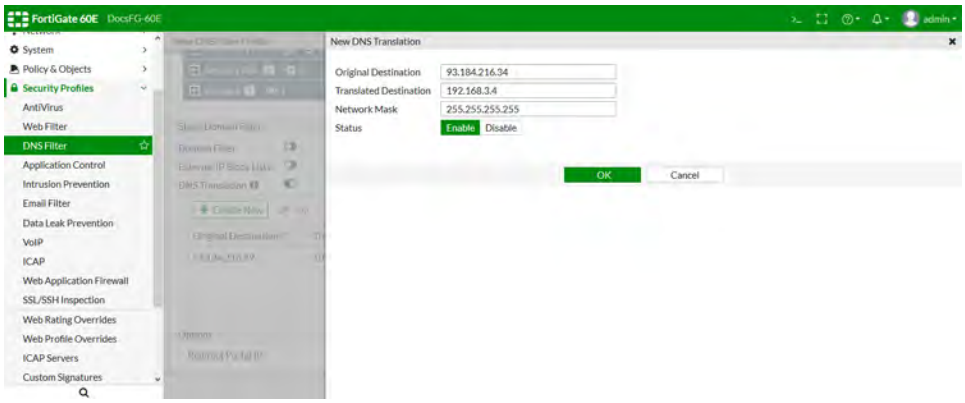
Netskope Exhibit 1017

Example

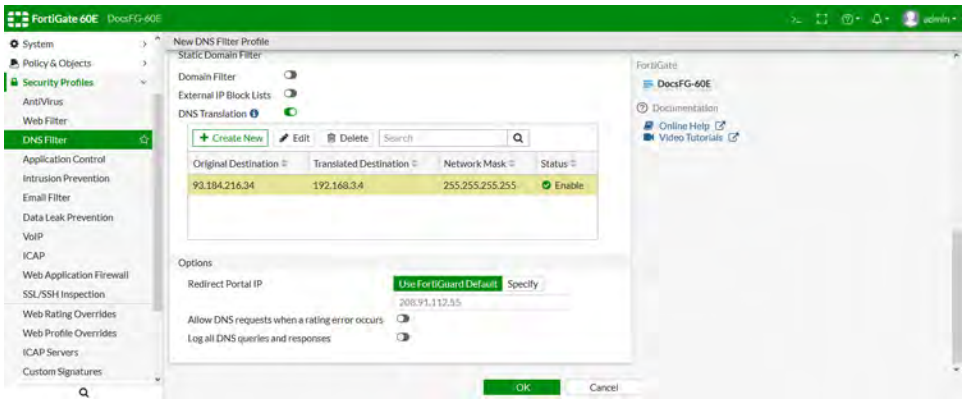
This example configuration forces the DNS Filter profile to translate 93.184.216.34 (www.example.com) to 192.168.3.4. When internal network users do a DNS query for www.example.com, they do not get the original www.example.com IP address of 93.184.216.34. Instead, it is replaced with 192.168.3.4.

To configure DNS translation in the GUI:

1. Go to *Security Profiles* > *DNS Filter* and edit or create a DNS Filter profile.
2. Enable *DNS Translation* and click *Create New*.
3. Enter the *Original Destination* (the domain's original IP address), the *Translated Destination* IP address, and the *Network Mask*, and set *Status* to *Enable*.



4. Click *OK*.



5. Click *OK* to create or edit the DNS profile.

To configure DNS translation in the CLI:

```
config dnsfilter profile
    edit "demo"
        set comment ''
        ...
        config dns-translation
            edit 1
                set src 93.184.216.34
                set dst 192.168.3.4
                set netmask 255.255.255.255
            next
```

```
        end
        set redirect-portal 0.0.0.0
        set redirect-portal6 ::
        set youtube-restrict strict
    next
end
```

To check DNS translation using a command line tool before DNS translation:

```
# dig www.example.com
;; ->>HEADER<<- opcode: QUERY; status: NOERROR; id: 27030
;; Flags: qr rd ra; QUERY: 1; ANSWER: 1; AUTHORITY: 2; ADDITIONAL: 0

;; QUESTION SECTION:
;; www.example.com.             IN      A

;; ANSWER SECTION:
www.example.com.        33946   IN      A       93.184.216.34

;; AUTHORITY SECTION:
example.com.            18578   IN      NS      b.iana-servers.net.
example.com.            18578   IN      NS      a.iana-servers.net.

;; Received 97 B
;; Time 2019-04-08 10:47:26 PDT
;; From 172.16.95.16@53(UDP) in 0.5 ms
```

To check DNS translation using a command line tool after DNS translation:

```
# dig www.example.com
;; ->>HEADER<<- opcode: QUERY; status: NOERROR; id: 62060
;; Flags: qr rd ra; QUERY: 1; ANSWER: 1; AUTHORITY: 2; ADDITIONAL: 0

;; QUESTION SECTION:
;; www.example.com.             IN      A

;; ANSWER SECTION:
www.example.com.        32491   IN      A       192.168.3.4   <<<==== resolved IP translated
into 192.168.3.4

;; AUTHORITY SECTION:
example.com.            17123   IN      NS      b.iana-servers.net.
example.com.            17123   IN      NS      a.iana-servers.net.

;; Received 97 B
;; Time 2019-04-08 11:11:41 PDT
;; From 172.16.95.16@53(UDP) in 0.5 ms
```

DNS translation network mask

The following is an example of DNS translation and result:

```
config dns-translation
        edit 1
           set src 93.184.216.34
           set dst 1.2.3.4
```

```
            set netmask 255.255.224.0
        next
    end
```

```
# dig www.example.com
;; ->>HEADER<<- opcode: QUERY; status: NOERROR; id: 6736
;; Flags: qr rd ra; QUERY: 1; ANSWER: 1; AUTHORITY: 2; ADDITIONAL: 0

;; QUESTION SECTION:
;; www.example.com.             IN      A

;; ANSWER SECTION:
www.example.com.        29322   IN      A       1.2.24.34

;; AUTHORITY SECTION:
example.com.            13954   IN      NS      a.iana-servers.net.
example.com.            13954   IN      NS      b.iana-servers.net.

;; Received 97 B
;; Time 2019-04-08 12:04:30 PDT
;; From 172.16.95.16@53(UDP) in 2.0 ms
```

1) **AND src(Orginal IP) with negative netmask (93.184.216.34 & ~255.255.224.0)**

```
01011101.10111000.11011000.00100010 93.184.216.34 <-- ip
00000000.00000000.00011111.11111111 ~255.255.224.0 <-- ~netmask
----------------------------------------------------- &
00000000.00000000.00011000.00100010 0.0.24.34 <- right bits
```

2) **AND dst(Translated IP) with netmask**

```
00000001.00000010.00000011.00000100 1.2.3.4 <- dst
11111111.11111111.11100000.00000000 255.255.224.0 <- netmask
----------------------------------------------------- &
00000001.00000010.00000000.00000000 1.2.0.0 <- left bits
```

3) **Final step 2 bitwise-OR 3:**

```
00000000.00000000.00011000.00100010 0.0.24.34
00000001.00000010.00000000.00000000 1.2.0.0
----------------------------------------------------- |
00000001.00000010.00011000.00100010 1.2.24.34
```

Using a FortiGate as a DNS server

You can configure and use FortiGate as a DNS server in your network. When you enable DNS Service on a specific interface, FortiGate will listen for DNS Service on that interface.

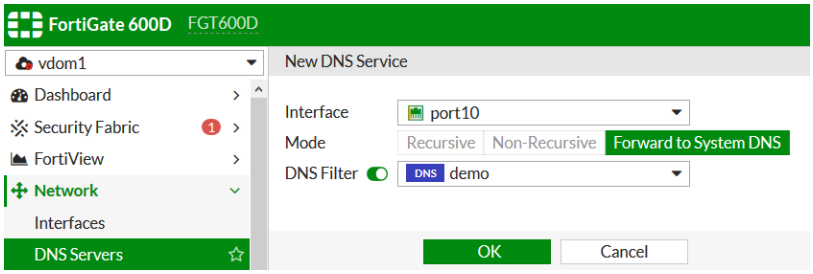
Depending on the configuration, DNS Service on FortiGate can work in three modes: *Recursive*, *Non-Recursive*, or *Forward to System DNS* (server). For details on how to configure DNS Service on FortiGate, see the FortiGate System Configuration Guide.

You can apply a DNS Filter profile to Recursive Mode and Forward to System DNS Mode. This is the same as FortiGate working as a transparent DNS Proxy for DNS relay traffic.

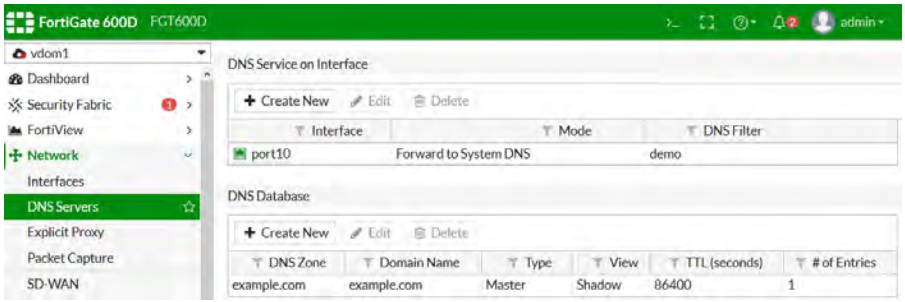
Netskope Exhibit 1017

To configure DNS Service on FortiGate using GUI:

1. Go to *Network* > *DNS Servers*.
2. In the *DNS Service on Interface*, click *Create New* and select an *Interface*.



The *Recursive* and *Non-Recursive Mode* is available only after you configure the DNS database.



To configure DNS Service on FortiGate using CLI:

```
config system dns-server
    edit "port10"   <<<==== Enable DNS Serive on Interface
        set mode forward-only
        set dnsfilter-profile "demo"   <<<==== apply DNS Filter Profile for the service
    next
end
```

Sample configuration

In this example, FortiGate port 10 is enabled as a DNS Service with the DNS Filter profile "demo". Suppose port 10 has an IP address 10.1.100.5 and DNS Filter profile "demo" is set to block category 52 (Information Technology), then from your internal network PC, use a command line tool such as dig or nslookup to do a DNS query. For example:

```
# dig @10.1.100.5 www.fortinet.com    <<<====Specify FortiGate interface address as DNS
Server
;; ->>HEADER<<- opcode: QUERY; status: NOERROR; id: 52809
;; Flags: qr rd; QUERY: 1; ANSWER: 1; AUTHORITY: 0; ADDITIONAL: 0

;; QUESTION SECTION:
;; www.fortinet.com.            IN      A

;; ANSWER SECTION:
www.fortinet.com.       60      IN      A       208.91.112.55   <<<==== DNS Filter profile
will filter the relay DNS traffic based on profile configuration. It blocked with redirect
portal IP

;; Received 50 B
```


Netskope Exhibit 1017

```
;; Time 2019-04-08 14:36:34 PDT
;; From 10.1.100.5@53(UDP) in 13.6 ms
```

Troubleshooting for DNS filter

If you have trouble with the DNS Filter profile in your policy, start with the following troubleshooting steps:

- Check the connection between FortiGate and FortiGuard DNS rating server (SDNS server).
- Check that FortiGate has a valid FortiGuard Web Filter license.
- Check the FortiGate DNS Filter configuration.

Troubleshooting connection between FortiGate and FortiGuard SDNS server

Ensure FortiGate can connect to the FortiGuard SDNS server. By default, FortiGate uses UDP port 53 to connect to the SDNS server.

To check the connection between FortiGate and the SDNS server in the CLI:

1. In the CLI Console, run the `diagnose test application dnsproxy 3` command to find the FortiGuard SDNS server.
2. Check the `FDG_SERVER` line. The SDNS server IP address might be different depending on location. For this example, it is:

   ```
   FDG_SERVER:208.91.112.220:53
   ```

3. In the CLI Console under the management VDOM, run the command `execute ping 208.91.112.220` to check the communication between the FortiGate and the SDNS server.
4. Optionally, you can also check the communication using a PC on the internal network.
   a. Disable the DNS Filter profile so that it does not affect your connection check.
   b. Ping your ISP or a public DNS service provides's DNS server, for example, Google's public DNS server of 8.8.8.8:

      ```
      #dig @8.8.8.8 www.fortinet.com
      ```

      Or specify the SDNS server as DNS server:

      ```
      #dig @208.91.112.220 www.fortinet.com
      ```

   c. Check that you can get domain www.fortinet.com A record from the DNS server which shows that UDP port 53 connection path is not blocked.

      ```
      #dig @8.8.8.8 www.fortinet.com
      ;; ->>HEADER<<- opcode: QUERY; status: NOERROR; id: 35121
      ;; Flags: qr rd ra; QUERY: 1; ANSWER: 3; AUTHORITY: 0; ADDITIONAL: 0

      ;; QUESTION SECTION:
      ;; www.fortinet.com.            IN      A

      ;; ANSWER SECTION:
      www.fortinet.com.       289     IN      CNAME   fortinet-prod4-858839915.us-west-
      1.elb.amazonaws.com.
      fortinet-prod4-858839915.us-west-1.elb.amazonaws.com. 51        IN      A
      52.8.142.247
      fortinet-prod4-858839915.us-west-1.elb.amazonaws.com. 51        IN      A
      ```


Netskope Exhibit 1017

```
13.56.55.78

;; Received 129 B
;; Time 2019-04-29 14:13:18 PDT
;; From 8.8.8.8@53(UDP) in 13.2 ms
```

Checking FortiGuard DNS Rating Service license

The FortiGuard DNS Rating Service shares the license with FortiGuard Web Filter so you must have a valid Web Filter license for the DNS Rating Service to work. While the license is shared, the DNS Rating Service uses a separate connection mechanism from the Web Filter Rating.

To check the DNS Rating Service license in the CLI:

1. In the CLI Console, run the command `diagnose test application dnsproxy 3`.
2. Look for the `LICENSE` line and check that the license has not expired, for example:
   `LICENSE: expiry=2029-08-21, expired=0, type=2`
3. Check the `dns-server` lines. Some `dns-server` lines show `secure=1 ready=1`. These lines show the functioning SDNS servers. For example:
   `dns-server:208.91.112.220:53 tz=-480 req=7 to=0 res=7 rt=1` **`secure=1 ready=1`** `timer=0 probe=0 failure=0 last_failed=0`

Checking FortiGate DNS Filter profile configuration

To check the FortiGate DNS Filter profile configuration:

1. Create a local domain filter and set the *Action* to *Redirect to Block Portal*.
   See Local domain filter on page 953.
2. Apply this DNS Filter profile to the policy.
3. From the client PC, DNS query this domain.

If you get the profile's redirected portal address, that shows that the DNS Filter profile works as expected.

More troubleshooting steps

To reload the DNS proxy in the CLI:

```
(global)#diagnose test application dnsproxy 99
```

To debug DNS proxy details:

These commands might create more output in your console.

```
#diagnose debug application dnsproxy -1

#diagnose debug enable/disable
```

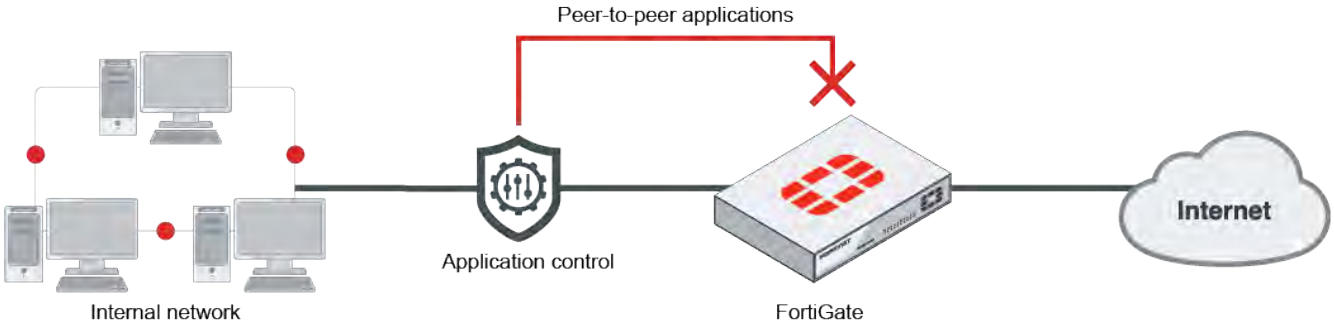
DNS proxy command reference

Use the following `diagnose test application dnsproxy` command line options to check DNS proxy status and help with troubleshooting.

```
(global) # diagnose test application dnsproxy ?
worker idx: 0
1. Clear DNS cache
2. Show stats
3. Dump DNS setting
4. Reload FQDN
5. Requery FQDN
6. Dump FQDN
7. Dump DNS cache
8. Dump DNS DB
9. Reload DNS DB
10. Dump secure DNS policy/profile
11. Dump Botnet domain
12. Reload Secure DNS setting
13. Show Hostname cache
14. Clear Hostname cache
15. Show SDNS rating cache
16. Clear SDNS rating cache
17. DNS debug bit mask
99. Restart dnsproxy worker
```

Application control

FortiGates can recognize network traffic generated by a large number of applications. Application control sensors specify what action to take with the application traffic. Application control uses IPS protocol decoders that can analyze network traffic to detect application traffic, even if the traffic uses non-standard ports or protocols. Application control supports traffic detection using the HTTP protocol (versions 1.0, 1.1, and 2.0).



FortiOS includes three preloaded application sensors:

- *default* (monitors all applications)
- *wifi-default* (default configuration for offloading WiFi traffic)
- *block-high-risk*

You can customize these sensors, or you can create your own to log and manage the applications on your network.

Once configured, you can add the application sensor to a firewall policy.



This functionality requires a subscription to FortiGuard Application Control.

The following topics provide information about application control:

- Basic category filters and overrides on page 964
- Port enforcement check on page 967
- Protocol enforcement on page 968
- SSL-based application detection over decrypted traffic in a sandwich topology on page 970
- Matching multiple parameters on application control signatures on page 971

Basic category filters and overrides

Once you have created an application sensor, you can define the applications that you want to control. You can add applications and filters using categories, application overrides, and/or filter overrides.

- **Categories**: Choose groups of signatures based on a category type.
- **Application overrides**: Choose individual applications.
- **Filter overrides**: Select groups of applications and override the application signature settings for them.

Categories

Categories allow you to choose groups of signatures based on a category type. Applications belonging to the category trigger the action that is set for the category.

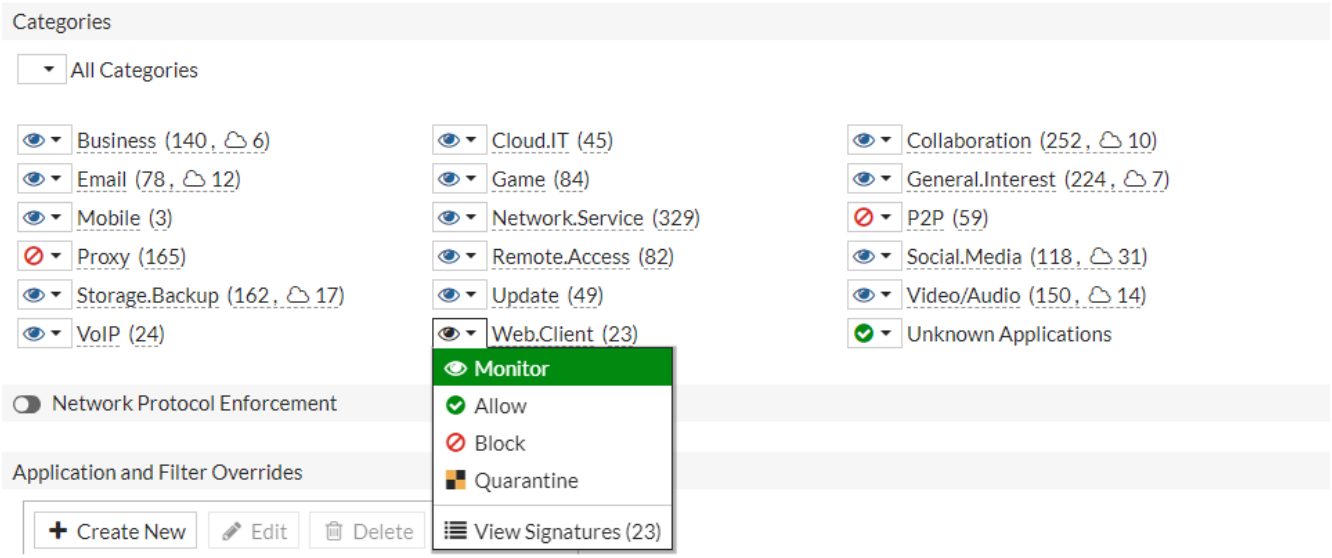
To set category filters in the CLI:

```
config application list
    edit {id}
        config entries
            edit 1
                set category <id>
                    ID          Select Category ID
                    2           P2P
                    3           VoIP
                    5           Video/Audio
                    6           Proxy
                    7           Remote.Access
                    8           Game
                    12          General.Interest
                    15          Network.Service
                    17          Update
                    21          Email
                    22          Storage.Backup
                    23          Social.Media
                    25          Web.Client
                    26          Industrial
                    28          Collaboration
                    29          Business
                    30          Cloud.IT
```

```
                31           Mobile
            set action {pass | block | reset}
                pass     Pass or allow matching traffic.
                block    Block or drop matching traffic.
                reset    Reset sessions for matching traffic.
            set log {enable | disable}
        next
    end
next
end
```

To set category filters in the GUI:

1. Go to *Security Profiles* > *Application Control*.
2. Under *Categories*, left click the icon next to the category name to view a dropdown of actions:
   - Allow
   - Monitor
   - Block
   - Quarantine
   - View signatures
3. Select *OK*.



Application and filter overrides

Override type	Setting
Application	**Type**: Choose *Application* for application overrides.
	Action: Can be set to Monitor/Allow/Block/Quarantine.
	Application: Multiple app signatures can be added for one entry. A slide-in presenting an application list will be shown to select specific app signatures, and the search box can be used to filter matched signatures.

Override type	Setting
Filter	**Type**: Choose *Filter* for filter overrides.
	Action: Can be set to Monitor/Allow/Block/Quarantine.
	Filter: Filters can be selected by behavior, application category, technology, popularity, protocol, risk, or vendor subtypes.
	Search box: Can be used to determine if the input signature is included in selected filters, where matched applications are shown at the bottom.

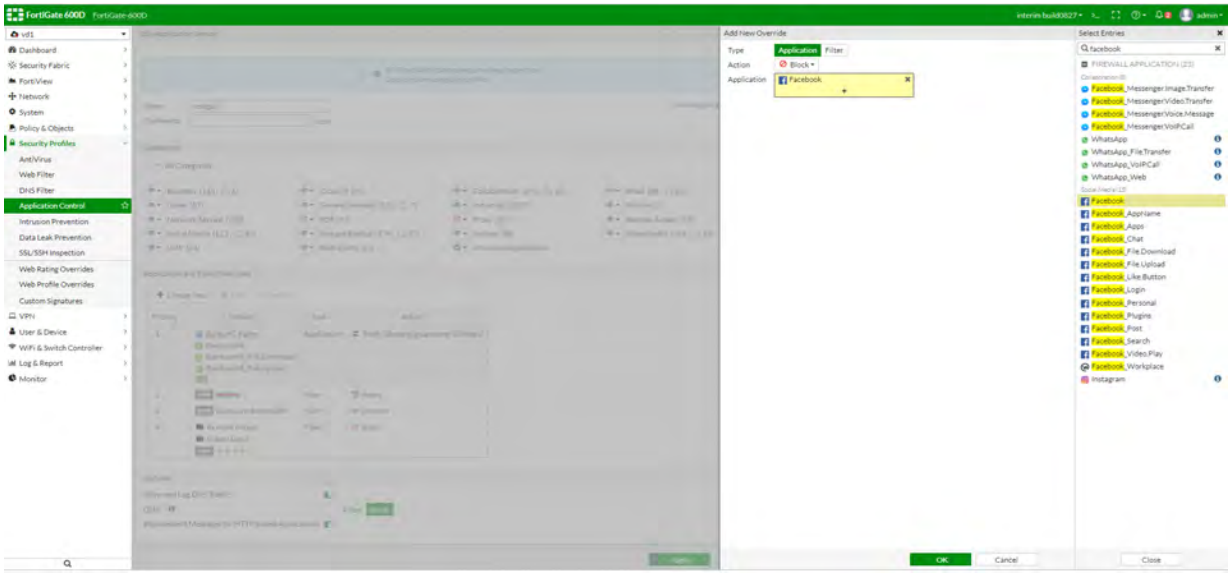
To set overrides in the CLI:

```
config application list
    edit {id}
        config entries
            edit 1
                set protocols <0-47>    #network protocol ID
                set risk <id>
                    *level    Risk of allowing traffic from this application to occur (1 -
5; Low, Elevated, Medium, High, and Critical).
                set vendor <0-25>       #vendor ID
                set technology <id>
                    All          All
                    0            Network-Protocol
                    1            Browser-Based
                    2            Client-Server
                    4            Peer-to-Peer
                set behavior <id>
                    All          All
                    2            Botnet
                    3            Evasive
                    5            Excessive-Bandwidth
                    6            Tunneling
                    9            Cloud
                set popularity <1-5>    #Popularity level 1-5
                set action {pass | block | reset}
                    pass     Pass or allow matching traffic.
                    block    Block or drop matching traffic.
                    reset    Reset sessions for matching traffic.
                set log {enable | disable}
            next
        end
    next
end
```

To set overrides in the GUI:

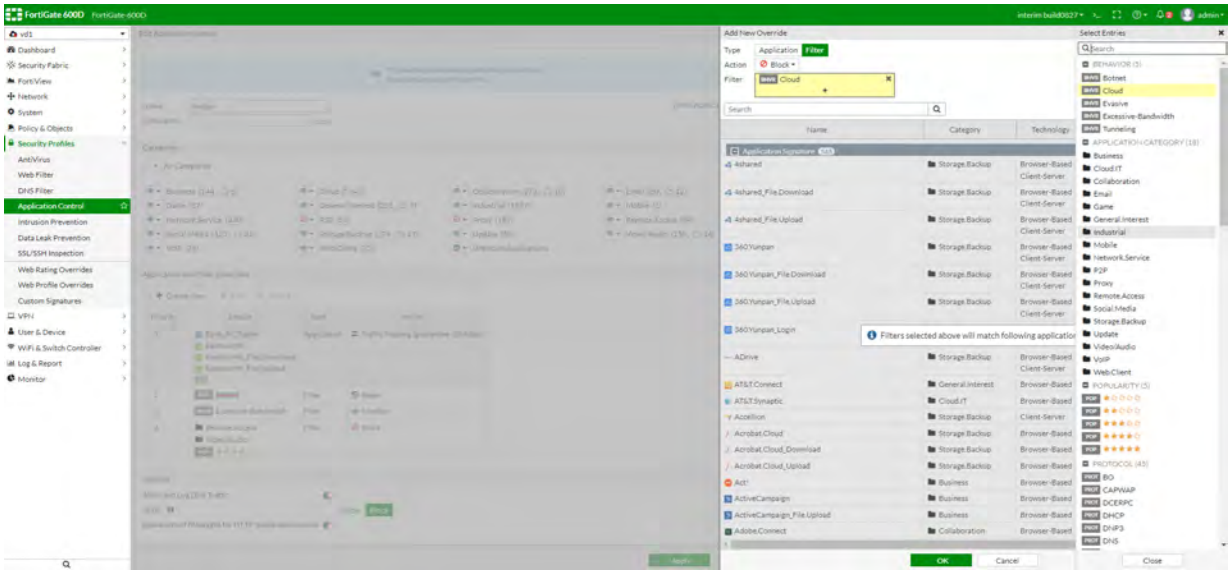
1. Go to *Security Profiles* > *Application Control*.
2. Under the *Application and Filter Overrides* table, click *Create New*.
3. To add individual applications:
   a. Select *Application* as the *Type*.
   b. Choose an action to be associated with the application.

c. Click the + button in the *Application* field and choose the specific applications from the list where app signatures are displayed. Multiple applications may be selected.



d. Click *OK*.

4. To add advanced filters:
   a. Create another entry in the *Application and Filter Overrides* table.
   b. Select *Filter* as the *Type*.
   c. Select *Cloud* under the behavior section from the *Select Entries* list.
   Matched signatures are shown along the bottom.



   d. Click *OK*.

Port enforcement check

Most networking applications run on specific ports. For example, SSH runs on port 22, and Facebook runs on ports 80 and 443.

If the default network service is enabled in the Application Control profile, a port enforcement check is done at the application profile level, and any detected application signatures running on the non-standard TCP/IP port are blocked. This means that each application allowed by the app control sensor is only run on its default port.

To set port enforcement check in the CLI:

```
config application list
    edit "default_port"
        set enforce-default-app-port {enable | disable}
            disable        Disable default application port enforcement.
            enable         Enable default application port enforcement.
        config entries
            edit 1
                set application 15896
                set action pass
            next
        end
    next
end
```

For example, when applying the above appctrl sensor, FTP traffic with the standard port (port 21) is allowed, while the non-standard port (port 2121) is blocked.

Protocol enforcement

Protocol enforcement allows you to configure networking services (e.g. FTP, HTTP, HTTPS) on known ports (e.g. 21, 80, 443). For protocols that are not allowlisted under select ports, the IPS engine performs the violation action to block, allow, or monitor that traffic.

This feature can be used in the following scenarios:

- When one protocol dissector confirms the service of network traffic, protocol enforcement can check whether the confirmed service is allowlisted under the server port. If it is not allowlisted, the traffic is considered a violation and IPS can take the action specified in the configuration (block or monitor it).
- When there is no confirmed service for the network traffic, the traffic is considered a service violation if IPS dissectors rule out all of the services enforced under its server port.

In an applicable profile, a default-network-service list can be created to associate well known ports with accepted services.

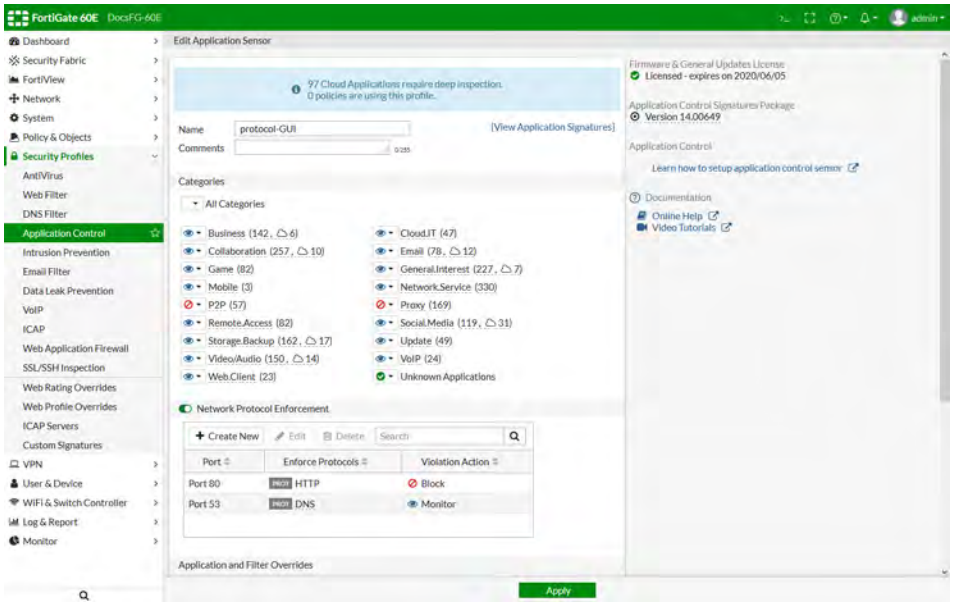
To setup protocol enforcement in the CLI:

```
config application list
   edit "protocol-GUI"
      set other-application-log enable
      set control-default-network-services {enable | disable}   # Enable/Disable enforcement
of protocols over select ports
      config default-network-services                           # Default network service
entries
         edit 1
           set port 80                        # Port number, enter an integer value from <0>
to <65535>
           set services http                  # Network protocols: http, ssh, ftp, dns,
smtp, pop3, imap, snmp, nntp, and https
```

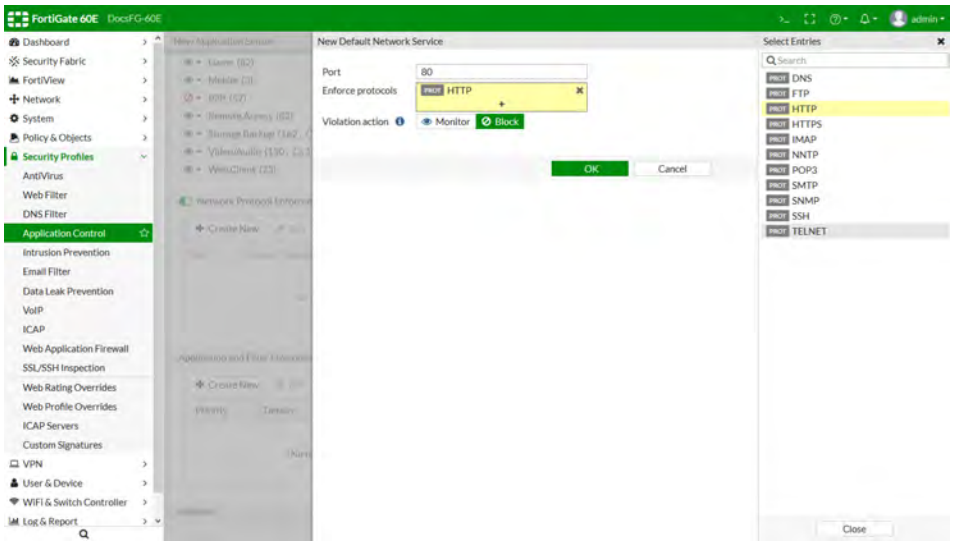
```
        next
        edit 2
          set port 53
          set services dns
          set violation-action {pass | monitor | block}      # Pass, Log, or block when
non-DNS traffic run over port 53
        next
      end
   next
end
```

To setup protocol enforcement in the GUI:

1. Go to *Security Profiles* > *Application Control*.
2. Create a new application sensor or edit an existing one.
3. Enable *Network Protocol Enforcement*.
   Enforcement entries can be created, edited, or deleted to configure network services on certain ports and determine the violation action.



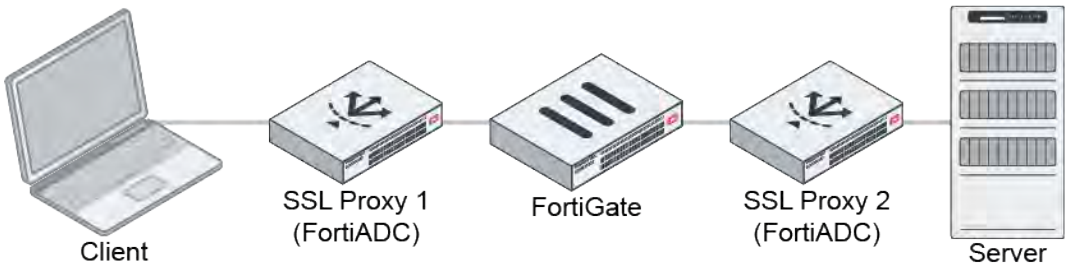
4. Click *Create New* in the *Network Protocol Enforcement* table.



5. In the *New Default Network Service* pane:
   a. Enter a *Port* number.
   b. Select *Enforced protocols*.
   c. Choose the *Violation action*.
   d. Click *OK*.
6. Click *OK*.

SSL-based application detection over decrypted traffic in a sandwich topology

When a FortiGate is sandwiched between SSL encryption and decryption devices, the FortiGate can process the decrypted traffic that passes between those devices. This feature adds support for decrypted traffic in application control. In some pre-defined signatures, the signature is pre-marked with the *require_ssl_di* tag. The `force-inclusion-ssl-di-sigs` option under `application list` allows users to control the inspection of dissected traffic. When this option is enabled, the IPS engine forces the pre-marked SSL-based signatures to be applied to the decrypted traffic of the respective applications. In the following topology, SSL Proxy 1 handles the client connection and SSL Proxy 2 handles the server connection, leaving the content unencrypted as traffic passes through the FortiGate.



To configure SSL-based application detection over decrypted traffic:

```
config application list
    edit "test"
        set force-inclusion-ssl-di-sigs {enable | disable}
```

```
    next
end
```

Example pre-marked SSL-based signature:

F-SBID(--vuln_id 15722; --attack_id 42985; --name "Facebook_Chat"; --group im; --protocol tcp; --default_action pass; --revision 4446; --app_cat 23; --vendor 3; --technology 1; --behavior 9; --pop 4; --risk 2; --language "Multiple"; --weight 20; --depend-on 15832; --depend-on 38468; ***--require_ssl_di "Yes"****; --casi 1; --casi 8; --parent 15832; --app_port "TCP/443"; --severity info; --status hidden; --service http; --flow from_client; --pattern "/pull?"; --context uri; --no_case; --pattern ".facebook.com"; --context host; --no_case; --tag set,Tag.Facebook.Pull; --tag quiet; --scan-range 10m,all; --date 20190301;)*

All signatures that include the *require_ssl_di* tag are pre-defined and cannot be customized.

Matching multiple parameters on application control signatures

Application control signatures that support parameters (such as SCADA protocols) can have multiple parameters grouped together and matched at the same time. Multiple application parameter groups can be added to an override. Traffic will be flagged if it matches at least one parameter group.

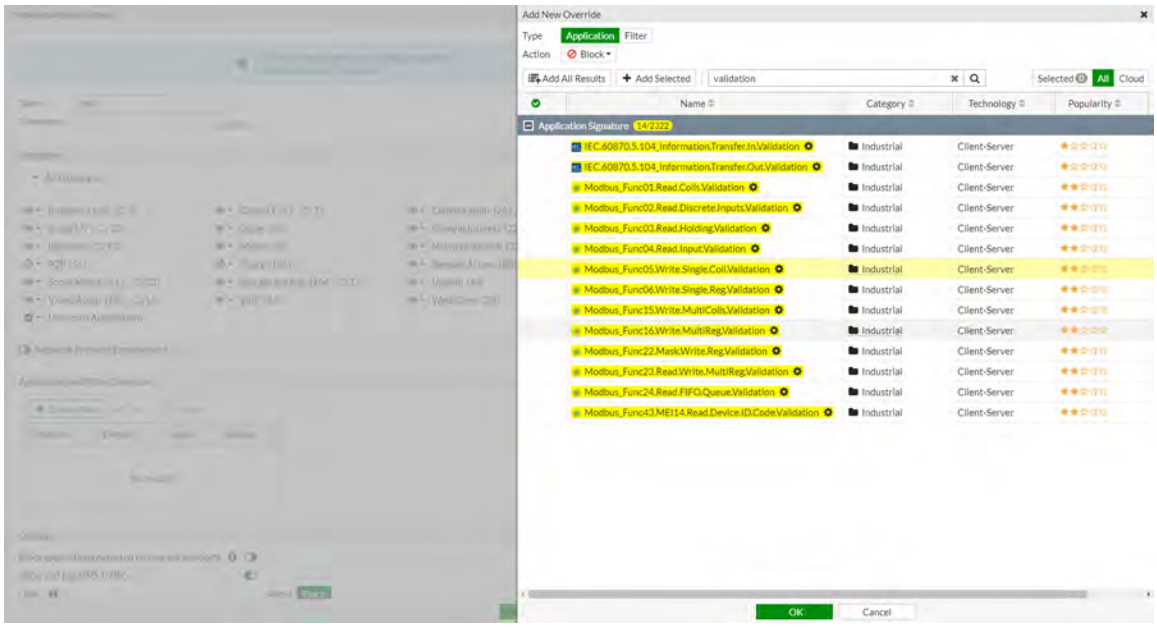
This example uses the Modbus_Func05.Write.Single.Coil.Validation signature. This is an industrial signature, so ensure that no signatures are excluded:

```
config ips global
    set exclude-signatures none
end
```

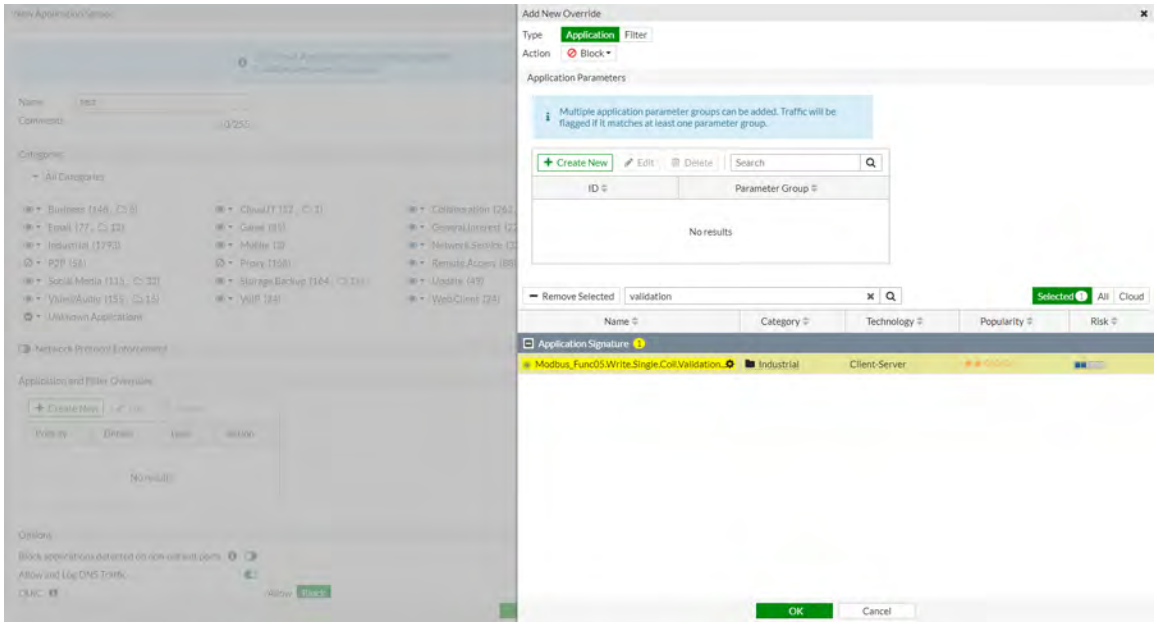
To configure an application sensor with multiple parameters in the GUI:

1. Go to *Security Profiles* > *Application Control* and click *Create New*, or edit an existing sensor.
2. In the *Application and Filter Overrides* table, click *Create New*.
3. Search for *Modbus_Func05.Write.Single.Coil.Validation* and press `Enter`. A gear icon beside the signature name indicates it has configurable application parameters.

4. In the search results, select *Modbus_Func05.Write.Single.Coil.Validation* and click *Add Selected*.

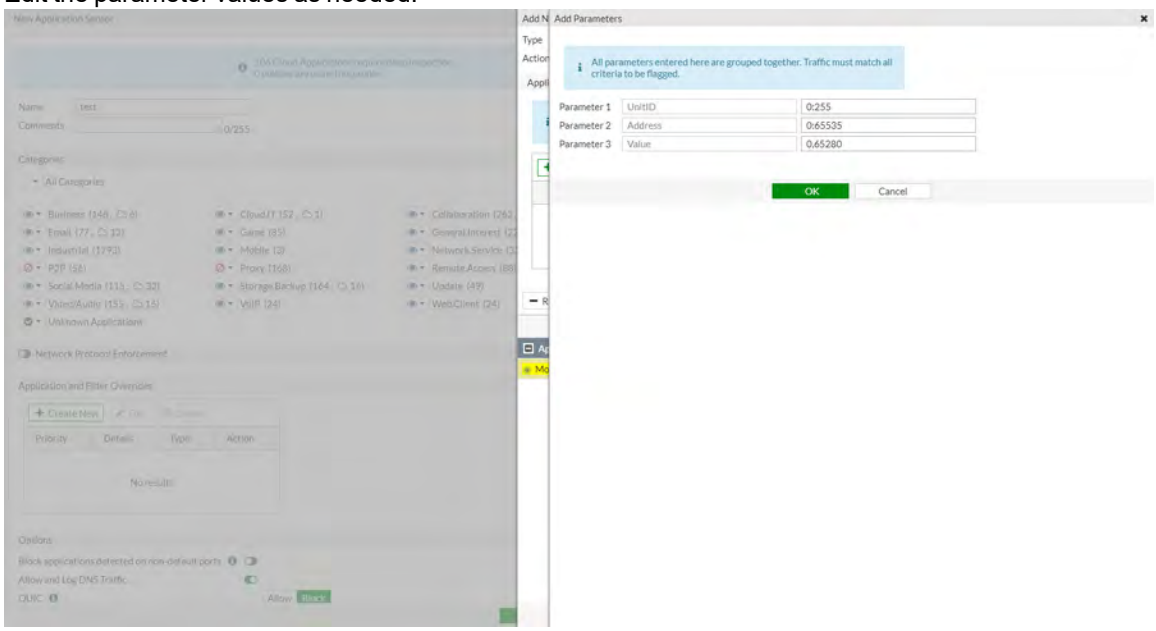


5. Click the *Selected* tab. In the *Application Parameters* section, click *Create New*.



Netskope Exhibit 1017

6. Edit the parameter values as needed.



7. Click *OK*.
8. Add more signatures if needed.
9. Click *OK*.

To configure an application sensor with multiple parameters in the CLI:

```
config application list
    edit "test"
        set other-application-log enable
        config entries
            edit 1

                set application 48885
                config parameters
                    edit 1
                        config members
                            edit 1
                                set name "UnitID"
                                set value "0:255"
                            next
                            edit 2
                                set name "Address"
                                set value "0:65535"
                            next
                            edit 3
                                set name "Value"
                                set value "0,65280"
                            next
                        end
                    next
                end
            next
            edit 2
```

```
                set category 2 6
            next
        end
    next
end
```

Intrusion prevention

With the FortiOS intrusion prevention system (IPS), you can detect and block network-based attacks. You can configure IPS sensors based on IPS signatures, IPS filters, outgoing connections to botnet sites, and rate-based signatures.

FortiOS includes eight preloaded IPS sensors:

- *all_default*
- *all_default_pass*
- *default*
- *high_security*
- *protect_client*
- *protect_email_server*
- *protect_http_server*
- *wifi-default*

You can customize these sensors, or you can create your own and apply it to a firewall policy.



This functionality requires a subscription to FortiGuard IPS Service.

The following topic provides information about IPS sensors:

- Botnet C&C IP blocking on page 974
- Detecting IEC 61850 MMS protocol in IPS on page 978

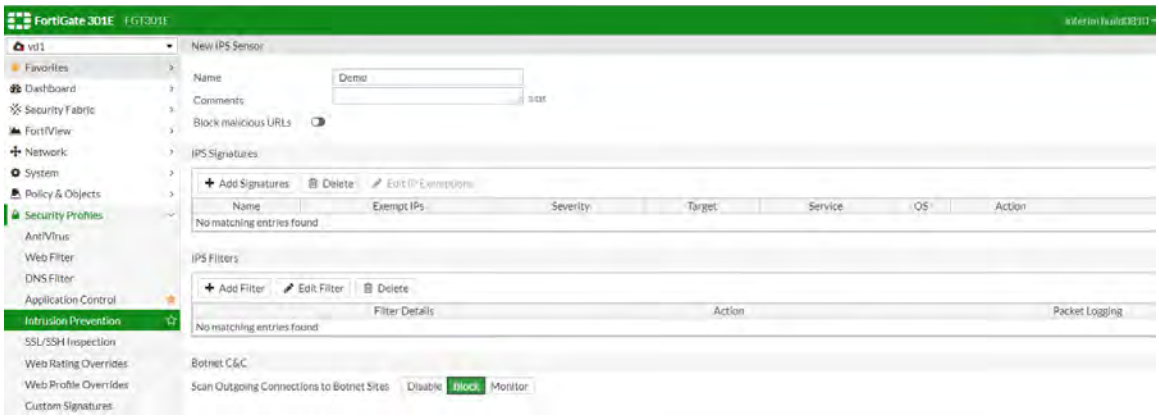
Botnet C&C IP blocking

The *Botnet C&C* section consolidates multiple botnet options in the IPS profile. This allows you to enable botnet blocking across all traffic that matches the policy by configuring one setting in the GUI, or by the `scan-botnet-connections` option in the CLI.

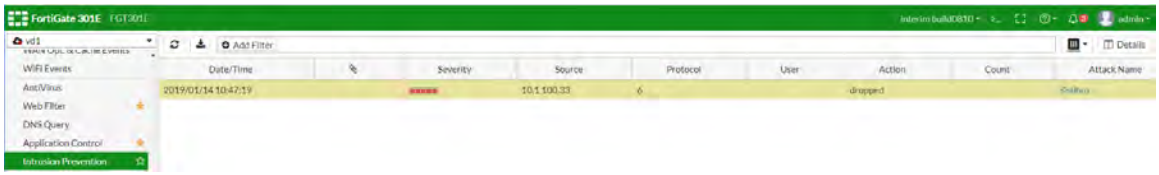
To configure botnet C&C IP blocking using the GUI:

1. Go to *Security Profiles* > *Intrusion Prevention*.
2. Edit an existing sensor, or create a new one.
3. Navigate to the *Botnet C&C* section.

4. For *Scan Outgoing Connections to Botnet Sites*, click *Block* or *Monitor*.



5. Configure other settings as needed.
6. Click *Apply*. Botnet C&C is now enabled for the sensor.
7. Add this sensor to the firewall policy.
   The IPS engine will scan outgoing connections to botnet sites. If you access a botnet IP, an IPS log is generated for this attack.
8. Go to *Log & Report* > *Intrusion Prevention* to view the log.



To configure botnet C&C IP blocking using the CLI:

```
config ips sensor
   edit "Demo"
      set scan-botnet-connections <disable | block | monitor>
   next
end
```

The `scan-botnet-connections` option is no longer available in the following CLI commands:

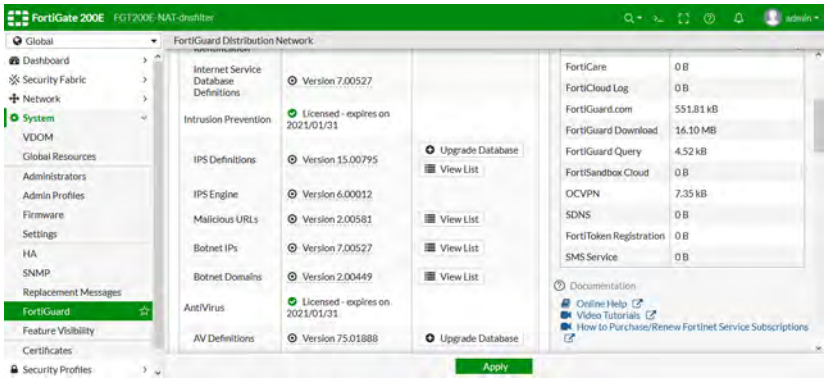
- `config firewall policy`
- `config firewall interface-policy`
- `config firewall proxy-policy`
- `config firewall sniffer`

Netskope Exhibit 1017

Botnet IPs and domains lists

To view botnet IPs and domains lists using the GUI:

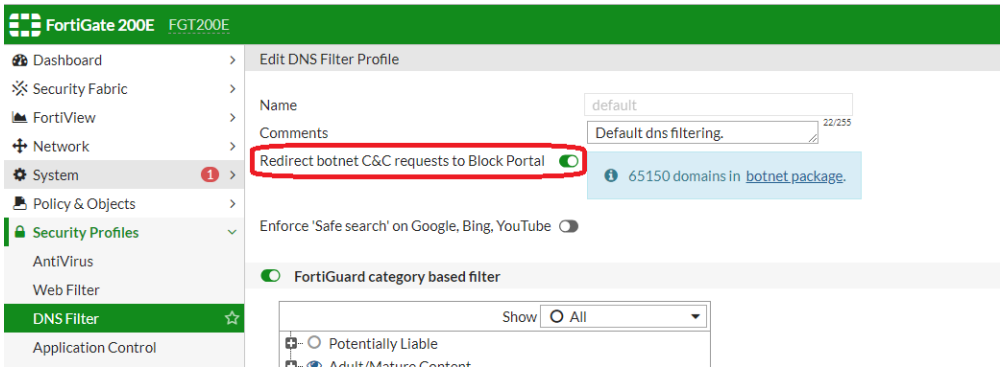
1. Go to *System* > *FortiGuard*. *Botnet IPs* and *Botnet Domains* are visible in the *Intrusion Prevention* section.
2. Click *View List* for more details.



Botnet C&C domain blocking

To block connections to botnet domains using the GUI:

1. Go to *Security Profiles* > *DNS Filter*.
2. Edit an existing filter, or create a new one.
3. Enable *Redirect botnet C&C requests to Block Portal*.



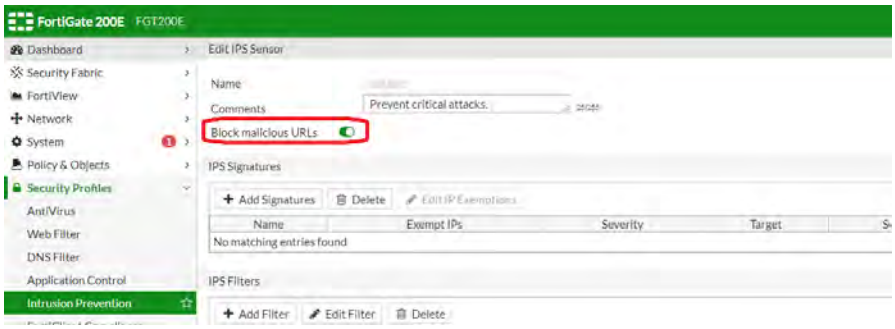
4. Configure other settings as needed.
5. Click *OK*.
6. Add this filter profile to a firewall policy.

Botnet C&C URL blocking

To block malicious URLs using the GUI:

1. Go to *Security Profiles* > *Intrusion Prevention*.
2. Edit an existing sensor, or create a new one.

3. Enable *Block malicious URLs*.

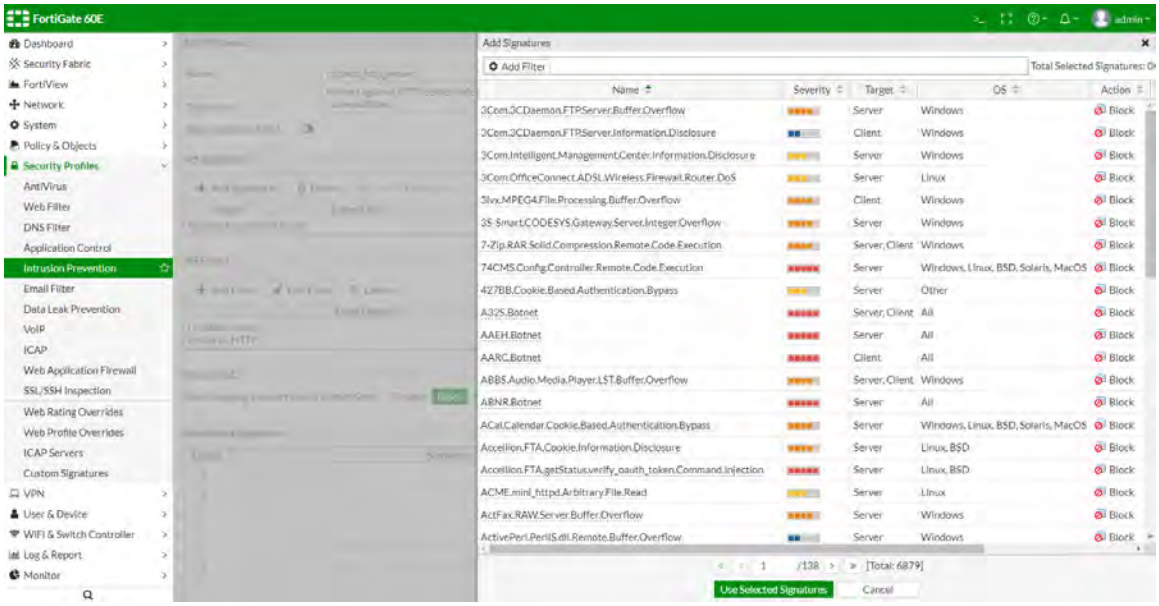


4. Configure other settings as needed.
5. Click *OK*.
6. Add this sensor to a firewall policy.

Botnet C&C signature blocking

To add IPS signatures to a sensor using the GUI:

1. Go to *Security Profiles* > *Intrusion Prevention*.
2. Edit an existing sensor, or create a new one.
3. In the *IPS Signatures* section, click *Add Signatures*. A list of available signatures appears.
4. Select the signatures you want to include from the list.
5. Click *Use Selected Signatures*.



6. Configure other settings as needed.
7. Add this sensor to a firewall policy to detect or block attacks that match the IPS signatures.

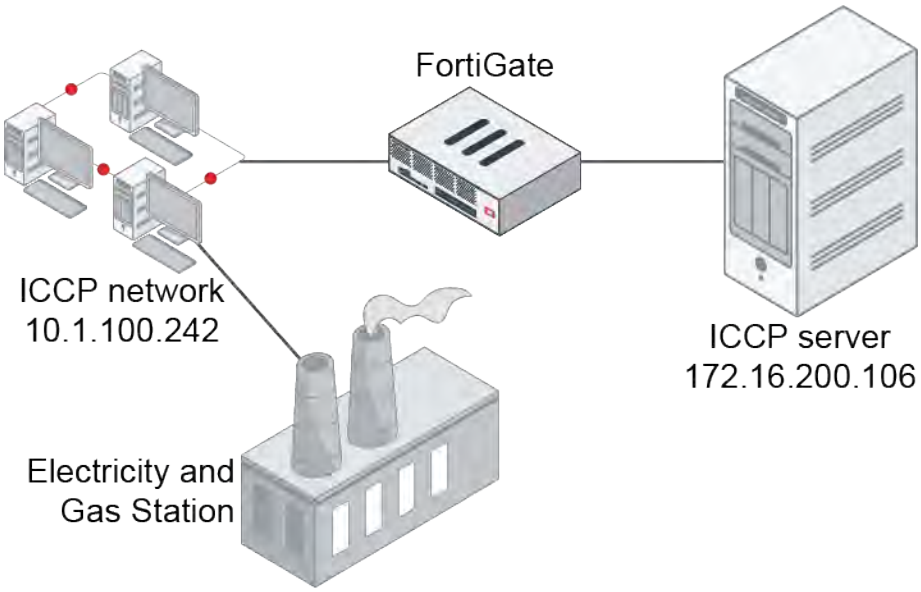
Netskope Exhibit 1017

Detecting IEC 61850 MMS protocol in IPS

IEC 61850 is a SCADA protocol whose services are mapped to a number of protocols, including MMS services. MMS/ICCP detection is supported in IPS. The purpose of the MMS dissectors is to identify every IEC 61850 service to distinguish different MMS/ICCP messages. IPS engine 6.0.12 and later support MMS dissectors.

The following scenarios are also supported:

- Multiple MMS PDUs are transferred in one TCP payload, and the IPS engine identifies individuals.
- An MMS message is split over multiple TCP segments, where MMS runs over COTP segments.
- ICCP/TASE.2 that also uses MMS transport (ISO transport over TCP for ICCP) is detected.



Industrial signatures must be enabled in the global IPS settings to receive MMS/ICCP signatures. By default, industrial signatures are excluded.

```
config ips global
    set exclude-signatures none
end
```

Below are some industrial signatures for MMS/ICCP messages that can be detected by the IPS engine. This is not an exhaustive list.

- MMS_GetNameList.Request
- MMS_GetNamedVariableListAttributes.Request
- MMS_GetVariableAccessAttributes.Request
- MMS_Identify.Request
- MMS_Initiate.Request
- MMS_Read.Request
- MMS_Reset.Request
- ICCP_Transfer.Reporting
- ICCP_Create.Dataset
- ICCP_Abort
- ICCP_Start.Transfer.DSTransferSet
- ICCP_Get.Dataset.Element.Values

- ICCP_Get.Next.DSTransfer.Set.Value
- ICCP_Delete.Dataset
- ICCP_Start.Transfer.IMTransferSet

Diagnose command

The COTP dissector adds support for identifying every MMS PDU, and let the IPS engine separate them, like the Modbus and IEC-104 services for example.

```
# diagnose ips debug enable all
# diagnose debug enable
[284@78]ips_l7_dsct_processor: serial=8142 create: cotp
[284@78]ips_l7_dsct_processor: serial=8142 create: iec104
[284@78]ips_l7_dsct_processor: serial=8142 create: modbus
```

Log samples

MMS dissectors can be triggered, and MMS/ICCP signatures can be monitored and logged.

Log samples:

```
date=2020-03-26 time=15:51:10 logid="1059028704" type="utm" subtype="app-ctrl"
eventtype="signature" level="information" vd="vd1" eventtime=1585263070836106492 tz="-0700"
appid=43699 srcip=10.1.100.242 dstip=172.16.200.106 srcport=50963 dstport=102
srcintf="port13" srcintfrole="undefined" dstintf="port14" dstintfrole="undefined" proto=6
service="tcp/26112" direction="outgoing" policyid=1 sessionid=2711 applist="test"
action="pass" appcat="Industrial" app="MMS_Read.Request" incidentserialno=376610508
msg="Industrial: MMS_Read.Request," apprisk="elevated"
```

```
date=2020-03-26 time=16:15:45 logid="1059028704" type="utm" subtype="app-ctrl"
eventtype="signature" level="information" vd="vd1" eventtime=1585091746264983273 tz="-0700"
appid=44684 srcip=10.1.100.242 dstip=172.16.200.106 srcport=41665 dstport=102
srcintf="port13" srcintfrole="undefined" dstintf="port14" dstintfrole="undefined" proto=6
service="tcp/26112" direction="incoming" policyid=1 sessionid=194463 applist="test"
action="pass" appcat="Industrial" app="ICCP_Transfer.Reporting" incidentserialno=762763993
msg="Industrial: ICCP_Transfer.Reporting," apprisk="elevated"
```

Email filter

Email filters can be configured to perform spam detection and filtering. You can customize the default profile, or create your own and apply it to a firewall policy.



Two kinds of filtering can be defined in a single profile, and they will act independent of one another.

Filter options can be organized according to the source of the decision:

- Local options: the FortiGate qualifies the email based on local conditions, such as block/allowlist, banned words, or DNS checks using FortiGuard Antispam.

- FortiGuard-based options: the FortiGate qualifies the email based on the score or verdict returned from FortiGuard Antispam.
- Third-party options: the FortiGate qualifies the email based on information from a third-party source (like an ORB list).

Local and FortiGuard block/allowlists can be enabled and combined in a single profile. When combined, the local block/allowlist has a higher priority than the FortiGuard blocklist during a decision making process. For example, if a client IP address is blocklisted in the FortiGuard server, but you want to override this decision and allow the IP to pass through the filter, you can define the IP address or subnet in a local block/allowlist with the *clear* action. Because the information coming from the local list has a higher priority than the FortiGuard service, the email will be considered clean.



Some features of this functionality require a subscription to FortiGuard Antispam.

Protocol comparison between Email filter inspection modes

The following table indicates which Email filters are supported by their designated inspection modes.

	SMTP	POP3	IMAP	MAPI
Proxy	Yes	Yes	Yes	Yes
Flow	Yes	Yes	Yes	No

The following topics provide information about email filter profiles:

- Local-based filters on page 980
- FortiGuard-based filters on page 985
- File type-based filters on page 986
- Protocols and actions on page 991
- Configuring webmail filtering on page 992

Local-based filters

You can make block/allowlists from emails or IP subnets to forbid or allow them to send or receive emails.

You can also define a list of banned words. Emails that contain any of these banned words are considered spam.

With the `spamhelodns` and `spamraddrdns` commands, the FortiGate performs a standard DNS check on the machine name used in the helo SMTP message, and/or the return-to field to determine if these names belong to a registered domain. The FortiGate does not check the FortiGuard service during these operations.

To configure a local-based email filter in the CLI:

1. Configure a BWL:

```
config emailfilter bwl
    edit 1
        set name "mtBWL"
```


Netskope Exhibit 1017

```
        config entries
            edit 1
                set status enable
                set type ip
                set action spam
                set addr-type ipv4
                set ip4-subnet 10.1.100.0 255.255.255.0
            next
        end
    next
end
```

2. Configure an email filter profile:

```
config emailfilter profile
    edit "myLocalEmailFilter"
        set spam-filtering enable
        set options spambwl spamhelodns spamraddrdns
        config smtp
            set action tag
        end
        set spam-bwl-table 1
    next
end
```

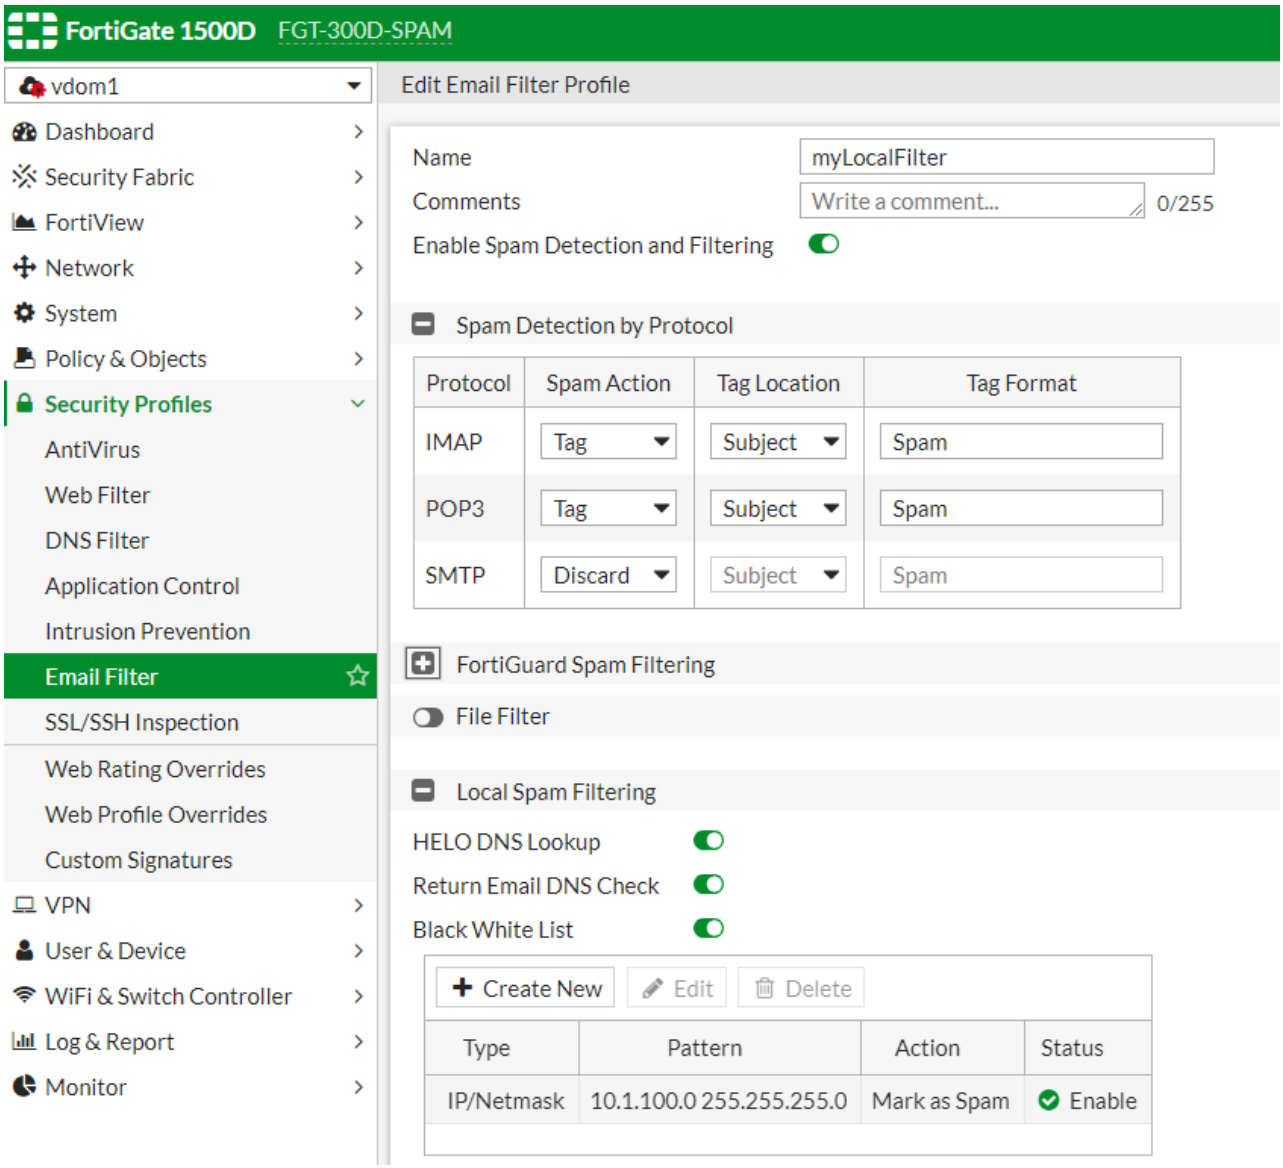
3. Use the profile in a firewall policy:

```
config firewall policy
    edit 1
        .....
        set inspection-mode proxy
        set emailfilter-profile "myLocalEmailFilter"
    next
end
```


Netskope Exhibit 1017

To configure a local-based email filter in the GUI:

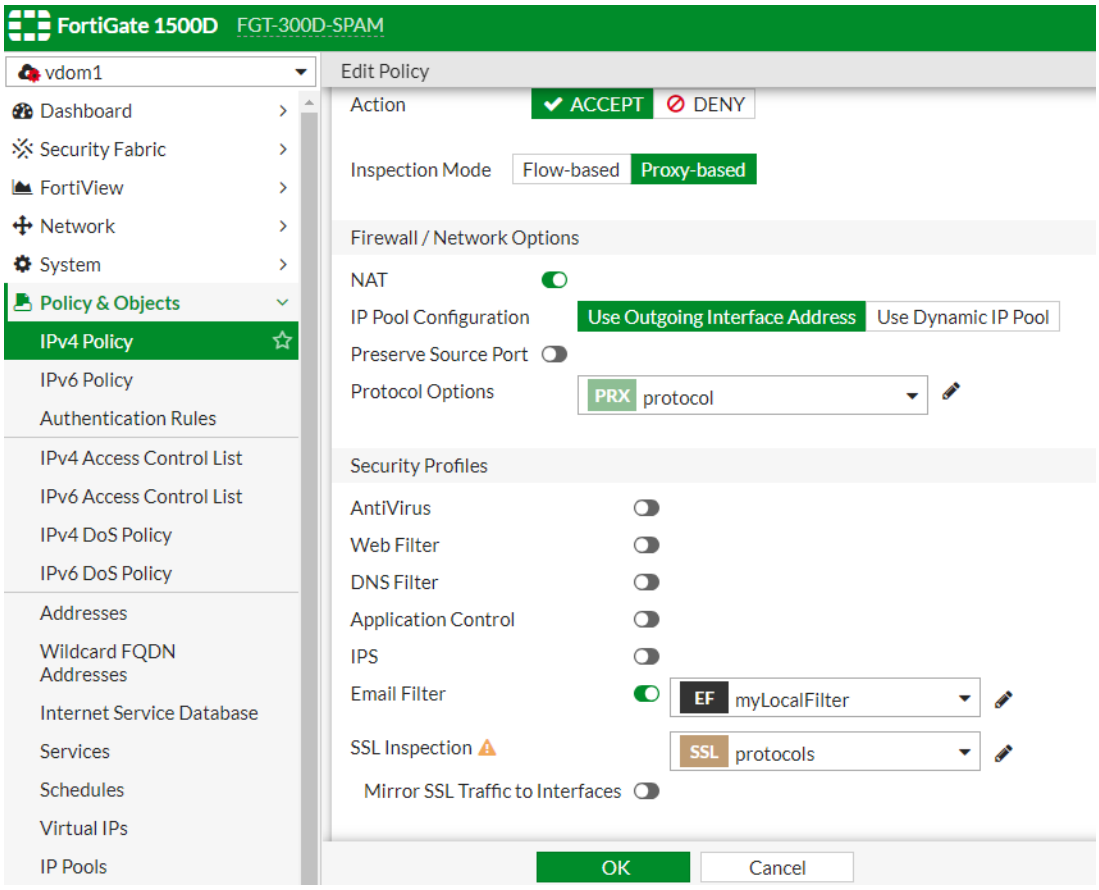
1. Go to *Security Profiles* > *Email Filter*.
2. Click *Create New* or select an existing profile and click *Edit*.



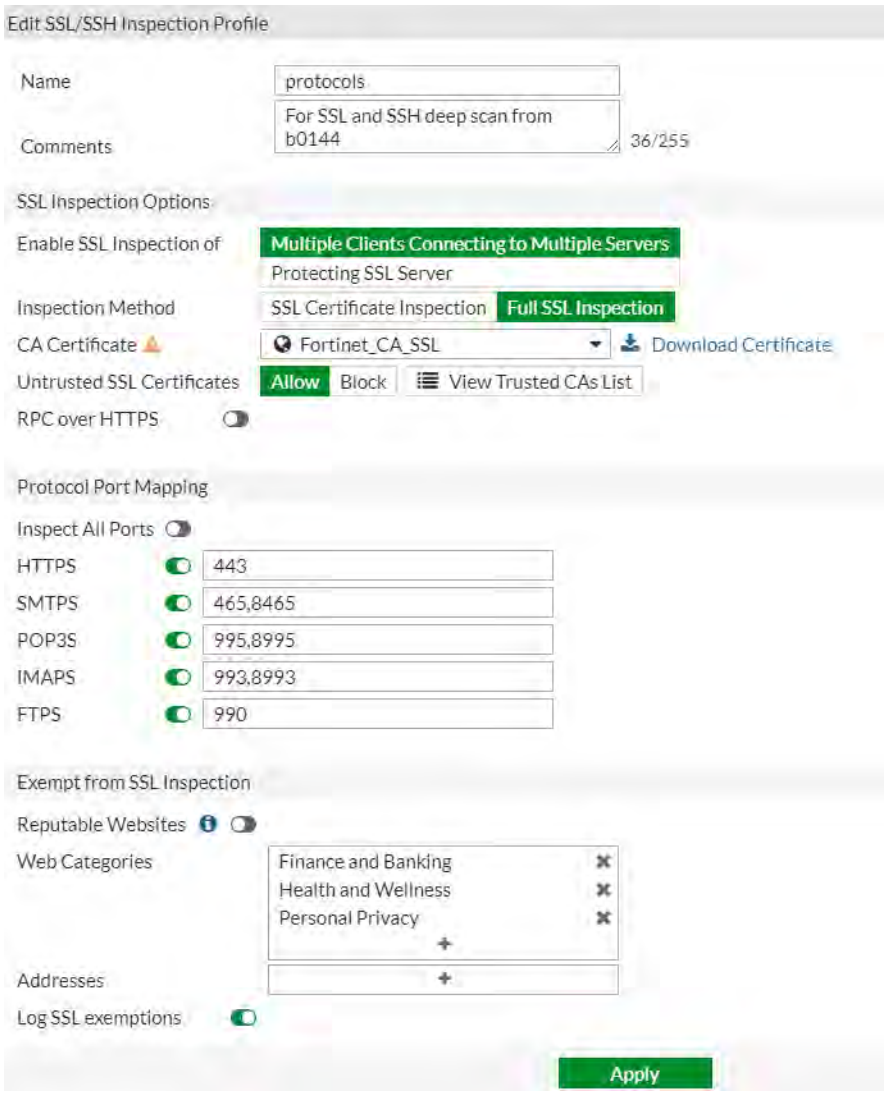
3. In the firewall policy, create or edit a rule.
4. Set the inspection-mode to *Proxy-based*.

Netskope Exhibit 1017

5. Enable the *Email Filter* option and select the previously created profile.



6. Set *SSL Inspection* to a profile that has deep SSL inspection enabled.
Deep inspection is required if you intend to filter SMTP, POP3, IMAP, or any SSL/TLS encapsulated protocol. The below SSL-inspection profile has deep inspection enabled:



7. Click *OK*.

To configure bannedwords in the CLI:

1. Configure a bannedwords list:

```
config emailfilter bword
    edit 1
        set name "banned"
        config entries
            edit 1
                set pattern "undesired_word"
            next
        end
    next
end
```

2. Configure an email filter profile:

```
config emailfilter profile
    edit "myBannedWordsProfile"
        config file-filter
            set status disable
        end
        set spam-filtering enable
        set options bannedword
        set spam-bword-table 1
    next
end
```

3. Use the profile in a firewall policy:

```
config firewall policy
    edit 1
        .....
        set inspection-mode proxy
        set emailfilter-profile "myBannedWordsProfile"
    next
end
```

Once created, this profile should be set in the firewall policy.



Bannedwords can only be configured through the CLI.

FortiGuard-based filters

The FortiGate consults FortiGuard servers to help identify spammer IP address or emails, known phishing URLs, known spam URLs, known spam email checksums, and others.

FortiGuard servers have maintained databases that contain blocklists, which are fed from Fortinet sensors and labs distributed all over the world.

To configure the FortiGuard filters in the CLI:

```
config emailfilter profile
    edit "myEmailFilterProfile"
        set spam-filtering enable
        set options spamfsip spamfssubmit spamfschksum spamfsurl spamrbl spamhdrcheck
spamfsphish
    next
end
```

To configure the FortiGuard filters in the GUI:

1. Go to *Security Profiles* > *Email Filter*.
2. In the *FortiGuard Spam Filtering Spam Filtering* section, you can enable or disable the following filters:
   - IP Address Check
   - URL Check

- Detect Phising URLs in Email
- Email Checksum Check
- Spam Submission

3. Click *OK*.

File type-based filters

With *File Filter*, you can define undesired file types within the email filter profile and associate an action to be taken for each file type, such as block or log.

For each entry, you can also specify the protocol to inspect (SMTP, POP3, or IMAP) and if only encrypted files should be matched. While file filter entries are ordered, the block action takes precedence over the log action.

File filtering in email filter profiles is based only on the file type (file meta data) and not on file size or content. You would need to configure a DLP sensor to block files based on size or content, such as SSN numbers, credit card numbers, or regexp.

File filtering only works in proxy mode policies. The traffic direction cannot be configured because it is implied by the protocol.

The following file types are supported:

File Type Name	Description
7z	Match 7-zip files
arj	Match arj compressed files
cab	Match Windows cab files
lzh	Match lzh compressed files
rar	Match rar archives
tar	Match tar files
zip	Match zip files
bzip	Match bzip files
gzip	Match gzip files
bzip2	Match bzip2 files
xz	Match xz files
bat	Match Windows batch files
msc	Match msc files
uue	Match uue files
mime	Match mime files
base64	Match base64 files
binhex	Match binhex files

File Type Name	Description
bin	Match bin files
elf	Match elf files
exe	Match Windows executable files
hta	Match hta files
html	Match html files
jad	Match jad files
class	Match class files
cod	Match cod files
javascript	Match javascript files
msoffice	Match MS-Office files. For example, doc, xls, ppt, and so on.
msofficex	Match MS-Office XML files. For example, docx, xlsx, pptx, and so on.
fsg	Match fsg files
upx	Match upx files
petite	Match petite files
aspack	Match aspack files
prc	Match prc files
sis	Match sis files
hlp	Match Windows help files
activemime	Match activemime files
jpeg	Match jpeg files
gif	Match gif files
tiff	Match tiff files
png	Match png files
bmp	Match bmp files
unknown	Match unknown files
mpeg	Match mpeg files
mov	Match mov files
mp3	Match mp3 files
wma	Match wma files
wav	Match wav files

File Type Name	Description
pdf	Match pdf files
avi	Match avi files
rm	Match rm files
torrent	Match torrent files
msi	Match Windows Installer msi bzip files
mach-o	Match Mach object files
dmg	Match Apple disk image files
.net	Match .NET files
xar	Match xar archive files
chm	Match Windows compiled HTML help files
iso	Match ISO archive files
crx	Match Chrome extension files

Example

In the following example, one file filter entry is created to block executable (exe) files from being sent or received, and a second entry logs any documents that are sent.

To configure a file-type based email filter in the CLI:

1. Configure an email filter profile:

```
config emailfilter profile
    edit "file-type-filter"
        config file-filter
            set status enable
            set log enable
            set scan-archive-contents enable
            config entries
                edit "filter1"
                    set comment "Block executable files"
                    set protocol smtp imap pop3
                    set action block
                    set encryption any
                    set file-type "exe"
                next
                edit "filter2"
                    set comment "Log document files"
                    set protocol smtp
                    set action log
                    set encryption any
                    set file-type "pdf" "msoffice" "msofficex"
                next
            end
```

```
        end
        set spam-filtering enable
    next
end
```

2. Use the profile in a firewall policy:

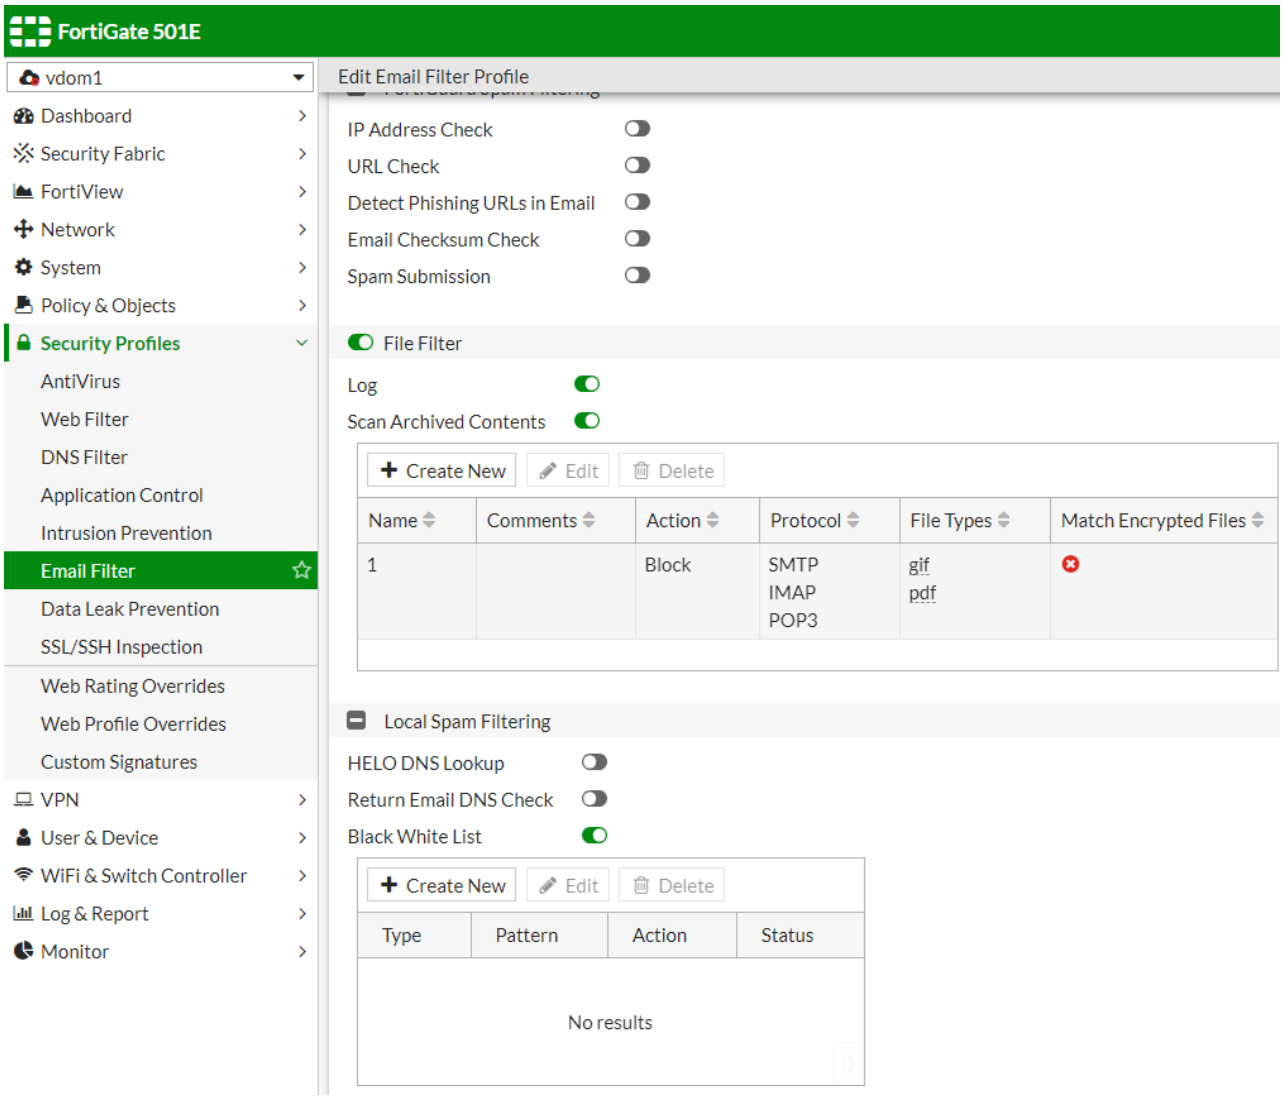
```
config firewall policy
    edit 1
        set name "client-to-internet"
        set srcintf "port2"
        set dstintf "port1"
        set srcaddr "all"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "ALL"
        set utm-status enable
        set utm-inspection-mode proxy
        set logtraffic all
        set emailfilter profile "file-type-filter"
        set profile-protocol-options "protocol"
        set ssl-ssh-profile "protocols"
        set nat enable
    next
end
```

To configure a file-type based email filter in the GUI:

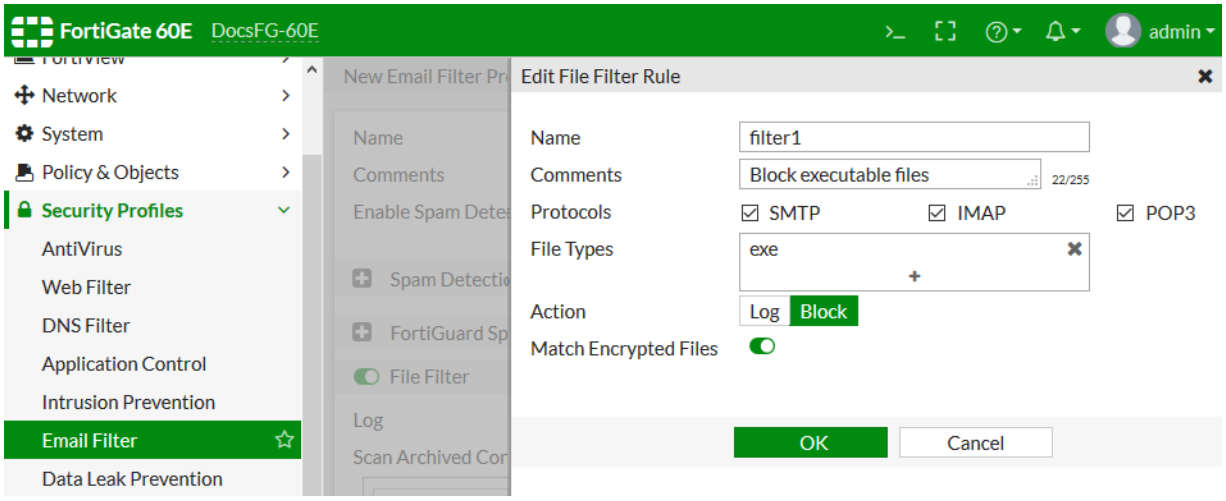
1. Go to *Security Profiles* > *Email Filter*.
2. Click *Create New*, or select an existing profile and click *Edit*.
3. Enable *Enable Spam Detection and Filtering*.
4. Enable *File Filter*.
5. Enable *Log* and *Scan Archived Contents*.

Netskope Exhibit 1017

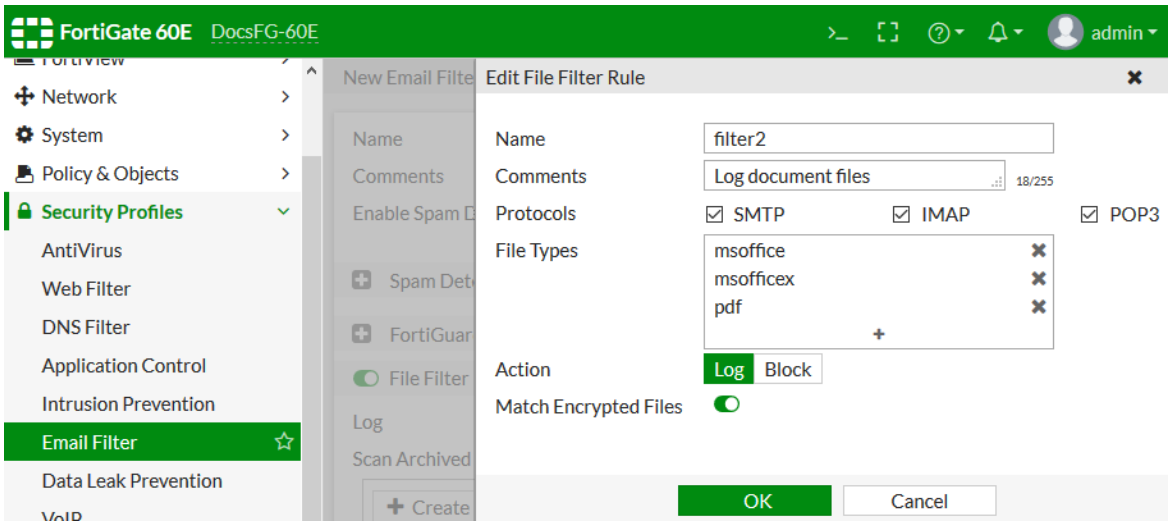
6. In the *File Filter* table, click *Create New*.



7. Configure the filters:
   a. *filter1* blocks all sent or received executable files.



b. *filter2* logs all sent documents.



8. Click *OK*.
9. Add the new email filter profile to a firewall policy.

Log examples

File filter block action:

```
1: date=2019-01-25 time=15:20:16 logid="0554020511" type="utm" subtype="file-filter"
eventtype="file-filter" level="warning" vd="vdom1" eventtime=1548458416 policyid=1
sessionid=2881 srcip=10.1.100.12 srcport=45974 srcintf="port2" srcintfrole="undefined"
dstip=172.16.200.56 dstport=143 dstintf="port1" dstintfrole="undefined" proto=6
service="IMAP" action="blocked" from="emailuser1@qa.fortinet.com"
to="emailuser2@qa.fortinet.com" recipient="emailuser2" direction="incoming" subject="EXE
file block" size="622346" attachment="yes" filename="putty.exe" filtername="filter1"
filetype="exe"
```

File filter log action:

```
1: date=2019-01-25 time=15:23:16 logid="0554020510" type="utm" subtype="file-filter"
eventtype="file-filter" level="notice" vd="vdom1" eventtime=1548458596 policyid=1
sessionid=3205 srcip=10.1.100.12 srcport=55664 srcintf="port2" srcintfrole="undefined"
dstip=172.16.200.56 dstport=25 dstintf="port1" dstintfrole="undefined" proto=6
service="SMTP" profile="emailfilter-file-filter" action="detected"
from="emailuser1@qa.fortinet.com" to="emailuser2@qa.fortinet.com"
sender="emailuser1@qa.fortinet.com" recipient="emailuser2@qa.fortinet.com"
direction="outgoing" subject="PDF file log" size="390804" attachment="yes"
filename="fortiauto.pdf" filtername="filter2" filetype="pdf"
```

Protocols and actions

In an email filtering profile, there are sections for SMTP, POP3, and IMAP protocols. In each section, you can set an action to either discard, tag, or pass the log for that protocol.

Protocol	Available action
SMTP	**Pass**: Allow spam email to pass through.
	Tag: Tag spam email with configured text in the subject or header.
	Discard: Discards (blocks) spam email.
POP3 & IMAP	**Pass**: Allow spam email to pass through.
	Tag: Tag spam email with configured text in the subject or header.
MAPI:	**Pass**: Allow spam email to pass through.
	Discard: Discards (blocks) spam email.

For example, to tag spam SMTP email, use the following commands:

```
config emailfilter profile
    edit "smtpFilter"
        set spam-filtering enable
        set options <options>
        ...
        config smtp
            set log enable
            set action tag
        end
    next
end
```

MAPI email filtering

MAPI is a proprietary protocol from Microsoft. It uses HTTPS to encapsulate email requests and responses between Microsoft Outlook clients and Microsoft Exchange servers. The configuration of MAPI email filters are only possible through the CLI.

To configure the MAPI email filter in the CLI:

```
config emailfilter profile
    edit "myMapiFilter"
        set spam-filtering enable
        set options spamfsip spamfssubmit spamfsurl spamfsphish
        config mapi
            set log enable
            set action "discard or pass"
        end
    next
end
```

Configuring webmail filtering

You can configure an email filter to detect and log emails sent via Gmail and Hotmail. These interfaces do not use standard email protocols (SMTP, POP3, or IMAP) and instead use HTTPS. However, you can still configure the email filter to detect emails that pass through the FortiGate.

The FortiGate only detects and logs the emails, it does not discard or tag them.

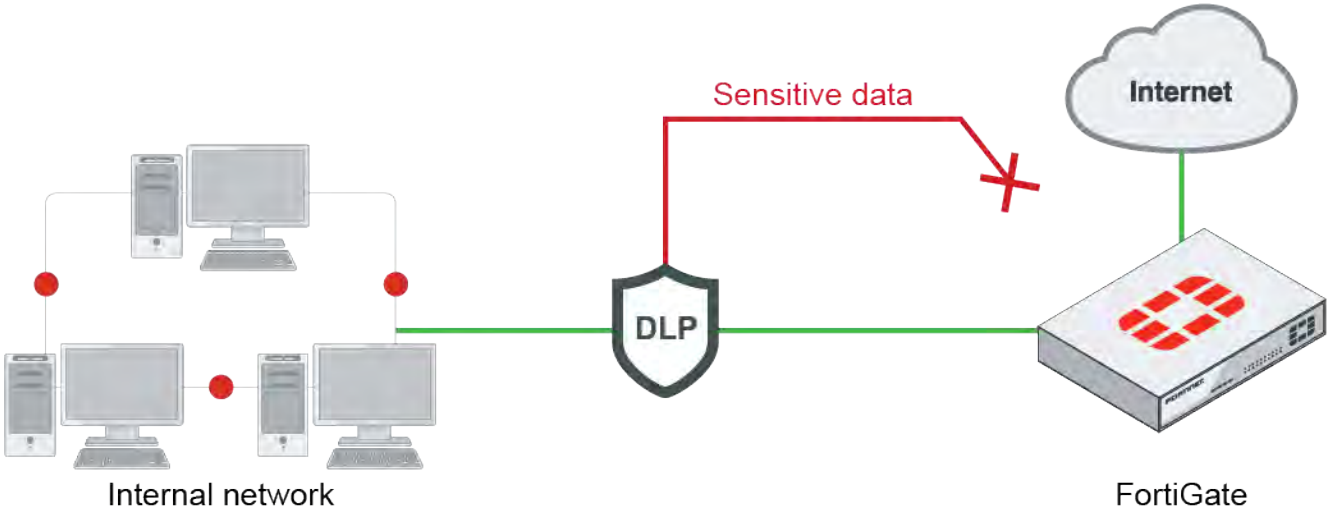
To configure webmail filtering in the CLI:

```
config emailfilter profile
    edit "myWebMailDetector"
        set spam-filtering enable
        config msn-hotmail
            set log enable
        end
        config gmail
            set log enable
        end
    next
end
```

Data leak prevention

The FortiGate data leak prevention (DLP) system prevents sensitive data from leaving or entering your network. You can customize the default sensor or create your own by adding individual filters based on file type, file size, a regular expression, an advanced rule, or a compound rule. Once configured, you can apply the DLP sensor to a firewall policy. Data matching defined sensitive data patterns is blocked, logged, or allowed when it passes through the FortiGate.

DLP can only be configured in the CLI.



The filters in a DLP sensor can examine traffic for the following:

- Known files using DLP fingerprinting
- Known files using DLP watermarking
- Particular file types
- Particular file names

- Files larger than a specified size
- Data matching a specified regular expression
- Credit card and social security numbers



Filters are ordered, but there is no precedence between the possible actions

DLP is primarily used to stop sensitive data from leaving your network. DLP can also be used to prevent unwanted data from entering your network and to archive some or all of the content that passes through the FortiGate. DLP archiving is configured per filter, which allows a single sensor to archive only the required data. You can configure the DLP archiving protocol in the CLI (see Configure DLP sensors).

There are two forms of DLP archiving:

- Summary only: a summary of all the activity detected by the sensor is recorded. For example, when an email message is detected, the sender, recipient, message subject, and total size are recorded. When a user accesses the web, every URL that they visit is recorded.
- Full: detailed records of all the activity detected by the sensor is recorded. For example, when an email message is detected, the message itself, including any attachments, is recorded. When a user accesses the web, every page that they visit is archived.

Protocol comparison between DLP inspection modes

The following table indicates which protocols can be inspected by DLP based on the specified inspection modes.

	HTTP	FTP	IMAP	POP3	SMTP	NNTP	MAPI	CIFS	SFTP/SCP
Proxy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	Yes
Flow	Yes	Yes	Yes	Yes	Yes	No	No	Yes	No

The following topics provide information about DLP:

- Basic DLP filter types on page 994
- DLP fingerprinting on page 997
- DLP watermarking on page 1001

Basic DLP filter types

Basic filter types can be configured in the CLI and include:

- File type and name
- File size
- Regular expression
- Credit card and SSN

File type and name

A file type filter allows you to block, allow, log, or quarantine based on the file type specified in the file filter list (see File filter on page 916).

To configure file type and name filtering:

1. Create a file pattern to filter files based on the file name patter or file type:

```
config dlp filepattern
    edit <filepatern_entry_integer>
        set name <string>
        config entries
            edit <file pattern>
                set filter-type <type | pattern>
                set file-type <file type>
            next
        end
    next
end
```

For example, to filter for GIFs and PDFs:

```
config dlp filepattern
    edit 11
        set name "sample_config"
        config entries
            edit "*.gif"
                set filter-type pattern
            next
            edit "pdf"
                set filter-type type
                set file-type pdf
            next
        end
    next
end
```

2. Attach the file pattern to a DLP sensor, and specify the protocols and actions:

```
config dlp sensor
   edit <string>
      config filter
         edit <integer>
            set name <string>
            set proto <smtp | pop3 | imap | http-get | http-post | ftp | nntp |
mapi>
            set filter-by file-type
            set file-type 11    <-- Previously configured filepattern
            set action <allow | log-only| block | quarantine-ip>
         next
      end
   next
end
```


Netskope Exhibit 1017

File size

A file size filter checks for files that exceed the specific size, and performs the DLP sensor's configured action on them.

To configure file size filtering:

```
config dlp sensor
   edit <string>
      config filter
         edit <integer>
            set name <string>
            set proto <smtp | pop3 | imap | http-get | http-post | ftp | nntp | mapi>
            set filter-by file-size  <-- Match any file over with a size over the
threshold
            set file-type 11   <-- Previously configured filepattern
            set action <allow | log-only| block | quarantine-ip>
         next
      end
   next
end
```

Regular expression

A regular expression filter is used to filter files or messages based on the configured regular expression pattern.

To configure regular expression filtering:

```
config dlp sensor
   edit <string>
      config filter
         edit <integer>
            set name <string>
            set type <file | message>   <-- Check contents of a file or of messages, web
pages, etc.
            set proto <smtp | pop3 | imap | http-get | http-post | ftp | nntp | mapi>
            set filter-by regexp   <-- Use a regular expression to match content
            set regexp <regexp>   <-- Input a regular expression pattern
            set action <allow | log-only| block | quarantine-ip>
         next
      end
   next
end
```

Credit card and SSN

The credit card sensor can match the credit card number formats used by American Express, Mastercard, and Visa. It can be used to filter files or messages.

The SSN sensor can be used to filter files or messages for Social Security Numbers.

To configure credit card or SSN filtering:

```
config dlp sensor
   edit <string>
      config filter
         edit <integer>
            set name <string>
            set type <file | message>  <-- Check contents of a file, or of messages, web
pages, etc.
            set proto <smtp | pop3 | imap | http-get | http-post | ftp | nntp | mapi>
            set filter-by < credit-card | ssn >   <-- Match credit cards or social security
numbers
            set action <allow | log-only| block | quarantine-ip>
         next
      end
   next
end
```

DLP fingerprinting

DLP fingerprinting can be used to detect sensitive data. The file that the DLP sensor will filter for is uploaded and the FortiGate generates and stores a checksum fingerprint. The FortiGate unit generates a fingerprint for all of the files that are detected in network traffic, and compares all of the checksums stored in its database. If a match is found, the configured action is taken.

Any type of file can be detected by DLP fingerprinting, and fingerprints can be saved for each revision of a file as it is updated.

To use fingerprinting:

- Select the files to be fingerprinted by targeting a document source.
- Add fingerprinting filters to DLP sensors.
- Add the sensors to firewall policies that accept traffic that the fingerprinting will be applied on.



The document fingerprint feature requires a FortiGate device that has internal storage.

To configure a DLP fingerprint document:

```
    config dlp fp-doc-source
        edit <name_str>
            set server-type smb
            set server <string>
            set period {none | daily | weekly | monthly}
            set vdom {mgmt | current}
            set scan-subdirectories {enable | disable}
            set remove-deleted {enable | disable}
            set keep-modified {enable | disable}
            set username <string>
            set password <password>
            set file-path <string>
            set file-pattern <string>
```

```
        set sensitivity <Critical | Private | Warning>
        set tod-hour <integer>
        set tod-min <integer>
        set weekday {sunday | monday | tuesday | wednesday | thursday | friday |
saturday}
        set date <integer>
    next
end
```

Command	Description
server-type smb	The protocol used to communicate with document server. Only Samba (SMB) servers are supported.
server <string>	IPv4 or IPv6 address of the server.
period {none | daily | weekly | monthly}	The frequency that the FortiGate checks the server for new or changed files.
vdom {mgmt | current}	The VDOM that can communicate with the file server.
scan-subdirectories {enable | disable}	Enable/disable scanning subdirectories to find files.
remove-deleted {enable | disable}	Enable/disable keeping the fingerprint database up to date when a file is deleted from the server.
keep-modified {enable | disable}	Enable/disable keeping the old fingerprint and adding a new one when a file is changed on the server.
username <string>	The user name required to log into the file server.
password <password>	The password required to log into the file server.
file-path <string>	The path on the server to the fingerprint files.
file-pattern <string>	Files matching this pattern on the server are fingerprinted.
sensitivity <Critical | Private | Warning>	The sensitivity or threat level for matches with this fingerprint database.
tod-hour <integer>	Set the hour of the day. This option is only available when `period` is not `none`.
tod-min <integer>	Set the minute of the hour. This option is only available when `period` is not `none`.
weekday {sunday | monday | tuesday | wednesday | thursday | friday | saturday}	Set the day of the week. This option is only available when `period` is `weekly`.
date <integer>	Set the day of the month. This option is only available when `period` is `monthly`.

To configure a DLP fingerprint sensor:

```
config dlp sensor
    edit <sensor name>
        config filter
            edit <id number of filter>
                set proto {smtp | pop3 | imap http-get | http-post | ftp | nntp | mapi}
```

```
                set filter-by fingerprint
                set sensitivity {Critical | Private | Warning}
                set match-percentage <integer>
                set action {allow | log-only | block | ban | quarantine-ip}
            next
        end
    next
end
```

Command	Description
proto {smtp | pop3 | imap http-get | http-post | ftp | nntp | mapi}	The protocol to inspect.
filter-by fingerprint	Match against a fingerprint sensitivity.
sensitivity {Critical | Private | Warning}	Select a DLP file pattern sensitivity to match.
match-percentage <integer>	The percentage of the checksum required to match before the sensor is triggered.
action {allow | log-only | block | ban | quarantine-ip}	The action to take with content that this DLP sensor matches.

View the DLP fingerprint database on the FortiGate

The CLI debug command `diagnose test application dlpfingerprint` can be used to display the fingerprint information that is on the FortiGate.

```
Fingerprint Daemon Test Usage;
-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-
  1 : This menu
  2 : Dump database
  3 : Dump all files
  5 : Dump all chunk
  6 : Refresh all doc sources in all VDOMs
  7 : Show the db file size and the limit
  9 : Display stats
 10 : Clear stats
 99 : Restart this daemon
```

For example, option 3 will dump all fingerprinted files:

```
DLP_WANOPT-CLT (global) # diagnose test application dlpfingerprint 3
DLPFP diag_test_handler called
File DB:
--------------------------------------
id, filename,                              vdom, archive, deleted, scanTime,   docSourceSrvr,
sensitivity, chunkCnt, reviseCnt,
1,  /fingerprint/upload/1.txt,             vdom1,  0,      0,      1494868196,   1,      2,
1,    0,
2,  /fingerprint/upload/30percentage.xls,  vdom1,  0,      0,      1356118250,   1,      2,
13,      0,
3,  /fingerprint/upload/50.pdf,            vdom1,  0,      0,      1356118250,   1,      2,
122,  0,
4,  /fingerprint/upload/50.pdf.tar.gz,     vdom1,  0,      0,      1356118250,   1,      2,
114,     0,
5,  /fingerprint/upload/check-list_AL-SIP_HA.xls,   vdom1,  0,      0,      1356118251,     1,
```

```
2,       32,     0,
6,  /fingerprint/upload/clean.zip,          vdom1,  0,      0,      1356118251,   1,      2,
1,    0,
7,  /fingerprint/upload/compare.doc,        vdom1,  0,      0,      1522097410,   1,      2,
18,      0,
8,  /fingerprint/upload/dlpsensor-watermark.pdf,    vdom1,  0,      0,      1356118250,     1,
2,       11,     0,
9,  /fingerprint/upload/eicar.com,          vdom1,  0,      0,      1356118250,   1,      2,
1,    0,
10, /fingerprint/upload/eicar.zip,          vdom1,  0,      0,      1356118250,   1,      2,
1,    0,
11, /fingerprint/upload/EMAIL-CONTENT-ARCHIVE.ppt,  vdom1,  0,      0,      1356118250,     1,
2,       11,     0,
12, /fingerprint/upload/encrypt.zip,        vdom1,  0,      0,      1356118250,   1,      2,
77,      0,
13, /fingerprint/upload/extension_7_8_1.crx,        vdom1,  0,      0,      1528751781,     1,
2,       2720,   0,
14, /fingerprint/upload/fingerprint.txt,    vdom1,  0,      0,      1498582679,   1,      2,
37,      0,
15, /fingerprint/upload/fingerprint90.txt,  vdom1,  0,      0,      1498582679,   1,      2,
37,      0,
16, /fingerprint/upload/fo2.pdf,            vdom1,  0,      0,      1450488049,   1,      2,
1,    0,
17, /fingerprint/upload/foo.doc,            vdom1,  0,      0,      1388538131,   1,      2,
9,    0,
18, /fingerprint/upload/fortiauto.pdf,      vdom1,  0,      0,      1356118251,   1,      2,
146,     0,
19, /fingerprint/upload/image.out,          vdom1,  0,      0,      1531802940,   1,      2,
5410, 0,
20, /fingerprint/upload/jon_file.txt,       vdom1,  0,      0,      1536596091,   1,      2,    1,
      0,
21, /fingerprint/upload/machotest,          vdom1,  0,      0,      1528751955,   1,      2,
19,   0,
22, /fingerprint/upload/nntp-server.doc,    vdom1,  0,      0,      1356118250,   1,      2,
17,      0,
23, /fingerprint/upload/notepad++.exe,      vdom1,  0,      0,      1456090734,   1,      2,
1061,    0,
24, /fingerprint/upload/nppIExplorerShell.exe,      vdom1,  0,      0,      1438559930,     1,
2,       5,      0,
25, /fingerprint/upload/NppShell_06.dll,    vdom1,  0,      0,      1456090736,   1,      2,
111,     0,
26, /fingerprint/upload/PowerCollections.chm,       vdom1,  0,      0,      1533336889,     1,
2,       728,    0,
27, /fingerprint/upload/reflector.dmg,      vdom1,  0,      0,      1533336857,   1,      2,
21117,   0,
28, /fingerprint/upload/roxio.iso,          vdom1,  0,      0,      1517531765,   1,      2,
49251,0,
29, /fingerprint/upload/SciLexer.dll,       vdom1,  0,      0,      1456090736,   1,      2,
541,     0,
30, /fingerprint/upload/screen.jpg,         vdom1,  0,      0,      1356118250,   1,      2,
55,   0,
31, /fingerprint/upload/Spec to integrate FASE into FortiOS.doc,    vdom1,  0,    0,
1356118251,    1,      2,      31,     0,
32, /fingerprint/upload/subdirectory1/subdirectory2/subdirectory3/hibun.aea,        vdom1,  0,
0,       1529019743,     1,      2,      1,      0,
33, /fingerprint/upload/test.pdf,           vdom1,  0,      0,      1356118250,   1,      2,
5,    0,
34, /fingerprint/upload/test.tar,           vdom1,  0,      0,      1356118251,   1,      2,
3,    0,
35, /fingerprint/upload/test.tar.gz,        vdom1,  0,      0,      1356118250,   1,      2,    1,
```


Netskope Exhibit 1017

```
      0,
36, /fingerprint/upload/test1.txt,          vdom1,  0,      0,      1540317547,   1,      2,
1,    0,
37, /fingerprint/upload/thousand-files.zip, vdom1,  0,      0,      1536611774,   1,      2,
241,     0,
38, /fingerprint/upload/Thumbs.db,          vdom1,  0,      0,      1445878135,   1,      2,
3,    0,
39, /fingerprint/upload/widget.pdf,         vdom1,  0,      0,      1356118251,   1,      2,
18,   0,
40, /fingerprint/upload/xx00-xx01.tar,      vdom1,  0,      0,      1356118250,   1,      2,    5,
      0,
41, /fingerprint/upload/xx02-xx03.tar.gz,   vdom1,  0,      0,      1356118251,   1,      2,    1,
      0,
```

DLP watermarking

Watermarking marks files with a digital pattern to designate them as proprietary to a specific company. A small pattern is added to the file that is recognized by the DLP watermark filter, but is invisible to the end user (except for text files).

FortiExplorer client, or a Linux-based command line tool, can be used to add a watermark to the following file types:

- .txt
- .doc and .docx
- .pdf
- .ppt and .pptx
- .xls and .xlsx

The following information is covered in this section:

- Watermarking a file with FortiExplorer.
- Watermarking a file with the Linux tool.
- Configuring a DLP sensor to detect watermarked files.

FortiExplorer

In this example, a watermark will be added to small text file. The content of the file is:

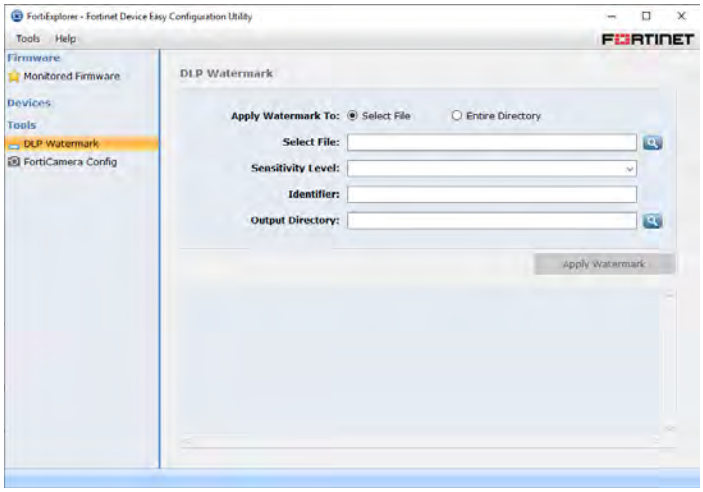
```
This is to show how DLP watermarking is done using FortiExplorer.
```

FortiExplorer can also be used to watermark an entire directory.

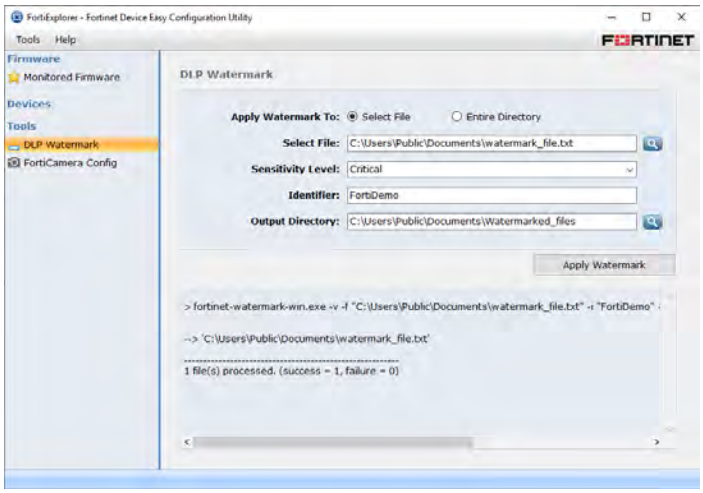
Netskope Exhibit 1017

To watermark the text file with FortiExplorer:

1. Open the FortiExplorer client.
2. Select DLP Watermark from the left side bar.



3. Set *Apply Watermark To* to *Select File*.
4. Browse for the file, copy the file's path into the *Select File* field.
5. Set the *Sensitivity Level*. The available options are: *Critical*, *Private*, and *Warning*.
6. Enter a company identifier in the *Identifier* field.
7. Select the *Output Directory* where the watermarked file will be saved.
8. Click *Apply Watermark*. The file is watermarked.



9. The watermarked file content is changed to:

```
This is to show how DLP watermarking is done using FortiExplorer.=-=-=-=-=-=-=-=-=-=-=-=-
=-=-=-=-=-=-=-=-=identifier=FortiDemo sensitivity=Critical=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-
=-=-=-=-=-=-=-=-
```



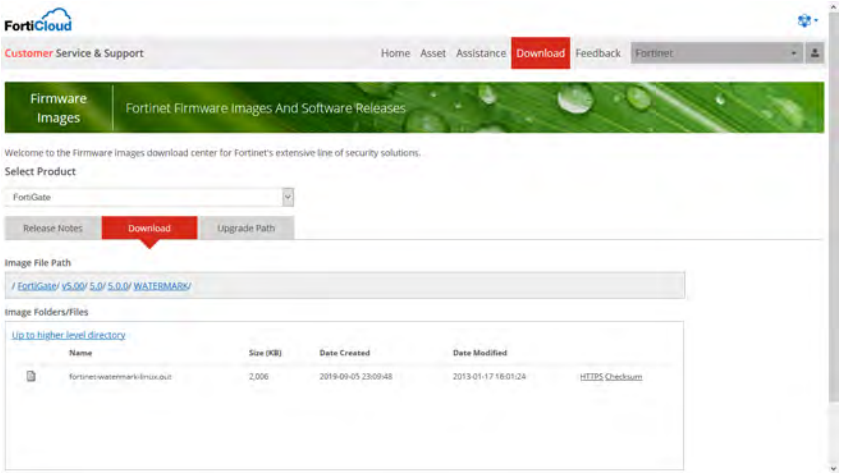
The watermark pattern is visible in text files. For all other supported file types, it is invisible.

Linux-based command line tool

A Linux-based command line tool can be used to watermark files. The tool can be executed is a Linux environment by passing in files or directories of files.

To download the tool:

1. Log in to Fortinet Service and Support. A valid support contract is required.
2. Go to *Download* > *Firmware Images*.
3. Select the *Download* tab, and go to *FortiGate/v5.00/5.0/5.0.0/WATERMARK*.



4. Download the *fortinet-watermark-linux.out* file.

To run the tool:

Enter the following to run the tool on a file:

```
watermark_linux_amd64 <options> -f <file name> -i <identifier> -l <sensitivity level>
```

Enter the following to run the tool on a directory:

```
watermark_linux_amd64 <options> -d <directory> -i <identifier> -l <sensitivity level>
```

The following options are available:

`-h`	Print this help.
`-I`	Watermark the file in place (don't make a copy of the file).
`-o`	The output file or directory.
`-e`	Encode <to non-readable>.
`-i`	Add a watermark identifier.
`-l`	Add a watermark sensitivity level.
`-D`	Delete a watermark identifier.
`-L`	Delete a watermark sensitivity level.

DLP watermark sensor

A DLP watermark sensor must be configured to detect watermarked files.

To configure a DLP watermark sensor:

```
config dlp sensor
   edit <sensor name>
      config filter
         edit <id number of filter>
           set proto {smtp | pop3 | imap http-get | http-post | ftp | nntp | mapi} <--
Protocol to inspect
           set filter-by watermark
           set sensitivity {Critical | Private | Warning}
           set company-identifier <string>
           set action {allow | log-only | block | ban | quarantine-ip}
         next
      end
   next
end
```

VoIP solutions

You can configure VoIP profiles to allow SIP and SCCP traffic and to protect your network from SIP- and SCCP-based attacks.

FortiOS includes two preloaded VoIP profiles:

- *default*
- *strict*

You can customize these profiles, or you can create your own and add them to firewall policies that allow VoIP.

The following topics provide information about VoIP profiles:

- General use cases on page 1004
- SIP message inspection and filtering on page 1008
- SIP pinholes on page 1010
- SIP over TLS on page 1011
- Custom SIP RTP port range support on page 1012
- Voice VLAN auto-assignment on page 1014

General use cases

There are three scenarios in which the FortiOS session initiation protocol (SIP) solution is usually deployed:

1. The SIP server is in a private network, protected from the internet by a FortiOS device.
2. The SIP clients are in a private network, protected from the internet by a FortiOS device.

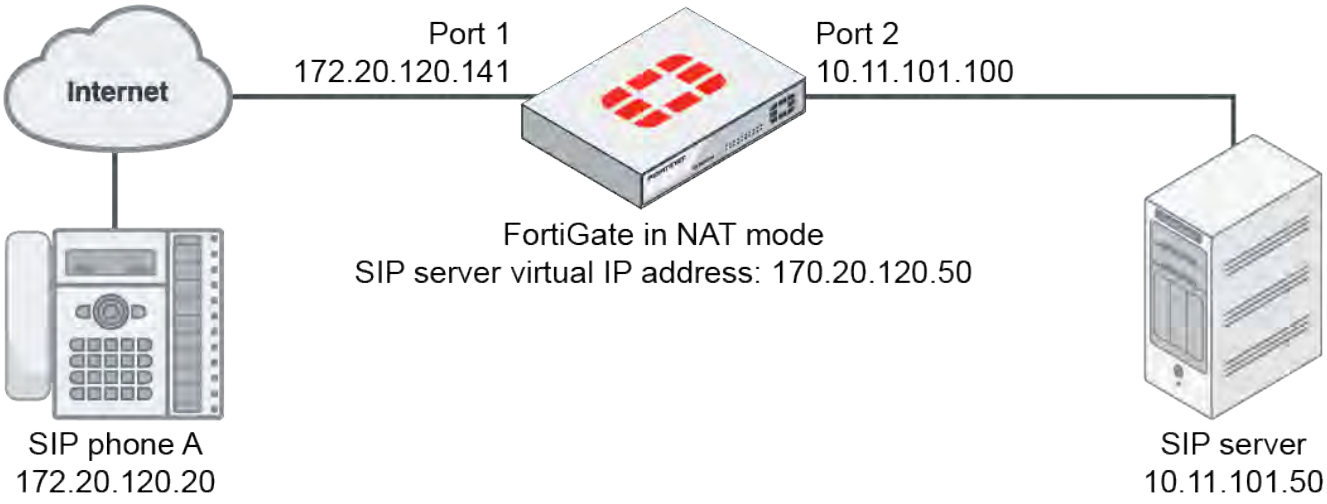
Netskope Exhibit 1017

3. The SIP server is in a private network, such as a corporation's internal network or an ISP's network, protected from the Internet by a FortiOS device. The SIP clients are in a remote private network, such as a SOHO network, and behind a NAT device that is not aware of SIP applications.

The following VIP, NAT, and HNT examples show configurations for each of the three common scenarios.

VIP

A FortiGate with SIP Application Layer Gateway (ALG) or SIP Session Helper protects the SIP server from the internet, while SIP phones from the internet need to register to the SIP server and establish calls through it.



A VIP needs to be configured for the SIP server, and the VIP must be applied in a firewall policy for the phones to send REGISTER messages through the FortiGate from port1 to port2.

Only one firewall policy needs to be configured for all SIP phones on both the internet and private network to register to the SIP server through Port1 and set up SIP calls.

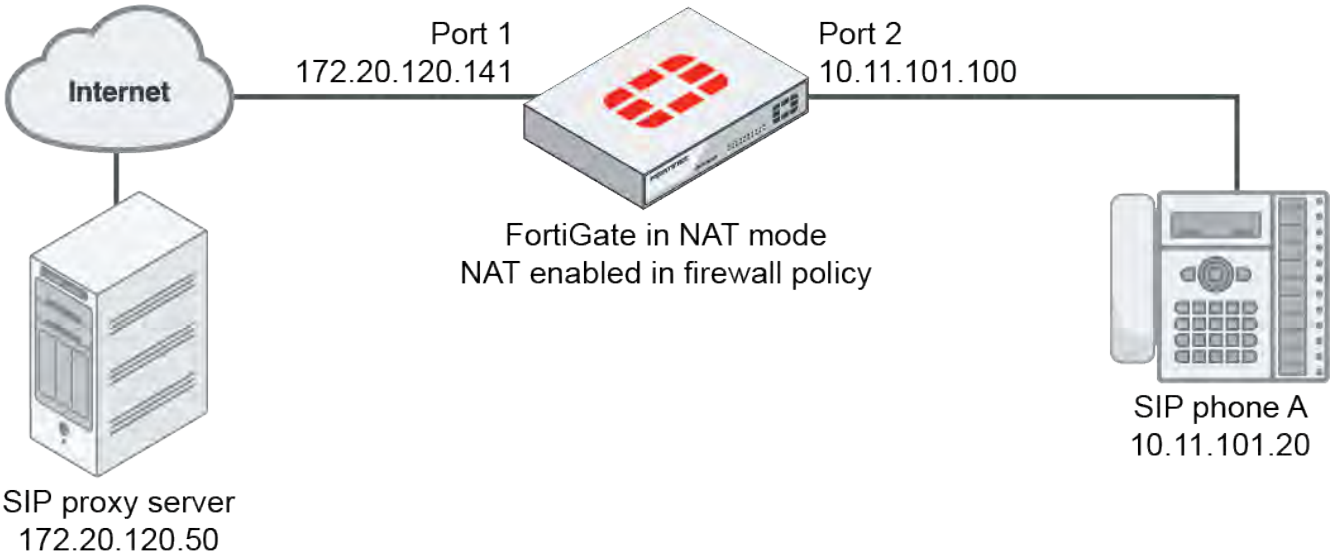
Assuming either SIP ALG or SIP Session Helper is enabled, configure the FortiGate with the following CLI commands:

```
config firewall vip
    edit "VIP_for_SIP_Server"
        set extip 172.20.120.50
        set extintf "port1"
        set mappedip "10.11.101.50"
    next
end
config firewall policy
    edit 1
        set srcintf "port1"
        set dstintf "port2"
        set srcaddr "all"
        set dstaddr "VIP_for_SIP_Server"
        set action accept
        set schedule "always"
        set service "SIP"
    next
end
```

Setting `service` to `SIP` and not `All` in the firewall policy can improve protection by restricting the data traffic passing through the FortiGate to the SIP call traffic only.

NAT

A FortiGate with SIP ALG or SIP Session Helper protects the SIP phones and the internal network from the internet, while SIP phones in the internal network need to register to the SIP server installed on the internet and establish calls through it.



One firewall policy needs to be configured with NAT enabled for SIP phones to send REGISTER messages through the FortiGate from port2 to port1.

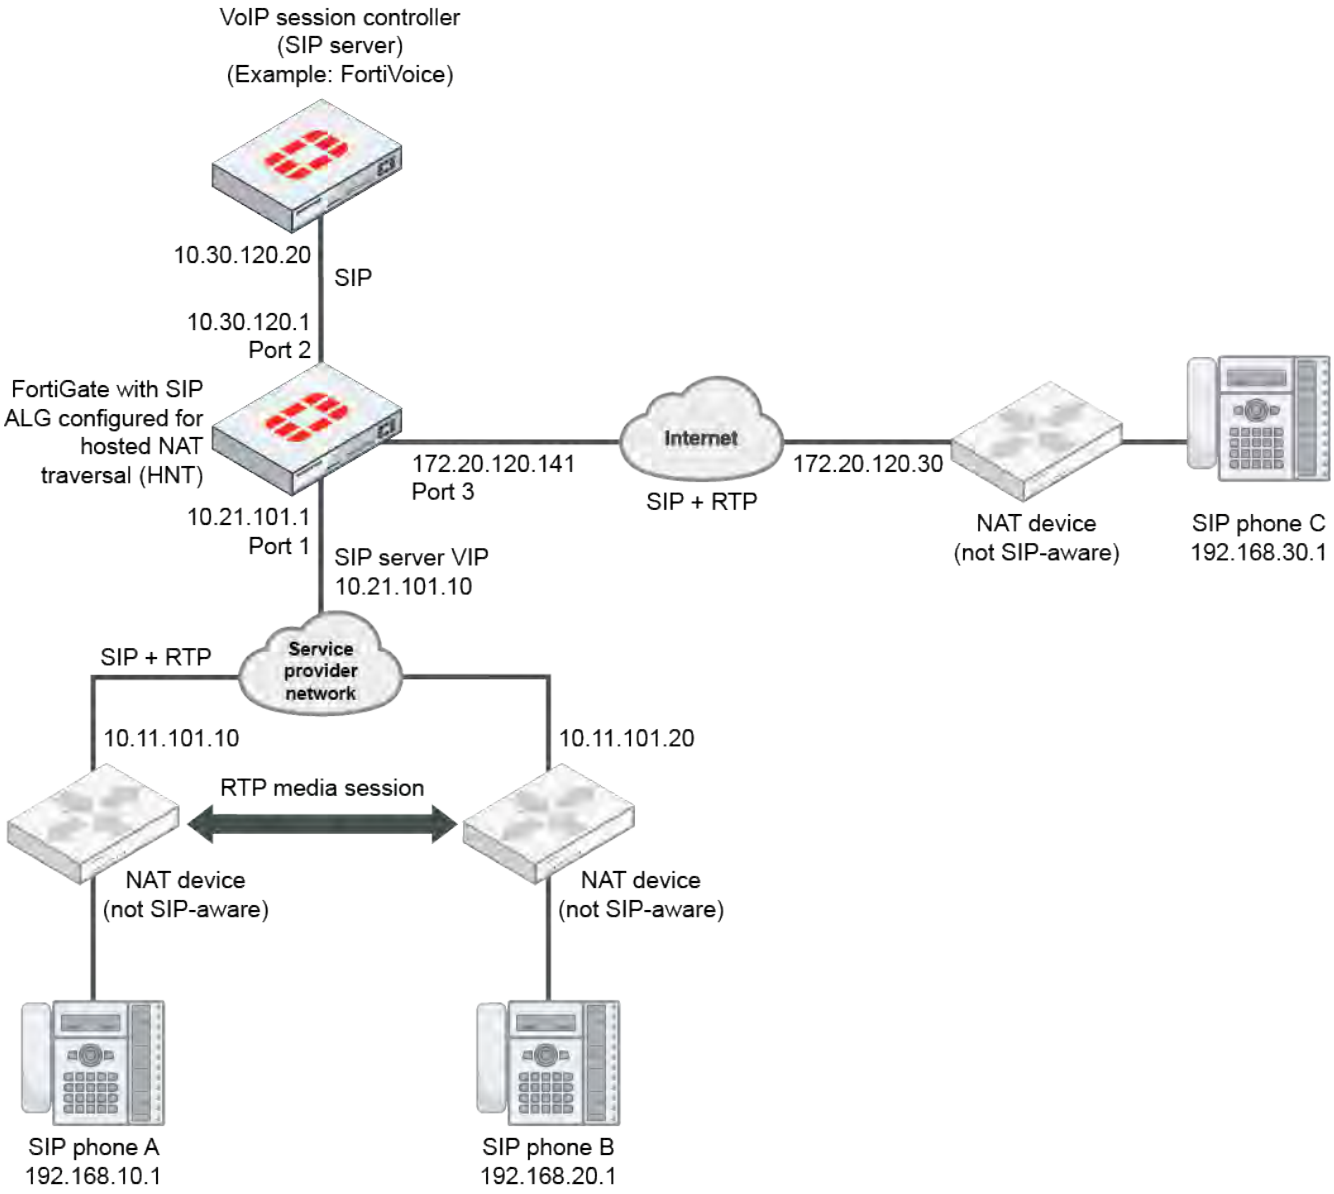
Assuming either SIP ALG or SIP Session Helper is enabled, configure the FortiGate with the following CLI commands:

```
config firewall policy
    edit 1
        set srcintf "port2"
        set dstintf "port1"
        set srcaddr "all"
        set dstaddr "all"
        set action accept
        set schedule "always"
        set service "SIP"
        set nat enable
    next
end
```

HNT

A FortiGate with SIP ALG or SIP Session Helper protects the SIP server from the internet, while SIP phones are in remote private networks behind NAT devices that are not aware of the SIP application.

For example, the SIP server is located in an ISP's service cloud that is protected by the FortiGate SIP ALG, and the SIP phones are installed in the home networks of the ISP's customers.



The SIP messages traversing the remote NAT devices might have their IP addresses translated by the NAT device at the network layer, but untranslated at the SIP application layer because those NAT devices are not aware of the SIP applications. This causes problems in a SIP session initiated process. Special configurations for the Hosted NAT Traversal (HNT) are required to resolve this issue.

To configure the FortiGate with HNT support for SIP phones A and B to set up calls with each other:

1. Identify port1 as the external interface:

```
config system interface
    edit "port1"
        set external enable
```

```
    next
end
```

2. Configure VIP for the SIP server:

```
config firewall vip
    edit "VIP_for_SIP_Server"
        set extip 10.21.101.10
        set extintf "port1"
        set mappedip "10.30.120.20"
    next
end
```

3. Configure a VoIP profile with HNT enabled:

```
config voip profile
    edit "hnt"
        config sip
            set hosted-nat-traversal enable
            set hnt-restrict-source-ip enable
        end
    next
end
```



`hosted-nat-traversal` must be enabled.

`hnt-restrict-source-ip` does not have to be enabled, but can be enabled to restrict the RTP packets' source IP to be the same as the SIP packets' source IP.

4. Apply the VoIP profile and VIP in a firewall policy for phone A and B to register and set up SIP calls through the FortiGate and SIP server:

```
config firewall policy
    edit 1
        set srcintf "port1"
        set dstintf "port2"
        set srcaddr "all"
        set dstaddr "VIP_for_SIP_Server"
        set action accept
        set schedule "always"
        set service "SIP"
        set utm-status enable
        set voip-profile "hnt"
        set nat enable
    next
end
```



`nat` must be enabled in the firewall policy.

SIP message inspection and filtering

SIP ALG provides users with security features to inspect and control SIP messages that are transported through FortiOS devices, including:

Netskope Exhibit 1017

- Verifying the SIP message syntax.
- Blocking particular types of SIP requests.
- Restricting the rate of particular SIP requests.

These features are configured in the VoIP profile:

```
config voip profile
    edit <voip_profile_name>
        config sip set ...
```

The VoIP profile can then be applied to a firewall policy to process the SIP call traffic.

SIP message syntax inspection

For syntax verification, the following attributes are available for configuration in the VoIP profile to determine what action is taken when a specific syntax error or attack based on invalid syntax is detected. For example, the action can be set to pass or discard it.

```
malformed-request-line
malformed-header-via
malformed-header-from
malformed-header-to
malformed-header-call-id
malformed-header-cseq
malformed-header-rack
malformed-header-rseq
malformed-header-contact
malformed-header-record-route
malformed-header-route
malformed-header-expires
malformed-header-content-type
malformed-header-content-length
malformed-header-max-forwards
malformed-header-allow
malformed-header-p-asserted-identity
malformed-header-sdp-v
malformed-header-sdp-o
malformed-header-sdp-s
malformed-header-sdp-i
malformed-header-sdp-c
malformed-header-sdp-b
malformed-header-sdp-z
malformed-header-sdp-k
malformed-header-sdp-a
malformed-header-sdp-t
malformed-header-sdp-r
malformed-header-sdp-m
```

SIP message blocking

The following options are available in the VoIP profile to block SIP messages:

```
block-long-lines
block-unknown
block-ack
```

```
block-bye
block-cancel
block-info
block-invite
block-message
block-notify
block-options
block-prack
block-publish
block-refer
block-register
block-subscribe
block-update
block-geo-red-options
```

SIP message rate limiting

The rate of certain types of SIP requests that are passing through the SIP ALG can be restricted :

```
register-rate
invite-rate
subscribe-rate
message-rate
notify-rate
refer-rate
update-rate
options-rate
ack-rate
prack-rate
info-rate
publish-rate
bye-rate
cancel-rate
```

SIP pinholes

When SIP ALG processes a SIP call, it usually opens pinholes for SIP signaling and RTP/RTCP packets. NAT usually takes place during the process at both the network and SIP application layers. SIP ALG ensures that, with NAT happening, corresponding SIP and RTP/RTCP pinholes are created during the process when it is necessary for call sessions to be established through FortiOS devices.

By default, SIP ALG manages pinholes automatically, but some special configurations can be used to restrict the pinholes if required.

SIP pinhole restriction

By default, the *strict-register* attribute is enabled. When enabled, after a SIP endpoint registers to the SIP server through a firewall policy on the FortiOS device, only the SIP messages sent from the same IP address as the SIP server are allowed to pass through the SIP pinhole that is created in the FortiOS device to reach the SIP endpoints. If the attribute is disabled, SIP messages from any IP addresses can pass through the pinhole created after the registration.

```
config voip profile
    edit "voip-profile-name"
        config sip
            set strict-register [enable|disable]
            ...
        end
    next
end
```

RTP/RTCP pinhole restriction

In a SIP call through SIP ALG, the NATed RTP/RTCP port range is 5117 to 65533 by default. If required, the port range can be restricted.

```
config voip profile
    edit "voip-profile-name"
        config sip
            set nat-port-range <start_port_number>-<end_port_number>
            ...
        end
    next
end
```

In a SIP call session, the RTP port number is usually an even number and the RTCP port number is an odd number that is one more than the RTP port number. It is best practice to configure `start_port_number` to an even number, and `end_port_number` to an odd number, for example:

```
config voip profile
    edit "voip-profile-name"
        conf sip
            set nat-port-range 30000-39999
        end
    next
end
```

SIP over TLS

Some SIP phones and servers can communicate using TLS to encrypt the SIP signaling traffic. To allow SIP over TLS calls to pass through the FortiGate, the encrypted signaling traffic must be unencrypted and inspected. The FortiGate SIP ALG intercepts, unencrypts , and inspects the SIP packets, which are then re-encrypted and forwarded to their destination.

The SIP ALG only supports full mode TLS. This means that the SIP traffic between SIP phones and the FortiGate, and between the FortiGate and the SIP server, is always encrypted. The highest TLS version supported by SIP ALG is TLS 1.2.

To enable SIP over TLS support, the SSL mode in the VoIP profile must be set to `full`. The SSL server and client certificates can be provisioned so that the FortiGate can use them to establish connections to SIP phones and servers, respectively.

Netskope Exhibit 1017

To configure SIP over TLS:

1. Configure a VoIP profile with SSL enabled:

```
config voip profile
    edit "tls"
        config sip
            set ssl-mode full
            set ssl-client-certificate "ssl_client_cert"
            set ssl-server-certificate "ssl_server_cert"
        end
    next
end
```

   The `ssl_server_cert`, `ssl_client_cert`, and key files can be generated using a certification tool, such as OpenSLL, and imported to the local certificate store of the FortiGate from *System* > *Certificates* in the GUI. Existing local certificates in the certificate store can also be used. As always for TLS connections, the certificates used must be verified and trusted at the other end of the connection when required.

   For example, the CA certificate of the SIP server's certificate should be imported to the FortiGate as an external CA certification, such that the FortiGate can use it to verify the SIP server's certificate when setting up the TLS connection. The CA certificate configured as the `ssl_server_cert` should be installed as the trusted certificate on the SIP phones. The deployment of the certificates across the network depends on the SIP client and server devices that are used in the system.

2. Apply the profile to the firewall policy:

```
config firewall policy
    edit 1
        set srcintf "port1"
        set dstintf "port2"
        set srcaddr "all"
        set dstaddr "vip_sip_server"
        set action accept
        set schedule "always"
        set service "SIP"
        set utm-status enable
        set voip-profile "tls"
    next
end
```

Custom SIP RTP port range support

The `nat-port-range` variable is used to specify a port range in the VoIP profile to restrict the NAT port range for real-time transport protocol/real-time transport control protocol (RTP/RTCP) packets in a session initiation protocol (SIP) call session that is handled by the SIP application layer gateway (ALG) in a FortiGate device.

When NAT is enabled, or VIP is used in a firewall policy for SIP ALG to handle a SIP call session established through a FortiGate device, the SIP ALG can perform NAT to translate the ports used for the RTP/RTCP packets when they are flowing through the device between the external and internal networks.

You can control the translated port range for RTP/RTCP packets using the CLI:

```
config voip profile
   edit <profile-name>
      config sip
         set nat-port-range <port range>
      end
```


Netskope Exhibit 1017

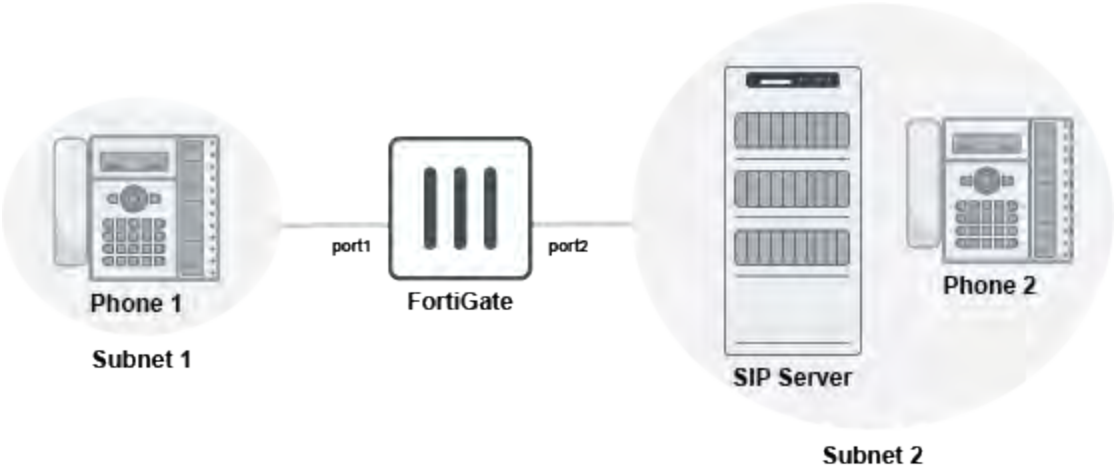
```
   next
end
```

Command	Description
nat-port-range <port range>	The NAT port range (minimum port number = 5117, default = 5117-65535).

Example

In this example, Phone1 is in subnet_1, and the SIP server and phone are in subnet_2. All SIP signaling messages and RTP/RTCP packets go through the SIP Server. The RTP/RTCP ports on Phone1 are configured as 17078/17079.

The FortiGate administrator wants to use NAT for the port 17078/17079 to 30000/30001. As a result, all RTP/RTCP packets going out of port2 have source ports of 30000/30001, and all RTP/RTCP packets going into port2 also have destination ports of 30000/30001, which is specified in `nat-port-range`.



To configure the custom port range:

```
config voip profile
   edit "natPortRange"
        config sip
           set nat-port-range 30000-30001
        end
   next
end
configure firewall policy
   edit 1
        set srcintf port1
        set  dstintf port2
        set srcaddr all
        set dstaddr all
        set service SIP
        set action accept
        set schedule always
        set voip-profile natPortRange
        set nat enable
end
```

If phone1 and phone2 are registered to the SIP server, and they establish a call session between them through the FortiGate and the SIP server, then the RTP/RTCP ports 17078/17079 of phone1 will be translated to ports 30000/30001 at the FortiGate unit based on the NAT port range setting. That is, the RTP/RTCP packets egressing port2 of the Fortigate will have source ports of 30000/30001, and the RTP/RTCP packets ingressing port2 will have destination ports of 30000/30001.

Voice VLAN auto-assignment

You can leverage LLDP-MED to assign voice traffic to the desired voice VLAN. After detection and setup, the IP phone on the network is segmented to its own VLAN for policy, prioritization, and reporting. The LLDP reception capabilities in FortiOS have been extended to support LLDP-MED assignment for voice, voice signaling, guest, guest voice signaling, softphone, video conferencing, streaming video, and video signaling.

You can configure this feature using the following steps:

1. Setting up the VLAN for the voice device
2. Setting up the DHCP server for the voice VLAN
3. Setting up the LLDP network policy
4. Enabling LLDP on the physical interface that the VLAN belongs to
5. Applying the LLDP network policy on the physical interface
6. Confirming that the VLAN was assigned

To set up the VLAN for the voice device:

```
config system interface
    edit "vlan_100"
        set vdom "root"
        set ip 192.168.1.99 255.255.255.0
        set alias "voice_vlan"
        set device-identification enable
        set role lan
        set snmp-index 25
        set interface "port10"
        set vlanid 100
    next
end
```

To set up the DHCP server for the voice VLAN:

```
config system dhcp server
    edit 1
        set dns-service default
        set default-gateway 192.168.1.99
        set netmask 255.255.255.0
        set interface "vlan_100"
        config ip-range
            edit 1
                set start-ip 192.168.1.110
                set end-ip 192.168.1.210
            next
        end
```

```
    next
end
```

To set up the LLDP network policy:

```
config system lldp network-policy
    edit "1"
        config voice
            set status enable
            set tag dot1q
            set vlan 100
        end
    next
end
```

To enable LLDP on the physical interface that the VLAN belongs to:

```
config system interface
    edit "port10"
        set vdom "root"
        set type physical
        set lldp-reception enable
        set lldp-transmission enable
        set snmp-index 14
    next
end
```

To apply the LLDP network policy on the physical interface:

```
config system interface
    edit "port10"
        set lldp-network-policy "1"
    next
end
```

To confirm that the VLAN was assigned as expected:

1. Connect an IP phone to the network.
2. Check the IP address on the phone.
   The IP address should belong to the voice VLAN.
3. Sniff on the FortiGate incoming interface to see if traffic from the IP phone has the desired VLAN tag.
   In the example commands above, the voice VLAN was configured as VLAN 100. Therefore, voice traffic from the IP phone should be in VLAN 100.

ICAP

Internet Content Adaptation Protocol (ICAP) is an application layer protocol that is used to offload tasks from the firewall to separate, specialized servers. For more information see RFC 3507.

ICAP profiles can only be applied to policies that use proxy-based inspection. If you enable ICAP in a policy, HTTP and HTTPS (if HTTPS inspection is supported) traffic that is intercepted by the policy is transferred to the ICAP server

specified by the selected ICAP profile. Responses from the ICAP server are returned to the FortiGate, and then forwarded to their destination.



By default, *ICAP* is not visible in the GUI. See Feature visibility on page 694 for instructions on making it visible.

To configure ICAP:

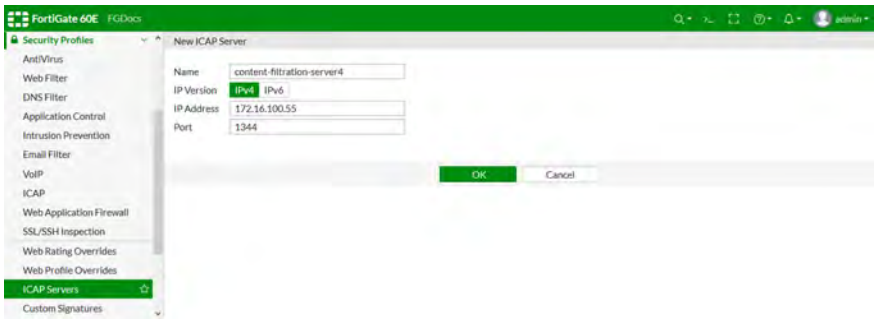
1. Set up your ICAP server.
2. On the FortiGate, add an ICAP server.
3. Create an ICAP profile.
4. Use the ICAP profile in a firewall policy that covers the traffic that needs to be offloaded to the ICAP server.

ICAP configuration example

In this example, the ICAP server performs proprietary content filtering on HTTP and HTTPS requests. If the content filter is unable to process a request, then the request is blocked. Streaming media is not considered by the filter, so it is allowed through and is not processed.

To add the ICAP server to the FortiGate in the GUI:

1. Go to *Security Profiles* > *ICAP Servers*.
2. Click *Create New*.
3. In the *Name* field, enter a name for the ICAP server, such as *content-filtration-server4*.
4. Select the *IP Version*.
5. In the *IP Address* field, enter the IP address of the ICAP server.
6. In the *Port* field, enter a new port number if required. The default value is *1344*.



7. Click *OK*.

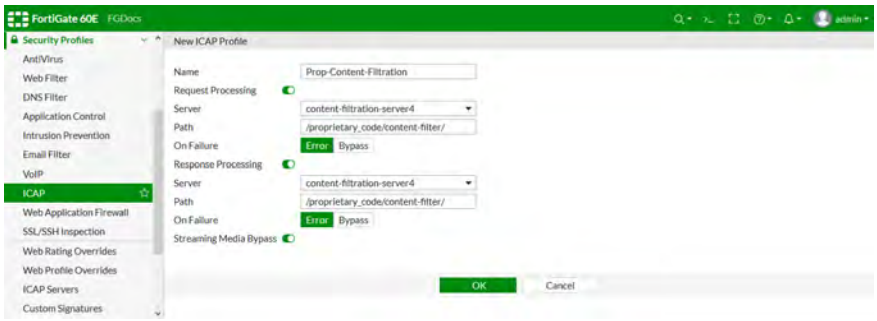


The maximum number of concurrent connections to ICAP server can be configured in the CLI. The default setting is 100 connections.

Netskope Exhibit 1017

To create an ICAP profile in the GUI:

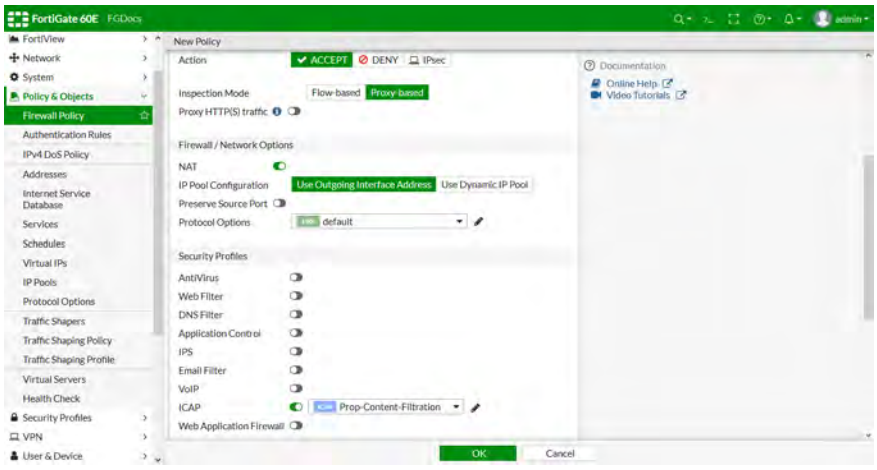
1. Go to *Security Profiles* > *ICAP*.
2. Click *Create New*.
3. In the *Name* field, enter a name for the ICAP profile, such as *Prop-Content-Filtration*.
4. Enable *Request Processing* then set the following:
   - *Server* - Select the ICAP server. In this example, select *content-filtration-server4*
   - *Path* - The path to the processing component on the server, such as */proprietary_code/content-filter/*.
   - *On Failure* - Select *Error* to block the request. If the message cannot be processed, it will not be blocked.
5. Enable *Response Processing* then set the following:
   - *Server* - Select the ICAP server: *content-filtration-server4*
   - *Path* - The path to the processing component on the server, such as */proprietary_code/content-filter/*.
   - *On Failure* - Select *Error* to block the request. If the message cannot be processed, it will not be blocked.
6. Enable *Streaming Media Bypass* to not offload streaming media to the ICAP server.



7. Click *OK*.

To add the ICAP profile to a policy in the GUI:

1. Go to *Policy & Objects* > *Firewall Policy*.
2. Click *Create New*.
3. Configure the policy as needed to apply to the required traffic.
4. Set *Inspection Mode* to *Proxy-based*.
5. Under *Security Profiles*, enable *ICAP* and select the ICAP server.



6. Click *OK*.

To configure the ICAP setup in the CLI:

1. Add the ICAP server:

```
config icap server
    edit "content-filtration-server4"
        set ip-version 4
        set ip-address 172.16.100.55
        set port 1344
        set max-connections 200
    next
end
```

2. Create the ICAP profile:

```
config icap profile
    edit "Prop-Content-Filtration"
        set request enable
        set response enable
        set streaming-content-bypass enable
        set request-server "content-filtration-server4"
        set response-server "content-filtration-server4"
        set request-failure error
        set response-failure error
        set request-path "/proprietary_code/content-filter/"
        set response-path "/proprietary_code/content-filter/"
        set methods delete get head options post put trace other
    next
end
```

3. Add the ICAP profile to a policy:

```
config firewall policy
    edit 5
        set name "icap_filter3"
        set srcintf "virtual-wan-link"
        set dstintf "virtual-wan-link"
        set srcaddr "FABRIC_DEVICE"
        set dstaddr "FABRIC_DEVICE"
        set dstaddr-negate enable
        set action accept
        set schedule "always"
        set service "ALL"
        set utm-status enable
        set inspection-mode proxy
        set ssl-ssh-profile "certificate-inspection"
        set icap-profile "Prop-Content-Filtration"
        set logtraffic disable
        set fsso disable
        set nat enable
    next
end
```

ICAP response filtering

ICAP HTTP responses can be forwarded or bypassed based on the HTTP header value and status code.

When configuring the ICAP profile, if `response` is enabled, the `respmod-default-action` option can be configured:

Netskope Exhibit 1017

- If `respmod-default-action` is set to forward, FortiGate will treat every HTTP response, and send ICAP requests to the ICAP server.
- If `respmod-default-action` is set to bypass, FortiGate will only send ICAP requests if the HTTP response matches the defined rules, and the rule's action is set to forward.

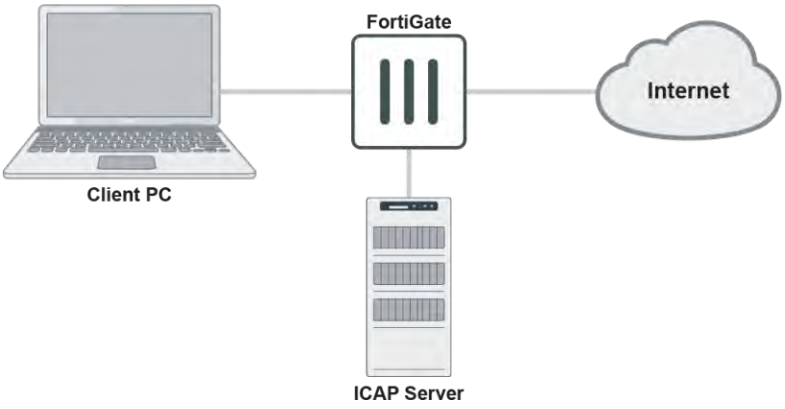
When configuring a response rule:

- The `http-resp-status-code` option is configured to specific HTTP response codes. If the HTTP response has any one of the configured values, then the rule takes effect.
- Multiple header value matching groups can be configured. If the header value matches one of the groups, then the rule takes effect.
- If both status codes and header values are specified in a rule, the response must match at least one of each.

The UTM ICAP log category is used for logging actions when FortiGate encounters errors with the ICAP server, such as no service, unreachable, error response code, or timeout. If an error occurs, a traffic log and an associated UTM ICAP log will be created.

Example

The FortiGate acts as a gateway for the client PC and connects to a reachable ICAP server. The ICAP server can be in NAT, transparent, or proxy mode.



In this example, client request HTTP responses will be forwarded to the ICAP server from all hosts if they have an HTTP status code of 200, 301, or 302, and have `content-type: image/jpeg` in the their header.

To configure an ICAP profile with HTTP response rules:

```
config icap profile
    edit "icap_profile2"
        set request disable
        set response enable
        set streaming-content-bypass disable
        set preview disable
        set response-server "icap_server1"
        set response-failure error
        set response-path ''
        set methods delete get head options post put trace other
        set response-req-hdr disable
        set respmod-default-action bypass
        config respmod-forward-rules
```

```
            edit "rule2"
                set host "all"
                set action forward
                set http-resp-status-code 200 301 302
                config header-group
                    edit 2
                        set header-name "content-type"
                        set header "image/jpeg"
                    next
                end
            next
        end
    next
end
```

To view the logs if an error occurs:

1. View the traffic log:

```
# execute log filter category 0
# execute log display
1 logs found.
1 logs returned.

 1: date=2019-10-25 time=17:43:47 logid="0000000013" type="traffic" subtype="forward"
level="notice" vd="vdom1" eventtime=1572050627037314464 tz="-0700" srcip=10.1.100.145
srcport=47968 srcintf="port1" srcintfrole="undefined" dstip=172.16.200.46 dstport=80
dstintf="port2" dstintfrole="undefined" poluuid="a4d5324e-f6c3-51e9-ce2d-f360994fb547"
sessionid=43549 proto=6 action="close" policyid=1 policytype="policy" service="HTTP"
dstcountry="Reserved" srccountry="Reserved" trandisp="snat" transip=172.16.200.1
transport=47968 duration=1 sentbyte=485 rcvdbyte=398 sentpkt=6 rcvdpkt=5
appcat="unscanned" wanin=478 wanout=165 lanin=165 lanout=165 utmaction="block"
counticap=1 crscore=5 craction=262144 crlevel="low" utmref=65532-0
```

2. View the UTM ICAP log:

```
# execute log filter category 20
# execute log display
1 logs found.
1 logs returned.

 1: date=2019-10-25 time=17:43:46 logid="2000060000" type="utm" subtype="icap"
eventtype="icap" level="warning" vd="vdom1" eventtime=1572050626010097145 tz="-0700"
msg="Request blocked due to ICAP server error" service="HTTP" srcip=10.1.100.145
dstip=172.16.200.46 srcport=47968 dstport=80 srcintf="port1" srcintfrole="undefined"
dstintf="port2" dstintfrole="undefined" policyid=1 sessionid=43549 proto=6
action="blocked" profile="icap_profile1" url="/icap_test/"
```

The logs show that, in this case, the ICAP services stopped before the access. When the client tried to access HTTP and ICAP took effect, the FortiGate sent the ICAP request to the ICAP server and received an error. The client sees a *502 Bad Gateway* message, and FortiGate writes the two logs. In the GUI, the logged traffic is displayed as *Result: Deny: UTM Blocked*.

Netskope Exhibit 1017

Web application firewall

Web application firewall (WAF) profiles can detect and block known web application attacks. You can configure WAF profiles to use signatures and constraints to examine web traffic. You can also enforce an HTTP method policy, which controls the HTTP method that matches the specified pattern.

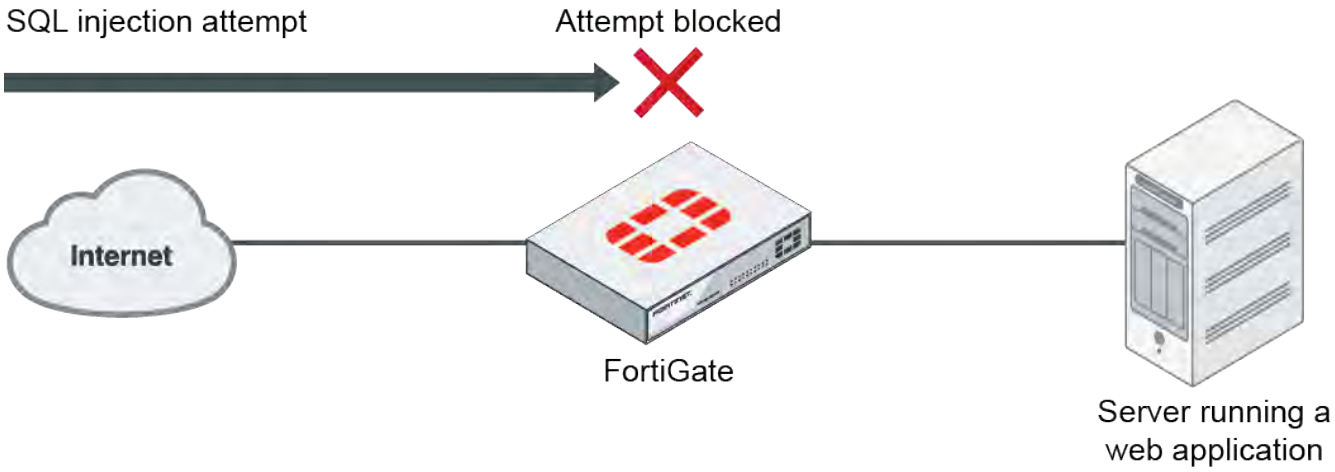
You can customize the default profile, or you can create your own profile to apply access rules and HTTP protocol constraints to traffic. You can apply WAF profiles to firewall policies when the inspection mode is set to proxy-based.

The following topic provides information about WAF profiles:

- Protecting a server running web applications on page 1021

Protecting a server running web applications

You can use a web application firewall profile to protect a server that is running a web application, such as webmail.



Web application firewall profiles are created with a variety of options called signatures and constraints. Once these options are enabled, the action can be set to allow, monitor, or block. The severity can be set to high, medium, or low.

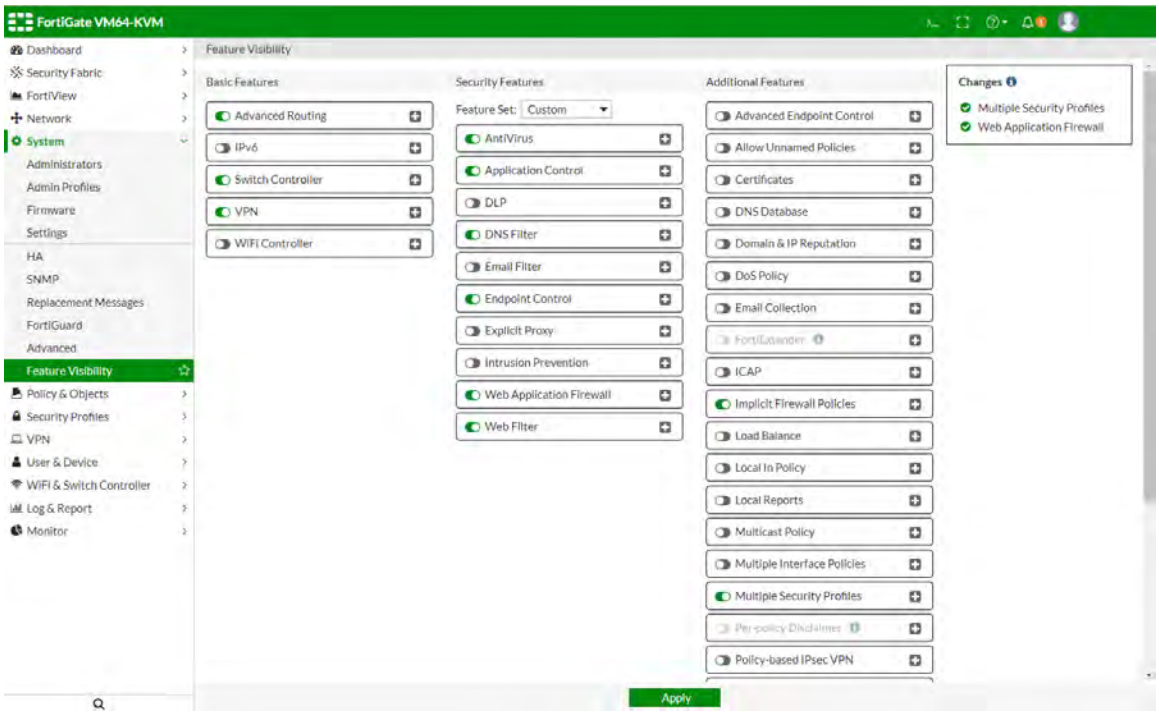
In the following example, the default profile will be targeted to block SQL injection attempts and generic attacks.

The web application firewall feature is only available when the policy inspection mode is proxy-based.

To protect a server running web applications:

1. Enable the web application firewall:
   a. Go to *System* > *Feature Visibility*.
   b. Under *Security Features*, enable *Web Application Firewall*.
   c. Under *Additional Features*, click *Show More* and enable *Multiple Security Profiles*.

d. Click *Apply*.



2. Edit the default web application firewall profile:
Trojans and *Known Exploits* are blocked by default.

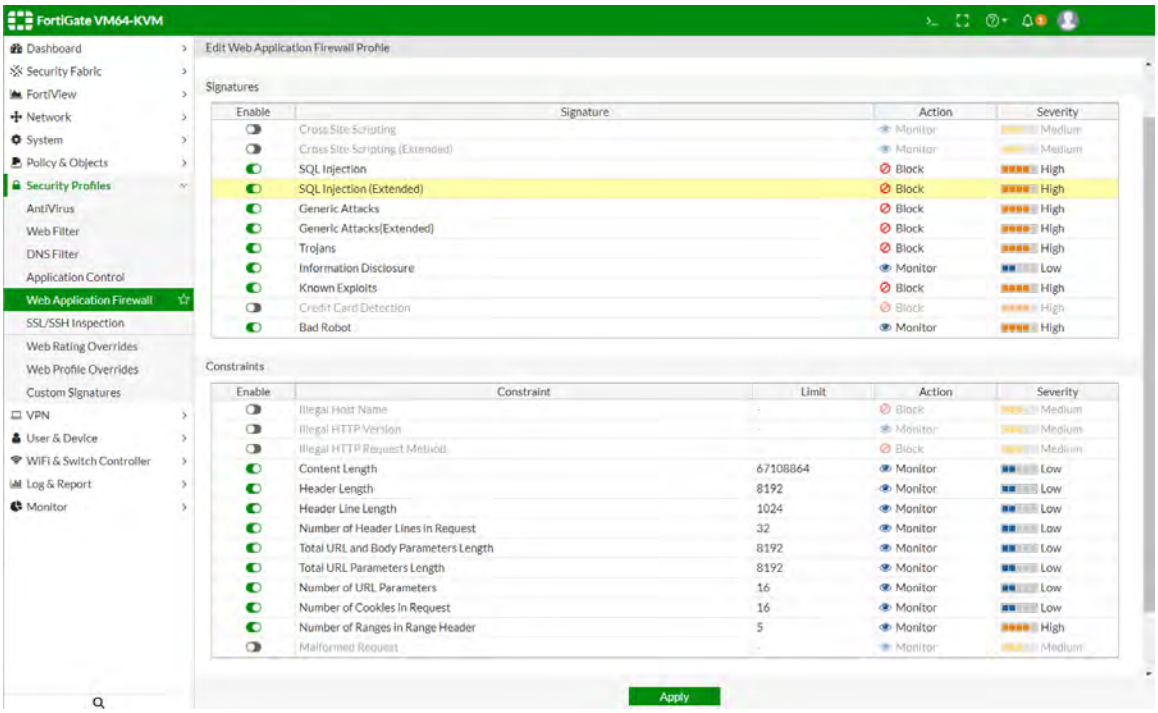
a. Go to *Security Profiles* > *Web Application Firewall*.

b. Edit the *default* profile signature:

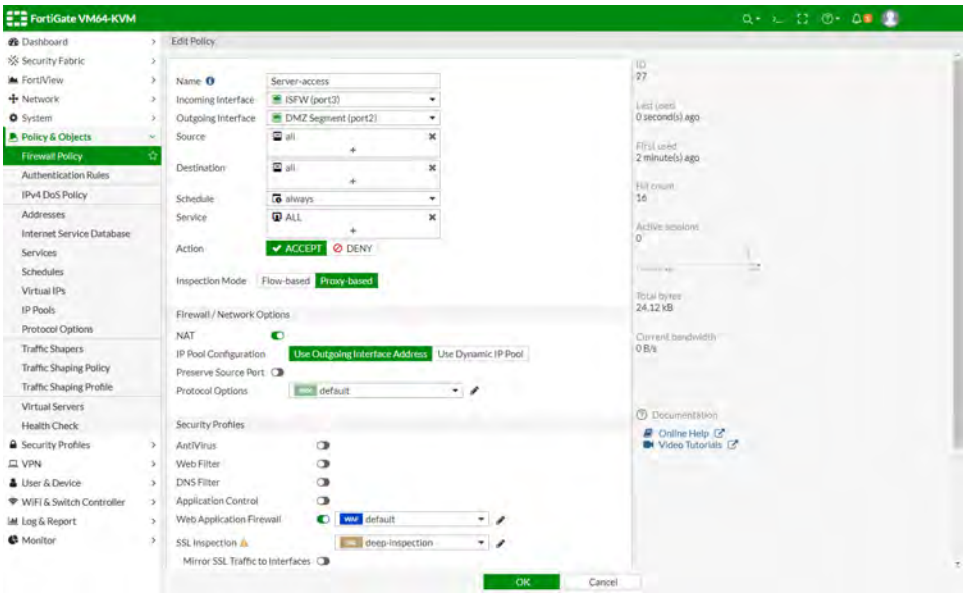
i. Enable *SQL Injection (Extended)* and *Generic Attacks (Extended)*.

ii. For both signatures, set the *Action* to *Block* and the *Severity* to *High*.

iii. Click *Apply*.



3. Apply the profile to a security policy:
   a. Go to *Policy & Objects* > *Firewall Policy*.
   b. Edit the policy that allows access to the web server:
      i. Under *Firewall / Network Options*, select the appropriate *Protocol Option*.
      ii. Under *Security Profiles*, enable *Web Application Firewall* and set it to use the *default* profile.
      iii. Set the *SSL Inspection* to use the *deep-inspection* profile.
      iv. Click *OK*.



4. Verify that the web application firewall blocks traffic:
   a. Use the following URL to simulate an attack on your web server and substitute the IP address of your server:
   `http://<server IP>/index.php?username=1'%20or%20'1'%20=%20'1&password=1'%20or%20'1'%20=%20'1`
   An error message appears, stating that the web application firewall has blocked the traffic:

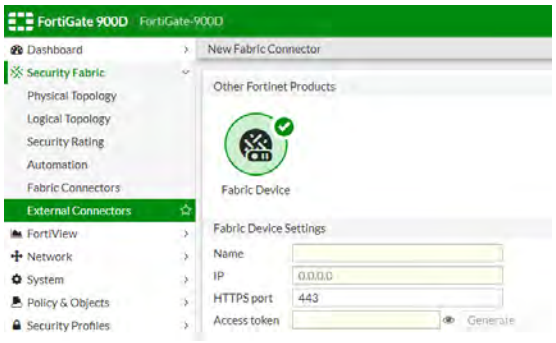


Offloading to a FortiWeb

If you have a FortiWeb, you may be able to offload the functions of the web application control to your FortiWeb. To find out if this option is available, refer to the FortiOS or FortiWeb Release Notes for information about device compatibility.

To offload to a FortiWeb:

1. Go to *Security Fabric* > *Fabric Connectors*.
2. Click *Create New*, and click *Fabric Device*.
3. Enter the following for the device:
   a. Name (`FortiWeb`)
   b. FortiWeb IP address
   c. HTTPS service port



4. Click *Generate*.
5. Enter your credentials to generate the access token.
6. Click *OK*.

SSL & SSH Inspection

Secure sockets layer (SSL) content scanning and inspection allows you to apply antivirus scanning, web filtering, and email filtering to encrypted traffic. You can apply SSL inspection profiles to firewall policies.

FortiOS includes four preloaded SSL/SSH inspection profiles, three of which are read-only and can be cloned:

- *certificate-inspection*
- *deep-inspection*
- *no-inspection*

The *custom-deep-inspection* profile can be edited, or you can create your own SSL/SSH inspection profiles.

Deep inspection (also known as SSL/SSH inspection) is typically applied to outbound policies where destinations are unknown. Depending on your policy requirements, you can configure the following:

- Which CA certificate will be used to decrypt the SSL encrypted traffic
- Which SSL protocols will be inspected
- Which ports will be associated with which SSL protocols for inspection
- Whether or not to allow invalid SSL certificates
- Whether or not SSH traffic will be inspected
- Which addresses or web category allowlists can bypass SSL inspection

The following topics provide information about SSL & SSH Inspection:

- Certificate inspection on page 1025
- Deep inspection on page 1026
- Protecting an SSL server on page 1029
- Ignoring the AUTH TLS command on page 1030
- SSH traffic file scanning on page 1031
- Redirect to WAD after handshake completion on page 1034

Certificate inspection

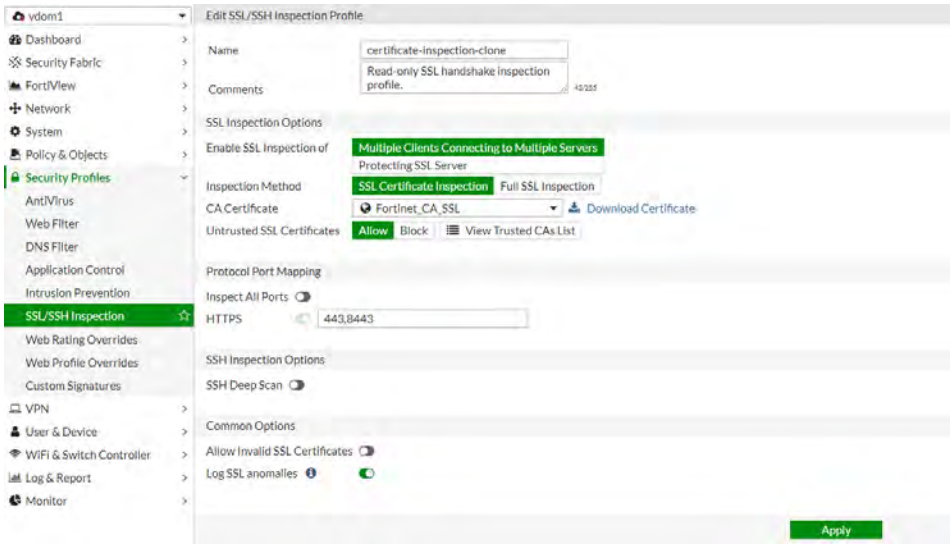
FortiGate supports certificate inspection. The default configuration has a built-in *certificate-inspection* profile which you can use directly. When you use certificate inspection, the FortiGate only inspects the headers up to the SSL/TLS layer.

If you do not want to deep scan for privacy reasons but you want to control web site access, you can use *certificate-inspection*.

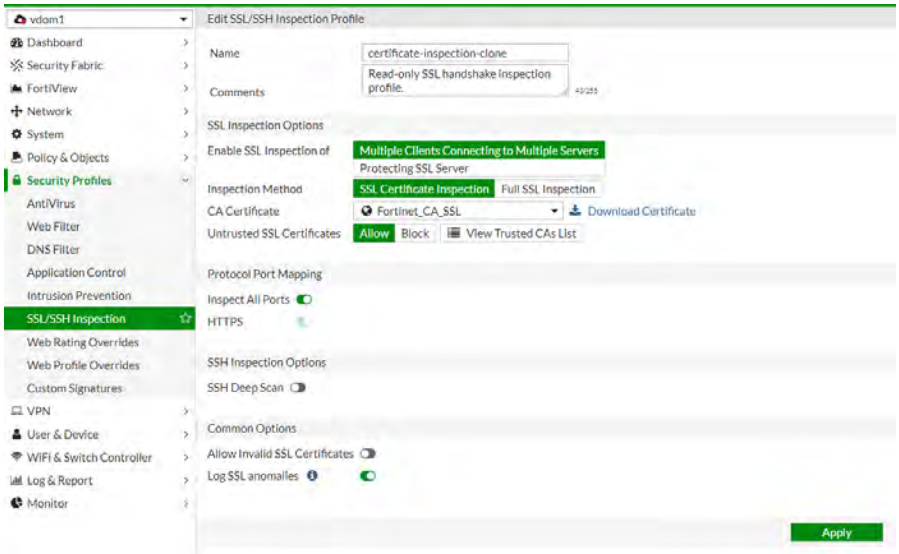
Inspect non-standard HTTPS ports

The built-in *certificate-inspection* profile is read-only and only listens on port 443. If you want to make changes, you must create a new certificate inspection profile.

If you know the non-standard port that the web server uses, such as port 8443, you can add this port to the *HTTPS* field.



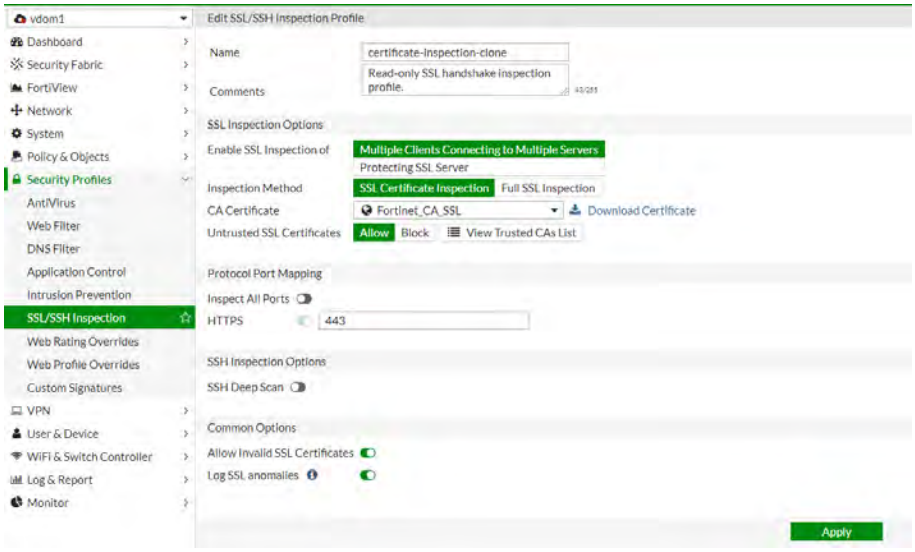
If you do not know which port is used in the HTTPS web server, you can select *Inspect All Ports*.



Block untrusted or allow invalid certificate

The default setting in the *certificate-inspection* profile is to block invalid certificates and allow untrusted certificates.

For example, the server certificate has expired but you still want to access this server until you have a new server certificate. But because certificate inspection cannot do an exemption, you have to allow the invalid certificate in your SSL profile. This means you need to create a new certificate inspection profile using the built-in read-only *certificate-inspection*.



Deep inspection

You can configure address and web category allowlists to bypass SSL deep inspection.

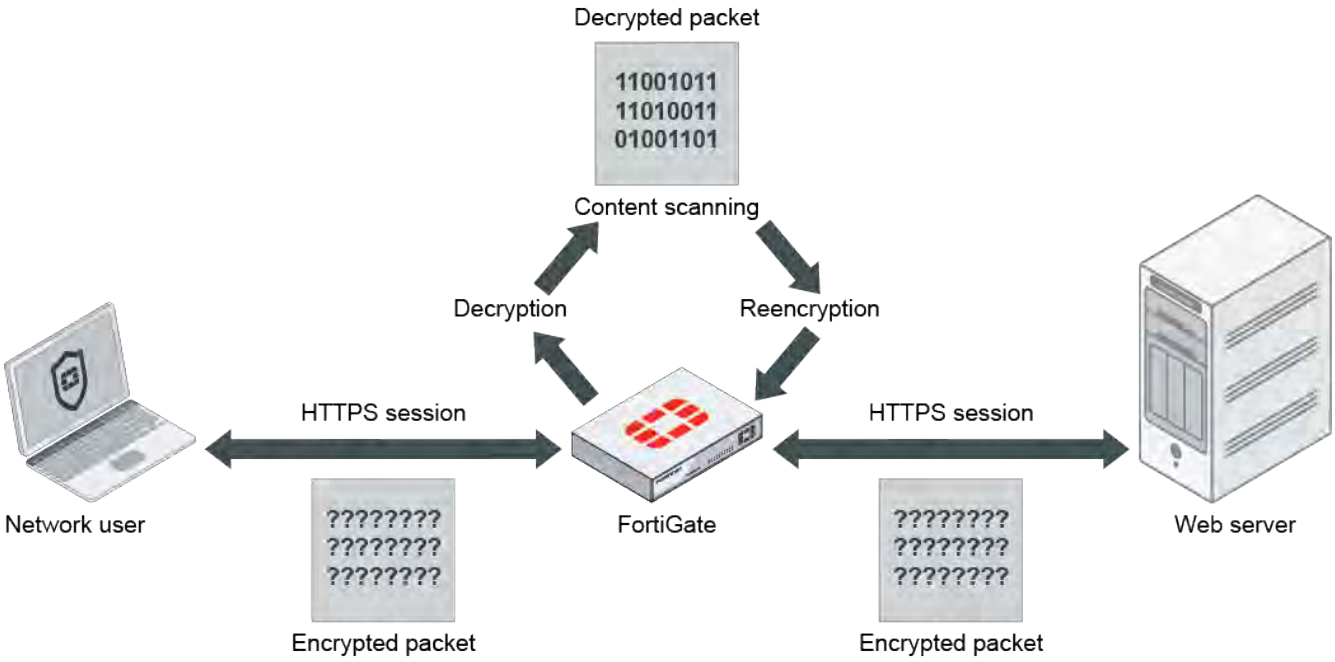
Reasons for using deep inspection

While Hypertext Transfer Protocol Secure (HTTPS) offers protection on the Internet by applying Secure Sockets Layer (SSL) encryption to web traffic, encrypted traffic can be used to get around your network's normal defenses.

For example, you might download a file containing a virus during an e-commerce session, or you might receive a phishing email containing a seemingly harmless download that, when launched, creates an encrypted session to a command and control (C&C) server and downloads malware onto your computer. Because the sessions in these attacks are encrypted, they might get past your network's security measures.

When you use deep inspection, the FortiGate impersonates the recipient of the originating SSL session, then decrypts and inspects the content to find threats and block them. It then re-encrypts the content and sends it to the real recipient.

Deep inspection not only protects you from attacks that use HTTPS, it also protects you from other commonly-used SSL-encrypted protocols such as SMTPS, POP3S, IMAPS, and FTPS.



Browser messages when using deep inspection

When FortiGate re-encrypts the content, it uses a certificate stored on the FortiGate such as *Fortinet_CA_SSL*, *Fortinet_CA_Untrusted*, or your own CA certificate that you uploaded.

Because there is no *Fortinet_CA_SSL* in the browser trusted CA list, the browser displays an untrusted certificate warning when it receives a FortiGate re-signed server certificate. To stop the warning messages, trust the FortiGate-trusted CA *Fortinet_CA_SSL* and import it into your browser.

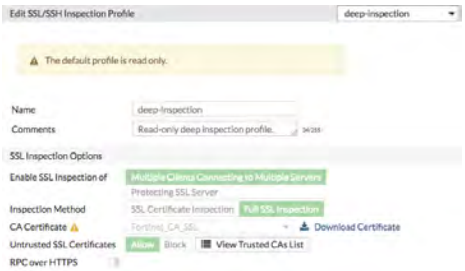
After importing *Fortinet_CA_SSL* into your browser, if you still get messages about untrusted certificate, it must be due to *Fortinet_CA_Untrusted*. **Never** import the *Fortinet_CA_Untrusted* certificate into your browser.

To import *Fortinet_CA_SSL* into your browser:

1. On the FortiGate, go to *Security Profiles* > *SSL/SSH Inspection* and select *deep-inspection*.
2. The default *CA Certificate* is *Fortinet_CA_SSL*.

Netskope Exhibit 1017

3. Select *Download Certificate*.

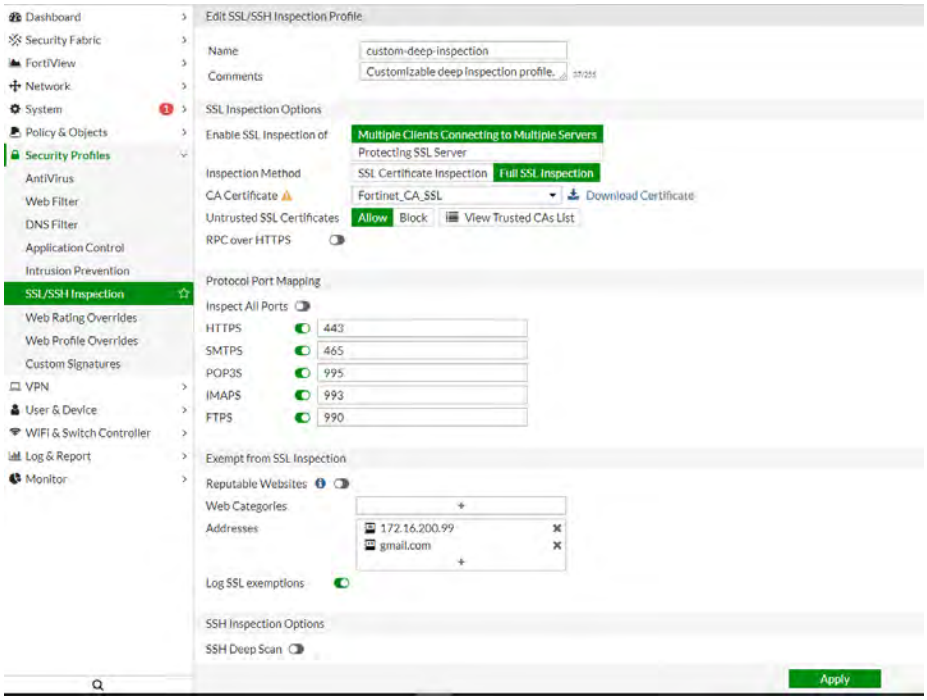


4. On the client PC, double-click the certificate file and select *Open*.
5. Select *Install Certificate* to launch the *Certificate Import Wizard* and use the wizard to install the certificate into the *Trusted Root Certificate Authorities* store.
   If a security warning appears, select *Yes* to install the certificate.

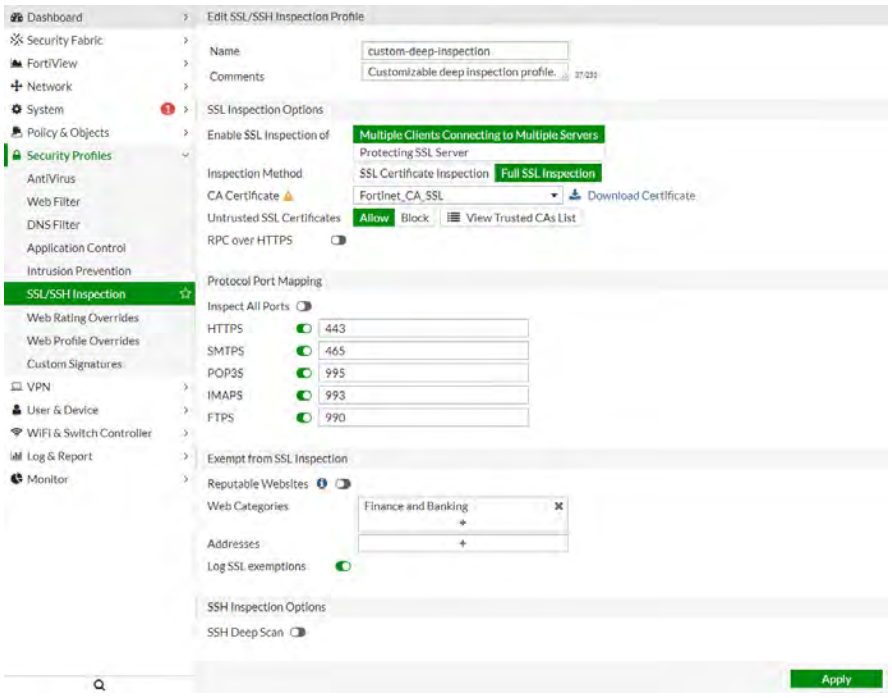
Exempt web sites from deep inspection

If you do not want to apply deep inspection for privacy or other reasons, you can exempt the session by address, category, or allowlist.

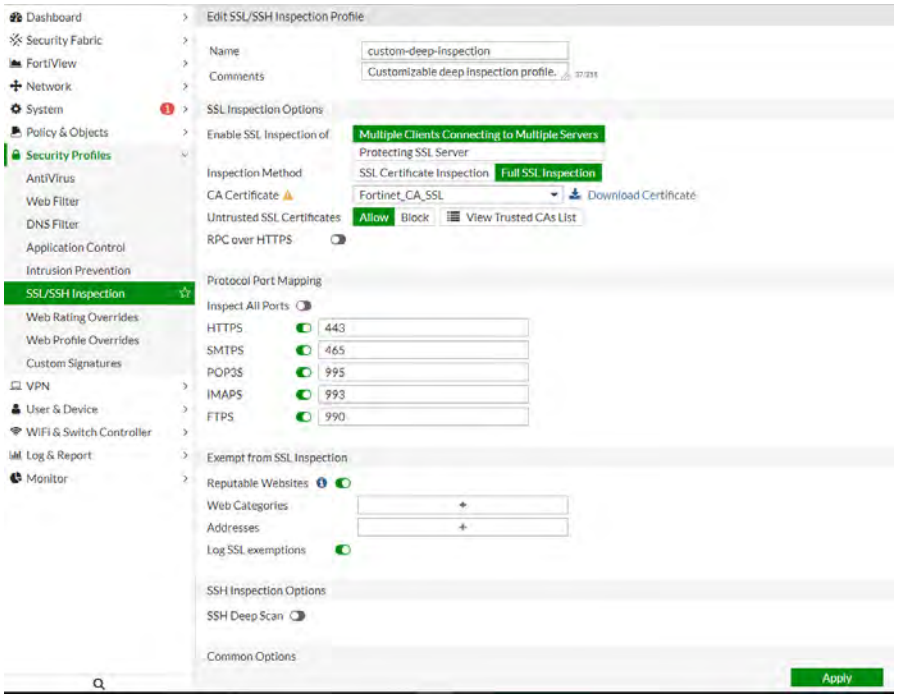
If you know the address of the server you want to exempt, you can exempt that address. You can exempt specific address type including IP address, IP address range, IP subnet, FQDN, wildcard-FQDN, and geography.



If you want to exempt all bank web sites, an easy way is to exempt the *Finance and Banking* category which includes all finance and bank web sites identified in FortiGuard.



If you want to exempt commonly trusted web sites, you can bypass the SSL allowlist in the SSL/SSH profile. The allowlist includes common web sites trusted by FortiGuard. Simply enable *Reputable Websites*.



Protecting an SSL server

You typically use the FortiGate *Protecting SSL Server* profile as an inbound policy for clients on the internet that access the server through the internal side of the FortiGate.