Administration Guide–Standalone Mode

FortiSwitchOS 7.4.4

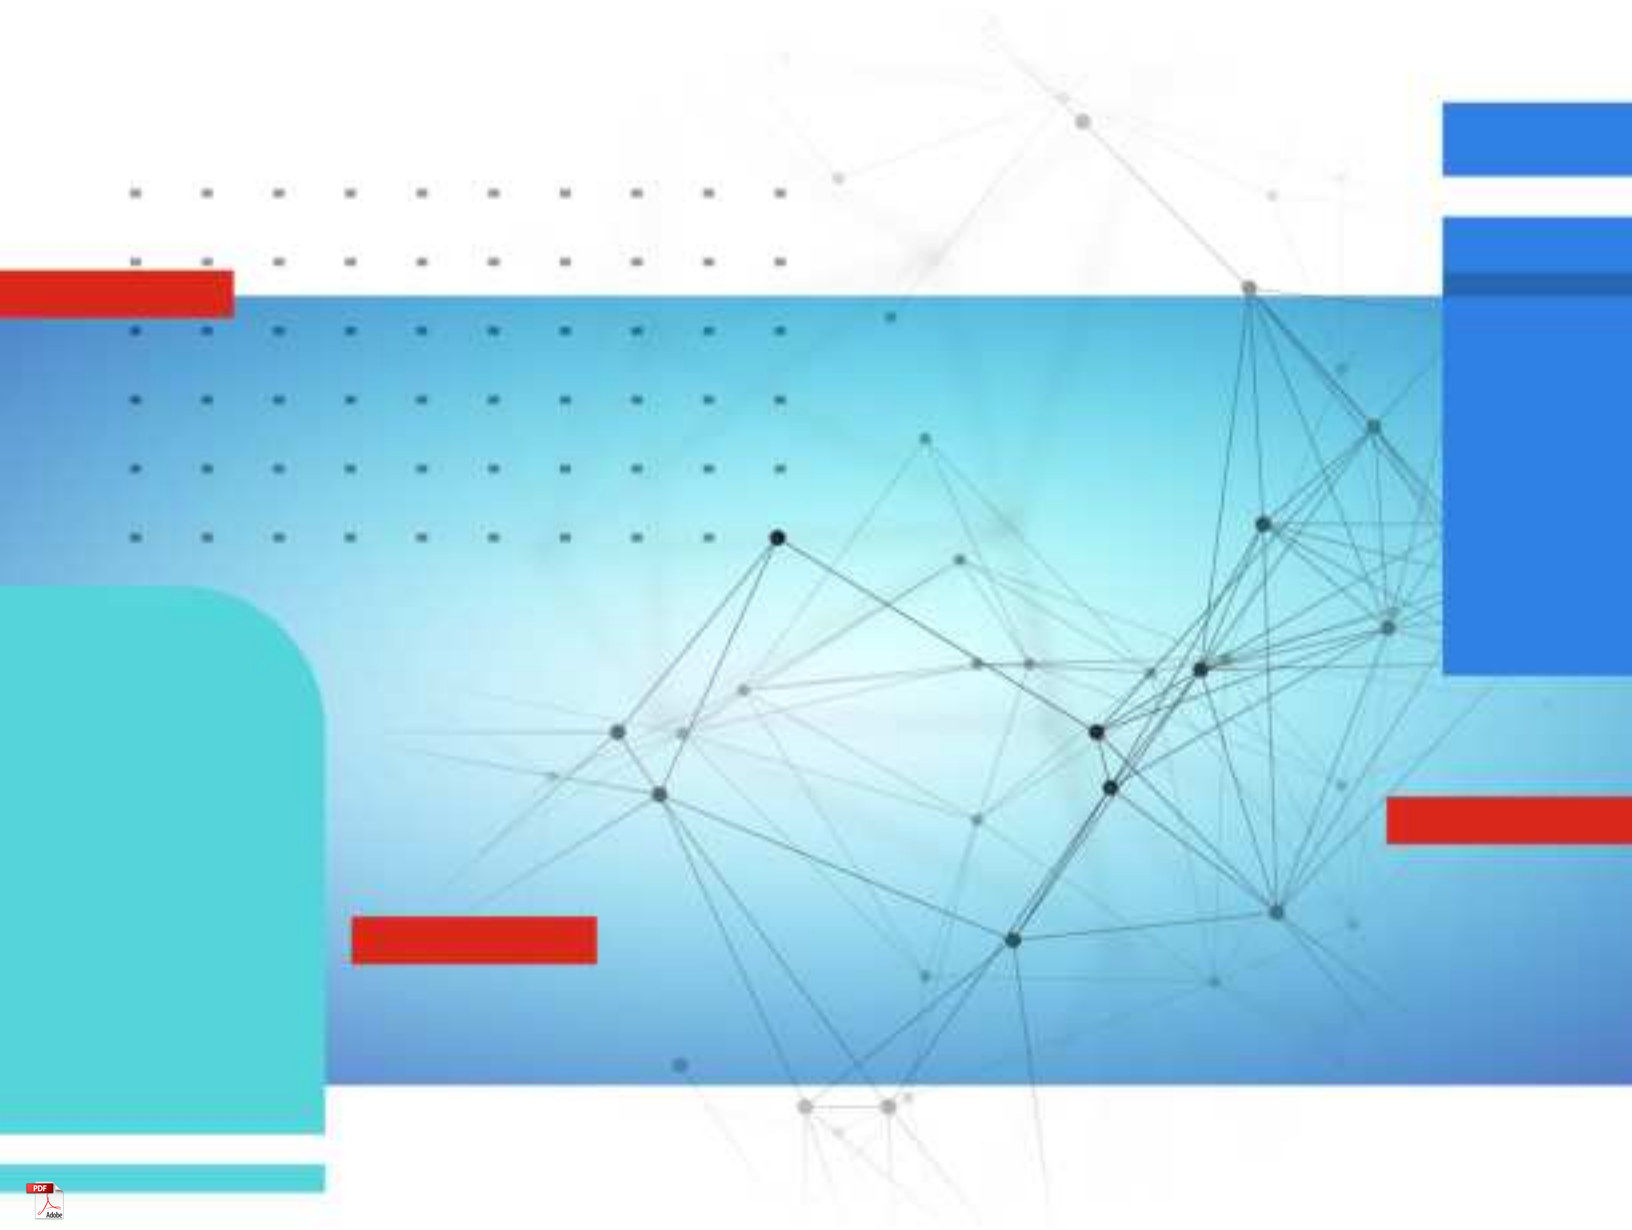

FORTINET DOCUMENT LIBRARY
https://docs.fortinet.com

FORTINET VIDEO LIBRARY
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

FORTINET TRAINING INSTITUTE
https://training.fortinet.com

FORTIGUARD LABS
https://www.fortiguard.com

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

FEEDBACK
Email: techdoc@fortinet.com

March 27, 2025
FortiSwitchOS 7.4.4 Administration Guide—Standalone Mode
11-744-1066320-20250327

TABLE OF CONTENTS

Change log **9**
What's new in FortiSwitchOS 7.4.4 ..... **10**
Introduction **11**
Supported models 11
Before you begin 11
System **12**
Dashboard 12
FortiLAN Cloud 13
CPU usage 14
RAM usage 14
Temperature 15
Bandwidth 15
Losses 16
Network 16
Management ports 16
Overlapping subnets 24
Switch virtual interfaces 24
VXLAN interfaces 26
Routed VLAN interfaces 43
VRRP 46
Loopback 48
IP conflict detection 49
ARP timeout value 50
Using SSH and the Telnet client 50
Config 50
SNMP 52
Firmware 57
Backup 60
Revisions 60
Licenses 61
Time 63
SSL 64
Configuring the temperature sensor 66
Admin 66
Administrators 66
Profiles 75
Access control 77
Monitor 80
Setting the idle timeout 80
Configuring system banners 80
Using Wake-on-LAN packets 82
Configuring automation stitches 82
Using the alias commands 88
User 95
User definition 95

Peer user 96
User groups 98
Authentication 99
RADIUS 99
TACACS 102
Certificate 105
Local 106
Remote 108
Authorities 109
CRLs 110
Flow export 111
Enabling packet sampling 115
Configuring flow export 116
Viewing the flow-export data 119
Deleting the flow-export data 120
DHCP 120
Configuring a DHCP server 120
Detailed operation of a DHCP relay 126
Configuring a DHCP relay 126
Packet capture 127
Creating a packet-capture profile 127
Starting the packet capture 128
Pausing or stopping the packet capture 129
Displaying or uploading the packet capture 129
Deleting the packet-capture file 130
Debug report 130
Fault relay support 130
Identifying a specific FortiSwitch unit 131
Using the Reset button on FortiSwitch units 131
Amber and red LEDs 131
Switch 133
Physical ports 133
Physical port settings 134
Switched interfaces 149
Dynamic MAC address learning 149
Layer-2 table 153
Loop guard 153
TFTP network port 154
Cable diagnostics 155
Link aggregation groups 156
MCLAG 159
Multi-stage load balance 165
Unicast hashing 167
LLDP-MED 167
Interfaces 181
Port security 181
MAC security 226

STP 240
MSTP overview and terminology 241
MSTP configuration 243
Interactions outside of the MSTP region 250
Viewing the MSTP configuration 251
Support for interoperation with Rapid per-VLAN RSTP (Rapid PVST+ or RPVST+) 251
Flap guard 252
Retaining the triggered state 253
Configuring the port flap guard 253
Resetting a port 255
Viewing the port flap guard configuration 255
DHCP snooping 255
Configuring DHCP snooping 256
Checking the DHCP-snooping configuration 262
Removing an entry from the DHCP-snooping binding database 263
IP source guard 264
Enabling IP source guard 264
Configuring IP source-guard static entries 265
Checking the IP source-guard entries 265
Checking the IP source-guard violation log 266
Dynamic ARP inspection 266
IPv6 router advertisement guard 268
ACL 271
ACL policy attributes 272
Configuring an ACL policy 273
Configuration examples 281
Selective packet sampling 283
Creating a schedule 284
IGMP snooping 285
Notes 286
Configuring IGMP snooping 287
Configuring the IGMP querier 292
Configuring mRouter ports 293
MLD snooping 293
Notes 294
Configuring MLD snooping 294
Configuring the MLD querier 298
PoE 299
sFlow 300
Configuring sFlow 300
Checking the sFlow configuration 302
Mirror 302
Configuring a SPAN mirror 304
Configuring an RSPAN mirror 307
Configuring an ERSPAN auto mirror 311
Configuring an ERSPAN manual mirror 312
VLAN 314
Native VLAN 314

Allowed VLAN list 314
Untagged VLAN list 315
Frame processing 315
Configuring VLANs 316
Example 1 316
Example 2 317
VLAN stacking (QnQ) 318
MAC/IP/protocol-based VLANs 324
Private VLANs 327
Virtual wires 329
Storm control 330
Configuring system-wide storm control 331
Configuring port-level storm control 331
Monitoring storm control 332
Checking the storm-control configuration 333
MAC entries 333
Persistent (sticky) MAC addresses 333
Static MAC addresses 334
IP-MAC binding 335
Configuring IP-MAC binding 335
Viewing IP-MAC binding configuration 336
QoS 336
Classification 337
Marking 338
Queuing 338
Determining the egress queue 339
Configuring FortiSwitch QoS 340
Checking the QoS statistics 346
Resetting and restoring QoS counters 346
Network monitor 346
Survey mode 347
Directed mode 348
Network-monitoring statistics 349
Configuring security checks 351
Syntax (for FS-108D-POE, FS-112D-POE, FS-224D-POE, and FSR-216F-POE) 352
Syntax (for FS-1xxE, FS-1xxF, and FS-110G-FPOE) 352
Syntax (for all other FortiSwitch models) 353
Cut-through switching mode 354
Enabling packet forwarding 354
Configuring auto-topology 355
Viewing port statistics 356
Media Redundancy Protocol 358
Configuring an MRP network 359
Viewing the MRP configuration 362
Precision Time Protocol 362
PTP node types 363
PTP message types 364
Packet flow in end-to-end mode 365

Packet flow in peer-to-peer mode 366
PTP profiles 366
PTP settings 368
PTP policies 368
Topology examples 369
Configuring PTP 369
Troubleshooting PTP 371
Configuration example 372
PTP operation details and limitations 373
High-Availability Seamless Redundancy 374
Configuring the HSR settings 375
Configuring the HSR ring 375
Clearing the HSR statistics 376
Troubleshooting HSR 376
Configuration examples 376
HSR operation details and limitations 380
Parallel Redundancy Protocol 381
Configuring the PRP settings 382
Configuring the PRP channel 382
Clearing the PRP statistics 383
Troubleshooting PRP 383
Configuration examples 384
PRP operation details and limitations 386
Router 388
Config 388
Layer-3 routing in hardware 388
Using layer-3 routing within an MCLAG 389
Unicast reverse-path forwarding (uRPF) 404
BGP routing 405
IS-IS routing 424
OSPF 429
RIP 437
Multicast 448
Access lists 451
Static and IPv6 static routing 453
Link probes 457
Virtual routing and forwarding 461
Policy-based routing 464
Key chains 468
Diagnostic 471
Multi-traceroute 473
ARP table 475
Monitor 476
Log 479
Syslog server 480
Deployment scenario 482
Working configuration for PC and phone for 802.1X authentication using MAC 482

Summary 482
A. Configure all devices 482
B. Authenticate phone using MAB 486
C. Authenticate the PC using EAP dot1x 488
Appendix A: FortiSwitch-supported RFCs 490
BFD 490
BGP 490
DHCP 491
IP/IPv4 491
IP multicast 491
IPv6 491
IS-IS 492
MIB 492
OSPF 493
Other protocols 493
RADIUS 493
RIP 494
SNMP 494
Syslog 494
VXLAN 494
Appendix B: Supported attributes for RADIUS CoA and RSSO 495
Appendix C: SNMP OIDs for FortiSwitch models 505

Change log

Date	Change Description
October 4, 2024	Initial release for FortiSwitchOS 7.4.4
January 9, 2025	Reworded the following sentence in Notes on page 160: “There is a maximum of two FortiSwitch models per MCLAG.”
March 12, 2025	Added more information to Fault relay support on page 130.
March 27, 2025	Added another item to Detailed deployment notes on page 225.

What's new in FortiSwitchOS 7.4.4

Release 7.4.4 provides the following new features:

- The FS-624F, FS-624F-FPOE, FS-648F, and FS-648F-FPOE models now support LLDP-MED Power Management TLVs. For more details, see LLDP-MED on page 167.
- The FS-1024E, FS-1048E, FS-T1024E, FS-T1024F-FPOE, FS-2048F, FS-3032E, and FS-6xxF models now support 2,550 switch virtual interfaces (SVIs). For more information about SVIs, see Switch virtual interfaces on page 24.
- Media Access Control security (MACsec) now supports AWS Direct Connect in the CLI with the FS-1024E model. For more information about MACsec, see MAC security on page 226.
- The FortiSwitch unit now supports two Media Redundancy Protocol (MRP) rings. For more details, see Media Redundancy Protocol on page 358.

Introduction

This guide provides information about configuring a FortiSwitch unit in standalone mode. In standalone mode, you manage the FortiSwitch unit by connecting directly to the unit, either using the web-based manager (also known as the GUI) or the CLI.

If you will be managing your FortiSwitch unit using a FortiGate unit, refer to the *FortiLink Guide (FortiOS 7.4.3)*.

If you will be managing your FortiSwitch unit using FortiLAN Cloud, see the *FortiLAN Cloud User Guide*.

If you will be managing your FortiSwitch unit using FortiSwitch Manager, see the *FortiSwitch Manager Administration Guide*.

This section covers the following topics:

- Supported models on page 11
- Before you begin on page 11

Supported models

This guide is for all FortiSwitch models that are supported by FortiSwitchOS, which includes all of the D-series, E-series, and F-series models.

Refer to the FortiSwitch feature matrix for details about the features supported by each FortiSwitch model.

Before you begin

Before you start administrating your FortiSwitch unit, it is assumed that you have completed the initial configuration of the FortiSwitch unit, as outlined in the QuickStart Guide for your FortiSwitch model and have administrative access to the FortiSwitch unit's GUI and CLI.

System

This section contains information about FortiSwitch administration and system configuration that you can do after installing the FortiSwitch unit in your network.

- Dashboard on page 12
- Network on page 16
- Config on page 50
- Admin on page 66
- User on page 95
- Authentication on page 99
- Certificate on page 105
- Flow export on page 111
- DHCP on page 120
- Packet capture on page 127
- Debug report on page 130
- Fault relay support on page 130
- Identifying a specific FortiSwitch unit on page 131
- Using the Reset button on FortiSwitch units on page 131
- Amber and red LEDs on page 131

Dashboard

Go to *System* > *Dashboard* to see your FortiSwitch operational status and data for the last day and last week of the switch's CPU usage, RAM usage, temperature, bandwidth, and losses. The operation status reports on the following:

- Port information
  Hover your cursor over the port to see the link status, port speed, maximum transmission unit (MTU), number of packets sent, and number of packets received.
- Status of power supply units (PSUs)
  NOTE: The PSU status is shown only on FortiSwitch models with redundant PSUs.
- FortiLink status
  The FortiLink widget shows whether the FortiSwitch unit is managed by a FortiGate device.
- FortiLAN Cloud status
  Click *FortiLAN Cloud* to go to the *System* > *FortiLAN Cloud* page
- How long the switch has been running (uptime)
- Time of last backups
  Click *Last Backup* to go to the *System* > *Config* > *Revisions* page.
- Current CPU usage
- Current RAM usage
- Current temperature for FortiSwitch models that have temperature sensors
- Current power over Ethernet (PoE) usage (on FortiSwitch PoE models)

FortiLAN Cloud

Click *FortiLAN Cloud* to go to the *System* > *FortiLAN Cloud* page.

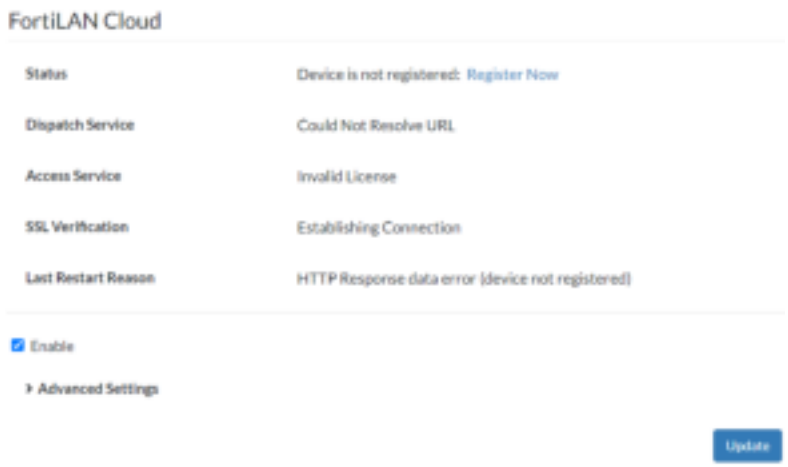

Select the *Enable* checkbox and then expand *Advanced Settings* to configure your FortiSwitch unit to be managed by FortiLAN Cloud.

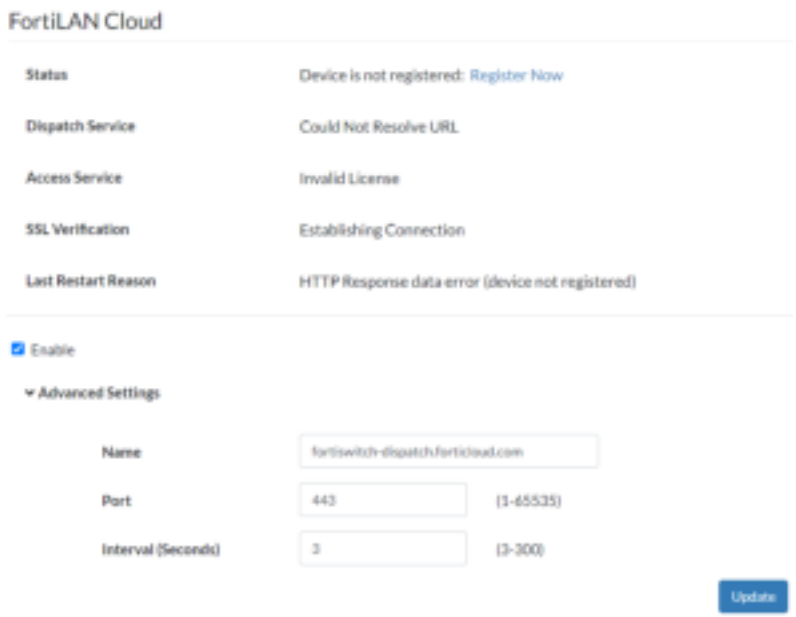

To switch to FortiLAN Cloud management:

1. On the *FortiLAN Cloud* page, select the *Enable* checkbox and then expand *Advanced Settings*.
2. Enter the domain name for FortiLAN Cloud.
3. Enter the port number used to connect to FortiLAN Cloud.
4. Enter the time in seconds allowed for domain name system (DNS) resolution.
5. Click *Update* to save your changes.

CPU usage

The *CPU Usage* graphs show how much of the CPU is used by the FortiSwitch unit over a day and over a week. Hover your cursor over the graph to display the data for a specific day and time.

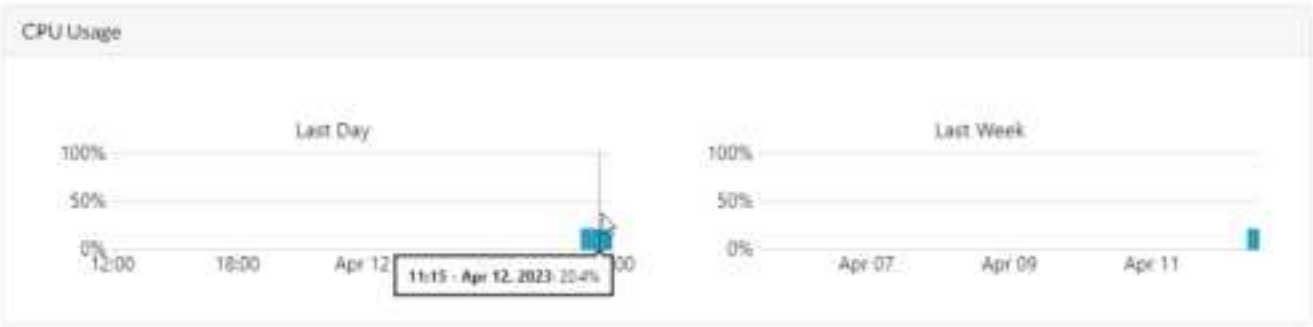

RAM usage

The *RAM Usage* graphs show how much of the RAM is used by the FortiSwitch unit over a day and over a week. Hover your cursor over the graph to display the data for a specific day and time.

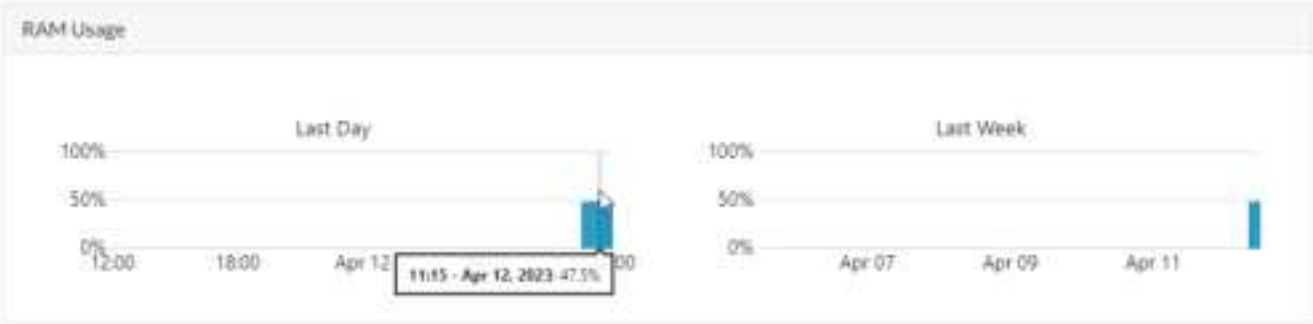

Temperature

The *Temperature* graphs show the printed circuit board (PCB) temperature over a day and over a week. Hover your cursor over the graph to display the data for a specific day and time.

If the FortiSwitch model has multiple temperature sensors, the temperature displayed is an average of the readings from all sensors.

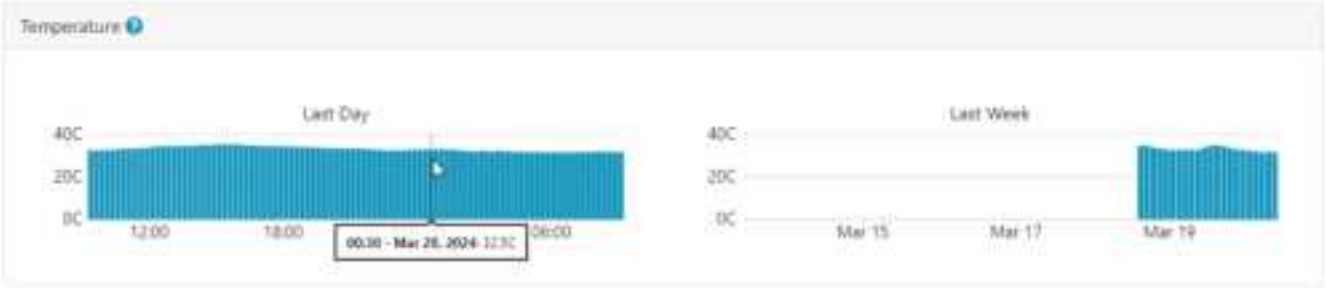

Bandwidth

The *Bandwidth* graphs show the inbound and outbound bandwidth for the entire FortiSwitch unit over a day and over a week. Hover your cursor over the graph to display the data for a specific day and time.

Losses

The *Losses* graphs show the inbound errors, outbound errors, inbound drops, and outbound drops for the entire FortiSwitch unit over a day and over a week. Hover your cursor over the graph to display the data for a specific day and time.

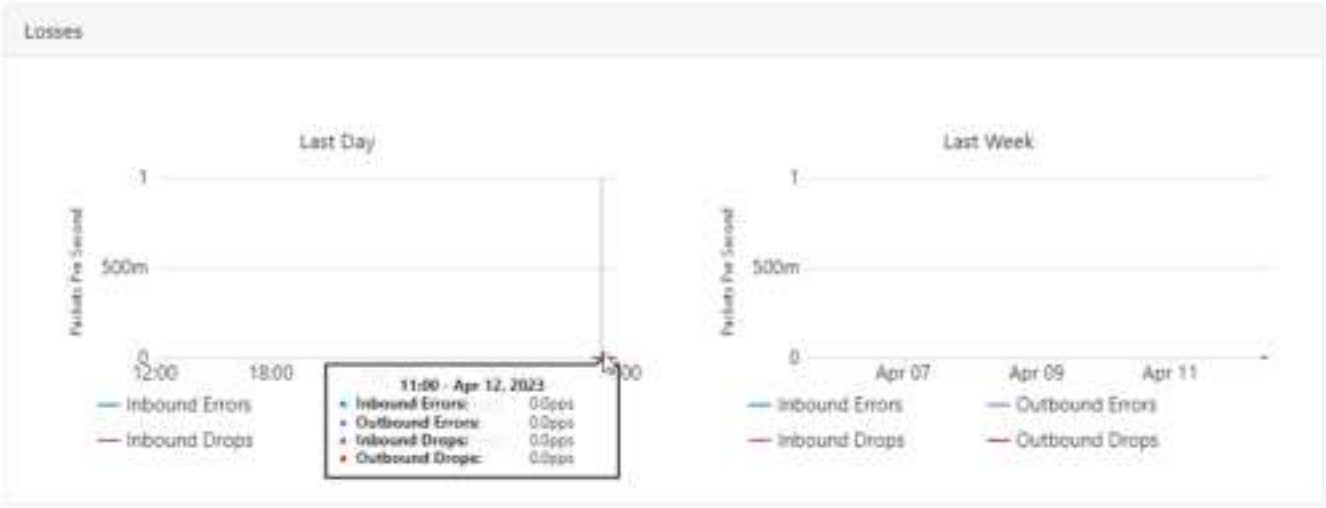

Network

The following topics provide information about network settings:

- Management ports on page 16
- Overlapping subnets on page 24
- Switch virtual interfaces on page 24
- VXLAN interfaces on page 26
- Routed VLAN interfaces on page 43
- VRRP on page 46
- Loopback on page 48
- IP conflict detection on page 49
- ARP timeout value on page 50
- Using SSH and the Telnet client on page 50

Management ports

This section describes how to configure management ports on the FortiSwitch unit:

- Models without a dedicated management port on page 17
- Models with a dedicated management port on page 20
- Example configurations on page 22

You can use HTTP, HTTPS, Telnet, and SSH to manage FortiSwitch units.

NOTE: SSHv2 is supported.

Models without a dedicated management port

For FortiSwitch models without a dedicated management port, configure the internal interface as the management port.

NOTE: For FortiSwitch models without a dedicated management port, the internal interface has a default VLAN ID of 1.

Using the GUI:

First start by editing the default *internal* interface's configuration.

1. Go to *System* > *Network* > *Interface* > *Physical*, select *Edit* for the *internal* interface.

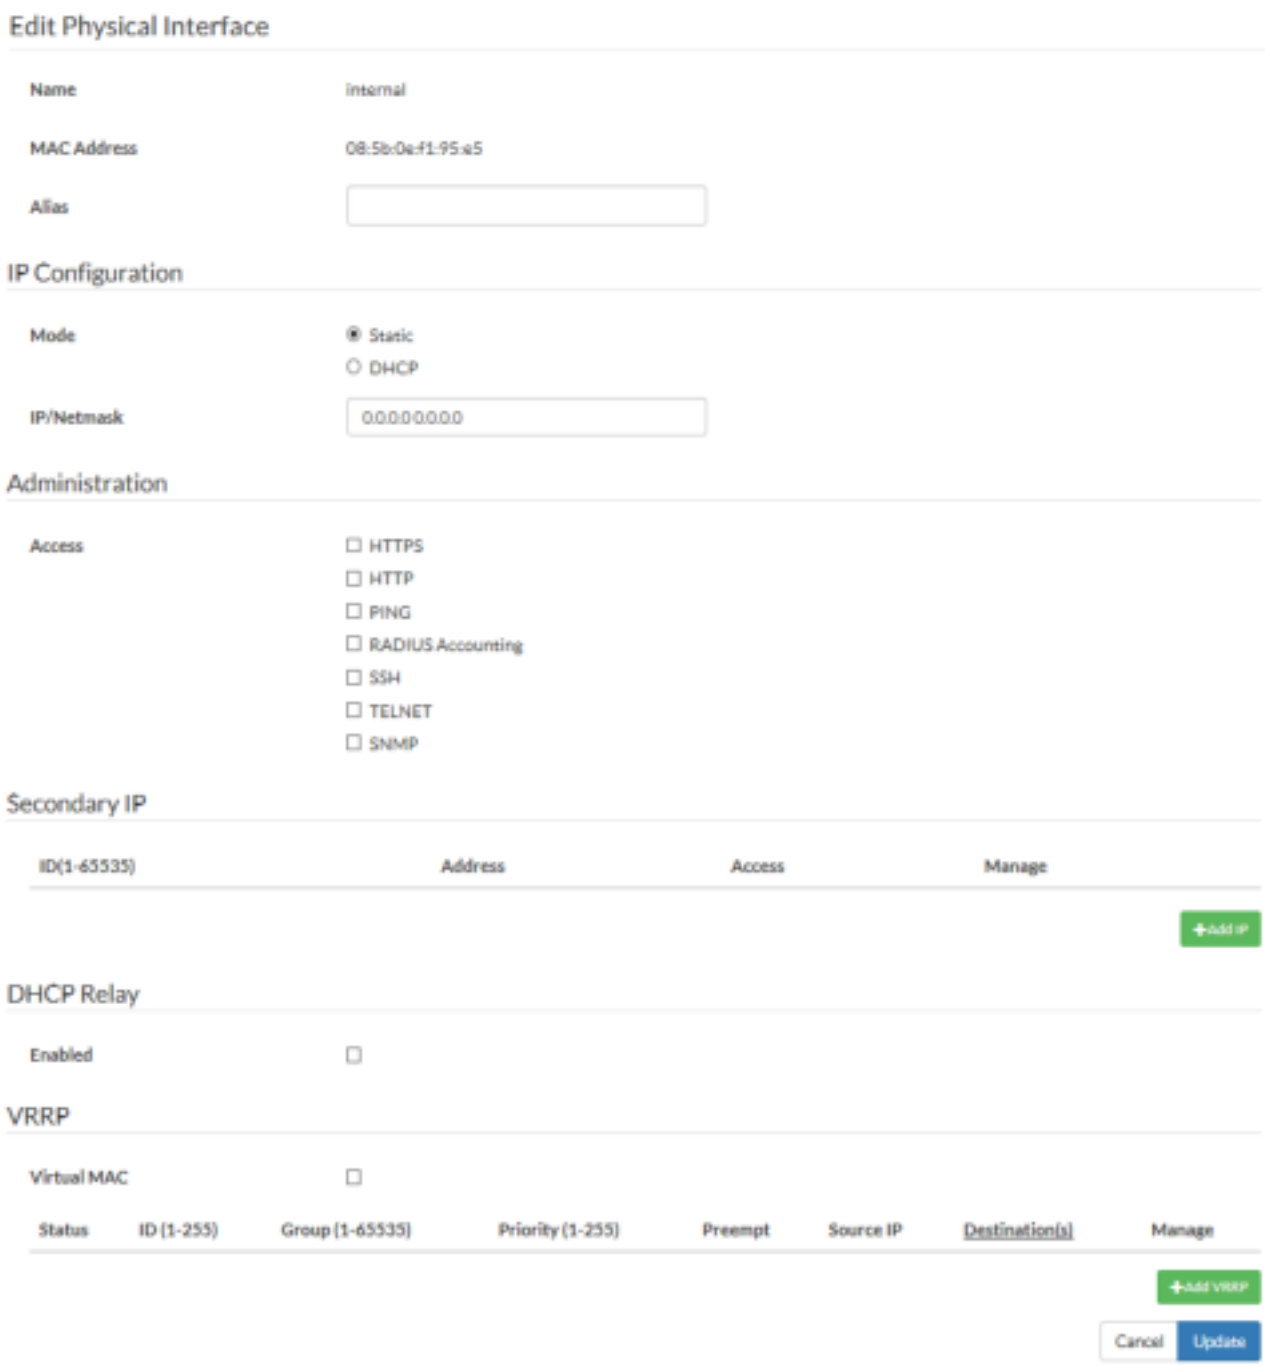

2. In the IP/Netmask field, enter the IP address and netmask.
3. Select the appropriate protocols to connect to the interface for administrative access.
4. Optional. Select *Add IP* to add a secondary IP address for the internal interface.
5. Select *Update* to save your changes.

Next, create a new interface to be used for management.

1. Go to *System* > *Network* > *Interface* > *VLAN* and select *Add VLAN* to create a management VLAN.

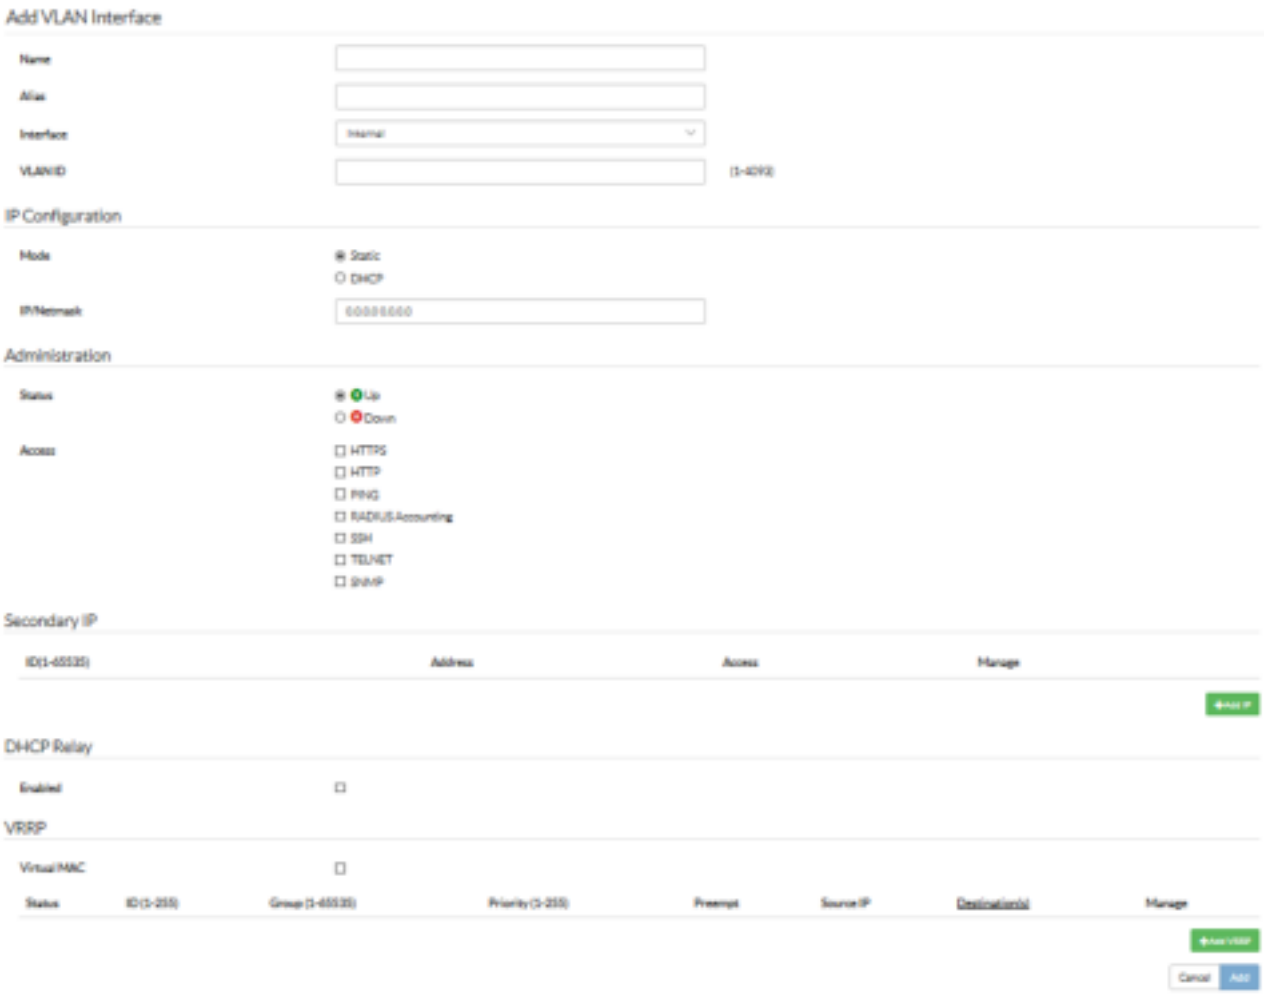

2. Give the interface an appropriate name.
3. Confirm that *Interface* is set to *internal*.
4. Set a *VLAN ID*.
5. In the IP/Netmask field, enter the IP address and netmask.
6. Select the appropriate protocols to connect to the interface for administrative access.
7. Optional. Select *Add IP* to add a secondary IP address for this VLAN.
8. Select *Add*.

Using the CLI:

```
config system interface
   edit internal
      set ip <IP_address_and_netmask>
      set allowaccess <access_types>
      set type physical
      set secondary-IP enable
         config secondaryip
            edit <id>
               set ip <IP_address_and_netmask>
               set allowaccess <access_types>
            next
```

```
      end
   next
edit <vlan name>
   set ip <IP_address_and_netmask>
   set allowaccess <access_types>
   set interface internal
   set vlanid <VLAN id>
   set secondary-IP enable
      config secondaryip
         edit <id>
            set ip <IP_address_and_netmask>
            set allowaccess <access_types>
      end
   end
```

Models with a dedicated management port

For FortiSwitch models with a dedicated management port, configure the IP address and allowed access types for the management port.

NOTE: For FortiSwitch models with a dedicated management port, the internal interface has a default VLAN identifier of 4094.

Using the GUI:

1. Go to *System* > *Network* > *Interface* > *Physical*, select *Edit* for the *mgmt* interface.

2. In the ID field, enter a unique identifier from 1 to 65525.
3. In the *IP/Netmask* field, enter the IP address and netmask.
4. Select the appropriate protocols to connect to the interface for administrative access.
5. Optional. You can select *Remove* if you want to delete the default secondary IP address or select *Add IP* to add a secondary IP address for the management interface.
6. Select *Update* to save your changes.

Using the CLI:

```
config system interface
   edit mgmt
      set ip <IP_address_and_netmask>
      set allowaccess <access_types>
      set type physical
      set secondary-IP enable
         config secondaryip
            edit <id>
               set ip <IP_address_and_netmask>
               set allowaccess <access_types>
            next
         end
```

```
      next
   edit internal
      set type physical
   end
end
```

Example configurations

In this example, the *internal* interface is used as an inbound management interface. Also, the FortiSwitch unit has a default VLAN across all physical ports and its internal port.

Using the internal interface of a FortiSwitch-524D-FPOE

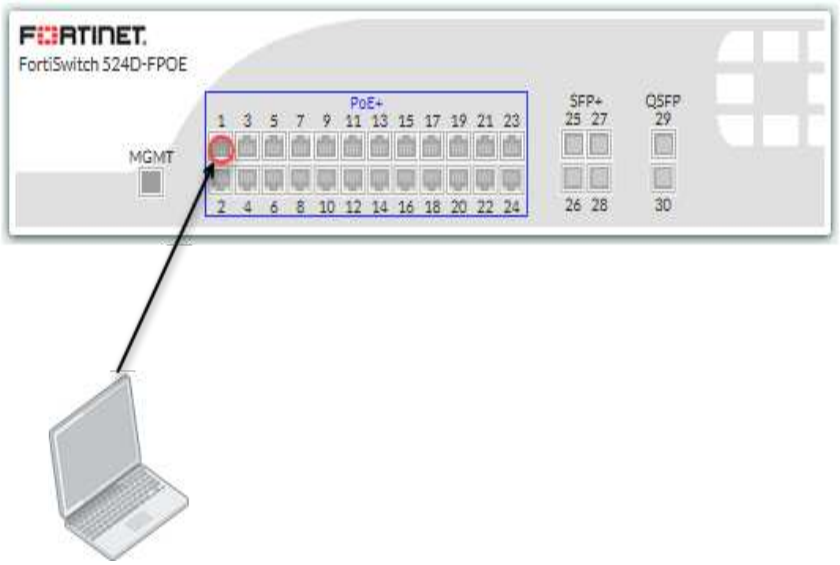

Syntax

```
config system interface
   edit internal
      set ip 192.168.1.99 255.255.255.0
      set allowaccess ping https http ssh
      set type physical
   end
end
```

In this example, an out-of-band management interface is used as the dedicated management port. You can configure the management port for local or remote access.

Out-of-band management on a FortiSwitch-1024D

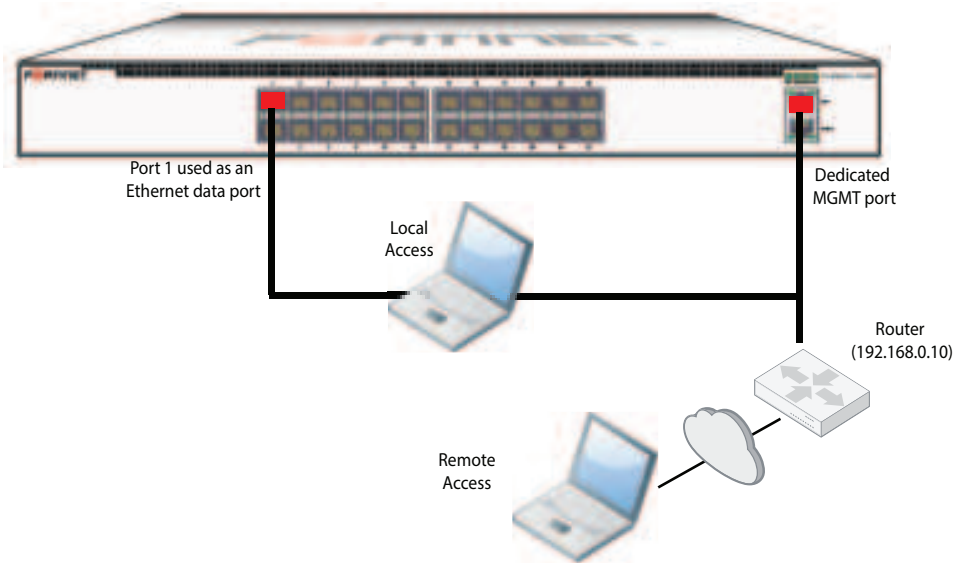

Option 1: management port with static IP

```
config system interface
   edit mgmt
      set mode static
      set ip 10.105.142.19 255.255.255.0
      set allowaccess ping https http ssh snmp telnet
      set type physical
   next
   edit internal
      set type physical
   end
end
// optional configuration to allow remote access to the management port

config router static
   edit 1
      set device mgmt
      set gateway 192.168.0.10
      set status enable
   end
```

Option 2: management port with IP assigned by DHCP

```
config system interface
   edit mgmt
      set mode dhcp
      set defaultgw enable // allows remote access
```

```
      set allowaccess ping https http ssh snmp telnet
      set type physical
   next
   edit internal
      set type physical
   end
```

Overlapping subnets

You can use the `set allow-subnet-inteface` command to allow two interfaces to include the same IP address in the same subnet. The command applies only between the mgmt interface and an internal interface.

NOTE: Different interfaces cannot have overlapping IP addresses or subnets. The same IP address can be used on different switches.

For example:

```
config system global
   set admintimeout 480
   set allow-subnet-overlap enable
end
config system interface
   edit "mgmt"
      set ip 172.16.86.112 255.255.255.0
      set allowaccess ping https http ssh snmp telnet
      set type physical
      set alias "test"
      set snmp-index 27
   next
   edit "internal"
      set ip 10.0.1.112 255.255.255.0
      set allowaccess ping
      set type physical
      set alias "testing-2"
      set snmp-index 26
   next
end
```

Switch virtual interfaces

A switch virtual interface (SVI) is a logical interface that is associated with a VLAN and supports routing and switching protocols.

You can assign an IP address to the SVI to enable routing between VLANs. For example, SVIs can route between two different VLANs connected to a switch (no need to connect through a layer-3 router).

Configuring a switch virtual interface

Using the GUI:

1. Go to *System* > *Network* > *Interface* > *VLAN*.
2. Select *Add VLAN*.
3. Enter a name for the interface.

4. Select *internal* from the *Interface* drop-down list.
5. Enter a VLAN identifier in the *VLAN ID* field.
6. Select *Static* for the mode and enter an IP address and netmask in the *IP/Netmask* field.
7. Select the administration status.
8. Select *PING*, *SSH*, and T*ELNET* for the *Access* options.
9. Select *Add*.

Using the CLI:

Create a system interface. Give it an IP subnet and an associated VLAN:

```
config system interface
   edit <system interface name>
      set ip <IP address and mask>
      set vlanid <vlan>
      set allowaccess ping ssh telnet
```

Example SVI configuration

The following is an example CLI configuration for SVI static routing.

In this configuration, Server-1 is connected to switch Port1, and Server-2 is connected to switch Port2. Port1 is a member of VLAN 4000, and Port2 is a member of VLAN 2. Port1 is the gateway for Server-1, and port2 is the gateway for Server-2.

NOTE: For simplicity, assume that both port1 and port are on same switch.

1. Configure the native VLANs for Port 1 and Port 2:

```
config switch interface
   edit port1
      set native-vlan 4000
   edit port2
      set native-vlan 2
   end
```

2. Create L3 system interfaces that correspond to Port 1 (VLAN 4000) and Port 2 (VLAN 2):

```
config system interface
   edit vlan4000
      set ip 192.168.11.1/24
      set vlanid 4000
      set allowaccess ping ssh telnet
   next
      edit vlan2
         set ip 192.168.10.1/24
         set vlanid 2
         set allowaccess ping ssh telnet
   end
```

Viewing the SVI configuration

To display the status of SVI configuration:

```
show system interface [ <system interface name> ]
```

VXLAN interfaces

This section covers the following topics:

- Configuring VXLAN interfaces on page 27
- Creating an STP virtual root on page 29
- Using BGP EVPN with VXLAN on page 32
- Using ECMP routing with VXLAN interfaces on page 36
- Using DHCP snooping with VXLAN on page 37
- Using QoS with VXLAN tunnels on page 42

You can use Virtual Extensible LAN (VXLAN) interfaces to send layer-2 traffic between FortiSwitch units over a layer-3 tunnel. VXLAN tunnels connect virtual tunnel endpoints (VTEPs) using VXLAN network identifiers (VNIs).

A FortiSwitch unit (VTEP) encapsulates traffic from a VNI and then sends it across the physical IP network using the VXLAN tunnel to another FortiSwitch unit (VTEP)

In the following configuration example, three VNIs connect four FortiSwitch units (VTEPs).

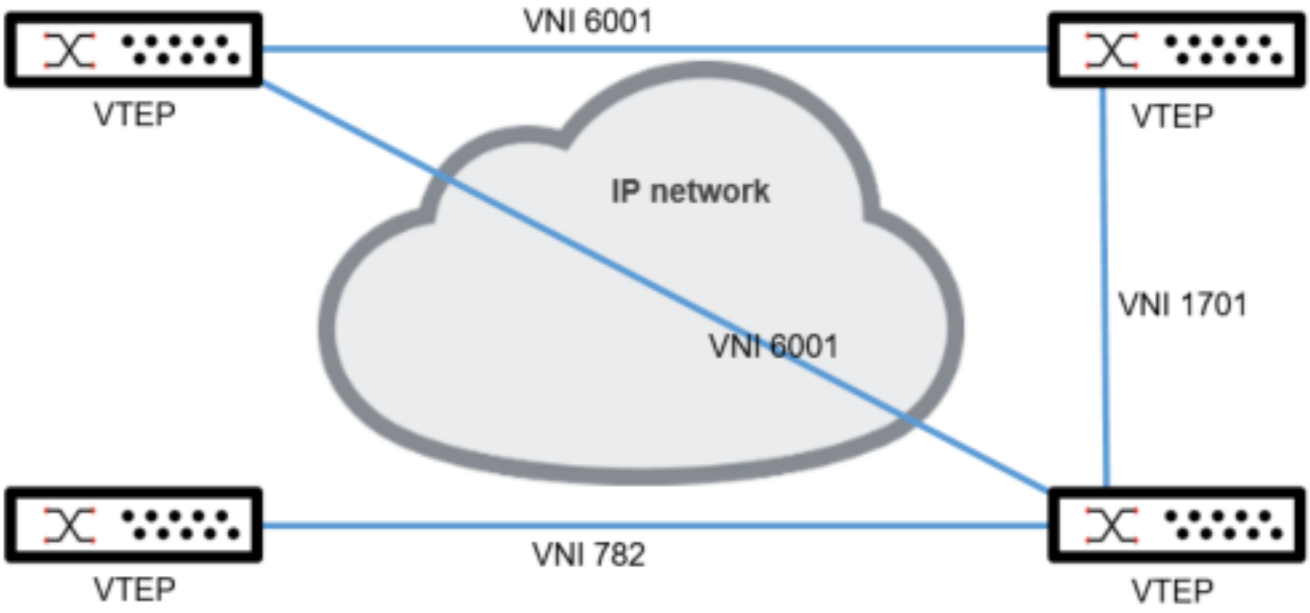

The FortiSwitch units learn remote MAC addresses by flooding broadcast, unicast, and multicast packets to each `remote-ip` address to find out the MAC address associated with the tunnel source.

The following requirements apply to VXLAN tunnels:

- When you configure the VXLAN interface, the system interface defines the VXLAN tunnel destination, and the VXLAN tunnel destination must match the `remote-ip` setting of the VXLAN tunnel initiator.
- The IP address used for the VXLAN tunnel must be a static IP address and must be the primary IP address on the interface. If the primary IP address is static but the IP address has not been configured, no VXLAN tunnel is created.
- The `mode` for `config system interface` cannot be set to `dhcp`; otherwise, the results are unreliable.
- If you are using VXLAN with FortiLink, refer to Managing FortiSwitch units on VXLAN interfaces.
- The routing functionalities of the SVI created for the VXLAN's access VLANs are not supported.

Configuring VXLAN interfaces

To create a VXLAN tunnel:

1. Set the UDP port for the VXLAN tunnel destination.
   The range of values is 1-65535. The default port is 4789.
2. Configure the VXLAN interface.
3. Check the VXLAN configuration.

To set the VXLAN tunnel destination:

```
config switch global
   set vxlan-dport <1-65535>
end
```

For example:

```
config switch global
   set vxlan-dport 100
end
```

To configure the VXLAN interface:

```
config system vxlan
   edit <VXLAN_interface_name>
      set vni <1-16777215>
      set vlanid <1-4094>
      set interface <interface_name>
      set ip-version {ipv4-multicast | ipv4-unicast}
      set remote-ip <IPv4_address>
   next
end
```

Variable	Description	Default
<VXLAN_interface_name>	Enter a name for the VXLAN interface	No default
vni <integer>	Required. Set the VXLAN network identifier (VNI). The range of values is 1-16777215.	0
vlanid <integer>	Required. Set the VLAN identifier that is mapped to the VNI. When `tunnel-loopback` is set, VLAN 4087 is reserved.	0

Variable	Description	Default
interface <interface_name>	Required. Enter the name of the outgoing interface for the VXLAN tunnel. Starting in FortiSwitchOS 7.2.1, you can specify a routed VLAN interface (RVI).	No default
ip-version {ipv4-multicast | ipv4-unicast}	Required. Select the type of IPv4 address to use to communicate over the VXLAN tunnel. • `ipv4-multicast`—Use IPv4 multicast addressing over the VXLAN tunnel. • `ipv4-unicast`—Use IPv4 unicast addressing over the VXLAN tunnel.	ipv4-unicast
remote-ip <IPv4_address>	Required. Enter the source and destination IPv4 addresses of the VXLAN interface. The VXLAN tunnel destination must match the `remote-ip` setting of the VXLAN tunnel initiator. Starting in FortiSwitchOS 7.2.1, you can specify an RVI as the source or destination IPv4 address.	No default

For example, if you want to create the following two VXLAN tunnels:

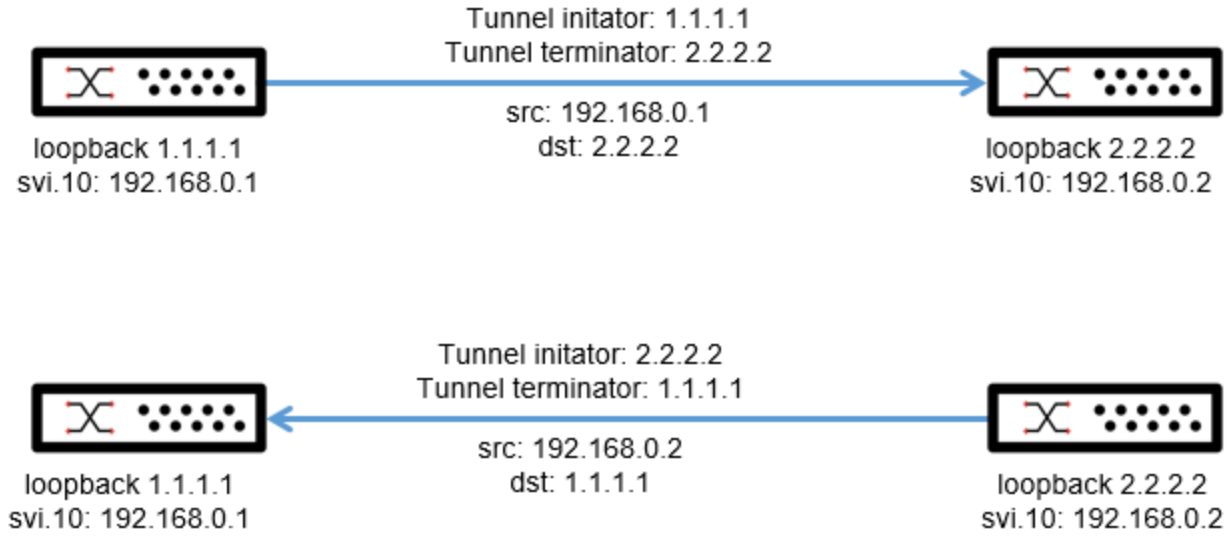

To configure loopback 1.1.1.1:

```
config system vxlan
   edit "vni.4094"
      set vni 4094
      set vlanid 4094
      set ip-version ipv4-unicast
      set remote-ip "2.2.2.2"
      set interface "loopback"
   next
end

config system interface
   edit "svi.10"
      set ip 192.168.0.1
   next
   edit "loopback"
      set ip 1.1.1.1/32
   next
end
```

To configure loopback 2.2.2.2:

```
config system vxlan
   edit "vni.4094"
      set vni 4094
      set vlanid 4094
      set ip-version ipv4-unicast
      set remote-ip "1.1.1.1"
      set interface "loopback"
   next
end

config system interface
   edit "svi.10"
      set ip 192.168.0.2
   next
   edit "loopback"
      set ip 2.2.2.2/32
   next
end
```

To check the VXLAN configuration:

- `diagnose switch vxlan mac-address list <VXLAN_interface_name>`
- `get system interface vxlan`

Creating an STP virtual root

Starting in FortiSwitchOS 7.2.1, you can prevent layer-2 loops between VTEPs. When the STP virtual root feature is enabled on all VTEPs in a VXLAN tunnel, the FortiSwitch units act as a single STP root so that no loops can form between any of the switches.

For example, in the following topology, the user has accidentally configured a loop between switch E1 and switch E3:

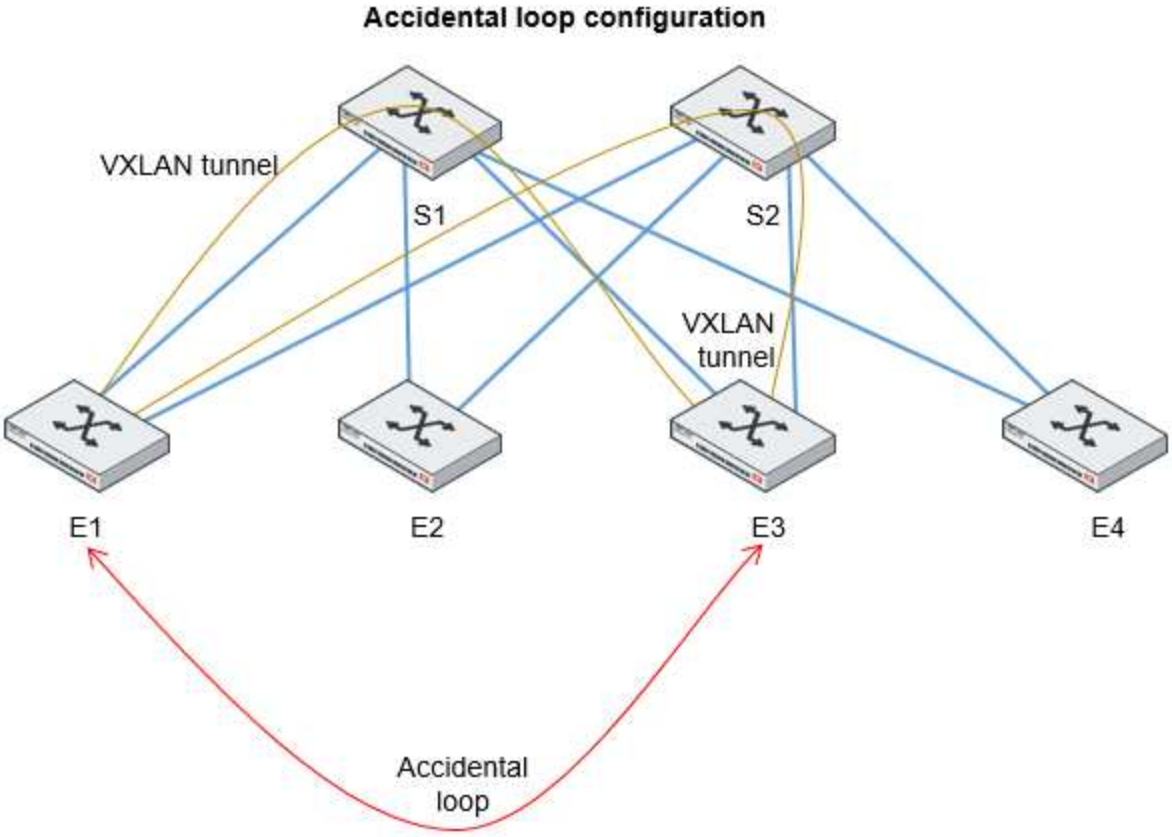

Using the STP virtual root feature, the loop between switch E1 and switch E3 is prevented:

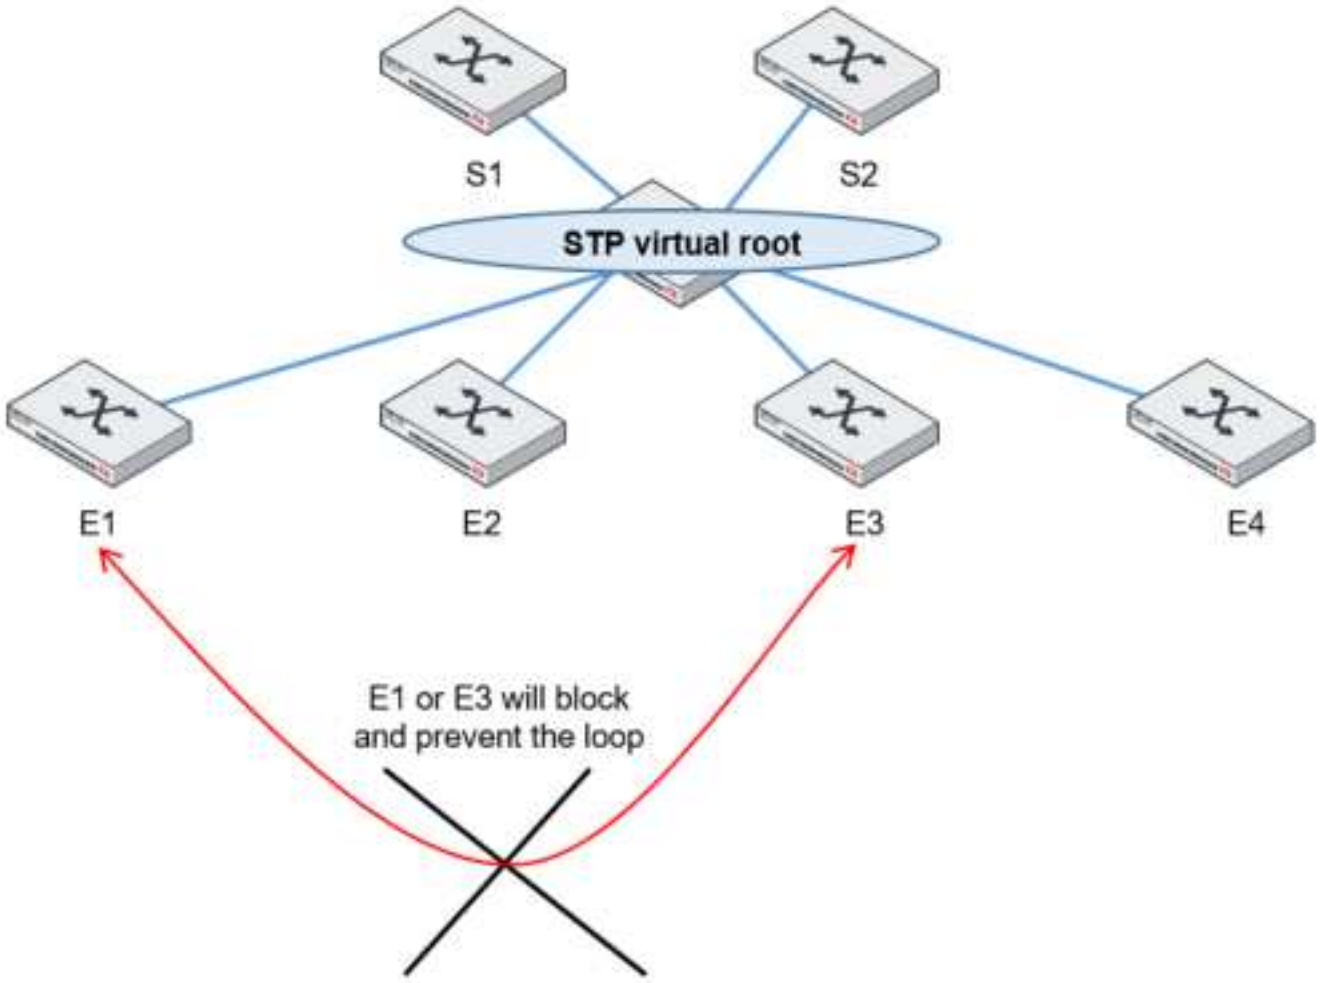

For the STP virtual root feature to work correctly, the core of the network must be a routed layer-3 network that is not participating in the Spanning Tree Protocol. Commonly, the network is using routed interfaces instead that terminate the layer-2 network.

By default, the STP virtual root feature is disabled. After you enable this feature, the MAC address for the virtual STP root is set to 08:5B:0E:00:00:00 by default, and the STP instance priority is set to 0. If you want to use a different MAC address for the virtual STP root, you can configure any unicast MAC address, but the same MAC address must be configured on all VTEPs in the VXLAN tunnel. If there are different MAC addresses configured on the VTEPs, there will be an “ERROR: virtual-root enable, not root!” listed on the *Switch* > *STP* > *Instances* page.

The VTEPs must meet one of the following requirements to become an STP virtual root:

- Run IEEE 802.1s multiple Spanning Tree Protocol (MSTP) and belong to a Common and Internal Spanning Tree (CIST).
- Run IEEE 802.1s MSTP and are in the same MSTP region.
- Run IEEE 802.1D Spanning Tree Protocol (STP).

- Run IEEE 802.1w Rapid Spanning Tree Protocol (RSTP).
- Support interoperation with per-VLAN Rapid Spanning Tree (RPVST) with their roots within FortiSwitch units.

If you are using an SVI that is associated with one or more VLANs on the network side, Fortinet recommends locating the network-side VLAN and the access-side VLAN on different STP instances.

To create the STP virtual root, configure the following commands on all VTEPs in a VXLAN tunnel:

```
config switch global
   set vxlan-stp-virtual-root enable
   set vxlan-stp-virtual-mac <MAC_address>
end
```

Using BGP EVPN with VXLAN

NOTE: You must have an advanced features license to use BGP routing.

Starting with FortiSwitchOS 7.4.0, you can use the Border Gateway Protocol (BGP) Ethernet Virtual Private Network (EVPN) with VXLAN.

BGP is a gateway protocol that enables the Internet to exchange routing information between autonomous systems. An EVPN extends BGP to advertise layer-2 MAC addresses and layer-3 IP addresses. EVPN uses overlay networks to add scalability and performance in the data center. This ensures workloads can be dynamically placed anywhere, removing traditional layer-3 boundaries of the physical infrastructure.

RFC 7432 describes BGP EVPN. For more information about BGP routing, see BGP routing on page 405.

To use a BGP EVPN with VXLAN:

1. Enable the `set activate-evpn` command (under `config neighbor`, which is under `config router bgp`) to allow FortiSwitchOS to exchange layer-2 VNI information with its neighbor.
2. Configure EVPN. There are two mandatory settings for EVPN; all other settings are optional.
   - Enable the `set activate-evpn` command (under `config neighbor`, which is under `config router bgp`).
   - Enable the `set advertise-vni` command (under `config evpn`, which is under `config router bgp`).
3. Enable EVPN for the VXLAN using the `set evpn` command (under `config system vxlan`).

To enable the exchange of layer-2 VNI information and to configure EVPN:

```
config router bgp
   set as <MANDATORY_router_AS_number>
   set router-id <MANDATORY_IP_address>
   config neighbor
      edit "<IPv4_address>"
         set activate-evpn enable
         set attribute-unchanged-evpn as-path
         set remote-as <MANDATORY_1-4294967295>
         set route-reflector-client-evpn {disable | enable}
         set route-map-in-evpn <string>
         set route-map-out-evpn <string>
         set soft-reconfiguration-evpn {disable | enable}
         set update-source <interface_name>
```

```
      next
   end
   config evpn
      set advertise-vni enable
      config vni
         edit <VNI_ID>
            set export-rt <ASN:VNI_number or A.B.C.D.VNI_number>
            set import-rt <ASN:VNI_number or A.B.C.D.VNI_number>
            set rd <ASN:VNI_number or A.B.C.D.VNI_number>
         next
      end
   end
end
```

For example:

```
config router bgp
   set as 60000
   set router-id 1.1.1.1
   config neighbor
      edit "2.2.2.2"
         set activate-evpn enable
         set remote-as 60000
         set update-source "loop1"
      next
   end
   config evpn
      set advertise-vni enable
   end
```

To enable EVPN for the VXLAN:

```
config system vxlan
   edit <VXLAN_interface_name>
      set vni <integer>
      set vlanid <integer>
      set interface <interface_name>
      set evpn enable
   next
end
```

For example:

```
config system vxlan
   edit "vx1"
      set vni 500
      set vlanid 50
      set interface "vlan12"
      set evpn enable
   next
end
```

To check the BGP EVPN configuration:

- `get router info evpn vni detail`
- `get router info evpn vni <VNI_number>`
- `get router info bgp evpn statistics`

- `get router info bgp evpn summary`
- `get router info bgp evpn route {detail | type | vni}`
- `get router info bgp evpn vni <VNI_number>`

ARP and ND suppression

Address Resolution Protocol (ARP) suppression works with BGP EVPN to allow a VTEP to reduce the flooding of ARP messages over a VXLAN tunnel with IPv4 addresses.

Neighbor Discovery (ND) suppression works with BGP EVPN to allow a VTEP to reduce the flooding of ND messages over a VXLAN tunnel with IPv6 addresses.

When ARP/ND suppression is enabled, the ARP or ND messages are filtered. The BGP EVPN sends this information to peers. When an ARP/ND request is received and a MAC-IP binding is available for the remote host, an ARP/ND reply is sent to the host. If no ARP/ND reply is sent, the ARP/ND traffic continues to its destination.

ARP/ND suppression only reduces the flooding of messages to known remote hosts. Unknown hosts are still flooded.

You need to configure an SVI for the VXLAN access VLAN before ARP/ND suppression can be enabled. This SVI does not support routing functionalities. Fortinet does not recommend changing this SVI's property.

To enable ARP/ND suppression:

```
config system vxlan
   edit <VXLAN_interface_name>
      set vni <integer>
      set vlanid <integer>
      set interface <interface_name>
      set evpn enable
      set arp-nd-supression enable
   next
end
```

For example:

```
config system vxlan
   edit "vx1"
      set vni 500
      set vlanid 50
      set interface "vlan12"
      set evpn enable
      set arp-nd-supression enable
   next
end
```

To debug ARP/ND suppression:

- `diagnose switch vxlan arp-nd-cache show [<VLAN_ID>]`
- `diagnose switch vxlan arp-nd-cache clear-stats <VLAN_ID>`
- `get router info evpn arp-nd-cache vni [<VNI_number>]`
- `get router info evpn arp-nd-proxy-stats vni <VNI_number>`

Duplicate address detection

Starting in FortiSwitchOS 7.4.1, FortiSwitchOS can detect duplicate MAC addresses in a BGP EVPN with VXLAN interfaces. Duplicate address detection is enabled by default for EVPN configurations.

You can use CLI commands to define how many times the same MAC address is detected moving in the network within a specified time period before it is identified as a duplicate MAC address. By default, MAC addresses that are detected moving five times within 300 seconds are identified as duplicate MAC addresses.

You can specify whether duplicate MAC addresses are locked (frozen) permanently or for a specific number of seconds. By default, duplicate MAC addresses are not locked (frozen).

FortiSwitchOS logs duplicate MAC addresses as errors, making it quicker to find and resolve problems in the network configuration.

To configure BGP EVPN with VXLAN interfaces, see Using BGP EVPN with VXLAN on page 32.

To configure duplicate address detection:

```
config router bgp
   set as <MANDATORY_router_AS_number>
   set router-id <MANDATORY_IP_address>
   config evpn
      set advertise-vni enable
      set dup-addr-detection {enable | disable}
      set dup-addr-freeze {disable | permanent | time}
      set dup-addr-freeze-time <30-3600 seconds>
      set dup-addr-max-moves <1-1000>
      set dup-addr-window-time <2-1800 seconds>
   end
end
```

For example:

```
config router bgp
   set as 100
   set router-id 1.0.48.2
   config neighbor
      edit "1.0.24.1"
         set activate-evpn enable
         set remote-as 100
         set update-source "loopback"
         set route-reflector-client-evpn enable
      next
      edit "1.0.48.1"
         set activate-evpn enable
         set remote-as 100
         set update-source "loopback"
         set route-reflector-client-evpn enable
      next
   end
   .
   .
   .
   config evpn
      set advertise-vni enable
      set dup-addr-freeze time
      set dup-addr-freeze-time 120
```

```
      set dup-addr-max-moves 3
      set dup-addr-window-time 120
   end
end

config system vxlan
   edit "vni1000"
      set vni 1000
      set vlanid 3
      set interface "loopback"
      set evpn enable
   next
end
```

To view all duplicate MAC addresses:

```
get router info evpn mac vni dup-addr
```

To clear all duplicate MAC addresses:

```
execute router clear evpn dup-addr vni all
```

To clear duplicate MAC addresses in a specific VNI:

```
execute router clear evpn dup-addr vni <VNI_number>
```

To clear a specified MAC addresses in a specific VNI:

```
execute router clear evpn dup-addr vni <VNI_number> mac <MAC_address>
```

Using ECMP routing with VXLAN interfaces

Starting in FortiSwitchOS 7.41, you can use flow-based Equal Cost Multi-Path (ECMP) routing with VXLAN interfaces for load balancing. By using multiple routing paths at the same time, ECMP provides redundancy in case of network congestion or failures.

At the ingress VTEP, ECMP is always enabled by default and is not dependent on the VXLAN UDP source port.

For transit routers, you must specify a nonzero source UDP port. The default source UDP port is 0; the range of values is 1-65535. This setting is applied only when the VXLAN interface is created.

ECMP routing with VXLAN interfaces is supported only on known unicast traffic. Broadcast, unknown-unicast (BUM) traffic is not supported.

To use ECMP routing:

```
config switch global
   set vxlan-sport <integer>
end
```

Using DHCP snooping with VXLAN

VXLAN primarily extends layer-2 networks over the layer-3 infrastructure to overcome the limitations of traditional VLANs. The layer-3 network is responsible for delivering the virtualized layer-2 packets encapsulated in the VXLAN header through a layer-3 tunnel. From a layer-2 bridge point of view, it connects with the remote layer-2 bridge using a virtual wire going over the layer-3 tunnel. The layer-2 bridge views the connecting point as a switch virtual port.

Starting in FortiSwitchOS 7.4.3, a virtual switch port is created automatically when there is a VXLAN (VNI) associated with a VXLAN tunnel. This virtual switch port is per VXLAN (VNI) per tunnel. Like DHCP-snooping properties on a normal switch interface, you can configure DHCP-snooping properties on a virtual switch port. Currently, the parameters described in this section can be configured on a virtual switch port. The automatically created name of the virtual switch port is in the following format:

```
vni.<VNI>.<remote_end_VTEP_IP_address>
```

For example, if the VNI is 100 and the remote end of the VXLAN tunnel is at 1.1.1.1, the virtual port name is `vni.100.1.1.1.1`.

A virtual port is deleted when the corresponding VXLAN configuration is deleted or when the corresponding VXLAN tunnel is down and the VXLAN configuration uses the default settings.

To configure DHCP snooping:

```
config switch virtual-port
   edit <virtual_port_name>
      set description <string>
      set dhcp-snooping {trusted | untrusted}
      set dhcp-snoop-learning-limit-check {enable | disable}
      set dhcp-snoop-learning-limit <1-16000>
   next
end
```

Variable	Description	Default
<virtual_port_name>	Enter a name for the virtual port.	No default
description <string>	Enter a description for the virtual port.	No default
dhcp-snooping {trusted | untrusted}	Set the virtual port to trusted or untrusted.	trusted
dhcp-snoop-learning-limit-check {enable | disable}	Enable or disable whether there is a limit for how many IP addresses are in the DHCP-snooping binding database for this virtual port. The `set dhcp-snoop-learning-limit-check` command is available only when `dhcp-snooping` has been set to `untrusted`.	disable
dhcp-snoop-learning-limit <1-16000>	Set the maximum number of IP addresses learned on this virtual port for the DHCP-snooping binding database. The `set dhcp-snoop-learning-limit` command is available only when `dhcp-snoop-learning-limit-check` is enabled.	5

For example:

```
config switch virtual-port
   edit vni.100.1.1.1.1
      set description "virtual port for VNI 100"
      set dhcp-snooping untrusted
      set dhcp-snoop-learning-limit-check enable
      set dhcp-snoop-learning-limit 100
   next
end
```

To display details about the DHCP-snooping client and server database:

```
get switch dhcp-snooping status
```

To display the details about the IPv6 DHCP-snooping server database:

```
get switch dhcp-snooping server6-db-details
```

To display the details about the IPv6 DHCP-snooping client database:

```
get switch dhcp-snooping client6-db-details
```

Configuration example

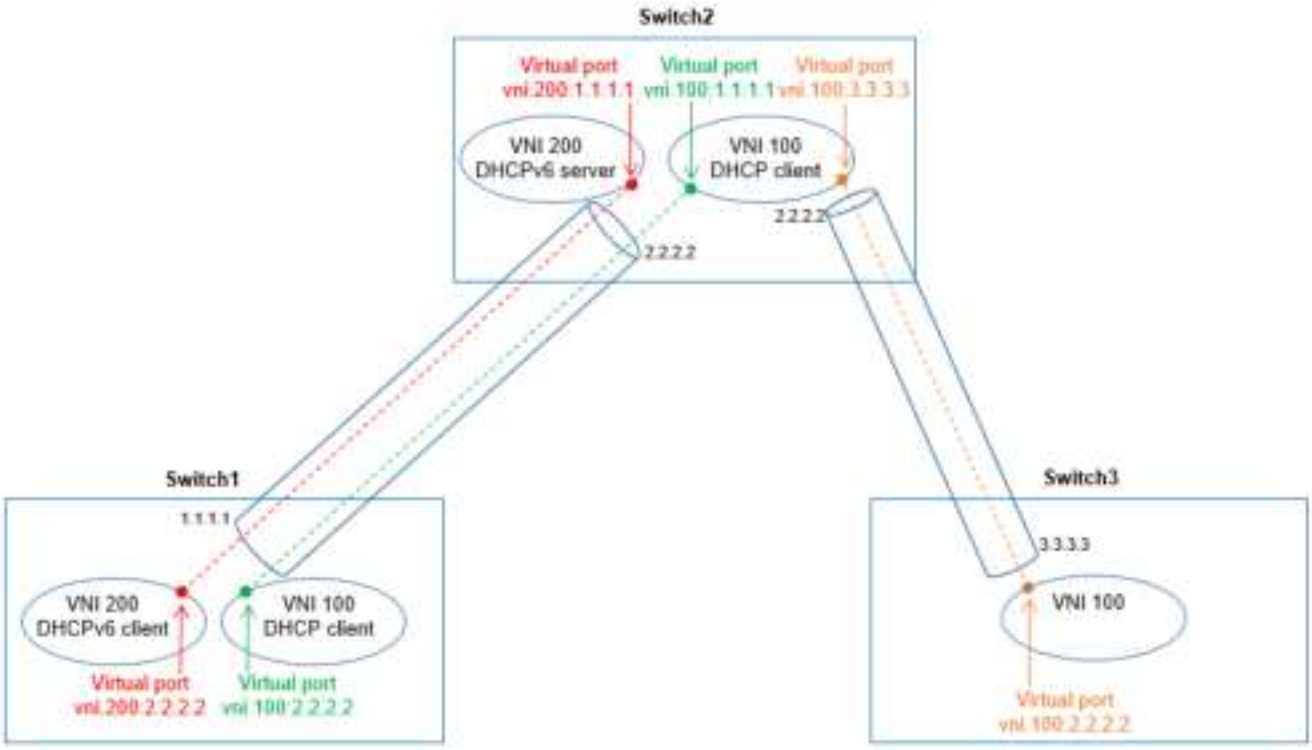

In this example:

- Switch1 has VNI 100 and VNI 200 with EVPN enabled.
- Switch2 has VNI 100 and VNI 200 with EVPN enabled.
- Switch3 has VNI 100 with EVPN enabled.

There are two VXLAN tunnels:

- One between Switch1 and Switch2
- One between Switch2 and Switch3

There are two virtual ports on Switch1:

- vni.100.2.2.2.2
- vni.200.2.2.2.2

There are three virtual ports on Switch2:

- vni.100.1.1.1.1
- vni.100.3.3.3.3
- vni.200.1.1.1.1

There is one virtual port on Switch3:

- vni.100.2.2.2.2

To configure the example topology:

1. Configure the loopbacks.

   To configure Switch1:

   ```
   config system interface
   ...
          edit "loopback1"
                  set ip 1.1.1.1 255.255.255.255
                  set allowaccess ping
                  set type loopback
                  set snmp-index 38
          next
   ...
   ```

   To configure Switch2:

   ```
   config system interface
   ...
          edit "loopback1"
                  set ip 2.2.2.2 255.255.255.255
                  set allowaccess ping
                  set type loopback
                  set snmp-index 38
          next
   ...
   ```

   To configure Switch3:

   ```
   config system interface
   ...
          edit "loopback1"
                  set ip 3.3.3.3 255.255.255.255
                  set allowaccess ping
                  set type loopback
   ```

```
                set snmp-index 38
        next
...
```

2. Configure the VXLAN.

   To configure Switch1:

   ```
   config system vxlan
      edit "vx100"
         set vni 100
         set vlanid 100
         set interface "loopback1"
         set evpn enable
      next
      edit "vx200"
         set vni 200
         set vlanid 200
         set interface "loopback1"
         set evpn enable
      next
   end
   ```

   To configure Switch2:

   ```
   config system vxlan
      edit "vx100"
         set vni 100
         set vlanid 100
         set interface "loopback1"
         set evpn enable
      next
      edit "vx200"
         set vni 200
         set vlanid 200
         set interface "loopback1"
         set evpn enable
      next
   end
   ```

   To configure Switch3:

   ```
   config system vxlan
      edit "vx100"
         set vni 100
         set vlanid 100
         set interface "loopback1"
         set evpn enable
      next
   end
   ```

3. Configure DHCP snooping.

   To configure Switch1:

   ```
   config switch virtual-port
   ```

```
   edit "vni.100:2.2.2.2"
      set dhcp-snooping untrusted
   next
   edit "vni.200:2.2.2.2"
   next
end
```

To configure Switch2:

```
config switch virtual-port
   edit "vni.100:3.3.3.3"
   next
   edit "vni.100:1.1.1.1"
      set dhcp-snooping untrusted
      set dhcp-snoop-learning-limit-check enable
   next
   edit "vni.200:1.1.1.1"
      set dhcp-snooping untrusted
   next
end
```

To configure Switch3:

```
config switch virtual-port
   edit "vni.100:2.2.2.2"
   next
end
```

4. Enable DHCP snooping for the VLANs used by the VXLAN.

To configure Switch1:

```
config switch vlan
   edit vlan100
      set dhcp-snooping enable
   next
   edit vlan200
      set dhcp-snooping enable
   next
end
```

To configure Switch2:

```
config switch vlan
   edit vlan100
      set dhcp-snooping enable
   next
   edit vlan200
      set dhcp-snooping enable
   next
end
```

To configure Switch3:

```
config switch vlan
   edit vlan100
```

```
      set dhcp-snooping enable
    next
end
```

5. Verify that your DHCP-snooping configuration is working correctly.

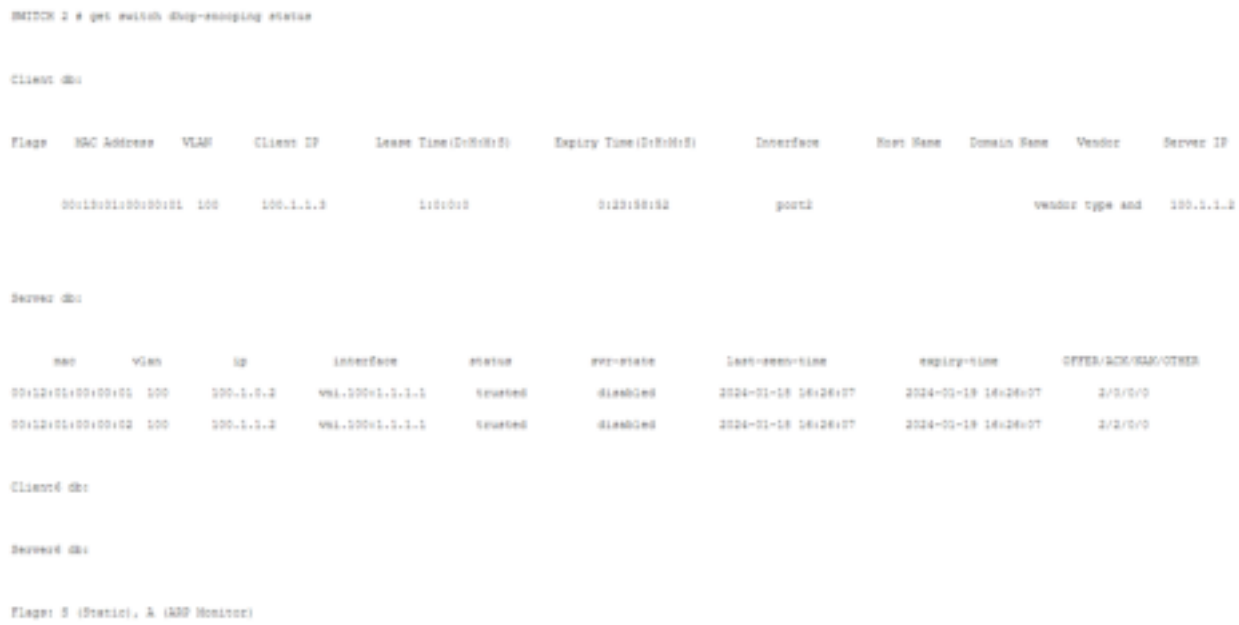

Using QoS with VXLAN tunnels

Starting in FortiSwitchOS 7.4.3, you can add quality of service (QoS) capabilities to VXLAN traffic.

In a VXLAN tunnel, the DSCP values in the outer IP header of VXLAN-encapsulated packets determine the end-to-end QoS on a per-hop basis:

- On the ingress VTEP, the DSCP value in the outer header can be copied from the original header or assigned a fixed value per configuration.
- On the egress VTEP, the priority of decapsulated packets is determined by the DSCP value in the outer header.

All VXLAN tunnels on a FortiSwitch unit must share the same QoS marking configuration for tunnel initiation.

On each FortiSwitch unit, you can configure the Dot1p map or DSCP map on the switch interfaces to determine the VXLAN traffic's queuing priorities. If both the Dot1p map and DSCP map are configured on the same switch interface, they apply to non-IP and IP traffic, respectively.

On the ingress VTEP, you can configure the mapping on the VXLAN access ports. You can configure different mappings on different access ports. Multiple VXLAN virtual ports under the same access port will inherit the same mapping.

On the transit node, VXLAN-encapsulated traffic is treated as regular IP packets for QoS.

On the egress VTEP, the Dot1p map or DSCP map, if configured, is enforced based on the outer header of VXLAN-encapsulated traffic.

To configure QoS for the VXLAN tunnel ingress:

```
config switch global
   set vxlan-qos-inner-to-outer {copy-to-outer | fixed}
   set vxlan-qos-dscp <0-63>
```

```
end
```

Variable	Description	Default
vxlan-qos-inner-to-outer {copy-to-outer | fixed}	Select how the differential service code point (DSCP) is determined: • `copy-to-outer`—Copy the DSCP value from the inner header to the outer header. • `fixed`—Use a fixed DSCP value in the IP header of the outer encapsulation. Specify the fixed value with the `set vxlan-qos-dscp` command.	copy-to-outer
vxlan-qos-dscp <0-63>	Specify the fixed DSCP value in the IP header of the outer encapsulation. This command is available only when `vxlan-qos-inner-to-outer` is set to `fixed`.	0

For example:

```
config switch global
   set vxlan-qos-inner-to-outer fixed
   set vxlan-qos-dscp 20
end
```

Routed VLAN interfaces

A routed VLAN interface (RVI) is a physical port or trunk interface that supports layer-3 routing protocols. When the physical port or trunk is administratively down, the RVI for that physical port or trunk goes down as well. All RVIs use the same VLAN, 4095.

RVIs support ECMP, VRF, multiple IP addresses, IPv4 addresses, IPv6 addresses, BFD, VRRP, DHCP server, DHCP relay, RIP, OSPF, ISIS, BGP, and PIM.

Layer-2 protocols and most switch interface features are disabled on RVIs. When RVI is enabled, the following features are not available:

- 802.1X port mode
- 802.1X MAC-based security mode
- User-based (802.1X) VLAN assignment
- 802.1X enhancements, including MAB
- MAB reauthentication
- open-auth mode
- Support of the RADIUS accounting server
- Support of RADIUS CoA and disconnect messages
- EAP pass-through
- Network device detection
- DHCP snooping
- DHCP blocking
- Dynamic ARP inspection
- Access VLANs

- VLAN tag by ACL
- IGMP snooping
- IGMP proxy
- IGMP querier
- Per-port maximum for learned MACs
- MAC learning limit
- Learning limit violation log
- set mac-violation-timer
- Sticky MAC
- Total MAC entries
- MSTP
- STP root guard
- STP BPDU guard
- 'forced-untagged' or 'force-tagged' setting on switch interfaces
- Private VLANs
- Multi-stage load balancing
- MAC/IP/protocol-based VLAN assignment
- Virtual wire
- Loop guard
- VLAN stacking (QnQ)
- VLAN mapping
- MCLAG
- STP support in MCLAGs
- IGMP snooping support in MCLAG
- Cut-through switching
- Edge port
- Host quarantine on switch port

Configuring an RVI

When you configure a trunk interface as an RVI, you must confgure a static MAC address to avoida disruption of adjacency when adding or removing a group of ports.

Using the CLI:

Create a system interface. Set the IP address and netmask, set the interface type to `physical`, and then assign the layer-2 interface.

```
config system interface
   edit <new_interface_name>
      set ip <IP_address_and_netmask>
      set type physical
      set l2-interface <existing_interface_name>
   next
```

```
end
```

For example:

```
config system interface
   edit RVInew
      set ip 10.1.1.1 255.255.255.0
      set allowaccess ping
      set type physical
      set l2-interface port2
   next
end
```

Configuring VRF for an RVI

Starting in FortiSwitchOS 7.2.1, you can configure port-based virtual routing and forwarding (VRF) for an RVI.

To configure VRF for an RVI:

```
config system interface
   edit <new_interface_name>
      set ip <IP_address_and_netmask>
      set type physical
      set l2-interface <port_name>
      set vrf <VRF_instance_name>
   next
end
```

For example:

```
config system interface
   edit "rvi1"
      set ip 192.168.10.1 255.255.255.0
      set allowaccess ping https http ssh telnet radius-acct
      set type physical
      set l2-interface "port15"
      set snmp-index 77
      set vrf "vrf2"
      config ipv6
         set ip6-address 192:168:10::1/64
         set ip6-allowaccess ping
         set dhcp6-information-request enable
      end
   next
end
```

Viewing the RVIs

Use the following command to list which ports and trunks are RVIs:

```
diagnose ip router fwd l3-rvi-info
```

Use the following command to list MAC addresses, priorities, source ports, and flags for RVIs:

```
diagnose hardware switchinfo l2-station-table
```

VRRP

NOTE: You must have an advanced features license to use VRRP.

The Virtual Router Redundancy Protocol (VRRP) uses virtual routers to control which physical routers are assigned to an access network. A VRRP group consists of a master router and one or more backup routers that share a virtual IP address. If the master router fails, the VRRP automatically assigns one of the backup routers without affecting network traffic. When the failed router is functioning again, it becomes the master router again. VRRP provides this redundancy without user intervention or additional configuration to any of the devices on the network.

To create a VRRP group, you need to create a VRRP virtual MAC address, which is a shared MAC address adopted by the VRRP master. The VRRP virtual MAC address feature is disabled by default. You must enable the VRRP virtual MAC address feature on all members of a VRRP group.

The VRRP master router sends VRRP advertisement messages to the backup routers. When the VRRP master router fails to send advertisement messages, the backup router with the highest priority takes over as the master router.

Configuring VRRP

Using the GUI:

1. Go to *System* > *Network* > *Interface* > *Physical*.
2. Click *Edit* for the appropriate interface.
3. Click *Add VRRP* to add a virtual router.
   - Enter the unique virtual router identifier (VRID).
   - Enter the VRRP group number.

     NOTE: Specifying the VRRP group number is optional. If you do not specify it, the value defaults to 0. If you want to change the VRRP group number to 0 for an existing VRRP entry, click *Remove* to delete the VRRP entry and re-enter the VRRP entry without specifying the VRRP group number.
   - Enter the priority. If the highest priority value of 255 is entered, the virtual router becomes the master router.
   - Select *Preempt* if you want the router to preempt the master virtual router if the priority changes.
   - Enter the source virtual IP address that will be shared across the VRRP group.
   - Enter one or two IP addresses that the master router must track. The maximum number of IP addresses is two. If these IP addresses cannot be reached by the master router, the priority of the master router changes to 0.
   - Select *Add VRRP* to add each additional virtual router.
4. After filling in the fields for the virtual routers, click *Update*.

Using the CLI:

```
config system interface
   edit <VLAN name>
      set ip <IP address> <netmask>
      set allowaccess <access_types>
      set vrrp-virtual-mac enable
         config vrrp
            edit <VRRP router identifier>
               set adv-interval <seconds>
               set preempt {enable | disable}
               set priority <priority_number>
               set start-time <seconds>
               set status {enable | disable}
```

```
            set version {2 | 3}
            set vrdst <IPv4_address>
            set vrgrp <VRRP_group_number>
            set vrip <IPv4_address>
         next
      end
   set snmp-index <index number>
   set vlanid <VLAN identifier>
   set interface "internal"
  next
end
```

NOTE: You can also configure VRRP using IPv6 with the `config ipv6` and `config vrrp6` commands under the `config system interface` command.

Example of configuring VRRP using IPv4

In this example, the two FortiSwitch units, FSW-1 and FSW-2, function as both master and backup routers. For VRRP 10, FSW-1 is the master router, and FSW-2 is the backup router. For VRRP, FSW-1 is that standby router, and FSW-2 is the master router. This configuration allows the switches to balance the load and provide redundancy to each other. The downstream clients can split their gateways into two virtual routers, 10.10.10.255 and 10.10.20.255.

For the FSW-1 switch, VRID 10 has the highest priority of 255, so it is the master router; VRID 20 is the backup router.

```
config system interface
  edit "vlan-8"
    set ip 10.10.1.1 255.255.0.0
    set allowaccess ping https http ssh telnet snmp
    set vrrp-virtual-mac enable
    config vrrp
      edit 10
        set priority 255
        set vrip 10.10.10.255
      next
      edit 20
        set vrip 10.10.20.255
      next
    end
    set snmp-index 20
    set vlanid 8
    set interface "internal"
  next
end
```

For the FSW-2 switch, VRID 10 is the backup router; VRID 20 has the highest priority of 255, so it is the master router.

```
config system interface
  edit "vlan-8"
    set ip 10.10.1.2 255.255.0.0
    set allowaccess ping https http ssh telnet snmp
    set vrrp-virtual-mac enable
    config vrrp
      edit 10
        set vrip 10.10.10.255
      next
      edit 20
        set priority 255
```

```
            set vrip 10.10.20.255
        next
    end
    set snmp-index 20
    set vlanid 8
    set interface "internal"
  next
end
```

Checking the VRRP configuration

Using the GUI:

Go to *Router* > *Config* > *Interface* to see which interfaces have VRRP configured.

Go to *Router* > *Monitor* > *VRRP* to see the interface, source virtual IP address that is shared across the VRRP group, MAC address for the interface, and virtual router identifier for each VRRP configuration, as shown in the following figure.

Using the CLI:

```
get router info vrrp
```

Loopback

A loopback interface is a special virtual interface created in software that is not associated with any hardware interface.

Dynamic routing protocols typically use a loopback interface as a reliable IP interface for routing updates. You can assign the loopback IP address to the router rather than the IP address of a specific hardware interface. Services (such as Telnet) can access the router using the loopback IP address, which remains available independent of hardware interfaces status.

No limit exists on the number of loopback interfaces you can create.

A loopback interface does not have an internal VLAN ID or a MAC addresses and usually has a /32 network mask.

Using the GUI:

1. Go to *System* > *Network* > *Interface* > *Loopback*.
2. Select *Add Interface*.
3. Enter a name for the loopback interface.
4. Select *Static* for the mode and then enter the IP address and netmask in the *IP/Netmask* field.
5. Select the protocols allowed to access the loopback interface.

6. Select the administration status.
7. Select *Add*.

Using the CLI:

```
config system interface
   edit "loopback"
      set ip 172.168.20.1 255.255.255.255
      set allowaccess ping https http ssh telnet
      set type loopback
      set snmp-index 28
   next
end
```

IP conflict detection

IP conflicts can occur when two systems on the same network are using the same IP address. The FortiSwitch unit monitors the network for conflicts and raises a system log message and an SNMP trap when it detects a conflict.

The IP conflict detection feature provides two methods to detect a conflict. The first method relies on a remote device to send a broadcast ARP (Address Resolution Protocol) packet claiming ownership of a particular IP address. If the IP address in the source field of that ARP packet matches any of the system interfaces associated with the receiving FortiSwitch system, the system logs a message and raises an SNMP trap.

For the second method, the FortiSwitch unit actively broadcasts gratuitous ARP packets when any of the following events occurs:

- System boot-up
- Interface status changes from down to up
- IP address change

If a system is using the same IP address, the FortiSwitch unit receives a reply to the gratuitous ARP. If it receives a reply, the system logs a message.

Configuring IP conflict detection

IP conflict detection is enabled on a global basis. The default setting is enabled.

Using the GUI:

1. Go to *Network* > *Settings*.
2. Select *Enable IP Conflict Detection*.
3. Select *Apply*.

Using the CLI:

```
config system global
   set detect-ip-conflict <enable|disable>
```

Viewing IP conflict detection

If the system detects an IP conflict, the system generates the following log message:

```
IP Conflict: conflict detected on system interface mgmt for IP address 10.10.10.1
```

ARP timeout value

By default, ARP entries in the cache are removed after 180 seconds. Use the following commands to change the default ARP timeout value:

```
config system global
   set arp-timeout <seconds>
end
```

For example, to set the ARP timeout to 1,000 seconds:

```
config system global
   set arp-timeout 1000
end
```

Using SSH and the Telnet client

Starting in FortiSwitchOS 6.2.0, you can use both IPv4 and IPv6 addresses with SSH and Telnet. If the IPv6 address is a link-local address, you must specify an output interface using %. For example:

```
execute ssh admin@fe80::926c:acff:fe7b:e059%vlan20 // vlan20 is the output interface.
execute ssh admin@172.20.120.122
execute ssh 1002::21
execute ssh 12.345.6.78
execute telnet fe80::926c:acff:fe7b:e059%vlan20 // vlan20 is the output interface.
execute telnet 1002::21
execute telnet 12.345.6.78
```

Starting in FortiSwitchOS 7.4.3, you can use the CLI to regenerate the SSH server keys:

```
execute ssh-regen-keys
```

After you enter the command, the SSH server restarts, and the current SSH connections are disconnected.

Config

The following topics provide information about system configuration:

- SNMP on page 52
- Firmware on page 57
- Backup on page 60
- Revisions on page 60

- Licenses on page 61
- Time on page 63
- SSL on page 64
- Configuring the temperature sensor on page 66

SNMP

Simple Network Management Protocol (SNMP) enables you to monitor hardware on your network.

The FortiSwitch SNMP implementation is read-only. SNMP v1-compliant and v2c-compliant SNMP managers have read-only access to FortiSwitch system information through queries and can receive trap messages from the FortiSwitch unit.

To monitor FortiSwitch system information and receive FortiSwitch traps, you must first compile the Fortinet and FortiSwitch management information base (MIB) files. A MIB is a text file that describes a list of SNMP data objects that are used by the SNMP manager. These MIBs provide information that the SNMP manager needs to interpret the SNMP trap, event, and query messages sent by the FortiSwitch SNMP agent.

You can download the Fortinet and FortiSwitch MIB files from Customer Service & Support.

To download the FortiSwitch MIB file from Customer Service & Support:

1. Go to https://support.fortinet.com.
2. Click *Login Now*.
3. Enter your email address and password and then click *LOG IN*.
4. Go to *Support* > *Firmware Download*.
5. From the *Select Product* dropdown list, select *FortiSwitch*.
6. Select the *Download* tab.
7. Select the directory for the image build that you are using.
8. Click the *HTTPS* link for `FORTINET-FORTISWITCH-MIB.mib` to download it.

To download the Fortinet core MIB file from Customer Service & Support:

1. Go to https://support.fortinet.com.
2. Click *Login Now*.
3. Enter your email address and password and then click *LOG IN*.
4. Go to *Support* > *Firmware Download*.
5. From the *Select Product* dropdown list, select *FortiGate* if it is not already selected.
6. Select the *Download* tab.
7. Select the directory for the image build that you are using.
8. Click the *MIB* directory to open it.
9. Click the *HTTPS* link for `FORTINET-CORE-MIB-buildxxxx.mib` to download it.

When you use the dot1dTpFdbTable table, the index provided does not contain the dot1dTpFdbAddress as defined by the standard. Instead, the index is an increasing numerical value.

This section covers the following topics:

- SNMP access on page 53
- SNMP agent on page 53

- SNMP on page 52
- SNMP v3 user on page 55
- Configuration example on page 56

SNMP access

Ensure that the management VLAN has SNMP added to the access-profiles.

Using the GUI:

1. Go to *System* > *Network* > *Interface* > *Physical*.
2. Select *Edit* for the *mgmt* interface.
3. Select *SNMP* in the access section.
4. Select *Update*.

Using the CLI:

```
config system interface
   edit <name>
      set allowaccess <access_types>
   end
end
```

NOTE: Re-enter the existing allowed access types and add `snmp` to the list.

SNMP agent

Create the SNMP agent.

Using the GUI:

1. Go to *System* > *Config* > *SNMP* > *Settings*.
2. Select *Agent Enabled*.
3. Enter a descriptive name for the agent.
4. Enter the location of the FortiSwitch unit.
5. Enter a contact or administrator for the SNMP agent or FortiSwitch unit.
6. Select *Apply*.

Using the CLI:

```
config system snmp sysinfo
   set status enable
   set contact-info <contact_information>
   set description <description_of_FortiSwitch>
   set location <FortiSwitch_location>
end
```

SNMP community

An SNMP community is a grouping of devices for network administration purposes. Within that SNMP community, devices can communicate by sending and receiving traps and other information. One device can belong to multiple communities, such as one administrator terminal monitoring both a FortiGate SNMP and a FortiSwitch SNMP community.

Add SNMP communities to your FortiSwitch unit so that SNMP managers can connect to view system information and receive SNMP traps.

You can add up to three SNMP communities. Each community can have a different configuration for SNMP queries and traps. Each community can be configured to monitor the FortiSwitch unit for a different set of events. You can also add the IP addresses of up to eight SNMP managers for each community.

Starting in FortiSwitchOS 7.0.0, you can set up one or more SNMP v3 notifications (traps) in the CLI. The following notifications are supported:

- The CPU usage is too high.
- The configuration of an entity was changed.
- The IP address for an interface was changed.
- The available log space is low.
- The available memory is low.

Starting in FortiSwitchOS 7.0.2, you can configure an SNMP trap so that you receive a message when the MAC learning limit is exceeded.

Starting in FortiSwitchOS 7.2.0, you can configure an SNMP trap so that you receive a message when a layer-2 MAC address has been added, deleted, or moved. This SNMP trap applies only to dynamic MAC addresses learned on the port.

Starting in FortiSwitchOS 7.2.1, you can configure SNMP traps for the following:

- The fan was detected, not detected, resumed, or failed.
- There is a conflict between IP addresses.
- The status of the power supply unit has changed.
- The sensor triggered an alarm.
- The sensor is faulty.
- The trunk member's heart beat is unsynchronized.

Starting in FortiSwitchOS 7.4.3, you can configure SNMP traps for the following:

- There has been a change in the storm-control status.
- Custom triggers for automation stitches.

By default, all SNMP notifications are enabled, except for `l2mac`. Notifications are sent to one or more IP addresses.

Adding an SNMP v1/v2c community

After you run the `execute factoryreset` command, FortiSwitchOS creates an SNMP community with the `name` set to `public`.

Using the GUI:

1. Go to *System* > *Config* > *SNMP* > *Communities*.
2. Select *Add Community*.
3. Enter a community name and identifier.
4. Select *Add Host* and enter the identifier, IP address and netmask, and interface for each host.
5. Select *V1*, *V2C*, or both and enter the port number that the SNMP managers in this community use for SNMP v1 and SNMP v2c queries to receive configuration information from the FortiSwitch unit.
6. Select *V1*, *V2C*, or both and enter the local and remote port numbers that the FortiSwitch unit uses to send SNMP v1 and SNMP v2c traps to the SNMP managers in this community.
7. Select which events to report.
8. Select *Add*.

Using the CLI:

```
config system snmp community
   edit <index_number>
      set events <cpu-high | ent-conf-change | fan-detect | fsTrapStitch1 | fsTrapStitch2 |
           fsTrapStitch3 | fsTrapStitch4 | fsTrapStitch5 | intf-ip | ip-conflict | l2mac |
           llv | log-full | mem-low | psu-status | sensor-alarm | sensor-fault | storm-
           control | tkmem-hb-oo-sync>
      set name <community_name>
      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 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}
      config hosts
         edit <host_number>
            set interface <interface_name>
            set ip <IPv4_address/mask>
            set source-ip <IPv4_address>
         end
      config hosts6
         edit <host_number>
            set interface <interface_name>
            set ip6 <IPv6_address>
            set source-ip6 <IPv6_address>
         end
      end
```

SNMP v3 user

Using the GUI:

1. Go to *System* > *Config* > *SNMP* > *Users*.
2. Select *Add User*.

3. Enter a user name.
4. Select a security level to specify the authentication and privacy settings.
5. Enter the port number that the SNMP managers in this community use to receive configuration information from the FortiSwitch unit.
6. Make certain that *Enable Queries* is enabled.
7. Select *Add*.

Using the CLI:

```
config system snmp user
   edit <index_number>
      set queries enable
      set query-port <port_number>
      set security-level [auth-priv | auth-no-priv | no-auth-no-priv}
      set auth-proto {md5 | sha1 | sha224 | sha256 | sha384 | sha512}
      set events {cpu-high | ent-conf-change | fan-detect | fsTrapStitch1 | fsTrapStitch2 |
           fsTrapStitch3 | fsTrapStitch4 | fsTrapStitch5 | intf-ip | ip-conflict | l2mac |
           llv | log-full | mem-low | psu-status | sensor-alarm | sensor-fault | storm-
           control | tkmem-hb-oo-sync}
      set notify-hosts <IP_address>
      set auth-pwd <password>
      set priv-proto {aes128 | aes192 | aes192c | aes256 | aes256c | des}
      set priv-pwd <password>
   end
```

Configuration example

In the following example, the SNMP trap is triggered when there is a change in the storm-control status on 1.2.3.4. When the storm-control status changes, the FortiSwitch unit generates a log message and sends the SNMP v1/v2c traps to 1.2.3.4.

NOTE: If you are using the mgmt interface, you must allow access to SNMP, as well as enabling SNMP.

```
config system interface
   edit mgmt
      set allowaccess ping https ssh snmp
   end
end

config system snmp sysinfo
   set status enable
end

config system snmp community
   edit 1
      set events storm-control
      config hosts
         edit 1
            set ip 1.2.3.4 255.255.255.0
         next
      end
      set name "public"
   next
end
```

```
config system snmp user
   edit "1"
      set notify-hosts 1.2.3.4
      set events storm-control
      set auth-pwd ENC y0rlV4kXyEf5/YWnFgj5B+pazUL
      set priv-pwd ENC YuyfrHV30EwgRJNscY
   next
end
```

Firmware

Starting in FortiSwitchOS 7.4.0, the GUI provides OS image signature verification. To see which models support this feature, refer to the FortiSwitch feature matrix. If the BIOS version does not support OS image signature verification, the GUI displays a warning message when you log in.

- If you upload an unverified firmware image, the GUI displays a “WARNING: This firmware failed signature validation.” message.
- If you log in to a FortiSwitch unit running an unverified firmware image, the GUI displays an “Unverified Image Detected” message.
- After you log in to a FortiSwitch unit running an unverified firmware image, the GUI displays a triangle with a red exclamation mark in the title bar.

This section covers the following topics:

- Upgrading the firmware on page 57
- Verifying image integrity on page 58
- Setting the boot partition on page 59
- Restoring or upgrading the BIOS on page 59

Upgrading the firmware

Use these procedures to upgrade your FortiSwitch firmware.

Using the GUI

1. Go to *System* > *Config* > *Firmware*.
2. Click *Choose File* and then navigate to the firmware image.

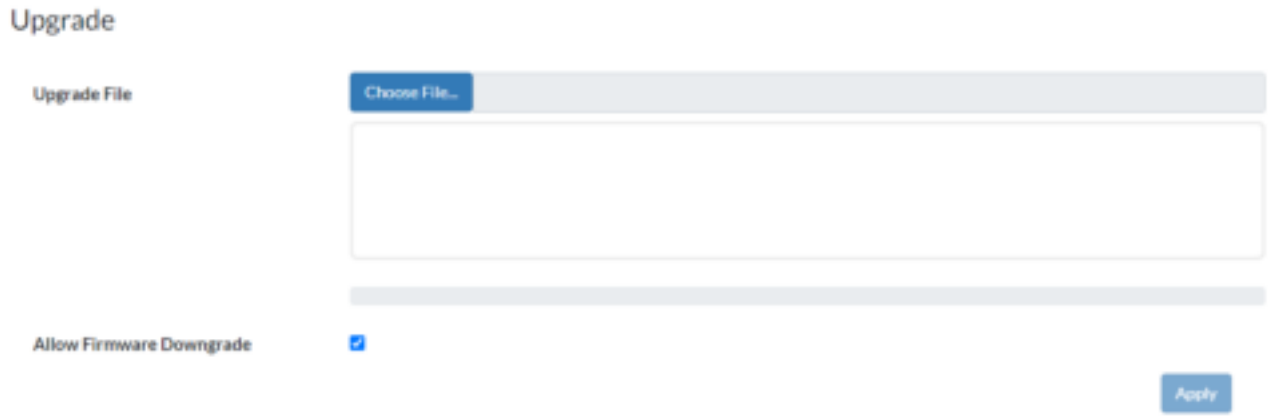

3. Select *Apply*.
4. If the firmware image is unverified, the GUI displays a “WARNING: This firmware failed signature validation.” message. You must click *Continue* if you still want to upgrade to this firmware image.

Using the CLI

You can download a firmware image from a FortiManager unit or from an FTP, SFTP, or TFTP server. The FortiSwitch unit reboots and then loads the new firmware.

```
execute restore image management-station <version_int>
execute restore image ftp <filename_str> <server_ipv4_ipv6_fqdn[:port_int]> [<username_str>
     <password_str>]
execute restore image sftp <filename_str> <server_ipv4_ipv6_fqdn[:port_int]> [<username_str>
     <password_str>]
execute restore image tftp <filename_str> <server_ipv4_ipv6_fqdn> [<source_ipv4_ipv6>]
```

You can download a previously backed-up configuration file from a flash disk or from an FTP, SFTP, or TFTP server. The FortiSwitch unit reboots and then loads the new firmware.

```
execute restore config flash <revision>
execute restore config ftp <filename_str> <server_ipv4_ipv6_fqdn[:port_int]> [<username_str>
     <password_str>] [<backup_password_str>]
execute restore config sftp <filename_str> <server_ipv4_ipv6_fqdn[:port_int]> [<username_
     str> <password_str>] [<backup_password_str>]
execute restore config tftp <filename_str> <server_ipv4_ipv6_fqdn> [<backup_password_str>]
```

The following example shows how to download a configuration file from a TFTP server to the FortiSwitch unit and restart the FortiSwitch unit with this configuration. The name of the configuration file on the TFTP server is `backupconfig`. The IPv4 address of the TFTP server is 192.168.1.23. The source IPv4 address is 1.2.3.4.

```
execute restore config tftp backupconfig 192.168.1.23 1.2.3.4
```

The following example shows how to download a configuration file from an SFTP server to the FortiSwitch unit and restart the FortiSwitch unit with this configuration. The name of the configuration file on the SFTP server is `backupconfig`. The IPv6 address of the SFTP server is 6001:7:7:7::2, and the port number is 2222. To access the SFTP server, you need to add the user name, `admin`, and the password, `adminpassword`.

```
execute restore config sftp backupconfig [6001:7:7:7::2]:2222 admin adminpassword
```

You can also download a firmware image from an FTP, SFTP, or TFTP server and stage it without restarting the FortiSwitch unit:

```
execute stage image ftp <file_name> <ftp server>[:ftp port] [<FTP_user_name> <FTP password>]
execute stage image sftp <file_name> <sftp server>[:sftp port] <SFTP_user_name> <SFTP
     password>
execute stage image tftp <file_name> <tftp server> [<source_IPv4_IPv6_address>]
```

Verifying image integrity

To check if the firmware image is verified:

```
get system status
```

To verify the integrity of the images in the primary and secondary (if applicable) flash partitions, use the following commands:

```
execute verify image primary
execute verify image secondary
```

If the image is corrupted or missing, the command fails with a return code of -1.

For example:

```
execute verify image primary

Verifying the image in flash......100%
No issue found!

execute verify image secondary

Verifying the image in flash......100%
Bad/corrupted image found in flash!
Command fail. Return code -1
```

Setting the boot partition

You can specify the flash partition for the next reboot. The system can use the boot image from either the primary or the secondary flash partition:

```
execute set-next-reboot <primary | secondary>
```

NOTE: You must disable image rotation before you can use the `execute set-next-reboot` command.

If your FortiSwitch model has dual flash memory, you can use the primary and backup partitions for image rotation. By default, this feature is enabled.

```
config system global
   set image-rotation <enable | disable>
end
```

To list all of the flash partitions:

```
diagnose sys flash list
```

Restoring or upgrading the BIOS

You can restore or upgrade the basic input/output system (BIOS) if needed. After a BIOS upgrade, passwords for all FortiSwitch local users must be reconfigured using the `config user local` setting.

CAUTION: Only restore or upgrade the BIOS if Customer Support recommends it.

To upgrade or restore the BIOS from the CLI:

```
execute restore bios tftp <filename_str> <server_ipv4_ipv6_fqdn>
```

For example:

```
execute restore bios tftp PPC/FS-3032D/04000009/FS3D323Z14000004.bin 10.105.2.201
```

The example downloads the BIOS file from the TFTP server at the specified IPv4 address.

NOTE: If the BIOS upgrade fails, do not restart the FortiSwitch unit. Instead, try the CLI command again. If repeating the CLI command does not work, the FortiSwitch unit might require a return merchandise authorization (RMA).

Backup

You can set preferences for saving configuration files:

1. Go to *System* > *Config* > *Backup*.
2. Select one of the *Configuration Save* options:
   - *Automatically Save*—The system automatically saves the configuration after each change.
   - *Manually Save*—You must manually save configuration changes from the *Backup* link on the *System* > *Dashboard*.
   - *Manually Save and Revert Upon Timeout*—You must manually save configuration changes. The FortiSwitch unit reboots and reverts to the saved configuration after the timeout and a restart. You can set the timeout using the CLI:
     ```
     config system global
     set cfg-revert-timeout <integer>
     ```
3. If you select *Revision Backup on Logout*, the FortiSwitch unit creates a configuration file each time a user logs out.
4. If you select *Revision Backup on Upgrade*, the FortiSwitch unit creates a configuration file before starting a system upgrade.
5. Select *Update*.

Revisions

You can revert your FortiSwitch configuration to a previous revision.

The following are the maximum numbers of saved configuration revisions:

FortiSwitch models	Maximum number of saved configuration revisions
FS-1xx-FS-2xx	20 revisions
FS-4xx-FS-6xx	40 revisions
FS-1xxx-FS-3xxx	80 revisions

Using the GUI:

1. Go to *System* > *Config* > *Revisions*.
   The system displays a new page with an entry for each configuration file revision.
2. Select a revision and then click *Revert* to revert the system configuration to the selected revision.
3. In the *Confirm Revert* dialog, click *Revert*.

Using the CLI to revert to the last saved configuration:

```
config system global
   set cfg-save revert
   set cfg-revert-timeout <10-2147483647 seconds>
end
```

For example:

```
config system global
   set cfg-save revert
   set cfg-revert-timeout 20
end
```

This example saves the current configuration, waits 20 seconds, restarts the FortiSwitch unit, and reverts to the last saved configuration if the configuration was not manually saved within the timeout period. Before FortiSwitchOS 7.2.1, there was no restart before the configuration was reverted.

Using the CLI to check the configuration file revisions:

Use the following command to display the list of configuration file revisions:

```
execute revision list config
```

The FortiSwitch unit assigns a numerical ID to each configuration file. To display a particular configuration file contents, use the following command and specify the ID of the configuration file:

```
execute revision show config id <ID number>
```

The following example displays the list of configuration file revisions:

```
# execute revision list config

ID TIME ADMIN FIRMWARE VERSION COMMENT
1 2015-08-31 11:11:00 admin V3.0.0-build117-REL0 Automatic backup (session expired)
2 1969-12-31 16:06:29 admin V3.0.0-build150-REL0 baseline
3 2015-08-31 15:19:31 admin V3.0.0-build150-REL0 baseline
4 2015-08-31 15:28:00 admin V3.0.0-build150-REL0 with admin timeout
```

The following example displays the configuration file contents for revision ID 62:

```
# execute revision show config id 62

#config-version=FS1D24-3.04-FW-build171-160201:opmode=0:vdom=0:user=admin
#conf_file_ver=1784779075679102577
#buildno=0171
#global_vdom=1
config system global
   set admin-concurrent enable
       ...
   (output truncated)
```

Licenses

Advanced features (such as dynamic routing protocols) require a feature license.

Each feature license is tied to the serial number of the FortiSwitch unit. Therefore, a feature license is valid on one system.

This section covers the following topics:

- Checking the license status on page 62
- Adding a license on page 62
- Downloading a license file on page 62
- Removing a license on page 63

Checking the license status

Using the GUI:

Go to *System* > *Config* > *Licenses*.

Using the CLI:

```
execute license status
```

Adding a license

NOTE: Adding license keys causes the system to log you out.

Using the GUI:

1. Go to *System* > *Config* > *Licenses*.
2. Click *Add License*.
3. Enter your license key.
4. Select *Add*.

Using the CLI:

```
execute license add <key>
```

Downloading a license file

Starting in FortiSwitchOS 7.4.3, the Advanced Features License is a text file signed by the Fortinet certificate authority (CA) for better security and includes the license key. The licensing SKUs remain the same. The updated license file is backwards compatible if FortiSwitchOS is downgraded.

Starting in FortiSwitchOS 7.4.3, FortiSwitchOS requires the license file to be uploaded using the CLI, instead of requiring the license key. You can download the license file from an FTP , SFTP, or TFTP server to the FortiSwitch unit. You can use an IPv4 address, IPv6 address, or FQDN to specify the FTP, SFTP, or TFTP server.

If you already have an Advanced Features License, you can get the new license file from the FortiCare team. The new license key in the updated license file will not match the previous license key but will still be accepted by older versions of FortiSwitchOS.

To download the Advanced Features License file:

```
execute restore license ftp <filename_str> <server_ipv4_ipv6_fqdn[:port_int]> [<username_
     str> <password_str>]
execute restore license sftp <filename_str> <server_ipv4_ipv6_fqdn[:port_int]> [<username_
     str> <password_str>]
execute restore license tftp <filename_str> <server_ipv4_ipv6_fqdn>
```

The following example shows how to download a license file from a TFTP server to the FortiSwitch unit. The name of the license file on the TFTP server is `newlicense`. The IP address of the TFTP server is 192.168.1.23.

```
execute restore license tftp newlicense 192.168.1.23
```

The following example shows how to download a license file from an SFTP server to the FortiSwitch unit. The name of the license file on the SFTP server is `newlicense`. The IPv6 address of the SFTP server is 6001:7:7:7::2, and the port number is 2222. To access the SFTP server, you need to add the user name, `admin`, and the password, `adminpassword`.

```
execute restore license sftp newlicense 6001:7:7:7::2]:2222 admin adminpassword
```

Removing a license

Using the GUI:

1. Go to *System* > *Config* > *Licenses*.
2. Click *Delete* for the license to remove
3. Click *Delete* to acknowledge the warning.

NOTE: Deleting license keys causes the system to log you out before rebooting. You will lose all configurations related to the license.

Using the CLI:

```
execute license type <type> clear
```

Time

For effective scheduling and logging, the system date and time must be accurate. You can either manually set the system date and time or configure the system to automatically keep its time correct by synchronizing with a Network Time Protocol (NTP) server.

NOTE: Some FortiSwitch models do not have a battery-backup real-time clock. For FortiSwitch models without a real-time clock, the time is reset when the switch is rebooted. These models must be connected to an NTP server if you want to maintain the correct system date and time.

The Network Time Protocol enables you to keep the system time synchronized with other network systems. This will also ensure that logs and other time-sensitive settings are correct.

When the system time is synchronized, polling occurs every 2 minutes. When the system time is not synchronized but the NTP server can be reached, polling is attempted every 2 seconds to synchronize quickly. If the NTP server cannot be reached, polling occurs up to every 64 seconds. If DNS cannot resolve the host name, polling occurs up to every 60 seconds.

Starting in FortiSwitchOS 6.4.0, the default Sync Interval is 10 minutes. The polling interval is one-fifth of the configured Sync Interval.

Using the GUI:

1. Go to *System* > *Dashboard*.

2. Next to the *System Time* field, select *Change*.

3. Select your *Time Zone*.
4. Either select *Manual Setting* and enter the system date and time or select *Synchronize with NTP Server*. If you select synchronization, you can either use the default FortiGuard server or specify a different server. You can also set the *Sync Interval*.
5. Select *Update*.

Using the CLI:

If you use an NTP server, you can identify the IPv4 or IPv6 address for this self-originating traffic with the `set source-ip` or `set source-ip6` command. For example, you can set the source IPv4 address of NTP to be on the DMZ1 port with an IP of 192.168.4.5:

```
config system ntp
   set authentication enable
   set ntpsyn enable
   set syncinterval 5
   set source-ip 192.168.4.5
end
```

SSL

You can set strong cryptography and select which certificates are used by the FortiSwitch unit.

When you enable strong cryptography, the following ciphers and algorithms are supported:

- Ciphers (encryption algorithms):
  - chacha20-poly1305@openssh.com
  - aes128-ctr
  - aes192-ctr
  - aes256-ctr
  - aes128-gcm@openssh.com
  - aes256-gcm@openssh.com
- Key-exchange algorithms:
  - curve25519-sha256@libssh.org
  - diffie-hellman-group-exchange-sha256
- Host-key algorithm:
  - ssh-ed25519
- Message authentication code algorithms:

- umac-128-etm@openssh.com
- hmac-sha2-256-etm@openssh.com
- hmac-sha2-512-etm@openssh.com

Using the GUI:

1. Go to *System* > *Config* > *SSL*.
2. By default, the *Strong Crypto* checkbox is selected so that FortiSwitchOS uses strong cryptography for HTTPS and SSH access.

   If you clear the *Strong Crypto* checkbox, FortiSwitchOS displays a warning that the switch will reboot and then requires you to confirm before rebooting the switch.
3. Select one of the 802.1X certificate options:
   - *Entrust_802.1x*—This certificate is embedded in the firmware and is the same on every unit (not unique). It has been signed by a public CA. This is the default certificate for 802.1X authentication.
   - *Fortinet_Factory*—This certificate is embedded in the hardware at the factory and is unique to this unit. It has been signed by a proper CA.
   - *Fortinet_Factory2*—This certificate is embedded in the hardware at the factory and is unique to this unit. It has been signed by a proper CA.
   - *Fortinet_Firmware*—This certificate is embedded in the firmware and is the same on every unit (not unique). It has been signed by a proper CA. It is not recommended to use it for server-type functionality since any other unit could use this same certificate to spoof the identity of this unit.
4. Select one of the 802.1X certificate authority (CA) options:
   - *Entrust_802.1x_CA*—Select this CA if you are using 802.1X authentication.
   - *Entrust_802.1x_G2_CA*—Select this CA if you want to use the Google Internet Authority G2.
   - *Entrust_802.1x_L1K_CA*—Select this CA if you want to use http://ocsp.entrust.net.
   - *Fortinet_CA*—Select this CA if you want to use the factory-installed certificate.
   - *Fortinet_CA2*—Select this CA if you want to use the factory-installed certificate.
5. Select one of the GUI HTTPS certificate options:
   - *Entrust_802.1x*—This certificate is embedded in the firmware and is the same on every unit (not unique). It has been signed by a public CA.
   - *Fortinet_Factory*—This certificate is embedded in the hardware at the factory and is unique to this unit. It has been signed by a proper CA.
   - *Fortinet_Factory2*—This certificate is embedded in the hardware at the factory and is unique to this unit. It has been signed by a proper CA.
   - *Fortinet_Firmware*—This certificate is embedded in the firmware and is the same on every unit (not unique). It has been signed by a proper CA. It is not recommended to use it for server-type functionality since any other unit could use this same certificate to spoof the identity of this unit.
6. Select *Update*.

Using the CLI:

```
config system global
   set strong-crypto {enable | disable}
   set 802.1x-certificate {Entrust_802.1x | Fortinet_Factory | Fortinet_Factory2 | Fortinet_
        Firmware}
   set 802.1x-ca-certificate {Entrust_802.1x_CA | Entrust_802.1x_G2_CA | Entrust_802.1x_L1K_
        CA | Fortinet_CA | Fortinet_CA2}
   set admin-server-cert {self-sign | Entrust_802.1x | Fortinet_Factory | Fortinet_Factory2
        | Fortinet_Firmware}
end
```

Configuring the temperature sensor

If your FortiSwitch unit has a temperature sensor, you can set a warning and an alarm for when the system temperature reaches specified temperatures. When these thresholds are exceeded, a log message and SNMP trap are generated. The warning threshold must be lower than the alarm threshold.

Use the following commands to set warning and alarm thresholds:

```
config system snmp sysinfo
   set status enable
   set trap-temp-warning-threshold <temperature in degrees Celsius>
   set trap-temp-alarm-threshold <temperature in degrees Celsius>
end
```

By default, the FortiSwitch unit generates an alert (in the form of an SNMP trap and a SYSLOG entry) every 30 minutes when the temperature sensor exceeds its set threshold. You can change this interval with the following commands:

```
config system global
   set alertd-relog enable
   set alert-interval <1-1440 minutes>
end
```

Admin

The following topics provide information about FortiSwitch administration:

- Administrators on page 66
- Profiles on page 75
- Access control on page 77
- Monitor on page 80
- Setting the idle timeout on page 80
- Configuring system banners on page 80
- Using Wake-on-LAN packets on page 82
- Configuring automation stitches on page 82
- Using the alias commands on page 88

Administrators

You can use the default “admin” account to configure administrator accounts, adjust system settings, upgrade firmware, create backup files, and configure security features.

This section covers the following topics:

- Setting the administrator password on page 67
- Setting the password retries and lockout time on page 69
- Using PKI on page 69
- Adding administrators on page 70
- Administrators on page 66
- Restricting logins from local administrator accounts when remote servers are available on page 74

Setting the administrator password

By default, your system has an administrator account set up with the user name `admin` and no password. On your first login to the GUI or CLI of a new FortiSwitch unit, you must create an admin password. You are also forced to create an admin password after resetting the FortiSwitch configuration to the factory default settings with the `execute factory reset` or `execute factoryresetfull` command.

To set the admin password in the GUI:

1. From the admin menu in the page banner, select *Change Password*.

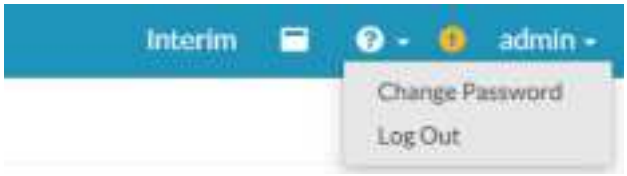

2. Enter the new password in the *Password* and *Confirm Password* fields. Passwords can be up to 64 characters in length.
3. Select *Change*.

Specifying the hash algorithm

Starting in FortiSwitchOS 7.4.0, you can use the CLI to specify which hash algorithm is used to encode passwords for new administrator accounts and updated passwords. You can select the PBKDF2 (with a lower or higher iteration count), SHA1, or SHA256 hash algorithm. By default, the SHA256 hash algorithm is used. Starting in FortiSwitchOS 7.4.3, you can use the GUI to specify the hash algorithm.

Using the GUI:

1. Go to *System* > *Admin* > *Settings*.

2. From the *Hash Algorithm* dropdown list, select which hash algorithm is used to encode passwords for new administrator accounts.
3. To save your changes, click *Update*.

Using the CLI:

```
config system global
   set admin-password-hash {pbkdf2 | pbkdf2-high | sha1 | sha256}
end
```

Variable	Description
pbkdf2	PBKDF2 hash algorithm with a lower iteration count.
pbkdf2-high	PBKDF2 hash algorithm with a higher iteration count.
sha1	SHA1 hash algorithm.
sha256	SHA256 hash algorithm.

Downgrading your FortiSwitchOS version

Before downgrading to a FortiSwitchOS version earlier than 7.0.0, you need to ensure that the administrator password is in SHA1 format. Use the `execute system admin account-convert-sha1` command to convert the administrator password to SHA1 encryption.

Before downgrading to FortiSwitchOS 7.0.0 or later, you need to ensure that the administrator password is in SHA1 or SHA256 format.

- Use the `execute system admin account-convert-sha1` command to convert the administrator password to SHA1 encryption.
- Use the `execute system admin account-convert-sha256` command to convert the password for a system administrator account to SHA256 encryption.

If you do not convert the admin password before downgrading, the admin password will not work after the switch reboots with the earlier FortiSwitchOS version.

To convert the format of the admin password to SHA1 format:

1. Enter the following CLI command to convert the admin password to SHA1 encryption:

   ```
   execute system admin account-convert-sha1 <admin_name>
   ```

2. Downgrade your firmware.

To convert the format of the admin password to SHA256 format:

1. Enter the following CLI command to convert the admin password to SHA256 encryption:

   ```
   execute system admin account-convert-sha256 <admin_name>
   ```

2. Downgrade your firmware.

Setting the password retries and lockout time

By default, the system includes a set number of three password retries, allowing the administrator a maximum of three attempts to log into their account before they are locked out for a set amount of time (by default, 60 seconds).

The number of attempts can be set to an alternate value, as well as the default wait time before the administrator can try to enter a password again. You can also change this value to make it more difficult to hack. Both settings are must be configured with the CLI

To configure the lockout options:

```
config system global
   set admin-lockout-threshold <failed_attempts>
   set admin-lockout-duration <seconds>
end
```

For example, to set the lockout threshold to one attempt and the duration before the administrator can try again to log in to five minutes, enter these commands:

```
config system global
   set admin-lockout-threshold 1
   set admin-lockout-duration 300
end
```

Using PKI

You can use Public Key Infrastructure (PKI) to require administrators to provide a valid certificate when logging in with HTTPS.

Use the following steps to configure PKI:

1. Configure a peer user.
2. Add the peer user to a user group.
3. Configure the administrator account.
4. Configure the global settings.

To configure a peer user:

```
config user peer
   edit <peer_name>
      set ca <name_of_certificate_authority>
   next
end
```

For example:

```
config user peer
   edit pki_peer_1
      set ca Fortinet_CA
   next
end
```

To add the peer user to a user group:

```
config user group
   edit <group_name>
      set member <peer_name>
   next
end
```

For example:

```
config user group
   edit pki_group_1
      set member pki_peer_1
   next
end
```

To configure the administrator account:

```
config system admin
   edit <admin_name>
      set peer-auth enable
      set peer-group <group_name>
   next
end
```

For example:

```
config system admin
   edit pki_admin_1
      set peer-auth enable
      set peer-group pki_group_1
   next
end
```

To configure the global settings:

```
config system global
   set clt-cert-req enable
end
config system web
   set https-pki-required enable
end
```

Adding administrators

Only the default “admin” account can create a new administrator account. If required, you can add an additional account with read-write access control to add new administrator accounts.

If you log in with an administrator account that does not have the super_admin admin profile, the administrators list will show only the administrators for the current virtual domain.

When adding administrators, you are setting up the administrator’s user account. An administrator account comprises an administrator’s basic settings as well as their access profile. The access profile is a definition of what the administrator is capable of viewing and editing.

You can add regular or remote administrators in both the CLI and GUI.

Adding a regular administrator

Using the GUI:

1. Go to *System* > *Admin* > *Administrators*.
2. Click *Add Administrator*.

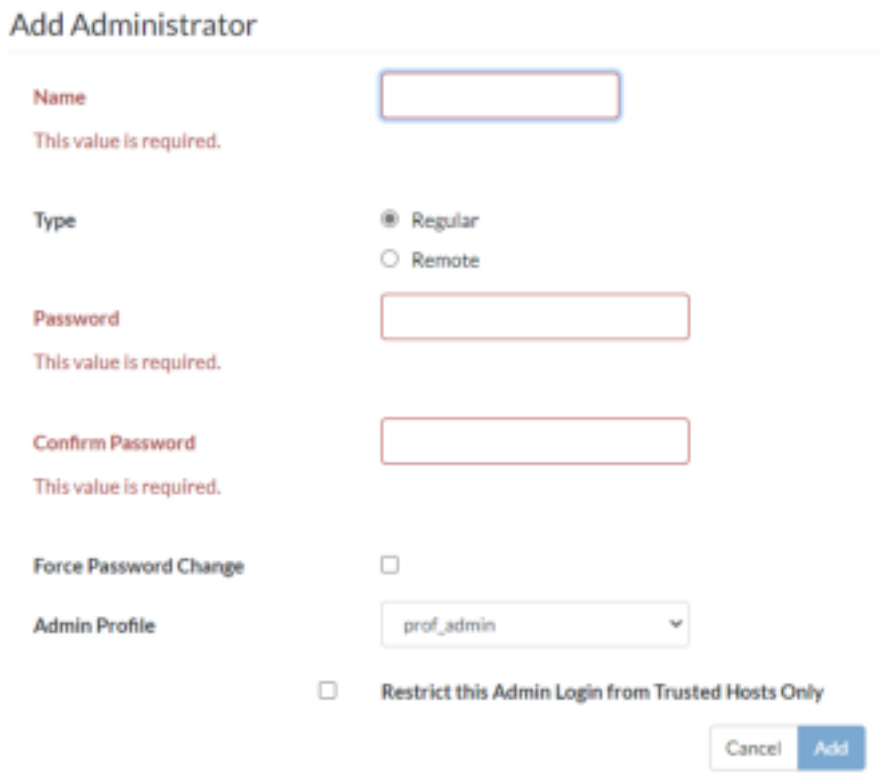

3. In the Name field, enter the administrator name.
4. Click *Regular* for the type of administrator.
5. In the *Password* field and the *Confirm Password* field, enter the password for the administrator.
   Passwords can be up to 64 characters in length.
6. Select the *Force Password Change* checkbox to force the administrator to change the password when next logging in.
7. Select an admin profile from the *Admin Profile* dropdown list.
   Admin profiles control administrator access to FortiSwitch features. To create an admin profile, see Profiles on page 75.
8. Select the *Restrict this Admin Login from Trusted Hosts Only* checkbox if you want to specify which hosts that the administrator can use to connect to the system.

   You can specify up to 10 IPv4 addresses or subnet addresses and netmasks in the *Trusted Host* fields. If you want the administrator to be able to access the system from any address, set the trusted hosts to 0.0.0.0 and the netmask to 0.0.0.0.
9. Click *Add*.

Using the CLI:

```
config system admin
```

```
   edit <admin_name>
      set remote-auth disable
      set password <admin_password>
      set force-password-change{enable | disable}
      set accprofile <profile-name>
      set {trusthost1 | trusthost2 | trusthost3 | trusthost4 |
      trusthost5 | trusthost6 | trusthost7 | trusthost8 | trusthost9
      | trusthost10} <address_ipv4mask>
   next
end
```

Adding a remote administrator

You can configure the RADIUS server to set the access profile. This process uses RADIUS vendor-specific attributes (VSAs) passed to the FortiSwitch unit for authorization. The RADIUS access profile override is mainly used for administrative logins.

Starting in FortiSwitchOS 7.4.0, you can add multiple remote administrators with wildcards in their names. When there is a wildcard in the administrator name, FortiSwitchOS examines possible matches based on their order in the Configuration Management Database (CMDB) and allows the first match to log in.

Starting in FortiSwitchOS 7.4.0, you can enable `wildcard-fallback` when `set remote-auth` is enabled and `set wildcard` is disabled. When `wildcard-fallback` is enabled, FortiSwitchOS first tries to match the login name; if the login name matches the system administrator's credentials, the login is successful. If FortiSwitchOS cannot match the login name with the system administrator's credentials, it will try to match the login name with wildcard system administrator names based on their order in the CMDB and allows the first match to log in.

Using the GUI:

1. Go to *System* > *Admin* > *Administrators*.
2. Click *Add Administrator*.

3. Click *Remote* for the type of administrator. .

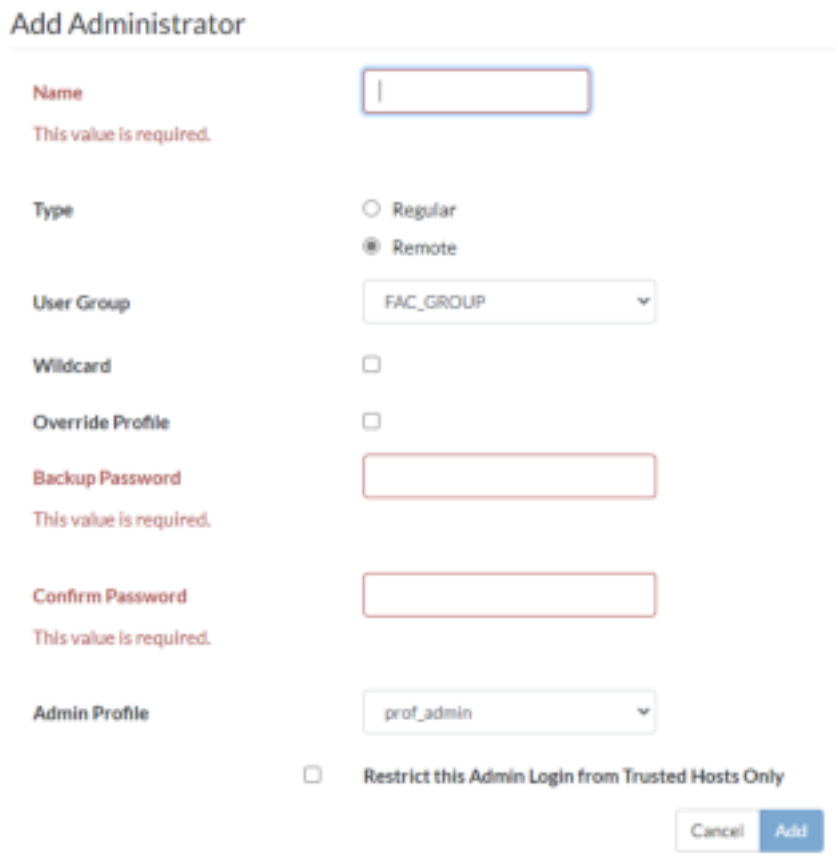

4. In the *Name* field, enter a name for the RADIUS system administrator.
5. Select the user group from the *User Group* dropdown list.
   To add a user group to the list, see User groups on page 98.
6. Select the *Wildcard* checkbox if you want to use wildcard characters in the name of the RADIUS system administrator.
7. Click the *Accprofile Override* checkbox if you want the remote authentication server to be able to override the access profile.
8. In the *Backup Password* field and the *Confirm Password* field, enter the password for the administrator.
   Passwords can be up to 64 characters in length.
9. Select an admin profile from the *Admin Profile* dropdown list.
   Admin profiles control administrator access to FortiSwitch features. To create an admin profile, see Profiles on page 75.
10. Select the *Restrict this Admin Login from Trusted Hosts Only* checkbox if you want to specify which hosts that the administrator can use to connect to the system.

    You can specify up to 10 IPv4 addresses or subnet addresses and netmasks in the *Trusted Host* fields. If you want the administrator to be able to access the system from any address, set the trusted hosts to 0.0.0.0 and the netmask to 0.0.0.0.
11. Click *Add*.

Using the CLI:

```
config system admin
   edit <admin_name>
      set remote-auth enable
      set remote-group <name>
      set wildcard {enable | disable}
      set wildcard-fallback {enable | disable}
      set accprofile-override {enable | disable}
      set password <admin_password>
      set accprofile <profile-name>
      set {trusthost1 | trusthost2 | trusthost3 | trusthost4 |
      trusthost5 | trusthost6 | trusthost7 | trusthost8 | trusthost9
      | trusthost10} <address_ipv4mask>
   next
end
```

For example, the following commands create a RADIUS-system admin group with accprofile-override enabled:

```
   config system admin
      edit "RADIUS_Admins"
         set remote-auth enable
         set accprofile no_access
         set wildcard enable
         set remote-group "RADIUS_Admins"
         set accprofile-override enable
      next
   end
```

Ensure that the RADIUS server is configured to send the appropriate VSA.

To send an appropriate group membership and access profile, set VSA 1 and VSA 6, as in the following example:

```
VENDOR fortinet 12356
ATTRIBUTE Fortinet-Group-Name 1 <admin profile>
ATTRIBUTE Fortinet-Access-Profile 6 <access profile>
```

The value of VSA 1 must match the remote group, and VSA 6 must match a valid access profile.

Restricting logins from local administrator accounts when remote servers are available

You can restrict logins from local administrator accounts when remote servers (such as TACACS+, LDAP, or RADIUS) are available. When the CLI command is enabled, FortiSwitchOS checks if all of the remote servers used by administrators are down before allowing a local administrator to log in. This option is applied globally; it is disabled by default.

To restrict local administrator authentication when a remote authentication server available:

```
config system global
   set admin-restrict-local enable
end
```

Profiles

In addition to the default “admin” account, you might want to set up other administrators with different levels of system access.

Administer profiles define what the administrator user can do when logged into the FortiSwitch unit. When you set up an administrator user account, you also assign an administrator profile, which dictates what the administrator user will see. 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 required.

The super_admin administrator is the administrative account that the primary administrator should have to log into the FortiSwitch unit. The profile cannot be deleted or modified to ensure there is always a method to administer the FortiSwitch unit. This user profile has access to all components of the system, including the ability to add and remove other system administrators. For some administrative functions, such as backing up and restoring the configuration using SCP, super_admin access is required.

To configure administrator profiles, go to *System* > *Admin* > *Profiles*. You can only assign one profile to each administrator user.

On the *Add Profile* page, you define the components of the FortiSwitch unit that will be available to view and/or edit. For example, if you configure a profile so that the administrator can only access system configuration, this admin will not be able to change the network configuration. For more detail about what is covered by each access control, see Access control on page 77.

Using the GUI:

1. Go to *System* > *Admin* > *Profiles* and select *Add Profile*.

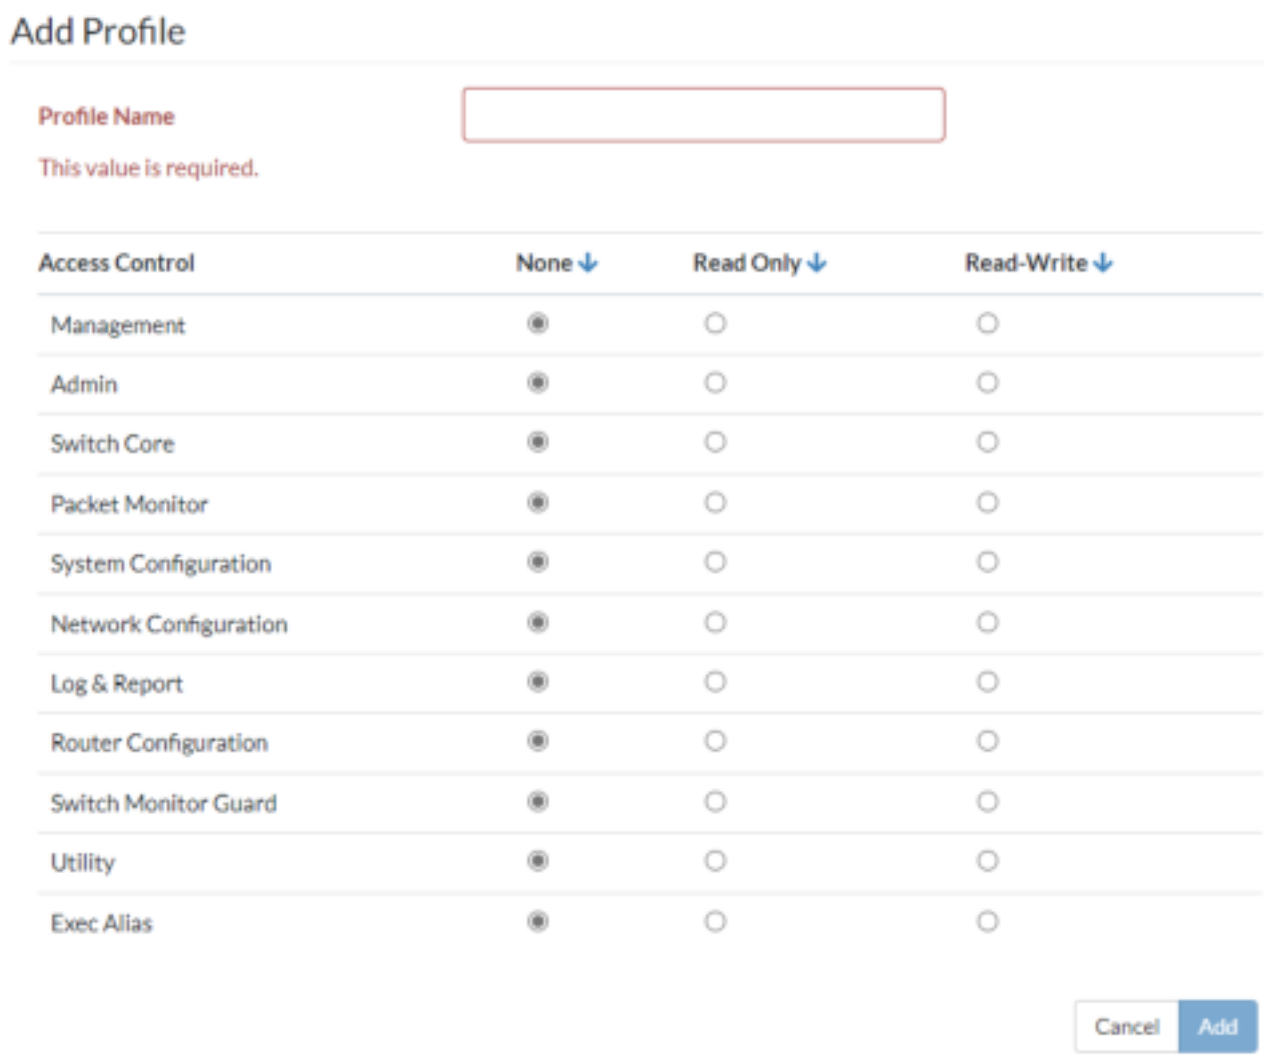

2. Give the profile an appropriate name.
3. Set *Access Control* as required, selecting *None*, *Read Only*, or *Read-Write* for each line.
4. Select *Add*.

Using the CLI:

```
config system accprofile
   edit <name>
      set admingrp {none | read | read-write}
      set alias-commands {all | group} <alias-command>
      set exec-alias-grp {none | read | read-write}
      set loggrp {none | read | read-write}
      set mntgrp {none | read | read-write}
      set netgrp {none | read | read-write}
      set pktmongrp {none | read | read-write}
      set routegrp {none | read | read-write}
      set swcoregrp {none | read | read-write}
      set swmonguardgrp {none | read | read-write}
      set sysgrp {none | read | read-write}
      set utilgrp {none | read | read-write}
```

```
    next
end
```

Access control

The *Management* access control applies to following menus:

- *System > Dashboard*
- *System > Config > Firmware*
- *System > Revisions*
- *System > Licenses*

The *Admin* access control applies to the following menus:

- *System > Dashboard*
- *System > Admin > Administrators*
- *System > Admin > Profiles*
- *System > Admin > Monitor*
- *System > Authentication > LDAP*
- *System > Authentication > RADIUS*
- *System > Authentication > TACACS*

The *Switch Core* access control applies to the following menus:

- *System > Dashboard*
- *Switch > Physical Ports*
- *Switch > Trunks*
- *Switch > Interfaces*
- *Switch > Port Security*
- *Switch > STP > Settings*
- *Switch > STP > Instances*
- *Switch > Flap Guard*
- *Switch > MAC Limit*
- *Switch > ACL > Ingress*
- *Switch > ACL > Egress*
- *Switch > ACL > Prelookup*
- *Switch > ACL > Policer*
- *Switch > ACL > Service*
- *Switch > POE*
- *Switch > VLAN*
- *Switch > Virtual Wires*
- *Switch > MAC Entries*
- *Switch > QoS > 802.1p*
- *Switch > QoS > IP/DSCP*
- *Switch > QoS > Egress Policy*
- *Switch > Monitor > ACL Counters*

- *Switch > Monitor > Forwarding Table*
- *Switch > Monitor > Trunks*
- *Switch > Monitor > Port Stats*
- *Switch > Monitor > Spanning Tree*
- *Switch > Monitor > Modules*
- *Switch > Monitor > 802.1x Status*
- *Switch > Monitor > IGMP Snooping*

The *Packet Monitor* access control applies to the following menus:

- *System > Dashboard*
- *System > Flow Export > Configure*
- *System > Flow Export > Monitor*
- *System > Packet Capture*
- *Switch > sFlow*
- *Switch > Mirror*

The *System Configuration* access control applies to the following menus:

- *System > Dashboard*
- *System > Network > Settings*
- *System > Config > Backup*
- *System > Config > Hostname*
- *System > Config > Time*
- *System > Config > SSL*
- *System > Admin > Settings*
- *System > User > Definition*
- *System > User > Group*
- *System > Certificate > Local*
- *System > Certificate > Remote*
- *System > Certificate > Authorities*
- *System > Certificate > CRLs*
- *System > Link Monitor*
- *System > FortiLAN Cloud*
- *System > Locations*
- *Switch > DHCP Snooping*
- *Router > Config > Link Probes*

The *Network Configuration* access control applies to the follow menus:

- *System > Dashboard*
- *System > Network > Interface > Physical*
- *System > Network > Interface > VLAN*
- *System > Network > Interface > Loopback*
- *System > Network > DNS*
- *System > DHCP*

- *Router > Config > Interface*
- *Router > ARP Table*

The *Log & Report* access control applies to the follow menus:

- *System > Dashboard*
- *System > Config > SNMP > Communities*
- *System > Config > SNMP > Users*
- *System > Config > SNMP > Settings*
- *Log > Entries*
- *Log > Config*

The *Router Configuration* access control applies to the following menus:

- *System > Dashboard*
- *Router > Config > OSPF > Settings*
- *Router > Config > OSPF > Areas*
- *Router > Config > OSPF > Networks*
- *Router > Config > OSPF > Interfaces*
- *Router > Config > RIP > Settings*
- *Router > Config > RIP > Distances*
- *Router > Config > RIP > Networks*
- *Router > Config > RIP > Interfaces*
- *Router > Config > Multicast > Settings*
- *Router > Config > Multicast > Interfaces*
- *Router > Config > Multicast > Flows*
- *Router > Config > Access Lists*
- *Router > Config > Static*
- *Router > Config > IPv6 Static*
- *Router > Config > VRF*
- *Router > Monitor > Routing*
- *Router > Monitor > IPv6 Routing*
- *Router > Monitor > BFD Neighbor*
- *Router > Monitor > VRRP*

The *Switch Monitor Guard* access control applies to the following menus:

- *System > Dashboard*
- *Switch > LLDP-MED > Profiles*
- *Switch > LLDP-MED > Settings*
- *Switch > Storm Control*
- *Switch > IP-MAC Binding*
- *Switch > Monitor > DHCP Snooping > Clients*
- *Switch > Monitor > DHCP Snooping > Servers*
- *Switch > Monitor > LLDP*
- *Switch > Monitor > Loop Guard*

- *Switch > Monitor > Flap Guard*
- *Switch > Monitor > BPDU Guard*

The *Utility* access control applies to the following menus:

- GUI CLI
- *System > Dashboard*
- *System > Debug Report*
- *Router > Diagnostic*

The *Exec Alias* access control applies to the following menus:

- *System > Dashboard*

Monitor

You can find out which administrators are logged in by looking at the *System Information* section of the *Dashboard*. The *Admin Sessions* row shows how many administrators are logged in.

Clicking on the number of sessions goes to *Admin > Monitor* for details about which admins are logged in, when they logged in, the connection type, and the IP address.

You can select an administrator and click *Disconnect* to log out the administrator.

Setting the idle timeout

By default, the GUI disconnects administrative sessions if no activity occurs for five minutes. This prevents someone from using the GUI if the management PC is left unattended.

To change the idle timeout:

1. Go to *System > Admin > Settings*.
2. Enter the time in minutes in the *Idle Timeout (Minutes)* field.
3. Update other settings as required:
   - TCP/UDP port values for HTTP, HTTPS, Telnet, SSH
   - Display language
4. Select *Update*.

Configuring system banners

You can specify system banner messages in the CLI that will appear when users log in using either the CLI or the GUI.

You can enter up to 2,048 characters for each system banner. Currently, only text is supported. By default, no system banners are displayed.

The GUI displays the pre-login banner before you enter your user name or password:

The GUI displays the post-login banner after you enter your user name and password and select *Log In*:

You cannot finish logging in until you select *I Agree*.

The CLI displays the pre-login banner before you enter your user name. The CLI displays the post-login banner after you enter your password; you cannot finish logging in until you press `a` to accept the message.

To configure system banners:

```
config system global
   set pre-login-banner "<string>"
   set post-login-banner "<string>"
end
```

For example:

```
S548DF5018000776 # config system global
S548DF5018000776 (global) # set pre-login-banner "All systems will be unavailable,
> starting at midnight. Please exit all applications by then."
S548DF5018000776 (global) # set post-login-banner "Remember to exit before midnight."
S548DF5018000776 (global) # end
```

NOTE: For multi-line messages, just press the Return key between lines.

Using Wake-on-LAN packets

You can use the CLI to send Wake-on-LAN (WoL) packets to a specific MAC address to remotely turn on a computer. You can either use a system interface or a switch port to send the WoL packets. The WoL packets are sent by UDP by default, but you can use the WoL protocol to send the WoL packets instead. You can also specify a password if the remote computer is protected by SecureOn.

To send WoL packets:

```
execute wake-on-lan <interface_type> <interface_or_port> <host_MAC_address> <protocol>
     <port> <IP_address> <password>
```

Variable	Description
<interface_type>	Select the interface type that will send the WoL packets. Select 1 if to use the system interface or 2 to use the switch port. The default is 1.
<interface_or_port>	If you selected 1 for the interface type, specify which system interface to use (required). If you selected 2 for the interface type, specify which switch port to use (optional).
<host_MAC_address>	Required. Enter the MAC address (XX:XX:XX:XX:XX:XX) of the computer that needs to be turned on.
<protocol>	Optional. Select which protocol to use to send the WoL packets. Select 1 for WoL or 2 for UDP. The default is 2.
<port>	Optional. If you selected 2 for the protocol, select which port the WoL packets will use. You can select 0, 7, or 9. The default is 9.
<IP_address>	Optional. If you selected 2 for the protocol, enter the broadcast IP address used by the WoL packets.
<password>	Optional. Enter the password if a 6-byte SecureOn password is enabled on the destination host. The password can be a string or 0x plus a hexadecimal value.

If you are sending the WoL packets by UDP from the FortiSwitch port3 to a MAC address of aa:bb:cc:00:11:22:

```
execute wake-on-lan 2 port3 aa:bb:cc:00:11:22 2 9 1.2.3.4
```

If you are sending the WoL packets by UDP from the FortiSwitch port10 to a MAC address of 10:20:30:40:50:60 and the destination host is protected by a SecureOn password:

```
execute wake-on-lan 2 port10 10:20:30:40:50:60 2 9 10.10.10.10 passwd
```

Configuring automation stitches

To configure an automation stitch, you specify a trigger and the action that is performed when the trigger occurs and then set the status to `enable`.

You can specify one of the following triggers:

- The configuration changed.
- There was a warm or cold reboot of the switch.
- The scheduled time occurred.
- An event was logged.

  NOTE: When you specify the log ID, the range of values is 1-65535. If you use the full 10-digit entry, the first four digits are truncated.

You can specify one of the following actions:

- Run a CLI script.
- Send an email message.
- Display an alert in the console.
- Generate an SNMP trap.
- Send data to a uniform resource identifier (URI), such as an IP address or URL.

Starting in FortiSwitchOS 7.2.2, you can use the following wildcard characters in the `set value` command for the automation trigger:

- Use an asterisk to match any character string of any length, including 0-characters long. For example, use `set value "*1567*"` to match values of `81567` and `156789`.
- Use square brackets to match one of the multiple characters. For example, use `set value "[aA]dmin"` to match values of `admin` and `Admin`.

Starting in FortiSwitchOS 7.2.2, you can configure multiple fields for the automation trigger when the `event-type` is `event-log` and the `logid` is set. The action is only performed if all conditions are valid (using AND logic). For example, the following automation trigger requires both the log message to include `VRRP` and the interface to be `svi777` before the action is performed.

```
config system automation-trigger
   edit "VRRPlogtrigger"
      set event-type event-log
      set logid 102003209
      config fields
         edit 1
            set name "msg"
            set value "*VRRP*"
         next
         edit 2
            set name "interface"
            set value "svi777"
         next
      end
   next
end
```

Starting in FortiSwitchOS 7.4.3, you can use an automation stitch to shut down a port when the storm-control dropped-packet rate is too high and bring up the port when the dropped-packet rate is below the specified threshold.

Starting in FortiSwitchOS 7.4.3, you can use five custom SNMP traps (`fsStitchTrap1`, `fsStitchTrap2`, `fsStitchTrap3`, `fsStitchTrap4`, and `fsStitchTrap5`) for automation actions.

To configure an automation stitch:

1.
```
config system automation-trigger
   edit <trigger_name>
      set trigger-type {event-based | scheduled}
      set event-type {config-change | event-log | reboot}
      set logid <log_ID>
      set trigger-frequency {daily | hourly | monthly | weekly}
      set trigger-hour <0-23>
      set trigger-minute <0-59>
      set trigger-day <1-31>
      set trigger-weekday <friday | monday | saturday | sunday | thursday | tuesday |
           wednesday>
      config fields
         edit <entry_ID>
            set name <string>
            set value <string>
         next
      end
   next
end
```
2. Create an automation action.
```
config system automation-action
   edit <name>
      set action-type {alert | cli-script | email | snmp-trap | webhook}
      set accprofile <string>
      set email-body <string>
      set email-from <string>
      set email-subject <string>
      set email-to <email_address>
      set headers <request_headers>
      set http-body <request_body>
      set method {delete | get | patch | post | put}
      set minimum-interval <0-2592000>
      set port <1-65535>
      set protocol {http | https}
      set script <string>
      set snmp-trap {fsStitchTrap1 | fsStitchTrap2 | fsStitchTrap3 | fsStitchTrap4 |
           fsStitchTrap5}
      set uri <request_API_URI>
   next
end
```
3. Create the automation stitch.
```
config system automation-stitch
   edit <name>
      set status {enable | disable}
      set trigger <trigger_name>
      set action <action_name>
   next
end
```
4. Test the automation stitch.
```
diagnose automation test <automation-stitch-name> [<log_ID>]
```

Example 1

The following example shows how to create an automation stitch that will display an alert in the console every hour.

```
config system automation-trigger
   edit testtrigger
      set trigger-type scheduled
      set trigger-frequency hourly
      set trigger-minute 30
   next
end

config system automation-action
   edit testaction
      set action-type alert
      set minimum-interval 1200
   next
end

config system automation-stitch
   edit teststitch
      set status enable
      set trigger testtrigger
      set action testaction
   next
end

diagnose automation test teststitch 0
```

Example 2

In the following example, the specified log identifier (32002) causes the FortiSwitch unit to send the log message to the server.

```
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

config system automation-trigger
   edit "badLogin"
      set event-type event-log
      set logid 32002
   next
end

config system automation-stitch
   edit "webhookstitch"
      set trigger "badLogin"
      set action "Send log to server"
   next
end
```

Example 3

The log message with ID 103042599 is generated when the storm-control dropped-packet rate is too high. The log message with ID 103042600 is generated when the storm-control dropped-packet rate goes below the threshold.

In the following example, log ID 103042599 triggers a CLI script that shuts down the affected port, and log ID 103042600 triggers another CLI script that waits 5 miniutes and then brings up the affected port.

```
config system automation-trigger
   edit "PortStormControlDrop"
      set event-type event-log
      set logid 103042599
   next
   edit "PortStormControlClear"
      set event-type event-log
      set logid 103042600
   next
end

config system automation-action
   edit "ShutdownPort"
      set action-type cli-script
      set script "config switch physical-port
         edit %%log.switch.physical-port%%
            set status down
         next
      end"
      set accprofile "super_admin"
   next
   edit "BringupPort"
      set action-type cli-script
      set script "sleep 300
         config switch physical-port
            edit %%log.switch.physical-port%%
               set status up
            next
         end"
      set accprofile "super_admin"
   next
end

config system automation-stitch
   edit "DisablePortOnStormControlDrop"
      set trigger "PortStormControlDrop"
      set action "ShutdownPort"
   next
   edit "EnablePortOnStormControlClear"
      set trigger "PortStormControlClear"
      set action "BringupPort"
   next
end
```

Example 4

The following example shows how to use a custom automation action.

1. Create an automation trigger.

   There are two triggers in this example, one trigger for the switch restarting and one trigger for a specific log message.

```
config system automation-trigger
   edit "trigger-reboot"
      set event-type reboot
   next
   edit "trigger-log"
      set event-type event-log
      set logid 103032003
   next
end
```

2. Create the automation action.

   There are two actions in this example, one action to send the custom trap fsStitchTrap1 and one action to send the custom trap fsStitchTrap2.

```
config system automation-action
   edit "action-trap1"
      set action-type snmp-trap
      set snmp-trap fsStitchTrap1
   next
   edit "action-trap2"
      set action-type snmp-trap
      set snmp-trap fsStitchTrap2
   next
end
```

3. Create the automation stitch.

   For the first automation stitch, when the switch restarts, custom trap fsStitchTrap1 is sent. For the second automation stitch, when a specific log message is generated, custom trap fsStitchTrap2 is sent.

```
config system automation-stitch
   edit "stitch1"
      set trigger "trigger-reboot"
      set action "action-trap1"
   next
   edit "stitch2"
      set trigger "trigger-log"
      set action "action-trap2"
   next
end
```

4. Configure the SNMP user if your SNMP trap receiver is expecting SNMP v3 traps. By default, all SNMP notifications are enabled, except for `l2mac`.

   In this example, the SNMP v3 notifications are sent to 1.2.3.4. If your SNMP trap receiver is expecting SNMP v1 or v2c traps, configure the SNMP community with the `config system snmp community` command.

```
config system snmp user
   edit "1"
      set notify-hosts 1.2.3.4
      set events fsStitchTrap1 fsStitchTrap2
   next
end
```

Using the alias commands

Previously, you could use the *Add Profile* page or the `config system accprofile` command to control the view access, edit access, and no access to groups of menu commands for an administrator account. Starting in FortiSwitchOS 7.0.0, you can use the alias CLI commands to grant an administrator access to individual configuration attributes or CLI commands, instead of having to grant access to large groups of CLI commands and configuration attributes.

Notes:

- Configuration-type aliases cannot create or delete table entries. For example, under the `config switch interface` command, you cannot create a new interface name with the `edit <interface_name>` command.
- Configuration-type aliases cannot act on child tables or child objects. For example, the configuration-type alias for `config system interface` can affect attributes (`set` commands) under `config system interface` but not `config ipv6` under `config system interface` or attributes (`set` commands) under both `config ipv6` and `config system interface`.
- The super_admin administrator profile has access to all command aliases.
- You can use the `sleep <1-172800 seconds>` command to add a delay in a script.

Procedure

1. Do one of the following:
   - Specify a configuration-type alias for each configuration attribute you want to control access to.
   - Specify a script-type alias for CLI commands or groups of CLI commands that you want to control access to.
   - Specify alias groups to bundle different alias commands together for easy assignment.
2. Create an access profile that uses the aliases and alias groups that you created.
3. Create an administrator account and assign the access profile that you created.
4. Run the alias command or script.

Step 1: Create a configuration-type alias or script-type alias

To specify a configuration-type alias for a command that you want to control access to:

```
config system alias command
   edit <alias_name>
      set description <string>
      set type configuration
      set path <path>
      set attribute <attribute-name>
      set table-listing {allow | deny}
      set permission {read | read-write}
      set limit-shown-attributes {disable | enable}
      set read-only-attributes <attribute-name>
      set table-ids-allowed <table-ID-value>
   end
```

Variable	Description	Default
<alias_name>	Enter an alias name for the command in this configuration.	No default

Variable	Description	Default
	The alias name cannot be `all` or match an alias group name.	
description <string>	Enter a description of the command or a help message. It can be up to 80-characters long. The description is displayed with the alias name when you enter `execute alias configure {get | show | show-full-configuration | set | unset} ?`.	No default
type configuration	The `configuration` type provides configuration-specific functionality to control `get`, `show`, `show-full-configuration`, `set`, and `unset` commands. You can also use the `configuration` type to limit accessible table entries and limit displayed attributes.	configuration
path <path>	Required. Enter the period-separated path to the CLI command. For example, enter `set path switch.lldp.profile` to apply the configuration to the `config switch lldp profile` command. Enter `set path system.interface` to apply the configuration to the `config system interface` command. You can specify only top-level objects, such as `system.interface`, `router.bgp`, or `system.snmp.settings`. If you specify child objects or child tables (such as `system.interface.ipv6`, `router.bgp.neighbor`, or `switch.lldp.profile.custom-tlv`), FortiSwitch returns an error.	No default
attribute <attibute-name>	Required. Enter the attribute that can be retrieved or modified. Enter `set attribute ?` to see the list of valid attributes. If you enter an invalid value, FortiSwitchOS returns an error. This option is available only when `path` has been set.	No default
permission {read | read-write}	Select `read` to allow this alias to be used by the `execute alias configure {get | show | show-full-configuration}` command. Select `read-write` to allow this alias to be used by the `execute alias configure {get | show | show-full-configuration | set | unset}` command.	read
table-listing {allow | deny}	Allow or prevent the listing of all entries by the `execute alias configure {get | show | show-full-configuration}` command commands. • Select `allow` to permit all entries to be listed. • Select `deny` to prevent the entries from being listed except for the entries specified in the `table-ids-allowed` setting. If `table-ids-allowed` is empty, a valid entry must be provided for listing.	deny

Variable	Description	Default
	This option is available only when `path` has been set.	
limit-shown-attributes {disable | enable}	Enable or disable whether to limit the attributes displayed with the `show` and `get` commands. Selecting `disable` displays all attributes for the show and get commands. Selecting `enable` displays only the attributes listed in `attributes` and `read-only-attributes`.	enable
read-only-attributes <attribute-name>	When `limit-shown-attributes` is enabled, you can enter additional attributes to display with the `show` and `get` commands. When you enter `read-only-attributes ?` to see a list of valid attributes, more attributes are available than when you enter `set attribute ?`. Read-only attributes can include child tables, child objects, and get-only attributes. You can list up to 31 attributes.	No default
table-ids-allowed <table-ID-value>	Specify which entries can be accepted by the `execute alias configure {get | show | show-full-configuration | set | unset}` command. Enter `set table-ids-allowed ?` to see a list of valid entries. You can specify entries that do not currently exist; they can be created later. If `table-listing` is set to `deny`, the `table-ids-allowed` entries are displayed when the user runs the `execute alias configure {get | show | show-full-configuration}` command without specifying any entry. This option is available only when `path` has been set.	No default

The following example creates two aliases for the `config switch physical-port` command.

- The `port-description` alias allows an administrator to change the `set description` value; when running a `get` or `show` command, the administrator will see only the description configuration.
- The `port-status` alias allows an administrator to change the `set status` value; the administrator will see both the description and port status configuration when running `get` or `show` commands.

```
config system alias command
   edit "port-status"
      set description "View or change the port status."
      set type configuration
      set path "switch.physical-port"
      set attribute "status"
      set permission read-write
      set limit-shown-attributes enable
      set read-only-attributes "description"
   next
   edit "port-description"
      set description "View or change the port description."
      set type configuration
      set path "switch.physical-port"
      set attribute "description"
      set permission read-write
```

```
      set limit-shown-attributes enable
   next
end
```

To create a script:

```
config system alias command
   edit <script_name>
      set description <string>
      set type script
      set command <string>
      set table-entry-create {allow | deny}
      config script-arguments
         edit <argument_ID>
            set type {integer | string | table-id}
            set name <string>
            set help <string>
            set optional {enable | disable}
            set range {enable | disable}
            set range-delay <0-172800>
            set allowed-values <string>
         next
      end
   next
end
```

Variable	Description	Default
<script_name>	Enter a script name. The script name cannot be `all` or match an alias group name.	No default
description <string>	Enter a description of the script. It can be up to 80-characters long. The description is displayed with the script name when you enter `execute alias script ?`.	No default
type script	The `script` type allows the administrator to create a list of CLI commands to run.	configuration
command <string>	Enter the script command (within quotation marks) to be run. You can use the Enter key to separate command lines. Enter `set command ?` for formatting details. This option is available only when `type` has been set to `script`.	No default
table-entry-create {allow | deny}	Allow or deny the creation of new table (or sub-table) entries. This option is available only when `type` has been set to `script`. When `type` has been set to `configuration`, you cannot create any new table entries.	deny
config script-arguments		
<argument_ID>	Enter an identifier for the argument. The identifier must match the identifier used in the script.	No default

Variable	Description	Default
type {integer | string | table-id}	Enter the data type that the argument accepts.	string
name <string>	Enter the display name for the argument. You can use uppercase and lowercase letters, numbers, and hyphens. The display name is shown when the user runs the `execute alias script` command.	No default
help <string>	Enter a help message for the argument. You can use uppercase and lowercase letters, numbers, slashes, parentheses, brackets, commas, underscores, and hyphens. The help message is displayed when the user runs the `execute alias script` command.	No default
optional {enable | disable}	Enable this option to allow the user to omit entering a value for this argument. Disable this option to force the user to specify a value for this argument.	disable
range {enable | disable}	Enable this option to allow a range of integers, a range of table identifiers, or a comma-separated list of strings. Disable this option to allow only a single value for this argument.	disable
range-delay <0-172800>	Enter the number of seconds to delay between values when executing. This option is available only when `range` has been set to `enable`.	0
allowed-values <string>	Enter the values allowed for this argument. • If `type` is set to `string`, separate values with a space. For example: `set allowed-values port1 port3 port7` • If `type` is set to `integer`, you can use ranges and comma-separated values, such as “1-10” or “1-10,3,11,55”. • If `type` is set to `table-id` and the table identifiers are integers, you can use both ranges and comma-separated values, such as “1-10” or “1-10,3,11,55”.	No default

The following example creates two scripts. Both scripts list the switch mac-address table.

- The `mac-list` script is more flexible because it requires that the user specify the VLANs to list the MAC addresses from.
- The `list-mac-by-port-and-vlan-customer-AAA` script is more controlled because it allows the user to see the MAC addresses learned on the specified VLANs.

```
config system alias command
   edit "list-mac-by-port-and-vlan-customer-AAA"
      set description "List MAC addresses on your VLANs and ports."
      set type script
      set command "diag switch mac-address filter clear
diag switch mac-address filter port-id-map 3-8
diag switch mac-address filter vlan-map 1000-1010
diag switch mac-address list
```

```
diag switch mac-address filter clear"
   next
   edit "mac-list"
      set description "List MAC addresses learned on the provided VLANs"
      set type script
      set command "diag switch mac-address filter clear
diag switch mac-address filter vlan-map $1
diag switch mac-address list | grep -i mac
diag switch mac-address filter clear"
      config script-arguments
         edit 1
            set name "VLAN-ID-map"
            set help "List of VLANs to check"
         next
      end
   next
end
```

To create a group of configuration-type aliases:

```
config system alias group
   edit <alias_group_name>
      set description <string>
      set commands <alias_command_list>
   end
```

Variable	Description	Default
<alias_group_name>	Enter a name for the alias group. The name cannot be `all` or match an alias name.	No default
description <string>	Enter a description of the command alias group. It can be up to 80-characters long.	No default
commands <alias_command_name>	Enter a list of command aliases. Use a space to separate them.	No default

The following example creates a group of two command aliases.

```
config system alias group
   edit aliasgroup1
      set description "Alias group for config switch physical-port."
      set commands port-status port-description
   end
```

Step 2: Create the access profile

To create an access profile for aliases or alias groups:

```
config system accprofile
   edit <profile_name>
      set alias-commands {all | <list>}
      set exec-alias-grp {none | read | read-write}
   end
```

The following example creates an access profile with read-write access to all the `execute alias` commands for the alias commands from the `aliasgroup1` alias group and for the `list-mac-by-port-and-vlan-customer-AAA` script:

```
config system accprofile
   edit newaccprofile
      set alias-commands list-mac-by-port-and-vlan-customer-AAA
      set exec-alias-grp read-write
   end
```

Step 3: Create the administrator account

Using the GUI:

Go to *System* > *Admin* > *Administrators*, click *Add Administrator*, and select the access profile that you created.

Using the CLI:

```
config system admin
   edit <admin_name>
      set password <password>
      set accprofile <profile_name>
   end
```

For example:

```
config system admin
   edit newadmin
      set password newpassword
      set accprofile newaccprofile
   end
```

Step 4: Run the alias command or script

To run an alias command:

```
execute alias configure set <alias_name> <table-entry-id-if-needed> <attribute-value>
```

Note: If the alias has a multi-value attribute (for example, `set allowaccess` under the `config system interface` command or `set members` under the `config switch trunk` command), you can enter up to 31 values for it. If the alias does not reference a table and no table entry ID is needed, you can enter up to 32 values.

The following example changes the value for the `port2` table entry to `up`.

```
S548DF5018000776 # execute alias configure set port-status port2 up
Command to be run:
--------------------------------------------------------------------------------
config switch physical-port
edit "port2"
set status "up"
next
end
--------------------------------------------------------------------------------
Do you want to continue? (y/n)y
```

To run a script:

```
execute alias script <script_name> <values…>
```

The following example shows how to run the `mac-list` script for VLAN 4092.

```
S524DF4K15000024 # execute alias script mac-list 4092
Command to be run:
-------------------------------------------------------------------------------
diag switch mac-address filter clear
diag switch mac-address filter vlan-map "4092"
diag switch mac-address list | grep -i mac
diag switch mac-address filter clear
-------------------------------------------------------------------------------
Do you want to continue? (y/n)y

MAC: 08:5b:0e:f1:95:e5 VLAN: 4092 Port: internal(port-id 31)
```

User

The FortiSwitch unit provides authentication mechanisms to control user access to the system (based on the user group associated with the user). The members of user groups are user accounts. Local users and peer users are defined on the FortiSwitch unit. User accounts can also be defined on remote authentication servers.

This section describes how to configure local users, peer users and user groups. For information about configuring the authentication servers, see RADIUS on page 99.

This section covers the following topics:

- User definition on page 95
- Peer user on page 96
- User groups on page 98

User definition

A user account consists of a user name, password, and potentially other information, configured in a local user database or on an external authentication server.

Users can access resources that require authentication only if they are members of an allowed user group.

Using the GUI:

1. Go to *System* > *User* > *Definition*.
2. Click *Add User*.
3. Enter the user name.
4. Select *Enable* to make the user account active.
5. Enter the password for the user account. Passwords can be up to 64 characters in length.
6. Click *Add*.

Using the CLI:

```
config user local
   edit <user_name>
      set ldap-server <server_name>
      set passwd <password_string>
      set radius-server <server_name>
      set tacacs+-server <server_name>
      set status {enable | disable}
      set type <auth-type>
   end
```

Field	Description
user_name	Identifies the user
password_string	A password for the local user. Passwords can be up to 64 characters in length.
ldap-server <server_name>	To authenticate this user using a password stored on a remote authentication server, select the type of server and then select the server from the list. You can select only a server that has already been added to the FortiSwitch configuration.
radius-server <server_name>	To authenticate this user using a password stored on a remote authentication server, select the type of server and then select the server from the list. You can select only a server that has already been added to the FortiSwitch configuration.
tacacs+-server <server_name>	To authenticate this user using a password stored on a remote authentication server, select the type of server and then select the server from the list. You can select only a server that has already been added to the FortiSwitch configuration.
status	Enable or disable this user.

Peer user

A peer user is a digital certificate holder that authenticates using a client certificate.

Using the GUI:

1. Go to *System* > *User* > *Peer*.
2. Click *Add Peer*.
3. In the *Name* field, enter the name of the peer user.
4. In the *Subject* field, enter any limitations on the peer certificate name.
5. Select the type of common name for the peer certificate from the *Common Type* dropdown list. You can select *Fully Qualified Domain Name*, *Email*, *IPv4*, *IPv6*, or *String*.
6. In the *Common Name* field, enter the common name for the peer certificate.
7. Select which certificate authority (CA) certificate to use from the *Certificate* dropdown list.
8. Select the *Mandatory Verify* checkbox for mandatory CA verification.

9. Select the *Two-Factor* checkbox for two-factor authentication. When two-factor authentication is selected, the certificate and password are required.
10. If you selected the *Two-Factor* checkbox, enter a password to use in the *Password* field.
11. If you want to use an LDAP server to check access permission:
    a. Select the server name from the Server dropdown list. If no server name is available, go to *System* > *Authentication* > *LDAP* to add an LDAP server.
    b. Select the authentication mode from the *Mode* dropdown list, either *Username and Password* or *Principal Name*.
    c. If you selected *Username and Password*, enter the user name in the *Username* field and enter the password in the *Password* field.
12. Click *Add*.

Using the CLI:

```
config user peer
   edit <peer_name>
      set ca {Entrust_802.1x_CA | Entrust_802.1x_G2_CA | Entrust_802.1x_L1K_CA | Fortinet_CA
           | Fortinet_CA2}
      set cn <string>
      set cn-type {FQDN | email | ipv4 | ipv6 | string}
      set ldap-mode {password | principal-name}
      set ldap-password <password>
      set ldap-server <string>
      set ldap-username <string>
      set mandatory-ca-verify {enable | disable}
      set passwd <password>
      set subject <string>
      set two-factor {enable |disable}
   next
end
```

The following table describes the parameters:

Variable	Description	Default
<peer_name>	Enter the name of the peer user.	No default
ca {Entrust_802.1x_CA | Entrust_802.1x_G2_CA | Entrust_802.1x_L1K_CA | Fortinet_CA | Fortinet_CA2}	Select a certificate authority (CA) for the peer certificate.	No default
cn <string>	Enter the common name for the peer certificate.	No default
cn-type {FQDN | email | ipv4 | ipv6 | string}	Enter the type of common name for the peer certificate: fully qualified domain name, email address, IPv4 address, IPv6 address, or a text description.	string
ldap-mode {password | principal-name}	Select whether the peer LDAP requires a password or an email address. The password is specified with the `set ldap-password` command.	password
ldap-password <password>	Enter the password for the peer LDAP.	No default

Variable	Description	Default
	This option is available only when the `ldap-mode` is set to `password`.	
ldap-server <string>	Enter the name of the LDAP server used for checking access permission.	No default
ldap-username <string>	Enter the user name for the LDAP server.	No default
mandatory-ca-verify {enable | disable}	Enable or disable whether there is mandatory CA verification.	disable
passwd <password>	Enter the user password for two-factor authentication. This option is available only when `two-factor` is enabled.	No default
subject <string>	Enter any limitations on the peer certificate name.	No default
two-factor {enable |disable}	Enable or disable two-factor authentication. When this option is enabled, the certificate and password are required. Specify the password in the `set passwd` command.	disable

User groups

A user group contains a list of local and remote users.

Security policies allow access to specified user groups only. This restricted access enforces Role Based Access Control (RBAC) to your organization's network and its resources. Users must be in a group and that group must be part of the security policy.

Using the GUI:

1. Go to *System* > *User* > *Group*.
2. Click *Add Group*.
3. Enter the group name.
4. Select which available users will be members of the new user group.
5. *Enable* to make the user account active.
6. If you want to use an authentication server, select *Add Server*.
   - Select the server name. If no server name is available, go to *System* > *Authentication* to add an authentication server.
   - Enter a group name or select *Any*.
7. Click *Add*.

Using the CLI:

```
config user group
   edit <groupname>
      set authtimeout <timeout>
      set group-type <grp_type>
      set http-digest-realm <attribute>
      set member <names>
      config match
         edit <match_id>
```

```
        set group-name <gname_str>
        set server-name <srvname_str>
    end
end
```

The following table describes the parameters:

Field	Description
groupname	Identifies the user group.
authtimeout <timeout>	Sets the authentication timeout for the user group. The range is 1 to 480 minutes. If this field is set to 0, the global authentication timeout value is used.
group-type <grp_type>	Enter the group type. <grp_type> determines the type of users and is one of the following: • `firewall`—FortiSwitch users defined in user local, user ldap, or user radius • `fsso-service`—Directory Service users
http-digest-realm <attribute>	Enter the realm attribute for MD5-digest authentication.
member <names>	Enter the names of users, peers, LDAP servers, or RADIUS servers to add to the user group. Separate the names with spaces. To add or remove names from the group, you must re-enter the whole list with the additions or deletions required.
config match fields	
<match_id>	Enter an ID for the entry.
group-name <gname_str>	Identifies the matching group on the remote authentication server.
server-name <srvname_str>	Specifies the remote authentication server.

Authentication

This section covers the following topics:

- RADIUS on page 99
- TACACS on page 102

RADIUS

The information you need to configure the system to use a RADIUS server includes:

- The RADIUS server's domain name or IP address
- The RADIUS server's shared secret key

The default port for RADIUS traffic is 1812. Some RADIUS servers use port 1645. You can configure the FortiSwitch unit to use port 1645:

```
config system global
   set radius-port 1645
end
```

To configure RADIUS authentication with the GUI:

1. Go to *System* > *Authentication* > *RADIUS* and click *Add Server*.

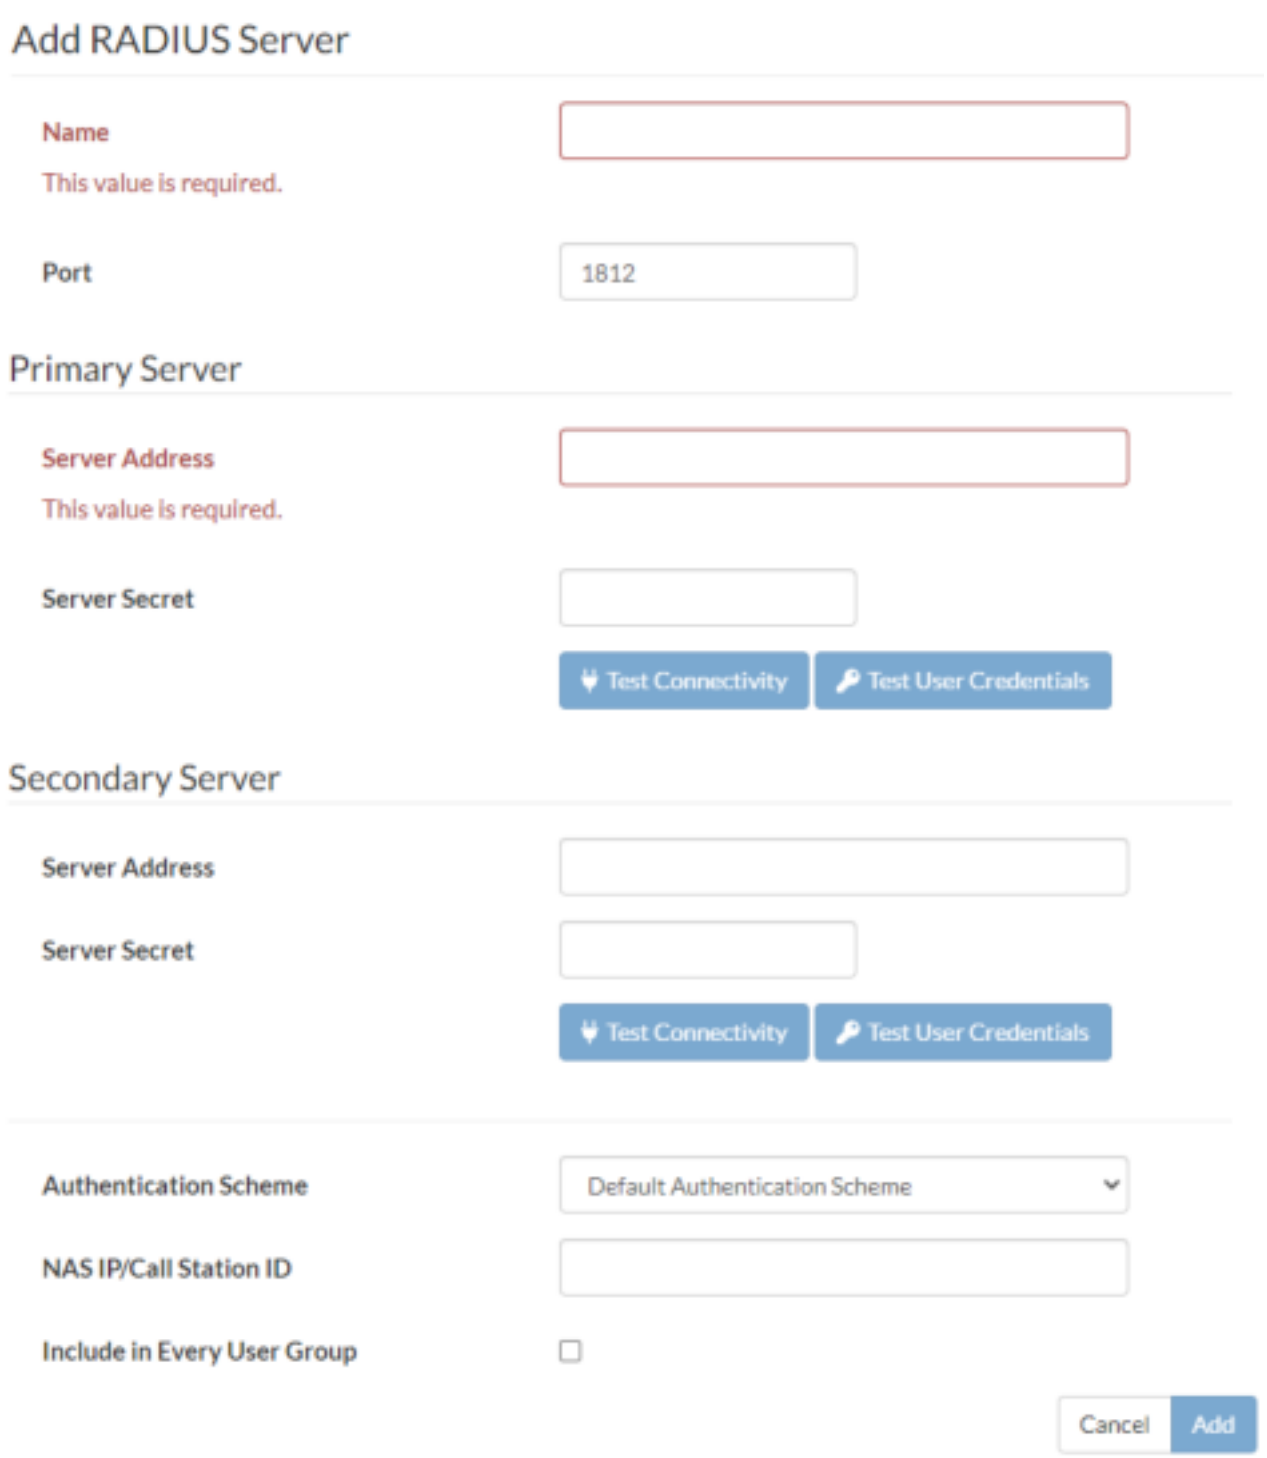

2. Enter the following information.

Field	Description
Name	Enter a name to identify the RADIUS server on the FortiSwitch unit.

Field	Description
Primary Server Address	Enter the IPv4 address of the RADIUS server.
Primary Server Secret	Enter the server secret key, such as radiusSecret. This key can be a maximum of 16 characters long. This value must match the secret on the RADIUS primary server.
Secondary Server Address	Optionally enter the IPv4 address of the secondary RADIUS server.
Secondary Server Secret	Optionally, enter the secondary server secret key, such as radiusSecret2. This key can be a maximum of 16 characters long. This value must match the secret on the RADIUS secondary server.
Authentication Scheme	If you know the RADIUS server uses a specific authentication protocol, select that protocol from the dropdown list. Otherwise, select *Use Default Authentication Scheme*. The default authentication scheme will usually work.
NAS IP/Called Station ID	Enter the IP address to be used as an attribute in RADIUS access requests. The NAS IP address is a RADIUS setting or IP address of the FortiSwitch interface used to talk to the RADIUS server, if not configured. The Called Station ID is the same value as the NAS IP address but in text format.
Include in Every User Group	When this option is enabled, this RADIUS server is automatically included in all user groups. This option is useful if all users will be authenticating with the remote RADIUS server.

3. Click *Test Connectivity* to check if the RADIUS server address is valid.
4. Click *Test User Credentials*, enter the user name and password for the RADIUS server, and then click *Test* to check if the user name and password are valid.
5. Click *Add*.

To configure the FortiSwitch unit for RADIUS authentication, see Port security on page 181.

TACACS

This section contains information on using Terminal Access Controller Access-Control System (TACACS+) authentication with your FortiSwitch unit.

This section covers the following topics:

- TACACS+ server on page 102
- Administrative accounts on page 103
- User accounts on page 104
- Example configuration on page 105

TACACS+ server

TACACS+ is a remote authentication protocol that provides access control for routers, network access servers, and other networked computing devices using one or more centralized servers. TACACS+ allows a client to accept a user

name and password and send a query to a TACACS+ authentication server. The server host determines whether to accept or deny the request and sends a response back that allows or denies the user access to the network.

TACACS+ offers fully encrypted packet bodies and supports both IP and AppleTalk protocols. TACACS+ uses TCP port 49, which is seen as more reliable than RADIUS's UDP protocol.

To configure TACACS+ authentication using the GUI:

1. Go to *System* > *Authentication* > *TACACS* and select *Add Server*.

2. Enter the following information and select *Add*.

Field	Description
Name	Enter a name to identify the TACACS server on the FortiSwitch unit.
Server Address	Enter the domain name (such as fgt.example.com) or the IP address of the TACACS server.
Server Key	Enter the server key for the TACACS server.
Authentication Type	Select the authentication type to use for the TACACS+ server. *Auto* tries PAP, MSCHAP, and CHAP (in that order).

To configure the FortiSwitch unit for TACACS+ authentication, see TACACS on page 102.

Administrative accounts

Administrative, or admin, accounts allow access to various aspects of the FortiSwitch configuration. The level of access is determined by the admin profile that is assigned to the admin account.

See Admin on page 66 for the steps to create an admin profile.

TACACS+ is a remote authentication protocol that provides access control for routers, network access servers, and other network computing devices using one or more centralized servers. If you have configured TACACS+ support and

an administrator is required to authenticate using a TACACS+ server, the FortiSwitch unit contacts the TACACS+ server for authentication.

Using the GUI:

1. Go to *System* > *Admin* > *Administrators* and select *Add Administrator*.
2. Give the administrator account an appropriate name.
3. Select *Remote* for the administrator type.
4. Select a user group for remote users.
5. Enable *Wildcard*.
6. Select an administrator profile.
7. Select *Add*.

Using the CLI:

```
config system admin
   edit tacuser
      set remote-auth enable
      set wildcard enable
      set remote-group <group>
      set accprofile <profile>
   end
end
```

User accounts

User accounts identify a network user and determine what parts of the network the user is allowed to access.

To configure a user account:

```
config user tacacs+
   edit <tacserver>
      set authen-type {ascii | auto | chap | ms_chap | pap}
      set authorization enable
      set key <authorization_key>
      set server <server>
   end
end
```

To configure a user group:

```
config user group
   edit <tacgroup>
      set member <tacserver>
      config match
         edit 1
            set server-name <server>
            set group-name <group>
         end
      end
   end
end
```

Example configuration

The following is an example configuration of a TACACS+ user account, with the CLI syntax shown to create it:

1. Configuring a TACACS user account for login authentication:

```
config user tacacs+
   edit tacserver
   set authen-type ascii
   set authorization enable
   set key temporary
   set server tacacs_server
end
```

2. Configuring a TACACS+user group:

```
config user group
   edit tacgroup
      set member tacserver
      config match
         edit 1
            set server-name tacserver
            set group-name tacgroup
         end
      end
   end
end
```

3. Configuring a TACACS+ system admin user account:

```
config system admin
   edit tacuser
      set remote-auth enable
      set wildcard enable
      set remote-group tacgroup
      set accprofile noaccess
   end
end
```

Certificate

This section contains information about importing local and remote certificates, certificate authority (CA) certificates, and certificate revocation lists (CRLs):

- Local on page 106
- Remote on page 108
- Authorities on page 109
- CRLs on page 110

Local

Use the *Local Certificates* page to import or generate a local certificate.

Importing a local certificate

There are three options for importing a local certificate in the GUI:

- Local certificate

  This option allows you to upload a single file and no key. You must upload a .CER file.
- PKCS12 certificate

  This option takes a specific certificate file type that contains the private key. The certificate is encrypted and a password must be supplied with the certificate file. PKCS #12 certificates are .PFX files. The following sizes are supported: 1024, 2048, and 4096 bits.
- Certificate

  This option is similar to PKCS #12 certificate, but the certificate and key file are separate files, usually .CER and .PEM files.

You cannot import a PKCS12 certificate if the password is missing. To work around this issue, extract the certificate and key from the `.p12` file and then use the GUI to import the certificate and key.

In the CLI, you can import a local certificate from a TFTP server using an IPv4 or IPv6 address or fully qualified domain name.

You can also generate a Rivest–Shamir–Adleman (RSA) certificate or elliptic curve (ECDSA) certificate using a certificate signing request (CSR).

Import a local certificate using the GUI:

1. Go to *System* > *Certificate* > *Local*.
2. Click *Import*.
3. Select the type of certificate that you want to import: *Local Certificate*, *PKCS12 Certificate*, or *Certificate*.
4. In the *Certificate File* field, click *Choose File* and browse to your certificate file.
5. If you selected *Certificate*, click *Choose File* and browse to your key file.
6. If you selected *PKCS12 Certificate* or *Certificate*, enter a password in the *Password* field.
7. Click *Import*.

Import a local certificate using the CLI:

```
execute system certificate local import tftp <local_certificate_file_name> <TFTP_server_
     IPv4_IPv6_FQDN> <cer | p12> [password_for_PKCS12_file]
```

For example:

```
execute system certificate local import tftp p12certificate.p12 10.105.17.77 p12 mypassword
```

Importing a PKCS12 certificate without a password ("")

Because the default algorithm for certificate encryption has changed, there is a different procedure for generating a PKCS12 certificate if you are running OpenSSL 1.x.x.

To generate and import the PKCS12 certificate if you are running OpenSSL 1.x.x:

1. Generate the PKCS12 certificate.
2. Run this command:
   ```
   openssl pkcs12 -in <PKCS12_certificate>.p12 -out <PKCS12_certificate>.pem
   ```
3. Run this command:
   ```
   openssl pkcs12 -keypbe PBE-SHA1-3DES -certpbe PBE-SHA1-3DES -export -in <PKCS12_
        certificate>.pem -out <new_PKCS12_certificate>.p12
   ```
4. Use the CLI or GUI to import `<new_PKCS12_certificate>.p12` into FortiSwitchOS.

Generating a local certificate

Generate a local certificate using the GUI:

1. Go to *System* > *Certificate* > *Local*.
2. Click *Generate*.
3. In the *Certificate Name* field, enter the certificate name, which will appear in the *Local Certificates* table.
4. In the *Key Type* dropdown list, select *RSA* or *Elliptic Curve*.
5. If you selected *RSA* for the key type, select *1024 Bit*, *1536 Bit*, *2048 Bit*, or *4096 Bit* from the *Key Size* dropdown list.
6. If you selected *Elliptic Curve* for the key type, select *SECP256R1*, *SECP384R1*, or *SECP521R1* from the *Curve Name* dropdown list.
7. From the *ID Type* dropdown list, select *Host IP*, *Domain Name*, or *Email*.
8. In the *IP*, *Domain Name*, or *Email* field, enter the IP address, domain name, or email address.
9. In the *Country/Region* dropdown list, select the country or region where the FortiSwitch unit is located.
10. In the *State/Province* field, enter the state or province where the FortiSwitch unit is located.
11. In the *Locality (City)* field, enter the city where the FortiSwitch unit is located.
12. In the *Organization Name* field, enter the name of your organization.
13. In the *Organization Unit* field, enter the business unit.
14. In the *Email* field, enter your email address.
15. In the *Subject Alternative Name* field, enter multiple domains to be used in an SSL certificate.
16. Under *Enrollment Method*, select *File-Based* or *Online SCEP*.
17. If you selected *Online SCEP*, enter the CA server URL and challenge password.
18. Click *Add*.

Generate an elliptic curve local certificate using the CLI:

```
execute system certificate local generate ec <local_certificate_name> {secp256r1 | secp384r1
     | secp521r1} <IP_address_domain_name_email> <country> <state> <city> <organization>
     <business_unit> <email> [<subject alternative name>] [<URL_of_CA_server>] [<challenge_
     password>] [<source_IP_address>] [<CA_identifier_of_CA_server>] [<password_for_
     private_key>
```

For example:

```
execute system certificate local generate ec localcertificate secp256r1 1.2.3.4 northamerica
     CA Sunnyvale Fortinet "R&D" "admin@fortinet.com"
```

Generate an RSA local certificate using the CLI:

```
execute system certificate local generate rsa <local_certificate_name> {1024 | 1536 | 2048 |
     4096} <IP_address_domain_name_email> <country> <state> <city> <organization>
     <business_unit> <email> [<subject alternative name>] [<URL_of_CA_server>] [<challenge_
     password>] [<source_IP_address>] [<CA_identifier_of_CA_server>] [<password_for_
     private_key>]
```

For example:

```
execute system certificate local generate rsa localcertificate 1024 1.2.3.4 northamerica CA
     Sunnyvale Fortinet "R&D" "admin@fortinet.com"
```

Remote

Use the *Remote Certificates* page to import a remote certificate. Remote certificates are public certificates and contain only the public key. They are used as OCSP (Online Certificate Status Protocol) server certificates. They are used to identify a remote device.

In the CLI, you can import a remote certificate from a TFTP server or export a remote certificate from the FortiSwitch unit to a TFTP server.

Import a remote certificate using the GUI:

1. Go to *System* > *Certificate* > *Remote*.
2. Click *Import*.
3. In the *Local PC* field, click *Choose File* and browse to your certificate file.
4. Click *Import*.

Import a remote certificate using the CLI:

```
execute system certificate remote import tftp <remote_certificate_name> <TFTP_server_IP_
     address>
```

For example:

```
execute system certificate remote import tftp remotecertificate 1.2.3.4
```

Export a remote certificate using the CLI:

```
execute system certificate remote export tftp <remote_certificate_name> <file_name_on_TFTP_
     server> <TFTP_server_IP_address>
```

For example

```
execute system certificate remote export tftp remotecertificate remote_cert.cer 1.2.3.4
```

Authorities

FortiSwitch units come with many CA certificates from well-known certificate authorities pre-installed. Use the *Certificate Authorities* page to add private CA certificates to the FortiSwitch unit so that certificates signed by the private CA are trusted by the FortiSwitch unit.

In the CLI, you can import a CA certificate from a TFTP or SCEP server to the FortiSwitch unit or export a CA certificate from the FortiSwitch unit to a TFTP server. Before using either CLI command, you must obtain a CA certificate issued by a Certificate Authority.

Import a CA certificate using the GUI:

1. Go to *System* > *Certificate* > *Authorities*.
2. Click *Import*.
3. In the *Type* dropdown list, select *SCEP* or *Local PC.*
4. If you selected *SCEP*, enter the CA server URL and the CA identifier.
5. If you selected *Local PC*, click *Choose File* and browse to your certificate file.
6. Click *Import*.

Import a CA certificate using the CLI:

- `execute system certificate ca import auto <CA_certificate_server_URL> [ca_identifier]`
- `execute system certificate ca import tftp <file_name_on_TFTP_server> <TFTP_server_IP_address>`

For example:

```
execute system certificate ca import tftp cacert.cer 1.2.3.4
```

Export a CA certificate using the CLI:

```
execute system certificate ca export tftp <CA_certificate_name> <file_name_on_TFTP_server>
     <TFTP_server_IP_address>
```

For example:

```
execute system certificate ca export tftp cacertificate ca_cert.cer 1.2.3.4
```

CRLs

Because it is not possible to recall a certificate, the CRL lists certificates signed by valid CAs that should no longer be trusted. Certificates might be revoked for many reasons, such as if the certificate was issued erroneously or if the private key of a valid certificate has been compromised.

In the CLI, you can get a certificate revocation list using LDAP, HTTP, or SCEP, depending on the `autoupdate` configuration.

Import a CRL using the GUI:

1. Go to *System* > *Certificate* > *CRLs*.
2. Click *Import*.
3. In the *Type* dropdown list, select *No Local* or *Local PC*.
4. If you selected *No Local*, enter the URL of the HTTP server, select the LDAP server and SCEP server, and enter the URL for the SCEP server.
5. If you selected *Local PC*, click *Choose File* and browse to your certificate file.
6. Click *Import*.

Import a CRL using the CLI:

```
execute system certificate crl import auto <CRL_name>
```

For example:

```
execute system certificate crl import auto CRL1
```

Flow export

NOTE:

- To see which models support this feature, refer to the FortiSwitch feature matrix.
- Starting in FortiSwitchOS 7.0.0, you can use the CLI to configure multiple flow-export collectors, control how often the template is exported, and specify a Berkeley packet filter (BPF).
- Layer-2 flows for NetFlow version 1 and NetFlow version 5 are not supported.
- For 2xxE models and higher, flow export uses pseudorandom sampling (approximately 1 of *x* packets).

You can sample IP packets on a FortiSwitch unit and then export the data in NetFlow format or Internet Protocol Flow Information Export (IPFIX) format. Specifying the flow-tracking level controls which fields are exported:

IPv4 or IPv6	Flow-tracking level	Fields that are exported
IPv4 and IPv6	MAC	• Source MAC Address • Destination MAC Address • First Switched Time • Last Switched Time • Bytes • Packets • Input Interface Index • Output Interface Index
IPv4	IP	• Source IPv4 Address • Destination IPv4 Address • First Switched Time • Last Switched Time • Bytes • Packets • Input Interface Index • Output Interface Index
IPv6	IP	• Source IPv6 Address • Destination IPv6 Address • First Switched Time • Last Switched Time • Bytes • Packets • Input Interface Index • Output Interface Index • IP Protocol Version

IPv4 or IPv6	Flow-tracking level	Fields that are exported
IPv4	Protocol	• Source IPv4 Address • Destination IPv4 Address • First Switched Time • Last Switched Time • Bytes • Packets • Input Interface Index • Output Interface Index • IP Protocol version • Protocol • TOS
IPv6	Protocol	• Source IPv6 Address • Destination IPv6 Address • First Switched Time • Last Switched Time • Bytes • Packets • Input Interface Index • Output Interface Index • IP Protocol Version • Protocol • TOS
IPv4	Port	• Source IPv4 Address • Destination IPv4 Address • First Switched Time • Last Switched Time • Bytes • Packets • Input Interface Index • Output Interface Index • IP Protocol version • Protocol • TOS • Source Port • Destination Port • TCP Flags

IPv4 or IPv6	Flow-tracking level	Fields that are exported
IPv6	Port	• Source IPv6 Address • Destination IPv6 Address • First Switched Time • Last Switched Time • Bytes • Packets • Input Interface Index • Output Interface Index • IP Protocol Version • Protocol • TOS • Source Port • Destination Port • TCP Flags
IPv4	VLAN	• Source IPv4 Address • Destination IPv4 Address • First Switched Time • Last Switched Time • Bytes • Packets • Input Interface Index • Output Interface Index • IP Protocol version • Protocol • TOS • Source Port • Destination Port • TCP Flags • ICMP Type • Source VLAN ID

IPv4 or IPv6	Flow-tracking level	Fields that are exported
IPv6	VLAN	• Source IPv6 Address • Destination IPv6 Address • First Switched Time • Last Switched Time • Bytes • Packets • Input Interface Index • Output Interface Index • IP Protocol Version • Protocol • TOS • Source Port • Destination Port • TCP Flags • ICMP Type • Source VLAN ID

The maximum number of concurrent flows is defined by the FortiSwitch model. When this limit is exceeded, the oldest flow expires and is exported.

The following table describes the NetFlow template:

Template number	Template field	Template key	Template size	IP (IPv4)	Protocol (IPv4)	Ports (IPv4)	VLAN (IPv4)	Aggregate (IPv4)	IP (IPv6)	Protocol (IPv6)	Ports (IPv6)	VLAN (IPv6)	Aggregate (IPv6)
1	Source IPv4 Address	✓	4	✓	✓	✓	✓	✓	—	—	—	—	—
2	Destination IPv4 Address	✓	4	✓	✓	✓	✓	✓	—	—	—	—	—
3	Source IPv6 Address	✓	16	—	—	—	—	—	✓	✓	✓	✓	—
4	Destination IPv6 Address	✓	16	—	—	—	—	—	✓	✓	✓	✓	—
5	Source IPv4 Mask	—	1	—	—	—	—	✓	—	—	—	—	—
6	Destination IPv4 Mask	—	1	—	—	—	—	✓	—	—	—	—	—
7	IP Protocol Version	—	1	✓	✓	✓	✓	✓	✓	✓	✓	✓	—
8	Protocol	✓	1	—	✓	✓	✓	—	—	✓	✓	✓	—
9	TOS	—	1	—	✓	✓	✓	—	—	✓	✓	✓	—
10	Source Port Number	✓	2	—	—	✓	✓	—	—	—	✓	✓	—
11	Destination Port Number	✓	2	—	—	✓	✓	—	—	—	✓	✓	—
12	TCP Flags	—	1	—	—	✓	✓	—	—	—	✓	✓	—
13	VLAN Id	✓	2	—	—	—	✓	—	—	—	—	✓	—
14	Flow Start Uptime	—	4	✓	✓	✓	✓	✓	✓	✓	✓	✓	—
15	Flow End Uptime	—	4	✓	✓	✓	✓	✓	✓	✓	✓	✓	—
16	Packet Count	—	4	✓	✓	✓	✓	✓	✓	✓	✓	✓	—
17	Packet Bytes	—	4	✓	✓	✓	✓	✓	✓	✓	✓	✓	—
18	Ingress Interface Index	✓	4	✓	✓	✓	✓	✓	✓	✓	✓	✓	—
19	Egress Interface Index	✓	4	✓	✓	✓	✓	✓	✓	✓	✓	✓	—

The following table describes the IPFIX template:

Template number	Template field	Template key	Template size	IP (IPv4)	Protocol (IPv4)	Ports (IPv4)	VLAN (IPv4)	Aggregate (IPv4)	IP (IPv6)	Protocol (IPv6)	Ports (IPv6)	VLAN (IPv6)	Aggregate (IPv6)
1	Source IPv4 Address	✓	4	✓	✓	✓	✓	—	—	—	—	—	—
2	Destination IPv4 Address	✓	4	✓	✓	✓	✓	—	—	—	—	—	—
3	Source IPv6 Address	✓	16	—	—	—	—	—	✓	✓	✓	✓	—
4	Destination IPv6 Address	✓	16	—	—	—	—	—	✓	✓	✓	✓	—
5	Source IPv4 Prefix	✓	4	—	—	—	—	✓	—	—	—	—	—
6	Destination IPv4 Prefix	✓	4	—	—	—	—	✓	—	—	—	—	—
7	Source IPv4 Prefix Length	—	1	—	—	—	—	✓	—	—	—	—	—
8	Destination IPv4 Prefix Length	—	1	—	—	—	—	✓	—	—	—	—	—
9	IP Protocol Version	—	1	✓	✓	✓	✓	✓	✓	✓	✓	✓	—
10	Protocol	✓	1	—	✓	✓	✓	—	—	✓	✓	✓	—
11	TOS	—	1	—	✓	✓	✓	—	—	✓	✓	✓	—
12	Source Port Number	✓	2	—	—	✓	✓	—	—	—	✓	✓	—
13	Destination Port Number	✓	2	—	—	✓	✓	—	—	—	✓	✓	—
14	TCP Flags	—	1	—	—	✓	✓	—	—	—	✓	✓	—
15	VLAN Id	✓	2	—	—	—	✓	—	—	—	—	✓	—
16	Flow Start Uptime	—	4	✓	✓	✓	✓	✓	✓	✓	✓	✓	—
17	Flow End Uptime	—	4	✓	✓	✓	✓	✓	✓	✓	✓	✓	—
18	Packet Count	—	4	✓	✓	✓	✓	✓	✓	✓	✓	✓	—
19	Packet Bytes	—	4	✓	✓	✓	✓	✓	✓	✓	✓	✓	—
20	Ingress Interface Index	✓	4	✓	✓	✓	✓	✓	✓	✓	✓	✓	—
21	Egress Interface Index	✓	4	✓	✓	✓	✓	✓	✓	✓	✓	✓	—

To use flow export:

1. Enabling packet sampling on page 115
2. Configuring flow export on page 116
3. Viewing the flow-export data on page 119
4. Deleting the flow-export data on page 120

Enabling packet sampling

To use flow export, you must first enable packet sampling for each switch port and trunk:

```
config switch interface
   edit <interface>
      set packet-sampler enabled
      set packet-sample-rate <0-99999>
   next
end
```

Variable	Description	Default
packet-sampler {enabled | disabled}	Enable or disable packet sampling for flow export.	disabled
packet-sample-rate <0-99999>	If packet-sampler is set to enabled, you can change the packet sample rate.	512

Configuring flow export

Using the GUI:

1. Go to *System* > *Flow Export* > *Configure*.
2. Configure the collectors.
   a. Click +.
   b. In the *Name* field, enter the name of the collector.
   c. Required. In the *IP* field, enter the IPv4 address for the collector. When the value is “0.0.0.0” or blank, the feature is disabled.
   d. In the *Port* field, enter the port number for the collector. The default port for NetFlow is 2055; the default port for IPFIX is 4739.
   e. In the *Transport* dropdown list, select *SCTP*, *TCP*, or *UDP* for the transport of exported packets.
3. Configure the flow export options.
   a. In the Format drop-down list, select the format of the exported flow data as NetFlow version 1, NetFlow version 5, NetFlow version 9, or IPFIX sampling.
   NOTE: When the export format is NetFlow version 5, the sample rate used in the exported packets is derived from the lowest port number where sampling is enabled. Fortinet recommends that administrators using NetFlow version 5 set the sample rate consistently across all ports.
   b. In the Identity field, enter a unique number to identify which FortiSwitch unit the data originates from. If the identity is not specified, the “Burn in MAC” value is used instead (from the `get system status` command output).
   c. In the Level field, select the flow-tracking level from one of the following:
   —When you select *IP*, the FortiSwitch unit collects the source IP address and destination IP address from the sample packet.
   —When you select *MAC*, the FortiSwitch unit collects the source MAC address and destination MAC address from the sample packet.
   —When you select *Port,* the FortiSwitch unit collects the source IP address, destination IP address, source port, destination port, and protocol from the sample packet.
   —When you select *Protocol*, the FortiSwitch unit collects the source IP address, destination IP address, and protocol from the sample packet.
   —When you select *VLAN*, the FortiSwitch unit collects the source IP address, destination IP address, source port, destination port, protocol, and VLAN from the sample packet.
   d. In the Max Export Packet Size (Bytes) field, enter the maximum size of exported packets in the application level.
4. Configure the timeouts.
   a. In the General field, enter the general timeout in seconds for the flow session.
   b. In the ICMP field, enter the ICMP timeout for the flow session.
   c. In the Max field, enter the maximum number of seconds before the flow session times out.
   d. In the TCP field, enter the TCP timeout for the flow session.
   e. In the TCP FIN field, enter the TCP FIN flag timeout for the flow session.
   f. In the TCP RST field, enter the TCP RST flag timeout for the flow session.
   g. In the UDP field, enter the UDP timeout for the flow session.
5. Configure the aggregates.
   a. Select +.
   b. In the ID field, enter a number to identify the entry or use the default value.
   c. Required. In the IP/Netmask field, enter the IPv4 address and mask to match. All matching sessions are

aggregated into the same flow.

d. To add another entry, select +.

6. Select *Update*.

Using the CLI:

```
config system flow-export
   set filter <BPF_filter>
   set format {netflow1 | netflow5 | netflow9 | ipfix}
   set identity <hexadecimal>
   set level {ip | mac | port | proto | vlan}
   set max-export-pkt-size <512-9216 bytes>
   set template-export-period <1-60 minutes>
   set timeout-general <60-604800 seconds>
   set timeout-icmp <60-604800 seconds>
   set timeout-max <60-604800 seconds>
   set timeout-tcp <60-604800 seconds>
   set timeout-tcp-fin <60-604800 seconds>
   set timeout-tcp-rst <60-604800 seconds>
   set timeout-udp <60-604800 seconds>
   config collectors
      edit <collector_name>
         set ip <IPv4_address>
         set port <port_number>
         set transport {sctp | tcp | udp}
      end
   config aggregates
      edit <aggregate_ID>
         set ip <IPv4_address_mask>
      end
   end
```

Variable	Description	Default
filter <string>	Specify the Berkeley packet filter (BPF) to use. For example, `set filter "host 33.33.33.2"`.	No default
format {netflow1 | netflow5 | netflow9 | ipfix}	You can set the format of the exported flow data as NetFlow version 1, NetFlow version 5, NetFlow version 9, or IPFIX sampling. **NOTE:** When the export format is NetFlow version 5, the sample rate used in the exported packets is derived from the lowest port number where sampling is enabled. Fortinet recommends that administrators using NetFlow version 5 set the sample rate consistently across all ports.	netflow9
identity <hexadecimal>	Required. Enter a unique number to identify which FortiSwitch unit the data originates from. The range of values is 0x00000000-0xFFFFFFFF. If `identity` is not specified, the “Burn in MAC” value is used instead (see `get system status`).	0x00000000

Variable	Description	Default
level {ip | mac | port | proto | vlan}	You can set the flow-tracking level to one of the following: - `ip`—The FortiSwitch unit collects the source IP address and destination IP address from the sample packet. • `mac`—The FortiSwitch unit collects the source MAC address and destination MAC address from the sample packet. • `port`—The FortiSwitch unit collects the source IP address, destination IP address, source port, destination port, and protocol from the sample packet. • `proto`—The FortiSwitch unit collects the source IP address, destination IP address, and protocol from the sample packet. • `vlan`—The FortiSwitch unit collects the source IP address, destination IP address, source port, destination port, protocol, and VLAN from the sample packet.	ip
max-export-pkt-size <integer>	Set the maximum size in bytes of exported packets in the application level. The range of values is 512-9216.	512
template-export-period <1-60>	Set the number of minutes before the template is exported.	5
timeout-general <integer>	Set the general timeout in seconds for the flow session. The range of values is 60-604800.	3600
timeout-icmp <integer>	Set the ICMP timeout for the flow session. The range of values is 60-604800.	300
timeout-max <integer>	Set the maximum number of seconds before the flow session times out. The range of values is 60-604800.	604800
timeout-tcp <integer>	Set the TCP timeout for the flow session. The range of values is 60-604800.	3600
timeout-tcp-fin <integer>	Set the TCP FIN flag timeout for the flow session. The range of values is 60-604800.	300
timeout-tcp-rst <integer>	Set the TCP RST flag timeout for the flow session. The range of values is 60-604800.	120
timeout-udp <integer>	Set the UDP timeout for the flow session. The range of values is 60-604800.	300
config collectors		
<collector_name>	Enter the name of the flow-export collector.	No default
ip <IPv4_address>	Enter the IP address for the collector. The default is 0.0.0.0. Setting the value to “0.0.0.0” or “” disables this feature. The format is xxx.xxx.xxx.xxx.	0.0.0.0
port <port_number>	Enter the port number for the collector.	0

Variable	Description	Default
	The range of values is 0-65535. The default port for NetFlow is 2055; the default port for IPFIX is 4739.	
transport {sctp | tcp | udp}	You can set exported packets to use UDP, TCP, or SCTP for transport.	udp
config aggregates		
<id>	Enter the identifier.	No default
<IPv4_address_mask>	Enter the IPv4 address and mask to match. All matching sessions are aggregated into the same flow.	No default

For example:

```
config system flow-export
   config collectors
      edit "netflow-col-1"
         set ip 192.168.3.100
         set port 0
      next
      edit "netflow-col-2"
         set ip 192.168.4.100
         set port 2004
      next
   end
end
```

Viewing the flow-export data

Using the GUI:

Go to *System* > *Flow Export* > *Monitor*.

Using the CLI:

You can display the flow-export data or raw data for a specified number of records or for all records. You can also display statistics for flow-export data.

```
get system flow-export-data flows {all | <count>} {ip | subnet | mac | all} <switch_
     interface_name>
get system flow-export-data flows-raw {all | <count>} {ip | subnet | mac | all} <switch_
     interface_name>
get system flow-export-data statistics
```

Layer-2 flows for netflow1 and netflow5 are not supported. For the output of the `get system flow-export-data statistics` command, the *Incompatible Type* field displays how many flows are not exported because they are not supported.

Deleting the flow-export data

Use the following commands to delete or expire all flow-export data:

```
diagnose sys flow-export delete-flows-all
diagnose sys flow-export expire-flows-all
```

DHCP

A DHCP server provides an address, from a defined address range, to a client on the network that requests it.

You can configure one or more DHCP servers on any FortiSwitch 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 FortiSwitch 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 the appropriate routing so that its response packets to the DHCP clients arrive at the unit.

NOTE:

- DHCP snooping and the DHCP server can be enabled at the same time.
- The DHCP server and DHCP relay cannot be enabled at the same time.

This section covers the following topics:

- Configuring a DHCP server on page 120
- Detailed operation of a DHCP relay on page 126
- Configuring a DHCP relay on page 126

Configuring a DHCP server

NOTE: To see which models support this feature, refer to the FortiSwitch feature matrix. The following table lists the maximum number of clients for the supported FortiSwitch models:

FortiSwitch models	Maximum number of clients
FS-1xx	250
FS-2xx	500
FS-4xx	15,000
FS-5xx	20,000
FS-1024D, FS-1048D, FS-3032D	30,000
FS-1048E, FS03032E	50,000

Using the GUI:

1. Go to *System* > *DHCP*.
2. Select *Add DHCP Server*.
3. Required. In the ID field, enter a number to identify the entry.
4. Select the Enable checkbox to make the DHCP server active.
5. Select the Auto-Configuration checkbox if you want the DHCP server to dynamically assign IP addresses to hosts on the network connected to the interface.
6. Required. In the Netmask field, enter the netmask of the addresses that the DHCP server assigns.
7. In the Interface drop-down list, select an interface. The DHCP server assigns IP configurations to clients connected to this interface.
8. Required. In the Lease Time field, enter the lease time in seconds. The lease time determines the length of time an IP address remains assigned to a client.
9. Required. In the Conflicted IP Timeout field, enter the number of seconds before a conflicted IP address is removed from the DHCP range and is available to be reused.
10. In the Default Gateway field, enter the IP address of the default gateway that the DHCP server assigns to DHCP clients.
11. In the Domain field, enter the domain name suffix for the IP addresses that the DHCP server assigns to the clients.
12. In the Next Server field, enter the IPv4 address of a server (for example, a TFTP sever) that DHCP clients can download a boot file from.
13. In the Filename field, enter the name of the boot file on the TFTP server.
14. In the DNS Service Type drop-down list, select how DNS servers are assigned to DHCP clients.
    - Select *Default* for clients to be assigned the FortiSwitch unit's configured DNS servers.
    - Select *Local* to use the IP address of the DHCP server interface for the client's DNS server IP address.
    - Select *Specify* to enter IPv4 addresses for up to three DNS servers.
15. In the Controller 1, Controller 2, and Controller 3 fields, enter the IPv4 addresses for the WiFi access controllers.
16. In the NTP Service Type drop-down list, select how Network Time Protocol (NTP) servers are assigned to DHCP clients.
    - Select *Default* for clients to be assigned the FortiSwitch unit's configured NTP servers.
    - Select *Local* to use the IP address of the DHCP server interface for the client's NTP server IP address.
    - Select *Specify* to enter the IPv4 address for up to three NTP servers.
17. In the WINS Server section, enter the IPv4 addresses for the Windows Internet Name Service (WINS) servers.
18. In the Timezone Mode drop-down list, select how the DHCP server sets the client's time zone.
    - Select *Default* for clients to be assigned the FortiSwitch unit's configured time zone.
    - Select *Disable* for the DHCP server to not set the client's time zone.
    - Select *Specify* to choose which time zone is assigned to DHCP clients.
19. In the VCI area, select the Enable checkbox to enter the vendor class identifier (VCI) to match. When enabled, only DHCP requests with a matching VCI are served.
20. In the IP Ranges section, you can configure the IP address range.
    a. In the ID field, enter a unique number to identify the entry or use the default value.
    b. Required. In the Start IP field, enter the start of the DHCP IP address range.
    c. Required. In the End IP field, enter the end of the DHCP IP address range.
    d. To add another IP address range, select *Add IP Range*.
21. In the Exclusion Ranges section, you can block a range of addresses that will not be included in the available addresses for the connecting users.
    a. Select *Add Exclusion Range*.
    b. In the ID field, enter a number to identify the entry or use the default value.

c. In the Start IP field, enter the start of the IP address range that will not be assigned to clients.
d. In the End IP field, enter the end of the IP address range that will not be assigned to clients.
e. To add another exclusion range, select *Add Exclusion Range*.

22. In the Reserved Addresses section, you can reserve IP addresses for the DHCP server to use to assign IP addresses to specific MAC addresses.
    a. Select *Add IP*.
    b. In the ID field, enter a number to identify the entry or use the default value.
    c. In the Type drop-down list, select whether to match the IP address with the MAC address or DHCP option 82.
    d. In the Action drop-down list, select how the DHCP server configures the client with the reserved MAC address. Select *Reserved* for the DHCP server to assign the reserved IP address to the client with this MAC address. Select *Assign* for the DHCP server to configure the client with this MAC address like any other client. Select *Block* to prevent the DHCP server from assigning IP settings to the client with this MAC address.
    e. In the Description field, enter a description of this entry.
    f. In the IP field, enter the IPv4 address to be reserved for the MAC address. This value is required when the action is *Reserved* and the type is *MAC*.
    g. In the MAC field, enter the MAC address of the client that will get the reserved IP address. This value is required when the type is *MAC* and the action is *Assign* or *Block*.
    h. In the Circuit Type drop-down list, select whether the format of the Circuit ID is hexadecimal or string. This option is only available when the type is *Option-82*.
    i. In the Circuit ID field, enter the DHCP option-82 Circuit ID of the client that will get the reserved IP address. The Circuit ID format is controlled by the Circuit Type setting. This value is required when the type is *Option-82*.
    j. In the Remote Type drop-down list, select whether the format of the Remote ID is hexadecimal or string. This option is only available when the type is *Option-82*.
    k. In the Remote ID field, enter the DHCP option-82 Remote ID of the client that will get the reserved IP address. This value is required when the type is *Option-82*.
    l. To add another reserved address, select *Add IP*.
23. In the Options section, you can add up to 30 DHCP custom options.
    a. Select *Add Option*.
    b. In the ID field, enter a number to identify the entry or use the default value.
    c. In the Type drop-down list, select the format of the DHCP option: fully qualified domain name (FQDN), hexadecimal, IP address, or string.
    d. In the Code field, select the DHCP option code. The range is 0-255.
    e. In the Value field, enter the DHCP option value. This value is required when the type is set to *FQDN*, *Hex*, or *String*.
    f. In the IP field, enter the IP address. This value is required when the type is set to *IP*.
    g. To add another DHCP custom option, select *Add Option*.
24. Select *Add* to save the new DHCP server.

Using the CLI:

```
config system dhcp server
   edit <id>
      set auto-configuration {enable | disable}
      set conflicted-ip-timeout <integer>
      set default-gateway <xxx.xxx.xxx.xxx>
      set dns-server1 <xxx.xxx.xxx.xxx>
      set dns-server2 <xxx.xxx.xxx.xxx>
      set dns-server3 <xxx.xxx.xxx.xxx>
      set dns-service {default | local | specify
      set domain <string>
```

```
      set filename <string>
      set interface <string>
      set lease-time <integer>
      set netmask <xxx.xxx.xxx.xxx>
      set next-server <xxx.xxx.xxx.xxx>
      set ntp-server1 <xxx.xxx.xxx.xxx>
      set ntp-server2 <xxx.xxx.xxx.xxx>
      set ntp-server3 <xxx.xxx.xxx.xxx>
      set ntp-service {default | local | specify}
      set status {enable | disable}
      set tftp-server <xxx.xxx.xxx.xxx>
      set timezone <00-75>
      set timezone-option {default | disable | specify}
      set vci-match {enable | disable}
      set vci-string <VCI_strings>
      set wifi-ac1 <xxx.xxx.xxx.xxx>
      set wifi-ac2 <xxx.xxx.xxx.xxx>
      set wifi-ac3 <xxx.xxx.xxx.xxx>
      set wins-server1 <xxx.xxx.xxx.xxx>
      set wins-server2 <xxx.xxx.xxx.xxx>
   next
end
```

For example:

```
config system dhcp server
   edit 1
      set default-gateway 50.50.50.2
      set domain "FortiswitchTest.com"
      set filename "text1.conf"
      set interface "svi10"
      config ip-range
         edit 1
            set end-ip 50.50.0.10
            set start-ip 50.50.0.5
         next
      end
      set lease-time 360
      set netmask 255.255.0.0
      set next-server 60.60.60.2
      config options
         edit 1
            set value "dddd"
         next
      end
      set tftp-server "1.2.3.4"
      set timezone-option specify
      set wifi-ac1 5.5.5.1
      set wifi-ac2 5.5.5.2
      set wifi-ac3 5.5.5.3
      set wins-server1 6.6.6.1
      set wins-server2 6.6.6.2
      set dns-server1 7.7.7.1
      set dns-server2 7.7.7.2
      set dns-server3 7.7.7.3
      set ntp-server1 8.8.8.1
      set ntp-server2 8.8.8.2
```

```
      set ntp-server3 8.8.8.3
   next
end
```

Configuring the IP address range

By default, the FortiSwitch unit assigns an address range based on the address of the interface for the complete scope of the address. For example, if the interface address is 172.20.120.230, the default range created is 172.20.120.231 to 172.20.120.254.

To configure the IP address range:

```
config system dhcp server
   edit <id>
      config ip-range
         edit <id>
            set end-ip <xxx.xxx.xxx.xxx>
            set start-ip <xxx.xxx.xxx.xxx>
         next
      end
   next
end
```

Excluding addresses in DHCP

If you have a large address range for the DHCP server, you can block a range of addresses that will not be included in the available addresses for the connecting users.

To exclude addresses in DHCP:

```
config system dhcp server
   edit <id>
      config exclude-range
         edit <id>
            set end-ip <xxx.xxx.xxx.xxx>
            set start-ip <xxx.xxx.xxx.xxx>
         next
      end
   next
end
```

Assigning IP settings to specific MAC addresses

If you want the DHCP server to assign IP addresses to specific MAC addresses, you need to reserve the IP addresses.

To reserve IP addresses:

```
config system dhcp server
   edit <id>
      config reserved-address
         edit <id>1
            set action {assign | block | reserved}
```

```
            set circuit-id {<string> | <hex>}
            set circuit-id-type {hex | string}
            set description <string>
            set ip <xxx.xxx.xxx.xxx>
            set mac <xx:xx:xx:xx:xx:xx>
            set remote-id {<string> | <hex>}
            set remote-id-type {hex | string}
            set type {mac | option82}
        next
    end
  next
end
```

Configuring DHCP custom options

The DHCP server maintains a table for the potential options. The FortiSwitch DHCP server supports up to a maximum of 30 custom options.

To configure the DHCP custom options:

```
config system dhcp server
  edit <id>
    config options
      edit <id>
        set code <integer>
        set ip <IP_addresses>
        set type {fqdn | hex | ip | string}
        set value <string>
      next
    end
  next
end
```

Listing DHCP leases

The lease time determines the length of time an IP address remains assigned to a client. After the lease expires, the address is released for allocation to the next client that requests an IP address. Use one of the following commands to check the DHCP leases:

```
execute dhcp lease-list
execute dhcp lease-list <interface>
```

Breaking DHCP leases

If you need to end an IP address lease, you can break the lease. This is useful if you have limited addresses and longer lease times when some leases are no longer necessary, for example, with corporate visitors. Use one of the following commands to break the DHCP leases:

```
execute dhcp lease-clear all
execute dhcp lease-clear <xxx.xxx.xxx.xxx,yyy.yyy.yyy.yyy,...>
```

Detailed operation of a DHCP relay

A DHCP relay operates as follows:

1. DHCP client C broadcasts a DHCP/BOOTP discover message on its subnet.
2. The relay agent examines the gateway IP address field in the DHCP/BOOTP message header. If the field has an IP address of 0.0.0.0, the agent fills it with the relay agent's or router's IP address and forwards the message to the remote subnet of the DHCP server.
3. When DHCP server receives the message, it examines the gateway IP address field for a DHCP scope that can be used by the DHCP server to supply an IP address lease.
4. If DHCP server has multiple DHCP scopes, the address in the gateway IP address field (GIADDR) identifies the DHCP scope from which to offer an IP address lease.
5. DHCP server sends an IP address lease offer (DHCPOFFER) directly to the relay agent identified in the gateway IP address (GIADDR) field.
6. The router then relays the address lease offer (DHCPOFFER) to the DHCP client.

NOTE:

- DHCP relay service supports up to 8 relay targets per interface.
- Each target is sent a copy of the DHCP message.

Configuring a DHCP relay

You can configure a DHCP relay on any layer-3 interface.

Using the GUI:

1. Go to *System* > *Network* > *Interface* > *Physical*.
2. Select *Edit* for an interface.
3. Select *Enabled* under *DHCP Relay*.
4. Enter the IP addresses for the relay servers, separated by a space.
5. If you want to include Option-82 data, select *Option-82*.
6. Select *Update*.

Using the CLI:

```
config system interface
   edit <interface-name>
      set dhcp-relay-service (enable | disable)
      set dhcp-relay-ip <ip-address1> [<ip-address2> ... <ip-address8>]
      set dhcp-relay-option82 (enable | disable)
   next
end
```

In the following example, the DHCP server has address 192.168.23.2:

```
config system interface
   edit "v15-p15"
      set dhcp-relay-service enable
      set dhcp-relay-ip "192.168.23.2" -> the DHCP server address
      set ip 192.168.15.1 255.255.255.0 -> the DHCP client subnet
      set allowaccess ping ssh snmp telnet set snmp-index 53
```

```
    set vlanid 15
    set interface "internal"
  next
end
```

Packet capture

When troubleshooting networks, it helps to look inside the header of the packets. This helps to determine if the packets, route, and destination are all what you expect. Packet capture is also called a network tap, packet sniffing, or logic analyzing.

To capture packets:

1. Creating a packet-capture profile on page 127.
2. Starting the packet capture on page 128.
3. Pausing or stopping the packet capture on page 129.
4. Displaying or uploading the packet capture on page 129.
5. Deleting the packet-capture file on page 130.

The maximum number of packet-capture profiles and the RAM disk size allotted for packet captures are different for the various platforms:

Platform	Maximum number of profiles	RAM disk size in MB
1xx (124F/148E/148F)	8	20
2xx	8	50
4xx	16	75
5xx	16	100
1xxx	16	100
3xxx	16	100

Creating a packet-capture profile

To specify which packets to capture, define a filter and select a switch or system interface on which to capture the packets. You cannot select both a switch interface and a system interface.

The filter uses flexible logic. For example, if you want packets using UDP port 1812 between hosts named `forti1` and either `forti2` or `forti3`:

```
'udp and port 1812 and host forti1 and \( forti2 or forti3 \)'
```

You can specify the number of packets to capture and the maximum packet length to be captured. The maximum number of packets that can be captured depends on the RAM disk size.

Using the GUI:

1. Go to *System* > *Packet Capture*.
2. Select *Add Packet Capture*.
3. Enter a name for the packet-capture profile.
4. Select the switch or system interface that you want to capture packets on.
5. Enter how many packets to capture on the selected interface.
6. Enter the maximum packet length in bytes to capture on the interface.
7. If you want to use a filter to select which packets to capture, select the *Filter* checkbox.
   a. If you want to filter by hosts, enter the IP addresses, separated with commas.
   b. If you want to filter by ports, enter port numbers or ranges, separated with commas.
   c. If you want to filter by VLANs, enter VLAN numbers, separated with commas.
   d. If you want to filter by protocols, enter the numbers, separated with commas.
8. Select *Add*.

Using the CLI:

```
config system sniffer-profile
   edit <profile_name>
      set filter {<string> | none}
      set max-pkt-count <1-maximum>
      set max-pkt-len <64-1534>
      set switch-interface <switch_interface_name>
      set system-interface <system_interface_name>
   end
```

For example:

```
config system sniffer-profile
   edit profile1
      set filter none
      set max-pkt-count 100
      set max-pkt-len 100
      set system-interface mgmt
   end
```

Starting the packet capture

After you create a packet-capture profile, you can start the packet capture.

Using the GUI:

1. Go to *System* > *Packet Capture*.
2. Select .

Using the CLI:

```
execute system sniffer-profile start <profile-name>
```

For example:

```
execute system sniffer-profile start profile1
```

Pausing or stopping the packet capture

A packet capture continues to run until the `max-pkt-cnt` value is reached, or the packet capture is paused or stopped. You can restart a paused packet capture.

Using the GUI:

Go to *System* > *Packet Capture*.

- To pause a running packet capture, select .
- To resume a paused packet capture, select .

Using the CLI:

To pause a running packet capture:

```
execute system sniffer-profile pause <profile_name>
```

To restart a paused packet capture:

```
execute system sniffer-profile start <profile-name>
```

To stop a running packet capture:

```
execute system sniffer-profile stop <profile-name>
```

Displaying or uploading the packet capture

You can display parsed information from the packet capture or upload the .pcap file to a TFTP or FTP server for further analysis.

Using the GUI:

1. Go to *System* > *Packet Capture*.
2. Select .
   The `.pcap` file is saved in your Downloads folder.

Using the CLI:

To display the packet capture from a specific packet-capture profile:

```
get system sniffer-profile capture <profile_name>
```

To upload the `.pcap` file for a specific packet-capture profile to an FTP server:

```
execute system sniffer-profile upload ftp <profile_name> <packet_capture_file_name.pcap>
     <FTP_server_IP_address:<optional_port>>
```

To upload the `.pcap` file for a specific packet-capture profile to a TFTP server:

```
execute system sniffer-profile upload tftp <profile_name> <packet_capture_file_name.pcap>
     <TFTP_server_IP_address:<optional_port>>
```

Deleting the packet-capture file

After you have examined the packet capture, you can manually delete the `.pcap` file. You can only delete the `.pcap` after the packet capture is stopped. You cannot delete the `.pcap` file if the packet capture is paused or running. All `.pcap` files are deleted when you power cycle the switch.

Using the GUI:

1. Go to *System* > *Packet Capture*.
2. Select .

To delete all packet-capture files, select *Select All* and then select *Delete*.

Using the CLI:

```
execute system sniffer-profile delete-capture <profile_name>
```

For example:

```
execute system sniffer-profile delete-capture profile1
```

Debug report

Technical support uses debug reports to diagnose system problems.

Click *Generate Debug Logs* to geneate a detailed debugging report. Click *Download* to download the generated report so that you can send it to technical support. The report is identical to the output of the `diagnose debug report` command.

Fault relay support

Fault relays are available in FortiSwitch Rugged units:

- FSR-112D-POE: normally closed (NC)
- FSR-216F-POE and FSR-424F-POE: normally closed (NC) and normally open (NO)

“Normally” applies to the condition when the relay is not energized. A normally closed relay opens when energized (the switch is operating properly), and a normally open relay closes. When the FortiSwitch Rugged unit loses power, the NC relay contact is in a closed state, and the NO relay contact is in an open state. In addition, the alarm circuit is triggered.

You can change how the ALARM LED functions. Use the following command to have the ALARM LED turn red when only one power supply unit (PSU) is connected:

```
config system global
   set single-psu-fault enable
end
```

By default, the `set single-psu-fault` command is disabled.

Identifying a specific FortiSwitch unit

When you have multiple FortiSwitch units and need to locate a specific switch, use the following command to flash all port LEDs on and off for a specified number of minutes:

```
diagnose switch physical-ports led-flash <disable | time>
```

You can flash the port LEDs for 5, 15, 30, or 60 minutes. After you locate the FortiSwitch unit, you can use `disable` to stop the LEDs from flashing.

NOTE: For the FS-5xx switches, the `diagnose switch physical-ports led-flash` command flashes only the SFP port LEDs, instead of all the port LEDs.

Using the Reset button on FortiSwitch units

Except for the FS-1024D model, all FortiSwitch units have a Reset button. The Reset button is recessed in a small unlabeled hole in the FortiSwitch faceplate, except for the FS-1048E model, which has the Reset button in back of the switch.

To reset the FortiSwitch unit to the factory default configuration, press the Reset button for about 10 seconds and then release it.

Starting in FortiSwitchOS 7.2.7 and 7.4.3, you can use a CLI command to disable the FortiSwitch hardware Reset button while the OS is running. By default, you can use the FortiSwitch hardware Reset button, even while the OS is running.

To disable the FortiSwitch hardware Reset button while the OS is running:

```
config system global
   set reset-button disable
end
```

To enable the FortiSwitch hardware Reset button:

```
config system global
   set reset-button enable
end
```

Amber and red LEDs

Depending on your FortiSwitch model, there are several LEDs on the FortiSwitch faceplate.

- When the PWR, PWR1, or PWR2 LED flashes amber, the power supply has failed.
- When the FAN or ALARM LED is amber, the fan has failed.

- When the ALARM LED is amber on the FSR-112D or FSR-216F-POE model, the PoE power budget threshold has been reached.
- When the ALARM LED is red on the FSR-124D model, the system has a fault, or only one power supply is providing power.

Switch

The following topics provide information about switching functionality:

- Physical ports on page 133
- Interfaces on page 181
- STP on page 240
- Flap guard on page 252
- DHCP snooping on page 255
- IP source guard on page 264
- LLDP-MED on page 167
- ACL on page 271
- IGMP snooping on page 285
- MLD snooping on page 293
- PoE on page 299
- sFlow on page 300
- Mirror on page 302
- VLAN on page 314
- Virtual wires on page 329
- Storm control on page 330
- MAC entries on page 333
- IP-MAC binding on page 335
- QoS on page 336
- Network monitor on page 346
- Configuring security checks on page 351
- Cut-through switching mode on page 354
- Enabling packet forwarding on page 354
- Configuring auto-topology on page 355
- Viewing port statistics on page 356
- DHCP snooping on page 255
- Media Redundancy Protocol on page 358
- Precision Time Protocol on page 362
- High-Availability Seamless Redundancy on page 374
- Parallel Redundancy Protocol on page 381

Physical ports

This section covers how to configure ports;

- Physical port settings on page 134
- Switched interfaces on page 149
- Dynamic MAC address learning on page 149

- Layer-2 table on page 153
- Loop guard on page 153
- TFTP network port on page 154
- TFTP network port on page 154
- Link aggregation groups on page 156
- MCLAG on page 159
- Multi-stage load balance on page 165
- Unicast hashing on page 167

Physical port settings

The following sections describe the configuration settings that are associated with FortiSwitch physical ports:

- Configuring general port settings on page 134
- Configuring flow control, priority-based flow control, and ingress pause metering on page 135
- Auto-module speed detection on page 136
- Setting the port speed (autonegotiation) on page 136
- Setting the port speed (FS-2048F) on page 137
- Configuring power over Ethernet on a port on page 138
- Energy-efficient Ethernet on page 142
- Diagnostic monitoring interface module status on page 143
- Configuring split ports on page 144
- Configuring QSFP low-power mode on page 148
- Configuring physical port loopbacks on page 148

Configuring general port settings

Using the GUI:

1. Go to *Switch* > *Physical Ports*.
2. Select the port to update and then select *Edit*.
3. Enter an optional description of the port in the *Description* field.
4. Select *Up* or *Down* for the *Administrative Status*.
5. Click *Update* to save your changes.

Using the CLI:

```
config switch physical-port
   edit <port_name>
      set status {up | down}
      set description <string>
      set max-frame-size <bytes_int>
   next
end
```

General port settings include:

- `status`—Administrative status of the port
- `description`—Text description for the port

- `max-frame-size`—Maximum frame size in bytes. **NOTE:** For the FS-1xxE and FS-1xxF and FS-110G-FPOE models, the `max-frame-size` command is under `config switch global`.

Refer to the FortiSwitch feature matrix for details about the maximum frame size for each FortiSwitch model.

Configuring flow control, priority-based flow control, and ingress pause metering

Flow control allows you to configure a port to send or receive a “pause frame” (that is, a special packet that signals a source to stop sending flows for a specific time interval because the buffer is full). By default, flow control is disabled on all ports.

```
config switch physical-port
   edit <port_name>
      set flow-control {both | rx | tx | disable}
   end
```

Parameters enable flow control to do the following:

- `rx`—receive pause control frames
- `tx`—transmit pause control frames
- `both`—transmit and receive pause control frames

Priority-based flow control allows you to avoid frame loss by stopping incoming traffic when a queue is congested.

After you enable priority-based flow control, you then configure whether a port sends or receives a priority-based control frame:

```
config switch physical-port
   edit <port_name>
      set priority-based-flow-control enable
      set flow-control {both | rx | tx | disable}
   end
```

When priority-based flow control is disabled, 802.3 flow control can be used.

NOTE: Priority-based flow control does not support half-duplex speed. When FortiSwitch ports are set to autonegotiate the port speed (the default), priority-based flow control is available if the FortiSwitch model supports it. Lossless buffer management and traffic class mapping are not supported.

If you enable flow control to transmit pause control frames (with the `set flow-control tx` command), you can also use ingress pause metering to limit the input bandwidth of an ingress port. Because ingress pause metering stops the traffic temporarily instead of dropping it, ingress pause metering can provide better performance than policing when the port is connected to a server or end station. To use ingress pause metering, you need to set the ingress metering rate in kilobits and set the percentage of the threshold for resuming traffic on the ingress port.

```
config switch physical-port
   edit <port_name>
      set flow-control tx
      set pause-meter-rate <64-2147483647; set to 0 to disable>
      set pause-resume {25% | 50% | 75%}
   next
end
```

For example:

```
config switch physical-port
   edit port29
      set flow-control tx
      set pause-meter-rate 900
      set pause-resume 50%
   next
end
```

Auto-module speed detection

When you enable auto-module speed detection, the system reads information from the module and sets the port speed to the maximum speed that is advertised by the module. If the system encounters a problem when reading from the module, it sets the default speed (default value is platform specific).

When auto-module sets the speed, the system creates a log entry noting this speed.

NOTE: Auto-speed detection is supported on 1/10G ports, but not on higher speed ports (such as 40G).

Setting the port speed (autonegotiation)

By default, all of the FortiSwitch user ports are set to autonegotiate the port speed. You can also manually set the port speed. The port speeds available differ, depending on the port and switch.

NOTE: The `set speed 1000auto` command is required when FN-TRAN-GC is used with a FortiSwitch unit.

Using the GUI:

1. Go to *Switch* > *Physical Ports* and select the port.
2. Select *Edit*.
3. Select *Auto-Negotiation* or the appropriate port speed.
4. Select *Update*.

Using the CLI:

```
config switch physical-port
   edit <port>
      set speed {1000auto | 100full | 100half | 10full | 10half | auto | 10000cr | 10000full
           | 10000sr | 1000full | 40000auto | auto-module}
   end
```

Viewing auto-module configuration

Display the status of auto-module using following command:

```
config switch physical-port
   edit port47
      show
   end
config switch physical-port
   edit "port47"
```

```
      set max-frame-size 16360
      set speed 10000full
   get
      name : port47
      description : (null)
      flow-control : both
      link-status : down
      lldp-transmit : disable
      max-frame-size : 16360
      port-index : 47
      speed : 10000full
      status : up
   end
```

Setting the port speed (FS-2048F)

Ports 1 to port 48 have speeds of 1G, 10G, or 25G. The default speed is 25G.

To set the speed for the FS-2048F ports:

1. Specify the speed for a range of 12 ports on page 137.
2. Enable the auto-module for each port on page 137.

Specify the speed for a range of 12 ports

You can change the speed of 12 ports at the same time.

To specify the speed for a range of 12 ports:

```
config switch phy-mode
   set {port1-port12-phy-mode | port13-port24-phy-mode | port25-port36-phy-mode | port37-
        port48-phy-mode} {1G/10G | 25G}
end
```

For example:

```
config switch phy-mode
   set port1-port12-phy-mode 1G/10G
end
```

Enable the auto-module for each port

For ports set to 1G or 10G with the `config switch phy-mode` command, you can configure the port speed as 1G or 10G using the auto-module. For ports set to 25G with the `config switch phy-mode` command, you can only configure the port speed as 25G using the auto-module.

To enable the auto-module for each port:

```
config switch physical-port
   edit <port_name>
      set speed auto-module
   next
end
```

For example:

```
config switch physical-port
   edit port1
      set lldp-profile "default-auto-isl"
      set speed auto-module
   next
   edit port2
      set lldp-profile "default-auto-isl"
      set speed auto-module
   next
   edit port3
      set lldp-profile "default-auto-isl"
      set speed auto-module
   next
   edit port4
      set lldp-profile "default-auto-isl"
      set speed auto-module
   next
   edit port5
      set lldp-profile "default-auto-isl"
      set speed auto-module
   next
   edit port6
      set lldp-profile "default-auto-isl"
      set speed auto-module
   next
   edit port7
      set lldp-profile "default-auto-isl"
      set speed auto-module
   next
   edit port8
      set lldp-profile "default-auto-isl"
      set speed auto-module
   next
   edit port9
      set lldp-profile "default-auto-isl"
      set speed auto-module
   next
   edit port10
      set lldp-profile "default-auto-isl"
      set speed auto-module
   next
   edit port11
      set lldp-profile "default-auto-isl"
      set speed auto-module
   next
   edit port12
      set lldp-profile "default-auto-isl"
      set speed auto-module
   next
end
```

Configuring power over Ethernet on a port

You can enable PoE, configure dynamic guard band, and set the priority power allocation for a specific port.

The dynamic guard band is set automatically to the expected power of a port before turning on the port. So, when a PoE device is plugged in, the dynamic guard band is set to the maximum power of the device type based on the AF or AT mode. The AF mode DGB is 15.4 W, and the AT mode DGB is 36 W. When the FortiSwitch unit is fully loaded, the dynamic guard band prevents a new PoE device from turning on.

When power to PoE ports is allocated by priority, lower numbered ports have higher priority so that port1 has the highest priority. When more power is needed than is available, higher numbered ports are disabled first.

When power to PoE ports is allocated by first-come, first-served (FCFS), connected PoE devices receive power, but new devices do not receive power if there is not enough power.

If both priority power allocation and FCFS power allocation are selected, the physical port setting takes precedence over the global setting.

When you connect one PoE port to another PoE port, you must connect two PoE switches with fiber. If you have to connect two PoE ports together, you need to disable the PoE function on both ports before inserting the RJ45 cable. For example, use the `set poe-status disable` command under `config switch physical-port` for both ports before connecting them.

This section covers the following topics:

- Enabling or disabling PoE in the GUI on page 139
- Configuring PoE in the CLI on page 139
- Configuring perpetual PoE on page 140
- Determining the PoE power capacity on page 141
- Resetting the PoE power on page 141
- Displaying PoE information on page 141

Enabling or disabling PoE in the GUI

1. Go to *Switch* > *Physical Ports* .
2. Select a port and then select *Edit*.
3. For the POE Status, select *Enable* or *Disable*.
4. Select a power priority for the port. You can select *High Priority*, *Critical Priority*, or *Low Priority*. If there is not enough power, power is allotted first to Critical Priority ports, then to High Priority ports, and then to Low Priority ports.
5. Select *Update*.

Configuring PoE in the CLI

```
config switch physical-port
   edit <port>
      set poe-status {enable | disable}
      set poe-port-mode {IEEE802_3AF | IEEE802_3AT}
      set poe-port-priority {critical-priority | high-priority | low-priority}
      set poe-pre-standard-detect {disable | enable}
   end
```

Starting in FortiSwitchOS 7.2.2, you can select how a FortiSwitch unit with PoE disconnects from a powered device with the `set poe-disconnection-type {AC | DC | DC-delay}` command.

- `AC`—AC disconnect.
- `DC`—DC disconnect.
- `DC-delay`—DC disconnect with an extra 500-millisecond delay.

The new command is available on the following FortiSwitch models:

- FS-224D-FPOE
- FS-224E-POE
- FS-248E-FPOE
- FS-248E-POE
- FS-424E-FPOE
- FS-424E-POE
- FS-M426E-FPOE
- FS-448E-FPOE
- FS-448E-POE

PoE pre-standard detection is a global setting for the following FortiSwitch models: FSR-112D-POE, FS-548D-FPOE, FS-524D-FPOE, FS-108D-POE, FS-224D-POE, FS-108E-POE, FS-108E-FPOE, FS-124E-POE, FS-124E-FPOE, 148F-POE, and 148F-FPOE.

For the other FortiSwitch PoE models, PoE pre-standard detection is set on each port.

Before FortiSwitchOS 7.0.0, `poe-pre-standard-detect` was set to `enable` by default. Starting in FortiSwitchOS 7.0.0, `poe-pre-standard-detect` is set to `disable` by default.

Configuring perpetual PoE

Starting in FortiSwitchOS 7.2.1, some FortiSwitch PoE models provide perpetual PoE so that a FortiSwitch unit has uninterrupted power while restarting. By default, PoE power is not provided while a FortiSwitch unit restarts.

Refer to the FortiSwitch feature matrix for details about which FortiSwitch models support this feature.

To configure perpetual PoE:

```
config switch physical-port
   edit <PoE_port_name>
      set poe-port-power {normal | perpetual | perpetual-fast}
   next
end
```

Variable	Description
normal	PoE power is not provided while a switch restarts.
perpetual	PoE power is provided during a soft reboot (switch is restarted while powered up).
perpetual-fast	PoE power is provided during a hard reboot (the switch's power is physically turned off and then on again).

Determining the PoE power capacity

Using the GUI:

Go to *Switch* > *Physical Ports*. The Power column displays the power capacity for each PoE port.

Using the CLI:

```
get switch poe inline
```

Resetting the PoE power

Using the GUI:

1. Go to *Switch* > *Physical Ports* .
2. Select a port and then select *POE Reset*.
3. In the confirmation dialog box, select *Reset*.

Using the CLI:

```
execute poe-reset <port>
```

Displaying PoE information

Using the GUI:

Go to *Switch* > *Physical Ports* to see information about each PoE port. Hover over the traffic column to get specific values.

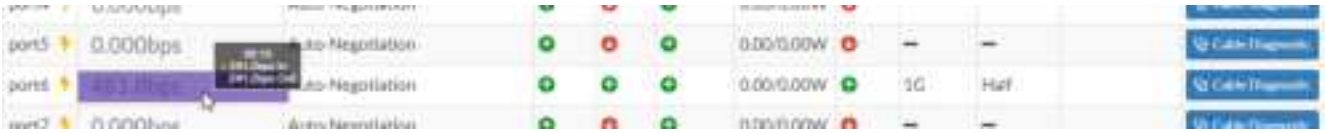

Using the CLI:

```
diagnose switch poe status <port_name>
```

The following example displays the information for port 6:

```
diagnose switch poe status port6

Port(6) Power:4.20W, Power-Status: Delivering Power
Power-Up Mode: Normal Mode
Remote Power Device Type: IEEE802.3AT PD
Power Class: 4
Defined Max Power: 30.0W, Priority:3
Voltage: 54.00V
Current: 71mA
```

Energy-efficient Ethernet

When no data is being transferred through a port, energy-efficient Ethernet (EEE) puts the data link in sleep mode to reduce the power consumption of the FortiSwitch unit. When data flows through the port, the port resumes using the normal amount of power. EEE works over standard twisted-pair copper cables and supports 10 Mbps, 100 Mbps, 1 Gps, and 10 Ge. EEE does not reduce bandwidth or throughput.

If you are using the CLI, you can also specify the number of microseconds that circuits are turned off to save power and the number of microseconds during which no data is transmitted while the circuits that were turned off are being restarted.

In addition, you can use the LLDP 802.3 TLV to advertise the EEE configuration.

NOTE: EEE is not supported on SFP and QSFP modules.

Using the GUI:

1. Go to *Switch* > *Physical Ports*.
2. Select a port and then select *Edit*.
3. Under Energy-Efficient Ethernet, select *Enable*.
4. To save your changes, select *Update*.

To check which ports have EEE enabled, go to *Switch* > *Physical Ports*. A green arrow in the EEE column indicates that EEE is enabled for that port. A red arrow in the EEE column indicates that EEE is disabled for that port.

Using the CLI:

NOTE: When you change the `eee-tx-wake-time` value, the port resets, and the connection is lost briefly.

```
config switch physical-port
   edit <port_name>
      set energy-efficient-ethernet {enable | disable}
      set eee-tx-idle-time <0-2560>
      set eee-tx-wake-time <0-2560>
   end
```

For example, to use EEE on port 7:

```
config switch physical-port
   edit port7
      set energy-efficient-ethernet enable
      set eee-tx-idle-time 500
      set eee-tx-wake-time 200
   end
```

To check that EEE is enabled on port 7:

```
diagnose switch physical-ports eee-status port7
```

To check which ports have EEE enabled:

```
diagnose switch physical-ports eee-status
```

To advertise the EEE configuration in the LLDP 802.3 TLV:

```
config switch lldp profile
   edit <profile_name>
      set 802.3-tlvs eee-config
   next
end
```

To check that the EEE configuration is being advertised:

```
diagnose switch physical-ports eee-status
```

Diagnostic monitoring interface module status

With the diagnostic monitoring interface (DMI), you can view the following information

- Module details (detail)
- Eeprom contents (eeprom)
- Module limits (limit)
- Module status (status)
- Summary information of all a port's modules (summary)

Using the GUI:

Go to *Switch* > *Monitor* > *Modules*.

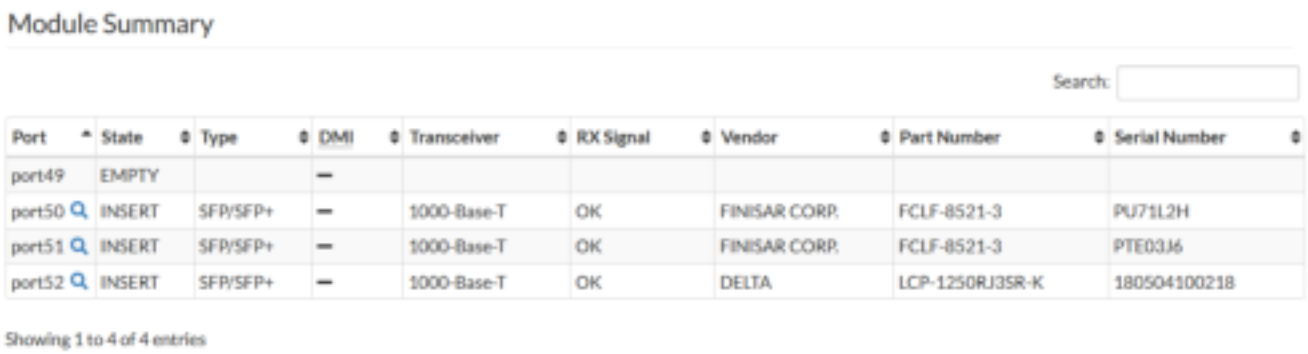

Using the CLI:

Use the following commands to enable or disable DMI status for the port. If you set the status to `global`, the port setting will match the global setting:

```
config switch physical-port
   edit <interface>
      set dmi-status {disable | enable | global}
   end
```

Use the `get switch modules detail/status` command to display DMI information:

```
S148FNTF20000098 # get switch modules detail port50
_____________________________________________________________
Port(port50)
```

```
identifier       SFP/SFP+
connector        Unk(0x00)
transceiver      1000-Base-T
encoding         8B/10B
Length Decode Common
       length_smf_1km  N/A
       length_cable    100 meter
SFP Specific
       length_smf_100m N/A
       length_50um_om2 N/A
       length_62um_om1 N/A
       length_50um_om3 N/A
vendor           FINISAR CORP.
vendor_oid       0x009065
vendor_pn        FCLF-8521-3
vendor_rev
vendor_sn        PU71L2H
manuf_date       08/15/2015
```

The following is an example of the output for the `get switch modules status` command:

```
FS1E48T419000004 # get switch modules status port50
_____________________________________________________________
Port(port50)
temperature      23.957031 C
voltage          3.293100 volts
alarm_flags[0]   0x0000
warning_flags[0] 0x0000
laser_bias[0]    0.761600 mAmps
tx_power[0]      -2.246809 dBm
rx_power[0]      -2.926854 dBm
alarm_flags[1]   0x0000
warning_flags[1] 0x0000
laser_bias[1]    0.755200 mAmps
tx_power[1]      -1.993517 dBm
rx_power[1]      -3.300326 dBm
alarm_flags[2]   0x0000
warning_flags[2] 0x0000
laser_bias[2]    0.761600 mAmps
tx_power[2]      -2.105603 dBm
rx_power[2]      -2.486439 dBm
alarm_flags[3]   0x0000
warning_flags[3] 0x0000
laser_bias[3]    0.748800 mAmps
tx_power[3]      -2.128939 dBm
rx_power[3]      -2.641617 dBm
options          0x000F ( TX_DISABLE TX_FAULT RX_LOSS TX_POWER_LEVEL1 )
options_status   0x0008 ( TX_POWER_LEVEL1 )
```

Configuring split ports

On FortiSwitch models that provide 40G/100G QSFP (quad small form-factor pluggable) interfaces, you can install a breakout cable to convert one 40G/100G interface into four 10G/25G interfaces.

This section covers the following topics:

- Notes on page 145
- Configuring a split port on page 145
- Configuring forward error correction on page 147

Notes

- Splitting ports is supported on the following FortiSwitch models:
  - FS-3032D—Ports 5 to 28 are splittable.
  - FS-3032E—Ports can be split into 4 x 25G when configured in 100G QSFP28 mode or can be split into 4 x 10G when configured in 40G QSFP mode. Use the `set <port-name>-phy-mode disabled command` to disable some 100G ports to allow up to sixty-two 100G/25G/10G ports.
  - FS-524D and FS-524D-FPOE—Ports 29 and 30 are splittable as 4 x 10G.
  - FS-548D and FS-548D-FPOE—Ports 53 and 54 are splittable as 4 x 10G.
  - FS-1024E—Ports 25 and 26 have a maximum speed of 100G; each port can be split into four subports of 25G or 10G.
  - FS-T1024E and FS-T1024F-FPOE—Ports 25 and 26 have a maximum speed of 100G; each port can be split into four subports of 25G or 10G.
  - FS-1048E—In the 4 x 100G configuration, ports 49, 50, 51, and 52 are splittable as 4 x 25G, 4 x 10G, 4 x 1G, or 2 x 50G. Only two of the available ports can be split.
  - FS-1048E—In the 4 x 4 x 25G configuration, ports 49, 50, 51, and 52 are splittable as 4 x 4 x 25G or 2 x 50G. All four ports can be split, but ports 47 and 48 are disabled.
  - FS-1048E—In the 6 x 40G configuration, ports 49, 50, 51, 52, 53, 54 are splittable as 4 x 10G or 4 x 1G.

  Use the `set port-configuration ?` command to check which ports are supported for each model.
- Currently, the maximum number of ports supported in software is 64 (including the management port). Therefore, only 10 QSFP ports can be split. This limitation applies to all of the models, but only the 3032D, the 3032E, and the 1048E models have enough ports to encounter this limit.
- Starting in FortiOS 6.2.0, splitting ports is supported in FortiLink mode (that is, the FortiSwitch unit managed by a FortiGate unit).
- Use `10000full` for the general 10G interface configuration. If that setting does not work, use `10000cr` for copper connections (with copper cables such as 10GBASE-CR) or use `10000sr` for fiber connections (fiber optic transceivers such as 10GBASE-SR/-LR/-ER/-ZR).

Configuring a split port

Use the following commands to configure a split port:

```
config switch phy-mode
   set port-configuration {default | disable-port54 | disable-port41-48 | 4x100G | 6x40G |
        4x4x25G}
   set {<port-name>-phy-mode <single-port>} {4x25G | 4x10G | 4x1G | 2x50G}
   ...
   (one entry for each port that supports split port)
end
```

The following settings are available:

- `disable-port54`—For 548D and 548D-FPOE, only port53 is splittable; port54 is unavailable.
- `disable-port41-48`—For 548D and 548D-FPOE, port41 to port48 are unavailable, but you can configure port53 and port54 in split-mode.
- `4x100G`—For 1048E, enable the maximum speed (100G) of ports 49 through 52. Ports 53 and 54 are disabled.

- `6x40G`—For 1048E, enable the maximum speed (40G) of ports 49 through 54.
- `4x4x25G`—For 1048E, enable the maximum speed (100G) of ports 49 through 52; each split port has a maximum speed of 25G. Ports 47 and 48 are disabled.
- `single-port`—Use the port at the full base speed without splitting it.
- `4x25G`—For 100G QSFP only, split one port into four subports of 25 Gbps each.
  NOTE: For the FS-T1024E, FS-T1024F-FPOE, and FS-1024E models, the auto-module selects the correct speed for the subports. If you insert a 100G QSFP28 module, the subports are automatically changed to 4x25G. If you insert a 40G QSFP+ module, the subports are automatically changed to 4x10G.
- `4x10G`—For 40G or 100G QSFP only, split one port into four subports of 10Gbps each.
- `4x1G`—For 40G or 100G QSFP only, split one port into four subports of 1 Gbps each.
- `2x50G`—For 100G QSFP only, split one port into two subports of 50 Gbps each.

In the following example, a FortiSwitch 3032D model is configured with ports 10, 14, and 28 set to 4x10G:

```
config switch phy-mode
   set port5-phy-mode 1x40G
   set port6-phy-mode 1x40G
   set port7-phy-mode 1x40G
   set port8-phy-mode 1x40G
   set port9-phy-mode 1x40G
   set port10-phy-mode 4x10G
   set port11-phy-mode 1x40G
   set port12-phy-mode 1x40G
   set port13-phy-mode 1x40G
   set port14-phy-mode 4x10G
   set port15-phy-mode 1x40G
   set port16-phy-mode 1x40G
   set port17-phy-mode 1x40G
   set port18-phy-mode 1x40G
   set port19-phy-mode 1x40G
   set port20-phy-mode 1x40G
   set port21-phy-mode 1x40G
   set port22-phy-mode 1x40G
   set port23-phy-mode 1x40G
   set port24-phy-mode 1x40G
   set port25-phy-mode 1x40G
   set port26-phy-mode 1x40G
   set port27-phy-mode 1x40G
   set port28-phy-mode 4x10G
end
```

In the following example, a FortiSwitch 1048E model is configured so that each port is split into four subports of 25 Gbps each.

```
config switch phy-mode
   set port-configuration 4x4x25G
   set port49-phy-mode 4x25G
   set port50-phy-mode 4x25G
   set port51-phy-mode 4x25G
   set port52-phy-mode 4x25G
end
```

In the following example, the FS-T1024E or FS-T1024F-FPOE or FS-1024E model is configured so that each port is split into four subports. Use the auto-module to select the correct speed. If you insert a 100G QSFP28 module, the subports

are automatically changed to 4x25G. If you insert a 40G QSFP+ module, the subports are automatically changed to 4x10G.

```
config switch phy-mode
   set port25-phy-mode 4x25G
   set port26-phy-mode 4x25G
end
config switch physical-port
   edit "port25.3"
      set lldp-profile "default-auto-isl"
      set speed auto-module
   next
end
```

To configure one of the split ports, use the notation ".x" to specify the split port:

```
config switch physical-port
   edit "port1"
      set lldp-profile "default-auto-isl"
      set speed 40000full
   next
   edit "port2"
      set lldp-profile "default-auto-isl"
      set speed 40000full
   next
   edit "port3"
      set lldp-profile "default-auto-isl"
      set speed 40000full
   next
   edit "port4"
      set lldp-profile "default-auto-isl"
      set speed 40000full
   next
   edit "port5.1"
      set speed 10000full
   next
   edit "port5.2"
      set speed 10000full
   next
   edit "port5.3"
      set speed 10000full
   next
   edit "port5.4"
      set speed 10000full
   next
end
```

Configuring forward error correction

You can set the forward error correction (FEC) state on the 25G ports of the FS-1048E and FS-3032E models with the following commands:

```
config switch physical-port
   edit <split_port_name>
      set fec-state {cl74 | detect-by-module | disabled}
   end
```

You can set the FEC state on the 100G ports of the FS-1048E and FS-3032E models with the following commands:

```
config switch physical-port
   edit <split_port_name>
      set fec-state {cl91 | detect-by-module | disabled}
   end
```

Use the `diagnose switch physical-ports list <port_name>` command to verify the current FEC state.

Starting in FortiSwitchOS 6.4.0, `cl74` is enabled as the default setting for 25G ports, and `cl91` is enabled as the default setting for 100G ports.

Starting in FortiSwitchOS 7.0.0, by default, the 25G and 100G ports of the FS-1048E and FS-3032E models now automatically detect whether FEC is supported by the module.

Starting in FortiSwitchOS 7.0.1, you can use the `set fec-state detect-by-module` command to allow split ports of the FS-1048E and FS-3032E models to automatically detect whether forward error correction (FEC) is supported by the module.

Configuring QSFP low-power mode

On FortiSwitch models with QSFP (quad small form-factor pluggable) ports, you can enable or disable the low-power mode with the following CLI commands:

```
config switch physical-port
   edit <port_name>
      set qsfp-low-power-mode {enabled | disabled}
   end
```

For example:

```
config switch physical-port
   edit port12
      set qsfp-low-power-mode disabled
   end
```

Configuring physical port loopbacks

You can use the CLI to loop a physical port back on itself, either locally or remotely:

- The local loopback is a physical-layer loopback. If the hardware does not support a physical-layer loopback, a MAC-address loopback is used instead.
- The remote loopback is a physical-layer lineside loopback.

By default this feature is disabled.

To configure a physical port loopback:

```
config switch physical-port
   edit <port_name>
      set loopback {disable | local | remote}
   next
end
```

Switched interfaces

Default configuration will suffice for regular switch ports. By default, VLAN is set to 1, STP is enabled, and all other optional capabilities are disabled.

You can configure optional capabilities such as STP, sFlow , Port security, and Private VLANs. These capabilities are covered in subsequent sections of this document.

Using the GUI:

1. Go to *Switch* > *Interfaces*.
2. Select one or more interfaces to update and select *Edit*.
   If you selected more than one port, the port names are displayed in the name field, separated by commas.
3. Enter new values as required for the *Native VLAN* and *Allowed VLANs* fields.
4. Select *OK* to save your changes.

Using the CLI:

```
config switch interface
   edit <port>
      set native-vlan <vlan>
      set allowed-vlans <vlan> [<vlan>] [<vlan> - <vlan>]
      set untagged-vlans <vlan> [<vlan>] [<vlan> - <vlan>]
      set stp-state {enabled | disabled}
      set edge-port {enabled | disabled}
```

Viewing interface configuration

Using the GUI:

Go to *Switch* > *Interfaces*.

Using the CLI:

```
show switch interface <port>
```

Display port settings using following command:

```
config switch interface
   edit <port>
      get
```

Dynamic MAC address learning

You can enable or disable dynamic MAC address learning on a port. The existing dynamic MAC entries are deleted when you change this setting. If you disable MAC address learning, you can set the behavior for an incoming packet with an unknown MAC address (to drop or forward the packet).

You can limit the number of learned MAC addresses on an interface or VLAN. The limit ranges from 1 to 128. If the learning limit is set to zero (the default), no limit exists. When the limit is exceeded, the FortiSwitch unit adds a warning to the system log.

Configuring dynamic MAC address learning

Use the following CLI commands to configure dynamic MAC address learning:

```
config switch physical-port
   edit <port>
      set l2-learning (enable | disable)
      set l2-sa-unknown (drop | forward)
   end
config switch interface
   edit <port>
      set learning-limit <0-128>
   end
config switch vlan
   edit <VLAN_ID>
      set learning {enable | disable}
      set learning-limit <0-128>
   end
```

NOTE: If you enable 802.1X MAC-based authorization on a port, you cannot change the `l2-learning` setting.

Changing when MAC addresses are deleted

By default, each learned MAC address is deleted after 300 seconds. The value ranges from 10 to 1,000,000 seconds. Set the value to zero to not delete learned MAC addresses.

Use the following command to change this value:

```
config switch global
   set mac-aging-interval 200
end
```

Logging dynamic MAC address events

By default, dynamic MAC address events are not logged. When you enable logging for an interface, the following events are logged:

- When a dynamic MAC address is learned
- When a dynamic MAC address is moved
- When a dynamic MAC address is deleted

NOTE: Some dynamic MAC address events might take a long time to be logged. If too many events happen within a short period of time, some events might not be logged.

To enable the logging of dynamic MAC address events:

```
config switch interface
   edit <interface_name>
      set log-mac-event enable
   end
```

To view the log entries:

```
execute log display
```

Using the learning-limit violation log

If you want to see the first MAC address that exceeded a learning limit for an interface or VLAN, you can enable the learning-limit violation log for a FortiSwitch unit. Only one violation is recorded per interface or VLAN.

By default, the learning-limit violation log is disabled. The most recent violation that occurred on each interface or VLAN is logged. After that, no more violations are logged until the log is reset for the triggered interface or VLAN. Only the most recent 128 violations are displayed in the console.

Using the GUI:

1. Go to *Switch* > *MAC Limit*.
2. Enable or disable *Enable Learning Limit Violation recording globally*.

Using the CLI:

```
config switch global
   set log-mac-limit-violations {enable | disable}
end
```

NOTE: The `set log-mac-limit-violations` command is only displayed if your FortiSwitch model supports it.

To view the content of the learning-limit violation log, use one of the following commands:

- `get switch mac-limit-violations all`—to see the first MAC address that exceeded the learning limit on any interface or VLAN. An asterisk by the interface name indicates that the interface-based learning limit was exceeded. An asterisk by the VLAN identifier indicates the VLAN-based learning limit was exceeded.
- `get switch mac-limit-violations interface <interface_name>`—to see the first MAC address that exceeded the learning limit on a specific interface.
- `get switch mac-limit-violations vlan <VLAN_ID>`—to see the first MAC address that exceeded the learning limit on a specific VLAN. This command is only displayed if your FortiSwitch model supports it.

To reset the learning-limit violation log, use one of the following commands:

- `execute mac-limit-violation reset all`—Use this command to clear all learning-limit violation logs or to clear the shutdown state of a port caused by the `set learning-limit-action shutdown` command.
- `execute mac-limit-violation reset interface <interface_name>`—Use this command to clear the learning-limit violation log for a specific interface or to clear the shutdown state of a port caused by the `set learning-limit-action shutdown` command.
- `execute mac-limit-violation reset vlan <VLAN_ID>`—Use this command to clear the learning-limit violation log for a specific VLAN.

You can also specify how often the learning-limit violation log is reset. When the mac-violation-timer expires, it will also clear the shutdown state of a port caused by the `set learning-limit-action shutdown` command.

To specify how often the learning-limit violation log is rest:

```
config switch global
   set log-mac-limit-violations enable
   set mac-violation-timer <0-1500>
end
```

For example:

```
config switch global
   set log-mac-limit-violations enable
   set mac-violation-timer 60
end
```

Configuring learning-limit violation actions

Starting in FortiSwitchOS 7.0.2, when the MAC learning limit is exceeded, you can specify that the interface that it is configured on is disabled (`set learning-limit action shutdown`) or that no action is taken (`set learning-limit action none`). The `learning-limit-action` applies only to physical switch port interfaces, not to trunks or VLANs.

To configure the action for learning-limit violations:

```
config switch global
   set log-mac-limit-violations enable
end

config switch interface
   edit <port_name>
      set learning-limit <1-128>
      set learning-limit-action {none | shutdown}
   next
end
```

After shutting down the port with the `set learning-limit-action shutdown` command, you can bring it back up in two ways:

- With the `execute mac-limit-violation reset {interface <port_name> | all}` command.
- With the `set mac-violation-timer <integer>` command (under `config switch global`).

Starting in FortiSwitchOS 7.0.2, you can configure an SNMP trap so that you receive a message when the MAC learning limit is exceeded.

To configure the SNMP trap for learning-limit violations:

```
config switch global
   set log-mac-limit-violations enable
end

config system snmp community
   edit <index_number>
      set events llv
   next
end
```

Layer-2 table

FortiSwitchOS uses the layer-2 table to store static MAC addresses and dynamic MAC addresses. Starting in FortiSwitchOS 7.0.0, you can use the CLI to control whether the size of the layer-2 table is checked and how often. By default, this feature is disabled. When you enable this feature, two checks are performed for each platform:

- The first time that the layer-2 table size is more than 75-percent full for each platform, FortiSwitchOS adds a warning to the system log.
- The first time that the layer-2 table size is less than 70-percent full for each platform, FortiSwitchOS adds a warning to the system log.

When you enable this feature, FortiSwitchOS checks the layer-2 table every 2 minutes. You can change how often the layer-2 table is checked to 5-86,400 seconds.

To enable this feature and specify the interval:

```
config switch global
   set l2-memory-check enable
   set l2-memory-check-interval <number_of_seconds>
end
```

For example:

```
config switch global
   set l2-memory-check enable
   set l2-memory-check-interval 1000
end
```

To disable this feature:

```
config switch global
   set l2-memory-check disable
end
```

Loop guard

NOTE: This feature is different from STP loop protection.

A loop in a layer-2 network results in broadcast storms that have far-reaching and unwanted effects. Loop guard helps to prevent loops. When loop guard is enabled on a switch port, the port monitors its subtending network for any downstream loops.

The loop guard feature is designed to work in concert with STP rather than as a replacement for STP. Each port that has loop guard enabled will periodically broadcast loop guard data packets (LGDP) packets to its network. If a broadcast packet is subsequently received by the sending port, a loop exists downstream.

You can also have the port check for a high rate of MAC address moves per second, which indicates a physical loop only when the rate exceeds the threshold for 6 consecutive seconds.

NOTE: If a port detects a loop, the system takes the port out of service to protect the overall network. The port returns to service after a configured timeout duration. If the timeout value is zero, you must manually reset the port.

By default, loop guard is disabled on all ports. When loop guard is enabled, the default `loop-guard-timeout` is 45 minutes, and the default `loop-guard-mac-move-threshold` is 0, which means that the traditional loop guard is used instead of the MAC-move loop guard.

Configuring loop guard

Using the GUI:

1. Go to *Switch* > *Interfaces*.
2. Select one or more interfaces to update and then select *Edit*.
   If you selected more than one port, the port names are displayed in the name field, separated by commas.
3. Under *Loop Guard*, select *Enable*.
4. Select *OK* to save your changes.

Using the CLI:

```
config switch interface
   edit port <number>
      set loop-guard <enabled | disabled>
      set loop-guard-timeout <0-120 minutes>
      set loop-guard-mac-move-threshold <0-100 MAC address moves per second>
```

When loop guard takes a port out of service, the system creates the following log messages:

```
Loop Guard: loop detected on <port_name>. Shutting down <port_name>
```

Use the following command to reset a port that detected a loop:

```
execute loop-guard reset <port>
```

Viewing the loop guard configuration

Using the GUI:

Go to *Switch* > *Monitor* > *Loop Guard*.

Using the CLI:

```
diagnose loop-guard status
```

TFTP network port

When you power on the FortiSwitch unit, the BIOS performs basic device initialization. When this activity is complete, and before the OS starts to boot, you can click any key to bring up the boot menu.

From the menu, click the "I" key to configure TFTP settings. With newer versions of the BIOS, you can specify the network port (where you have connected your network cable). If you are not prompted to specify the network port, you must connect your network cable to the default network port:

- If the switch model has a WAN port, the WAN port is the network port.
- If the switch has no WAN port, the highest port number is the network port.

Cable diagnostics

NOTE: There are some limitations for cable diagnostics on the FS-108E, FS-124E, FS-108E-POE, FS-108E-FPOE, FS-124E-POE, FS-124E-FPOE, FS-148E, and FS-148E-POE models:

- Crosstalk cannot be detected.
- There is a 5-second delay before results are displayed.
- The value for the cable length is inaccurate.
- The results are inaccurate for open and short cables.

You can check the state of cables connected to a specific port. The following pair states are supported:

- Open
- Short
- Ok
- Open_Short
- Unknown
- Crosstalk

If no cable is connected to the specific port, the state is Open, and the cable length is 0 meters.

To see which models support this feature, refer to the FortiSwitch feature matrix.

Using the GUI:

1. Go to *Switch* > *Physical Ports*.
2. Select *Cable Diagnostic* for the appropriate port.
3. Select *Continue* to start the cable diagnostics.
   NOTE: Running cable diagnostics on a port that has the link up will interrupt the traffic for several seconds.
4. Select *Back to Physical Ports* to close the Cable Diagnostics window.

Using the CLI:

Use the following command to run a time domain reflectometry (TDR) diagnostic test on cables connected to a specific port:

```
diagnose switch physical-ports cable-diag <physical port name>
```

NOTE: Running cable diagnostics on a port that has the link up will interrupt the traffic for several seconds.

For example:

```
# diagnose switch physical-ports cable-diag port1

port1: cable (4 pairs, length +/- 10 meters)
pair A Open, length 0 meters
pair B Open, length 0 meters
pair C Open, length 0 meters
pair D Open, length 0 meters
```

Use the following command to check the medium dependent interface crossover (MDI-X) interface status for a specific port:

```
diagnose switch physical-ports mdix-status <physical port name>
```

For example:

```
# diagnose switch physical-ports mdix-status port1

port1: MDIX(Crossover)
```

Link aggregation groups

This section provides information on how to configure a link aggregation group (LAG). For LAG control, the FortiSwitch unit supports the industry-standard Link Aggregation Control Protocol (LACP). The FortiSwitch unit supports LACP in active and passive modes. In active mode, you can optionally specify the minimum and maximum number of active members in a trunk group.

If the trunk is in LACP mode and has ports with different speeds, the ports of the same negotiated speed are grouped in an aggregator.

If multiple aggregators exist, one and only one of the aggregators is used by the trunk.

You can use the CLI to specify how the aggregator is selected:

- When the `aggregator-mode` is set to `bandwidth`, the aggregator with the largest bandwidth is selected. This mode is the default.
- When the `aggregator-mode` is set to `count`, the aggregator with the largest number of ports is selected.

The FortiSwitch unit supports flap-guard protection for switch ports in a LAG.

Starting in FortiSwitchOS 7.4.0, LACP fallback mode is supported in the CLI. LACP fallback mode allows a selected port to stay up so that a device not running LACP can still connect to the network. LACP fallback mode is useful if you have a preboot execution environment (PXE) and need to download an image from the network before running LACP.

When you select the fallback port for a switch trunk, the aggregate interface will use the LACP fallback mode if the trunk does not receive any LACP protocol data units (PDUs). The fallback port is set to up, and all other ports are blocked. When the trunk starts receiving LACP PDUs again, the switch trunk changes from fallback mode to LACP.

When the switch trunk is running LACP and stops receiving LACP PDUs:

- There is a 90-second delay before LACP fallback mode if the `lacp-speed` for the switch trunk is set to `slow`.
- There is a 30-second delay before LACP fallback mode if the `lacp-speed` for the switch trunk is set to `fast`.

The following are the requirements and limitations for LACP fallback mode:

- If you are using MCLAG, do not configure fallback mode on more than one MCLAG switch. If you configure fallback mode on both MCLAG switches, the `diagnose switch mclag peer-consistency-check` command will report it as a mismatch.
- You cannot use fallback mode with the `min_bundle` or `max_bundle` setting.
- You cannot use fallback mode with an MCLAG split-brain state.

Configuring the trunk and LAG ports

It is important to configure the trunk to prevent loops.

Using the GUI:

1. Go to *Switch* > *Trunks* and select *Add Trunk*.
2. Give the trunk an appropriate name.
3. For the mode, select *Static*, *LACP Active*, *LACP Passive*, or *Fortinet Trunk*.
4. Add the required ports to the *Included* list.
5. Select *Create*.

Using the CLI:

```
config switch trunk
   edit <trunk name>
      set aggregator-mode {bandwidth | count}
      set description <description_string>
      set members <ports>
      set mode {lacp-active | lacp-passive | static}
      set member-withdrawal-behavior {block | forward}
      set lacp-speed {fast | slow}
      set bundle [enable|disable]
         set min_bundle <integer>
         set max_bundle <integer>
      set port-selection-criteria
           {src-ip | src-mac | dst-ip |dst-mac | src-dst-ip |src-dst-mac}
   end
end
```

The `set auto-isl` command (under `config switch trunk`) is used when a trunk is automatically managed by the system. Do not set this command on trunks that you want to manage, for example, when the FortiSwitch unit is in standalone mode without auto-topology (`config switch auto-network`) enabled.

Example configuration

The following is an example CLI configurations for trunk/LAG ports:

Trunk/LAG ports

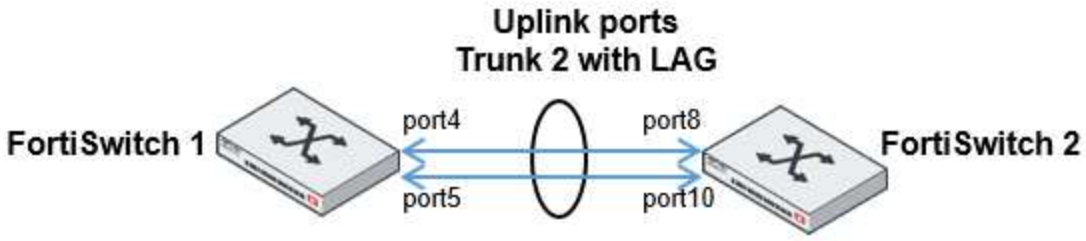

To configure Trunk 2 on FortiSwitch 1:

1. Configure the trunk 2 interface and assign member ports as a LAG group:

```
config switch trunk
   edit trunk2
      set members "port4" "port5"
      set description test
      set mode lacp-passive
      set port-selection criteria src-dst-ip
   end
end
```

2. Configure the trunks to allow the VLANs:

```
config switch interface
   edit trunk2
      set allowed-vlans 2,8
   next
end
```

To configure Trunk 2 on FortiSwitch 2:

1. Configure the trunk 2 interface and assign member ports as a LAG group:

```
config switch trunk
   edit trunk2
      set members "port8" "port10"
      set description test
      set mode lacp-active
      set port-selection criteria src-dst-ip
   end
end
```

2. Configure the trunks to allow the VLANs:

```
config switch interface
   edit trunk2
      set allowed-vlans 2,8
```

```
        next
    end
```

Configuring LACP fallback mode

To configure LACP fallback mode:

```
config switch trunk
   edit <trunk_name>
      set mode {lacp-active | lacp-passive}
      set fallback-port <port_name>
   next
end
```

For example:

```
config switch trunk
   edit LACPtrunk
      set mode lacp-active
      set fallback-port port5
   next
end
```

Checking the trunk configuration

Using the GUI:

Go to *Switch* > *Trunks* or *Switch* > *Monitor* > *Trunks*.

Using the CLI:

```
diagnose switch trunk list
```

MCLAG

A link aggregation group (LAG) provides link-level redundancy. A multichassis LAG (MCLAG) provides node-level redundancy by grouping two FortiSwitch models together so that they appear as a single switch on the network. If either switch fails, the MCLAG continues to function without any interruption, increasing network resiliency and eliminating the delays associated with the Spanning Tree Protocol (STP).

This section covers the following topics:

- Notes on page 160
- Example configuration on page 161
- Detecting a split-brain state on page 162
- Viewing the configured trunk on page 164
- Configuring an MCLAG with IGMP snooping on page 164
- Configuring an MCLAG with MLD snooping on page 164

Notes

For static routes in standalone, MCLAG, or layer-3 MCLAG network topologies, Fortinet recommends using a link monitor or BFD to detect whether the gateway is available.

- When you form a MCLAG from two switches, the trunk name must be the same in each switch configuration.
- When min_bundle or max_bundle is combined with MCLAG, the bundle limit properties are applied only to the local aggregate interface.
- Fortinet recommends that both peer switches be of the same hardware model and same software version. Mismatched configurations might work but are unsupported.
- There is a maximum of two FortiSwitch units per MCLAG peer group.
- Starting in FortiSwitchOS 3.6.4, by default, the MCLAG can use the STP.
- To use static MAC addresses within a MCLAG, you need to configure MAC addresses on both switches that form the LAG.
- When you run an MCLAG, Fortinet recommends but does not require that peers use the same hardware and software versions. Some hosts might not be dual-homed supported when MCLAG peers have different hardware; administrators need to size the layer-2 network to the MCLAG peer with the lowest capacity.
- From the STP tree's point of view, the MCLAG switches should not present themselves differently as a single MCLAG dual-homed virtual switch (accessed through an MCLAG trunk) and as a pair of STP running switches (accessed through asymmetric individual ports, typically in a ring topology). For example, the spanning tree with its root bridge outside of the MCLAG switches cannot connect to it through a dual-homed trunk on one side, while connecting to the MCLAG switches with asymmetric ports at the same time. Such configurations present a mixed view of the MCLAG switches to the STP instance and are not supported.
- Starting in FortiSwitchOS 7.4.4, the way that MCLAGs handle static MAC addresses has been improved. When an MCLAG trunk goes down, the static MAC addresses are removed from the hardware, and the traffic that was going to the static MAC addresses will flood over the interchassis link (ICL). When the MCLAG trunk goes up, the static MAC addresses are added to the hardware again.

Example configuration

The following is an example CLI configurations for a MCLAG:

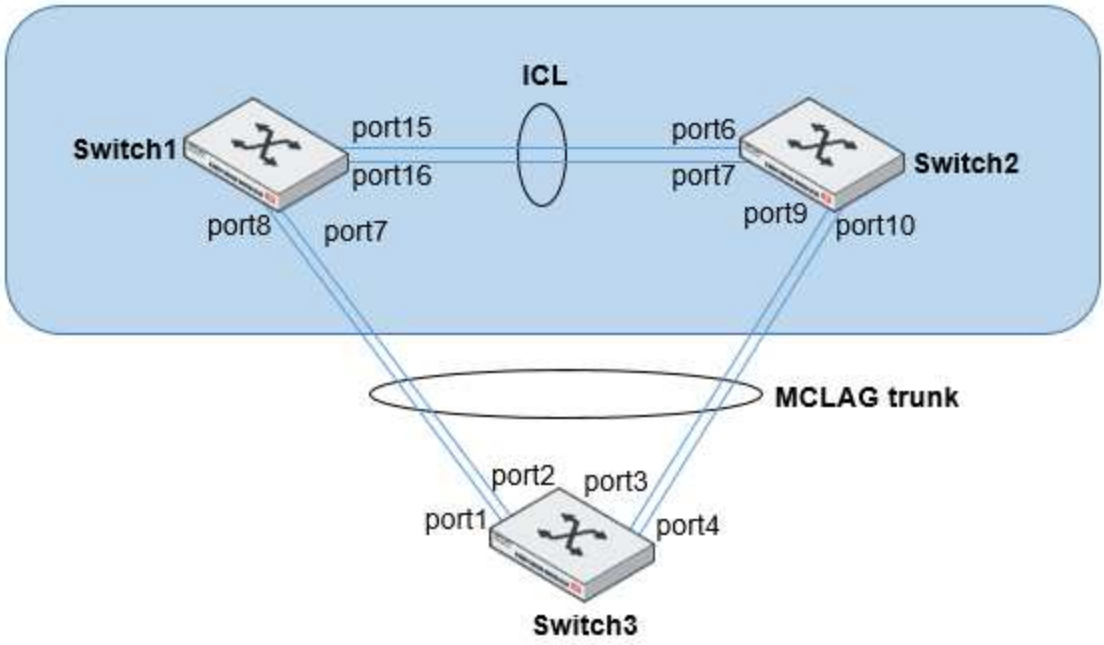

1. Create a LAG by configuring the ports for Switch1:

```
config switch trunk
   edit "MCLAG-ICL-trunk"
      set mode lacp-active
      set mclag-icl enable
      set members "port15" "port16"
   next
end
```

2. Set up the MCLAG for Switch1:

```
config switch trunk
   edit "first-mclag"
      set mode lacp-active
      set mclag enable
      set members "port7" "port8"
   next
end
```

3. Create a LAG by configuring the ports for Switch2:

```
config switch trunk
   edit "MCLAG-ICL-trunk"
      set mode lacp-active
      set mclag-icl enable
      set members "port6" "port7"
   next
end
```

4. Set up the MCLAG for Switch2:

```
config switch trunk
   edit "first-mclag"
      set mode lacp-active
      set mclag enable
      set members "port9" "port10"
   next
end
```

5. Set up the dual-homed trunk for Switch3:

```
config switch trunk
   edit "dht"
      set mode lacp-active
      set members "port1" "port2" "port3" "port4"
   next
end
```

Detecting a split-brain state

A split-brain state occurs when communication between the MCLAG peers over the ICL trunk is lost while both switches are still processing traffic, which could cause network issues. Enabling split-brain detection prevents this condition by making one MCLAG peer go dormant, while the other remains active and forwarding traffic. When communication over the ICL trunk is re-established, the dormant MCLAG peer becomes active again. Enabling split-brain detection does not prevent either MCLAG peer from going down. By default, split-brain detection is disabled.

Starting in FortiSwitchOS 6.2.2, you can use the CLI to detect when an MCLAG is in a split-brain state when the MCLAG ICL trunk is down. When the LACP is up again, the MCLAG trunk is reestablished. You can use this command in both one-tier and two-tier MCLAG topologies.

Starting in FortiSwitchOS 7.0.1, you can use the `set mclag-split-brain-priority` command to specify which switch goes dormant when the split-brain state occurs by setting the priority of each switch. The priority can be 0-100 and is 50 by default. The switch peer with the lowest priority value goes dormant when the split-brain state occurs. If both switch peers have the same priority, the switch with the lowest numerical MAC address goes dormant when the split-brain state occurs.

Starting in FortiSwitchOS 7.0.1, you can enable the `set mclag-split-brain-all-ports-down` command to force the switch going dormant to shut down all ports before going dormant. The state of the ICL trunk ports is not changed. By default, this option is disabled.

Starting in FortiSwitchOS 7.4.1, FortiSwitchOS can distinguish between the ICL being down and a peer switch being down or getting restarted. When a peer switch is down or restarted, the other switch does not mistakenly detect a split-brain state and shut down all ports. This feature requires at least one dual-homed FortiSwitch unit connected to the MCLAG peer group. For example:

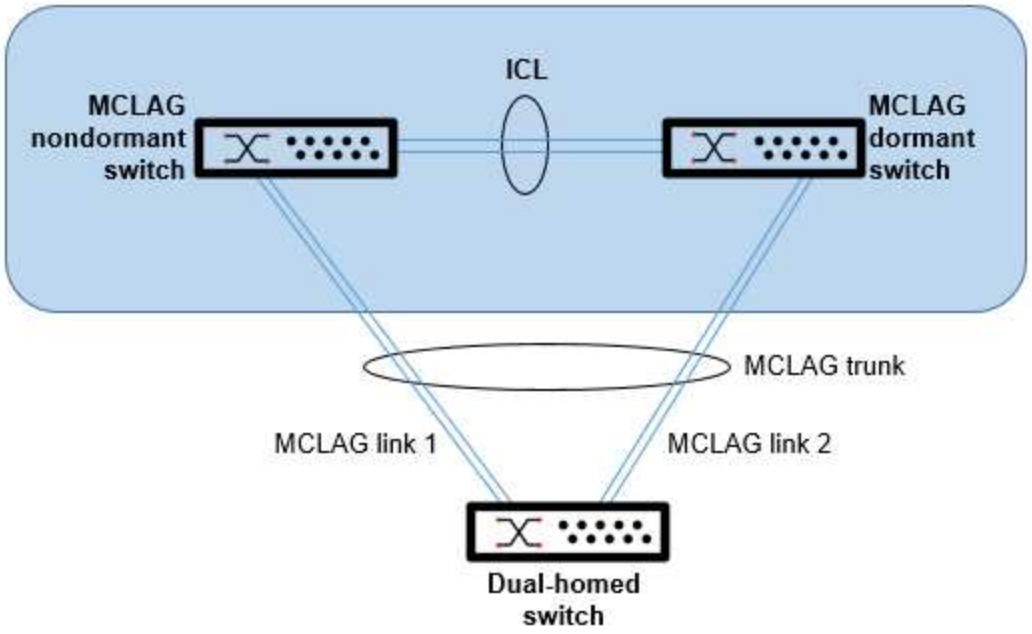

To detect a split-brain state:

1. Configure the detection of the split-brain state for Switch1:

```
config switch global
   set mclag-split-brain-detect enable
   set mclag-split-brain-all-ports-down {enable | disable}
   set mclag-split-brain-priority <0-100>
   set mclag-peer-info-timeout <30-600>
end
```

2. Configure the detection of the split-brain state for Switch2:

```
config switch global
   set mclag-split-brain-detect enable
   set mclag-split-brain-all-ports-down {enable | disable}
   set mclag-split-brain-priority <0-100>
   set mclag-peer-info-timeout <30-600>
end
```

3. Set up the dual-homed trunk for Switch3. **NOTE:** You must include the `set mclag enable` command on the dual-homed trunk.

```
config switch trunk
   edit "dht"
      set mode lacp-active
      set mclag enable
      set members "port1" "port2" "port3" "port4"
```

```
      next
    end
```

NOTE:

- You must configure `set mclag-split-brain-detect enable` on both MCLAG peer switches.
- Enabling split-brain detection can cause some traffic loss while the LACP is renegotiated.
- You can configure split-brain detection for multiple MCLAG pairs, but only one split-brain failure in a system is supported. You must fix the split-brain failure before proceeding.

Viewing the configured trunk

Using the GUI:

Go to *Switch* > *Monitor* > *Trunks*.

Using the CLI:

```
diagnose switch mclag icl
diagnose switch mclag list
```

Configuring an MCLAG with IGMP snooping

For IGMP snooping to work correctly in an MCLAG, you need to enable the `set mclag-igmpsnooping-aware` command on all FortiSwitch units in the network topology and enable the `set igmp-snooping-flood-reports` and `set mcast-snooping-flood-traffic` commands on each MCLAG core FortiSwitch unit. For example:

```
config switch global
   set mac-aging-interval 600
   set mclag-igmpsnooping-aware enable
   config port-security
      set max-reauth-attempt 3
   end
end
config switch interface
   edit "D483Z15000094-0"
      set native-vlan 4094
      set allowed-vlans 1-4094
      set dhcp-snooping trusted
      set stp-state disabled
      set edge-port disabled
      set igmp-snooping-flood-reports enable
      set mcast-snooping-flood-traffic enable
      set snmp-index 58
   next
end
```

Configuring an MCLAG with MLD snooping

For MLD snooping to work correctly in an MCLAG, you need to enable the `set mld-snooping-flood-reports` and `set mcast-snooping-flood-traffic` commands on each MCLAG core FortiSwitch unit.

Multi-stage load balance

You can use a FortiSwitch unit to configure multi-stage load balancing on a set of FortiGate units. This capability allows you to scale security processing while maintaining a simple basic architecture. This configuration is commonly referred to a “firewall sandwich.”

Because the FortiGate unit provides session-aware analysis, the load distribution algorithm must be symmetric (traffic for a given session, in both directions, must all traverse the same FortiGate unit).

For larger scale deployment, the topology uses multiple layers of load distribution to allow for far larger numbers of FortiGate devices.

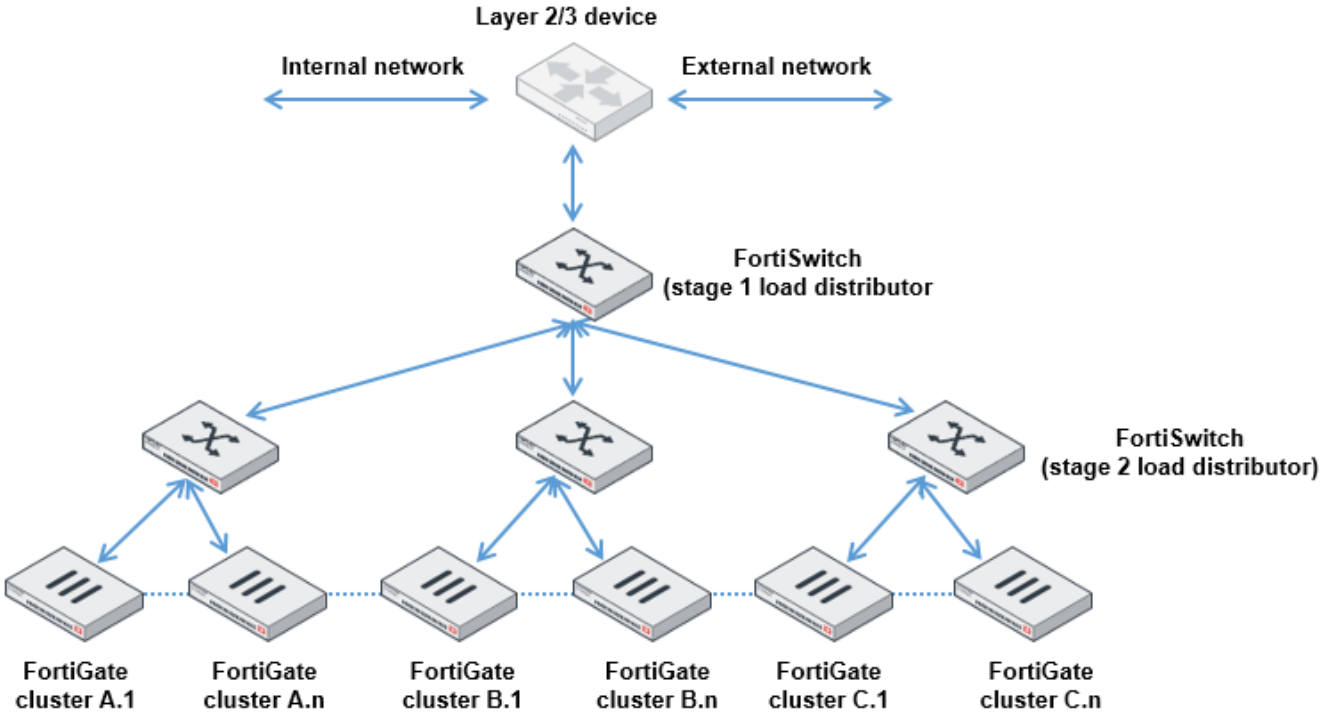

The hash at the first and second stages must be symmetric. The two stages must provide different hashing results.

Configuring the trunk ports

Use the following commands to configure the trunk members and set the port-selection criteria:

```
config switch trunk
   edit <trunk name>
      set description <description_string>
```

```
      set members <ports>
      set mode {fortinet-trunk | lacp-active | lacp-passive | static}
      set port-selection-criteria src-dst-ip-xor16
   end
end
```

Heartbeats

When in Fortinet-trunk mode, Heartbeat capability is enabled. Heartbeat messages monitor the status of FortiGate units. If one is unavailable, the FortiSwitch unit stops sending traffic to that FortiGate unit until the FortiGate unit becomes available.

If you enable `hb-verify`, each received heartbeat frame will be validated to match the signature (transmit-port plus switch serial number) and the following configured heartbeat parameters:

- hb-in-vlan
- hb-src-ip
- hb-dst-ip
- hb-src-upd-port
- hb-dst-udp-port

The destination MAC address of the heartbeat frame is set by default to 02:80:c2:00:00:02. You can change the value to any MAC address that is not a broadcast or multicast MAC address.

Configuring heartbeats

Configure the heartbeat fields using trunk configuration commands, as shown in this section. By default, all of the configurable values are set to zero, and hb-verify is disabled.

Set the mode to `forti-hb` and set the heartbeat loss limit to a value between 3 and 32.

The heartbeat will transmit at 1-second intervals on any link in the trunk that is up. This value is not configurable.

The heartbeat frame has configurable parameters for the layer-3 source and destination addresses and the layer-4 UDP ports. You must also specify the transmit and receive VLANs.

```
config switch trunk
   edit hb-trunk
      set mode fortinet-trunk
      set members <port> [<port>] ... [<port>]
      set hb-loss-limit <3-32>
      set hb-out-vlan <int>
      set hb-in-vlan <int>
      set hb-src-ip <x.x.x.x>
      set hb-dst-ip <x.x.x.x>
      set hb-src-udp-port <int>
      set hb-dst-udp-port <int>
      set hb-verify [ enable | disable ]
end
```

Use the following command to configure the destination MAC address:

```
config switch global
   set forti-trunk-dmac <mac address>
end
```

Example

The following example creates trunk tr1 with heartbeat capability:

```
config switch trunk
   edit "tr1"
      set mode fortinet-trunk
      set members "port1" "port2"
      set hb-out-vlan 300
      set hb-in-vlan 500
      set hb-src-ip 10.105.7.200
      set hb-dst-ip 10.105.7.199
      set hb-src-udp-port 12345
      set hb-dst-udp-port 54321
      set hb-verify enable
   next
end
```

Unicast hashing

You can configure the trunk hashing algorithm for unicast packets to use the source port:

```
config switch global
   set trunk-hash-unicast-src-port {enable | disable}
end
```

LLDP-MED

The Fortinet data center switches support the Link Layer Discovery Protocol (LLDP) for transmission and reception wherein the switch will multicast LLDP packets to advertise its identity and capabilities. A switch receives the equivalent information from adjacent layer-2 peers.

Fortinet data center switches support LLDP-MED (Media Endpoint Discovery), which is an enhancement of LLDP that provides the following facilities:

- Auto-discovery of LAN policies (such as VLAN, layer-2 priority, and differentiated services settings), to enable plug-and-play networking.
- Device location discovery to allow the creation of location databases and Enhanced 911 services for Voice over Internet Protocol (VoIP).
- Extended and automated power management for power over Ethernet (PoE) endpoints.
- Inventory management, allowing network administrators to track their network devices, and determine their characteristics (manufacturer, software and hardware versions, serial or asset number).

The switch will multicast LLDP packets to advertise its identity and capabilities. The switch receives the equivalent information from adjacent layer-2 peers.

Starting in FortiSwitch 6.2.0, you can use the CLI to configure the location table used by LLDP-MED for enhanced 911 emergency calls.

This section covers the following topics:

- Configuration notes on page 168
- LLDP global settings on page 169

- Configuring LLDP profiles on page 172
- Configuring an LLDP profile for the port on page 176
- Enabling LLDP on a port on page 177
- Checking the LLDP configuration on page 177
- Configuration deployment example on page 178
- Checking LLDP details on page 180
- LLDP OIDs on page 180

Configuration notes

Review the following notes before configuring LLDP-MED:

- When 802.1X and LLDP turn on at the same port, switching between LLDP profiles requires a manual reset of all authentication sessions.
- Fortinet recommends LLDP-MED-capable phones.
- The FortiSwitch unit functions as a Network Connectivity device (that is, NIC, switch, router, and gateway), and will only support sending TLVs intended for Network Connectivity devices.
- LLDP supports up to 16 neighbors per physical port.
- The FortiSwitch unit accepts and parses packets using the CDP (Cisco Discovery Protocol) and count CDP neighbors towards the neighbor limit on a physical port. If neighbors exist, the FortiSwitch unit transmits CDP packets in addition to LLDP.
- With release 3.5.1, CDP is independently controllable through the `set cdp-status` command on the physical port. The FortiSwitch unit no longer requires a neighbor to trigger it to transmit CDP; it will transmit provided cdp-status is configured as tx-only or tx-rx. The default configuration for CDP-status is disabled. It still uses values pulled from the lldp-profile to configure its contents.
- LLDP must be globally enabled under the `config switch lldp settings` command for CDP to be transmitted or received:
- If a port is added into a *virtual-wire* (connects two ends of a controlled system using a radio frequency [RF] medium), the FortiSwitch unit will disable the transmission and receipt of LLDP and CDP packets and remove all neighbors from the port. This virtual-wire state is noted in the `get switch lldp neighbor-summary` command output.
- If the combination of configured TLVs exceeds the maximum frame size on a port, that frame cannot be sent.
- If a port is configured with an LLDP profile that has `auto-isl` enabled, the LLDP transmit frequency (normally set under `config switch lldp settings` with the `set tx-interval` command) for that port is overridden by the profile's `auto-isl-hello-timer` setting (the default is 3 seconds).
- When the switch is in FortLink mode, all ports are changed to have profiles with `auto-isl` enabled by default, and the ports' normal transmit interval is overridden by the `auto-isl-hello-timer` setting in that profile (the default is 3 seconds).
- The default-auto-isl LLDP profile, which is one of the two default LLDP profiles, has `auto-isl` enabled. Any port configured with the default-auto-isl profile will transmit LLDP PDUs every 3 seconds when the `auto-isl-hello-timer` option in that profile is set at the default of 3 seconds.
- The Time to Live (TTL) value sent in the LLDP PDUs is still based on the `tx-interval` and `tx-hold` values under `config switch lldp settings`, even if the transmit interval has been overridden by the `auto-isl-hello-timer` setting.

LLDP global settings

Using the GUI:

1. Go to *Switch* > *LLDP MED* > *Settings*.
2. Select or clear *Enable LLDP Transmit/Receive*.
3. Select the management interface.
4. Enter a value in the *Transmit Hold* field.
5. Enter the number of seconds for the transmit interval.
6. Select or clear *Fast Start*. If you select *Fast Start*, enter the number of seconds.
7. Select *Update*.

Using the CLI:

```
config switch lldp settings
   set status {enable | disable}
   set tx-hold <integer>
   set tx-interval <integer>
   set fast-start-interval <integer>
   set management-interface <layer-3 interface>
   set management-address {ipv4 | ipv6 | none}
end
```

Variable	Description
status	Enable or disable
tx-hold	Number of tx-intervals before the local LLDP data expires (that is, the packet TTL (in seconds) is `tx-hold` times `tx-interval`). The range for tx-hold is 1 to 16, and the default value is 4.
tx-interval	Frequency of LLDP PDU transmission ranging from 5 to 4095 seconds (default is 30).
fast-start-interval	How often the FortiSwitch unit transmits the first four LLDP packets when a link comes up. The range is 2 to 5 seconds, and the default is 2 seconds. Set this variable to zero to disable fast start.
management-interface	Primary management interface advertised in LLDP and CDP PDUs.
management-address	Select whether to advertise an IPv4 address, an IPv6 address, or none. By default, both IPv4 and IPv6 addresses are advertised.

Setting the asset tag

To help identify the unit, LLDP uses the asset tag, which can be at most 32 characters. It will be added to the LLDP-MED inventory TLV (when that TLV is enabled):

```
config system global
   set asset-tag <string>
end
```

Configuring the location table

Because mobile phones have no fixed addresses associated with them, calls to 911 need the location information provided in emergency location identifier numbers (ELINs). You need to first configure the location table used by LLDP-MED for enhanced 911 emergency calls and then configure the LLDP profile to use the location table.

Using the GUI:

1. Go to *System* > *Locations*.
2. Select *Add Location*.
3. Required. In the Name field, enter a unique name for the location entry.
4. In the ELIN Number field, enter the ELIN, which is a unique phone number. The value must be no more than 31-characters long.
5. Enter the civic address.
   a. In the Additional field, enter additional location information, for example, `west wing`.
   b. In the Additional Code field, enter the additional country-specific code for the location. In Japan, use the Japan Industry Standard (JIS) address code.
   c. In the Block field, enter the neighborhood (Korea) or block
   d. In the Branch Road field, enter the branch road name. This value is used when side streets do not have unique names so that both the primary road and side street are used to identify the correct road
   e. In the Building field, enter the name of the building (structure) if the address includes more than one building, for example, `Law Library`.
   f. In the City field, enter the city (Germany), township, or shi (Japan).
   g. In the City Division field, enter the city division, borough, city district (Germany), ward, or chou (Japan).
   h. Required. In the Country field, enter the two-letter ISO 3166 country code in capital ASCII letters, for example, `US`, `CA`, `DK`, and `DE`.
   i. In the Country Subdivision field, enter the national subdivision (such as state, canton, region, province, or prefecture). In Canada, the subdivision is province. In Germany, the subdivision is state. In Japan, the subdivision is metropolis. In Korea, the subdivision is province. In the United States, the subdivision is state.
   j. In the County field, enter the county (Canada, Germany, Korea, and United States), parish, gun (Japan), or district (India).
   k. In the Direction field, enter `N`, `E`, `S`, `W`, `NE`, `NW`, `SE`, or `SW` for the leading street direction.
   l. In the Floor field, enter the floor number, for example, `4`.
   m. In the Landmark field, enter the nickname, landmark, or vanity address, for example, `UC Berkeley`.
   n. In the Language field, enter the ISO 639 language code used for the address information.
   o. In the Name field, enter the person or organization associated with the address, for example, `Fortinet` or `Textures Beauty Salon`.
   p. In the Number field, enter the street address, for example, `1560`.
   q. In the Number Suffix field, enter any modifier to the street address. For example, if the full street address is 1560A, enter `1560` for the number and `A` for the number suffix.
   r. In the Place Type field, enter the type of place, for example, `home`, `office`, or `street`.
   s. In the Post Office Box field, enter the post office box, for example, `P.O. Box 1543`. When the post-office-box value is set, the street address components are replaced with this value.
   t. In the Postal Community field, enter the postal community name, for example, `Alviso`. When the postal community name is set, the civic community name is replaced by this value.
   u. In the Primary Road field, enter the primary road or street name for the address.
   v. In the Road Section field, enter the specific section or stretch of a primary road. This field is used when the same street number appears more than once on the primary road.

w. In the Room field, enter the room number, for example, `7A`.
x. In the Script field, enter the script used to present the address information, for example, `Latn`.
y. In the Seat field, enter the seat number in a stadium or theater or a cubicle number in an office or a booth in a trade show.
z. In the Street field, enter the street (Canada, Germany, Korea, and United States).
aa. In the Street Name Post Mod field, enter an optional part of the street name that appears after the actual street name. If the full street name is East End Avenue Extended, enter `Extended`.
ab. In the Street Name Pre Mod field, enter an optional part of the street name that appears before the actual street name. If the full street name is Old North First Street, enter `Old`.
ac. In the Street Suffix field, enter the type of street, for example, `Ave` or `Place`. Valid values are listed in the United States Postal Service Publication 28 [18], Appendix C.
ad. In the Sub Branch Road field, enter the name of a street that branches off of a branch road. This value is used when the primary road, branch road, and subbranch road names are needed to identify the correct street.
ae. In the Trailing Str Suffix field, enter `N`, `E`, `S`, `W`, `NE`, `NW`, `SE`, or `SW` for the trailing street direction.
af. In the Unit field, enter the unit (apartment or suite), for example, `Apt 27`.
ag. In the ZIP field, enter the postal or zip code for the address, for example, `94089-1345`.

6. Enter the GPS coordinates.
   a. Required. In the Altitude field, enter the vertical height of a location in feet or meters. The format is +/- floating-point number, for example, `117.47`.
   b. Select *Feet* or *Meters* for the unit of measurement for the altitude.
   c. For the Datum drop-down list, select which map is used for the location: *WGS84*, *NAD83*, or *NAD83/MLLW*.
   d. Required. In the Latitude field, enter the latitude. The format is floating point starting with +/- or ending with N/S, for example, `+/-16.67` or `16.67N`.
   e. Required. In the Longitude field, enter the longitude. The format is floating point starting with +/- or ending with E/W, for example, `+/-26.789` or `26.789E`.
7. Select *Add*.

Using the CLI:

```
config system location
   edit <name>
      config address-civic
         set additional <string>
         set additional-code <string>
         set block <string>
         set branch-road <string>
         set building <string>
         set city <string>
         set city-division <string>
         set country <string>
         set country-subdivision <string>
         set county <string>
         set direction <string>
         set floor <string>
         set landmark <string>
         set language <string>
         set name <string>
         set number <string>
         set number-suffix <string>
         set place-type <string>
         set post-office-box <string>
         set postal-community <string>
```

```
      set primary-road <string>
      set road-section <string>
      set room <string>
      set script <string>
      set seat <string>
      set street <string>
      set street-name-post-mod <string>
      set street-name-pre-mod <string>
      set street-suffix <string>
      set sub-branch-road <string>
      set trailing-str-suffix <string>
      set unit <string>
      set zip <string>
   end
   config coordinates
      set altitude <string>
      set altitude-unit {f | m}
      set datum {NAD83 | NAD83/MLLW | WGS84}
      set latitude <string>
      set longitude <string>
   end
   config elin-number
      set elin-number <number>
   end
```

For example:

```
config system location
   edit Fortinet
      config address-civic
         set country "US"
         set language "English"
         set county "Santa Clara"
         set city "Sunnyvale"
         set street "Kifer"
         set street-suffix "Road"
         set number "899"
         set zip "94086"
         set building "1"
         set floor "1"
         set seat "1293"
      end
   next
   edit "Fortinet"
      config elin-number
         set elin-number "14082357700"
      end
   end
```

Configuring LLDP profiles

LLDP profile contains most of the port-specific configuration. Profiles are designed to provide a central point of configuration for LLDP settings that are likely to be the same for multiple ports.

Two static LLDP profiles, default and default-auto-isl, are created automatically. They can be modified but not deleted. The default-auto-isl profile always has auto-isl enabled and rejects any configurations that attempt to disable it.

LLDP-MED network policies

LLDP-MED network policies cannot be deleted or added. To use a policy, set the med-tlvs field to include `network-policy` and the desired network policy to `enabled`. The VLAN values on the policy are cross-checked against the VLAN native and untagged attributes for any interfaces that contain physical-ports using this profile. The cross-check determines if the policy Type Length Value (TLV) should be sent (VLAN must be native or allowed) and if the TLV should mark the VLAN as tagged or untagged (VLAN is native, or is in untagged). The network policy TLV is automatically updated when either a switch interface changes VLAN configuration or a physical port is added to, or removed from, a trunk.

The FortiSwitch unit supports the following LLDP-MED TLVs:

- Inventory Management TLVs
- Location Identification TLVs
- Network Policy TLV
- Power Management TLVs

Refer to the Configuration deployment example on page 178.

Custom TLVs (organizationally specific TLVs)

Custom TLVs are configured in their own subtable, available in each profile. They allow you to emulate the TLVs defined in various specifications by using their OUI and subtype and ensuring that the data is formatted correctly. You could also define a purely arbitrary custom TLV for some other vendor or for their company.

The “name” value for each custom TLV is neither used by nor has an effect on LLDP; it simply differentiates between custom TLV entries:

```
config custom-tlvs
   edit <TLVname_str>
      set information-string <hex-bytes>
      set oui <hex-bytes>
      set subtype <integer>
   next
```

The OUI value for each TLV must be set to three bytes. If just one of those bytes is nonzero it is accepted; any value other than "000" is valid. The subtype is optional and ranges from 0 (default) to 255. The information string can be 0 to 507 bytes, in hexadecimal notation.

The FortiSwitch unit does not check for conflicts either between custom TLV values or with standardized TLVs. That is, other than ensuring that the OUI is nonzero, the FortiSwitch unit does not check the OUI, subtype (or data) values entered in the CLI for conflicts with other Custom TLVs or with the OUI and subtypes of TLVs defined by the 802.1, 802.3, LLDP-MED, or other standards. While this behavior could cause LLDP protocol issues, it also allows a large degree of flexibility were you to substitute a standard TLV that is not supported yet.

802.1 TLVs

Two 802.1 TLVs are supported in the LLDP profile:

- Port VLAN ID
- VLAN Name

By default, no 802.1 TLVs are enabled.

The Port VLAN ID TLV sends the native VLAN of the port. This value is updated when the native VLAN of the interface representing the physical port changes or if the physical port is added to, or removed from, a trunk.

The VLAN Name TLV sends the VLAN descriptions that are configured in the `set description` command under `config switch vlan`.

The following are the requirements for using the VLAN Name TLV:

- The VLAN description is set in the `set description` command under `config switch vlan`.
- The `set 802.1-tlvs` command is set to `vlan-name`.
- The VLAN identifiers listed in the `set vlan-name-map` command are separated with commas and no spaces.
- The `vlan-name-map` configuration must be less than 4,096 characters.
- The port that uses the LLDP profile with the VLAN Name TLV is the same port advertising the VLAN names.
- The VLAN identifier is allowed on the port in the `config switch interface` command.
- A maximum of 10 VLAN names is supported.

To enable the VLAN Name TLV:

```
config switch lldp profile
   edit <LLDP_profile_name>
      set 802.1-tlvs vlan-name
      set vlan-name-map <single_VLANs_or_VLAN_ranges>
   next
end
```

For example:

```
config switch lldp profile
   edit newprofile
      set 802.1-tlvs vlan-name
      set vlan-name-map 1,5,10-15
   next
end
```

802.3 TLVs

There are three 802.3 TLVs that can be enabled or disabled:

- *Efficient Energy Ethernet Config*—This TLV sends whether energy-efficient Ethernet is enabled on the port. If this variable is changed, the sent value will reflect the updated value.
- *Maximum Frame Size*—This TLV sends the max-frame-size value of the port. If this variable is changed, the sent value will reflect the updated value.
- *PoE+ Classification*—This TLV sends whether there is software PoE negotiation on the port.

By default, no 802.3 TLVs are enabled.

In the following example, you need to specify that the TLV sends the PoE classification of the port to power up an IP phone with expansion modules.

```
config switch lldp profile
   edit "phone-with-expansion-modules"
      set 802.3-tlvs power-negotiation <--------- must have
      ...
```

Auto-ISL

The auto-ISL configuration that was formerly in the `switch physical-port` command has been moved to the `switch lldp-profile` command. All behavior and default values are unchanged.

Assigning a VLAN to a port in the LLDP profile

You can configure the network policy of an LLDP profile to assign the specified VLAN to ports that use the LLDP profile. The VLAN is added as though it were configured in the `set allowed-vlans` setting in the `config switch interface` configuration.

This feature has the following requirements:

- The port cannot belong to a trunk or virtual wire.
- The port must have `lldp-status` set to `rx-only`, `tx-only`, or `tx-rx`.
- The port must have `private-vlan` set to `disabled`.
- LLDP must be enabled under the `config switch lldp settings` command.
- The `set med-tlvs network-policy` option must be set under the `config switch lldp profile` configuration.
- The `assign-vlan` option must be enabled in the `med-network-policy` configuration under the `config switch lldp profile` configuration.
- The VLAN assigned in the LLDP profile must be a valid VLAN.

Note:

- If the VLAN added to the interface by the LLDP profile is also listed under the `set untagged-vlans` configuration in the `config switch interface` command, the VLAN is added as untagged.
- If the VLAN added to the interface by the LLDP profile is also the native VLAN of the port, no changes occur.
- The LLDP service determines the contents of the network-policy TLV being sent based on the current state of the switch interface. If the LLDP VLAN assignment does not happen or the assigned VLAN is changed by another configuration (such as the `set untagged-vlans` configuration in `config switch interface`), the LLDP network policy TLVs being sent will reflect the actual state of the interface, not the configured value.

To specify a VLAN in the network policy of an LLDP profile:

```
config med-network-policy
   edit <policy_type_name>
      set status enable
      set assign-vlan enable
      set dscp <0-63>
      set priority <0-7>
      set vlan <0-4094>
   next
```

For example:

```
config med-network-policy
   edit default
      set status enable
      set assign-vlan enable
      set vlan 15
      set dscp 30
      set priority 3
   next
```

Configuring an LLDP profile for the port

Configure an LLDP profile for the port. By default, the port uses the default LLDP profile.

Using the GUI:

1. Go to *Switch* > *LLDP-MED* > *Profiles*.
2. Select *Add Profile*.
3. Enter a name for your LLDP profile.
4. If needed, select *Port VLAN ID*.
5. If needed, select one or more of the 802.3 TLVs: *Efficient Energy Ethernet Config*, *PoE+ Classification*, and *Maximum Frame Size*.
6. If needed, select *Enable* for Auto-ISL.
7. Enter the number of seconds for the Auto-ISL Hello Timer.
8. Enter the port group number for the Auto-ISL Port Group.
9. Enter the number of seconds for the Auto-ISL Receive Timeout.
10. If needed, select one or more of the MED TLVs: *Inventory Management*, *Location Identification*, *Network Policy*, and *Power Management*.
11. Select *Add*.

Using the CLI:

```
config switch lldp profile
   edit <profile>
      set 802.1-tlvs port-vlan-id
      set 802.3-tlvs max-frame-size
      set auto-isl {active | inactive}
      set auto-isl-hello-timer <1-30>
      set auto-isl-port-group <0-9>
      set auto-isl-receive-timeout <3-90>
      set auto-mclag-icl {enable | disable}
      set med-tlvs (inventory-management | location-identification | network-policy | power-
           management)
      config custom-tlvs
         edit <TLVname_str>
            set information-string <hex-bytes>
            set oui <hex-bytes>
            set subtype <integer>
         next
      config med-location-service
         edit address-civic
            set status {enable | disable}
            set sys-location-id <string>
         next
         edit coordinates
            set status {enable | disable}
            set sys-location-id <string>
         next
         edit elin-number
            set status {enable | disable}
            set sys-location-id <string>
         next
      config med-network-policy
```

```
      edit <policy_type_name>
         set status {enable | disable}
         set assign-vlan {enable | disable}
         set dscp <0-63>
         set priority <0-7>
         set vlan <0-4094>
      next
   end
```

Enabling LLDP on a port

To enable LLDP MED on a port, set the LLDP status to receive-only, transmit-only, or receive and transmit. The default value is TX/RX.

Using the GUI:

1. Go to *Switch* > *Physical Ports*.
2. Select a port and select *Edit*.
3. Select *TX/RX*, *RX Only*, *TX Only*, or *Disable* for the LLDP-MED status.
4. Select an LLDP profile.
5. Select *Update*.

Using the CLI:

```
config switch physical-port
   edit <port>
      set lldp-status (rx-only | tx-only | tx-rx | disable)
      set lldp-profile <profile name>
   next
end
```

Checking the LLDP configuration

View the LLDP configuration settings using the GUI:

1. Go to *Switch* > *LLDP-MED* > *Settings*.
2. Make any changes that are needed.
3. Select *Update*.

View the LLDP configuration settings using the CLI:

```
get switch lldp settings
status : enable
tx-hold : 4
tx-interval : 30
fast-start-interval : 2
management-interface: internal
```

View the LLDP profiles using the GUI:

1. Go to *Switch* > *LLDP-MED* > *Profiles*.
2. Select a profile and then select *Edit*.
3. Make any changes that are needed.
4. Select *Update*.

View the LLDP profiles using the CLI:

```
get switch lldp profile
== [ default ]
name: default 802.1-tlvs: 802.3-tlvs: med-tlvs: inventory-management network-policy
== [ default-auto-isl ]
name: default-auto-isl 802.1-tlvs: 802.3-tlvs: med-tlvs:
```

Use the following commands to display the LLDP information about LLDP status or the layer-2 peers for this FortiSwitch unit:

```
get switch lldp (auto-isl-status | neighbors-detail | neighbors-summary | profile | settings
     | stats)
```

Configuration deployment example

To configure LLDP:

1. Configure LLDP global configuration settings using the `config switch lldp settings` command.
2. Create LLDP profiles using the `config switch lldp profile` command to configure Type Length Values (TLVs) and other per-port settings.
3. Assign LLDP profiles to physical ports.
4. Apply VLAN to interface. (**NOTE:** LLDP profile values that are tied to VLANs will only be sent if the VLAN is assigned on the switch interface.)
   a. Configure the profile.

```
show switch lldp profile Forti670i
config switch lldp profile
   edit "Forti670i"
      config med-network-policy
         edit "voice"
            set dscp 46
            set priority 5
            set status enable
            set vlan 400
         next
         edit "guest-voice"
         next
         edit "guest-voice-signaling"
         next
         edit "softphone-voice"
         next
         edit "video-conferencing"
         next
         edit "streaming-video"
            set dscp 40
```

```
            set priority 3
            set status enable
            set vlan 400
         next
         edit "video-signalling"
         next
      end
   set med-tlvs inventory-management network-policy
next
end
```

b. Configure the interface.

```
show switch interface port4
config switch interface
   edit "port4"
      set allowed-vlans 400
      set snmp auto
   next
end
```

c. Connect a phone with LLDP-MED capability to the interface. **NOTE:** Make certain the LLDP, Learning, and DHCP features are enabled.

```
show switch physical-port port4
config switch physical-port
   edit "port4"
      set lldp-profile "Forti670i"
      set speed auto
   next
end
```

d. Verify.

```
show switch lldp neighbor-det port4

Neighbor learned on port port4 by LLDP protocol
Last change 12 seconds ago
Last packet received 12 seconds ago
Chassis ID: 10.105.251.40 (ip)
System Name: FON-670i
System Description:
V12.740.335.12.B
Time To Live: 60 seconds
System Capabilities: BT
Enabled Capabilities: BT
MED type: Communication Device Endpoint (Class III)
MED Capabilities: CP
Management IP Address: 10.105.251.40
Port ID: 00:a8:59:d8:f1:f6 (mac)
Port description: WAN Port 10M/100M/1000M
IEEE802.3, Power via MDI:
Power devicetype: PD
PSE MDI Power: Not Supported
PSE MDI Power Enabled: No
PSE Pair Selection: Can not be controlled
```

```
PSE power pairs: Signal
Power class: 1
Power type: 802.3at off
Power source: Unknown
Power priority: Unknown
Power requested: 0
Power allocated: 0
LLDP-MED, Network Policies:
voice: VLAN: 400 (tagged), Priority: 5 DSCP: 46
voice-signaling: VLAN: 400 (tagged), Priority: 4 DSCP: 35
streaming-video: VLAN: 400 (tagged), Priority: 3 DSCP: 40
```

Checking LLDP details

Using the GUI:

Go to *Switch* > *Monitor* > *LLDP*.

LLDP OIDs

Starting in FortiSwitchOS 6.2.2, the following object identifiers (OIDs) are supported by the LLDP management information base (MIB) file:

- .1.0.8802.1.1.2.1.1 (lldpConfiguration)
  - lldpMessageTxInterval
  - lldpMessageTxHoldMultiplier
  - lldpReinitDelay
  - lldpTxDelay
  - lldpNotificationInterval
- .1.0.8802.1.1.2.1.4.1 (lldpRemoteSystemsData.lldpRemTable)
  - lldpRemChassisIdSubtype
  - lldpRemChassisId
  - lldpRemPortSubtype
  - lldpRemPortId
  - lldpRemPortDesc
  - lldpRemSysName
  - lldpRemSysDesc
  - lldpRemSysCapSupported
  - lldpRemSysCapEnabled
- .1.0.8802.1.1.2.1.4.2 (lldpRemoteSystemsData.lldpRemManAddrTable)
  - lldpRemManAddrIfSubtype
  - lldpRemManAddrIfId
  - lldpRemManAddrOID

Interfaces

Interfaces refer to the layer-2 properties of FortiSwitch ports, including VLAN assignment, port security, and MAC security. Interfaces can be ports or trunks (such as link aggregation groups). To assign VLANs to an interface, see Configuring VLANs on page 316. Other layer-2 features are described in their respective chapters.

Go to *Switch* > *Interfaces* to see a list of switch interfaces and to see the type of interface and types of VLANs configured.

Except as mentioned in this manual, do not configure the internal interface.

Hover over the *Traffic* column to view the incoming and outgoing traffic rates:

At the bottom of the page is the *Average Bandwidth Per Interface* chart:

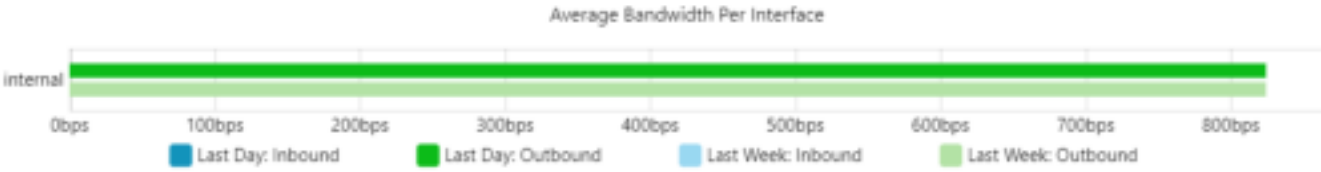

The following topics provide information about interfaces:

- Port security on page 181
- MAC security on page 226

Port security

To control network access, the FortiSwitch unit supports IEEE 802.1X authentication. A supplicant connected to a port on the switch must be authenticated by a RADIUS server to gain access to the network. The supplicant and the authentication server communicate using the switch using the Extensible Authentication Protocol (EAP). The FortiSwitch unit supports EAP-PEAP, EAP-TTLS, and EAP-TLS. Starting in FortiSwitchOS 7.4.0, EAP-FAST is supported.

To use the RADIUS server for authentication, you must configure the server before configuring the users or user groups on the FortiSwitch unit.

The FortiSwitch unit implements MAC-based authentication. The switch saves the MAC address of each supplicant's device. The switch provides network access only to devices that have successfully been authenticated.

The maximum number of MAC sessions per port is 20 for all FortiSwitch models. The following table lists the maximum number of MAC sessions per switch for each FortiSwitch model.

Model	Maximum number of MAC sessions per switch
108	80
112	60
124/224/424/524/1024	240
148/248/448/548/1048	480
3032	320

You can enable the MAC Authentication Bypass (MAB) option for devices (such as network printers) that cannot respond to the 802.1X authentication request. With MAB enabled on the port, the system will use the device MAC address as the user name and password for authentication.

Optionally, you can configure a guest VLAN for unauthorized users, a VLAN for users whose authentication was unsuccessful, and a VLAN for users when the authentication server is unavailable.

When the authentication server is unavailable after the server timeout period expires:

- You can control how many seconds the authentication server tries to authenticate users for before assigning them to an untagged VLAN:

```
config switch interface
   edit <interface_name>
      config port-security
         set port-security-mode {802.1X | 802.1X-mac-based}
         set authserver-timeout-period <3-15 seconds>
         set authserver-timeout-vlan {enable | disable}
         set authserver-timeout-vlanid <1-4094>
      end
      set security-groups <security-group-name>
   next
end
```

- If you are using 802.1x MAC-based authentication and FortSwitchOS 7.2.7 or later or FortiSwitchOS 7.4.3 or later, you can control how many seconds the authentication server tries to authenticate users for before assigning them to a tagged VLAN. Select `set authserver-timeout-tagged disable` if you do not want users to be assigned to a tagged VLAN when the authentication server times out. Select `set authserver-timeout-tagged lldp-voice` if you want users to be assigned to the VLAN specified in the `set lldp-profile` command (under `config switch physical-port`). Select `set authserver-timeout-tagged static` if you want users to be assigned to the VLAN specified in the `set authserver-timeout-tagged-vlanid` command.

```
config switch interface
   edit <interface_name>
      config port-security
         set port-security-mode 802.1X-mac-based
         set authserver-timeout-period <3-15 seconds>
         set authserver-timeout-tagged {disable | lldp-voice | static}
         set authserver-timeout-tagged-vlanid <1-4094>
      end
   set security-groups <security-group-name>
   next
end
```

- You can control how often the server checks if the RADIUS server is available:

```
config user radius
   edit <RADIUS_user_name>
      set link-monitor {enable | disable}
      set link-monitor-interval <5-120 seconds>
   next
end
```

Starting in FortiSwitchOS 7.2.1, you use the CLI to change the priority of MAB authentication and EAP 802.1X authentication.

When you are testing your system configuration for 802.1X authentication, you can use the monitor mode to allow network traffic to flow, even if there are configuration problems or authentication failures.

This section covers the following topics:

- Dynamic VLAN assignment on page 183
- Dynamic access control lists on page 187
- MAC authentication bypass (MAB) on page 192
- Configuring global settings on page 197
- Configuring the 802.1X settings on an interface on page 199
- Viewing the 802.1X details on page 203
- Clearing authorized sessions on page 205
- Authenticating users with a RADIUS server on page 206
- Authenticating an admin user with RADIUS on page 215
- RADIUS accounting and FortiGate RADIUS single sign-on on page 218
- RADIUS change of authorization (CoA) on page 220
- Use cases on page 223
- Detailed deployment notes on page 225

Dynamic VLAN assignment

You can configure the RADIUS server to return a VLAN in the authentication reply message:

1. On the FortiSwitch unit, select port-based authentication or MAC-based authentication and a security group.
2. On the RADIUS server, configure the attributes.

Using the GUI:

1. Go to *Switch* > *Interfaces*.
2. Select a port and then select *Edit*.
3. Select *802.1X* for port-based authentication or select *802.1X-MAC-based* for MAC-based authentication.

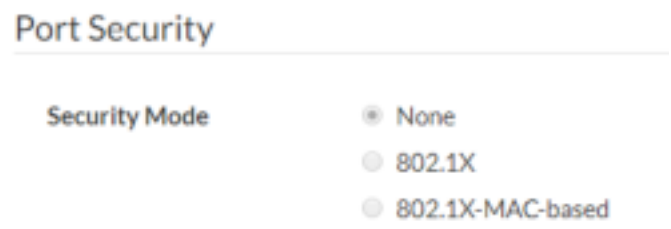

4. Select one or more security groups.
5. Select *OK*.

Using the CLI:

To select port-based authentication and the security group on the FortiSwitch unit:

```
config switch interface
   edit <interface_name>
      config port-security
         set port-security-mode 802.1X
      end
      set security-groups <security-group-name>
   end
```

The FortiSwitch unit will change the native VLAN of the port to that of the VLAN from the server.

To select MAC-based authentication and the security group on the FortiSwitch unit:

```
config switch interface
   edit <interface_name>
      config port-security
         set port-security-mode 802.1X-mac-based
      end
      set security-groups <security-group-name>
   end
```

Here, the switch assigns the returned VLAN only to this user's MAC address. The native VLAN of the port remains unchanged.

Use the following configuration command to view the MAC-based VLAN assignments:

```
diagnose switch vlan assignment mac list [sorted-by-mac | sorted-by-vlan]
```

Configure the following attributes in the RADIUS server:

- Tunnel-Private-Group-Id—VLAN ID or name (10)
- Tunnel-Medium-Type—IEEE-802 (6)
- Tunnel-Type—VLAN (13)

If the Tunnel-Private-Group-Id attribute is set to the VLAN name, the same string must be specified in the `set description` command under the `config switch vlan` command. For example:

```
config switch vlan
   edit 100
      set description "local_vlan"
   next
end
```

Starting in FortiSwitchOS 7.4.3, when the Tunnel-Private-Group-ID attribute has a Tag field, FortiSwitchOS will ignore the Tag field so that the VLAN string is parsed correctly.

Setting the priority for egress VLAN assignment

Starting in FortiSwitchOS 7.4.2, you can change how FortiSwitchOS searches for VLANs with names (specified in the `set description` command) that match the Egress-VLAN-Name attribute.

Before FortiSwitchOS 7.4.2, if there was more than one VLAN with the same name (specified in the `set description` command), FortiSwitchOS selected the VLAN with the lowest VLAN ID that matched the Egress-VLAN-Name attribute.

In the following example, the Egress-VLAN-Name attribute is set to `testVLAN`, and three VLANs have the same name of `testVLAN`. FortiSwitchOS matches the Egress-VLAN-Name attribute with the VLAN with the lowest ID, VLAN 4.

VLAN ID	VLAN name
4	testVLAN
5	testVLAN
6	testVLAN

In FortiSwitchOS 7.4.2, you can assign a priority to each VLAN. If there is more than one VLAN with the same name (specified in the `set description` command), FortiSwitchOS selects the VLAN with the lowest `assignment-priority` value (which is the highest priority) of the VLANs with names (specified in the `set description` command) that match the RADIUS Egress-VLAN-Name attribute. The `assignment-priority` value can be 1-255. By default, the `assignment-priority` is 128. The lowest `assignment-priority` value gets the highest priority.

In the following example, the Egress-VLAN-Name attribute is set to `localVLAN`, and four VLANs have the same name of `localVLAN`. FortiSwitchOS matches the Egress-VLAN-Name attribute with the VLAN with the lowest priority, VLAN 5.

VLAN ID	VLAN name	VLAN priority
4	localVLAN	50
5	localVLAN	25
6	localVLAN	75
7	localVLAN	100

To set the priority for matching VLAN names:

```
config switch vlan
   edit <VLAN_ID>
      set assignment-priority <1-255>
   next
end
```

For example:

```
config switch vlan
   edit 1
      set assignment-priority 200
   next
end
```

Using forced priority tagging

Starting in FortiSwitchOS 7.4.3, you can use forced priority tagging on the egress ports of the FS-1xxE, FS-1xxF, and FS-110G-FPOE models. When the `allowed-vlans` command is set on a port, all egress traffic will have the priority tag of vlan=0. This command is most useful when the port is acting as an access port for native traffic only.

To use forced priority tagging:

```
config switch interface
   edit <interface_name>
      set force-egr-prio-tag enable
   next
end
```

Configuring dynamic non-native VLANs

Starting in FortiSwitchOS 7.0.0, you can use the following RADIUS attributes to configure dynamic non-native VLANs:

- Egress-VLANID—Provides the VLAN identifier and controls whether egress packets are tagged (56).
  To set the VLAN ID value, use 0x31 for a tagged VLAN or 0x32 for an untagged VLAN. For example, to indicate that VLAN 16 is untagged, the Egress-VLANID is 0x32000010 or 838860816.
- Egress-VLAN-Name—Provides the VLAN name and controls whether egress packets are tagged (58).
  To provide the VLAN name as the VLAN description string defined under the `config switch vlan` command, use '1' for a tagged VLAN or '2' for an untagged VLAN. For example:
  - To assign the description "VLAN_8" to VLAN 8, which is tagged, use the following string: "1VLAN_8"
  - To assign the description "SALES_1772" to VLAN 1772, which is untagged, use the following string: "2SALES_1772"
- Ingress-Filters—Enables the use of ingress filters (57). The use of ingress filters cannot be disabled.

NOTE: The VLAN name in the Egress-VLAN-Name attribute must match the string specified in the `set description` command under the `config switch vlan` command. For example:

```
config switch vlan
   edit 100
      set description "local_vlan"
   next
end
```

You can verify your configuration with the `diagnose switch 802-1x status <port_name>` command. In the following example, the lines in boldface show the dynamic non-native VLANs:

```
S448DF3X15000026 # diagnose switch 802-1x status port1

port1 : Mode: port-based (mac-by-pass enable)
        Link: Link up
        Port State: authorized: (  )
        Dynamic Authorized Vlan : 101
        Dynamic Allowed Vlan list: 30-31,40-41
        Dynamic Untagged Vlan list: 40-41
        EAP pass-through : Enable
        EAP egress-frame-tagged : Enable
        EAP auto-untagged-vlans : Enable
        Allow MAC Move : Disable
        Quarantine VLAN (4093) detection : Enable
        Native Vlan : 101
```

```
        Allowed Vlan list: 4-7,30-31,40-41,101
        Untagged Vlan list: 40-41
        Guest VLAN :
        Auth-Fail Vlan :
        AuthServer-Timeout Vlan :

        Sessions info:
        00:00:00:01:01:02     Type=802.1x,PEAP,state=AUTHENTICATED,etime=0,eap_cnt=11
params:reAuth=3600
```

Dynamic access control lists

Starting in FortiSwitchOS 7.0.2, you can use RADIUS attributes to configure dynamic access control lists (DACLs) on 802.1X ports. DACLs are configured on a switch or saved on a RADIUS server. You can use DACLs to control traffic per user session, per port, or per MAC address for switch ports directly connected to user clients. DACLs apply to hardware only when 802.1X authentication is successful.

You can use DACLs with 802.1X port-based authentication and 802.1X MAC-based authentication. IPv4 is supported, but IPv6 is not supported. You can use DACLs with monitor mode (`open-auth`) and with static ACLs.

DACLs are disabled by default.

The maximum number of ACL entries per port is 45. The maximum number of entries includes both static ACL entries and DACL entries. Duplicate entries might cause an error.

FortiSwitch models	Maximum number of static ACL and DACL entries
108E	640 (ingress IPv4)
108F	640 (ingress IPv4)
124D	896
124E/148E	640 (ingress IPv4)
124F/148F	640 (ingress IPv4)
2xxD/2xxE	349 (ingress IPv4) 128 (ingress IPv6)
4xxD	896
424E/426E	732 (ingress IPv4) 256 (ingress IPv6) egress ipv4: 256; prelookup ipv4: 256
448E/424E-Fiber/FSR-424F-POE	1,500
5xxD	1,000
1024D	2,000
1024E/T1024E/T1024F-FPOE/2048F	3,000
1048E	4,000

FortiSwitch models	Maximum number of static ACL and DACL entries
3032D	3,072
3032E	1,000

To use the maximum number of DACL entries, you must enable the density mode:

```
config switch acl settings
   set density-mode enable
end
```

Two RADIUS attributes are supported:

- Filter-Id —The Filter-Id attribute defines the name of a access control list (ACL) predefined in FortiSwitchOS. With 802.1X port-based authentication, the DACL applies to the physical interface. With 802.1X MAC-based authentication, the DACL applies to the source MAC address of the authenticated client. If the Filter-Id cannot be found, the entire DACL fails.
- NAS-Filter-Rule—The NAS-Filter-Rule attribute defines the filter rules at the RADIUS server. After authentication, the DACL applies to the port.
  - The NAS-Filter-Rule supports a maximum of 80 characters, and you can specify a maximum of 45 entries per authentication session or a maximum of 45 entries per port.
  - Do not include blank spaces in the NAS-Filter-Rule. Commas and dashes are allowed.
  - A syntax error in one NAS-Filter-Rule causes the entire DACL to fail.

The following is the Filter-Id format:

```
Filter-Id += "<filter-name>"
```

For example:

```
Filter-Id += "filter-id-service1"
```

Changing the name of Filter-Id after authentication causes errors in the output of the `diagnose switch 802-1x status-dacl` command when the session is using Filter-Id.

The following is the NAS-Filter-Rule format:

```
NAS-Filter-Rule = " <deny|permit> in <ip|ip-protocol-value> from <any|<ip-addr>|ipv4-
     addr/mask> [<tcp/udp-port|tcp/udp min-max port>] to <any|<ip-addr>|ipv4-addr/mask>
     [<tcp/udp-port|tcp/udp min-max port>] [cnt] "
```

The following table explains the syntax of the NAS-Filter-Rule:

Option	Description
<deny|permit>	Select one of the following: • `permit`—Allow packets that match the rule. • `deny`—Drop packets that match the rule.
in	The `in` keyword specifies that the ACL applies only to the inbound traffic from the authenticated client.
<ip|ip-protocol-value>	Specify one of the following for the type of traffic to filter: • `ip`—Any protocol will match.

Option	Description
	• `ip-protocol-value`—IP traffic specified by either a protocol number or by `tcp`, `udp`, `icmp`, or (for IPv4 only) `igmp`. The range of protocol numbers is 0-255.
from <any|<ip-addr>|ipv4-addr/mask>	Required. Specify one of the following for the authenticated client source: • `any`—Specifies any IPv4 source address • `<ip-addr>|ipv4-addr/mask>`—Enter a series of contiguous source addresses or all source addresses in a subnet. The <mask> is the number of leftmost bits in a packet's source IPv4 address that must match the corresponding bits in the source IPv4 address. For example, `10.100.24.1/24` will match an inbound traffic from the authenticated client that has a source IPv4 address where the first three octets are 10.100.24.
[<tcp/udp-port|tcp/udp min-max port>] to	Specify the TCP or UDP port or range of ports. Used when the access control entry is intended to filter client TCP or UDP traffic with one or more specific TCP or UDP source port numbers. You can specify a single port or a single port range, such as 10.105.0.1/24 80 or 10.105.0.1/24 80-100.
<any|<ip-addr>|ipv4-addr/mask>	Specify one of the following: • `any`—Specifies any IPv4 destination address • `<ip-addr>|ipv4-addr/mask>`—Enter a series of contiguous destination addresses or all destination addresses in a subnet. The <mask> is the number of leftmost bits in a packet's destination IPv4 address that must match the corresponding bits in the destination IPv4 address. For example, `10.100.24.1/24` will match an inbound traffic from the authenticated client that has a destination IPv4 address where the first three octets are 10.100.24.
[<tcp/udp-port|tcp/udp min-max port>]	Specify the TCP or UDP port or range of ports. Used when the access control entry is intended to filter client TCP or UDP traffic with one or more specific TCP or UDP destination port numbers. You can specify a single port or a single port range, such as 10.105.0.1/24 80 or 10.105.0.1/24 80-100. For example, to deny any UDP traffic from an authenticated client that has a destination address of any address and a UDP destination port of 357-457: `deny in udp from any to any 357-457`
[cnt]	Specify the counter for a RADIUS-assigned access control entry.

For example:

- `NAS-Filter-Rule += "permit in 20 from any to any cnt"`
- `NAS-Filter-Rule += "deny in tcp from any to 10.10.10.1 23"`
- `NAS-Filter-Rule += "permit in tcp from any to any 23"`

When you use the NAS-Filter-Rule attribute, follow these guidelines:

- You can use 8 port ranges (source or destination ports) on the FS-148E, FS-148E-POE, and FS-148E-FPOE models.
- You can use 16 port ranges (source or destination ports) on the FS-108E, FS-108E-POE, FS-108E-FPOE, FS-108F, FS-108F-POE, FS-108F-FPOE, FS-124E, FS-124E-POE, FS-124E-FPOE, FS-124F, FS-124F-POE, FS-124F-FPOE, FS-148F, FS-148F-POE, and FS-148F-FPOE models.
- You can use up to 32 port ranges (source or destination ports) on the FS-1024D, FS-1024E, FS-T1024E, FS-T1024F-FPOE, FS-1048E, FS-3032E, FS-424E, FS-424E-POE, FS-424E-FPOE, FS-M426E-FPOE, FSR-124D, FS-224D-FPOE, FS-248D, FS-224E, FS-224E-POE, FS-248E-POE, FS-248E-FPOE, FS-424E-Fiber, FS-448E, FS-448E-POE, FS-448E-FPOE, FS-524D, FS-524D-FPOE, FS-548D, and FS-548D-FPOE models.
- Port ranges must have the smaller port number as the first number in the range and the larger port number as the second number in the range. For example, you can specify a port range of `8-10` but not `10-8`.
- If you specify a layer-4 port or layer-4 port range (for example, `permit in TCP from any to any 100-200 cnt`) when defining the source or destination in a dynamic ACL entry, FortiSwitchOS discards any port configurations made after the layer-4 configuration.

To enable DACL on an interface:

```
config switch interface
   edit <interface_name>
      config port-security
         set port-security-mode {802.1X | 802.1X-mac-based}
         set dacl enable
      end
   next
end
```

For example:

```
config switch interface
   edit port11
      config port-security
         set port-security-mode 802.1X
         set dacl enable
      end
   next
end
```

To configure a value for NAS-Filter-Rule or Filter-Id:

```
config switch acl service custom
   edit <ACL_service>
      set comment <string>
      set color <0-32>
      set protocol {ICMP | IP | TCP/UDP/SCTP}
      set protocol-number <IP protocol number>
      set tcp-portrange <port_number>-<port_number>
      set udp-portrange <port_number>-<port_number>
   next
```

```
end
```

For example:

```
config switch acl service custom
   edit filter-id-service1
      set comment "filter ID for service 1"
      set udp-portrange 10000-20000
   next
end
```

To create a template for the Filter-Id RADIUS attribute:

```
config switch acl 802-1X
   edit <policy_ID>
      set description <string>
      set filter-id <string>
      config access-list-entry
         edit <ingress_policy_ID>
            set description <string>
            set group <integer>
            config action
               set count {enable | disable}
               set drop {enable | disable}
            end
            config classifier
               set dst-ip-prefix <IP_address_and_netmask>
               set dst-mac <MAC_address>
               set ether-type <integer>
               set service <service_name>
               set src-ip-prefix <IP_address_and netmask>
               set src-mac <MAC_address>
            end
         next
      end
   next
end
```

For example:

```
config switch acl 802-1X
   edit 1
      set description "Test Filter-Id"
      set filter-id “Testing”
      config access-list-entry
         edit 1
            set description "Test ACL entry”
            config action
               set count enable
               set drop enable
            end
            config classifier
               set dst-ip-prefix 192.168.0.0 255.255.255.0
               set ether-type 0x0800
               set service "filter-id-service1"
               set src-ip-prefix 192.168.0.0 255.255.255.0
               set src-mac 00:00:00:00:00:00
            end
         next
```

```
      end
   next
end
```

To display the status of DACLs on a specified 802.1X port or on all ports:

```
diagnose switch 802-1x status-dacl [<port_name>]
```

To clear the DACLs from a specified interface or from all interfaces:

```
execute 802-1x dacl-clr-stat [<interface_name>]
```

To reinstall the DACLs on a specified interface or on all interfaces:

```
execute 802-1x dacl-reinstall [<interface_name>]
```

MAC authentication bypass (MAB)

Devices such as network printers, cameras, and sensors might not support 802.1X authentication. If you enable the MAB option on the port, the system will use the device MAC address as the user name and password for authentication.

MAB retries authentication three times before the device is assigned to a guest VLAN for unauthorized users. By default, reauthentication is disabled. Use the following commands if you want to change the default behavior:

```
config switch global
   config port-security
      set mab-reauth enable
   end
```

You must provision the RADIUS server to authenticate the devices that use MAB, either by adding the MAC addresses as regular users or by implementing additional logic to resolve the MAC addresses in a network inventory database.

The following flowchart shows the FortiSwitch 802.1X port-based authentication with MAB enabled and with an authentication priority of `auth-priority legacy`:

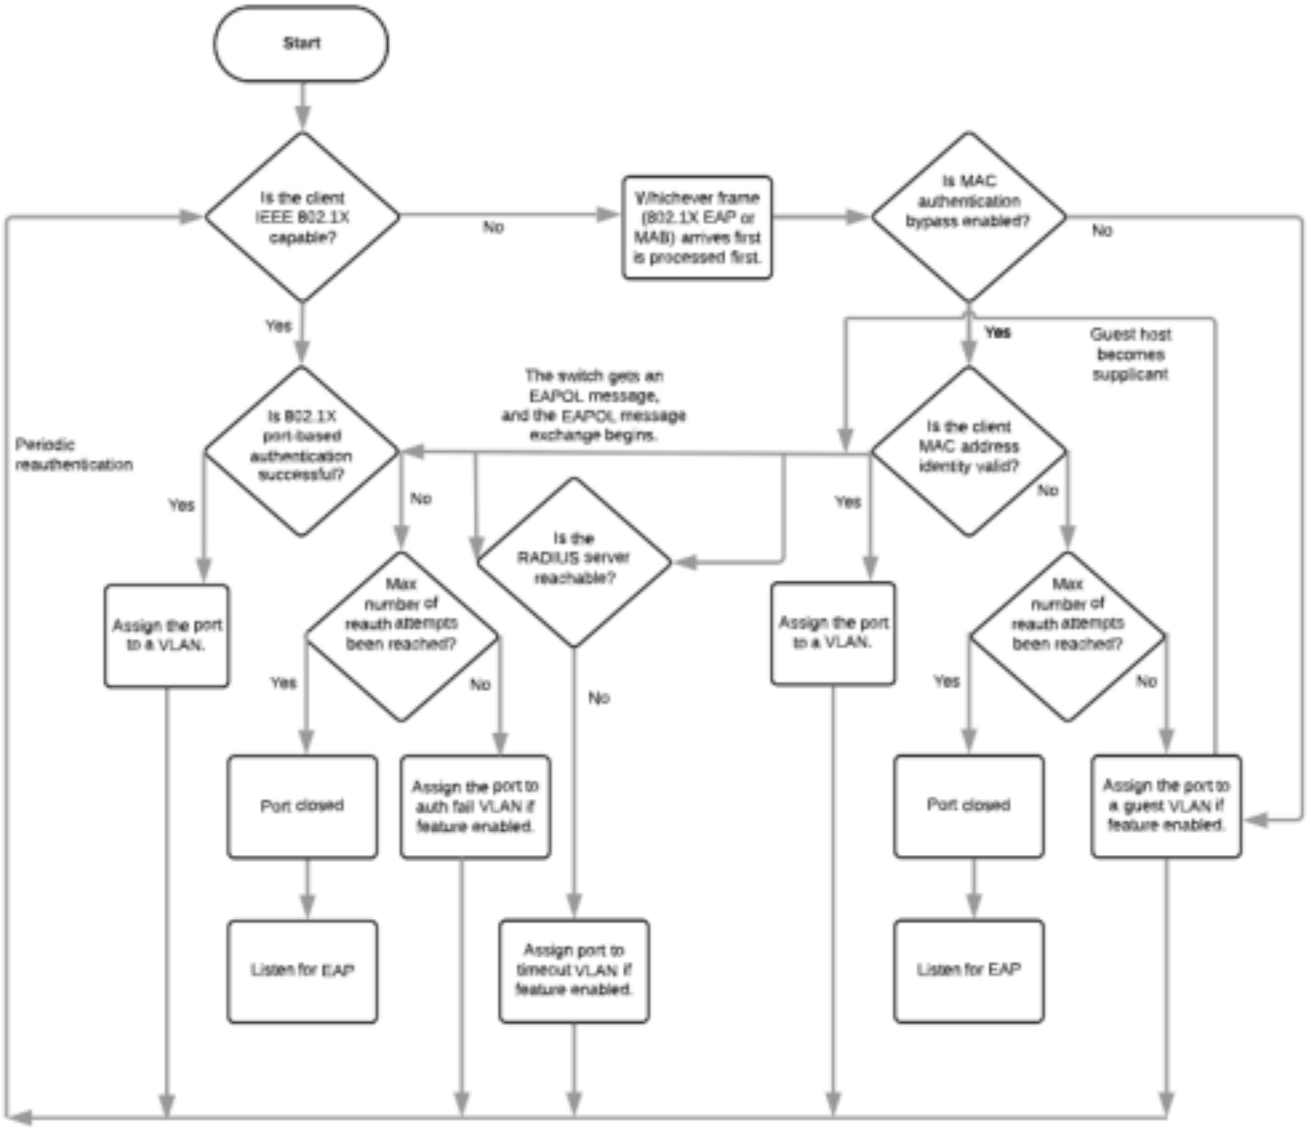

The following flowchart shows the FortiSwitch 802.1X MAC-based authentication with MAB enabled and with an authentication priority of `auth-priority legacy`:

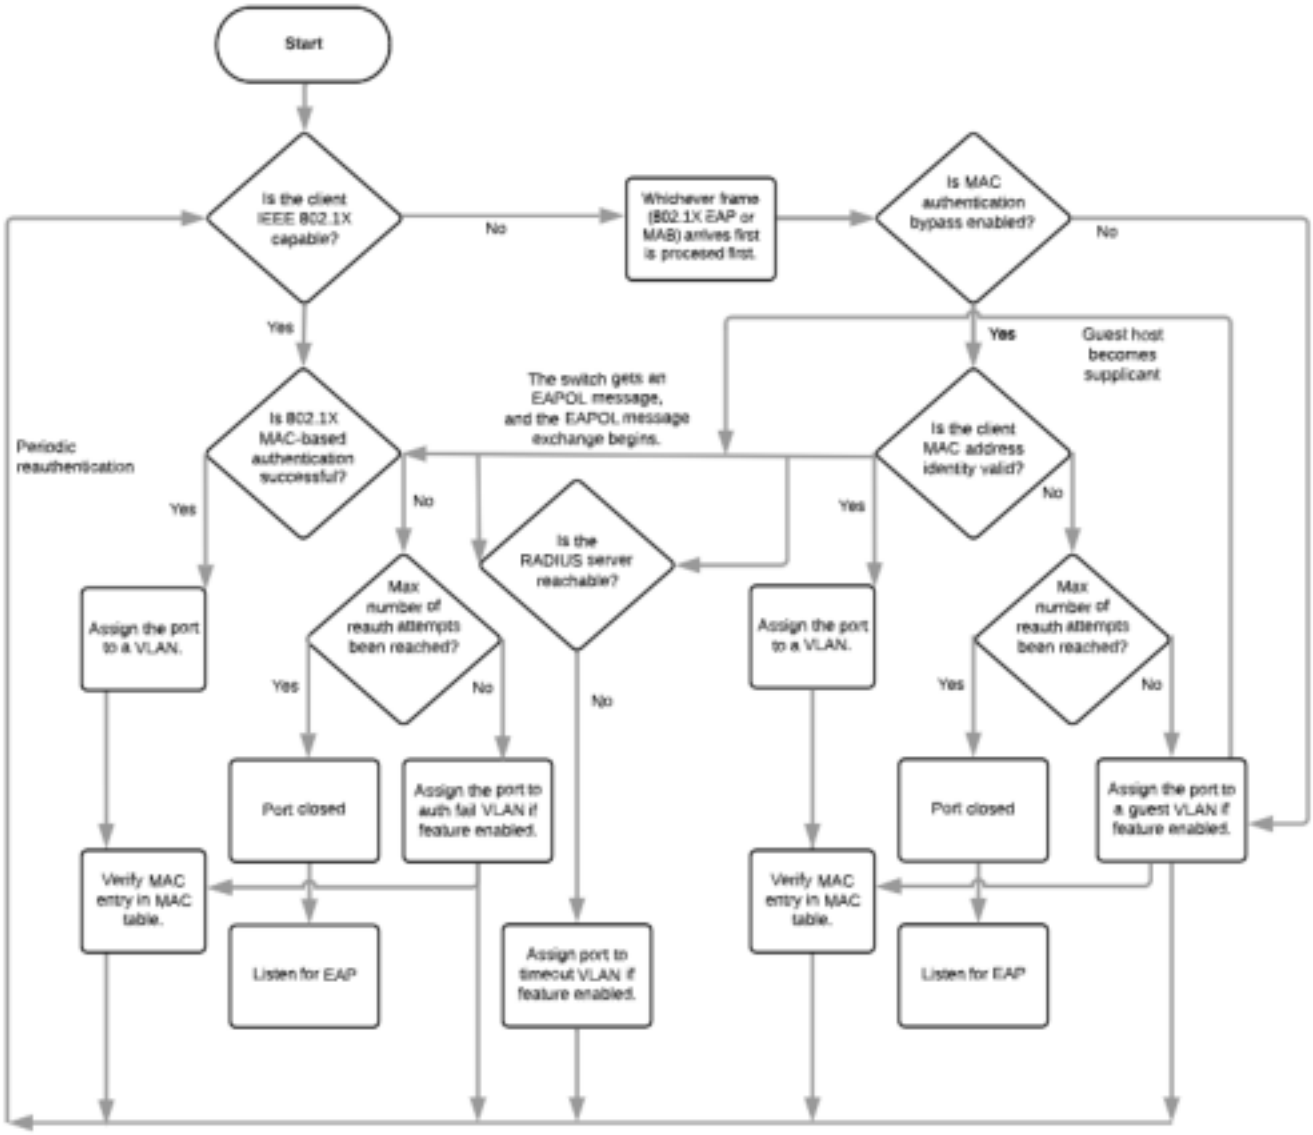

You use the CLI to change the priority of MAB authentication and EAP 802.1X authentication.

NOTE: MAB authentication must be enabled before you can change the priority of MAB authentication and EAP 802.1X authentication.

- Before FortiSwitchOS 7.2.1, the switch tried EAP 802.1X authentication and MAB authentication in the order that they were received with EAP 802.1X authentication having absolute priority. If authentication failed, users were assigned to the `auth-fail-vlanid` VLAN if it had been configured. There was no time delay. Starting in FortiSwitchOS 7.2.1, use the `set auth-priority legacy` command to keep this priority. After an upgrade, `auth-priority` is set to `legacy` by default.
- Starting in FortiSwitchOS 7.2.1, if you want the switch to try EAP 802.1X authentication first and then MAB authentication if EAP 802.1X fails, use the `set auth-priority dot1x-MAB` command. If MAB authentication also fails, users are assigned to the `auth-fail-vlanid` VLAN if it is configured.

- Starting in FortiSwitchOS 7.2.1, if you want the switch to try MAB authentication first and then EAP 802.1X authentication if MAB authentication fails, use the `set auth-priority MAB-dot1x` command. If EAP 802.1X authentication also fails, users are assigned to the `auth-fail-vlanid` VLAN if it is configured.
- Starting in FortiSwitchOS 7.2.3, MAB-only authentication is supported. In this mode, the FortiSwitch unit performs MAB authentication without performing EAP authentication. EAP packets are not sent. To enable MAB-only authentication:

```
config switch interface
   edit interface_name
      config port-security
         set port-security-mode {802.1X | 802.1X-mac-based}
         set mac-auth-bypass enable
         set auth-order MAB
      end
   next
end
```

In the following flowchart, the authentication priority is `dot1x-MAB`. If both EAP 802.1X authentication and MAB authentication fail, the user is assigned to the `auth-fail-vlanid` VLAN. If an EAPoL-Start packet is received after MAB authentication, the switch changes to EAP 802.1X authentication.

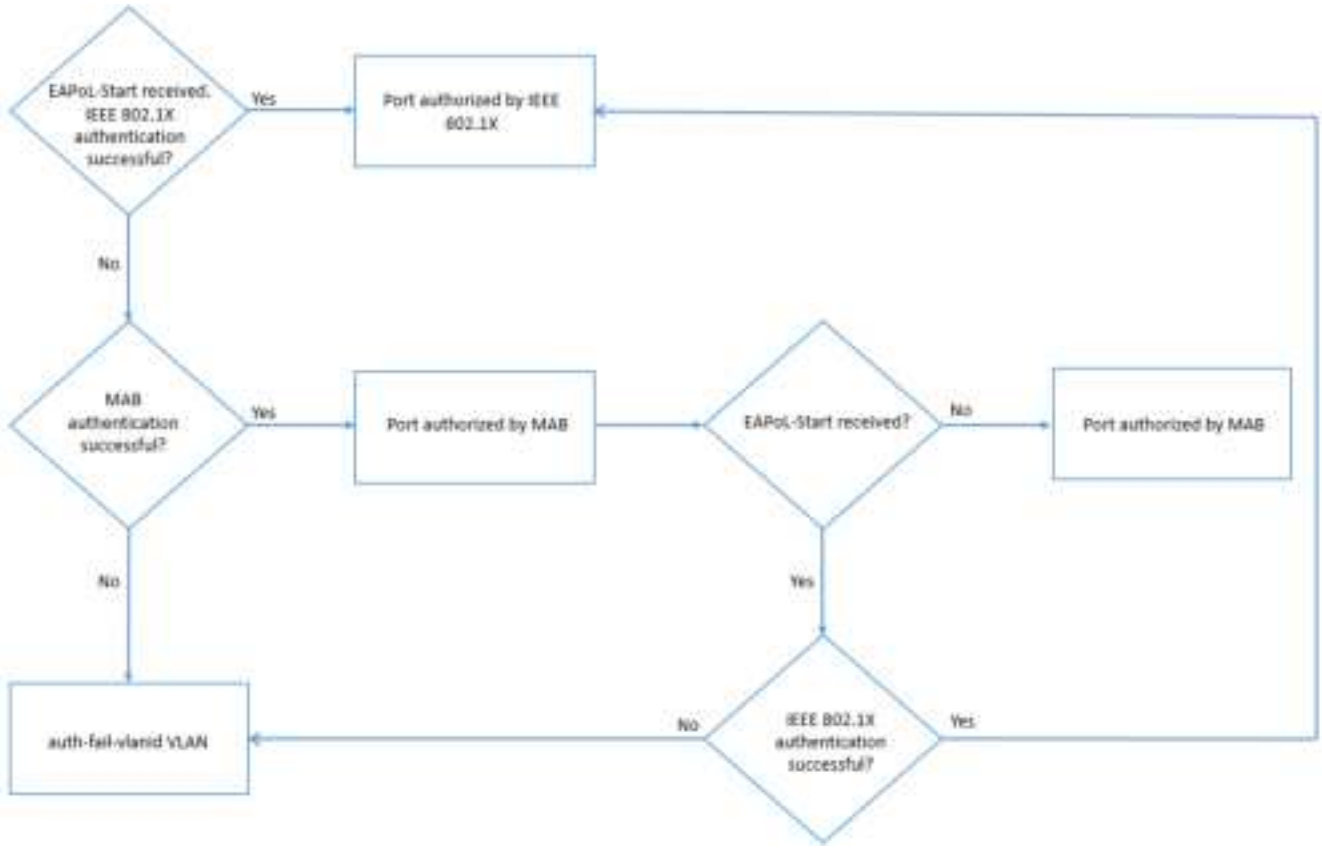

In the following flowchart, the authentication priority is `MAB-dot1x`. If MAB authentication fails, the switch attempts EAP 802.1X authentication. If an EAPoL-Start packet is received after MAB authentication, the switch attempts EAP 802.1X authentication without any time delay or processing impact.

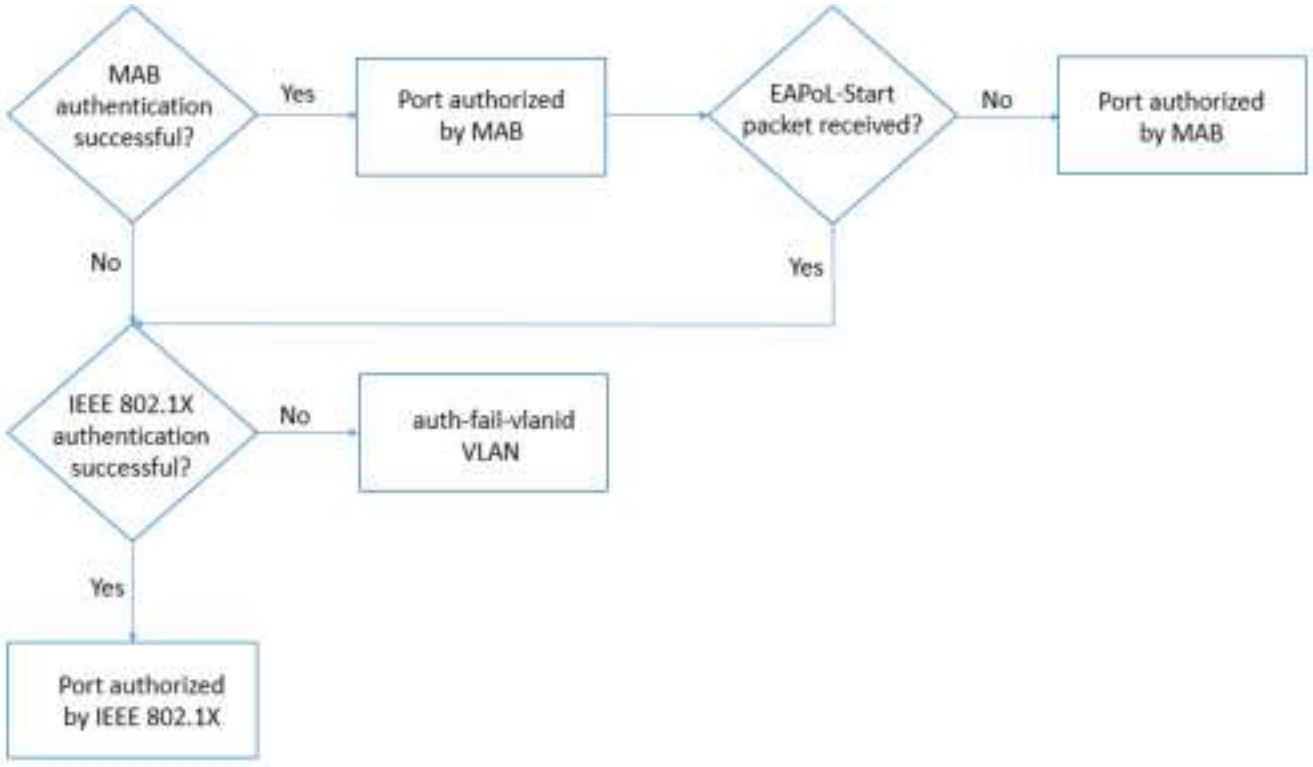

Configuring how long MAB sessions are kept

Starting in FortiSwitchOS 7.4.1, you can configure in the CLI how long MAC authentication bypass (MAB) sessions are kept:

- In static mode, MAB sessions are kept until the link goes down or the MAB sessions are manually deleted with the CLI. Before FortiSwitchOS 7.4.1, all MAB sessions behaved this way.

  Static mode is the default.
- In dynamic mode, MAB sessions are treated the same way as dynamically learned MAC addresses.

  In dynamic mode, you specify how long MAB sessions are kept with the `set mac-aging-interval <90-1000000> seconds` command (under `config switch global`). The minimum value is 90 seconds. By default, the `mac-aging-interval` is set to 300 seconds.

Whenever you change the `mac-aging-interval` or `mab-entry-as` setting, all MAB sessions are deleted, but the EAP sessions are unchanged.

This feature is supported on all FortiSwitch models.

To configure how long MAB sessions are kept:

```
config switch global
   config port-security
```

```
      set mab-entry-as {dynamic | static}
   end
end
```

Specifying how RADIUS request attributes are formatted

Starting in FortiSwitchOS 7.4.1, you can specify how the following RADIUS request attributes are formatted when they are sent to the RADIUS server:

- User-Name

  You can select a colon, hyphen, or single hyphen to use as a delimiter, or you can select `none` for no delimiter. By default, you can use a hyphen as the delimiter.
- User-Password

  You can select a colon, hyphen, or single hyphen to use as a delimiter, or you can select `none` for no delimiter. By default, you can use a hyphen as the delimiter.
- Called-Station-Id

  You can select a colon, hyphen, or single hyphen to use as a delimiter, or you can select `none` for no delimiter. By default, you can use a hyphen as the delimiter.
- Calling-Station-Id

  You can select a colon, hyphen, or single hyphen to use as a delimiter, or you can select `none` for no delimiter. By default, you can use a hyphen as the delimiter.

The following are examples of MAC addresses with the different delimiters:

- Using a colon as a delimiter: `00:11:22:33:44:55`
- Using a hyphen as a delimiter: `00-11-22-33-44-55`
- Using a single hyphen as a delimiter: `001122-334455`
- Using `none` for no delimiter: `001122334455`

You can also select whether to use lowercase or uppercase letters in MAC addresses. By default, lowercase letters are used.

To specify how RADIUS request attributes are formatted:

```
config switch global
   config port-security
      set mac-called-station-delimiter {colon | hyphen | none | single-hyphen}
      set mac-calling-station-delimiter {colon | hyphen | none | single-hyphen}
      set mac-case {lowercase | uppercase}
      set mac-password-delimiter {colon | hyphen | none | single-hyphen}
      set mac-username-delimiter {colon | hyphen | none | single-hyphen}
   end
```

Configuring global settings

To select which 802.1X certificate and certificate authority that the FortiSwitch unit uses, see SSL on page 64.

If a link goes down, you can select whether the impacted devices must reauthenticate. If reauthentication is unnecessary, select *Do Not Require Re-Authentication*. To revert all devices to the unauthenticated state and force each device to reauthenticate, select *Require Re-Authentication*.

MAB retries authentication before assigning a device to a guest VLAN for unauthorized users. MAB is disabled by default in the CLI.

The *Re-Authentication Period (Minutes)* field defines how often the device needs to reauthenticate (that is, if a session remains active beyond this number of minutes, the system requires the device to reauthenticate). Set the value to 0 to disable reauthentication. **NOTE:** For MAB authentication, the host entry is automatically re-authenticated after the re-authentication period. To clear the host entry, you need to clear the entry manually.

If 802.1X authentication fails, the Maximum Re-Authentication Attempts field caps the number of attempts that the system will initiate. Set the value to 0 to disable the reauthentication attempts.

Using the GUI:

1. Go to *Switch* > *Port Security*.

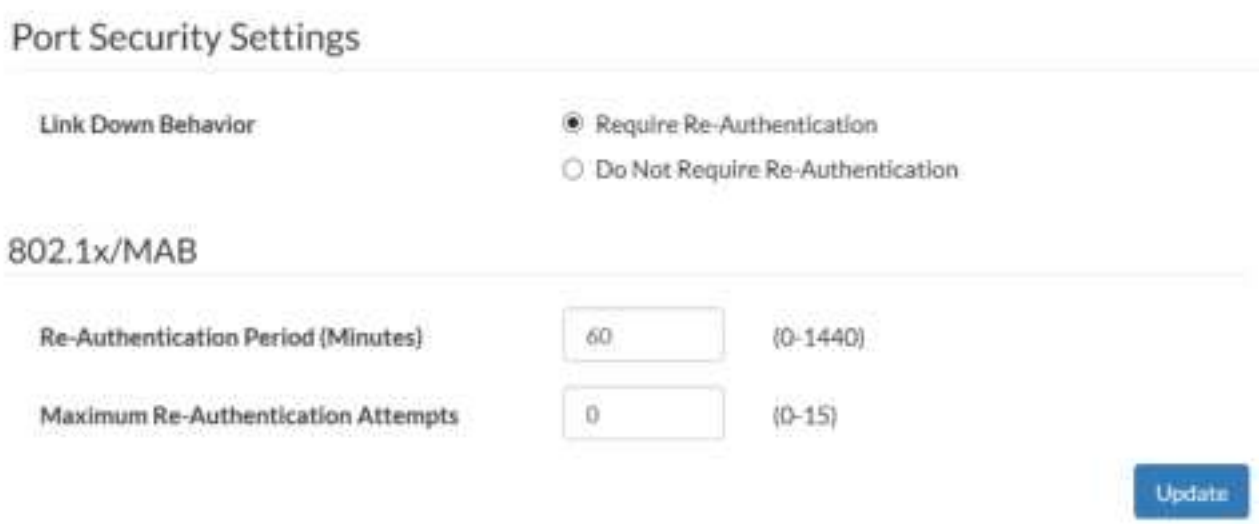

2. Select *Require Reauthentication* to revert all devices to the unauthenticated state if the link goes down or select *Do Not Require Reauthentication* if reauthentication is unnecessary if the link goes down.
3. In the *Re-Authentication Period (Minutes)* field, enter the number of minutes before the system requires the device to reauthenticate.
4. In the *Maximum Re-Authentication Attempts* field, enter the maximum number of times that the system tries to reauthorize the session.
5. Select *Update*.

Using the CLI:

```
config switch global
   config port-security
      set link-down-auth {no-action | set-unauth}
      set mab-reauth {enable | disable}
      set max-reauth-attempt <0-15>
      set reauth-period <0-1440>
   end
```

NOTE: Changes to global settings only take effect when new 802.1X/MAB sessions are created.

Configuring the 802.1X settings on an interface

Starting in FortiSwitchOS 7.0.0, you can use the CLI to allow an 802.1X client to move between ports that are not directly connected to the FortiSwitch unit without having to delete the 802.1X session. For example, you can move an 802.1X client PC that connects through an IP phone to port1 of the FortiSwitch unit to a port of a third-party switch that connects to port2 of the FortiSwitch unit.

This feature is available for 802.1X port-based authentication, 802.1X MAC-based authentication, MAB enabled or disabled, and EAP pass-through mode enabled or disabled. To see which models support this feature, refer to the FortiSwitch feature matrix.

NOTE: MAC-move tagged EAP is not supported.

To use this feature, enable `allow-mac-move` on the destination port (port2 in the example). When you are using the MAC move feature with EAP authentication, you can disable `eap-egress-tagged` to force the switch to always use the untagged EAP response.

Starting in FortiSwitchOS 7.2.3, the MAC move command has changed in the CLI:

- *For FSR-124D, 200 Series, FS-4xxE, 500 Series, FS-1024D, FS-1024E, FS-T1024E, FS-T1024F-FPOE, FS-1048E, and FS-3032E:* Use the `set allow-mac-move-to {enable | disable}` command under `config switch interface`. If you want to move an 802.1X client from interface 1 to interface 2, enable the MAC move command on interface 2.
- *For the FS-108E, FS-108E-POE, FS-108E-FPOE, FS-108F, FS-108F-POE, FS-108F-FPOE, FS-124E, FS-124E-POE, FS-124E-FPOE, FS-148E, and FS-148E-POE models:* Use the `set allow-mac-move-from {enable | disable}` command under `config switch interface`. If you want to move an 802.1X client from interface 1 to interface 2, enable the MAC move command on interface 1.
- *For the FS-124F, FS-124F-POE, FS-124F-FPOE, FS-148F, FS-148F-POE, and FS-148F-FPOE models:* Use the `set allow-mac-move {enable | disable}` command under `config switch global`. Instead of configuring the MAC move command on an interface, configure it globally.

Using the GUI:

1. Go to *Switch* > *Interfaces*.

2. Select a port and then select *Edit*.

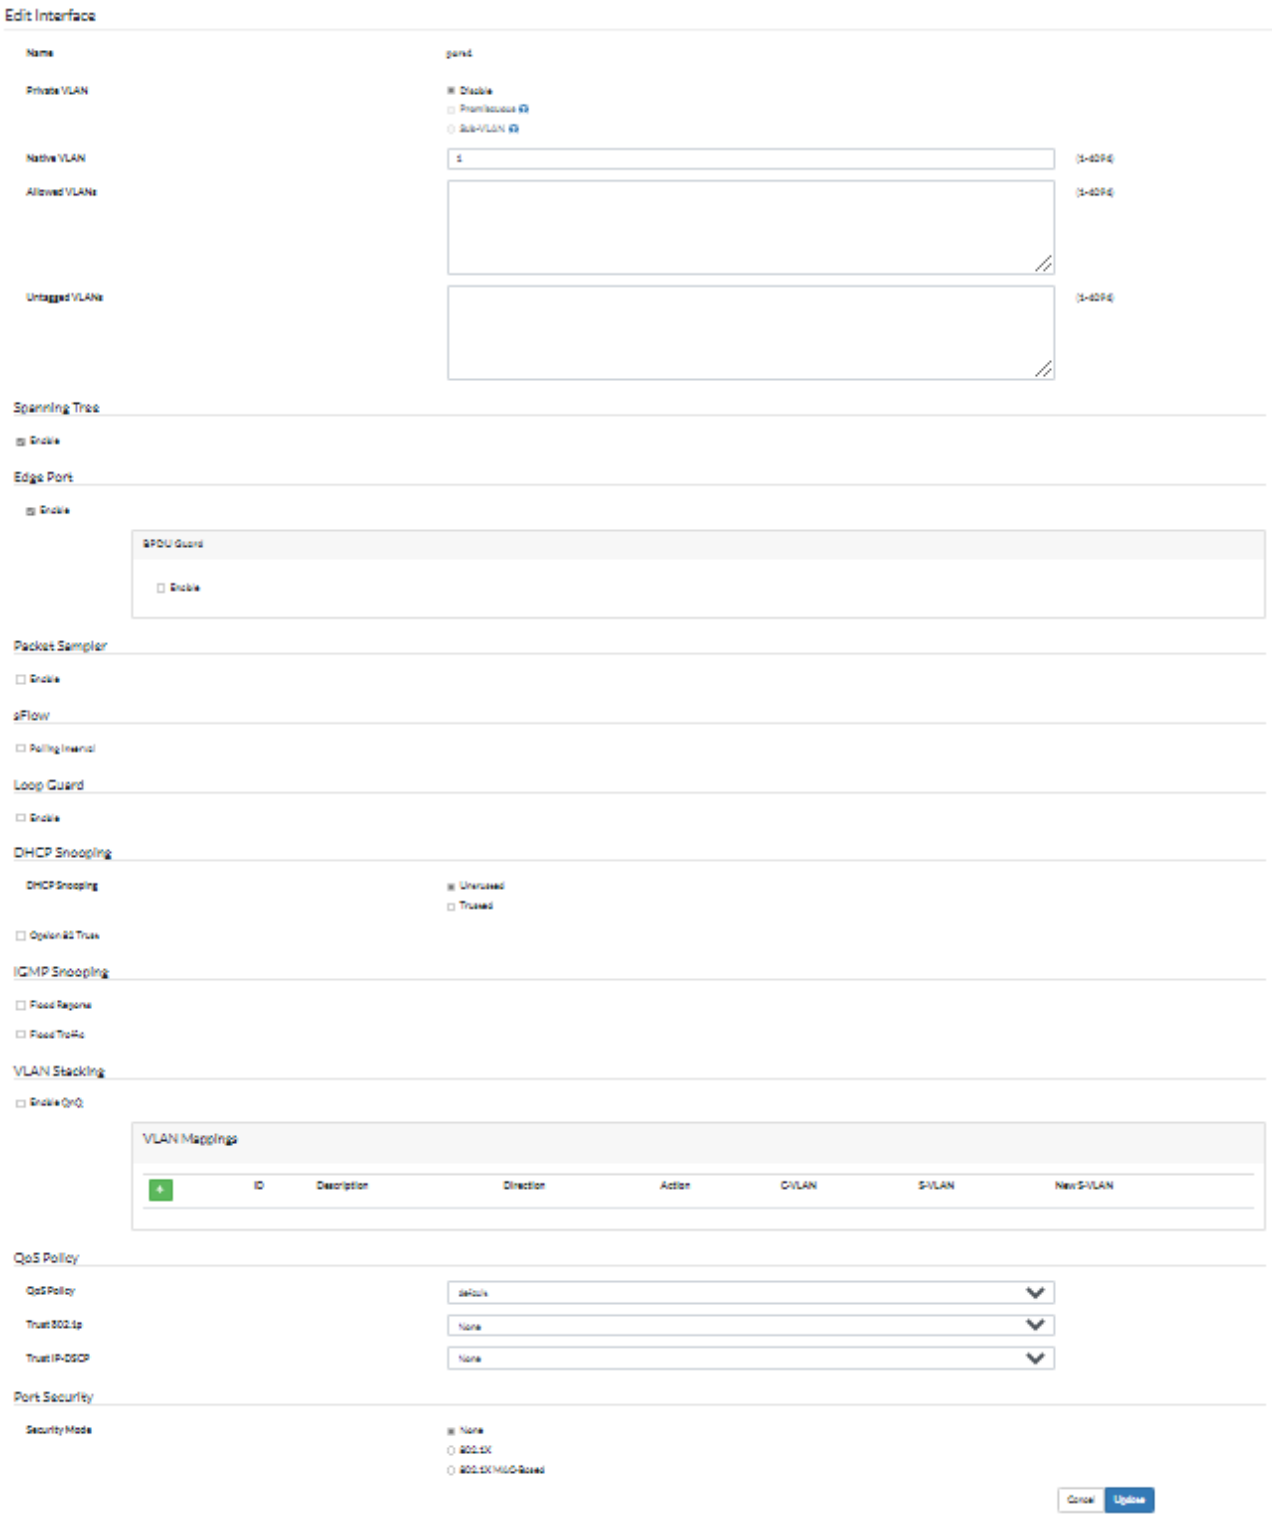

3. Select *802.1X* for port-based authentication or select *802.1X-MAC-based* for MAC-based authentication. The Port Security section displays additional options.

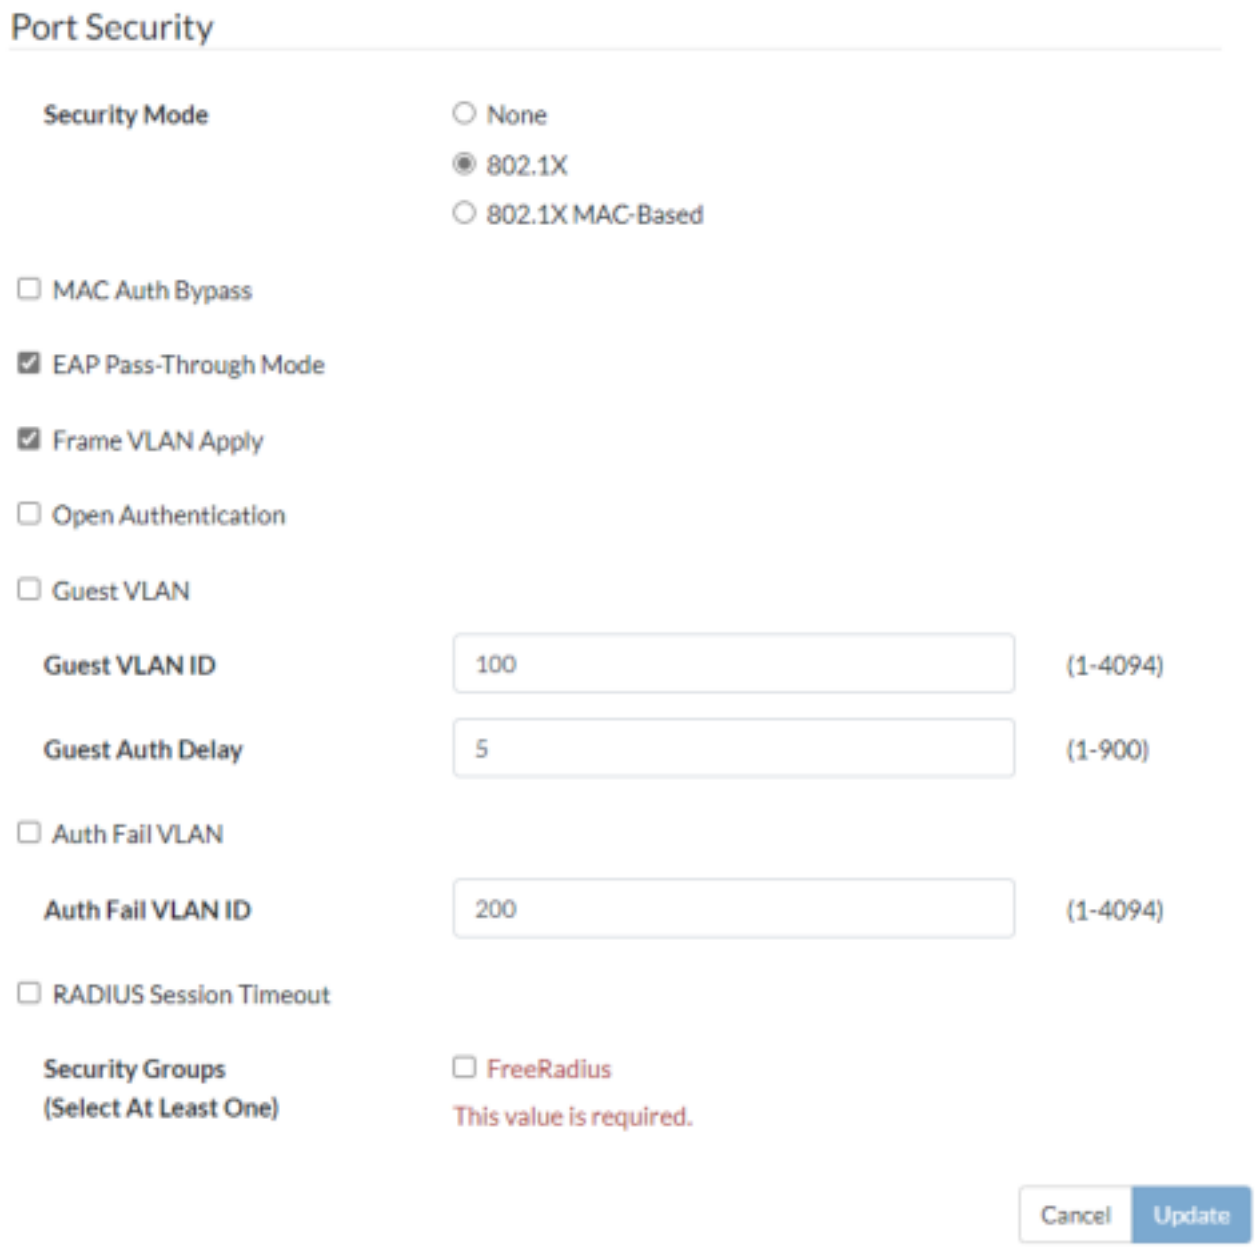

4. Select *MAC Auth Bypass*.
5. Select *EAP Pass-Through Mode*.
   NOTE: *EAP Pass-Through Mode* is enabled by default, which is the recommended setting. If the RADIUS authentication server does not support EAP-TLS, the *EAP Pass-Through Mode* needs to be disabled.
6. Select *Frame VLAN Apply* to apply the EAP/MAB frame VLAN to the port native VLAN.
   NOTE: For phone and PC configuration only, clear the checkbox to preserve the native VLAN when the data traffic is expected to be untagged.
7. Select *Open Authentication* to enable open authentication (monitor mode) on this interface. Use the monitor mode to test your system configuration for 802.1X authentication. You can use monitor mode to test port-based authentication, MAC-based authentication, EAP pass-through mode, and MAC authentication bypass. After you enable monitor mode, the network traffic will continue to flow, even if the users fail authentication.

8. Select *Guest VLAN* if you want to assign a VLAN to unauthorized users. If you select *Guest VLAN*, enter the guest VLAN identifier in the *Guest VLAN ID* field and enter the number of seconds for an unauthorized user to have access as a guest before authorization fails in the *Guest Auth Delay* field.
9. Select *Auth Fail VLAN* if you want to assign a VLAN to users who attempted to authenticate but failed to provide valid credentials. If you select *Auth Fail VLAN*, enter the VLAN identifier in the *Auth Fail VLAN ID* field.
10. If you want to use the RADIUS-provided reauthentication time, select *RADUS Session Timeout*.
11. If you are using port-based authentication or MAC-based authentication, select one or more security groups.
12. Select *Update*.

Using the CLI (for FSR-124D, 200 Series, FS-4xxE, 500 Series, FS-1024D, FS-1024E, FS-T1024E, FS-T1024F-FPOE, FS-1048E, and FS-3032E):

```
config switch interface
   edit <port>
      config port-security
         set allow-mac-move-to {disable | enable}
         set eap-egress-tagged {disable | enable}
         set port-security-mode {none | 802.1X | 802.1X-mac-based}
            set framevid-apply {disable | enable}
            set auth-fail-vlan {enable | disable}
            set auth-fail-vlanid <vlanid>
            set auth-priority {MAB-dot1x | dot1x-MAB | legacy}
            set authserver-timeout-period <3-15>
            set authserver-timeout-vlan {enable | disable}
            set authserver-timeout-vlanid <1-4094>
            set eap-passthru {enable | disable}
            set guest-auth-delay <integer>
            set guest-vlan {enable | disable}
            set guest-vlanid <vlanid>
            set mac-auth-bypass {enable | disable}
            set open-auth {enable | disable}
            set radius-timeout-overwrite {enable | disable}
      end
      set security-groups <security-group-name>
   end
```

Using the CLI (for FS-108E, FS-108E-POE, FS-108E-FPOE, FS-108F, FS-108F-POE, FS-108F-FPOE, FS-124E, FS-124E-POE, FS-124E-FPOE, FS-148E, and FS-148E-POE):

```
config switch interface
   edit <port>
      config port-security
         set allow-mac-move-from {disable | enable}
         set eap-egress-tagged {disable | enable}
         set port-security-mode {none | 802.1X | 802.1X-mac-based}
            set framevid-apply {disable | enable}
            set auth-fail-vlan {enable | disable}
            set auth-fail-vlanid <vlanid>
            set auth-priority {MAB-dot1x | dot1x-MAB | legacy}
            set authserver-timeout-period <3-15>
            set authserver-timeout-vlan {enable | disable}
            set authserver-timeout-vlanid <1-4094>
            set eap-passthru {enable | disable}
            set guest-auth-delay <integer>
            set guest-vlan {enable | disable}
```

```
            set guest-vlanid <vlanid>
            set mac-auth-bypass {enable | disable}
            set open-auth {enable | disable}
            set radius-timeout-overwrite {enable | disable}
      end
      set security-groups <security-group-name>
   end
```

Using the CLI (for FS-124F, FS-124F-POE, FS-124F-FPOE, FS-148F, FS-148F-POE, and FS-148F-FPOE):

```
config switch global
   config port-security
      set allow-mac-move {disable | enable}
   end
end

config switch interface
   edit <port>
      config port-security
         set eap-egress-tagged {disable | enable}
         set port-security-mode {none | 802.1X | 802.1X-mac-based}
            set framevid-apply {disable | enable}
            set auth-fail-vlan {enable | disable}
            set auth-fail-vlanid <vlanid>
            set auth-priority {MAB-dot1x | dot1x-MAB | legacy}
            set authserver-timeout-period <3-15>
            set authserver-timeout-vlan {enable | disable}
            set authserver-timeout-vlanid <1-4094>
            set eap-passthru {enable | disable}
            set guest-auth-delay <integer>
            set guest-vlan {enable | disable}
            set guest-vlanid <vlanid>
            set mac-auth-bypass {enable | disable}
            set open-auth {enable | disable}
            set radius-timeout-overwrite {enable | disable}
      end
      set security-groups <security-group-name>
   end
```

Viewing the 802.1X details

You can view the 802.1X details in the GUI and CLI.

Using the GUI:

Go to *Switch* > *Monitor* > *802.1x* > *Interfaces* to see which interfaces have been configured for 802.1x authentication.

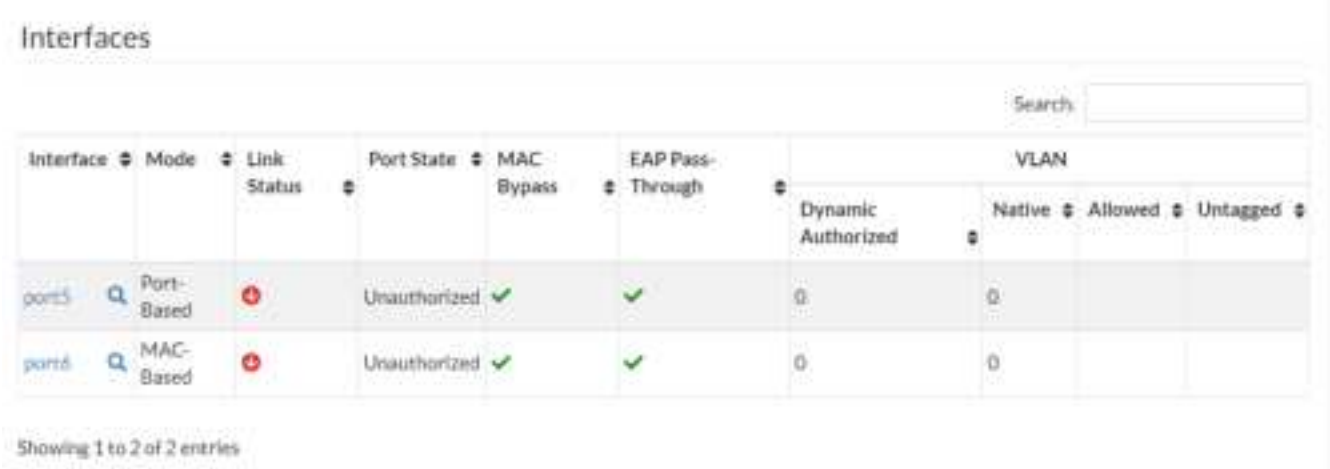

Go to *Switch* > *Monitor* > *802.1x* > *Sessions* to see a list of 802.1X MAC authenticated sessions.

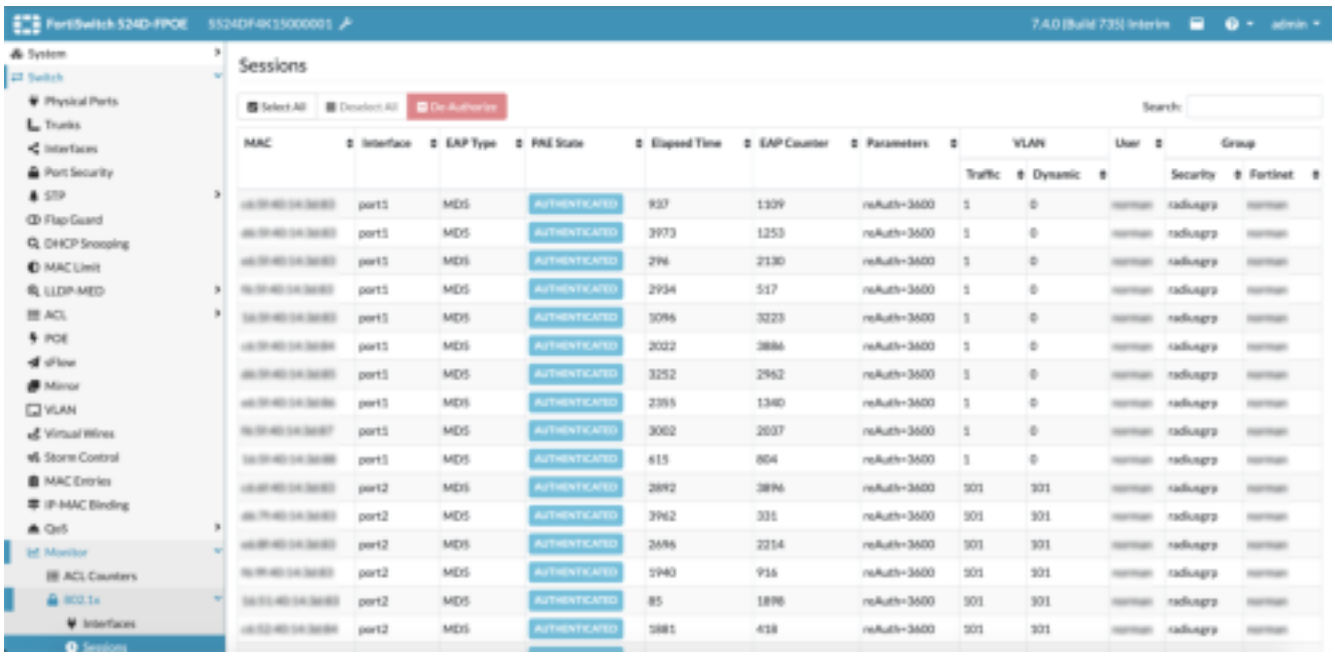

Enter a partial or complete MAC address in the *Search* field to find matching sessions.

Select one or more sessions and click *De-Authorize* to de-authorize the clients at the selected MAC addresses. You must click *De-Authorize* in the *Confirm De-Authorization* dialog before the clients are de-authorized.

Using the CLI:

Use the following command to show diagnostics on one or all ports:

```
diagnose switch 802-1x status [<port>]
```

For example:

```
diagnose switch 802-1x status port3

port3: Mode: mac-based (mac-by-pass enable)
       Link: Link up
```

```
Port State: authorized: ( )
Dynamic Allowed Vlan list: 101
Dynamic Untagged Vlan list: 101
EAP pass-through : Enable
Auth Order : MAB-dot1x
Auth Priority : Legacy
EAP egress-frame-tagged : Enable
EAP auto-untagged-vlans : Enable
Allow MAC Move : Disable
Dynamic Access Control List : Disable
Quarantine VLAN (4093) detection : Enable
Native Vlan : 101
Allowed Vlan list: 1-200
Untagged Vlan list: 101
Guest VLAN :
Auth-Fail Vlan :
AuthServer-Timeout Vlan :

Switch sessions 1/240, Local port sessions:1/20
Client MAC Type Traffic-Vlan Dynamic-Vlan
f0:4d:a2:be:a3:31 802.1x 101 101

Sessions info:
f0:4d:a2:be:a3:31 Type=802.1x,TTLS,state=AUTHENTICATED,etime=0,eap_cnt=9 params:reAuth=60
user="local-RADIUS",security_grp="radiusgrp",fortinet_grp="Radius_Admins"
```

Clearing authorized sessions

You can clear authorized sessions associated with a specific interface or a specific MAC address. When MAB is enabled, the authorized sessions are removed immediately. For EAP authentication, FortiSwitchOS will try to re-authenticate after clearing authorized sessions.

To use the GUI to clear the authorized sessions associated with an interface:

1. Go to *Switch* > *Interfaces*.
2. Select one or more ports that you want to clear the authorization from.
3. Select *Clear Auth*.

To use the CLI to clear the authorized sessions associated with an interface:

```
execute 802-1x clear interface {internal | <port_name>}
```

For example:

```
execute 802-1x clear interface port3
```

To use the CLI to clear the authorized session associated with a MAC address:

```
execute 802-1x clear mac <MAC_address>
```

For example:

```
execute 802-1x clear mac 00:21:cc:d2:76:72
```

Authenticating users with a RADIUS server

Using the GUI:

1. Define the RADIUS server:
   a. Go to *System* > *Authentication* > *RADIUS*.
   b. Click *Add Server*.

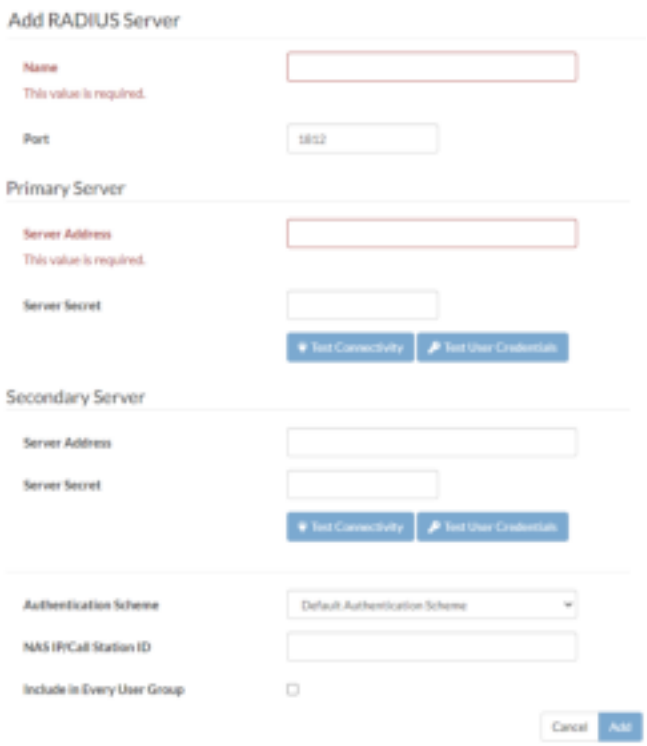

   c. In the *Name* field, enter a name for the RADIUS server.
   d. In the *Primary Server Address* field, enter the IP address for the RADIUS server.
   e. In the *Primary Server Secret* field, enter a password to use as a RADIUS key.
   f. Select *Add*.
2. Create a user group:
   a. Go to *System* > *User* > *Group*.

b. Select *Add Group*.

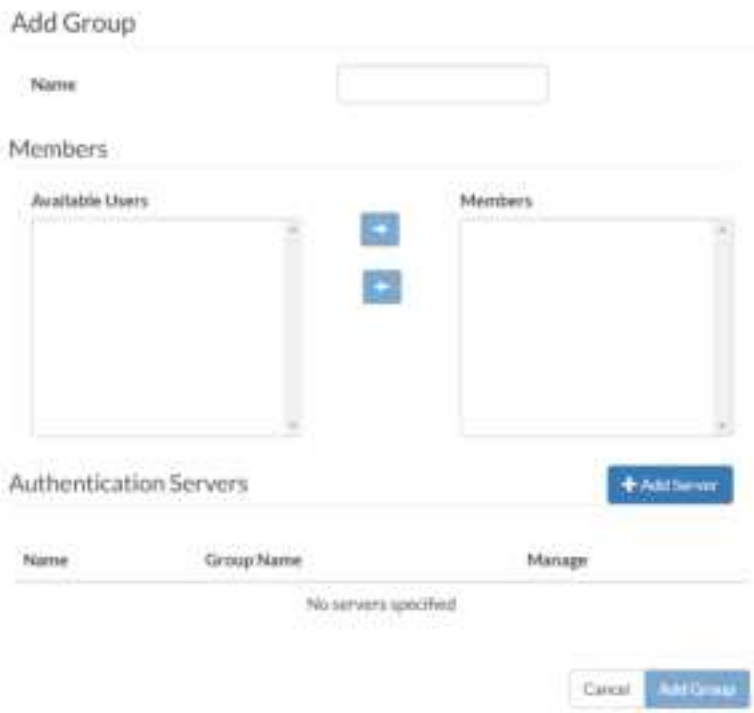

c. In the *Name* field, enter a name for the user group.
d. Select *Add Server*.
e. Select the name of the RADIUS server that you configured in step 1.
f. Select *Add Group*.

3. Configure the port security:
a. Go to *Switch* > *Interfaces*.
b. Select a port and then select *Edit*.
c. Select *802.1X* for port-based authentication or select *802.1X-MAC-based* for MAC-based authentication.

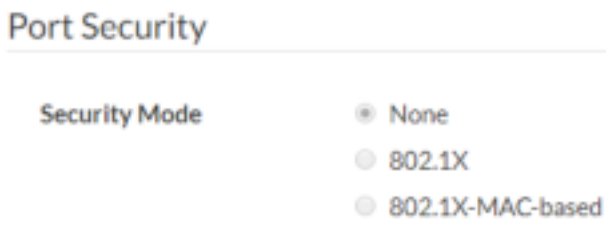

d. Select the user group that you configured in step 2.

e. Select *OK*.

Using the CLI:

1. Define an IPv4 or IPv6 RADIUS server:

```
config user radius
   edit <name>
      set addr-mode ipv4
      set server <IPv4_address>
      set source-ip <ipv4_address>
      set radius-port <radius_port_num>
      set secret <server_password>
      set auth-type {auto | chap | ms_chap | ms_chap_v2 | pap}
      set nas-ip <IPv4_address>
      set all-usergroup {enable | disable}
      set link-monitor {enable | disable}
      set link-monitor-interval <5-120 seconds>
   end
end

config user radius
   edit <name>
      set addr-mode ipv6
      set server <IPv6_address>
      set source-ip6 <ipv6_address>
      set radius-port <radius_port_num>
      set secret <server_password>
      set auth-type {auto | chap | ms_chap | ms_chap_v2 | pap}
      set nas-ip6 <IPv6_address>
      set all-usergroup {enable | disable}
      set link-monitor {enable | disable}
      set link-monitor-interval <5-120 seconds>
   end
end
```

2. Create a user group:

```
config user group
   edit <name>
      set member <list>
      config match
         edit 1
```

```
            set group-name <name>
            set server-name <name>
         end
      end
   end
end
```

3. Configure the switch interface for port-based or MAC-based 802.1X authentication:

```
config switch interface
   edit <interface>
      config port-security
         set port-security-mode 802.1X
      end
      set security-groups <security-group-name>
   end
end

config switch interface
   edit <interface>
      config port-security
         set port-security-mode 802.1X-mac-based
      end
      set security-groups <security-group-name>
   end
end
```

Example: RADIUS user group

Using the GUI:

1. Define the RADIUS server:
   a. Go to *System* > *Authentication* > *RADIUS*.
   b. Click *Add Server*.
   c. In the *Name* field, enter `FortiAuthenticator`.
   d. In the *Server Address* field under *Primary Server*, enter `10.160.36.190`.
   e. In the *Server Secret* field under *Primary Server*, enter `6rF7O4/Zf3p2TutNyeSjPbQc73QrS21wNDmNXd/rg9k6nTR6yMhBRsJGpArhle6UOCb7b8InM3nrCeuVETr/a02LpILmIltBq5sUMCNqbR6zp2fS3r35Eyd3IIrzmve4Vusi52c1MrCqVhzzy2EfxkBrx5FhcRQWxStvnVt4+dzLYbHZ`.

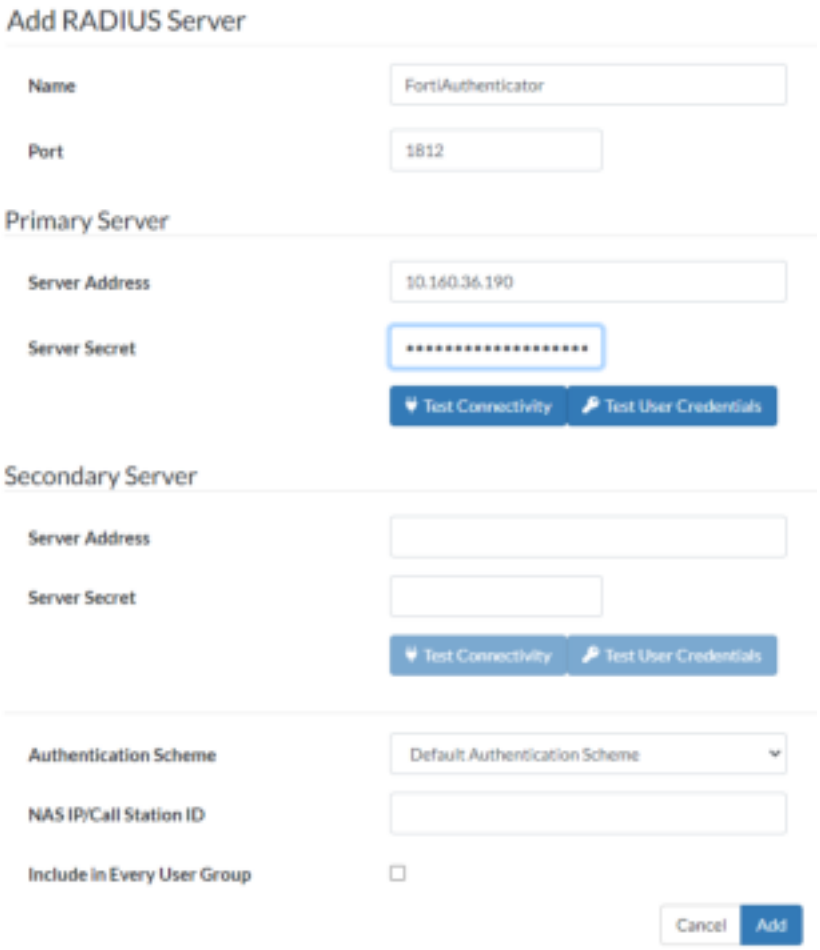

   f. Click *Add*.

2. Create a user group:
   a. Go to *System* > *User* > *Group*.
   b. Click *Add Group*.
   c. In the *Name* field, enter `Radius_group`.
   d. Click *Add Server*.

e. Select *FortiAuthenticator* as the authentication server.

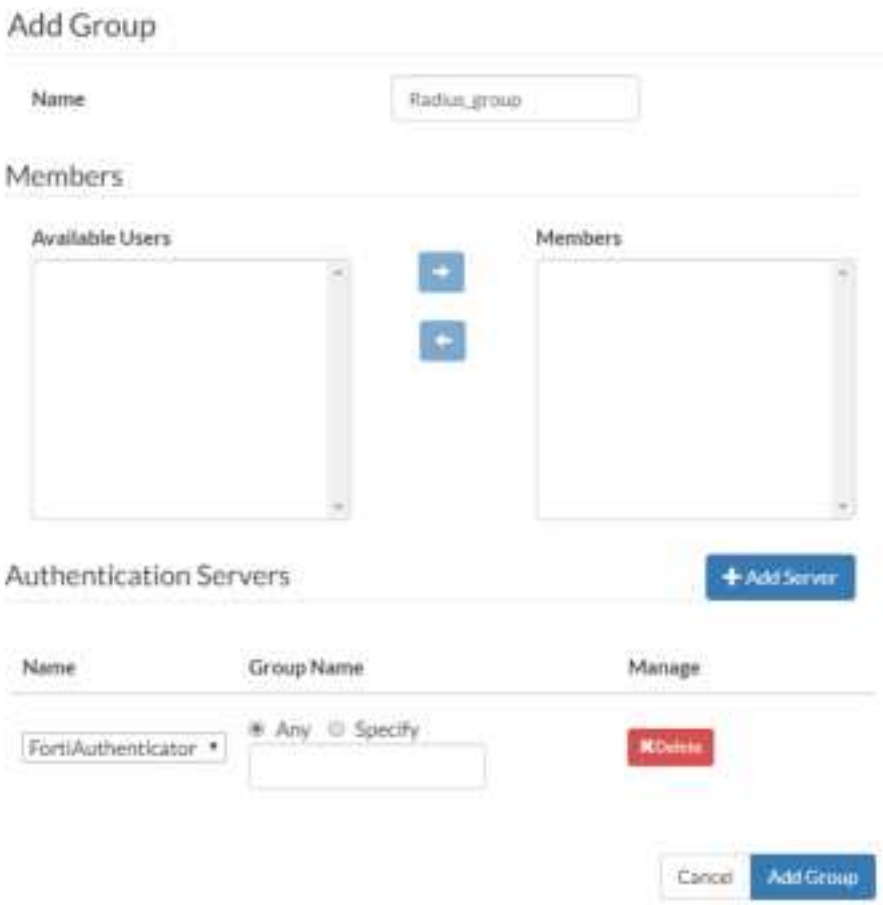

f. Click *Add*.

3. Configure the port security:

a. Go to *Switch* > *Interfaces*.

b. Select the *port1* row and then select *Edit*.

c. In the *Allowed VLANs* field, enter `1`.

d. Select *802.1X*.

e. Select *Radius_group*.

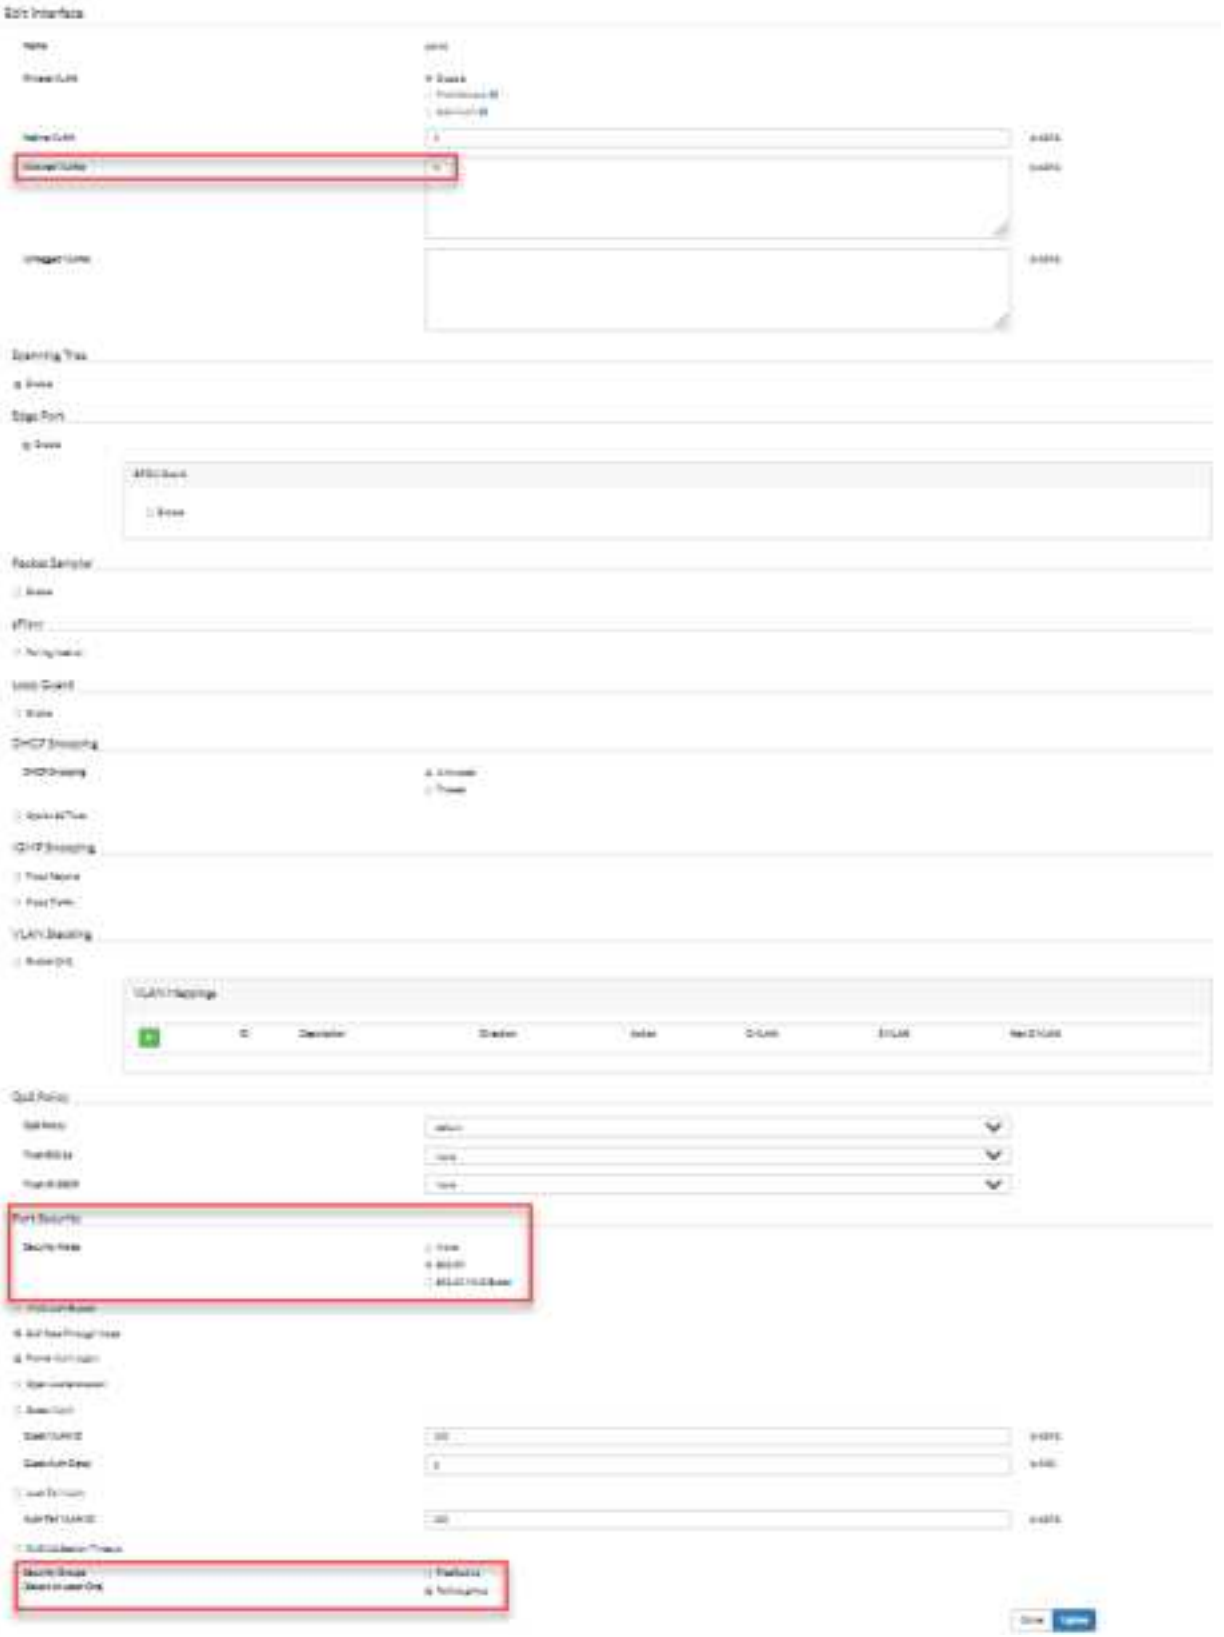

f. Click *Update*.

Using the CLI:

1. Define the RADIUS server:

```
config user radius
   edit "FortiAuthenticator"
      set secret ENC
           6rF7O4/Zf3p2TutNyeSjPbQc73QrS21wNDmNXd/rg9k6nTR6yMhBRsJGpArhle6UOCb7b8In
           M3nrCeuVETr/a02LpILmIltBq5sUMCNqbR6zp2fS3r35Eyd3IIrzmve4Vusi52c1MrCqVhzz
           y2EfxkBrx5FhcRQWxStvnVt4+dzLYbHZ
      set server "10.160.36.190"
      set addr-mode ipv4
   next
end
```

2. Create a user group:

```
config user group
   edit "Radius_group"
      set member "FortiAuthenticator"
   end
end
```

3. Configure the port security:

```
config switch interface
   edit "port1"
      set allowed-vlans 1
      config port-security
         set port-security-mode 802.1X
      end
      set security-groups "Radius_group"
   end
end
```

Example: dynamic VLAN

To assign VLAN dynamically for a port on which a user is authenticated, configure the RADIUS server attributes to return the VLAN ID when the user is authenticated. Assuming that the port security mode is set to 802.1X, the FortiSwitch unit will change the native VLAN of the port to the value returned by the server.

Ensure that the following attributes are configured on the RADIUS server:

- Tunnel-Private-Group-Id <integer or string> (the VLAN ID or VLAN name)
- Tunnel-Medium-Type IEEE-802 (6)
- Tunnel-Type VLAN (13)

NOTE: If the Tunnel-Private-Group-Id is set to the VLAN name, the same string must be specified in the `set description` command under the `config switch vlan` command.

Authenticating an admin user with RADIUS

If you want to use a RADIUS server to authenticate administrators, you must configure the authentication before you create the administrator accounts. Do the following:

1. Configure the FortiSwitch unit to access the RADIUS server.
2. Configure an administrator to authenticate with a RADIUS server and match the user secret to the RADIUS server entry.
3. Create the RADIUS user group.

Using the GUI:

1. Create a RADIUS system admin group:
   a. Go to *System* > *Admin* > *Administrators*.
   b. Click *Add Administrator*.
   c. In the *Name* field, enter `RADIUS_Admins`.
   d. Select *Remote*.
   e. For the user group, select *Radius_group*.
   f. Select *Wildcard*.
   g. For the admin profile, select *super_admin*.

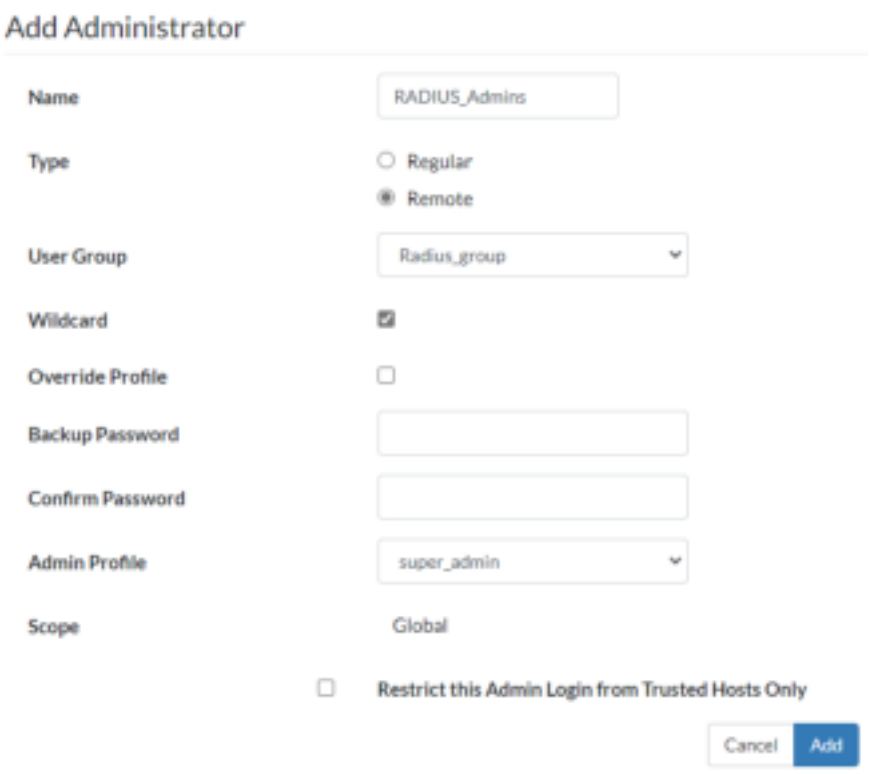

   h. Click *Add*.
2. Create a user:
   a. Go to *System* > *User* > *Definition*.
   b. Click *Add User*.

c. In the *Username* field, enter `RADIUS1`.
d. Select *Password* from the *Type* field.
e. In the *Password* field and *Confirm Password* field, enter `6rF7O4/Zf3p2TutNyeSjPbQc73QrS21wNDmNXd/rg9k6nTR6yMhBRsJGpArhle6UOCb7b8InM3nrCeuVETr/a02LpILmIltBq5sUMCNqbR6zp2fS3r35Eyd3IIrzmve4Vusi52c1MrCqVhzzy2EfxkBrx5FhcRQWxStvnVt4+dzLYbHZ`.

f. Click *Add*.

3. Create a user group:
a. Go to *System* > *User* > *Group*.
b. Click *Add Group*.
c. In the *Name* field, enter *RADIUS_Admins*.
d. Select *RADIUS1* in the Available Users box and select the right arrow to move it to the Members box.

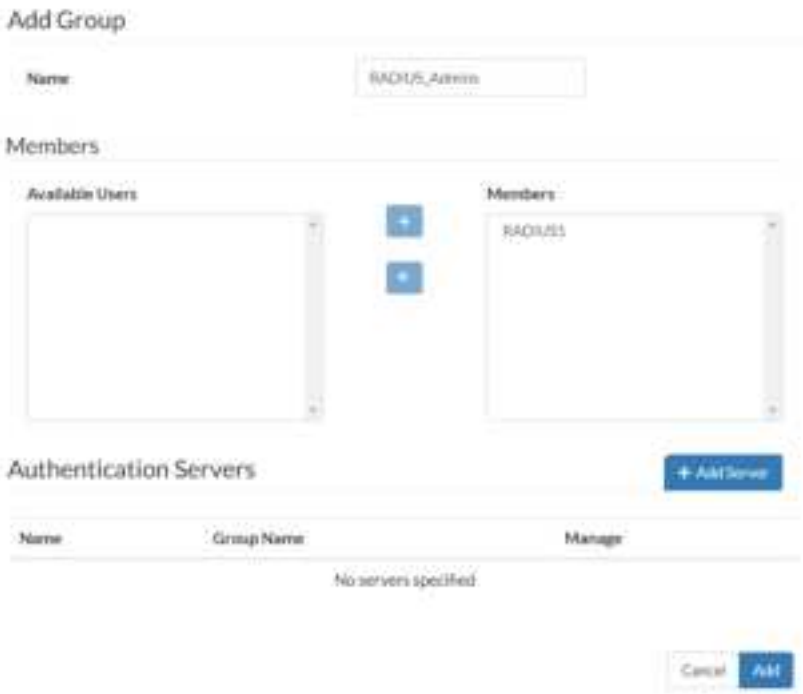

e. Click *Add*.

Using the CLI:

1. Create a RADIUS system admin group:

```
config system admin
   edit "RADIUS_Admins"
      set remote-auth enable
      set accprofile "super_admin"
      set wildcard enable
      set remote-group "RADIUS_Admins"
   next
end
```

2. Create a user:

```
config user radius
   edit "RADIUS1"
      set secret ENC
           6rF7O4/Zf3p2TutNyeSjPbQc73QrS21wNDmNXd/rg9k6nTR6yMhBRsJGpArhle6UOCb7b8InM3n
           rCeuVETr/a02LpILmIltBq5sUMCNqbR6zp2fS3r35Eyd3IIrzmve4Vusi52c1MrCqVhzzy2Efxk
           Brx5FhcRQWxStvnVt4+dzLYbHZ
      set addr-mode ipv4
   next
end
```

3. Create a user group:

```
 config user group
  edit "RADIUS_Admins"
     set member "RADIUS1"
```

```
      next
    end
```

RADIUS accounting and FortiGate RADIUS single sign-on

NOTE: To obtain a valid Framed-IP-Address attribute value, you need to manually configure DHCP snooping in the 802.1X-authenticated ports of your VLAN network for both port and MAC modes.

You can use your FortiSwitch unit for RADIUS single sign-on (RSSO) in two modes:

- Standalone mode
- FortiLink mode (FortiSwitch unit managed by FortiGate unit)

The FortiSwitch unit uses 802.1X-authenticated ports to send five types of RADIUS accounting messages to the RADIUS accounting server to support FortiGate RADIUS single sign-on:

- START—The FortiSwitch unit has been successfully authenticated, and the session has started.
- STOP—The FortiSwitch session has ended.
- INTERIM—Periodic messages sent based on the value set using the `set acct-interim-interval` command.
- ON—The FortiSwitch unit will send this message when the switch is turned on.
- OFF—The FortiSwitch unit will send this message when the switch is shut down.

NOTE: Starting in FortiSwitchOS 6.2.1, RADIUS accounting and CoA now support EAP and MAB 802.1X authentication.

Configuring the RADIUS accounting server and FortiGate RADIUS single sign-on

Use the following commands to set up RADIUS accounting and enable a FortiSwitch unit to receive CoA and disconnect messages from the RADIUS server:

```
config user radius
  edit <RADIUS_server_name>
    set acct-interim-interval <seconds>
    set secret <secret_key>
    set server <domain_ipv4_ipv6>
    set addr-mode {ipv4 | ipv6}
    set source-ip <ipv4_addr>
    set source-ip6 <ipv6_addr>
    config acct-server
      edit <entry_ID>
        set status {enable | disable}
        set server <accounting_server>
        set secret <secret_key>
        set port <port_number>
      next
    end
  next
end
```

Variable	Description
<RADIUS_server_name>	Enter the name of the RADIUS server that will be sending CoA and disconnect messages to the FortiSwitch unit. By default, the messages use port 3799.

Variable	Description
acct-interim-interval <seconds>	Enter the number of seconds between each interim accounting message sent to the RADIUS server. The value range is 60-86400. The default is 600.
addr-mode {ipv4 | ipv6}	Select whether to connect to the RADIUS server with IPv4 or IPv6. The default is IPv4.
secret <secret_key>	Enter the shared secret key for authentication with the RADIUS server.
server <domain_ipv4_ipv6>	Enter the domain name, IPv4 address, or IPv6 address for the RADIUS server. There is no default.
source-ip <ipv4_addr>	If the `addr-mode` was set to `ipv4`, enter the IPv4 address of the server that will be sending accounting messages. The default is 0.0.0.0.
source-ip6 <ipv6_addr>	If the `addr-mode` was set to `ipv6`, enter the IPv6 address of the server that will be sending accounting messages. There is no default.
<entry_ID>	Enter the entry identifier. The value range is 0-20.
status {enable | disable}	Enable or disable RADIUS accounting. The default is disable.
server <accounting_server>	Enter the domain name, IPv4 address, or IPv6 address of the RADIUS server that will be receiving the accounting messages. There is no default value.
secret <secret_key>	Enter the shared secret key for the RADIUS accounting server.
port <port_number>	Enter the port number for the RADIUS accounting server to receive accounting messages from the FortiSwitch unit. The default is 1813.

Example: RADIUS accounting and single sign-on

Use the following commands to set up RADIUS accounting:

```
config user radius
   edit "local-RADIUS"
      set server 10.0.23.5
      set addr-mode ipv4
      set secret ENC
           LE8xetYYGiE0bkQpBDdH6acilwkYROCos7XK2q5cNPhu8sUDW9/fvkgE+fVURgZGEzTsndt41gb+K+zV9
           m+nXCnoUXqivzQdt1UNlMxgKXADnCpXuiY966aJsYigmW/AZ1IM5kweUxvuHK8eqJkkT0nl64c8DID/LM
           AcCTx6JMapRCBS
      set auth-type ms_chap_v2
      set acct-interim-interval 1200
      set source-ip 10.105.142.19
      config acct-server
         edit 1
            set status enable
            set server 10.0.23.5
            set secret ENC
                 LE8xetYYGiE0bkQpBDdH6acilwkYROCos7XK2q5cNPhu8sUDW9/fvkgE+fVURgZGEzTsndt41gb+
                 K+zV9m+nXCnoUXqivzQdt1UNlMxgKXADnCpXuiY966aJsYigmW/AZ1IM5kweUxvuHK8eqJkkT0nl
                 64c8DID/LMAcCTx6JMapRCBS
```

```
            set port 1813
        next
    end
  next
end
```

RADIUS change of authorization (CoA)

NOTE: For increased security, each subnet interface that will be receiving CoA requests must be configured with the `set allowaccess radius-acct` command.

NOTE: Starting in FortiSwitchOS 6.2.1, RADIUS accounting and CoA support EAP and MAB 802.1X authentication.

The FortiSwitch unit supports two types of RADIUS messages:

- CoA messages to change session authorization attributes (such as data filters and the session-timeout setting) during an active session. To change the session timeout for an authenticated session, the CoA-Request message needs to use the IEEE session-timeout attribute.
- Disconnect messages (DMs) to flush an existing session. For MAC-based authentication, all other sessions are unchanged, and the port stays up. For port-based authentication, only one session is deleted.

RADIUS CoA messages use the following Fortinet proprietary attribute:

```
Fortinet-Host-Port-AVPair 42 string
```

The format of the value is as follows:

Attribute	Value	Description
Fortinet-Host-Port-AVPair	action=bounce-port	The FortiSwitch unit disconnects all sessions on a port. The port goes down for 10 seconds and then up again.
Fortinet-Host-Port-AVPair	action=disable-port	The FortiSwitch unit disconnects all session on a port. The port goes down until the user resets it.
Fortinet-Host-Port-AVPair	action=reauth-port	The FortiSwitch unit forces the reauthentication of the current session.

In addition, RADIUS CoA uses the following attributes:

Attribute	Value	Description
session-timeout	<session_timeout_value>	The FortiSwitch unit disconnects a session after the specified number of seconds of idleness. This value must be more than 60 seconds. **NOTE:** To use the session-timeout attribute, you must enable the `set radius-timeoutoverwrite` command first.
Tunnel-Private-Group-Id	VLAN ID or name (10)	This attribute requires FortiSwitchOS 6.4.12 or 7.2.2 or higher.

Attribute	Value	Description
Tunnel-Medium-Type	IEEE-802 (6)	This attribute requires FortiSwitchOS 6.4.12 or 7.2.2 or higher.
Tunnel-Type	VLAN (13)	This attribute requires FortiSwitchOS 6.4.12 or 7.2.2 or higher.

The FortiSwitch unit sends the following Error-Cause codes in RADIUS CoA-NAK and Disconnect-NAK messages:

Error Cause	Error Code	Description
Unsupported Attribute	401	This error is a fatal error, which is sent if a request contains an attribute that is not supported.
NAS Identification Mismatch	403	This error is a fatal error, which is sent if one or more NAS-Identifier Attributes do not match the identity of the NAS receiving the request.
Invalid Attribute Value	407	This error is a fatal error, which is sent if a CoA-Request or Disconnect-Request message contains an attribute with an unsupported value.
Session Context Not Found	503	This error is a fatal error if the session context identified in the CoA-Request or Disconnect-Request message does not exist on the NAS.

Configuring CoA and disconnect messages

Use the following commands to enable a FortiSwitch unit to receive CoA and disconnect messages from a RADIUS server:

```
config system interface
   edit "mgmt"
      set ip <address> <netmask>
      set allowaccess <access_types>
      set type physical
   next
config user radius
   edit <RADIUS_server_name>
      set radius-coa {enable | disable}
      set radius-port <port_number>
      set secret <secret_key>
      set server <server_name_ipv4_ipv6>
      set addr-mode {ipv4 | ipv6}
   end
```

Variable	Description
config system interface	
ip <address> <netmask>	Enter the interface IP address and netmask.

Variable	Description
allowaccess <access_types>	Enter the types of management access permitted on this interface. Valid types are as follows: `http https ping snmp ssh telnet radius-acct`. Separate each type with a space. You must include `radius-acct` to receive CoA and disconnect messages.
<RADIUS_server_name>	Enter the name of the RADIUS server that will be sending CoA and disconnect messages to the FortiSwitch unit. By default, the messages use port 3799.
config user radius	
radius-coa {enable | disable}	Enable or disable whether the FortiSwitch unit will accept CoA and disconnect messages. The default is disable.
radius-port <port_number>	Enter the RADIUS port number. By default, the value is 1812.
secret <secret_key>	Enter the shared secret key for authentication with the RADIUS server.
server <server_name_ipv4_ipv6>	Enter the domain name, IPv4 address, or IPv6 address for the RADIUS server. There is no default.
addr-mode {ipv4 | ipv6}	Select whether to connect to the RADIUS server with IPv4 or IPv6.

Example: RADIUS CoA

The following example enables the FortiSwitch unit to receive CoA and disconnect messages from the specified RADIUS server:

```
config system interface
   edit "mgmt"
      set ip 10.105.4.14 255.255.255.0
      set allowaccess ping https http ssh snmp telnet radius-acct
      set type physical
   next
config user radius
   edit "Radius-188-200"
      set radius-coa enable
      set secret ENC
           +2NyBcp8JF3/OijWl/w5nOC++aDKQPWnlC8Ug2HKwn4RcmhqVYE+q07yI9eSDhtiIw63kR/oMBLGwFQoe
           ZfOQWengIlGTb+YQo/lYJn1V3Nwp9sdkcblfyayfc9gTeqe+mFltKl5IWNI7WRYiJC8sxaF9Iyr2/l4hp
           CiVUMiPOU6fSrj
      set server "10.105.188.200"
      set addr-mode ipv4
   next
end
```

Viewing the CoA configuration

Use the following command to check the CoA settings:

```
S524DF4K15000024 # diagnose user radius coa

90075.874 DAS: :radius_das_diag_handler:
RADIUS DAS Server List:
radius2:
```

```
Type: RADIUS_8021X, IP: 10.105.252.79,
Last CoA/DM Client IP Addr    : 10.105.252.79
Disc Reqs     : 2
Disc ACKs     : 1
Disc NAKs     : 1
CoA  Reqs     : 0
CoA  ACKs     : 0
CoA  NAKs     : 0
radius3:
Type: RADIUS_8021X, IP: 10.105.252.76,
Last CoA/DM Client IP Addr    :
Disc Reqs     : 0
Disc ACKs     : 0
Disc NAKs     : 0
CoA  Reqs     : 0
CoA  ACKs     : 0
CoA  NAKs     : 0
```

Use cases

Here are three use cases for 802.1X authentication.

Use case 1

In this use case, a Cisco phone uses MAB and uses LLDP-MED to assign the voice VLAN. A PC behind the Cisco phone uses 802.1X authentication with or without dynamic VLAN assignment.

The following is an example configuration:

```
config switch lldp profile
   edit "lldp-cisco-104"
      set 802.1-tlvs port-vlan-id
      set 802.3-tlvs power-negotiation
         config med-network-policy
            edit "voice"
               set assign-vlan enable
               set status enable
               set vlan 104
            next
         set med-tlvs inventory-management network-policy
      next
   end

config switch physical-port
   edit "port1"
      set lldp-profile "lldp-cisco-104"
   next
end

config switch interface
   edit "port1"
      set native-vlan 20
      set security-groups "CISEGRP"
      set snmp-index 1
         config port-security
            set mac-auth-bypass enable // Required. You need to enable MAB.
```

```
            set port-security-mode 802.1X-mac-based // Required
         end
      next
   end
```

Use case 2

In this use case, the Cisco phone uses 802.1X authentication and uses LLDP-MED to assign the voice VLAN. A PC behind the Cisco phone uses 802.1X authentication without dynamic VLAN assignment.

RADIUS dynamic VLAN assignment for the voice VLAN must match the voice VLAN configured in the LLDP-MED profile for Cisco phone 802.1X authentication.

The following is an example configuration:

```
config switch lldp profile
   edit "lldp-cisco-104"
      set 802.1-tlvs port-vlan-id
      set 802.3-tlvs power-negotiation
         config med-network-policy
            edit "voice"
               set assign-vlan enable
               set status enable
               set vlan 104
            next
         set med-tlvs inventory-management network-policy
      next
   end

config switch physical-port
   edit "port1"
      set lldp-profile "lldp-cisco-104"
   next
end

config switch interface
   edit "port1"
      set native-vlan 20
      set security-groups "CISEGRP"
      set snmp-index 1
         config port-security
            set mac-auth-bypass disable // Optional
            set eap-auto-untagged-vlans disable // Required. Needed to allow voice traffic
                 with voice VLAN tag at egress
            set port-security-mode 802.1X-mac-based // Required
         end
      next
   end
```

Use case 3

In this use case, the Cisco phone uses 802.1X authentication and uses LLDP-MED to assign the voice VLAN. The PC behind the Cisco phone uses 802.1X authentication with dynamic VLAN assignment.

RADIUS dynamic VLAN assignment for the voice VLAN has to match the voice VLAN configured in the LLDP-MED profile for Cisco phone 802.1X authentication.

The VLAN ID from the RADIUS dynamic VLAN assignment for the PC has to be added in the untagged VLAN list on the port.

The following is an example configuration:

```
config switch lldp profile
   edit "lldp-cisco-104"
      set 802.1-tlvs port-vlan-id
      set 802.3-tlvs power-negotiation
         config med-network-policy
            edit "voice"
               set assign-vlan enable
               set status enable
               set vlan 104
            next
         set med-tlvs inventory-management network-policy
      next
   end

config switch physical-port
   edit "port1"
      set lldp-profile "lldp-cisco-104"
   next
end

config switch interface
   edit "port1"
      set native-vlan 20
      set allowed-vlans 50 60 70 // Assume that VLANs 50, 60, and 70 are a part of the
           dynamic VLANs configured on RADIUS for PCs in different groups.
      set untagged-vlans 50 60 70
      set security-groups "CISEGRP"
      set snmp-index 1
         config port-security
            set mac-auth-bypass disable // Optional
            set eap-auto-untagged-vlans disable // Required. Needed to allow voice traffic
                 with voice VLAN tag at egress
            set port-security-mode 802.1X-mac-based // Required
         end
      next
   end
```

Detailed deployment notes

- The learning-limit feature (`set learning-limit` under the `config switch interface` command) cannot be used with 802.1X authentication on the same port.
- Using more than one security group (with the `set security-groups` command) per security profile is not supported.
- CoA and single sign-on are supported only by the CLI in this release.
- RADIUS CoA is supported in standalone mode and in non-NAT FortiLink mode.
- The FortiSwitch unit supports using FortiAuthenticator, FortiConnect, Microsoft Network Policy Server (NPS), Aruba ClearPass, and Cisco Identity Services Engine (ISE) as the RADIUS server for CoA and RSSO.
- Each RADIUS CoA server can support only one accounting manager in this release.
- RADIUS accounting/CoA/VLAN-by-name features are supported only with `eap-passthru enable`.

- Fortinet recommends a unique secret key for each accounting server.
- For CoA to correctly function with FortiAuthenticator or FortiConnect, you must include the User-Name attribute (you can optionally include the Framed-IP-Address attribute) *or* the User-Name and Calling-Station-ID attributes in the CoA request.
- To obtain a valid Framed-IP-Address attribute value, you need to manually configure DHCP snooping in the 802.1X-authenticated ports of your VLAN network for both port and MAC modes.
- Port-based basic statistics for RADIUS accounting messages are supported in the Accounting Stop request.
- By default, the accounting server is disabled. You must enable the accounting server with the `set status enable` command.
- The default port for FortiAuthenticator single sign-on is 1813 for the FortiSwitch unit.
- In MAC-based authentication, the maximum number of client MAC addresses is 20. Each model has its own maximum limit.
- Static MAC addresses and sticky MAC addresses are mechanisms for manual/local authorization; 802.1X is a mechanism for protocol-based authorization. Do not mix them.

You can use static/sticky MAC addresses or 802.1X authentication but not both on the same port at the same time. If you do need to use both, you must ensure that the MAC addresses/devices authorized by 802.1X authentication are not included in the static-mac table.

- Fortinet recommends an 802.1X setup rate of 5 to 10 sessions per second.
- Starting in FortiSwitch 6.2.0, when 802.1X authentication is configured, the EAP pass-through mode (`set eap-passthru`) is enabled by default.
- For information about RADIUS attributes supported by FortiSwitchOS, refer to the “Supported attributes for RADIUS CoA and RSSO” appendix.
- The authentication and accounting server configuration must be in the same address mode within the same member. The address mode is either IPv4 or IPv6, no matter what the address mode is in the FQDN or raw IP address. The address mode cannot be mixed.
- When a client is authorized with the RADIUS timeout VLAN enabled, the client is placed in the authorization VLAN. If the RADIUS server becomes unavailable afterward and the reauthentication timer expires for the session, the device keeps the client in the authorization VLAN but the state changes from AUTHENTICATED to SERVER_TIMEOUT.
- In general for 802.1X deployment, Fortinet suggests disabling STP in the 802.1X security ports. If STP is enabled on the ports, the ports must be assigned to STP instances that belong to a dynamic VLAN, guest VLAN, or auth-fail VLAN; otherwise, the network connectivity fails after the ports are authorized and assigned to a dynamic VLAN, guest VLAN, or auth-fail VLAN.
- EAP-MD5 is not supported.

MAC security

Media Access Control security (MACsec) secures each switch-to-switch link by encrypting all network traffic within an Ethernet LAN.

You can either use the pre-shared key (PSK) mode or the dynamic connectivity association key (CAK) mode.

For the *PSK mode*, you specify the CAK and the connectivity association name (CKN) for the PSK in the MACsec profile and then apply the profile to a switch port.

For the *dynamic-CAK mode*, you use a port access entity (PAE) to allow the interface to act as a supplicant or an authenticator:

- When the interface acts as a supplicant, the PAE requests authentication from the RADIUS server before the interface can be accessed.
- When the interface acts as an authenticator, the PAE enforces the authentication, which is provided by the RADIUS server, before the interface can be accessed.

The RADIUS server generates the master session key (MSK), and the CAK is derived from the MSK. You can use the same MACsec profile for both the supplicant and the authenticator.

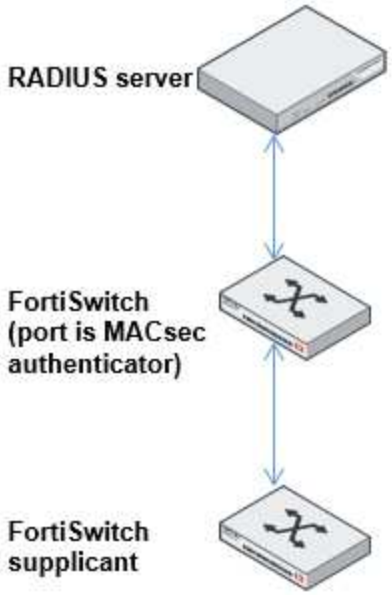

Starting in FortiSwitchOS 7.4.1, there is flexibility to exclude one or more protocols that will not be secured by the MACsec traffic policy. By default, all protocols are encrypted. You can use the CLI to exclude ARP, 802.1q VLAN, FortiLink, IPv4, IPv6, LACP, LLDP, 802.1ad QinQ, and STP packets.

Starting in FortiSwitchOS 7.4.4, MACsec supports AWS Direct Connect in the CLI.

Starting in FortiSwitchOS 7.4.4, you can use the `set mka-sak-rekey-time` command to set the number of seconds before a new secure association key (SAK) is generated. You can set the timer from 60 to 1,000,000 seconds. Set the command to 0 to disable the timer.

To apply a MACsec profile to a port, you need to specify the port and MACsec profile, set the port security mode to `macsec`, and select the MACsec PAE mode:

- If you want to use the PSK mode, select `none` because you do not need the PAE.
- If you want to use the dynamic-CAK mode, select `supp` to have the PAE request authentication from the RADIUS server before the interface can be accessed or select `auth` to have the PAE enforce authentication before the interface can be accessed.

Fortinet recommends that, if you want to change the MACsec profile and MACsec is up and running, shut down the port first. You do not need to shut down the port during the initial MACsec configuration.

Notes:

- SNMP is not supported.
- The `security-mode` must be set to `macsec` for each interface that you want to apply MACsec to.
- The MACsec profile must be applied at the port level.
- FortiSwitchOS supports PSK mode and dynamic-CAK mode. Static secure association key (SAK) mode is not supported.
- The FS-5xxD models only support the GCM-AES-128 cipher suite.

This section covers the following topics:

- Configuring PSK-mode MACsec on page 228
- Configuring dynamic-CAK MACsec on page 232
- Viewing the MACsec details on page 237
- Clearing the MACsec statistics on page 237
- Resetting the MACsec statistics on page 237
- Changing the link status on page 237
- Avoiding traffic oversubscription on page 238

Configuring PSK-mode MACsec

To configure PSK-mode MACsec:

1. Create the MACsec profile, either in the GUI or CLI.
2. Apply the MACsec profile to a port, either in the GUI or CLI.

To create a MACsec profile for PSK mode using the GUI:

1. Go to *Switch* > *MAC Security*.

2. Click *Add MACsec Profile*.

3. In the *Name* field, enter a name for your MACsec profile.
4. By default, the MACsec profile is enabled. If you want to disable the profile, clear the *Enabled* checkbox.
5. By default, the *MACsec Mode* is set to *Static CAK*.
6. By default, MACsec traffic is encrypted. If you do not want to encrypt the MACsec traffic, clear the *Encrypt Traffic* checkbox.
7. In the *MKA Priority* field, enter the MACsec MKA priority.
8. In the *Confidential Offset* dropdown list, select the number of bytes for the MACsec traffic confidentiality offset.

   Selecting *0 Bytes* means that all of the MACsec traffic is encrypted. Selecting *30 Bytes* or *50 Bytes* means that the first 30 or 50 bytes of MACsec traffic are not encrypted. If you selected *GCM-AES-XPN-128* or *GCM-AES-XPN-256* for the cipher suite, you must select *0 Bytes* for the confidentiality offset.
9. Select the *Replay Protect* checkbox if you want to drop packets that arrive out of sequence.
10. If you selected the *Replay Protect* checkbox, enter the number of packets for the MACsec replay window size in the *Replay Window* field.

    If two packets arrive with the difference between their packet identifiers more then the replay window size, the most recent packet of the two is dropped. Enter 0 to ensure that all packets arrive in order without any repeats.
11. In the *Cipher Suite* dropdown list, select *GCM-AES-128*, *GCM-AES-256*, *GCM-AES-XPN-128*, or *GCM-AES-XPN-256*.
12. By default, the MACsec Key Agreement (MKA) integrity check value (ICV) indicator is always included.

13. By default, the MACsec transmit secure channel identifier (SCI) is included. If you do not want the MACsec transmit SCI included, clear the *Include MACsec SCI* checkbox.
14. By default, the MACsec validation is always strict.
15. If you want to configure the MACsec MKA pre-shared key, click + in the *MKA PSK* area.
    a. Enter a name for the MACsec MKA pre-shared key configuration.
    b. Enter a 32-byte or 64-byte hexadecimal string for the CAK.
    c. Enter the string of hexadecimal digits for the CKN.
       The string can be 2-bytes to 64-bytes long.
    d. By default, the AES_128_CMAC algorithm is used for encrypting the pre-shared key.
16. In the *Traffic Policies* area, enter a name for this MACsec traffic policy.
17. By default, the traffic policy must secure traffic for MACsec, and the status of this MACsec traffic policy is enabled.
18. Click *Add*.

To create a MACsec profile for PSK mode using the CLI:

```
config switch macsec profile
   edit <MACsec_profile_name>
      set cipher_suite {GCM-AES-128 | GCM-AES-256 | GCM-AES-XPN-128 | GCM-AES-XPN-256}
      set confident-offset {0 | 30 | 50}
      set encrypt-traffic {enable | disable}
      set include-macsec-sci {enable | disable}
      set include-mka-icv-ind enable
      set macsec-mode static-cak
      set macsec-validate strict
      set mka-priority <0-255>
      set mka-sak-rekey-time {0 | 60-1000000}
      set replay-protect {enable | disable}
      set replay-window <0-16777215>
      set status {enable | disable}
      config mka-psk
         edit <pre-shared key name>
            set crypto-alg {AES_128_CMAC | AES_256_CMAC}
            set mka-cak <string>
            set mka-ckn <string>
            set status active
         next
      end
      config traffic-policy
         edit <traffic_policy_name>
            set exclude-protocol {arp | dot1q | fortilink | ipv4 | ipv6 | lacp | lldp | qinq
                 | stp}
            set security-policy must-secure
            set status enable
         next
      end
   next
end
```

For example:

```
config switch macsec profile
   edit "staticcak"
      set cipher_suite GCM-AES-128
      set confident-offset 0
```

```
      set encrypt-traffic enable
      set include-macsec-sci enable
      set include-mka-icv-ind enable
      set macsec-mode static-cak
      set macsec-validate strict
      set mka-priority 199
      config mka-psk
         edit "2"
            set crypto-alg AES_128_CMAC
            set mka-cak "0123456789ABCDEF0123456789ABCDEE"
            set mka-ckn "6162636465666768696A6B6C6D6E6F707172737475767778797A303132333436"
            set status active
         next
      end
      set replay-protect disable
      set replay-window 32
      set status enable
      config traffic-policy
         edit "2"
            set exclude-protocol stp
            set security-policy must-secure
            set status enable
         next
      end
   next
end
```

To apply a PSK-mode MACsec profile to a port in the GUI:

1. Go to *Switch* > *Physical Ports*.
2. Select a port and then click *Edit*.
3. Under *MACsec Profile*, click the name of the MACsec profile that you want to apply to this port.
4. Make certain that the *MACsec PAE Mode* is *None*.
5. Click *Update*.

To apply a PSK-mode MACsec profile to a port in the CLI:

```
config switch physical-port
   edit <port_name>
      set security-mode macsec
      set macsec-pae-mode none
      set macsec-profile <MACsec_profile_name>
   next
end
```

For example:

```
config switch physical-port
   edit port49
      set security-mode macsec
      set macsec-pae-mode none
      set macsec-profile "macsec_profile1"
   next
end
```

Configuring dynamic-CAK MACsec

For the dynamic-CAK mode, you need to do the following before creating the MACsec profile:

- Specify the certificate authority (CA) for the MACsec CAK.
  This is used in the `set eap-tls-ca-cert <CA_certificate>` command in the MACsec profile.
- Import the client certificate to use for the MACsec CAK.
  This is used in the `set eap-tls-cert <client_certificate>` command in the MACsec profile.
- Configure the RADIUS server to use for MACsec CAK.
  This is used in the `set eap-tls-radius-server <name_of_RADIUS_server>` command in the MACsec profile.

To specify the CA for the MACsec CAK:

```
config system certificate ca
   edit <CA_name>
   next
end
```

For example:

```
config system certificate ca
   edit "MACsec_CA"
   next
end
```

To import the client certificate to use for the MACsec CAK:

```
config system certificate local
   edit <certificate_name>
      set password <string>
   next
end
```

For example:

```
config system certificate local
   edit "MACsec_certificate"
      set password ENC
           jVXPqKiU35+clW0peV401S3G3y1wIKPnU0203VPqOou5bZn3uUGem6YUipSHPqME5Lb83KS9n9AmFHhIf
           6AkZgaiXBLSWcEczSk95MuZcAl/rS1fl3DIJZ1ev3scj35gANo7bZZq16n+ufqP1QIE2RcUBmqF/ctCdA
           Uxn4BbUthahvj1
   next
end
```

To configure the RADIUS server used for MACsec CAK:

```
config user radius
   edit <RADIUS_server_name>
      set secret <server_password>
      set server <domain_ipv4_ipv6>
   next
end
```

For example:

```
config user radius
```

```
    edit "radiusserver"
      set secret ENC
           mAKp/cPwUvJktZnfY4aT3Xlz6n+hZEdhO1safouKrY2Vousxu9kGUGx9NEZWargxMQOfkF8GtVqjgrs0p
           GoS+dHZohqwK4HDtmHmoC9AYsgen9Vmzp1CI0N/5uMr+jjAHYFPdUdW6VBawCAGUeYtXOFL5l74Y5H+Q6
           zju3qqhE84DO0k
      set server "10.105.252.125"
    next
end
```

To create a MACsec profile for dynamic-CAK mode using the GUI:

1. Go to *Switch* > *MAC Security*.

2. Click *Add MACsec Profile*.
3. In the *Name* field, enter a name for your MACsec profile.
4. By default, the MACsec profile is enabled. If you want to disable the profile, clear the *Enabled* checkbox.
5. In the *MACsec Mode* dropdown list, select *Dynamic CAK*.

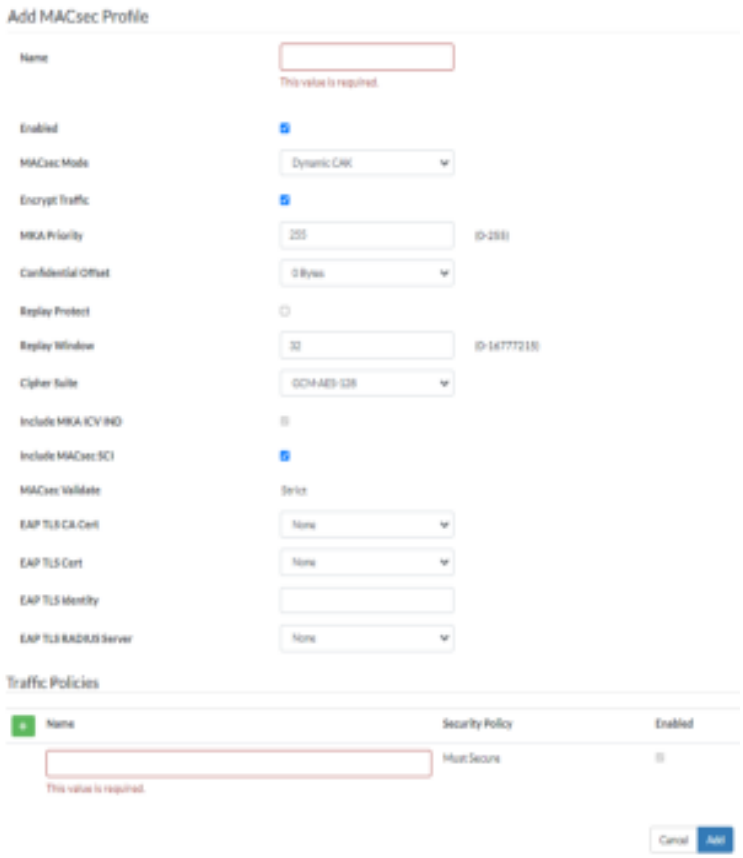

6. By default, MACsec traffic is encrypted. If you do not want to encrypt the MACsec traffic, clear the *Encrypt Traffic* checkbox.
7. In the *MKA Priority* field, enter the MACsec MKA priority.
8. In the *Confidential Offset* dropdown list, select the number of bytes for the MACsec traffic confidentiality offset.

   Selecting *0 Bytes* means that all of the MACsec traffic is encrypted. Selecting *30 Bytes* or *50 Bytes* means that the first 30 or 50 bytes of MACsec traffic are not encrypted. If you selected *GCM-AES-XPN-128* or *GCM-AES-XPN-256* for the cipher suite, you must select *0 Bytes* for the confidentiality offset.
9. Select the *Replay Protect* checkbox if you want to drop packets that arrive out of sequence.
10. If you selected the *Replay Protect* checkbox, enter the number of packets for the MACsec replay window size in the *Replay Window* field.

    If two packets arrive with the difference between their packet identifiers more then the replay window size, the most recent packet of the two is dropped. Enter 0 to ensure that all packets arrive in order without any repeats.
11. In the *Cipher Suite* dropdown list, select *GCM-AES-128*, *GCM-AES-256*, *GCM-AES-XPN-128*, or *GCM-AES-XPN-256*.
12. By default, the MKA ICV indicator is always included.
13. By default, the MACsec transmit SCI is included. If you do not want the MACsec transmit SCI included, clear the *Include MACsec SCI* checkbox.
14. By default, the MACsec validation is always strict.
15. In the *EAP TLS CA Cert* dropdown list, select the CA to use for the MACsec CAK.
16. In the *EAP TLS Cert* dropdown list, select the client certificate that you imported for the MACsec CAK.
17. In the *EAP TLS Identity* field, enter the name of the client for the MACsec CAK.
18. In the *EAP TLS RADIUS Server* dropdown list, select the name of the RADIUS server to use for the MACsec CAK.

19. In the *Traffic Policies* area, enter a name for this MACsec traffic policy.
20. By default, the traffic policy must secure traffic for MACsec, and the status of this MACsec traffic policy is enabled.
21. Click *Add*.

To create a MACsec profile for dynamic-CAK mode using the CLI:

```
config switch macsec profile
   edit <MACsec_profile_name>
      set cipher_suite {GCM-AES-128 | GCM-AES-256 | GCM-AES-XPN-128 | GCM-AES-XPN-256}
      set confident-offset {0 | 30 | 50}
      set eap-tls-ca-cert <CA_certificate>
      set eap-tls-cert <client_certificate>
      set eap-tls-identity <name_of_client>
      set eap-tls-radius-server <name_of_RADIUS_server>
      set encrypt-traffic {enable | disable}
      set include-macsec-sci {enable | disable}
      set include-mka-icv-ind enable
      set macsec-mode dynamic-cak
      set macsec-validate strict
      set mka-priority <0-255>
      set mka-sak-rekey-time {0 | 60-1000000}
      set replay-protect {enable | disable}
      set replay-window <0-16777215>
      set status {enable | disable}
      config traffic-policy
         edit <traffic_policy_name>
            set exclude-protocol {arp | dot1q | fortilink | ipv4 | ipv6 | lacp | lldp | qinq
                 | stp}
            set security-policy must-secure
            set status enable
         next
      end
   next
end
```

For example:

```
config switch macsec profile
   edit "dynamiccak"
      set cipher_suite GCM-AES-128
      set confident-offset 0
      set eap-tls-ca-cert "MACsec_CA"
      set eap-tls-cert "MACsec_certificate"
      set eap-tls-identity "macsecclient"
      set eap-tls-radius-server "radiusserver"
      set encrypt-traffic enable
      set include-macsec-sci enable
      set include-mka-icv-ind enable
      set macsec-mode dynamic-cak
      set macsec-validate strict
      set mka-priority 215
      set replay-protect disable
      set replay-window 32
      set status enable
      config traffic-policy
         edit "trafficpolicy1"
            set exclude-protocol qinq
```

```
            set security-policy must-secure
            set status enable
        next
      end
   next
end
```

To apply a dynamic-CAK MACsec profile to a port that will act as the PAE supplicant in the GUI:

1. Go to *Switch* > *Physical Ports*.
2. Select a port and then click *Edit*.
3. Under *MACsec Profile*, click the name of the MACsec profile that you want to apply to this port.
4. For the *MACsec PAE Mode*, click *Supplicant*.
5. Click *Update*.

To apply a dynamic-CAK MACsec profile to a port that will act as the PAE supplicant in the CLI:

```
config switch physical-port
   edit <port_name>
      set security-mode macsec
      set macsec-pae-mode supp
      set macsec-profile <MACsec_profile_name>
   next
end
```

For example:

```
config switch physical-port
   edit "port25"
      set macsec-pae-mode supp
      set macsec-profile "dynamiccak"
      set security-mode macsec
   next
end
```

To apply a dynamic-CAK MACsec profile to a port that will act as the PAE authenticator in the GUI:

1. Go to *Switch* > *Physical Ports*.
2. Select a port and then click *Edit*.
3. Under *MACsec Profile*, click the name of the MACsec profile that you want to apply to this port.
4. For the *MACsec PAE Mode*, click *Authenticator*.
5. Click *Update*.

To apply a dynamic-CAK MACsec profile to a port that will act as the PAE authenticator in the CLI:

```
config switch physical-port
   edit <port_name>
      set security-mode macsec
      set macsec-pae-mode auth
      set macsec-profile <MACsec_profile_name>
   next
end
```

For example:

```
config switch physical-port
   edit "port25"
      set macsec-pae-mode auth
      set macsec-profile "dynamiccak"
      set security-mode macsec
   next
end
```

Viewing the MACsec details

To view which MACsec profiles are available:

```
show switch macsec profile
```

To view a specific MACsec profile:

```
show switch macsec profile <MACsec_profile_name>
```

To view the MACsec status for a specific port:

```
diagnose switch macsec status <port_name>
```

To view the MACsec traffic statistics for a specific port:

```
diagnose switch macsec statistics <port_name>
```

Clearing the MACsec statistics

To clear all MACsec statistics on a single port:

```
execute macsec clearstat physical-port <port_name>
```

For example:

```
execute macsec clearstat physical-port port15
```

Resetting the MACsec statistics

To reset the MACsec session on a single port on the server side *or* the client side:

```
execute macsec reset physical-port <port_name>
```

For example:

```
execute macsec reset physical-port port15
```

Changing the link status

This command applies to the dynamic-CAK mode.

To change the link status and reset the MACsec session on a single port for both the server side *and* the client side:

```
execute macsec toggle physical-port <port_name>
```

For example:

```
execute macsec toggle physical-port port5
```

Avoiding traffic oversubscription

When the amount of traffic being sent on a network link is greater than the link's capacity, packets might be dropped or delayed, and MACsec sessions might be disrupted. In addition, traffic oversubscription can cause dropped MACsec sessions. For example:

- Traffic oversubscription might occur if traffic is sent from a 100G port to a 10G port and the 10G port is connected to the 10G MACsec physical signaling layer.
- A third-party switch port (10G or 100G) is connected to a FortiSwitch port (10G or 100G), which is connected to the MACsec physical signaling layer. The MACsec physical signaling layer adds a 32-byte header to every packet, causing traffic oversubscription and dropped egress packets.

There are three ways to avoid traffic oversubscription:

- Avoiding traffic oversubscription with a QoS egress policy on page 238
- Avoiding traffic oversubscription with network design on page 239
- Avoiding traffic oversubscription with egress traffic shaping on page 239

Avoiding traffic oversubscription with a QoS egress policy

You can avoid MACsec sessions being dropped by enabling a QoS egress policy. MACsec control packets go to the highest priority queue in the CPU port, which is the queue least likely to be dropped during traffic congestion.

In the following example, the user configures a QoS policy, `qos-strict`, and applies it to port1.

```
config switch qos qos-policy
   edit "qos-strict"
      config cos-queue
      edit "queue-0"
      next
      edit "queue-1"
      next
      edit "queue-2"
      next
      edit "queue-3"
      next
      edit "queue-4"
      next
      edit "queue-5"
      next
      edit "queue-6"
      next
      edit "queue-7"
      next
      end
      set rate-by kps
      set schedule strict
```

```
   next
end

config switch interface
   edit port1
      set native-vlan 100
      set stp-state disabled
      set auto-discovery-fortilink enable
      set snmp-index 25
      config port-security
         set allow-mac-move-to disable
         set macsec-pae-mode none
         set macsec-profile "static"
         set port-security-mode macsec
      end
      set qos-policy "qos-strict"
   next
end
```

Avoiding traffic oversubscription with network design

You can design your network to avoid traffic congestion. Because the MACsec physical signaling layer adds a 32-byte header to every packet, you cannot send the full line rate to the FortiSwitch unit.

For example, if you are sending packets from a third-party switch to a FortiSwitch unit:

Packet size sent by the third-party switch	Maximum percentage of line rate that packets from third-party switch are transmitted	Packets per second (received by the FortiSwitch 100G port)
64 bytes	72%	107.14 Mbps
128 bytes	82%	69.27 Mbps
256 bytes	89%	40.31 Mbps
1500 bytes	97%	7.98 Mbps

Avoiding traffic oversubscription with egress traffic shaping

You can avoid traffic oversubscription by applying egress traffic shaping on the switch chip. Egress traffic shaping drops traffic before the packets reach the MACsec physical signaling layer, which prevents traffic oversubscription.

In the following example, the user sets the maximum traffic rate (`set max-rate 52413793`) for queue 0 in the QoS policy, `qos-shaper`, and then applies the QoS policy to port1. The maximum traffic rate depends on the port rate and packet size. You might need to experiment with different values for the `set max-rate` command to find the best setting for your network.

```
config switch qos qos-policy
   edit "qos-shaper"
      config cos-queue
         edit “queue-0”
         set max-rate 52413793
         next
         edit “queue-1”
         next
```

```
            edit "queue-2"
            next
            edit "queue-3"
            next
            edit "queue-4"
            next
            edit "queue-5"
            next
            edit "queue-6"
            next
            edit "queue-7"
            next
        end
        set rate-by kps
        set schedule strict
    next
end

config switch interface
    edit port1
        set native-vlan 100
        set stp-state disabled
        set auto-discovery-fortilink enable
        set snmp-index 25
        config port-security
            set allow-mac-move-to disable
            set macsec-pae-mode none
            set macsec-profile "static"
            set port-security-mode macsec
        end
        set qos-policy "qos-shaper"
    next
end
```

STP

The FortiSwitch unit supports the following:

- Spanning Tree Protocol, a link-management protocol that ensures a loop-free layer-2 network topology
- Multiple Spanning Tree Protocol (MSTP), which is based on the IEEE 802.1s standard
- Per-VLAN Rapid Spanning Tree Protocol (also known as Rapid PVST or RPVST); RSTP is defined in the IEEE 802.1w standard

This section covers the following topics:

- MSTP overview and terminology on page 241
- MSTP configuration on page 243
- Interactions outside of the MSTP region on page 250
- Viewing the MSTP configuration on page 251
- Support for interoperation with Rapid per-VLAN RSTP (Rapid PVST+ or RPVST+) on page 251

MSTP overview and terminology

MSTP supports multiple spanning tree instances, where each instance carries traffic for one or more VLANs (the mapping of VLANs to instances is configurable).

MSTP is backward-compatible with STP and Rapid Spanning Tree Protocol (RSTP). A layer-2 network can contain switches that are running MSTP, STP, or RSTP.

MSTP is built on RSTP, so it provides fast recovery from network faults and fast convergence times.

Regions

A region is a set of interconnected switches that have the same multiple spanning tree (MST) configuration (region name, MST revision number, and VLAN-to-instance mapping). A network can have any number of regions. Regions are independent of each other because the VLAN-to-instance mapping is different in each region.

The FortiSwitch unit supports 15 MST instances in a region. Multiple VLANs can be mapped to each MST instance. Each switch in the region must have the identical mapping of VLANs to instances.

The MST region acts like a single bridge to adjacent MST regions and to non-MST STPs.

IST

Instance 0 is a special instance, called the internal spanning-tree instance (IST). IST is a spanning tree that connects all of the MST switches in a region. All VLANs are assigned to the IST.

IST is the only instance that exchanges bridge protocol data units (BPDUs). The MSTP BPDU contains information for each MSTP instance (captured in an M-record). The M-records are added to the end of a regular RSTP BPDU. This allows MSTP region to inter-operate with an RSTP switch.

CST

The common spanning tree (CST) interconnects the MST regions and all instances of STP or RSTP that are running in the network.

Hop count and message age

MST does not use the BPDU message age within a region. The message-age and maximum-age fields in the BPDU are propagated unchanged within the region.

Within the region, a hop-count mechanism is used to age out the BPDU. The IST root sends out BPDUs with the hop count set to the maximum number of hops. The hop count is decremented each time the BPDU is forwarded. If the hop count reaches zero, the switch discards the BPDU and ages out the information on the receiving port.

STP port roles

STP assigns a port role to each switch port. The role is based on configuration, topology, relative position of the port in the topology, and other considerations. Based on the port role, the port either sends or receives STP BPDUs and forwards or blocks the data traffic. Here is a brief summary of each STP port role:

- **Designated**—One designated port is elected per link (segment). The designated port is the port closest to the root bridge. This port sends BPDUs on the link (segment) and forwards traffic towards the root bridge. In an STP converged network, each designated port is in the STP forwarding state.
- **Root**—The bridge can have only one root port. The root port is the port that leads to the root bridge. In an STP converged network, the root port is in the STP forwarding state.
- **Alternate**—Alternate ports lead to the root bridge but are not root ports. The alternate ports maintain the STP blocking state.
- **Backup**—This is a special case when two or more ports of the same switch are connected together (either directly or through shared media). In this case, one port is designated, and the remaining ports are backup (in the STP blocking state).

STP loop protection

NOTE: This feature is different from loop guard.

When an STP blocking port in a redundant topology starts to incorrectly forward traffic, a layer-2 forwarding loop might form. You can use STP loop protection to help prevent these STP loops, but they still might be formed in unique cases.

A port remains in blocking state only if it continues to receive BPDU messages. If it stops receiving BPDUs (for example, due to unidirectional link failure), the blocking port (alternate or backup port) becomes designated and transitions to a forwarding state. In a redundant topology, this situation may create a loop.

If the loop-protection feature is enabled on a port, that port is forced to remain in blocking state, even if the port stops receiving BPDU messages. It will not transition to forwarding state and does not forward any user traffic.

The loop-protection feature is enabled on a per-port basis. Fortinet recommends that you enable loop protection on all nondesignated ports (all root, alternate, and backup ports).

STP root guard

Root guard protects the interface on which it is enabled from becoming the path to root. When enabled on an interface, superior BPDUs received on that interface are ignored or dropped. Without using root guard, any switch that participates in STP maintains the ability to reroute the path to root. Rerouting might cause your network to transmit large amounts of traffic across suboptimal links or allow a malicious or misconfigured device to pose a security risk by passing core traffic through an insecure device for packet capture or inspection. By enabling root guard on multiple interfaces, you can create a perimeter around your existing paths to root to enforce the specified network topology.

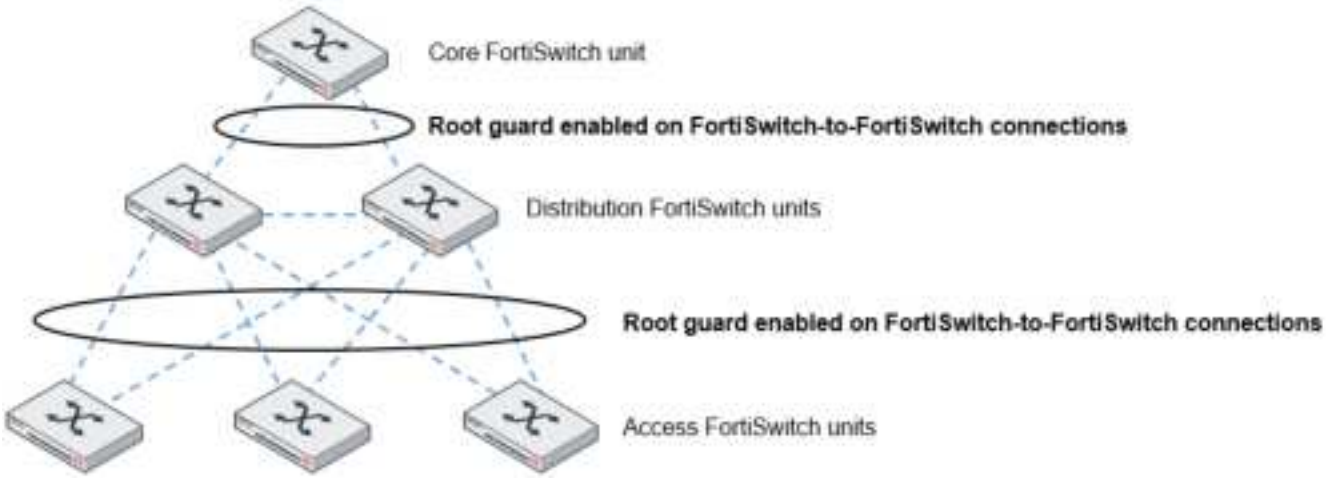

STP BPDU guard

Similar to root guard, BPDU guard protects the designed network topology. When BPDU guard is enabled on STP edge ports, any BPDUs received cause the ports to go down for a specified number of minutes. The BPDUs are not forwarded, and the network edge is enforced.

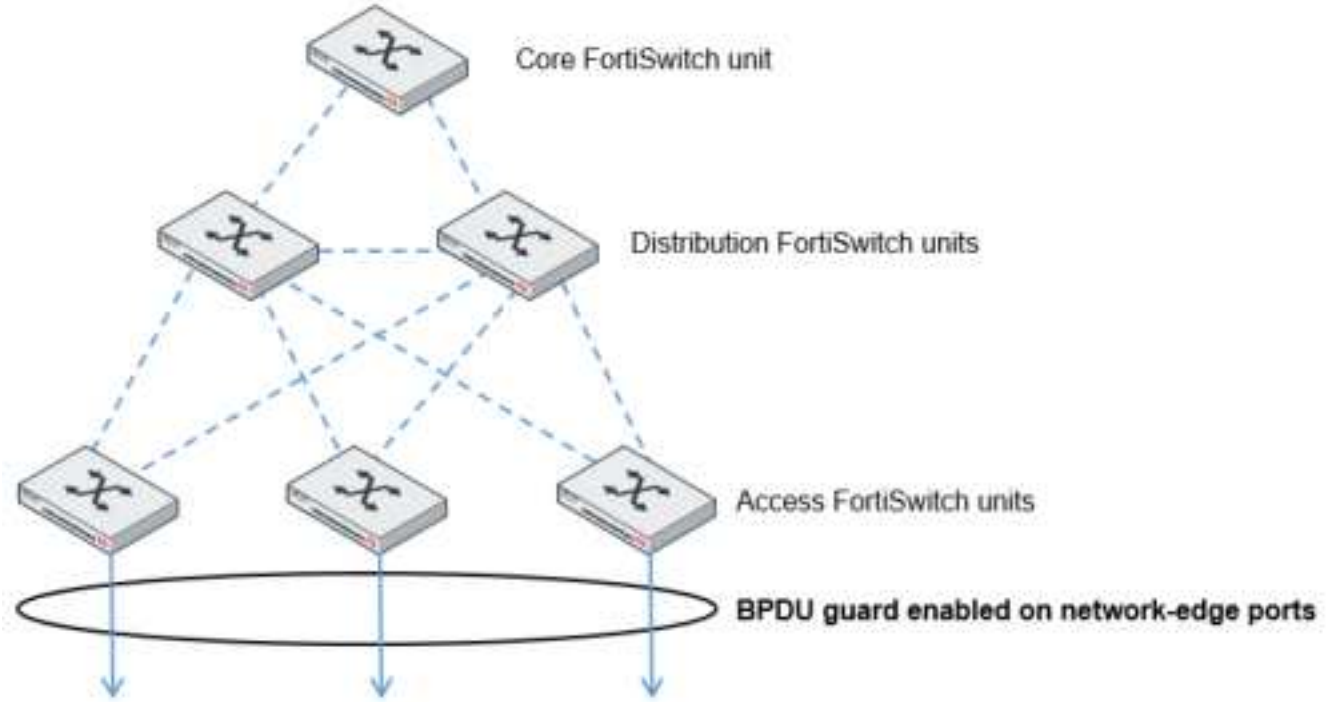

MSTP configuration

MSTP configuration consists of the following steps:

1. Configure STP settings that are common to all MST instances.
2. Configure settings that are specific to each MST instance.

3. Configure loop-protection on all nondesignated ports.

Configuring STP settings

Some STP settings (region name and MST revision number) are common to all MST instances. Also, protocol timers are common to all instances because only the IST sends out BPDUs.

Using the GUI:

1. Go to *Switch* > *STP* > *Settings*.

2. Update the settings as described in the following table.
3. Select *Update* to save the settings.

Settings	Guidelines
Disabled	Disables MSTP for this switch.
Flood BPDU Packets	Select this checkbox if you want the STP packets arriving at any port to pass through the switch without being processed. If you do not select this checkbox, STP packets arriving at any port are blocked. This option is only available when MSTP is disabled.
Enabled	Enables MSTP for this switch.
Name	Region name. All switches in the MST region must have the identical name.
Revision	The MSTP revision number. All switches in the region must have the same revision number. The range of values is 0 to 65535. The default value is 0.
Hello Time (Seconds)	Hello time is how often (in seconds) that the switch sends out a BPDU. The range of values is 1 to 10. The default value is 2.

Settings	Guidelines
Forward Time (Seconds)	Forward time is how long (in seconds) a port will spend in the listening-and-learning state before transitioning to forwarding state. The range of values is 4 to 30. The default value is 15.
Max Age (Seconds)	The maximum age before the switch considers the received BPDU information on a port to be expired. Max-age is used when interworking with switches outside the region. The range of values is 6 to 40. The default value is 20.
Max Hops	Maximum hops is used inside the MST region. Hop count is decremented each time the BPDU is forwarded. If max-hops reaches zero, the switch discards the BPDU and ages out the information on the receiving port. The range of values is 1 to 40. The default value is 20.

Using the CLI:

```
config switch stp settings
   set flood {enable | disable}
   set forward-time <fseconds_int>
   set hello-time <hseconds_int>
   set max-age <age>
   set max-hops <hops_int>
   set mclag-stp-bpdu {both | single}
   set name <name_str>
   set revision <rev_int>
   set status {enable | disable}
end
```

Configuring an MST instance

The STP topology is unique for each MST instance in the region. You can configure a different bridge priority and port parameters for each instance.

Using the GUI:

1. Go to *Switch* > *STP* > *Instances*.

2. Select *Add Instance* to create a new MST instance or select an existing instance and then select *Edit*.
3. Update the instance parameters as described in the following table.
4. Select *Add* or *Update* to save the settings.

Settings	Guidelines
ID	Instance identifier. The range differs for the various FortiSwitch models.
Priority	Priority is a component of bridge ID. The switch with the lowest bridge ID becomes the root switch for this MST instance. Allowed values: 0, 4096, 8192, 12288, 16384, 20480, 24576, 28672, 32768, 36864, 40960, 45056, 49152, 53248, 57344, and 61440. The default value is 32768.
VLAN Range	The VLANs that map to this MST instance. You can specify individual VLAN numbers or a range of numbers. **NOTE:** Do not assign any VLAN to more than one MST instance. Each VLAN number is in the range 1-4094.
Port Configuration	
Name	Port that will participate in this MST instance.
Cost	The switch uses port cost to select designated ports. Port cost is added to the received BPDU root cost in any BPDU sent on this port. A lower value is preferred. The range of values is 1 to 200,000,000. The default value depends on the interface speed: - 10 Gigabit Ethernet: 2,000 - Gigabit Ethernet: 20,000 - Fast Ethernet: 200,000 - Ethernet: 2,000,000

Settings	Guidelines
Priority	The switch uses port priority to choose among ports of the same cost. The port with the lowest priority is put into forwarding state. The valid values are: 0, 32, 64, 96, 128, 160, 192, and 224. The default value is 128.

Using the CLI:

```
config switch stp instance
   edit <instance number>
      set priority <>
      config stp-port
         edit <port name>
            set cost <>
            set priority <>
         next
      set vlan-range <vlan range>
   end
```

Example:

```
config switch stp instance
   edit "1"
      set priority 8192
      config stp-port
         edit "port18"
            set cost 0
            set priority 128
         next
         edit "port19"
            set cost 0
            set priority 128
         next
      end
      set vlan-range 5 7 11-20
   end
```

Configuring an STP edge port

You can use the edge-port setting when a device connected to a FortiSwitch port is not an STP bridge. When this setting is enabled, the FortiSwitch port immediately moves to a forwarding state rather than passing through listening and learning states.

By default, STP (and edge port) is enabled on all ports.

Using the GUI:

1. Go to *Switch* > *Interfaces*.
2. Select a port and then click *Edit*.

3. Under *Edge Port*, select *Enable*.
4. Select *Update* to save the settings.

Using the CLI:

```
config switch interface
   edit <port_name>
      set edge-port <enabled | disabled>
   next
end
```

Configuring STP loop protection

By default, STP loop protection is disabled on all ports.

Using the GUI:

1. Go to *Switch* > *Interfaces*.
2. Select a port and then click *Edit*.
3. Under *Loop Guard*, select *Enable*.
4. Select *Update* to save the settings.

Using the CLI:

```
config switch interface
   edit <port_name>
      set stp-loop-protection <enabled | disabled>
   next
end
```

Configuring STP root guard

Enable root guard on all ports that should not be root bridges. Do not enable root guard on the root port. You must have STP enabled to be able to use root guard.

Using the CLI:

```
config switch interface
   edit <port_name>
      set stp-root-guard <enable | disable>
   next
end
```

For example, to enable root guard on port 20:

```
config switch interface
   edit port20
      set stp-state enabled
      set stp-root-guard enable
   next
end
```

Configuring STP BPDU guard

There are three prerequisites for using BPDU guard:

- You must define the port as an edge port with the `set edge-port enabled` command.
- You must enable STP on the switch interface with the `set stp-state enabled` command.
- You must enable STP on the global level with the `set status enable` command.

You can set how long the port will go down for when a BPDU is received for a maximum of 120 minutes. The default port timeout is 5 minutes. If you set the timeout value to 0, the port will not go down when a BPDU is received, but you will have manually reset the port.

Using the GUI:

1. Go to *Switch* > *Interfaces*.
2. Select a port and then click *Edit*.
3. Under *Edge Port*, select *Enable*.
4. Under *BPDU Guard*, select *Enable*.
5. In the *Automatically, After Minutes* field, enter how many minutes the port will go down for when a BPDU is received.
6. Select *Update* to save the settings.

To check if BPDU guard has been triggered and on which ports, go to *Switch* > *Monitor* > *BPDU Guard*.

Using the CLI:

```
config switch interface
   edit <port_name>
      set stp-bpdu-guard <enabled | disabled>
      set stp-bpdu-guard-timeout <0-120>
   next
end
```

For example, to enable BPDU guard on port 30 with a timeout value of 1 hour:

```
config switch stp settings
   set status enable
end
config switch interface
   edit port30
      set stp-state enabled
      set edge-port enabled
      set stp-bpdu-guard enabled
      set stp-bpdu-guard-timeout 60
   next
end
```

If you set the port timeout to 0, you will need to reset the port after it receives BPDUs and goes down. Use the following command to reset the port:

```
execute bpdu-guard reset <port_name>
```

To check if BPDU guard has been triggered and on which ports, use the following command:

```
diagnose bpdu-guard display status
```

```
Portname              State      Status       Timeout(m)     Count    Last-Event
____________________  _______    _________    ___________    _____   ____________________

port1                 disabled     -              -             -           -
port2                 disabled     -              -             -           -
port3                 disabled     -              -             -           -
port4                 disabled     -              -             -           -
port5                 disabled     -              -             -           -
port6                 disabled     -              -             -           -
port7                 disabled     -              -             -           -
port8                 disabled     -              -             -           -
port9                 disabled     -              -             -           -
port10                disabled     -              -             -           -
port11                disabled     -              -             -           -
port12                disabled     -              -             -           -
port13                disabled     -              -             -           -
port14                disabled     -              -             -           -
port15                disabled     -              -             -           -
port16                disabled     -              -             -           -
port17                disabled     -              -             -           -
port18                disabled     -              -             -           -
port19                disabled     -              -             -           -
port20                disabled     -              -             -           -
port21                disabled     -              -             -           -
port22                disabled     -              -             -           -
port23                disabled     -              -             -           -
port25                disabled     -              -             -           -
port26                disabled     -              -             -           -
port27                disabled     -              -             -           -
port28                disabled     -              -             -           -
port29                disabled     -              -             -           -
port30                enabled      -              60            0           -
__FoRtI1LiNk0__       disabled     -              -             -           -
```

You can also check BPDU guard by going to the *Monitor* > *BPDU Guard* page.

Interactions outside of the MSTP region

A boundary port on an MST switch is a port that receives an STP (version 0) BPDU, an RSTP (version 2) BPDU, or a BPDU from a different MST region.

If the port receives a version 0 BPDU, it will only send version 0 BPDUs on that port. Otherwise, it will send version 3 (MST) BPDUs because the RSTP switch will read this as an RSTP BPDU.

Viewing the MSTP configuration

To view the MSTP configuration details:

- `get switch stp instance`
- `get switch stp settings`

To display information about the MSTP instances in the network:

- `diagnose stp instance list`
- `diagnose stp vlan list`
- `diagnose stp mst-config list`

Support for interoperation with Rapid per-VLAN RSTP (Rapid PVST+ or RPVST+)

Starting in FortiSwitchOS 6.2.2, FortiSwitch units can now interoperate with a network that is running RPVST+. The existing network's configuration can be maintained while adding FortiSwitch units as an extended region.

When an MSTP domain is connected with an RPVST+ domain, FortiSwitch interoperation with the RPVST+ domain works in two ways:

- If the root bridge for the CIST is within an MSTP region, the boundary FortiSwitch unit of the MSTP region duplicates instance 0 information, creates one BPDU for every VLAN, and sends the BPDUs to the RPVST+ domain.

  In this case, follow this rule: If the root bridge for the CIST is within an MSTP region, VLANs other than VLAN 1 defined in the RPVST+ domains must have their bridge priorities worse (numerically greater) than that of the CIST root bridge within MSTP region.

- If the root bridge for the CIST is within an RPVST+ domain, the boundary FortiSwitch unit processes only the VLAN 1 information received from the RPVST+ domain. The other BPDUs (VLANs 2 and above) sent from the connected RPVST+ domain are used only for consistency checks.

  In this case, follow this rule: If the root bridge for the CIST is within the RPVST+ domain, the root bridge priority of VLANs other than VLAN 1 within that domain must be better (numerically less) than that of VLAN 1.

Configuring Rapid PVST or RPVST+ interoperation support

Using the CLI:

Enable the RPVST+ interoperation support on the appropriate switch port or trunk.

```
config switch interface
   edit <interface_name>
      set allowed-vlans <one or more VLANs> // The VLANs must be configured for RSTP.
      set rpvst-port enabled
   next
end
```

For example, to enable RPVST+ interoperation support on port 9:

```
config switch interface
```

```
   edit "port9"
      set allowed-vlans 10,20
      set rpvst-port enabled
   next
end
```

For example, to enable RPVST+ interoperation support on trunk 1:

```
config switch interface
   edit "trunk1"
      set allowed-vlans 10,20
      set rpvst-port enabled
   next
end
```

Note: Refer to the FortiSwitch feature matrix for details about how many VLANs are supported by each FortiSwitch model. The maximum number of VLANs includes native VLANs. You must configure the same VLANs as those used in the RPVST+ domain.

Viewing the configuration

Use one of the following commands to check your configuration and to diagnose any problems.

- `diagnose stp instance list`

  If either rule is violated, the RPVST port is flagged with “IC” in the command output, and the port is in the Discard state.

  If the VLANs used by the RPVST+ domain are not all within the VLAN range configured on the RPVST port, an “MV” flag is displayed in the command output. **NOTE:** Only the ports in instance 0 show this flag.

- `diagnose stp rapid-pvst-port list`

  This command shows the status of one port or all ports. If any of the ports is in the “IC” state, the command output gives the reason: VLAN priority inconsistent, VLAN configuration mismatch, or both.

- `diagnose stp rapid-pvst-port clear`

  This command clears all flags and timers on the RPVST+ port.

Flap guard

A flapping port is a port that changes status rapidly from up to down. A flapping port can create instability in protocols such as STP. If a port is flapping, STP must continually recalculate the role for each port. Flap guard also prevents unwanted access to the physical ports.

The port flap guard detects how many times a port changes status during a specified number of seconds, and the system shuts down the port if necessary. You can manually reset the port and restore it to the active state.

This section covers the following topics:

- Retaining the triggered state on page 253
- Configuring the port flap guard on page 253
- Resetting a port on page 255
- Viewing the port flap guard configuration on page 255

Retaining the triggered state

When the flap guard is triggered, the status for the port is shown as “triggered” in the output of the `diagnose flapguard status` command. By default, rebooting the switch resets the state of the flap guard and removes the “triggered” state. You can change the setting so that the triggered state remains after a switch is rebooting until the port is reset. See Resetting a port on page 255.

Using the GUI:

1. Go to *Switch* > *Flap Guard*.

2. Select *Retain Triggered State Across Reboot*.
3. Select *Update* to save the change.

Using the CLI:

```
config switch global
   set flapguard-retain-trigger enable
end
```

Configuring the port flap guard

The port flap guard is configured and enabled on each port. The default setting is disabled.

The flap rate counts how many times a port changes status during a specified number of seconds. The range is 1 to 30 with a default setting of 5.

The flap duration is the number of seconds during which the flap rate is counted. The range is 5 to 300 seconds with a default setting of 30 seconds.

The flap timeout (CLI only) is the number of minutes before the flap guard is reset. The range is 0 to 120 minutes. The default setting of 0 means that there is no timeout.

NOTE:

- If a triggered port times out while the switch is in a down state, the port is initially in a triggered state until the switch has fully booted up and calculated that the timeout has occurred.
- The following models do not store time across reboot; therefore, any triggered port is initially in a triggered state until the switch has fully booted up—at which point the trigger is cleared:

- FS-1xxE
- FS-2xxD/E
- FS-4xxD
- FS-4xxE

Using the GUI:

1. Go to *Switch* > *Physical Ports*.
2. Select a port.
3. Select *Edit*.
4. Under Flap Guard, select *Enable*.

5. Enter values for *Flap Duration (Seconds)* and *Flap Rate*.
6. Select *Update* to save the changes.

Using the CLI:

```
config switch physical-port
   edit <port_name>
      set flapguard {enabled | disabled}
      set flap-rate <1-30>
      set flap-duration <5-300 seconds>
      set flap-timeout <0-120 minutes>
   end
```

For example:

```
config switch physical-port
   edit port10
      set flapguard enabled
      set flap-rate 15
      set flap-duration 100
      set flap-timeout 30
   end
```

Resetting a port

After the flap guard detects that a port is changing status rapidly and the system shuts down the port, you can reset the port and restore it to service.

Using the GUI:

1. Go to *Switch* > *Physical Ports*.
2. Select the port that was shut down.
3. Select *Reset*.

Using the CLI:

```
execute flapguard reset <port_name>
```

For example:

```
execute flapguard reset port15
```

Viewing the port flap guard configuration

Use the following command to check if the flap guard is enabled on a specific port:

```
show switch physical-port <port_name>
```

For example:

```
show switch physical-port port10
```

Use the following command to display the port flap guard information for all ports:

```
diagnose flapguard status
```

DHCP snooping

The DHCP-snooping feature monitors the DHCP traffic from untrusted sources (for example, typically host ports and unknown DHCP servers) that might initiate traffic attacks or other hostile actions. To prevent this, DHCP snooping filters messages on untrusted ports by performing the following activities:

- Validating DHCP messages received from untrusted sources and filtering out invalid messages. For example, a request to decline an DHCP offer or release a lease is ignored if the request is from a different interface than the one that created the entry.
- Building and maintaining a DHCP-snooping binding database, which contains information about untrusted hosts with leased IP addresses.

Other security features like dynamic ARP inspection (DAI), a security feature that rejects invalid and malicious ARP packets, also use information stored in the DHCP-snooping binding database.

In the FortiSwitch unit, all ports are untrusted by default, and DHCP snooping is disabled on all untrusted ports. You indicate that a source is trusted by configuring the trust state of its connecting interface.

When DHCP snooping is enabled and a DHCP server is detected on an untrusted interface, a log entry is generated, either “A rogue DHCPv6 server has been detected on the interface” or “A rogue DHCP server has been detected on the interface.” **NOTE:** FortiSwitchOS logs only the first DHCPv4 or DHCPv6 server packet that arrives from an untrusted switch interface and is added to the server database.

For additional security, you can specify which DHCP servers that DHCP snooping will include in the allowed server list.

Configuring DHCP snooping

DHCP snooping is enabled per VLAN and, by default, DHCP snooping is disabled.

Configuring DHCP snooping consists of the following steps:

1. Setting the system-wide DHCP-snooping options on page 256
2. Configuring the VLAN settings on page 258
3. Specifying any DHCP-snooping static entries on page 260
4. Configuring the interface settings on page 261

Setting the system-wide DHCP-snooping options

Before you use DHCP snooping, you need to enable the trusted DHCP server list.

NOTE: The maximum number of DHCP servers that can be added to the list is 2,048. This maximum is a global limit and applies across all VLANs.

Using the GUI:

1. Go to *Switch* > *DHCP Snooping*.
2. Enable *Only Allow DHCP from Whitelisted Servers*.

Using the CLI:

```
config system global
   set dhcp-server-access-list {enable | disable}
end
```

For example:

```
config system global
   set dhcp-server-access-list enable
end
```

Including option-82 data

You can include option-82 data in the DHCP request. (DHCP option 82 provides additional security by enabling a controller to act as a DHCP relay agent to prevent DHCP client requests from untrusted sources.) You can select a fixed format for the Circuit ID and Remote ID fields or select which values appear in the Circuit ID and Remote ID fields.

The following is the fixed format for the option-82 Circuit ID field:

Circuit-ID: vlan-mod-port

- vlan—[2 bytes]
- mod—[(1 Byte) -> Snoop - 1 , Relay - 0]
- port—[1 byte]

The following is the fixed format for the option-82 Remote ID field:

Remote-ID: mac [6 bytes]

If you want to select which values appear in the Circuit ID and Remote ID fields:

- For the Circuit ID field, you can include the interface description, host name, interface name, mode, and VLAN.
- For the Remote ID field, you can include the host name, IP address, and MAC address.

To configure the option-82 data:

```
config system global
   set dhcp-option-format {ascii | legacy}
   set dhcp-client-location {description | hostname | intfname | mode | vlan}
   set dhcp-remote-id {hostname | ip | mac}
end
```

Overriding the option-82 settings for a specific VLAN on a port

If you have included option-82 data in the DHCP request, it applies globally. Starting in FortiSwitchOS 7.2.2, you can override the global option-82 setting to specify plain text strings for the Circuit ID field and the Remote ID field for a specific VLAN on a port. If `dhcp-snoop-option82-override` is not configured for the incoming VLAN and switch interface, the settings for the Circuit ID and Remote ID fields are taken from the global option-82 configuration.

NOTE: The values for the Circuit ID and Remote ID field are either both taken from the global option-82 configuration or both taken from the `dhcp-snoop-option82-override` settings. The system cannot take one value at the global level and the other value from the override settings.

Each plain text string can be a maximum of 256 characters long. Together, the combined length of both plain text strings can be a maximum of 256 characters long.

You can only select a VLAN that was configured with the `config switch vlan` command. To include option-82 data in the DHCP request, you must enable `dhcp-snooping` and `dhcp-snooping-option82` under the `config switch vlan` command. The syntax is shown in Configuring the VLAN settings on page 258.

NOTE: You can override the option-82 settings for DHCP snooping but not for DHCP relay.

To override the option-82 global settings:

```
config switch interface
   edit <interface-name>
      config dhcp-snoop-option82-override
         edit <VLAN_ID>
            set remote-id <string>
            set circuit-id <string>
         next
      end
   next
end
```

For example:

```
config switch interface
   edit "port5"
```

```
      config dhcp-snoop-option82-override
         edit 100
            set remote-id "remote-id test"
            set circuit-id "circuit-id test"
         next
      end
   next
end
```

Configuring the VLAN settings

You need to select a VLAN that is configured as a native VLAN or allowed VLAN for a switch interface.

Static IP address are not included as DHCP-snooping entries, so DAI does not analyze them. Starting in FortiSwitchOS 7.2.2, you can specify static entries for DHCP snooping and DAI by manually associating a single IPv4 address with a single MAC address.

Starting in FortiSwitchOS 7.4.3, you can use the DHCP-snooping monitor mode to collect DHCP information from untrusted interfaces in the DHCP client or server database. The monitor mode is available for IPv4 addresses only.

When you change `set dhcp-snooping enable` to `set dhcp-snooping monitor`, the entries in the DHCP client and server database are kept. When you change `set dhcp-snooping monitor` to `set dhcp-snooping enable`, the entries in the DHCP client database are deleted. When you change `set dhcp-snooping monitor` to `set dhcp-snooping disable`, the entries in the DHCP client and server database are deleted.

Starting in FortiSwitchOS 7.4.3, you can monitor ARP packets for a specific VLAN and save the VLAN ID, MAC addresses, and IP addresses in the DHCP-snooping database. By default, the information learned from ARP packets is kept for 24 hours. You can configure how long the information is kept from 5 minutes to 7 days or specify that the information is never removed from the DHCP-snooping database. The static IP addresses can be used in RADIUS accounting.

Using the GUI:

1. Go to *Switch* > *VLAN*.
2. Select *Add VLAN*.
3. Enter the VLAN identifier.
4. Enter a description for the new VLAN.
5. Under DHCP Snooping, select *Enable*.
6. If needed, select *Verify Source MAC*, *Insert Option 82*, and *Dynamic ARP Inspection*.
7. Under the DHCP Server Whitelist, select + to add the name and IP address of an approved DHCP server.
8. In the Members by MAC Address section, select *Add* to add a MAC address.
9. In the Members by IP Address section, select *Add* to add an IPv4 address and netmask.
10. To save your changes, select *Add* at the bottom of the page.

Using the CLI:

```
config switch vlan
   edit <VLAN_ID>
      set dhcp-snooping enable
      set dhcp-snooping-verify-mac {enable | disable>}
      set dhcp-snooping-option82 {enable | disable}
      set dhcp6-snooping enable
      set arp-inspection {enable | disable | monitor}
```

```
      config member-by-mac
         edit <id>
            set mac XX:XX:XX:XX:XX:XX
            set description <128 byte string>
         next
      end
      config member-by-ipv4
         edit <id>
            set address a.b.c.d/e
            set description <128-byte string>
         next
      end
      config dhcp-server-access-list
         edit <string>
            set server-ip <xxx.xxx.xxx.xxx>
            set server-ip6 <xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx>
         next
      end
   next
end
```

NOTE: If you set the `server-ip6` under the `config dhcp-server-access-list` command, the source IPv6 address of the server you are allowing should be the DHCP server's IPv6 link-local address, such as `fe80::213:1ff:fe00:1`.

NOTE: If you enable `dhcp-snooping-verify-mac`, the system will verify that the source MAC address in the DHCP request from an untrusted port matches the client hardware address.

NOTE: If you enable `dhcp-snooping-option82`, the system inserts option-82 data into the DHCP messages for this VLAN.

For example, to configure IPv4 DHCP snooping:

```
config switch vlan
   edit 10
      set dhcp-snooping enable
         config dhcp-server-access-list
            edit "list1"
               set server-ip 100.1.0.2
            next
         end
      next
   end
```

For example, to configure IPv6 DHCP snooping:

```
config switch vlan
   edit 10
      set dhcp6-snooping enable
         config dhcp-server-access-list
            edit "list1"
               set server-ip6 3f2e:6a8b:78a3:0d82:1725:6a2f:0370:6234
            next
         end
      next
   end
```

To use the DHCP-snooping monitor mode:

```
config switch vlan
   edit <VLAN_ID>
      set dhcp-snooping monitor
   next
end
```

To monitor ARP packets:

1. Enable DHCP snooping and enable the monitoring of ARP packets.
   ```
   config switch vlan
      edit <VLAN_ID>
         set dhcp-snooping enable
         set arp-inspection monitor
      next
   end
   ```
2. Specify how long the information learned from ARP packets is kept (from 5 minutes to 7 days) or specify that the information is never removed from the DHCP-snooping database (`set arp-inspection-monitor-timeout 0`). By default, the information learned from ARP packets is kept for 24 hours.
   ```
   config system global
      set arp-inspection-monitor-timeout <5-10080 minutes>
   end
   ```
3. View the information learned from ARP packets in the DHCP-snooping database. The entries have an A flag.
   ```
   get switch dhcp-snooping status
   ```

Specifying any DHCP-snooping static entries

After you enable DHCP snooping for a VLAN, you can configure static entries by binding an IPv4 address with a MAC address for a specific switch interface:

- Specify the MAC address in the form of `xx:xx:xx:xx:xx:xx`.
- To find out which switch interfaces are valid, type `set switch-interface ?`.
- Bind a single MAC address to a single IPv4 address. Multiple IP addresses cannot be bound to the same MAC address. The MAC address cannot be used in more than one static entry.

You can specify a maximum of 64 DHCP static entries for the entire FortiSwitch unit.

Using the CLI:

```
config switch vlan
   edit <vlan-id>
      set dhcp-snooping enable
         config dhcp-snooping-static-client
            set mac-addr <MAC_address>
            set switch-interface <interface_name>
            set ip-addr <IPv4_address>
         next
      end
   next
end
```

For example:

```
config switch vlan
   edit 10
      set dhcp-snooping enable
         config dhcp-snooping-static-client
            set mac-addr 00:01:00:00:00:01
            set switch-interface port20
            set ip-addr 10.1.1.1
         next
      end
   next
end
```

- You cannot use a DHCP trusted switch interface or an 802.1X interface for the static entry's switch interface.
- You cannot configure a DHCP static entry for a private VLAN.
- After you configure a DHCP-snooping static entry for a VLAN, you cannot remove that VLAN from the switch interface.
- After you configure a DHCP-snooping static entry for a switch interface, the switch interface cannot be included as a member of a trunk until the DHCP-snooping static entry is deleted.
- If you configure a DHCP-snooping static entry for a trunk, the trunk cannot be deleted until the DHCP-snooping static entry is deleted.

Configuring the interface settings

After you enable DHCP snooping on a VLAN, all interfaces are in an untrusted state by default, and DHCP snooping is disabled on all untrusted interfaces. You must explicitly configure the trusted interfaces and enable DHCP snooping for each interface.

In addition, you can set a limit for how many IP addresses are in the DHCP snooping binding database for each interface by enabling the `dhcp-snoop-learning-limit-check` and setting the `learning-limit`. By default, `dhcp-snoop-learning-limit-check` is disabled, and the number of entries for an untrusted ports is 5. You can set the number of entries to 0. The maximum number of entries depends on which FortiSwitch unit you are using. For example:

```
S548DN4K16000313 # show switch vlan 1
config switch vlan
   edit 1
      set learning-limit 100
      set dhcp-snooping enable
   next
end
```

NOTE: If the FortiSwitch unit has already learned more IP addresses than the `dhcp-snoop-learning-limit` before the limit is set, the configuration is rejected because the FortiSwitch unit cannot select which IP addresses should be kept. If the FortiSwitch unit has learned fewer IP address or the same number of IP addresses as the `dhcp-snoop-learning-limit` before the limit is set, the configuration is accepted.

NOTE: The per-VLAN learning limit is not supported on dual-chip platforms (448 series).

Using the GUI:

1. Go to *Switch* > *Interfaces*.
2. Select an interface.
3. Click *Edit*.
4. Select *Trusted* or *Untrusted* for DHCP snooping.
5. If you want to accept DHCP messages with option-82 data from an untrusted interface, select the *Option 82 Trust* check box.
6. Click *Update*.

Using the CLI:

```
config switch {interface | trunk}
   edit <interface-name>
      set native-vlan <VLAN-ID>
      set dhcp-snooping {trusted | untrusted}
      set dhcp-snoop-learning-limit-check {enable | disable}
      set learning-limit <integer>
      set dhcp-snoop-option82-trust {enable | disable}
   next
end
```

For example:

```
config switch interface
   edit "port5"
      set native-vlan 10
      set dhcp-snooping untrusted
      set dhcp-snoop-learning-limit-check enable
      set learning-limit 7
      set dhcp-snoop-option82-trust enable
      set snmp-index 5
   next
end
```

Set `dhcp-snooping` to reflect the trust state of the interface. Where DHCP servers are located, you must configure interfaces as trusted.

If you enable `dhcp-snoop-option82-trust`, the system accepts DHCP messages with option-82 data from an untrusted interface.

Checking the DHCP-snooping configuration

Starting in FortiSwitchOS 7.4.3, there are three flags added to the CLI output of the DHCP client database:

- `D`—Indicates that the entry was added when the VLAN was in DHCP-snooping monitor mode.
- `S`—Indicates that the entry is a DHCP-snooping static entry
- `A`—Indicates that the entry was added by DAI.

To view the detailed status of IPv4 and IPv6 DHCP-snooping VLANs and ports:

```
get switch dhcp-snooping database-summary
```

An entry in the DHCP snooping binding database that contains an * after the IP address indicates a temporary or incomplete entry. For example:

```
08:00:27:13:16:51 2000 100.0.0.159* 10 4 port4
```

The DHCP server has not acknowledged this entry yet. If the DHCP server does not acknowledge the entry within 10 seconds, the entry is removed from the database. If the DHCP server does acknowledge the entry within 10 seconds, the entry will be considered “complete” (that is, no * after the IP address), and a proper expiration time is assigned to it.

To view the details of the IPv4 and IPv6 DHCP-snooping client and server databases, including DHCP-snooping static entries:

```
get switch dhcp-snooping status
```

To view the DHCP-snooping static entries:

```
get switch dhcp-snooping static-clients
```

To view the details of the IPv4 DHCP-snooping client database:

- Enter the following CLI command: `get switch dhcp-snooping client-db-details`
- Go to *Switch* > *Monitor* > *DHCP Snooping* > *Clients*.

To view the details of the IPv6 DHCP-snooping client database:

- Enter the following CLI command: `get switch dhcp-snooping client6-db-details`
- Go to *Switch* > *Monitor* > *DHCP Snooping* > *Clients*.

To view the details of the IPv4 DHCP-snooping server database:

- Enter the following CLI command: `get switch dhcp-snooping server-db-details`
- Go to *Switch* > *Monitor* > *DHCP Snooping* > *Servers*.

To view the details of the IPv6 DHCP-snooping server database:

- Enter the following CLI command: `get switch dhcp-snooping server6-db-details`
- Go to *Switch* > *Monitor* > *DHCP Snooping* > *Servers*.

If the `dhcp-server-access-list` is enabled globally and the server is configured for the `dhcp-server-access-list`, the `svr-list` column displays `allowed` for that server. If the `dhcp-server-access-list` is enabled globally and the server is not configured in the `dhcp-server-access-list`, the `svr-list` column displays `blocked` for that server.

Removing an entry from the DHCP-snooping binding database

You can remove an IP address from the DHCP-snooping binding database by specifying the associated VLAN ID and MAC address:

```
execute dhcp-snooping expire-client <1-4095> <xx:xx:xx:xx:xx:xx>
```

For example:

```
execute dhcp-snooping expire-client 100 01:23:45:67:89:01
```

IP source guard

IP source guard protects a network from IPv4 spoofing by only allowing traffic on a port from specific IPv4 addresses. Traffic from other IPv4 addresses is discarded. The discarded addresses are not logged.

IP source guard allows traffic from the following sources:

- Static entries—IP addresses that have been manually associated with MAC addresses.
- Dynamic entries—IP addresses that have been learned through DHCP snooping.

By default, IP source guard is disabled. You must enable it on each port that you want protected. If you enable IP source guard and then disable it, all static and dynamic entries are removed for that interface.

There is a maximum of 2,048 IP source guard entries. When there is a conflict between static entries and dynamic entries, static entries take precedence over dynamic entries.

To see which models support this feature, refer to the FortiSwitch feature matrix.

NOTE: IP source guard does not work with VLAN translation.

Configuring IP source guard consists of the following steps:

1. Enabling IP source guard on page 264
2. Configuring IP source-guard static entries on page 265
3. Checking the IP source-guard entries on page 265
4. (Optional) Checking the IP source-guard violation log on page 266

Enabling IP source guard

You must enable IP source guard before you can configure it.

To enable IP source guard:

```
config switch interface
   edit <port_name>
      set ip-source-guard enable
   end
```

For example:

```
config switch interface
   edit port6
      set ip-source-guard enable
   end
```

To reset IP source-guard violations for a specific switch interface:

```
execute source-guard-violation reset interface <interface_name>
```

Configuring IP source-guard static entries

After you enable IP source guard, you can configure static entries by binding IPv4 addresses with MAC addresses. For IP source-guard dynamic entries, you need to configure DHCP snooping. See DHCP snooping on page 255.

Using the GUI:

1. Go to *Switch* > *IP Source Guard*.
2. Select *Configure* for the interface that you want to add IP source guard to.
3. In the Description field, add a description of the configuration.
4. Select +.
5. Required. In the Name field, enter a name for the binding entry.
6. Required. In the IP address field, enter the IPv4 address to bind to the MAC address. Masks are not supported.
7. Required. In the MAC address field, enter the MAC address to bind to the IPv4 address.
8. Select *Configure* to save your configuration.

Using the CLI:

```
config switch ip-source-guard
   edit <port_name>
      config binding-entry
         edit <id>
            set ip <xxx.xxx.xxx.xxx>
            set mac <XX:XX:XX:XX:XX:XX>
         next
      end
   next
end
```

For example:

```
config switch ip-source-guard
   edit port4
      config binding-entry
         edit 1
            set ip 172.168.20
            set mac 00:21:cc:d2:76:72
         next
      end
   next
end
```

Checking the IP source-guard entries

After you configure IP source guard, you can check the database entries. Static entries are manually added by the `config switch ip-source-guard` command. Dynamic entries are added by DHCP snooping.

Using the GUI:

Go to *Switch* > *Monitor* > *IP Source Guard*.

Using the CLI:

```
diagnose switch ip-source-guard hardware entry list
```

Checking the IP source-guard violation log

If you want to see events that violate the IP source-guard settings, enable the IP source-guard violation log.

The IP source-guard violation log contains a maximum of 128 entries with a maximum of 5 entries per port, even if more violations have occurred. The maximum values cannot be changed.

To enable the IP source-guard violation log:

```
config switch global
   set log-source-guard-violations enable
   set source-guard-violation-timer <1-1500 minutes>
end
```

To display all IP source-guard violations:

```
get switch ip-source-guard-violations all
```

To display IP source-guard violations for a specific switch interface:

```
get switch ip-source-guard-violations interface <interface_name>
```

To reset all IP source-guard violations:

```
execute source-guard-violation reset all
```

To reset IP source-guard violations for a specific switch interface:

```
execute source-guard-violation reset interface <interface_name>
```

Dynamic ARP inspection

Dynamic ARP Inspection (DAI) prevents man-in-the-middle attacks and IP address spoofing by checking that packets from untrusted ports have valid IP-MAC-address binding. To use DAI, you must first enable the DHCP snooping feature and then enable DAI for each VLAN. See DHCP snooping on page 255.

Configuring DAI

Configuring DAI consists of the following steps:

1. Enable DAI for each VLAN. By default, it is disabled.
2. Enable DAI for the switch interface. By default, all interfaces are in an untrusted state. You must explicitly configure the trusted interfaces.

Enable DAI for each VLAN

Using the GUI:

1. Go to *Switch* > *VLAN*.
2. Select *Add VLAN*.
3. Enter the VLAN identifier.
4. Enter a description for the new VLAN.
5. Under DHCP Snooping, select *Enable*.
6. Select *Dynamic ARP Inspection*.
7. To save your changes, select *Add* at the bottom of the page.

Using the CLI:

```
config switch vlan
   edit <vlan-id>
      set arp-inspection {enable | disable}
   next
end
```

Enable DAI for the switch interface

Using the GUI:

1. Go to *Switch* > *Interfaces*.
2. Select an interface and click *Edit*.
3. Enter the VLAN identifier.
4. Enter a description for the new VLAN.
5. Select *Untrusted* or *Trusted* for *DHCP Snooping*.
6. Click *Update*.

Using the CLI:

```
config switch interface
   edit <interface-name>
      set arp-inspection-trust <untrusted | trusted>
   next
end
```

Checking ARP packets

Use the following command to see how many ARP packets have been dropped or forwarded:

```
#diagnose switch arp-inspection stats

vlan 100           arp-request               arp-reply
-------------------------------------------------------------------------
received                 0                         0
forwarded                0                         0
dropped                  0                         0
```

IPv6 router advertisement guard

IPv6-enabled routers send router advertisement (RA) messages to neighboring hosts in the local network. To prevent the spoofing of the RA messages, RA guard inspects RA messages to see if they meet the criteria contained in an RA-guard policy. If the RA messages match the criteria in the policy, they are forwarded. If the RA messages do not match the criteria in the policy, they are dropped.

The IPv6 RA-guard policy checks for the following criteria in each RA message:

- Whether it has been flagged with the M (managed address configuration) flag or O (other configuration) flag
- Whether the hop number is equal or more than the minimum hop limit
- Whether the hop number is equal or less than the maximum hop limit
- Whether the default router preference is set to high, medium, or low
- Whether the source IPv6 address matches an allowed address in an IPv6 access list (created with the `config router access-list6` command)
- Whether the IPv6 address prefix matches an allowed prefix in an IPv6 prefix list (created with the `config router prefix-list6` command)
- Whether the device is a host or a router. If the device is a host, all RA messages are dropped. If the device is a router, the other criteria in the policy are checked.

To see which models support this feature, refer to the FortiSwitch feature matrix.

Configuring IPv6 RA guard consists of the following steps:

1. (Optional) Creating an IPv6 access list on page 268
2. (Optional) Creating an IPv6 prefix list on page 269
3. Creating an IPv6 RA-guard policy on page 269
4. Applying the IPv6 RA-guard policy on page 270
5. (Optional) Viewing available IPv6 RA-guard policies on page 271

Creating an IPv6 access list

Create an IPv6 access list if you want to specify which source IPv6 address are allowed in RA messages. When no rule in the IPv6 access list is matched, the RA messages are dropped.

To create an IPv6 access list:

```
config router access-list6
   edit <name_of_IPv6_access_list>
      set comments <string>
      config rule
         edit <rule_ID>
            set action {deny | permit}
            set prefix6 {<xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx> | any}
            set exact-match {enable | disable}
         next
      end
   end
```

For example:

```
config router access-list6
   edit accesslist1
```

```
      set comments "IPv6 access list"
      config rule
         edit 1
            set action permit
            set prefix6 fe80::a5b:eff:fef1:95e5
            set exact-match disable
         next
      end
   end
```

Creating an IPv6 prefix list

Create an IPv6 prefix list if you want to specify which IPv6 prefixes in the RA option type 3 are allowed in RA messages. When no rule in the IPv6 prefix list is matched, the RA messages are dropped.

To create an IPv6 prefix list:

```
config router prefix-list6
   edit <name_of_IPv6_prefix_list>
      set comments <string>
      config rule
         edit <rule_ID>
            set action {deny | permit}
            set prefix6 {<IPv6_prefix> | any}
            set ge <0-128>
            set le <0-128>
         next
      end
   end
```

For example:

```
config router prefix-list6
   edit prefixlist1
      set comments "IPv6 prefix list"
      config rule
         edit 1
            set action permit
            set prefix6 any
            set ge 50
            set le 50
         next
      end
   end
```

Creating an IPv6 RA-guard policy

In the IPv6 RA-guard policy, you specify the criteria that RA messages must match before the RA messages are forwarded.

To create an IPv6 RA-guard policy:

```
config switch raguard-policy
   edit <RA-guard policy name>
```

```
      set device-role {host | router}
      set managed-flag {Off | On}
      set other-flag {Off | On}
      set max-hop-limit <0-255>
      set min-hop-limit <0-255>
      set max-router-preference {high | medium | low}
      set match-src-addr <name_of_IPv6_access_list>
      set match-prefix <name_of_IPv6_prefix_list>
   next
end
```

For example:

```
config switch raguard-policy
   edit RApolicy1
      set device-role router
      set managed-flag On
      set other-flag On
      set max-hop-limit 100
      set min-hop-limit 5
      set max-router-preference medium
      set match-src-addr accesslist1
      set match-prefix prefixlist1
   next
end
```

Applying the IPv6 RA-guard policy

After you create an IPv6 RA-guard policy, you need to apply it to the appropriate switch ports or trunks and VLANs. You can create and apply different policies to different VLANs.

To apply the IPv6 RA-guard policy:

```
config switch interface
   edit <interface_name>
   config raguard
      edit <ID>
         set raguard-policy <name_of_RA_guard_policy>
         set vlan-list <list_of_VLANs>
      next
   end
end
```

For example:

```
config switch interface
   edit <interface_name>
   config raguard
      edit 1
         set raguard-policy RApolicy1
         set vlan-list 1
      next
      edit 2
         set raguard-policy RApolicy2
         set vlan-list 2-5
```

```
      next
    end
end
```

Viewing available IPv6 RA-guard policies

Use the following command to list the available IPv6 RA-guard policies:

```
get switch raguard-policy
```

For example:

```
S524DF4K15000024 # get switch raguard-policy
== [ RApolicy1 ]
name: RApolicy1
```

ACL

You can use access control lists (ACLs) to configure policies for three different stages in the pipeline:

- Ingress stage for incoming traffic
- Prelookup stage for processing traffic
- Egress stage for outgoing traffic

You can use an ACL policy at the ingress or egress stage to mirror packets to another port, interface, or trunk. You can mirror traffic from VLANs as well.

This section covers the following topics:

- ACL policy attributes on page 272
- Configuring an ACL policy on page 273
- Configuration examples on page 281
- Selective packet sampling on page 283
- Creating a schedule on page 284

NOTES

- Before FortiSwitchOS 6.0.0, you used the `config switch acl policy` command to configure ACL policies only for the ingress stage. In FortiSwitchOS 6.0.0 and later, the `config switch acl` command has changed to specify which stage is being configured.
- Starting in FortiSwitchOS 6.2.0, you can create groups for multiple ingress ACLs.
- Starting in FortiSwitchOS 7.2.0, you can count ingress and egress packets by color:
  - Ingress packets are marked green if the traffic rate is within the guaranteed information rate. Ingress packets are marked yellow if they exceed the committed burst size but do not exceed the excess burst size. All other ingress packets are marked red.
  - Egress packets are marked green if the traffic rate is within the guaranteed information rate. All other egress packets are marked yellow.

  The colors are displayed in the *Switch* > *Monitor* > *ACL Counters* page and in the output of the `get switch acl counters {all | egress | ingress}` commands. To use this feature, you must configure the ACL policer

first.

- The FS-1024D and FS-524D-FPOE models do not support all action options on the ingress policy.
- For the FS-448E, FS-448E-FPOE, and FS-448E-POE models, the prelookup stage is limited to 504 ACL entries.
- Starting in FortiOS 7.4.0 with FortiSwitchOS 7.4.0, ACL configuration is supported in FortiLink mode.
- There are some limitations for ACL configuration on the FS-108E, FS-108E-POE, FS-108E-FPOE, FS-124E, FS-124E-POE, FS-124E-FPOE, FS-148E, FS-148E-POE, FS-108F, FS-108F-POE, FS-108F-FPOE, FS-124F, FS-124F-POE, FS-124F-FPOE, FS-148F, FS-148F-POE, and FS-148F-FPOE models:
  - The layer-4 port range is limited and might not be available in FortiSwitchOS 6.4.0.
  - Use the `get switch acl usage` command to find out how many counters are available on your switch model.
  - If a classifier was created with only layer-2 fields, layer-3 fields cannot be added later. If a classifier was created with only layer-3 fields, layer-2 fields cannot be added later.
  - You cannot use both drop and redirect actions in the same ACL policy.
  - Only the ingress policy can be configured. There are seven options (`dst-ip-prefix`, `dst-mac`, `ether-type`, `service`, `src-ip-prefix`, `src-mac`, and `vlan-id`) for configuring the classifier and five options (`cos-queue`, `count`, `drop`, `outer-vlan-tag`, and `redirect`) for configuring the action.
- The `set redirect` command works differently for the following switch models:
  - For the FS-108E, FS-108E-POE, FS-108E-FPOE, FS-124E, FS-124E-POE, FS-124E-FPOE, FS-148E, and FS-148E-POE models, the egress VLAN membership is *not* necessary.
  - For the FS-148F, FS-148F-POE, FS-148F-FPOE, FS-124F, FS-124F-POE, and FS-124F-FPOE models, the egress VLAN membership is necessary.
  - For the FS-6xxF models, the `set redirect` command only works if the ports are in the same VLAN.

ACL policy attributes

Key attributes of a policy include:

- *Interface*. The interface(s) on which traffic arrives at the switch. The interface can be a port, a trunk, or all interfaces. The policy applies to ingress traffic only (not egress traffic).
- *Classifier*. The classifier identifies the packets that the policy will act on. Each packet can be classified based on one or more criteria. Criteria include source and destination MAC address, VLAN id, source and destination IP address, or service (layer 4 protocol id and port number).
- *Marking* involves setting bits in the packet header to indicate the priority of this packet.
- *Actions*. If a packet matches the classifier criteria for a given ACL, the following types of action may be applied to the packet:
  - allow or block the packet, redirect the packet, mirror the packet
  - police the traffic
  - mirror the packet to another port, interface, or trunk
  - mirror the traffic (you can mirror the traffic from VLANs as well)
  - CoS queue assignment
  - outer VLAN tag assignment
  - egress mask to filter packets
  - specify a schedule when the ACL policy will be applied
  - make the ACL policy active or inactive

The switch uses specialized TCAM memory to perform ACL matching.

NOTE: Each model of the FortiSwitch unit provides different ACL-related capabilities. When you configure the ACL policy, the system will reject the request if the hardware cannot support it.

Configuring an ACL policy

You can configure ACL policies for each stage: ingress, egress, and prelookup.

The order of the classifiers provided during group creation (or during an ACL update in a group when new classifiers are added) matter. Hardware resources are allocated as best fit at the time of creation, which can cause some fragmentation and segmentation of hardware resources because not all classifiers are available at all times. Because the availability of classifiers is order dependent, some allocations succeed or fail at different times. Rebooting the switch or running the `execute acl key-compaction <acl-stage><group-id>` command can help reduce the classifier resource fragmentation.

This section covers the following topics:

- Creating an ACL ingress policy on page 273
- Creating an ACL egress policy on page 275
- Creating an ACL prelookup policy on page 276
- Creating or customizing a service on page 277
- Creating a policer on page 278
- Viewing counters on page 279
- Clearing counters on page 280
- Clearing unused classifiers on page 280

Creating an ACL ingress policy

Starting in FortiSwitchOS 7.2.3, the CLI supports IPv6 address in ACLs for ingress policies.

NOTE: You must set `group` to 3 or higher for the `set dst-ip6-prefix` and `set src-ip6-prefix` commands to be available. If you are going to use a dynamic ACL, set `group` to 4 or higher.

Using the GUI:

1. Go to *Switch* > *ACL* > *Ingress*.
2. Select *Add Ingress Policy*.
3. Required. In the ID field, enter a unique number to identify this policy.
4. By default, *Active* is selected. If you do not want this policy to be active, clear the *Active* checkbox.
5. Required. Select which interfaces the policy applies to or select the *All Interface* checkbox.
6. Select a schedule for when the ACL policy is enforced. To create a schedule, see Example 4 on page 283.
7. In the *Description* field, enter a description or other information about the policy. The description is limited to 63 characters.
8. Configure the classifier.
   a. Enter the VLAN identifier to be matched.
   b. Enter the 802.1Q cost of service (CoS) value to match.
   c. Enter the DSCP value to match.

d. Enter the Ethernet type to be matched.
e. Select the service type to be matched.
f. Enter the source MAC address to be matched.
g. Enter the destination MAC address to be matched.
h. Enter the source IP address and subnet mask to be matched.
i. Enter the destination IP address and subnet mask to be matched.

9. Configure the action.
   a. Select the *Count* checkbox if you want to track the number of matching packets. Select black to count all ingress packets; select green to count ingress packets if the traffic rate is within the guaranteed information rate; select yellow to count ingress packets if they exceed the committed burst size but do not exceed the excess burst size; select red to count all other ingress packets.
   b. Select the *Drop* checkbox if you want to drop matching packets.
   c. Select the *Redirect Broadcast CPU* checkbox if you want to redirect broadcast traffic to all ports including the CPU.
   d. Select the *Redirect Broadcast No CPU* checkbox if you want to redirect broadcast traffic to all ports excluding the CPU.
   e. In the *CPU COS Queue* field, enter the CPU CoS queue number. This CoS queue is only used if the packets reach the CPU.
   f. In the *COS Queue* field, enter the CoS queue number.
   g. In the *Remark COS* field, enter the CoS marking value.
   h. In the *Outer VLAN Tag* field, enter the outer VLAN tag.
   i. In the *Remark DSCP* field, enter the DSCP marking value.
   j. Select *Egress Mask* to configure which physical ports are included in the egress mask or select *Redirect Physical Port* to redirect packets to the selected physical ports.
   k. Select the physical ports to include in the egress mask or to redirect packets to.
   l. Select which policer to use from the Policer drop-down list. To create a policer, see Creating a policer on page 278.
   m. Select which redirect interface to use from the Redirect Interface drop-down list.
   n. Select the name of the mirror to use collect packets to analyze.
10. Select *OK* to save the ingress policy.

Using the CLI:

```
config switch acl ingress
edit <policy_ID>
   set description <string>
   set group <group_ID>
   set ingress-interface <port_name>
   set ingress-interface-all {enable | disable}
   set schedule <schedule_name>
   set status {active | inactive}
   config classifier
      set cos <0-7>
      set dscp <0-63.>
      set dst-ip-prefix <IPv4_address> <mask>
      set dst-ip6-prefix <IPv6_address> <prefix>
      set dst-mac <MAC_address>
      set ether-type <0-65535>
      set service <service-id>
      set src-ip-prefix <IPv4_address> <mask>
      set src-ip6-prefix <IPv6_address> <prefix>
```

```
      set src-mac <MAC_address>
      set vlan-id <1-4094>
   end
   config action
      set cos-queue <0-7>
      set count {enable | disable}
      set count-type {all | green | yellow | red}
      set cpu-cos-queue <integer>
      set drop {enable | disable}
      set egress-mask {<physical_port_name> | internal}
      set mirror <mirror_session>
      set outer-vlan-tag <integer>
      set policer <policer>
      set redirect <interface_name>
      set redirect-bcast-cpu {enable | disable}
      set redirect-bcast-no-cpu {enable | disable}
      set redirect-physical-port <list of physical ports to redirect>
      set remark-cos <0-7>
      set remark-dscp <0-63>
   end
end
```

Creating an ACL egress policy

Using the GUI:

1. Go to *Switch* > *ACL* > *Egress*.
2. Select *Add Egress Policy*.
3. Required. In the ID field, enter a unique number to identify this policy.
4. By default, *Active* is selected. If you do not want this policy to be active, clear the *Active* checkbox.
5. Select which interface the policy applies to.
6. Select a schedule for when the ACL policy is enforced. To create a schedule, see Example 4 on page 283.
7. In the Description field, enter a description or other information about the policy. The description is limited to 63 characters.
8. Configure the classifier.
   a. Enter the VLAN identifier to be matched.
   b. Enter the 802.1Q cost of service (CoS) value to match.
   c. Enter the DSCP value to match.
   d. Enter the Ethernet type to be matched.
   e. Select the service type to be matched.
   f. Enter the source MAC address to be matched.
   g. Enter the destination MAC address to be matched.
   h. Enter the source IP address and subnet mask to be matched.
   i. Enter the destination IP address and subnet mask to be matched.
9. Configure the action.
   a. Select the *Count* checkbox if you want to track the number of matching packets. Select black to count all egress packets; select green to count egress packets if the traffic rate is within the guaranteed information rate; select yellow to count all other egress packets.
   b. Select the *Drop* checkbox if you want to drop matching packets.
   c. In the Outer VLAN Tag field, enter the outer VLAN tag.

d. In the Remark DSCP field, enter the DSCP marking value.
e. Select which policer to use from the Policer drop-down list. To create a policer, see Creating a policer on page 278.
f. Select which redirect interface to use from the Redirect Interface drop-down list.
g. Select the name of the mirror to use collect packets to analyze.

10. Select *OK* to save the egress policy.

Using the CLI:

```
config switch acl egress
edit <policy_ID>
   set description <string>
   set interface <port_name>
   set schedule <schedule_name>
   set status {active | inactive}
   config classifier
      set src-mac <MAC_address>
      set dst-mac <MAC_address>
      set ether-type <integer>
      set src-ip-prefix <IP_address> <mask>
      set dst-ip-prefix <IP_address> <mask>
      set service <service_ID>
      set vlan-id <VLAN_ID>
      set cos <802.1Q CoS value to match>
      set dscp <DSCP value to match>
   end
   config action
      set count {enable | disable}
      set count-type {all | green | yellow}
      set drop {enable | disable}
      set mirror <mirror_session>
      set outer-vlan-tag <integer>
      set policer <policer>
      set redirect <interface_name>
      set remark-dscp <0-63>
   end
end
```

Creating an ACL prelookup policy

Using the GUI:

1. Go to *Switch* > *ACL* > *Prelookup*.
2. Select *Add Prelookup Policy*.
3. Required. In the ID field, enter a unique number to identify this policy.
4. By default, *Active* is selected. If you do not want this policy to be active, clear the *Active* checkbox.
5. Select which ingress interface the policy applies to.
6. Select a schedule for when the ACL policy is enforced. To create a schedule, see Example 4 on page 283.
7. In the Description field, enter a description or other information about the policy. The description is limited to 63 characters.

8. Configure the classifier.
   a. Enter the VLAN identifier to be matched.
   b. Enter the 802.1Q cost of service (CoS) value to match.
   c. Enter the DSCP value to match.
   d. Enter the Ethernet type to be matched.
   e. Select the service type to be matched.
   f. Enter the source MAC address to be matched.
   g. Enter the destination MAC address to be matched.
   h. Enter the source IP address and subnet mask to be matched.
   i. Enter the destination IP address and subnet mask to be matched.
9. Configure the action.
   a. Select the *Count* checkbox if you want to track the number of matching packets.
   b. Select the *Drop*checkbox if you want to drop matching packets.
   c. In the Outer VLAN Tag field, enter the outer VLAN tag.
   d. In the COS Queue field, enter the CoS queue number.
   e. In the Remark COS field, enter the CoS marking value.
10. Select *OK* to save the prelookup policy.

Using the CLI:

```
config switch acl prelookup
   edit <policy_ID>
      set description <string>
      set interface <port_name>
      set interface-all {enable | disable}
      set schedule <schedule_name>
      set status {active | inactive}
      config classifier
         set src-mac <MAC_address>
         set dst-mac <MAC_address>
         set ether-type <integer>
         set src-ip-prefix <IP_address> <mask>
         set dst-ip-prefix <IP_address> <mask>
         set service <service_ID>
         set vlan-id <VLAN_ID>
         set cos <802.1Q CoS value to match>
         set dscp <DSCP value to match>
      end
      config action
         set cos-queue <0-7>
         set count {enable | disable}
         set drop {enable | disable}
         set outer-vlan-tag <integer>
         set remark-cos <0-7>
      end
   end
```

Creating or customizing a service

Optionally, you can create or customize a service. When you create an ACL policy (ingress, egress, or prelookup), you select the service to use with the `set service <service_ID>` command under `config classifier`.

The FortiSwitch unit provides a set of pre-configured services that you can use. Use the following command to list the services:

```
show switch acl service custom
```

Using the GUI to create a service:

1. Go to *Switch* > *ACL* > *Service*.
2. Click *Add Service*.
3. Required. In the *Name* field, enter the name of the service.
4. If you want to change the icon color for the service in the *Service* page, click *Change* and then click the new color.
5. Optional. Enter a description of the service.
6. If you do not want the custom service to use TCP, select *ICMP*, *IP*, *UDP*, or *SCTP*.
7. If you selected TCP, UDP, or SCTP, enter the destination ports and source ports.
   You can enter a single port or a range of ports in each field.
8. If you selected ICMP or IP, enter the protocol number.
9. If you selected ICMP, enter the ICMP code.
10. Click *Add*.

Using the GUI to customize a service:

1. Go to *Switch* > *ACL* > *Service*.
2. Click *Edit* for the service that you want to customize.
3. Make any changes.
4. Click *Update*.

Using the CLI to create or customize a service:

```
config switch acl service custom
   edit <service_name>
      set comment <string>
      set color <0-32>
      set protocol {ICMP | IP | TCP/UDP/SCTP}
      set sctp-portrange <dstportlow_int>[-<dstporthigh_int>: <srcportlow_int>-<srcporthigh_
          int>]
      set tcp-portrange <dstportlow_int>[-<dstporthigh_int>:
      <srcportlow_int>-<srcporthigh_int>]
      set udp-portrange <dstportlow_int>[-<dstporthigh_int>:<srcportlow_int>-<srcporthigh_
          int>]
   end
```

Creating a policer

The ACL policer uses a single-rate three-color marker (RFC 2697) to mark packets as green, yellow, or red, based on the guaranteed bandwidth, guaranteed burst, and maximum burst settings. Traffic below the guaranteed bandwidth is allowed. Traffic above the guaranteed burst or maximum burst is dropped.

Optionally, you can create a policer if you are defining ACLs to police different types of traffic. When you create an ACL ingress or egress policy, you select the policer to use with it.

Using the GUI:

1. Go to *Switch* > *ACL* > *Policer*.
2. Select *Add Policer*.
3. Required. In the ID field, enter a unique number to identify this policer.
4. In the Type drop-down list, select whether the policer is for the egress policy or the ingress policy.
5. In the Guaranteed Bandwidth field, enter the amount of bandwidth guaranteed (in Kbits/second) to be available for traffic controlled by the policy.
6. In the Guaranteed Burst field, enter the guaranteed burst size in bytes.
7. In the Maximum Burst field, enter the maximum burst size in bytes
8. In the Description field, enter a description of the policer.
9. Select *OK* to save the policer.

Using the CLI:

```
config switch acl policer
   edit <1-2048>
      set description <string>
      set guaranteed-bandwidth <bandwidth_value>
      set guaranteed-burst <in_bytes>
      set maximum-burst <in_bytes>
      set type {egress | ingress}
   end
```

Each policy is assigned a unique policy ID that is automatically assigned. To view it, use the `get switch acl {egress | ingress | prelookup}` command.

Viewing counters

On the 4xxE, 1xxE, 1xxF, and FS-110G-FPOE platforms, the ACL byte counters are not available (they will always show as 0 on the CLI). The packet counters are available.

You can use the GUI and CLI to view the counters associated with the ingress, egress, and prelookup policies.

Using the GUI:

Go to *Switch* > *Monitor* > *ACL Counters*.

Using the CLI:

```
get switch acl counters {all | egress | ingress | prelookup}
```

For example:

```
FS3E32T419000041 # get switch acl counters all
ingress:

ID     Group    Color     Packets               Bytes          description
```

```
_____________________________________________________________________________
1          1           all    0                    0                    test1
1          1         green    0                    0                    test1
1          1        yellow    0                    0                    test1
1          1           red    0                    0                    test1

egress:
ID      Group    Color      Packets                Bytes           description
_____________________________________________________________________________
1          1         green    0                    0                    test2
1          1        yellow    0                    0                    test2

prelookup:
```

Clearing counters

You can use the GUI or CLI to clear the counters associated with all policies or the counters associated with just ingress, egress, or prelookup policies.

Using the GUI:

1. Go to *Switch* > *Monitor* > *ACL Counters*.
2. Select *Ingress*, *Egress*, *Prelookup*, or *All* to clear those counters.

Using the CLI:

```
execute acl clear-counter {all | egress | ingress | prelookup}
```

Clearing unused classifiers

Use the following command to clear the unused classifiers on ASIC hardware associated with ingress, egress, prelookup, or all policies for a particular group:

```
execute acl key-compaction {all | ingress | egress | prelookup} <group_ID>
```

NOTE: This command currently only works on the ingress policy.

Configuration examples

Example 1

In the following example, traffic from VLAN 3 is blocked to a specified destination IP subnet (10.10.0.0/16) but allowed to all other destinations:

```
config switch acl ingress
   edit 1
      config action
         set count enable
         set drop enable
      end
      config classifier
         set dst-ip-prefix 10.10.0.0 255.255.0.0
         set vlan-id 3
      end
      set ingress-interface-all enable
      set status active
   end
```

Example 2

In the following example, Server Message Block (SMB) traffic received on port 1 is mirrored to port 3. SMB protocol uses port 445:

```
config switch acl service custom
   edit "SMB"
      set tcp-portrange 445
   next
end
config switch acl ingress # apply policy to port 1 ingress and send to port 3
   edit 1
      set description "cnt_n_mirror_smb"
      set ingress-interface-all disable
      set ingress-interface "port1"
      set status active
      config action
         set count enable
         set mirror mirror-1
      end
      config classifier
         set service "SMB"
         set src-ip-prefix 20.20.20.100 255.255.255.255
         set dst-ip-prefix 100.100.100.0 255.255.255.0
      end
   next
end
```

Example 3

The FortiSwitch unit can map different flows (for example, based on source and destination IP addresses) to specific outgoing ports.

In the following example, flows are redirected (based on destination IP) to different outgoing ports, connected to separate FortiDDOS appliances. This allows you to apply different FortiDDOS service profiles to different types of traffic:

```
config switch acl ingress # apply policy to port 1 ingress and send to port 3
   edit 1
      config action
         set count enable
         set redirect "port3‟ # use redirect to shift selected traffic to new destination
      end
      config classifier
         set dst-ip-prefix 100.100.100.0 255.255.255.0
      end
      set description "cnt_n_mirror13"
      set ingress-interface "port1"
      set status active
   next
   edit 2
      config action # apply policy to port 3 ingress and send to port 1
         set count enable
         set redirect "port1"
      end
      config classifier
         set src-ip-prefix 100.100.100.0 255.255.255.0
      end
      set description "cnt_n_mirror31"
      set ingress-interface-all disable
      set ingress-interface "port3"
      set status active
   next
end

config switch acl ingress # apply policy to port 1 ingress and send to port 4
   edit 3
      config action
         set count enable
         set redirect "port4‟ # use redirect to shift selected traffic to new destination
      end
      config classifier
         set dst-ip-prefix 20.20.20.0 255.255.255.0
      end
      set description "cnt_n_mirror14"
      set ingress-interface "port1"
      set status active
   next
   edit 4
      config action # apply policy to port 4 ingress and send to port 1
         set count enable
         set redirect "port1"
      end
      config classifier
         set src-ip-prefix 20.20.20.0 255.255.255.0
      end
      set description "cnt_n_mirror41"
      set ingress-interface "port4"
      set status active
   next
end
```

Example 4

In the following example, a recurring schedule is created and then used to control when the ACL policy is active:

```
config system schedule recurring
   edit schedule2
      set day monday tuesday wednesday thursday friday saturday sunday
      set start 07:00
      set end 17:00
   end
config switch acl ingress
   edit 1
      config action
         set remark-cos 1
         set remark-dscp 23
      end
      config classifier
         set src-mac 00:21:cc:d2:76:72
         set dst-mac d6:dd:25:be:2c:43
      end
      set ingress-interface-all enable
      set schedule schedule2
      set status active
   next
end
```

Example 5

In the following example, the ACL policy at the ingress stage is used to mirror traffic from VLAN 100:

```
config switch mirror
   edit "m1"
      set status active
      set dst "port4"
   next
end

config switch acl ingress
   edit 1
      config action
         set mirror "m1"
      end
      config classifier
         set vlan-id 100
      end
      set ingress-interface-all enable
   next
end
```

Selective packet sampling

NOTE: This feature is not supported on FS-3032.

During debugging, you might want to see whether a particular type of packet was received on an interface on the switch.

1. Set up an access control list (ACL) on the switch with the interface that you want to monitor. See ACL on page 271. This ACL is the ingress interface.
2. Set up a mirror for the “internal” interface.

For example, if you want to monitor interface port17 for any IP packet (ether-type 0x800) with a destination subnet of 10.10.10/24 and a source subnet of 20.20.20/24, use the following commands.

```
# show switch acl ingress
config switch acl ingress
   edit 1
      config action
         set mirror "internal"
      end
      config classifier
         set dst-ip-prefix 10.10.10.0 255.255.255.0
         set ether-type 0x0800
         set src-ip-prefix 20.20.20.0 255.255.255.0
      end
      set ingress-interface "port17"
      set status active
   next
end
```

To examine the packets that have been sampled in the example, use the following command:

```
# diagnose sniffer packet sp17 none 6
```

Creating a schedule

Use schedules to control when policies are enforced. For example, you can use a schedule to control when an access control list policy is enforced.

NOTE: If the status of an ACL policy is inactive, the schedule is ignored.

You can create a one-time schedule, a recurring schedule, or a group schedule:

- Use a one-time schedule when you want a policy enforced for a specified period.
- Use a recurring schedule when you want a policy enforced for specified hours and days every week.
- Use a group schedule to combine one-time schedules and recurring schedules.

To create a one-time schedule:

```
config system schedule onetime
   edit <schedule_name>
      set start <time_date>
      set end <time_date>
   end
```

For example:

```
config system schedule onetime
   edit schedule1
      set start 07:00 2019/03/22
      set end 07:00 2019/03/29
   end
```

To create a recurring schedule:

```
config system schedule recurring
   edit <schedule_name>
      set day {monday | tuesday | wednesday | thursday | friday | saturday | sunday}
      set start <time>
      set end <time>
   end
```

For example:

```
config system schedule recurring
   edit schedule2
      set day monday wednesday friday
      set start 07:00
      set end 08:00
   end
```

To create a group schedule:

```
config system schedule group
   edit <schedule_group_name>
      set member <schedule_name1> <schedule_name2> ...
   end
```

For example:

```
config system schedule group
   edit group1
      set member schedule1 schedule2
   end
```

IGMP snooping

The FortiSwitch unit uses the information passed in IGMP messages to optimize the forwarding of IPv4 multicast traffic.

IGMP snooping allows the FortiSwitch unit to passively listen to the Internet Group Management Protocol (IGMP) network traffic between hosts and routers. The switch uses this information to determine which ports are interested in receiving each multicast feed. The FortiSwitch unit can reduce unnecessary multicast traffic on the LAN by pruning multicast traffic from links that do not contain a multicast listener.

Essentially, IGMP snooping is a layer-2 optimization for the layer-3 IGMP.

The current version of IGMP is version 3, and the FortiSwitch unit is also compatible with IGMPv1 and IGMPv2.

Starting in FortiSwitchOS 6.4.3, you can configure the IGMP-snooping querier version 2 or 3. When the IGMP querier version 2 is configured, the FortiSwitch unit will send IGMP queries version 2 when no external querier is present. When the IGMP querier version 3 is configured, the FortiSwitch unit will send IGMP queries version 3 when no external querier is present. The default IGMP querier version is 2.

Here is the basic IGMP snooping operation:

1. A host expresses interest in joining a multicast group. (Sends or responds to a join message).
2. The FortiSwitch unit creates an entry in the layer-2 forwarding table (or adds the host's port to an existing entry). The switch creates one table entry per VLAN per multicast group.
3. The FortiSwitch unit removes the entry when the last host leaves the group (or when the entry ages out).

In addition, you can configure the FortiSwitch unit to send periodic queries from all ports in a specific VLAN to request IGMP reports. The FortiSwitch unit uses the IGMP reports to update the layer-2 forwarding table.

NOTE: If you want to use IGMP snooping with an MCLAG, see Configuring an MCLAG with IGMP snooping on page 164.

Notes

When either IGMP snooping or MLD snooping is enabled in a VLAN, both unknown IPv4 and IPv6 multicast traffic, that is, unregistered multicast traffic, will share the same flooding behavior because of hardware limitations. Unregistered multicast traffic will only be forwarded to multicast IPv4 or IPv6 router ports or a switch interface with `mcast-snooping-flood-traffic` enabled.

If the network has both IPv4 and IPv6 IGMP/MLD hosts, you need to enable both IGMP and MLD snooping on the VLAN if snooping is required in the VLAN, or you need to disable both IGMP and MLD snooping on the VLAN if snooping is not required in the VLAN.

- Multicast addresses with a destination of 239.x.x.x will flood within the VLAN. This issue affects the FS-124F, FS-124F-POE, FS-124F-FPOE, FS-148F, FS-148F-POE, and FS-148F-FPOE models.
- Platforms that support MAC-based IGMP snooping might convert IGMP IP groups to MAC addresses in hardware, and sometimes the IGMP IP groups might collapse to the same MAC addresses. Even so, the maximum number of supported IGMP groups is still the maximum number of IGMP IP groups. Counting the number of IGMP IP groups allows consistency across all platforms, as well as reporting the conservative worst-case numbers. To view the maximum supported IGMP groups, use the `get switch igmp-snooping status` command and then check the value of the *Max IGMP snooping groups* field.
- To make well-known multicast packets, such as mDNS, flood to all ports when IGMP snooping is enabled on FSR-112D-POE or FSR-216F-POE, you need to make the following configuration change:

  In 6.2.x through 6.4.2 GA:

  ```
  config switch igmp-snooping globals
     set flood-unknown-multicast enable
  end
  ```

  In 6.4.3 GA and later:

  ```
  config switch global
     set flood-unknown-multicast enable
  end
  ```

- On the FS-100E series, IGMP snooping can be enabled on a maximum of 6 VLANs.
- Enabling the `set flood-unknown-multicast` command and then disabling it disrupts the forwarding of unknown multicast traffic to mRouter ports for a short period, depending on the query interval, because the mRouter ports need to be relearned.
- The IGMP group's source address(es) in the IGMPv3 report are not considered.

- The IGMP snooping entries are added based on multicast group MAC addresses.
- When IGMP snooping is enabled on a VLAN on the FSR-112D-POE or FSR-216F-POE model:
  - All IPv6 multicast and any non-IP multicast are forwarded to querier ports only instead of getting flooded on the VLAN. The forwarding of IPv6 to the CPU is unchanged.
  - IPv4 reserved multicast is flooded to the VLAN and not forwarded to the CPU, even if the CPU is part of the VLAN.
  - Unregistered IPv4 multicast is forwarded to querier ports only.
    If IPv6 multicast and/or non-IP multicast is expected to be forwarded to any ports other than querier ports, the `mcast-snooping-flood-traffic` setting can be enabled on the required ports.
- Starting with FortiSwitchOS 6.4.0, when an inter-switch link (ISL) is formed automatically, the `igmp-snooping-flood-reports` and `mcast-snooping-flood-traffic` options are disabled by default.
- Proxy reporting is not supported for IGMPv3.
- Explicit host tracking is not supported.
- Immediate leave for IGMPv3 is not supported.
- Starting with FortiSwitchOS 7.0.0, the following snooping table limits apply:

FortiSwitch Models	Snooping Table Limit (values have been rounded)
FS-108E and FS-124E	500
FSR-112D-POE, FS-124F, FS-148E, FS-148F, FSR-216F-POE, FS-224E, FS-248D, FS-248E, FS-424D, FS-424E, FS-424E-Fiber, FS-426E, FS-448D, FS-448E	1,000
FS-1024D and FS-1048D	4,000
FS-3032D	6,000
FS-524D, FS-548D, FS-1024E, FS-1048E, FS-T1024E, FS-T1024F-FPOE, and FS-3032E	8,000

  The listed snooping table limits are “best case” and might not be achievable in real-world environments. With a large number of groups and high activity or high join/leave rates, it takes longer to update the hardware. The default values for IGMP snooping are adequate for most environments. For larger scales, additional tuning might be required.

  Until FortiSwitchOS 3.5.1, the table limits were hardware only. The software limit for all platforms was 8192.
- When the IGMP proxy is enabled, the proxy report and proxy leave use the IP address 0.0.0.0. IGMP group-specific queries sent by the proxy use the internal querier’s IP address if it is configured.

Configuring IGMP snooping

Follow these steps to configure IGMP snooping:

1. Configuring IGMP snooping on a global level on page 288
2. (Optional) Enabling IGMP-snooping options on the interfaces on page 288
3. Configuring IGMP snooping on the VLANs on page 289
4. (Optional) Checking the IGMP-snooping configuration on page 291

Configuring IGMP snooping on a global level

By default, the maximum time (`aging-time`) that multicast snooping entries without any packets are kept is for 300 seconds. This value can be in the range of 15-3,600 seconds. By default, `flood-unknown-multicast` is disabled, and unregistered multicast packets are forwarded only to mRouter ports. If you enable `flood-unknown-multicast`, unregistered multicast packets are forwarded to all ports in the VLAN.

Using the CLI:

```
config switch igmp-snooping globals
   set aging-time <15-3600>
end

config switch global
   set flood-unknown-multicast {enable | disable}
end
```

For example:

```
config switch igmp-snooping globals
   set aging-time 500
end

config switch global
   set flood-unknown-multicast enable
end
```

Enabling IGMP-snooping options on the interfaces

Optional. You can flood IGMP reports and flood multicast traffic on a specified switch interface. By default, these options are disabled.

Using the GUI:

1. Go to *Switch* > *Interfaces*.
2. Select an interface.
3. Click *Edit*.
4. Under *IGMP Snooping*, select *Flood Reports*, *Flood Traffic*, or both if needed.
5. Click *Update*.

Using the CLI:

```
config switch interface
   edit <port>
      set native-vlan <vlan-id>
      set igmp-snooping-flood-reports {enable | disable}
      set mcast-snooping-flood-traffic {enable | disable}
   next
end
```

For example:

```
config switch interface
   edit port10
      set native-vlan 30
      set igmp-snooping-flood-reports enable
      set mcast-snooping-flood-traffic enable
   next
   edit port2
      set native-vlan 30
      set igmp-snooping-flood-reports enable
      set mcast-snooping-flood-traffic enable
   next
   edit port4
      set native-vlan 30
      set igmp-snooping-flood-reports enable
      set mcast-snooping-flood-traffic enable
   next
   edit port6
      set native-vlan 30
      set igmp-snooping-flood-reports enable
      set mcast-snooping-flood-traffic enable
   next
   edit port8
      set native-vlan 30
      set igmp-snooping-flood-reports enable
      set mcast-snooping-flood-traffic enable
   next
end
```

Use one of the following commands to clear the learned/configured multicast groups:

- `execute clear switch igmp-snooping all`
- `execute clear switch igmp-snooping group <multicast_IPv4_address>`
- `execute clear switch igmp-snooping interface <interface_name>`
- `execute clear switch igmp-snooping vlan <VLAN_ID>`

You can combine the commands for more control.

To clear one IGMP-snooping group from one VLAN for all interfaces:

```
execute clear switch igmp-snooping group 1.2.3.4 100
```

To clear one IGMP-snooping group from one VLAN on one interface:

```
execute clear switch igmp-snooping group 1.2.3.4 100 port1
```

To clear all IGMP-snooping groups from one interface for one VLAN:

```
execute clear switch igmp-snooping interface port1 100
```

Configuring IGMP snooping on the VLANs

Enable IGMP snooping on a specified VLAN and configure IGMP static groups. By default, IGMP snooping is disabled.

You can define static groups for particular multicast addresses in a VLAN that has IGMP snooping enabled. You can specify multiple ports in the static group, separated by a space. The trunk interface can also be included in a static group. There are two restrictions for IGMP static groups:

- The range of multicast addresses (mcast-addr) from 224.0.0.1 to 224.0.0.255 cannot be used.
- The VLAN must already be assigned as the native VLAN for a switch interface and be included in the range of allowed VLANs for a switch interface. You can check the Physical Port Interfaces page to see which VLANs can be used for IGMP static groups.

Starting in FortiSwitchOS 6.2.0, you can also use the CLI to enable IGMP proxy, which allows the VLAN to send IGMP reports. After you enable `igmp-snooping-proxy` on a VLAN, it will start suppressing reports and leave messages. For each multicast group, only one report is sent to the upstream interface. When a leave message is received, the FortiSwitch unit will only send the leave message to the upstream interface when there are no more members left in the multicast group. The FortiSwitch unit will also reply to generic queries and will send IGMP reports to the upstream interface.

Starting in FortiSwitchOS 7.2.0, you can now configure an IGMP static group to ignore requests from other ports to become members. Preventing other ports from joining means that administrators control which ports receive traffic. This option is available in the GUI and CLI; it is disabled by default, which allows other ports to dynamically join.

Using the GUI:

1. Go to *Switch* > *VLAN*.
2. Select *Add VLAN*.
3. In the *ID* field, enter the VLAN identifier.
4. In the *Description* field, enter a description for the new VLAN.
5. Under *IGMP Snooping*, select *Enable*.
6. Optionally, select *IGMP Proxy*.
7. Optionally, select *IGMP Querier*. If you select *IGMP Querier*, you must enter the primary server address in the *Primary Server* field. You can also select the IGMP-snooping querier version 2 or 3.
8. Under *IGMP Static Groups*, select + to add an IGMP static group.
   NOTE: If the VLAN identifier that you entered in step 3 is not already assigned as the native VLAN for an interface and is not included in the range of allowed VLANs for an interface, the + button is not displayed.
9. In the *Name* field, enter a name for the IGMP static group.
10. In the *Multicast Address* field, enter the multicast address.
11. Select the interfaces to include.
12. Enable *Ignore Reports* if you want to prevent other ports from becoming members.
13. Select *Update* to create the new VLAN.

Using the CLI:

```
config switch vlan
   edit <vlan-id>
      set igmp-snooping {enable |disable}
      set igmp-snooping-proxy {enable | disable}
      set igmp-snooping-fast-leave {enable | disable}
      config igmp-snooping-static-group
         edit <group-name>
            set mcast-addr <IPv4_multicast_address>
            set members <interface_name1> <interface_name2>...
            set ignore-reports {enable | disable}
         next
```

```
      end
   next
end
```

For example, to configure two static groups for the same VLAN:

```
config switch vlan
   edit 30
      set igmp-snooping enable
      config igmp-snooping-static-group
         edit g239-1-1-1
            set mcast-addr 239.1.1.1
            set members port2 port5 port28
            set ignore-reports enable
         next
         edit g239-2-2-2
            set mcast-addr 239.2.2.2
            set members port5 port10 trunk-1
            set ignore-reports enable
         next
      end
   next
end
```

Checking the IGMP-snooping configuration

To display information about IGMP snooping:

```
# get switch igmp-snooping {globals | group | static-group | status}
```

- `globals`: display the IGMP-snooping global configuration on the FortiSwitch unit
- `group`: display a list of learned multicast groups
- `static-group`: display the list of configured static groups
- `status`: display the status of IGMP-snooping VLANs and group

To view the learned multicast groups in the GUI:

Go to *Switch* > *Monitor* > *IGMP Snooping*.

To view the learned multicast groups in the CLI:

```
FS1D243Z13000023 # get switch igmp-snooping group
Number of Groups: 7
port of-port VLAN GROUP Age
(__port__9) 1 23 231.8.5.4 16
(__port__9) 1 23 231.8.5.5 16
(__port__9) 1 23 231.8.5.6 16
(__port__9) 1 23 231.8.5.7 16
(__port__9) 1 23 231.8.5.8 16
(__port__9) 1 23 231.8.5.9 16
(__port__9) 1 23 231.8.5.10 16
(__port__43) 3 23 querier 17
(__port__14) 8 --- flood-reports ---
(__port__10) 2 --- flood-traffic ---
```

To display the list of configured static groups:

```
FS1D243Z13000023 # get switch igmp-snooping static-group

VLAN ID Group-Name     Multicast-addr  Member-interface
_______ _____________ _______________ ____________________________
11      g239-1         239:1:1:1       port6 trunk-2
11      g239-11        239:2:2:11      port26 port48 trunk-2
40      g239-1         239:1:1:1       port5 port25 trunk-2
40      g239-2         239:2:2:2       port25 port26
```

Configuring the IGMP querier

To use the IGMP querier, you need to configure how often IGMP queries are sent and enable the IGMP querier for a specific VLAN. By default, IGMP queries are sent every 120 seconds. You must specify the address for the IGMP querier.

To specify how many seconds are between IGMP queries:

```
config switch igmp-snooping globals
   set query-interval <10-1200>
end
```

For example:

```
config switch igmp-snooping globals
   set aging-time 300
   set query-interval 125
end
```

To enable the IGMP querier for a specific VLAN and specify the address that IGMP reports are sent to:

```
config switch vlan
   edit 100
      set igmp-snooping {enable | disable}
      set igmp-snooping-querier {enable | disable}
      set igmp-snooping-querier-addr <IPv4_address>
```

```
      set igmp-snooping-querier-version {2 | 3}
   next
end
```

For example:

```
config switch vlan
   edit 100
      set igmp-snooping enable
      set igmp-snooping-querier enable
      set igmp-snooping-querier-addr 1.2.3.4
      set igmp-snooping-querier-version 3
   next
end
```

Configuring mRouter ports

NOTE: These settings are not per-VLAN, so the port will act as a querier/mRouter port for all of its associated VLANs.

To configure a FortiSwitch port as an mRouter port:

```
config switch interface
   edit <port>
      set igmp-snooping-flood-reports enable
      set mcast-snooping-flood-traffic enable
   next
end
```

MLD snooping

The FortiSwitch unit uses the information passed in Multicast Listener Discovery (MLD) messages to optimize the forwarding of IPv6 multicast traffic.

MLD snooping allows the FortiSwitch unit to passively listen to the MLD network traffic between hosts and multicast routers. The switch uses this information to determine which hosts are interested in receiving each multicast feed. The FortiSwitch unit can reduce unnecessary multicast traffic on the VLAN by pruning multicast traffic from links that do not contain a multicast listener.

FortiSwitch MLD snooping supports MLD version 1. RFC 2710 describes MLD snooping; RFC 4605 describes MLD proxy and MLD querier.

Here is the basic MLD-snooping operation:

1. A host expresses interest in joining a multicast group. (Sends or responds to a join message).
2. The FortiSwitch unit creates one table entry per VLAN per multicast group per port.
3. The FortiSwitch unit removes the entry when the last host leaves the group (or when the entry ages out).

In addition, you can configure the FortiSwitch unit to send periodic queries from all ports in a specific VLAN to request MLD reports. The FortiSwitch unit uses the MLD reports to update the layer-2 forwarding table.

Notes

When either IGMP snooping or MLD snooping is enabled in a VLAN, both unknown IPv4 and IPv6 multicast traffic, that is, unregistered multicast traffic, will share the same flooding behavior because of hardware limitations. Unregistered multicast traffic will only be forwarded to multicast IPv4 or IPv6 router ports or a switch interface with `mcast-snooping-flood-traffic` enabled.

If the network has both IPv4 and IPv6 IGMP/MLD hosts, you need to enable both IGMP and MLD snooping on the VLAN if snooping is required in the VLAN, or you need to disable both IGMP and MLD snooping on the VLAN if snooping is not required in the VLAN.

- Enabling the `set flood-unknown-multicast` command and then disabling it disrupts the forwarding of unknown multicast traffic to mRouter ports for a short period, depending on the query interval, because the mRouter ports need to be relearned.
- The MLD-snooping entries are added based on multicast group IP addresses.
- Starting with FortiSwitchOS 7.0.0, the following snooping table limits apply:

FortiSwitch Models	Snooping Table Limit (values have been rounded)
FS-1024D and FS-1048D	1,800
FS-3032D	3,000
FS-524D, FS-548D, and FS-3032E	6,000
FS-1024E, FS-1048E, FS-T1024E, and FS-T1024F-FPOE	8,000

The listed snooping table limits are “best case” and might not be achievable in real-world environments. With a large number of groups and high activity or high join/leave rates, it takes longer to update the hardware. The default values for MLD snooping are adequate for most environments. For larger scales, additional tuning might be required.

Configuring MLD snooping

Configuring MLD snooping consists of the following major steps:

1. Configuring MLD snooping on a global level on page 294
2. (Optional) Enabling MLD-snooping options on the interfaces on page 295
3. Configuring MLD snooping on the VLANs on page 296
4. (Optional) Checking the MLD-snooping configuration on page 298

Configuring MLD snooping on a global level

By default, the maximum time (`aging-time`) that multicast snooping entries without any packets are kept is for 300 seconds. This value can be in the range of 15-3,600 seconds. By default, `flood-unknown-multicast` is disabled, and unregistered multicast packets are forwarded only to mRouter ports. If you enable `flood-unknown-multicast`, unregistered multicast packets are forwarded to all ports in the VLAN.

Using the CLI:

```
config switch mld-snooping globals
   set aging-time <15-3600>
end

config switch global
   set flood-unknown-multicast {enable | disable}
end
```

For example:

```
config switch mld-snooping globals
   set aging-time 500
end

config switch global
   set flood-unknown-multicast enable
end
```

Enabling MLD-snooping options on the interfaces

Optional. You can flood MLD reports and flood multicast traffic on a specified switch interface. By default, these options are disabled.

Using the CLI:

```
config switch interface
   edit <port>
      set native-vlan <vlan-id>
      set mld-snooping-flood-reports {enable | disable}
      set mcast-snooping-flood-traffic {enable | disable}
   next
end
```

For example:

```
config switch interface
   edit port10
      set native-vlan 30
      set mld-snooping-flood-reports enable
      set mcast-snooping-flood-traffic enable
   next
   edit port2
      set native-vlan 30
      set mld-snooping-flood-reports enable
      set mcast-snooping-flood-traffic enable
   next
   edit port4
      set native-vlan 30
      set mld-snooping-flood-reportsenable
      set mcast-snooping-flood-traffic enable
   next
   edit port6
      set native-vlan 30
```

```
      set mld-snooping-flood-reports enable
      set mcast-snooping-flood-traffic enable
   next
   edit port8
      set native-vlan 30
      set mld-snooping-flood-reports enable
      set mcast-snooping-flood-traffic enable
   next
end
```

Use one of the following commands to clear the learned/configured multicast groups:

- `execute clear switch mld-snooping all`
- `execute clear switch mld-snooping group <multicast_IPv6_address>`
- `execute clear switch mld-snooping interface <interface_name>`
- `execute clear switch mld-snooping vlan <VLAN_ID>`

You can combine the commands for more control.

To clear one MLD-snooping group from one VLAN for all interfaces:

```
execute clear switch mld-snooping group ff3f::1 100
```

To clear one MLD-snooping group from one VLAN on one interface:

```
execute clear switch mld-snooping group ff3f::1 100 port1
```

To clear all MLD-snooping groups from one interface for one VLAN:

```
execute clear switch mld-snooping interface port1 100
```

Configuring MLD snooping on the VLANs

Enable MLD snooping on a specified VLAN and configure MLD static groups. By default, MLD snooping is disabled.

You can define static groups for particular multicast addresses in a VLAN that has MLD snooping enabled. You can specify multiple ports in the static group, separated by a space. The trunk interface can also be included in a static group. There are two restrictions for MLD static groups:

- The range of well-known IPv6 multicast addresses that cannot be used for static groups is FF00::/12.
- The VLAN must already be assigned as the native VLAN for a switch interface *or* be included in the range of allowed VLANs for a switch interface. You can check the Physical Port Interfaces page to see which VLANs can be used for MLD static groups.

You can also enable the MLD proxy, which allows the VLAN to send MLD reports. After you enable `mld-snooping-proxy` on a VLAN, it will start suppressing reports and leave messages. For each multicast group, only one report is sent to the upstream interface. When a leave message is received, the FortiSwitch unit will only send the leave message to the upstream interface when there are no more members left in the multicast group. The FortiSwitch unit will also reply to generic queries and will send MLD reports to the upstream interface. If `mld-snooping-fast-leave` is disabled, the FortiSwitch unit sends a group-specific query (GSQ) when a leave message is received.

Starting in FortiSwitchOS 7.2.0, you can now configure an MLD static group to ignore requests from other ports to become members. Preventing other ports from joining means that administrators control which ports receive traffic. This option is available in the CLI; it is disabled by default, which allows other ports to dynamically join.

Using the GUI:

1. Go to *Switch* > *VLAN*.
2. Click *Add VLAN*.
3. In the *ID* field, enter the VLAN identifier.
4. In the *Description* field, enter a description for the new VLAN.
5. In the *MLD Snooping* area, select *Enable*.
6. Optionally, select *MLD Proxy*.
7. Optionally, select *MLD Querier*. If you select *MLD Querier*, you must enter the querier address in the *Querier Address* field.

   NOTE: The querier address cannot be an IPv6 multicast or loopback address.
8. In the *MLD Static Groups* area, select + to add an MLD static group.

   NOTE: If the VLAN identifier that you entered in step 3 is not already assigned as the native VLAN for an interface and is not included in the range of allowed VLANs for an interface, the + button is not displayed.
9. In the *Name* field, enter a name for the MLD static group.
10. In the *Multicast Address* field, enter the multicast address.

    NOTE: The multicast address cannot be a reserved multicast address (ff0*x*::).
11. Select the interfaces to include.
12. Enable *Ignore Reports* if you want to prevent other ports from becoming members.
13. Select *Add* to create the new VLAN.

Using the CLI:

```
config switch vlan
   edit <vlan-id>
      set mld-snooping {enable |disable}
      set mld-snooping-fast-leave {enable |disable}
      set mld-snooping-querier {enable |disable}
      set mld-snooping-querier-addr <IPv6_MLD_querier_address>
      set mld-snooping-proxy {enable | disable}
      config mld-snooping-static-group
         edit <group-name>
            set mcast-addr <IPv6_multicast_address>
            set members <interface_name1> <interface_name2>...
            set ignore-reports {enable | disable}
         next
      end
   next
end
```

For example:

```
config switch vlan
   edit 30
      set mld-snooping enable
      set mld-snooping-fast-leave enable
      set mld-snooping-querier enable
      set mld-snooping-querier-addr 2001::1
      set mld-snooping-proxy disable
      config mld-snooping-static-group
         edit g239-1-1-1
```

```
            set mcast-addr FF3E::1
            set members port2 port5 port28
            set ignore-reports enable
         next
      end
   next
end
```

Checking the MLD-snooping configuration

Use the following commands to display information about MLD snooping:

```
# get switch mld-snooping {globals | group | static-group | status}
```

- `globals`: display the MLD-snooping global configuration on the FortiSwitch unit
- `group`: display a list of learned multicast groups
- `static-group`: display the list of configured MLD static groups
- `status`: display the status of MLD-snooping VLANs and group

Configuring the MLD querier

To use the MLD querier, you need to configure how often MLD queries are sent and enable the MLD querier for a specific VLAN.You must specify the address for the MLD querier.

To specify how many seconds are between MLD queries. The default is 125 seconds.

```
config switch mld-snooping globals
   set query-interval <10-1200>
end
```

For example:

```
config switch mld-snooping globals
   set aging-time 150
   set query-interval 200
end
```

To enable the MLD querier for a specific VLAN and specify the address that MLD reports are sent to:

```
config switch vlan
   edit 100
      set mld-snooping {enable | disable}
      set mld-snooping-querier {enable | disable}
      set mld-snooping-querier-addr <IPv6_address>
   next
end
```

For example:

```
config switch vlan
   edit 100
      set mld-snooping enable
      set mld-snooping-querier enable
```

```
      set mld-snooping-querier-addr fe80::a5b:eff:fef1:95e5
    next
end
```

PoE

Power over Ethernet (PoE) describes any system that passes electric power along with data on twisted pair Ethernet cabling. Doing this allows a single cable to provide both data connection and electric power to devices (for example, wireless access points, IP cameras, and VoIP phones).

PoE is only available on models with the POE suffix in the model number (for example, FS-108E-POE).

Using the GUI:

1. Go to *Switch* > *POE*.

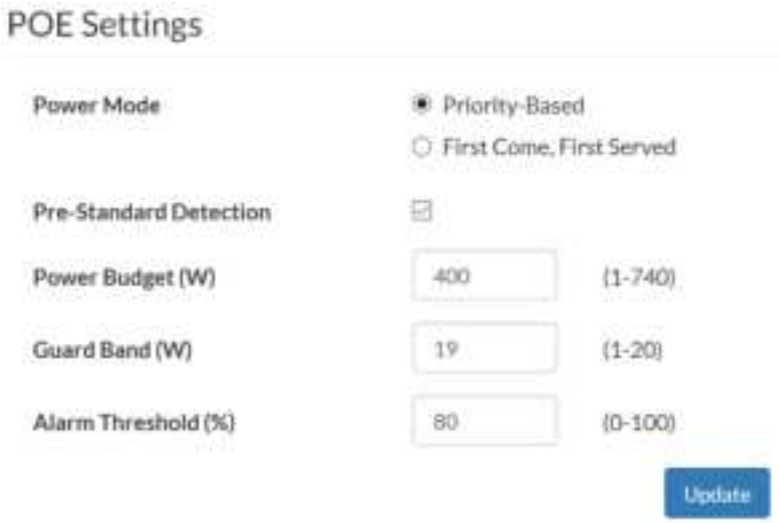

2. Set the PoE power mode to priority based or first-come, first-served.
   When power to PoE ports is allocated by priority, lower numbered ports have higher priority so that port 1 has the highest priority. When more power is needed than is available, higher numbered ports are disabled first.
   When power to PoE ports is allocated by first-come, first-served (FCFS), connected PoE devices receive power, but new devices do not receive power if there is not enough power.
   If both priority power allocation and FCFS power allocation are selected, the physical port setting takes precedence over the global setting.

3. Enable or disable PoE pre-standard detection.

PoE pre-standard detection is a global setting for the following FortiSwitch models:

FSR-112D-POE, FS-548D-FPOE, FS-524D-FPOE, FS-108D-POE, FS-224D-POE, FS-108E-POE, FS-108E-FPOE, FS-124E-POE, and FS-124E-FPOE.

For the other FortiSwitch PoE models, PoE pre-standard detection is set on each port.

4. Set the maximum power budget in Watts.
5. Enter the power in Watts to reserve in case of a spike in PoE consumption.
6. Enter the threshold (a specified percentage of the total power budget) above which an alarm event is generated. If your FortiSwitch unit has a PoE sensor, you can set an alarm for when the current power budget exceeds a specified percentage of the total power budget. When this threshold is exceeded, log messages and SNMP traps are generated. The default threshold is 80 percent.
7. Select *Update*.

Using the CLI:

```
config switch global
   set poe-alarm-threshold <0-100 percent>
   set poe-power-mode {first-come-first-serverved | priority}
   set poe-guard-band <1-20 Watts>
   set poe-pre-standard-detect {disable | enable}
   set poe-power-budget <1-740 Watts>
end
```

sFlow

sFlow is a method of monitoring the traffic on your network to identify areas on the network that may impact performance and throughput. With sFlow you can export truncated packets and interface counters. The FortiSwitch unit implements sFlow version 5 and supports trunks and VLANs.

sFlow uses packet sampling to monitor network traffic. The sFlow agent captures packet information at defined intervals and sends them to sFlow collectors for analysis, providing real-time data analysis. To minimize the impact on network throughput, the information sent is only a sampling of the data.

The sFlow collector is a central server running software that analyzes and reports on network traffic. The sampled packets and counter information, referred to as flow samples and counter samples, respectively, are sent as sFlow datagrams to collectors. Upon receiving the datagrams, the sFlow collectors provide real-time analysis and graphing to indicate the source of potential traffic issues. sFlow collector software is available from a number of third-party software vendors.

Configuring sFlow

Configuration consists of the following steps:

1. Configuring sFlow agents on page 301
2. Configuring the interfaces on page 301

Configuring sFlow agents

Starting in FortiSwitchOS 7.0.0, you can configure multiple collectors.

Using the GUI:

1. Go to *Switch* > *sFlow*.
2. Click *Add*.
3. Enter the collector name, IPv4 address, and port number.
   The collector port number is the destination port number for sFlow UDP packets. The default value is 6343.
4. Click *Apply*.

Using the CLI:

```
config system sflow
   config collectors
      edit <collector_name>
         set ip <collector_IPv4_address>
         set port <0-65535>
      next
   end
end
```

For example:

```
config system sflow
   config collectors
      edit collector1
         set ip 20.20.20.0
         set port 200
      next
   end
end
```

Configuring the interfaces

To configure sFlow on a port:

- Enable sFlow on the port (CLI only).
- Set the sample rate (CLI only). An average of one out of `count` packets is randomly sampled. The rate ranges from 0-99999; the default is 512. **NOTE:** The sample rate is 0-65535 on the FS-124F, FS-124F-POE, FS-124F-FPOE, FS-148E, FS-148E-POE, FS-148F, FS-148F-POE, and FS-148F-FPOE models.
- Set the direction for capturing the traffic (CLI only). sFlow can capture the ingress traffic (RX), the egress traffic (TX), or both (the default).
- Set the polling interval, which defines how often the switch sends interface counters to the collector. The range of values is 1-255 and default is 30.

Using the GUI:

1. Go to *Switch* > *Interfaces*.
2. Select one or more ports or a trunk to update and then click *Edit*.
3. Under *sFlow*, select *Polling Interval*.
4. In the *Interval (Seconds)* field, enter the number of seconds to use for the polling interval.
5. Click *Update* to save the changes.

Using the CLI:

```
config switch interface
   edit <port>
      set packet-sampler {enabled | disabled}
      set packet-sample-rate <count>
      set sample-direction {rx | tx | both}
      set sflow-counter-interval <interval>
   next
end
```

For example:

```
config switch interface
   edit "port20"
      set packet-sampler enabled
      set packet-sample-rate 4
      set sflow-counter-interval 3
      set snmp-index 58
   next
end
```

NOTE: Ensure that you can use the exec command `ping collector_ip_address` to ping the collector from the FortiSwitch unit. Then, use the built-in sniffer to trace sFlow packets (`diag sniff packet <vlan_interface_name> "udp port 6343"`).

Checking the sFlow configuration

Use the following command to display the sFlow configuration:

```
get system sflow
```

Mirror

Packet mirroring allows you to collect packets on specified ports and then send them to another port to be collected and analyzed. All FortiSwitch models support switched port analyzer (SPAN) mode, which mirrors traffic to the specified destination interface without encapsulation.

Using remote SPAN (RSPAN) or encapsulated RSPAN (ERSPAN) allows you to send the collected packets across layer-2 domains. You can have multiple RSPAN sessions but only one ERSPAN session. In RSPAN mode, traffic is

encapsulated in a VLAN. In ERSPAN mode, traffic is encapsulated in Ethernet, IPv4, and generic routing encapsulation (GRE) headers.

You can also use an ACL policy at the ingress or egress stage to mirror packets to another port, interface, or trunk. You can mirror traffic from VLANs as well. For more details, refer to ACL on page 271.

NOTE:

- Mirror sources cannot also be mirror destinations or members of mirror destinations if the destination is a trunk. When using RSPAN or ERSPAN in FortiLink mode, the destination ports or trunks are determined automatically (the automatically determined port can be viewed with the `diagnose switch-controller switch-info mirror status` command on the FortiGate device). The destination is often an ISL interface towards the FortiGate device. This destination can cause conflicts if the user tries to configure ports in the ISL as source ports. In the case of conflict, Fortinet recommends disabling the FortiLink traffic sniffer or omitting ports that are part of the ISL.
- Some models support setting the mirror destination to “internal.” This is intended only for debugging purposes and might prevent critical protocols from operating on ports being used as mirror sources.
- When there are multiple mirror sessions in the FS-108D-POE, FS-224D-POE, FSR-112D-POE, and FSR-216F-POE models, some traffic might not be mirrored to the destination ports.
- Some destination ports are not listed because those models (FSR-112D-POE, FS-108E, FS-124E, FS-108E-POE, FS-108E-FPOE, FS-124E-POE, FS-124E-FPOE, FS-148E, FS-148E-POE, FS-108F, FS-108F-POE, FS-108F-FPOE, FS-124F, FS-124F-POE, FS-124F-FPOE, FS-148F, FS-148F-POE, FS-148F-FPOE, and FSR-216F-POE) do not support mirroring to the software interface.
- You cannot select a destination interface for the ERSPAN auto mirror.
- In cases where the mirrored traffic is not unicast, or is flooded unicast, and the mirrored and non-mirrored packets both leave the mirror “dst” port, the `mirror-qos` value is overridden by the QoS value of the non-mirrored packet.
- You can use the following commands to specify the quality of service (QoS) priority for mirrored packets on the FortiSwitch unit doing the mirroring:
```
config switch global
   set mirror-qos <0-7>
end
```
- All FortiSwitch models can mirror STP (BPDU) and LLDP frames using SPAN and RSPAN.

Some of the platform differences are listed in the following table:

	1120-POE, 216E-POE	108E, 108E-FPOE, 108E-POE, 108F, 108F-POE, 108-FPOE, 124E, 124E-FPOE, 124E-POE, 124F, 124F-POE, 124F-FPOE, 148F, 148F-POE, 148F-FPOE	1240, 224D-FPOE, 224E, 224E-POE	248D, 248E-FPOE, 248E-POE	424D, 424D-FPOE, 424D-POE	448D, 448D-FPOE, 448D-POE	424E, 424E-POE, 424E-FPOE, M426-FPOE, 424F-POE	424E-Fiber, 448E, 448E-POE, 448E-FPOE	524D, 524D-FPOE, 548D, 548D-FPOE, 1048E	4xxF	1024D, 1024E, T1024E, 1048D, 2048E, M5112, M5132
"dot" values	Ports only (can be in trunk)	Ports only (can be in trunk)	Port or trunk (no trunk members)	Port or trunk (no trunk members)	Port or trunk (no trunk members)	Port or trunk (no trunk members)	Port or trunk (no trunk members)	Port or trunk (no trunk members)	Port or trunk (no trunk members)	Port only (can be in trunk)	Port or trunk (no trunk members)
Maximum sessions (active or inactive)	--	--	32	32	32	32	32	32	32	32	32
Maximum active sessions	7	4	8	8	8	8	8	8	8	7	4
Maximum sessions with src egress	6	4	1	1	1	1	1	1	1	7	1
Maximum sessions with src ingress	6	4	4	4	4	4	4	4	4	7	4
Same src ingress in multiple active sessions	No	Yes	No	No	No	No	No	No	Yes	No	Yes
Maximum sessions when one has src ingress + src egress and the rest are src ingress	N/A	N/A	5	5	5	5	5	5	5	7	5
VLAN CFI and priority can be configured in RSPAN	N/A	N/A	Yes	No	Yes	No	Yes	Yes	Yes	Yes	Yes
SPAN support	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
RSPAN and ERSPAN support	RSPAN	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	RSPAN	Yes
QoS support	No	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	Yes

This section covers the following topics:

- Configuring a SPAN mirror on page 304
- Configuring an RSPAN mirror on page 307
- Configuring an ERSPAN auto mirror on page 311
- Configuring an ERSPAN manual mirror on page 312

Configuring a SPAN mirror

You can use virtual wire ports as ingress and egress mirror sources. Egress mirroring of virtual wire ports will have an additional VLAN header on all mirrored traffic.

Using the GUI:

1. Go to *Switch* > *Mirror*.
2. Select *Add Port Mirror*.
3. Enter a name for the mirror.
4. Select *Enabled* to make the mirror active.
5. Select a destination interface.
   On FortiSwitch models that support RSPAN and ERSPAN, set the trunk or physical port that will act as a mirror. The physical port cannot be part of a trunk.
   On FortiSwitch models that do *not* support RSPAN and ERSPAN, set the physical port that will act as a mirror. The physical port can be part of a trunk.
6. Select from the excluded ports which ports to include for ingress mirroring and egress mirroring.
   NOTE: Only one active egress mirror session is allowed.
7. Select *Packet Switching When Mirroring* if the destination port is not a dedicated port. For example, enable this option if you connect a laptop to the switch and you are running a packet sniffer along with the management GUI on the laptop.
8. Select *SPAN* for the mode.
9. Select *Create* to create the mirror.

Using the CLI:

```
config switch mirror
   edit <mirror session name>
      set mode SPAN
      set dst <interface>
      set src-egress <interface_name>
      set src-ingress <interface_name>
      set switching-packet {enable | disable}
      set status active
   end
```

For example:

```
config switch mirror
   edit "m1"
      set mode SPAN
      set dst "port5"
      set src-egress "port2"
      set src-ingress "port3" "port4"
      set switching-packet enable
      set status active
end
```

Multiple mirror destination ports (MTPs)

With some FortiSwitch models, you can configure multiple mirror destination ports with the following guidelines and restrictions:

- Always set the destination port before setting the src-ingress or src-egress ports.
- Any port configured as a src-ingress or src-egress port in one mirror cannot be configured as a destination port in another mirror.
- The total number of active sessions depends on your configuration.

- For switch models 124D, 124D-POE, 224D-FPOE, 248D, 248D-POE, 248D-FPOE, 224E, 224E-POE, 248E-POE, 248E-FPOE, 424D, 424D-POE, 424D-FPOE, 448D, 448D-POE, and 448D-FPOE:
  - For access control lists, you can use a mirror destination that does not have src-ingress or src-egress configured or a mirror destination that has src-ingress or src-egress configured.
- For switch models 524D, 524D-FPOE, 548D, 548D-FPOE, 1024D, 1048D, 1048E, 3032D, and 3032E:
  - For access control lists, you can use a mirror destination that does not have src-ingress or src-egress configured or a mirror destination that has src-ingress or src-egress configured.
- For switch models FSR-112D-POE and FSR-216F-POE:
  - You can configure up to seven mirrors, each with a different destination port.
  - Multiple ingress or egress ports can be mirrored to the same destination port.
  - An ingress or egress port cannot be mirrored to more than one destination port.

These restrictions apply to active mirrors. If you try to activate an invalid mirror configuration, the system will display the `Hardware active mirror session limit reached. Please deactivate or delete another active session to make room.` error message.

The following example configuration is valid for FortiSwitch-3032D. This configuration includes three ingress ports, one egress port, and four destination ports. The port3 ingress and egress ports are mirrored to multiple destinations.

```
config switch mirror
   edit "m1"
      set mode SPAN
      set dst "port16"
      set status active
         set src-ingress "port3" "port5" "port7"
   next
   edit "m2"
      set mode SPAN
      set dst "port22"
      set status active
         set src-ingress "port3" "port5"
   next
   edit "m3"
      set mode SPAN
      set dst "port1"
      set status active
         set src-ingress "port3"
   next
   edit "m4"
      set mode SPAN
      set dst "port2"
      set status active
         set src-egress "port3"
end
```

The following example configuration includes three ingress ports, three egress ports and four destination ports. Each ingress and egress port is mirrored to only one destination port.

```
config switch mirror
   edit "m1"
      set mode SPAN
      set dst "port1"
      set status active
         set src-ingress "port2" "port7"
   next
   edit "m2"
```

```
      set mode SPAN
      set dst "port5"
      set status active
         set src-ingress "port2"
   next
   edit "m3"
      set mode SPAN
      set dst "port3"
      set status active
         set src-ingress "port6"
   next
   edit "m4"
      set mode SPAN
      set dst "port4"
      set status active
         set src-egress "port6" "port8"
end
```

Configuring an RSPAN mirror

NOTE: RSPAN traffic crossing a switch on a VLAN configured with “RSPAN-VLAN” enabled will appear as unknown unicast, multicast, or broadcast traffic. This traffic is not exempt from storm control and might be rate limited as a result. To avoid this issue, you can dedicate a port or ports to RSPAN and then disable storm control on those ports. Non-RSPAN VLANs can be used on those ports as well, but they will not be protected by storm control.

Using the GUI:

1. Go to *Switch* > *Mirror*.
2. Select *Add Port Mirror*.
3. Enter a name for the mirror.
4. Select *Enabled* to make the mirror active.
5. Select a destination interface.
   NOTE: The destination interface cannot be part of a trunk.
6. Select from the excluded ports which ports to include for ingress mirroring and egress mirroring.
   NOTE: Only one active egress mirror session is allowed.
7. Select *Packet Switching When Mirroring* if the destination port is not a dedicated port. For example, enable this option if you connect a laptop to the switch and you are running a packet sniffer along with the management GUI on the laptop.
8. Select *RSPAN* for the mode.
9. In the VLAN ID field, enter the VLAN identifier for the RSPAN VLAN header.
10. In the TPID field, enter the tag protocol identifier (TPID) for the encapsulating VLAN header.
    The default value, 0x8100, is for an IEEE 802.1Q-tagged frame.
11. In the Priority field, enter the class of service (CoS) bits in the RSPAN VLAN header.
    NOTE: This option is not available on the 248D, 248D-POE, 248D-FPOE, 248E, 248E-POE, 248E-FPOE, 448D, 448D-POE, and 448D-FPOE models.
12. In the CFI/DEI field, enter the canonical format identifier (CFI) or drop eligible indicator (DEI) bit in the RSPAN VLAN header.
    NOTE: This option is not available on the 248D, 248D-POE, 248D-FPOE, 248E, 248E-POE, 248E-FPOE, 448D, 448D-POE, and 448D-FPOE models.
13. Select *Create* to create the mirror.

Using the CLI:

```
config switch mirror
   edit <mirror session name>
      set mode RSPAN
      set dst <interface>
      set switching-packet {enable | disable}
      set src-ingress <interface_name>
      set src-egress <interface_name>
      set encap-vlan-tpid <0x0001-0xfffe>
      set encap-vlan-priority <0-7>
      set encap-vlan-cfi <0-1>
      set encap-vlan-id <1-4094>
      set status active
   end
```

Configuration example 1

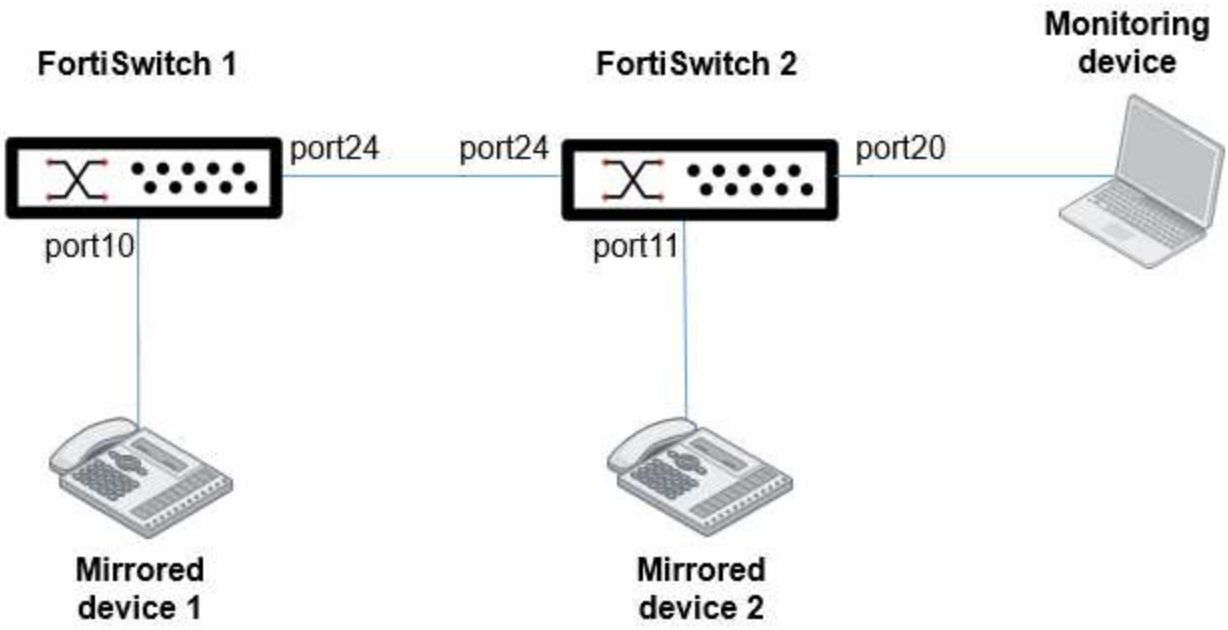

To configure FortiSwitch 1:

```
config switch vlan
   edit 300
      set rspan-mode enable
   next
end

config switch mirror
   edit "rspan"
      set status active
      set mode RSPAN-manual
      set dst "<auto-isl -switch2 >"
      set switching-packet enable
      set src-ingress "port10"
```

```
      set src-egress "port10"
      set encap-vlan-id 300
   next
end

config switch interface
   edit "port10"
      set native-vlan 101
   next
   edit "<auto-isl -switch2 >"
      set allowed-vlans 101,300
   next
end
```

To configure FortiSwitch2:

```
config switch vlan
   edit 300
      set rspan-mode enable
   next
end

config switch mirror
   edit "rspan"
      set status active
      set mode RSPAN-manual
      set dst "port20"
      set switching-packet enable
      set src-ingress "port11"
      set src-egress "port11"
      set encap-vlan-id 300
   next
end

config switch interface
   edit "port11"
      set native-vlan 101
   next
   edit "port20"
      set native-vlan 300
   next
   edit "<auto-isl -switch1 >"
      set allowed-vlans 101,300
   next
end
```

Configuration example 2

This example mirrors all ingress and egress traffic on VLAN 101 on FortiSwitch 1 and FortiSwitch 2. The monitoring device is connected to FortiSwitch 2. VLAN 101 is the VLAN used for voice traffic. VLAN 300 is being used to carry RSPAN mirrored traffic between the two switches. Auto-ISL trunks are disabled in this topology.

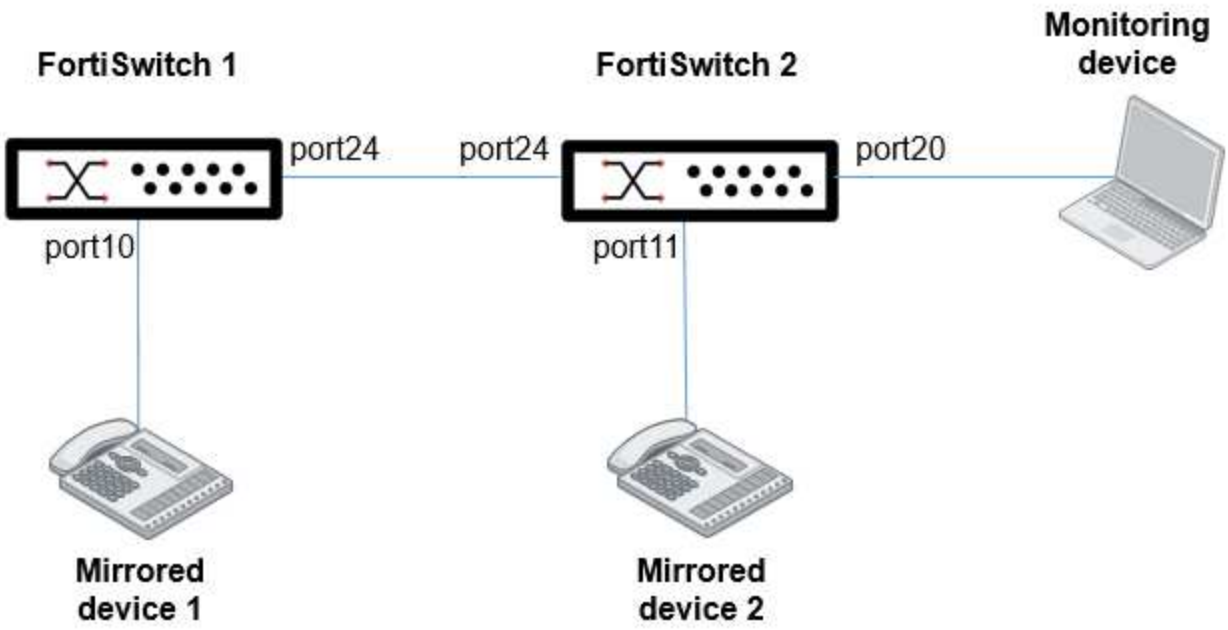

To configure FortiSwitch 1:

```
config switch vlan
   edit 300
      set rspan-mode enable
   next
end

config switch mirror
   edit "rspan"
      set status active
      set mode RSPAN-manual
      set dst "port24"
      set switching-packet enable
      set encap-vlan-id 300
   next
end

config switch acl ingress
   edit 1
      set ingress-interface port10
      config classifier
         set vlan-id 101
      end
      config action
      set count enable
         set mirror "rspan"
      end
   next
end

config switch interface
   edit "port10"
      set allowed-vlans 101
```

```
   next
   edit "port24"
      set allowed-vlans 101,300
   next
end
```

To configure FortiSwitch 2:

```
config switch vlan
   edit 300
      set rspan-mode enable
   next
end

config switch mirror
   edit "rspan"
      set status active
      set mode RSPAN-manual
      set dst "port20"
      set switching-packet enable
      set encap-vlan-id 300
   next
end

config switch acl ingress
   edit 1
      set ingress-interface port11
      config classifier
         set vlan-id 101
      end
      config action
         set count enable
         set mirror "rspan"
      end
   next
end

config switch interface
   edit "port11"
      set allowed-vlans 101
   next
   edit "port20"
      set native-vlan 300
   next
   edit "port24"
      set allowed-vlans 101,300
   next
end
```

Configuring an ERSPAN auto mirror

For an ERSPAN auto mirror, traffic on specified ports is mirrored to the specified destination interface using ERSPAN encapsulation. The header contents are automatically configured; you only need to specify the ERSPAN collector address.

Using the GUI:

1. Go to *Switch* > *Mirror*.
2. Select *Add Port Mirror*.
3. Enter a name for the mirror.
4. Select *Enabled* to make the mirror active.
5. Select from the excluded ports which ports to include for ingress mirroring and egress mirroring.
   NOTE: Only one active egress mirror session is allowed.
6. Select *ERSPAN Auto* for the mode.
7. Enable *Strip VLAN Tags from Mirrored Traffic* if you want to remove VLAN tags from mirrored traffic.
8. In the Collector IP field, enter the IP address for the ERSPAN collector.
9. In the IPv4 TTL field, enter the IPv4 time-to-live (TTL) value in the ERSPAN IP header.
10. In the IPv4 TOS field, enter the type of service (ToS) value or enter the DSCP and ECN values in the ERSPAN IP header.
11. In the GRE Protocol field, enter the protocol value in the ERSPAN GRE header.
12. In the TPID field, enter the TPID for the encapsulating VLAN header.
    The default value, 0x8100, is for an IEEE 802.1Q-tagged frame.
13. In the Priority field, enter the CoS bits in the ERSPAN VLAN header.
14. In the CFI/DEI field, enter the CFI or DEI bit in the ERSPAN VLAN header.
15. Select *Create* to create the mirror.

Using the CLI:

```
config switch mirror
   edit <mirror session name>
      set mode ERSPAN-auto
      set encap-gre-protocol <hexadecimal_integer>
      set encap-ipv4-tos <hexadecimal_integer>
      set encap-ipv4-ttl <0-255>
      set encap-vlan-cfi <0-1>
      set encap-vlan-priority <0-7>
      set encap-vlan-tpid <0x0001-0xfffe>
      set erspan-collector-ip <0.0.0.1-255.255.255.255>
      set src-egress <interface_name>
      set src-ingress <interface_name>
      set strip-mirrored-traffic-tags {disable | enable}
      set status active
   end
```

Configuring an ERSPAN manual mirror

For an ERSPAN manual mirror, traffic on specified ports is mirrored to the specified destination interface using ERSPAN encapsulation. You need to manually configure the header contents with layer-2 and layer-3 addresses.

Using the GUI:

1. Go to *Switch* > *Mirror*.
2. Select *Add Port Mirror*.
3. Enter a name for the mirror.
4. Select *Enabled* to make the mirror active.

5. Select a destination interface.
   NOTE: The destination interface cannot be part of a trunk.
6. Select from the excluded ports which ports to include for ingress mirroring and egress mirroring.
   NOTE: Only one active egress mirror session is allowed.
7. Select *Packet Switching When Mirroring* if the destination port is not a dedicated port. For example, enable this option if you connect a laptop to the switch and you are running a packet sniffer along with the management GUI on the laptop.
8. Select *ERSPAN Manual* for the mode.
9. Enable *Strip VLAN Tags from Mirrored Traffic* if you want to remove VLAN tags from mirrored traffic.
10. Select *Add ERSPAN Headers* if you want to add the VLAN header to the encapsulated traffic.
11. In the Collector IP field, enter the IP address for the ERSPAN collector.
12. In the IPv4 Source Address field, enter the IPv4 source address in the ERSPAN IP header.
13. In the IPv4 TTL field, enter the IPv4 TTL value in the ERSPAN IP header.
14. In the IPv4 TOS field, enter the ToS value or enter the DSCP and ECN values in the ERSPAN IP header.
15. In the GRE Protocol field, enter the protocol value in the ERSPAN GRE header.
16. In the VLAN ID field, enter the VLAN identifier in the ERSPAN VLAN header.
    This field is available only if *Add ERSPAN Headers* is selected.
17. In the TPID field, enter the TPID for the encapsulating VLAN header.
    This field is available only if *Add ERSPAN Headers* is selected.
18. In the Priority field, enter the CoS bits in the ERSPAN VLAN header.
    This field is available only if *Add ERSPAN Headers* is selected.
19. In the CFI/DEI field, enter the CFI or DEI bit in the ERSPAN VLAN header.
    This field is available only if *Add ERSPAN Headers* is selected.
20. In the Source MAC Address field, enter the source MAC address in the ERSPAN Ethernet header.
    This field is available only if *Add ERSPAN Headers* is selected.
21. In the Destination MAC Address field, enter the MAC address of the next-hop or gateway on the path to the ERSPAN collector IP address.
    This field is available only if *Add ERSPAN Headers* is selected.
22. Select *Create* to create the mirror.

Using the CLI:

```
config switch mirror
   edit <mirror session name>
      set mode ERSPAN-manual
      set dst <interface>
      set encap-gre-protocol <hexadecimal_integer>
      set encap-ipv4-src IPv4_address>
      set encap-ipv4-tos <hexadecimal_integer>
      set encap-ipv4-ttl <0-255>
      set encap-mac-dst <MAC_address>
      set encap-mac-src <MAC_address>
      set encap-vlan {tagged | untagged}
         set encap-vlan-cfi <0-1>
         set encap-vlan-id <1-4094>
         set encap-vlan-priority <0-7>
         set encap-vlan-tpid <0x0001-0xfffe>
      set erspan-collector-ip <IPv4_address>
      set src-egress <interface_name>
      set src-ingress <interface_name>
      set strip-mirrored-traffic-tags {disable | enable}
      set switching-packet {enable | disable}
```

```
    set status active
  end
```

VLAN

FortiSwitch ports process tagged and untagged Ethernet frames. Untagged frames do not carry any VLAN information.

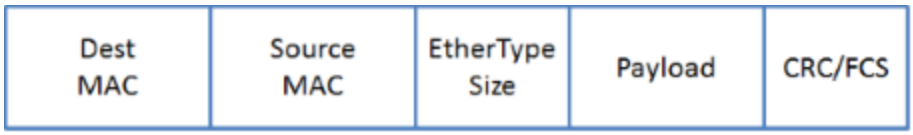

Tagged frames include an additional header (the 802.1Q header) after the Source MAC address. This header includes a VLAN ID. This allows the VLAN value to be transmitted between switches.

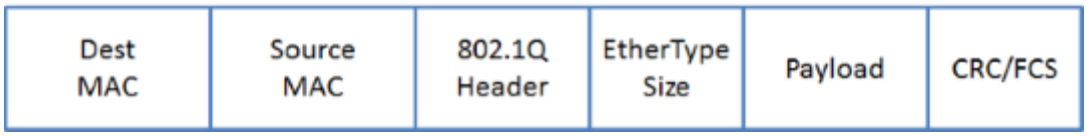

The FortiSwitch unit provides port parameters to configure and manage VLAN tagging.

This section covers the following topics:

- Native VLAN on page 314
- Allowed VLAN list on page 314
- Untagged VLAN list on page 315
- Frame processing on page 315
- Configuring VLANs on page 316
- Example 1 on page 316
- Example 2 on page 317
- VLAN stacking (QnQ) on page 318
- MAC/IP/protocol-based VLANs on page 324
- Private VLANs on page 327

Native VLAN

You can configure a native VLAN for each port. The native VLAN is like a default VLAN for untagged incoming frames. Outgoing frames for the native VLAN are sent as untagged frames.

The native VLAN is assigned to any untagged frame arriving at an ingress port.

At an egress port, if the frame tag matches the native VLAN, the frame is sent out without the VLAN header.

Allowed VLAN list

The allowed VLAN list for each port specifies the VLAN tag values for which the port can transmit or receive frames.

For a tagged frame arriving at an ingress port, the tag value must match a VLAN on the allowed VLAN list or the native VLAN.

At an egress port, the frame tag must match the native VLAN or a VLAN on the allowed VLAN list.

Untagged VLAN list

The untagged VLAN list on a port specifies the VLAN tag values for which the port will transmit frames without the VLAN tag. Any VLAN in the untagged VLAN list must also be a member of the allowed VLAN list.

The untagged VLAN list applies only to egress traffic on a port.

Frame processing

Ingress processing ensures that the port accepts only frames with allowed VLAN values (untagged frames are assigned the native VLAN, which is implicitly allowed). At this point, all frames are now tagged with a valid VLAN.

The frame is sent to each egress port that can send the frame (because the frame tag value matches the native VLAN or an Allowed VLAN on the port).

Ingress port

For an untagged frame:

- The frame is tagged with the native VLAN and allowed to proceed.
- The Allowed VLAN list is ignored.

For a tagged frame:

- The tag VLAN value must match an Allowed VLAN or the native VLAN.
- The frame retains the VLAN tag and is allowed to proceed.

To control what types of frames are accepted by the port:

```
config switch interface
   edit <interface>
      set discard-mode <all-tagged | all-untagged | none>
   end
```

Variable	Description
all-tagged	Tagged frames are discarded, and untagged frames can enter the switch.
all-untagged	Untagged frames are discarded, and tagged frames can enter the switch.
none	By default, all frames can enter the switch, and no frames are discarded.

Egress port

All frames that arrive at an egress port are tagged frames.

If the frame tag value is on the Allowed VLAN list, the frame is sent out with the existing tag.

If the frame tag value is the native VLAN or on the Untagged VLAN list, the tag is stripped, and then the frame is sent out.

Otherwise, the frame is dropped.

Configuring VLANs

Use the following steps to add VLANs to a physical port interface.

Using the GUI:

1. Go to *Switch* > *Interfaces*.
2. Select a port and then click *Edit*.
3. Give the VLAN an appropriate name.
4. In the Native VLAN field, enter the identifier for the native VLAN of the port.
5. In the Allowed VLANs field, enter one or more identifiers for the allowed VLANs for the port. Separate multiple numbers with commas without any space. For example, `2,4,8-10`.
6. In the Untagged VLANs field, enter one or more identifiers for the untagged VLANs for the port. Separate multiple numbers with commas without any space. For example, `2,4,8-10`.
7. Click *Update*.

Using the CLI:

```
config switch interface
   edit <port>
      set native-vlan <vlan>
      set allowed-vlans <vlan> [<vlan>] [<vlan> - <vlan>]
      set untagged-vlans <vlan> [<vlan>] [<vlan> - <vlan>]
   end
```

Example 1

The following example shows the flows for tagged and untagged frames.

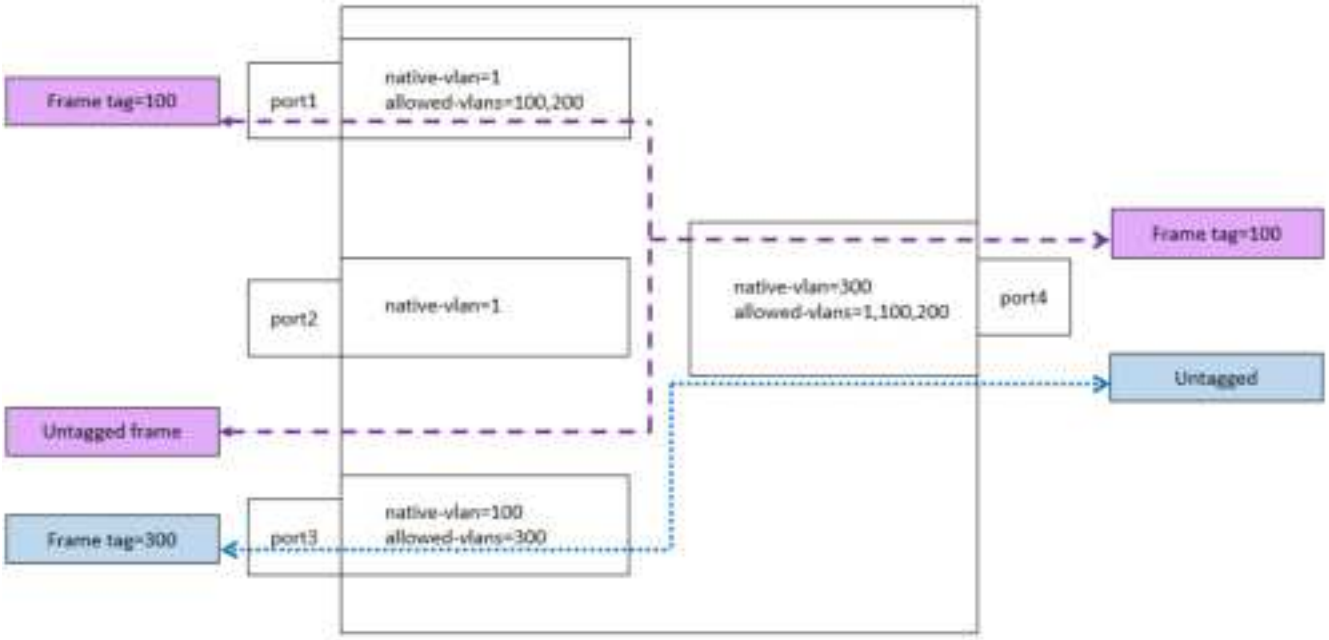

Purple (dashed) flow

An untagged frame arriving at port3 is assigned VLAN 100 (the native VLAN) and flows to all egress ports that will send VLAN 100 (port1 and port4).

A tagged frame (VLAN 100) arriving at port4 is allowed (VLAN 100 is allowed). The frame is sent out from port1 and port3. On port3, VLAN 100 is the native VLAN, so the frame is sent without a VLAN tag.

Blue (dotted) flow

An untagged frame arriving at port4 is assigned VLAN 300 (the native VLAN). Then it flows out all ports that will send VLAN 300 (port3).

A tagged frame (VLAN 300) arriving at port3 is allowed. The frame is sent to egress from port4. VLAN 300 is the native VLAN on port4, so the frame is sent without a VLAN tag.

Example 2

The following is an example of an invalid tagged VLAN.

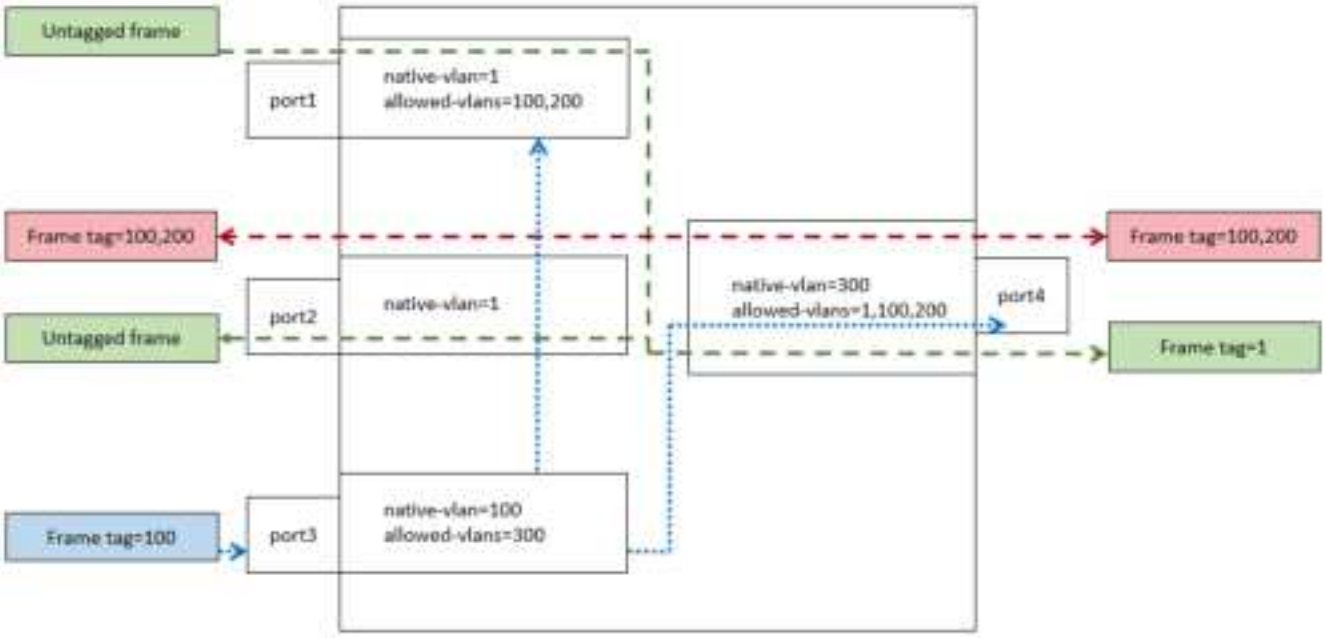

Green (dashed) flow

Between port1 and port2, frames are assigned to VLAN 1 at ingress, and then the tag is removed at egress.

Blue (dotted) flow

Incoming on port3, a tagged frame with VLAN value 100 is allowed because 100 is the port3 native VLAN (the hardware VLAN table accepts a tagged or untagged match to a valid VLAN).

The frame will be sent on port1 and port4 (with frame tag 100).

VLAN stacking (QnQ)

VLAN stacking allows you to have multiple VLAN headers in an Ethernet frame. The value of the EtherType field specifies where the VLAN header is placed in the Ethernet frame.

Use the VLAN TPID profile to specify the value of the EtherType field. The FortiSwitch unit supports a maximum of four VLAN TPID profiles, including the default (0x8100). The default VLAN TPID profile (0x8100) cannot be deleted or changed.

To see which models support this feature, refer to the FortiSwitch feature matrix.

The following features are not supported with VLAN stacking:

- DHCP relay
- DHCP snooping
- IGMP snooping
- IP source guard
- PVLAN
- STP

NOTE: Settings under `config qnq` are for customer VLANs (C-VLANs). Other settings such as `set allowed-vlans`, `set native-vlan`, and `set vlan-tpid` are for service-provider VLANs (S-VLANs).

Configuring VLAN stacking

Starting in FortiSwitchOS 7.4.3, you can use the CLI to specify the native customer VLAN (`native-c-vlan`) and allowed customer VLAN (`allowed-c-vlan`) when configuring VLAN stacking.

Using the GUI:

1. Go to *Switch* > *Interfaces*.
2. Select the interface that you want to configure and click *Edit*.
3. Select the *Enable QnQ* checkbox.
4. Select the *Drop Packets on VLAN Miss* checkbox if you want to drop the frame if the VLAN ID in the frame's tag is not defined in the VLAN-mapping configuration.
5. Select the *Remove Inner* checkbox if you want to remove the inner tag upon egress.
6. By default, the *STP QnQ Admin* checkbox is selected. You can clear the *STP QnQ Admin* checkbox if you are not using the options under it.
7. In the *Add Inner* field, enter the inner tag number for untagged frames upon ingress.
8. Click *Follow S-Tag* or *Follow C-Tag* to follow the priority of the S-tag (service tag) or C-tag (customer tag).
   NOTE: This command is not available on the 224D-FPOE, 248D, 424D, 424D-POE, 424D-FPOE, 448D, 448D-POE, 448D-FPOE, 224E, 224E-POE, 248E-POE and 248E-FPOE models.
9. Click + to add a VLAN mapping.
   a. In the *ID* field, enter a mapping entry identifier.
   b. In the *Description* field, enter a description of the mapping entry.
   c. In the *C-VLAN* field, enter a matching customer (inner) VLAN.
   d. In the *New C-VLAN* field, enter a new customer (inner) VLAN.
      NOTE: The VLAN must be in the port's allowed VLAN list.
10. Click *Update*.

Using the CLI (asterisks indicate the default setting):

```
config switch interface
   edit <interface_name>
      set vlan-tpid <default | string>
      config qnq
         set status {enable | *disable}
         set edge-type customer
         set vlan-mapping-miss-drop {enable | *disable}
         set add-inner <1-4095>
         set remove-inner {enable | *disable}
         set native-c-vlan <1-4094>
         set allowed-c-vlan <list_of_VLANs>
         set priority {follow-c-tag | *follow-s-tag}
         set s-tag-priority <0-7>
         config vlan-mapping
            edit <id>
               set description <string>
               set match-c-vlan <1-4094>
               set new-s-vlan <1-4094>
            end
```

```
    end
  next
end
```

Variable	Description	Default
<interface_name>	Enter the name of the interface.	No default
vlan-tpid <default | string>	Select which VLAN TPID profile to use. The default VLAN TPID profile has a value of 0x8100 and cannot be deleted or changed. This setting is only for service-provider VLANs (S-VLANs). **NOTE:** If you are not using the default VLAN TPID profile, you must have already defined the VLAN TPID profile with the `config switch vlan-tpid` command.	default
config qnq		
status {enable | *disable}	Enable this setting to use the VLAN stacking (QnQ) mode.	disable
edge-type customer	If the QnQ mode is enabled, the edge type is set to customer.	customer
vlan-mapping-miss-drop {enable | *disable}	If the QnQ mode is enabled, enable or disable whether a frame is dropped if the VLAN ID in the frame's tag is not defined in the vlan-mapping configuration. This option is available only when `allowed-c-vlan` has not been set.	disable
add-inner <1-4095>	If the QnQ mode is enabled, add the inner tag for untagged frames upon ingress.	No default
remove-inner {enable | *disable}	If the QnQ mode is enabled, enable or disable whether the inner tag is removed upon egress.	disable
native-c-vlan <1-4094>	Specify the native C VLAN (1-4094) for untagged packets. When you specify a value for `native-c-vlan`, FortiSwitchOS adds the native inner tag to untagged frames upon ingress and removes the native inner tag at egress.	No default
allowed-c-vlan <list_of_ VLANs>	Specify single VLANs or ranges of VLANs. Use a comma to separate values without any spaces. The `allowed-c-vlan` applies to both ingress and egress. You must use less than 4,096 characters to list the VLANs. This option is available only when `vlan-mapping-miss-drop` is disabled.	No default

Variable	Description	Default
priority {follow-c-tag | *follow-s-tag}	If the QnQ mode is enabled, select whether to follow the priority of the S-tag (service tag) or C-tag (customer tag). **NOTE:** This command is not available on the 224D-FPOE, 248D, 424D, 424D-POE, 424D-FPOE, 448D, 448D-POE, 448D-FPOE, 224E, 224E-POE, 248E-POE and 248E-FPOE models.	follow-s-tag
s-tag-priority <0-7>	If frames follow the priority of the S-tag (service tag), enter the priority value. This option is available only when the priority is set to `follow-s-tag`. **NOTE:** This command is not available on the 224D-FPOE, 248D, 424D, 424D-POE, 424D-FPOE, 448D, 448D-POE, 448D-FPOE, 224E, 224E-POE, 248E-POE and 248E-FPOE models.	0
<id>	Enter a mapping entry identifier.	No default
description <string>	Enter a description of the mapping entry.	No default
match-c-vlan <1-4094>	Enter a matching customer (inner) VLAN.	0
new-s-vlan <1-4094>	Enter a new service (outer) VLAN. **NOTE:** The VLAN must be in the port's allowed VLAN list. This option is only available after you set the value for `match-c-vlan`.	No default

Example configuration

```
config switch interface
   edit "port1"
      set native-vlan 3000
         config qnq
            set status enable
            set native-c-vlan 1
            set allowed-c-vlan 10,20
         end
      next
   end
```

Configuring VLAN mapping on an interface

Starting in FortiSwitchOS 7.0.2, partial VLAN mapping is supported by the FS-148F, FS-148F-POE, and FS-148F-FPOE models. Starting in FortiSwitchOS 7.0.3, partial VLAN mapping is supported by the FS-124F, FS-124F-POE, and FS-124F-FPOE models. Starting in FortiSwitchOS 7.2.0, partial VLAN mapping is supported by the FSR-112D-POE model. Starting in FortiSwitchOS 7.4.4, partial VLAN mapping is supported by the FSR-216F-POE model. Use the following syntax for partial VLAN mapping:

```
config switch interface
   edit <interface>
      config vlan-mapping
         edit <instance>
            set match-s-vlan <segment VLAN>
            set action replace
            set new-s-vlan <primary VLAN>
         next
      end
```

The FS-148F, FS-148F-POE, and FS-148F-FPOE models can map up to 1,024 physical or trunk ports. The FS-124F, FS-124F-POE,and FS-124F-FPOE models can map up to 512 physical or trunk ports. The FSR-112D-POE and FSR-216F-POE models can map up to 4,096 entries, but one VLAN can only be mapped to another VLAN; egress VLAN mapping can be enabled or disabled on individual ports.

Using the GUI:

1. Go to *Switch* > *Interfaces*.
2. Select the interface that you want to configure and click *Edit*.
3. In the *ID* field, enter a mapping entry identifier.
4. In the *Description* field, enter a description of the mapping entry.
5. In the *Direction* dropdown list, select *Ingress* or *Egress*.
6. If you selected *Ingress* for the direction:
   a. In the *Action* dropdown list, select *Add S-VLAN* or *Replace C-VLAN or S-VLAN*.
   b. In the *C-VLAN* field, enter a matching customer (inner) VLAN.
   c. In the *New S-VLAN* field, enter the new service (outer) VLAN.
   NOTE: The VLAN must be in the port's allowed VLAN list.
7. If you selected *Egress* for the direction:
   a. In the *Action* dropdown list, select *Delete S-VLAN* or *Replace C-VLAN or S-VLAN*.
   b. In the *S-VLAN* field, enter the matching service (outer) VLAN.
8. Click *Update*.

Using the CLI (asterisks indicate the default setting):

```
config switch interface
   edit <interface_name>
      set vlan-tpid <default | string>
      set vlan-mapping-miss-drop {enable | *disable}
      config vlan-mapping
         edit <id>
            set description <string>
```

```
            set direction ingress // ingress example
            set match-c-vlan <1-4094>
            set action {add | replace}
            set new-s-vlan <1-4094>
         next
         edit <id>
            set description <string>
            set direction egress // egress example
            set match-s-vlan <1-4094>
            set action {delete | replace}
            set new-s-vlan <1-4094>
         next
      end
   next
end
```

Variable	Description	Default
<interface_name>	Enter the name of the interface.	No default
vlan-tpid <default | string>	Select which VLAN TPID profile to use. The default VLAN TPID profile has a value of 0x8100 and cannot be deleted or changed. This setting is only for service-provider VLANs (S-VLANs). **NOTE:** If you are not using the default VLAN TPID profile, you must have already defined the VLAN TPID profile with the `config switch vlan-tpid` command.	default
vlan-mapping-miss-drop {enable | *disable}	Enable or disable whether a frame is dropped if the VLAN ID in the frame's tag is not defined in the vlan-mapping configuration.	disable
config vlan-mapping		
<id>	Enter an identifier for the VLAN mapping entry.	No default
description <string>	Enter a description of the VLAN mapping entry.	No default
direction {egress | ingress}	Select the ingress or egress direction.	No default
match-s-vlan <1-4094>	If the direction is set to egress, enter the service (outer) VLAN to match.	0
match-c-vlan <1-4094>	If the direction is set to ingress, enter the customer (inner) VLAN to match.	0
action {add | delete | replace}	Select what happens when the frame is matched: • `add`—When the frame is matched, add the service VLAN. You cannot set the `action` to `add` for the egress direction. • `delete`—When the frame is matched, delete the service VLAN. You cannot set the `action` to	No default

Variable	Description	Default
	`delete` for the ingress direction. • - `replace`—When the frame is matched, replace the customer VLAN or service VLAN. This option is only available after you set a value for `match-c-vlan` or `match-s-vlan`.	
new-s-vlan <1-4094>	Set the new service (outer) VLAN. This option is only available after you set the action to `add` or `replace` for the ingress direction or after you set the `action` to `replace` for the egress direction.	No default

Configuring the VLAN TPID profile

Use the CLI to specify the value of the EtherType field in the VLAN TPID profile:

```
config switch vlan-tpid
   edit <VLAN_TPID_profile_name>
      set ether-type <0x0001-0xfffe>
   next
end
```

Variable	Description	Default
<VLAN_TPID_profile_name>	Enter a name for the VLAN TPID profile name.	No default
ether-type <0x0001-0xfffe>	Enter a hexadecimal value for the EtherType field.	0x8100

Checking the VLAN stacking configuration

Use the CLI to check that VLAN stacking is configured correctly:

```
diagnose switch qnq dtag-cfg
```

MAC/IP/protocol-based VLANs

The FortiSwitch unit assigns VLANs to packets based on the incoming port or the VLAN tag in the packet. The MAC/IP/protocol-based VLAN feature enables the assignment of VLANs based on specific fields in an ingress packet (MAC address, IP address, or layer-2 protocol).

Overview

When a MAC/IP/protocol-based VLAN is assigned to a port, the default behavior is for egress packets with that VLAN value to include the VLAN tag. Use the `set untagged-vlans <vlan>` configuration command to remove the VLAN tag from egress packets. For an example of the command, see the Example configuration on page 326.

The MAC/IP/protocol-based VLAN feature assigns the VLAN based on MAC address, IP address, or layer-2 protocol.

MAC based

In MAC-based VLAN assignment, the FortiSwitch unit associates a VLAN with each packet based on the originating MAC address.

IP based

In IP-based VLAN assignment, the FortiSwitch unit associates a VLAN with each packet based on the originating IP address or IP subnet. IPv4 is supported with prefix masks from 1 to 32. IPv6 is also supported, depending on hardware availability, with prefix lengths from 1 to 64.

Protocol based

In protocol-based VLAN assignment, the FortiSwitch unit associates a VLAN with each packet based on the Ethernet protocol value and the frame type (ethernet2, 802.3d/SNAP, LLC).

Configuring MAC/IP/protocol-based VLANs

Note the following prerequisites:

- The VLAN must be created in the FortiSwitch unit
- The VLAN needs to be allowed on the ingress port

Using the GUI:

1. Go to *Switch* > *VLAN*.
2. Select *Add VLAN* for a new VLAN or select *Edit* for an existing VLAN.
3. To configure a MAC-based VLAN:
   a. Select *Add* under Members by MAC Address.
   b. Enter a description and the MAC address.
4. To configure an IP-based VLAN:
   a. Select *Add* under Members by IP Address.
   b. Enter a description and the IP address.
5. Select *Add* or *Update* to save the settings.

Using the CLI:

```
config switch vlan
   edit <vlan-id>
      config member-by-mac
         edit <id>
            set mac xx:xx:xx:xx:xx:xx
            set description <128 byte string>
         next
      end
      config member-by-ipv4
         edit <id>
            set address a.b.c.d/e #subnet mask must 1-32
            set description <128 byte string>
         next
      end
```

```
      config member-by-ipv6
         edit <id>
            set prefix xx:xx:xx:xx::/prefix #prefix must 1-64
            set description <128 byte string>
         next
      end
      config member-by-proto
         edit <id>
            set frametypes ethernet2 802.3d llc #default is all
            set protocol 0xXXXX
         next
      end
   next
end
```

NOTE: There are hardware limits regarding how many MAC/IP/protocol-based VLANs that you can configure. If you try to add entries beyond the limit, the CLI will reject the configuration:

- Editing an existing VLAN—when you enter `next` or `end` on the `config member-by` command
- Adding a new VLAN— when you enter `next` or `end` on the `edit vlan` command
- When VLANS are defined by `config member-by-ipv4` or `config member-by-ipv6` on some FortiSwitch platforms (2xx and higher), matching ARP traffic is included in the assigned VLANs. For example, if the ARP target IP address or the ARP sender IP address match the member-by-ipv4 or member-by-ipv6 IP address, those ARP packets are included in the assigned VLANs.

Example configuration

The following example shows a CLI configuration for MAC-based VLAN where a VOIP phone and a PC share the same switch port.

In this example, a unique VLAN is assigned to the voice traffic, and the PC traffic is on the default VLAN for the port.

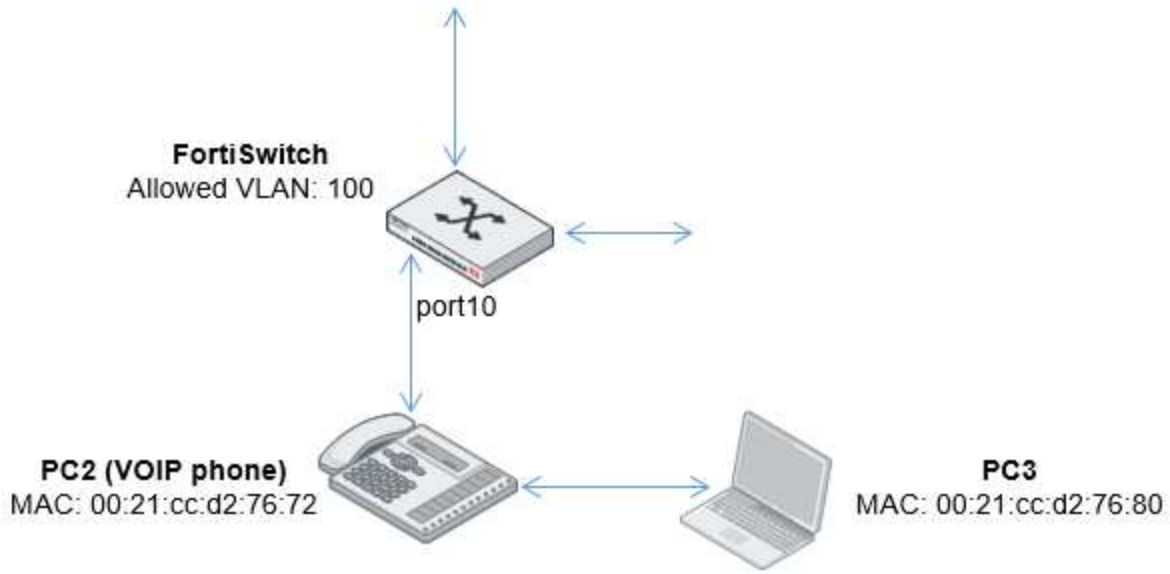

1. The FortiSwitch Port 10 is connected to PC2 (a VOIP phone), with MAC address 00:21:cc:d2:76:72.
2. The phone also sends traffic from PC3 (MAC= 00:21:cc:d2:76:80).

3. Assign the PC3 traffic to the default VLAN (1) on port 10.
4. Assign the voice traffic to VLAN 100.

Configure the voice VLAN

```
config switch vlan
   edit 100
      config member-by-mac
         edit 1
            set description "pc2"
            set mac 00:21:cc:d2:76:72
         next
      end
   end
end
```

Configure switch port 10

```
config switch interface
   edit "port10"
      # allow vlan=100 on this port
      # treat this as untagged on egress
      set allowed-vlans 100
      set untagged-vlans 100
      set snmp-index 10
   end
end
```

Checking the configuration

To view the MAC-based VLAN assignments, use the following command:

```
diagnose switch vlan assignment mac list sorted-by-mac

    00:21:cc:d2:76:72   VLAN: 100 Installed: yes
    Source: Configuration (entry 1)
    Description: pc2
```

Private VLANs

A private VLAN (PVLAN) divides the original VLAN (termed the primary VLAN) into sub-VLANs (secondary VLANs), while retaining the existing IP subnet and layer-3 configuration. Unlike a regular VLAN, which is a single broadcast domain, a PVLAN partitions one broadcast domain into multiple smaller broadcast subdomains.

After a PVLAN VLAN is configured, the primary VLAN forwards frames downstream to all secondary VLANs.

There are two main types of secondary VLANs:

- **Isolated**: Any switch ports associated with an isolated VLAN can reach the primary VLAN, but not any other secondary VLAN. In addition, hosts associated with the same isolated VLAN cannot reach each other. Only one isolated VLAN is allowed in one PVLAN domain.
- **Community**: Any switch ports associated with a common community VLAN can communicate with each other and with the primary VLAN but not with any other secondary VLAN. You might have multiple distinct community VLANs within one PVLAN domain.

There are mainly two types of ports in a PVLAN: promiscuous (P-Port) and host.

- **Promiscuous Port (P-Port)**: The switch port connects to a router, firewall, or other common gateway device. This port can communicate with anything else connected to the primary or any secondary VLAN. In other words, it is a type of a port that is allowed to send and receive frames from any other port on the VLAN.
- **Host Ports** further divides into two types – isolated port (I-Port) and community port (C-port).
- **Isolated Port (I-Port)**: Connects to the regular host that resides on isolated VLAN. This port communicates only with P-Ports.
- **Community Port (C-Port)**: Connects to the regular host that resides on community VLAN. This port communicates with P-Ports and ports on the same community VLAN.

Creating and enabling a PVLAN

Using the GUI:

1. Go to *Switch* > *VLAN*.
2. Select *Add VLAN* to create a new PVLAN.
3. Enter the VLAN identifier.
4. Enter a description for the new PVLAN.
5. Select *Enabled* to enable the new Private VLAN.
6. Enter a single VLAN identifier for the isolated subVLAN.
7. If needed, enter one VLAN identifier or multiple VLAN identifiers for a common community subVLAN.
8. To save your changes, select *Add* at the bottom of the page.

Configuring the PVLAN ports

Using the GUI:

1. Go to *Switch* > *Interfaces*.
2. Select the port to configure.
3. Click *Edit*.
4. Select if the Private VLAN port is a promiscuous port or part of a sub-VLAN.
5. For a promiscuous port, select the primary VLAN identifier.
6. For a port that is part of a sub-VLAN, select the primary VLAN identifier and the sub-VLAN identifier.
7. Click *Update*.

Private VLAN example

1. Enable a PVLAN:

```
config switch vlan
   edit 1000
      set private-vlan enable
      set isolated-vlan 101
      set community-vlans 200-210
   end
end
```

2. Configure the PVLAN ports:

```
config switch interface
   edit "port2"
      set private-vlan promiscuous
      set primary-vlan 1000
   next
   edit "port3"
      set private-vlan sub-vlan
      set primary-vlan 1000
      set sub-vlan 200
   next
   edit "port7"
      set private-vlan sub-vlan
      set primary-vlan 1000
      set sub-vlan 101
   next
   edit "port19"
      set private-vlan promiscuous
      set primary-vlan 1000
   next
   edit "port20"
      set private-vlan sub-vlan
      set primary-vlan 1000
      set sub-vlan 101
   next
   edit "port21"
      set private-vlan sub-vlan
      set primary-vlan 1000
      set sub-vlan 101
   end
end
```

Virtual wires

Some testing scenarios might require two ports to be wired 'back-to-back'. Instead of using a physical cable, you can configure a virtual wire between two ports. The virtual wire forwards traffic from one port to the other port with minimal filtering or modification of the packets.

Notes:

- ACL mirroring is not supported.
- You can select ports that are already ingress and egress mirror sources.

Using the GUI:

1. Go to *Switch* > *Virtual Wires*.
2. Select *Add Virtual Wire* to create a new virtual wire.
3. Enter a name and select the ports for first member and second member.
4. Select *Add* to save the changes.

Using the CLI:

Use the following commands to configure a virtual wire:

```
config switch virtual-wire
   edit <virtual-wire-name>
      set first-member <port-name>
      set second-member <port-name>
      set vlan <vlan-id>
   next
end
```

Virtual wire ports set a special Tag Protocol Identifier (TPID) in the VLAN header. The default value is 0xdee5, a value that real network traffic never uses.

Use the following commands to configure a value for the TPID:

```
config switch global
   set virtual-wire-tpid <hex value from 0x0001 to 0xFFFE>
end
```

Use the following command to display the virtual wire configuration:

```
diagnose switch physical-ports virtual-wire list


  port1(1) to port2(2) TPID: 0xdee5 VLAN: 4011
  port3(3) to port4(4) TPID: 0xdee5 VLAN: 4011
  port5(5) to port25(25) TPID: 0xdee5 VLAN: 4011
  port7(7) to port8(8) TPID: 0xdee5 VLAN: 4011
```

NOTE:

- Ports have ingress and egress VLAN filtering disabled. All traffic (including VLAN headers) is passed unchanged to the peer. All egress traffic is untagged.
- Ports have L2 learning disabled.
- Ports have their egress limited to their peer and do no allow egress from any other ports.
- The system uses TCAM to force forwarding from a port to its peer.
- The TCAM prevents any copy-to-cpu or packet drops.

Storm control

Storm control protects a LAN from disruption by traffic storms, which stem from mistakes in network configuration or denial-of-service attacks. A traffic storm, which can consist of broadcast, multicast, or unicast traffic, creates excessive traffic on the LAN and degrades network performance.

By default, storm control is disabled on a FortiSwitch unit. When enabled, it measures the data rate (in packets-per-second) for unknown unicast, unknown multicast, and broadcast traffic. You can enable and disable storm control for each of these traffic types individually. If the traffic rate for any of the types exceeds the configured threshold, the FortiSwitch unit drops the excess traffic.

By default, storm control configuration is global. Starting in FortiSwitchOS 6.2.0, you can configure storm control on a port level.

Starting in FortSwitchOS 6.4.3, you can configure the maximum burst size allowed by storm control. Select the burst-size level from 0 to 4 with the highest number for the highest maximum burst size allowed. The maximum number of packets or bytes allowed for each burst-size level depends on the switch model.

NOTE: The burst-size level cannot be controlled on a port level for the FS-108E, FS-108E-POE, FS-108-FPOE, FS-124E, FS-124E-POE, and FS-124E-FPOE models.

This section covers the following topics:

- Configuring system-wide storm control on page 331
- Configuring port-level storm control on page 331
- Monitoring storm control on page 332
- Checking the storm-control configuration on page 333

Configuring system-wide storm control

If you set the rate to zero, the system drops all packets (for the enabled traffic types).

Using the GUI:

1. Go to *Switch* > *Storm Control*.
2. Select *Restrict Traffic*.
3. Select *Broadcast*, *Unknown Unicast*, and *Unknown Multicast* as required.
4. Select the action to take, either *Drop Packets* or *Rate Limit*.
5. If you selected *Rate Limit*, enter the number of packets per second.
6. Select a *Fixed Level* or *Custom* burst size.
7. If you selected *Custom* for the burst size, enter the burst-size level from 1-4.
8. Click *Update* to save the changes.

Using the CLI:

```
config switch storm-control
   set broadcast {enable | disable}
   set burst-size-level <0-4>
   set rate [0 | 2-10000000]
   set unknown-unicast {enable | disable}
   set unknown-mcast {enable | disable}
end
```

Configuring port-level storm control

Using the GUI:

1. Go to *Switch* > *Physical Ports*.
2. Select a port and then select *Edit*.
3. In the Storm Control area, select *Configure Manually*.
4. Select one or more of the packet types: *Broadcast*, *Unknown Multicast*, and *Unknown Unicast*.
5. Select the action to take, either *Drop Packets* or *Rate Limit*.
6. If you selected *Rate Limit*, enter the number of packets per second.

7. Select a *Fixed Level* or *Custom* burst size.
8. If you selected *Custom* for the burst size, enter the burst-size level from 1-4.
9. Click *Update* to save the changes.

Using the CLI:

```
config switch physical-port
   edit <port_name>
      set storm-control-mode override
         config storm-control
            set broadcast {enable | disable}
            set burst-size-level <0-4>
            set rate [0 | 2-10000000]
            set unknown-multicast {enable | disable}
            set unknown-unicast {enable | disable}
         end
      end
```

Monitoring storm control

Starting in FortiSwitchOS 7.4.3, you can use the CLI to monitor the rate at which packets are dropped when storm control is enabled and generate a log message when a specified threshold is exceeded.

To monitor storm control:

```
config switch global
   set storm-control-monitor {enable | disable}
   set storm-control-high-rate <0-65536>
   set storm-control-rate-filter <0-100>
end
```

storm-control-monitor {enable | disable}	Enable or disable storm-control monitoring.	disable
storm-control-high-rate <0-65536>	When this rate (in dropped packets per second) is exceeded, a log message is generated. This command is only available when `storm-control-monitor` is enabled.	300
storm-control-rate-filter <0-100>	Set the percentage for how sensitive storm-control monitoring is to changes in the `storm-control-high-rate`. Higher percentages mean that the storm-control monitoring is more sensitive to changes in the `storm-control-high-rate`. This command is only available when `storm-control-monitor` is enabled.	20

In the following example, storm-control monitoring is enabled, and a log message is generated when the rate (100 dropped packets per second) is exceeded. The sensitivity filter is set at 50 percent.

```
config switch global
   set storm-control-monitor enable
   set storm-control-high-rate 100
   set storm-control-rate-filter 50
end
```

Checking the storm-control configuration

Use the following command to display the system-wide storm-control configuration:

```
get switch storm-control
```

Use the following command to display whether storm-control monitoring is enabled, the drop threshold, the drop filter percentage, and the port rate:

```
diagnose switch storm-control
```

MAC entries

This section covers the following topics:

- Persistent (sticky) MAC addresses on page 333
- Static MAC addresses on page 334
- Network monitor on page 346

Persistent (sticky) MAC addresses

You can make dynamically learned MAC addresses persistent when the status of a FortiSwitch port changes (goes down or up). By default, MAC addresses are not persistent.

NOTE:

- If you move a device within your network that has a sticky MAC address entry on the switch, remove the sticky MAC address entry from the interface. If you move the device and do not clear the sticky MAC address from the original port it was learned on, the new port will not learn the MAC address of the device.
- You cannot use persistent MAC addresses with 802.1X authentication.

You can use static/sticky MAC addresses or 802.1X authentication but not both on the same port at the same time. If you do need to use both, you must ensure that the MAC addresses/devices authorized by 802.1X authentication are not included in the static-mac table.

Using the GUI:

1. Go to *Switch* > *MAC Entries*.
2. Select *Add MAC Entry* to create a new item.
3. Select an interface and enter a value for *MAC Address* and *VLAN*.

4. Select *Sticky*.
5. Select *Add* to create the MAC entry.

To delete the persistent MAC addresses instead of saving them in the FortiSwitch configuration file:

1. Go to *Switch* > *Monitor* > *Forwarding Table*.
2. In the *Unsaved sticky MACs on* field, select an interface or select *All*.
3. Select *Delete*.

Using the CLI:

Use the following command to configure the persistence of MAC addresses on an interface:

```
config switch interface
   edit <port>
      set sticky-mac <enable | disable>
   next
end
```

You can also save persistent MAC addresses to the FortiSwitch configuration file so that they are automatically loaded when the FortiSwitch unit is rebooted. By default, persistent entries are lost when a FortiSwitch unit is rebooted. Use the following command to save persistent MAC addresses for a specific interface or all interfaces:

```
execute sticky-mac save {all | interface <interface_name>}
```

Use the following command to delete the persistent MAC addresses instead of saving them in the FortiSwitch configuration file:

```
execute sticky-mac delete-unsaved {all | interface <interface_name>}
```

Static MAC addresses

You can configure one or more static MAC addresses on an interface.

Starting in FortiSwitchOS 7.2.0, you can configure in the CLI whether packets with specific source static MAC address are allowed or dropped. By default, they are allowed.

You can use static/sticky MAC addresses or 802.1X authentication but not both on the same port at the same time. If you do need to use both, you must ensure that the MAC addresses/devices authorized by 802.1X authentication are not included in the static-mac table.

Using the GUI:

1. Go to *Switch* > *MAC Entries*.
2. Select *Add MAC Entry* to create a new item.
3. Select an interface and enter a value for *MAC Address* and *VLAN*.
4. Select the *Sticky* checkbox if you want the MAC address to be persistent, even when the status of a FortiSwitch port changes (goes down or up).
5. Select *Add* to create the MAC entry.

Using the CLI:

```
config switch static-mac
   edit <sequence_number>
      set action {allow | drop}
      set description <optional_string>
      set interface <interface_name>
      set mac <static_MAC_address>
      set type {sticky | static}
      set vlan-id <VLAN_ID>
   end
```

For example:

```
config switch static-mac
   edit 1
      set action drop
      set description "first static MAC address"
      set interface port10
      set mac d6:dd:25:be:2c:43
      set type static
      set vlan-id 10
   end
```

IP-MAC binding

Use IP-MAC binding to prevent ARP spoofing.

The port accepts a packet only if the source IP address and source MAC address in the packet match an entry in the IP-MAC binding table.

You can enable or disable IP-MAC binding for the whole switch, and you can override this global setting for each port.

Configuring IP-MAC binding

Use the following steps to configure IP-MAC binding:

1. Enable the IP-MAC binding global setting.
2. Create the IP-MAC bindings. You can activate each binding individually.
3. Set each port to follow the global setting. You can also override the global setting for individual ports by enabling or disabling IP-MAC binding for the port.

Using the GUI:

Create the IP-MAC binding:

1. Go to *Switch* > *IP MAC Binding*.
2. Select *Add IP MAC Binding* to create a new binding.
3. Select *Status*.
4. Enter the IP address and netmask.

5. Enter the MAC address.
6. Select *Add*.

Using the CLI:

```
config switch global
   set ip-mac-binding [enable| disable]

config switch ip-mac-binding
   edit 1
      set ip <IP address and network mask>
      set mac <MAC address>
      set status (enable| disable)
   next
end
config switch interface
   edit <port>
      set ip-mac-binding (enable| disable | global)
   edit <trunk name>
      set ip-mac-binding (enable| disable | global)
```

Notes:

- For a switch port, the default IP-MAC binding value is disabled.
- When you configure a trunk, the trunk follows the global value by default. You can also explicitly enable or disable IP-MAC binding for a trunk, as shown in the CLI configuration.
- When you add member ports to the trunk, all ports take on the trunk setting. If you later remove a port from the trunk group, the port is reset to the default value (disabled).
- No duplicate entries are allowed in the mapping table.
- Rules are disabled by default. You need to explicitly enable each rule.
- The mapping table holds up to 1024 rules.

Viewing IP-MAC binding configuration

Display the status of IP-MAC binding using the following command:

```
show switch ip-mac-binding <entry number>
```

QoS

Quality of service (QoS) provides the ability to set particular priorities for different applications, users, or data flows.

QoS involves the following elements:

- **Classification** is the process of determining the priority of a packet. This can be as simple as trusting the QoS markings in the packet header when it is received and so accept the packet. Alternatively, it can hinge on criteria (such as incoming port, VLAN, or service) that are defined by the network administrator.
- **Marking** involves setting bits in the packet header to indicate the priority of this packet.

- **Queuing** involves defining priority queues to ensure that packets marked as high priority take precedence over those marked as lower priority. If network congestion becomes so severe that packet drops are inevitable, the queuing process will also select the packets to drop.

The FortiSwitch unit supports the following QoS configuration capabilities:

- Mapping the IEEE 802.1p and layer-3 QoS values (Differentiated Services and IP Precedence) to an outbound QoS queue number.
- Providing eight egress queues on each port.
- Policing the maximum data rate of egress traffic on the interface.

NOTE: There are some differences in QoS configuration on the FS-108E, FS-108E-POE, FS-108E-FPOE, FS-124E, FS-124E-POE, FS-124E-FPOE, FS-148E, FS-148E-POE, FS-124F, FS-124F-POE, FS-124F-FPOE, FS-148F, FS-148F-POE, and FS-148F-FPOE models:

- You can configure only one dot1p-map per switch.
- You can configure only one ip-dscp-map per switch.
- You cannot set `min-rate`, `min-rate-percent`, `drop-policy`, or `wred-slope` under the `config switch qos qos-policy` command.
- Under the `config switch qos qos-policy` command, the switch rounds the `max-rate` value to the nearest multiple of 16 internally. If the rounding result is 0, `max-rate` is disabled internally.
- You cannot configure priority tagging on outgoing frames (`egress-pri-tagging`) under the `config switch qos dot1p-map` command.
- You can configure only one QoS drop policy per switch. You can configure the QoS drop policy under the `config switch global` command. You can specify random early detection (RED) with the `set qos-drop-policy random-early-detection` command on the FS-108E, FS-124E, FS-148E, FS-124F, and FS-148F models.
- You can set the QoS RED/WRED drop probability (`qos-red-probability`) under the `config switch global` command. The FS-108E, FS-108E-POE, FS-108E-FPOE, FS-124E, FS-124E-POE, and FS-124E-FPOE models support 0-100 percent. The FS-148E, FS-148E-POE, FS-148E-FPOE, FS-124F, FS-124F-POE, FS-124F-FPOE, FS-148F, FS-148F-POE, and FS-148F-FPOE models support 0-25 percent.
- Adaptive or active RED (ARED) and robust RED (RRED) are not supported.

This section covers the following topics:

- Classification on page 337
- Marking on page 338
- Queuing on page 338
- Determining the egress queue on page 339
- Configuring FortiSwitch QoS on page 340
- Checking the QoS statistics on page 346
- Resetting and restoring QoS counters on page 346

Classification

The IEEE 802.1p standard defines a class of service (CoS) value (ranging from 0-7) that is included in the Ethernet frame. The Internet Protocol defines the layer-3 QoS values that are carried in the IP packet (Differentiated Services, IP Precedence). The FortiSwitch unit provides configurable mappings from CoS or IP-DSCP values to egress queue values.

Fortinet recommends that you do not enable trust for both Dot1p and DSCP at the same time on the same interface. If you do want to trust both Dot1p and IP-DSCP, the switch uses the latter value (DSCP) to determine the queue. The

switch will use the Dot1p value and mapping only if the packet contains no DSCP value. For details, refer to Determining the egress queue on page 339.

Marking

FortiSwitchOS supports two ways to indicate the priority of outgoing packets:

- **CoS marking**: The priority is set with the CoS value of the 802.1Q tag. The range of CoS values is 0-7.
- **Differential service code point (DSCP) marking**: The priority is set with the DSCP value in the IP header. The range of DSCP values is 0-63.

You can use one of these methods or both methods.

Whether the CoS or DSCP values of inbound packets are remarked is subject to the classification by ACL rules for the ingress interfaces. When CoS or DSCP marking take place, the outbound queuing is not impacted, meaning it is still based on trust maps and the original CoS or DSCP values, as described in Determining the egress queue on page 339.

The following example shows how to use the CLI to configure an ACL policy to mark the CoS and DSCP values of inbound packets to 4 and 48 on port1 when their CoS values are 2:

```
config switch acl ingress
   edit 10
      config action
         set count enable
         set remark-cos 4
         set remark-dscp 48
      end
      config classifier
         set cos 2
      end
      set ingress-interface "port1"
      set status active
   next
end
```

Queuing

Queuing determines how queued packets on an egress port are served. Each egress port supports eight queues, and three scheduling modes are available:

- **Strict Scheduling**: The queues are served in descending order (of queue number), so higher number queues receive higher priority. Queue7 has the highest priority, and queue0 has the lowest priority. The purpose of the strict scheduling mode is to provide lower latency service to higher classes of traffic. However, if the interface experiences congestion, the lower priority traffic could be starved.
- **Simple Round Robin (RR)**: In round robin mode, the scheduler visits each backlogged queue, servicing a single packet from each queue before moving on to the next one. The purpose of round robin scheduling is to provide fair access to the egress port bandwidth.
- **Weighted Round Robin (WRR)**: Each of the eight egress queues is assigned a weight value ranging from 0 to 63. The purpose of weighted round robin scheduling is to provide prioritized access to the egress port bandwidth, such that queues with higher weight get more of the bandwidth, but lower priority traffic is not starved.

  When you set the `schedule` to `weighted` and set the queue `weight` to 0, then the queue will follow strict scheduling, which takes priority over all the queues with nonzero weights. If there are multiple queues with weight 0,

then the higher queue has the higher priority (as per the Scrict Scheduling definition). By using weight 0, you can have a combination of Strict Scheduling and WRR scheduling in the same QoS policy.

A drop policy determines what happens when a queue is full or exceeds a minimum threshold. Depending on your switch model, you can select from one of two drop policies:

- The **tail-drop** drop policy is the default. When a queue is full, additional incoming packets are dropped until there is space available in the queue. To see which models support this feature, refer to the FortiSwitch feature matrix.
- The **random early detection (RED)**. When the queue size exceeds the minimum threshold, packets are dropped at a constant rate until the queue is full. Using the RED drop policy helps improve the throughput during network congestion. To see which models support this feature, refer to the FortiSwitch feature matrix.
- The **weighted random early detection (WRED)** drop policy is an advanced version of RED. When the queue size exceeds the threshold, the WRED slope controls the rate at which packets are dropped until the queue is full. The drop rate increases when the queue buffer usage increases. If you select `weighted-random-early-detection` in the CLI, you can enable explicit congestion notification (ECN) marking to indicate that congestion is occurring without just dropping packets. To see which models support this feature, refer to the FortiSwitch feature matrix.

Determining the egress queue

To determine the egress queue value for the packet, the FortiSwitch unit uses the configured trust values (and mappings) on the port and the QoS/CoS fields in the packet.

Packets with DSCP and CoS values

If the port is set to trust DSCP and Dot1p, the switch uses the DSCP value to find the queue assignment in the DSCP map for the port.

If the port is set to trust Dot1p and **not** to trust DSCP, the switch uses the packet's CoS value to look up the queue assignment in the Dot1p map for the port.

If the port is **not** set to trust Dot1p, the switch uses the default queue 0.

Packets with a CoS value but no DSCP value

The switch ignores the trust DSCP value.

- If the port is set to trust Dot1p, the switch uses the packet's CoS value to look up the queue assignment in the Dot1p map for the port.
- If the port is **not** set to trust Dot1p, the switch uses the default queue 0.

Packets with a DSCP value but no CoS value

If the port is set to trust DSCP, the switch uses the packet's DSCP value to look up the queue assignment in the DSCP map for the port.

If the port is set to trust Dot1p but **not** to trust DSCP, the switch uses the default CoS value of the port to look up the queue assignment in the Dot1p map for the port.

If the port is **not** set to trust Dot1p, the switch uses the default queue 0.

Configuring FortiSwitch QoS

FortiSwitch uses “queue-7” for network control and critical management traffic. To avoid affecting critical network control and management traffic, do not oversubscribe queue-7 or avoid using queue-7 for data traffic when configuring QoS.

This section provides procedures for the following configuration tasks:

- Configure an 802.1p map on page 340
- Configure a DSCP map on page 341
- Configure the QoS egress policy on page 342
- Configure the egress drop mode on page 343
- Configure the switch ports on page 343
- Configure QoS on trunks on page 344
- Configure QoS on VLANs on page 345
- Configure CoS and DSCP markings on page 345

Configure an 802.1p map

Using the GUI:

1. Go to *Switch* > *QoS* > *802.1p*.
2. Select *Add Map*.
3. Enter the name of your 802.1p map.
4. Enter a description of your 802.1p map.
5. Select the queue number for each priority.
6. Select *Add Map*.

Values that are not explicitly included in the map will follow the default mapping, which maps each priority (0-7) to queue 0. If an incoming packet contains no CoS value, the switch assigns a CoS value of zero.

Using the CLI:

You can configure an 802.1p map, which defines a mapping between IEEE 802.1p CoS values (from incoming packets on a trusted interface) and the egress queue values.

If you want to enable priority tagging on outgoing frames, enable the `egress-pri-tagging` option. This option is disabled by default.

NOTE: “Priority tagging” refers to adding a VLAN tag to untagged traffic with with VLAN 0 and a valid priority value. If the port is configured to transmit packets with a valid VLAN, priority tagging is not applicable.

```
config switch qos dot1p-map
   edit <dot1p map name>
      set description <text>
      set [priority-0|priority-1|priority-2|....priority-7] <queue number>
      set egress-pri-tagging {disable | enable}
   next
end
```

For example:

```
config switch qos dot1p-map
   edit "test1"
      set priority-0 queue-2
      set priority-1 queue-0
      set priority-2 queue-1
      set priority-3 queue-3
      set priority-4 queue-4
      set priority-5 queue-5
      set priority-6 queue-6
      set priority-7 queue-7
      set egress-pri-tagging enable
   next
end
```

Values that are not explicitly included in the map will follow the default mapping, which maps each priority (0-7) to queue 0. If an incoming packet contains no CoS value, the switch assigns a CoS value of zero.

Use the `set default-cos` command to set a different default CoS value, ranging from 0 to 7:

```
config switch interface
   edit port1
      set default-cos <0-7>
```

NOTE: The `set default-cos` command is not available on the following FortiSwitch models: 224D-FPOE, 248D, 424D, 424D-POE, 424D-FPOE, 448D, 448D-POE, 448D-FPOE, 224E, 224E-POE, 248E-POE, and 248E-FPOE.

Configure a DSCP map

A DSCP map defines a mapping between IP precedence or DSCP values and the egress queue values.

Using the GUI:

1. Go to *Switch* > *QoS* > *IP/DSCP*.
2. Select *Add Map*.
3. Enter the name of your DCSP map.
4. Enter a description of your DCSP map.
5. Select which queue to configure.
6. Select the differentiated services to use.
7. Select the IP precedence to use.
8. Enter the raw values to use.
9. Select *Add Map*.

Using the CLI:

```
config switch qos ip-dscp-map
   edit <ip-dscp map name>
      set description <text>
      config map
      edit <entry-name1>
         set diffserv [ [ AF11 | AF12 | AF13 | AF21 | AF22 | AF23 | AF31 | AF32 | AF33 | AF41
               | AF42 | AF43 | CS0 | CS1 | CS2 | CS3 | CS4 | CS5 | CS6 | CS7 | EF ]
```

```
        set ip-precedence [ Network Control | Internetwork Control | Critic/ECP | Flash
             Override | Flash, Immediate | Priority | Routine ]
        set value <dscp raw value>
        set cos-queue <queue number>
      next
   end
end
```

The following example defines a mapping for two of the DSCP values:

```
config switch qos ip-dscp-map
   edit "m1"
      config map
         edit "e1"
            set cos-queue 0
            set ip-precedence Immediate
         next
         edit "e2"
            set cos-queue 3
            set value 13
         next
      end
   next
end
```

Configure the QoS egress policy

In a QoS egress policy, you set the scheduling mode (Strict, Round Robin, or Weighted Round Robin) for the policy, and configure one or more CoS queues.

The QoS egress policy includes the following settings:

- min-rate (minimum rate in kbps) or min-rate-percent (minimum percentage)
- max-rate (maximum rate in kbps) or max-rate-percent (maximum percentage)
- drop policy: tail drop, RED, or WRED
- weight value (applicable if the policy schedule is weighted)

Using the GUI:

1. Go to *Switch* > *QoS* > *Egress Policy*.
2. Select *Add Policy*.
3. Enter the name of your QoS egress policy.
4. Select the scheduling mode to use.
5. For each queue, enter a description, select the drop policy to use, and enter the minimum rate in kbps, maximum rate in kbps, weight value, and WRED slope. If you select *Weighted Random Early Detection Drop Policy*, you can use ECN marking by selecting the *ECN* checkbox.
6. Select *Add*.

Using the CLI:

```
config switch qos qos-policy
   edit <policy_name>
      set rate-by {kbps | percent}
      set schedule {strict | round-robin | weighted}
```

```
      config cos-queue
         edit [queue-0 ... queue-7]
            set description <text>
            set drop-policy {taildrop | weighted-random-early-detection}
            set ecn {enable | disable}
            set max-rate <rate kbps>
            set min-rate <rate kbps>
            set max-rate-percent <percentage>
            set min-rate-percent <percentage>
            set weight <value>
            set wred-slope <value>
         next
      end
   next
end
```

Configure the egress drop mode

NOTE: To see which models support this feature, refer to the FortiSwitch feature matrix.

When there are too many packets going through the same egress port, you can choose whether packets are dropped on ingress or egress.

To set the egress drop mode:

```
config switch physical-port
   edit <port>
      set egress-drop-mode <disabled | enabled>
   end
```

Variable	Description
disabled	Drop packets on ingress.
enabled	Drop packets on egress.

NOTE: Because too many packets are going through the same egress port, you might want to use the pause frame for flow control on the ingress side. To see the pause frame on ingress, enable the flow control “tx” on the ingress interface and disable egress-drop-mode on the egress interface.

Configure the switch ports

You can configure the following QoS settings on a switch port or a trunk:

- trust dot1p values on ingress traffic and the dot1p map to use
- trust ip-dscp values on ingress traffic and the ip-dscp map to use. (**NOTE:** Trust the dot1p values **or** the ip-dscp values but not both.)
- an egress policy for the interface
- a default CoS value (for packets with no CoS value)

If neither of the trust policies is configured on a port, the ingress traffic is mapped to queue 0 on the egress port.

If no egress policy is configured on a port, the FortiSwitch unit applies the default scheduling mode (that is, round-robin).

Using the GUI:

1. Go to *Switch* > *Interfaces*.
2. Select the switch port to update and then click *Edit*.
3. Select the QoS egress policy in the *QoS Policy* drop-down list.
4. Select the 802.1p map in the *Trust 802.1p* drop-down list.
5. Select the DSCP map in the *Trust IP-DSCP* drop-down list.
6. Select *Update*.

Using the CLI:

```
config switch interface
   edit <port>
      set trust-dot1p-map <map-name>
      set trust-ip-dscp-map <map-name>
      set qos-policy < policy-name >
      set default-cos <default cos value 0-7>
   next
end
```

NOTE: The `set default-cos` command is not available on the following FortiSwitch models: 224D-FPOE, 248D, 424D, 424D-POE, 424D-FPOE, 448D, 448D-POE, 448D-FPOE, 224E, 224E-POE, 248E-POE, and 248E-FPOE.

Configure QoS on trunks

Configuring QoS on trunk interface follows the same configuration steps as for a switch port (configure a Dot1p/DSCP map and an egress policy).

When you add a port to a trunk, the port inherits the QoS configuration of the trunk interface. A port member reverts to the default QoS configuration when it is removed from the trunk interface.

Using the GUI:

1. Go to *Switch* > *Interfaces*.
2. Select the trunk to update and then click *Edit*.
3. Select the QoS egress policy in the *QoS Policy* drop-down list.
4. Select the 802.1p map in the *Trust 802.1p* drop-down list.
5. Select the DSCP map in the *Trust IP-DSCP* drop-down list.
6. Click *Update*.

Using the CLI:

The following example shows QoS configuration on a trunk interface:

```
config switch interface
   edit "tr1"
      set snmp-index 56
      set trust-dot1p-map "dot1p_map1"
      set default-cos 1
      set qos-policy "p1"
   next
end
```

When you configure an egress QoS policy with rate control on a trunk interface, that rate control value is applied to each port in the trunk interface. The FortiSwitch unit does not support an aggregate value for the whole trunk interface.

NOTE: The `set default-cos` command is not available on the following FortiSwitch models: 224D-FPOE, 248D, 424D, 424D-POE, 424D-FPOE, 448D, 448D-POE, 448D-FPOE, 224E, 224E-POE, 248E-POE, and 248E-FPOE.

Configure QoS on VLANs

You can configure a CoS queue value for a VLAN by creating an ACL policy:

```
config switch acl ingress
   edit 1
      config action
         set cos-queue 7
         set count enable
      end
      config classifier
         set vlan-id 200
      end
      set ingress-interface "port25"
      set status active
end
```

Configure CoS and DSCP markings

You can classify a packet by matching the CoS value, DSCP value, or both CoS and DSCP values. You can also configure the action to set the CoS marking value, DSCP marking value, or both.

```
config switch acl ingress
   edit <policy-id>
      config classifier
         set cos <802.1Q CoS value to match>
         set dscp <DSCP value to match>
      end
      config action
         set remark-cos <0-7>
         set remark-dscp <0-63>
      end
```

For example:

```
config switch acl ingress
   edit 1
      config classifier
         set src-mac 11:22:33:44:55:66
         set cos 2
         set dscp 10
      end
      config action
         set count enable
         set remark-cos 4
         set remark-dscp 20
      end
   set ingress-interface port2
   set status active
```

```
end
```

Checking the QoS statistics

To check the statistics for the QoS queues for all ports:

```
diagnose switch physical-ports qos-stats list
```

To check the statistics for the QoS queues for specific ports:

```
diagnose switch physical-ports qos-stats list <list_of_ports>
```

NOTE: The output differs depending on the FortiSwitch model.

To view the real-time egress QoS queue rates for specific ports:

```
diagnose switch physical-ports qos-rates list <list_of_ports>
```

To view the real-time egress QoS queue rates for all ports:

```
diagnose switch physical-ports qos-rates list
```

NOTE: To stop the output: press `CTRL+c`.

Resetting and restoring QoS counters

To reset the QoS counters to zero (applies to all applications except SNMP) for the specified ports:

```
diagnose switch physical-ports qos-stats set-qos-counter-zero [<port_list>]
```

To restore the QoS counters to the hardware values for the specified ports:

```
diagnose switch physical-ports qos-stats set-qos-counter-revert [<port_list>]
```

For example:

```
diagnose switch physical-ports qos-stats set-qos-counter-zero 2,4,7-9
diagnose switch physical-ports qos-stats set-qos-counter-revert 1,3-5,7
```

Network monitor

You can monitor specific unicast MAC addresses in directed mode, monitor all detected MAC addresses on a FortiSwitch unit in survey mode, or do both. The FortiSwitch unit gives the directed mode a higher priority than survey mode. The directed mode and survey mode are disabled by default.

This section covers the following topics:

- Survey mode on page 347
- Directed mode on page 348
- Network-monitoring statistics on page 349

Survey mode

In survey mode, the FortiSwitch unit detects MAC addresses to monitor for a specified number of seconds. You can specify network monitoring for 120 to 3,600 seconds. The default time is 120 seconds. The FortiSwitch unit detects various fields of the packet—such as MAC address, IP address, VLAN, and user name—and stores the data in either of two databases.

Using the GUI:

1. Go to *Switch* > *Network Monitor* > *Settings*.
2. Select the *Enabled* checkbox.
3. In the *DB Aging Interval (Seconds)* field, enter the number of seconds for entries in the network monitor database to be kept.

   By default, entries are kept for 1 hour.
4. Select the *Survey Mode* checkbox.
5. In the Survey Mode Interval (Seconds) field, enter the number of seconds that the network monitor will operate in survey mode.

   By default, the network monitor will operate for 2 minutes.
6. Click *Update*.

Using the CLI:

```
config switch network-monitor settings
   set status enable
   set db-aging-interval <3600-86400>
   set survey-mode enable
   set survey-mode-interval <120-3600>
end
```

For example:

```
config switch network-monitor settings
   set status enable
   set survey-mode enable
   set survey-mode-interval 480
```

```
end
```

Directed mode

In directed mode, you select which unicast MAC addresses that you want examined. The FortiSwitch unit detects various fields of the packet—such as MAC address, IP address, VLAN, and user name—and stores the data in either of two databases.

NOTE: You cannot specify broadcast or multicast MAC addresses.

The maximum number of MAC addresses that can be monitored depends on the FortiSwitch model.

Platform Series	Maximum Number of MAC Addresses Monitored	Maximum Number of Hosts
1xx, 2xx	10	250
4xx, 5xx	20	1,024
10xx, 30xx	30	4,096

Using the GUI:

1. Go to *Switch* > *Network Monitor* > *Settings*.
2. Select the *Enabled* checkbox.
3. Under *Directed*, click +.
4. In the *Monitor MAC* column, enter the MAC address (formatted like this: xx:xx:xx:xx:xx:xx) to monitor.
5. If you want to add another MAC address, click + under *Directed*.
6. Click *Update*.

Using the CLI:

1. Find out how many network monitors are available:

```
diagnose switch network-monitor cfg-stats

Network Monitor Configuration Statistics:
---------------------------------
Adds         : 0
Deletes      : 0
Free Entries : 20
```

2. Find out which network monitors are being used currently:

```
diagnose switch network-monitor dump-monitors

Entry ID       Monitor Type       Monitor MAC      Packet-count
=================================================================
1               directed-mode   00:01:02:03:04:05       10
2               directed-mode   10:01:02:03:04:05       0
3               survey-mode     08:5b:0e:c1:07:65       419
4               survey-mode     08:5b:0e:4f:af:38       101
5               survey-mode     08:5b:0e:ce:59:40       2347
6               survey-mode     08:5b:0e:4f:af:44       0
7               survey-mode     08:5b:0e:c1:07:65       0
8               survey-mode     08:5b:0e:4f:af:38       80
9               survey-mode     08:5b:0e:ce:59:40       117
10              survey-mode     08:5b:0e:4f:af:44       0
```

3. Enable network monitoring:

```
config switch network-monitor settings
   set status enable
end
```

4. Specify a single unicast MAC address to be monitored:

```
config switch network-monitor directed
   edit <unused network monitor>
      set monitor-mac <MAC address>
   next
end
```

For example:

```
config switch network-monitor directed
   edit 1
      set monitor-mac 00:25:00:61:64:6d
   next
end
```

Network-monitoring statistics

After you have enabled network monitoring, you can view the statistics for the number and types of packets.

Using the GUI:

1. Go to *Switch* > *Network Monitor* > *Settings*.
2. Make certain that the *Enabled* checkbox is selected and then click *Update*.
   NOTE: After the *Enabled* checkbox is selected, it takes about 10 seconds for results to be displayed.
3. Go to *Switch* > *Network Monitor* > *Statistics*.
4. If you want to delete the current statistics, click *Clear Cache*.
5. If you want to display the most recent results, click *Refresh*.

To view the type of packets going to and from monitored MAC addresses:

```
diagnose switch network-monitor parser-stats

Network Monitor Parser Statistics:
----------------------------------
Arp         : 0
Ip          : 1
Udp         : 46
Tcp         : 353
Dhcp        : 0
Eapol       : 0
Unsupported : 352
```

To view the number of packets going to and from monitored MAC addresses:

```
diagnose switch network-monitor dump-monitors
```

```
Entry ID       Monitor Type       Monitor MAC      Packet-count
=============================================================
1               directed-mode   00:01:02:03:04:05       10
2               directed-mode   10:01:02:03:04:05       0
3               survey-mode     08:5b:0e:c1:07:65       419
4               survey-mode     08:5b:0e:4f:af:38       101
5               survey-mode     08:5b:0e:ce:59:40       2347
6               survey-mode     08:5b:0e:4f:af:44       0
7               survey-mode     08:5b:0e:c1:07:65       0
8               survey-mode     08:5b:0e:4f:af:38       80
9               survey-mode     08:5b:0e:ce:59:40       117
10              survey-mode     08:5b:0e:4f:af:44       0
```

The FortiSwitch unit creates an entry in the layer-3 database using the exact packet contents when they were parsed. If the MAC address is then assigned to a different VLAN, this change might not be detected immediately. If there is a discrepancy in the output for the `diagnose switch network-monitor dump-l2-db` and `diagnose switch network-monitor dump-l3-db` commands, use the output with the more recent time stamp.

To view all detected devices from the layer-2 database:

```
diagnose switch network-monitor dump-l2-db

mac 00:01:02:03:04:05 vlan 1
created 19 secs ago, last seen 16 secs ago
user JoE sources: eapol
```

To view all detected devices from the IP address database:

```
diagnose switch network-monitor dump-l3-db


mac 08:5b:0e:c1:07:65 ip 169.254.2.2 vlan 4094
created 63614 secs ago, last seen 2 secs ago
sources: arp ip
mac 00:10:20:30:40:50 ip 10.10.10.111 vlan 123
created 75 secs ago, last seen 45 secs ago
sources: arp ip
mac 00:11:22:33:44:55 ip 30.30.30.115 vlan 1
created 53 secs ago, last seen 53 secs ago
sources: dhcp arp ip
```

Configuring security checks

You can enable various security checks for incoming TCP/UDP packets. The packet is dropped if it matches one of the security rules that have been enabled. Use the appropriate syntax for your FortiSwitch model:

- Syntax (for FS-108D-POE, FS-112D-POE, FS-224D-POE, and FSR-216F-POE) on page 352
- Syntax (for FS-1xxE, FS-1xxF, and FS-110G-FPOE) on page 352
- Syntax (for all other FortiSwitch models) on page 353

Syntax (for FS-108D-POE, FS-112D-POE, FS-224D-POE, and FSR-216F-POE)

```
config switch security-feature
   set tcp-syn-data {enable | disable}
   set tcp-udp-port-zero {enable | disable}
   set tcp_flag_zero {enable | disable}
   set tcp_flag_FUP {enable | disable}
   set tcp_flag_SF {enable | disable}
   set tcp_flag_SR {enable | disable}
   set tcp_frag_ipv4_icmp {enable | disable}
   set tcp_arp_mac_mismatch {enable | disable}
   set allow-mcast-sa {enable | disable}
end
```

Variable	Description	Default
tcp-syn-data	TCP SYN packet contains additional data (possible DoS attack).	disable
tcp-udp-port-zero	TCP or UDP packet has the source or destination port set to zero.	disable
tcp_flag_zero	TCP packet with all flags set to zero.	disable
tcp_flag_FUP	TCP packet with FIN, URG and PSH flags set.	disable
tcp_flag_SF	TCP packet with SYN and FIN flags set.	disable
tcp_flag_SR	TCP packet with SYN and RST flags set.	disable
tcp_frag_ipv4_icmp	Fragmented ICMPv4 packet.	disable
tcp_arp_mac_mismatch	ARP packet with MAC source address mismatch between the layer- 2 header and the ARP packet payload.	disable
allow-mcast-sa	Ethernet packet whose source MAC address is multicast.	disable

Syntax (for FS-1xxE, FS-1xxF, and FS-110G-FPOE)

```
config switch security-feature
   set tcp-flag-zero {enable | disable}
   set tcp-flag-FUP {enable | disable}
   set tcp-flag-SF {enable | disable}
   set tcp-flag-SR {enable | disable}
   set arp-mac-mismatch {enable | disable}
   set macsa-eq-macda {enable | disable}
   set sip-eq-dip {enable | disable}
   set tcp-port-eq {enable | disable}
   set udp-port-eq {enable | disable}
   set ip-pod {enable | disable}
   set icmp-frag {enable | disable}
   set tcp-frag-off-min {enable | disable}
   set tcp-syn-sp-less-1024 {enable | disable}
   set invalid-ipv4-hdr-len {enable | disable}
   set gratuitous-arp {enable | disable}
end
```

Variable	Description	Default
tcp-flag-zero	TCP packet with all flags set to zero.	disable
tcp-flag-FUP	TCP packet with FIN, URG, and PSH flags set.	disable
tcp-flag-SF	TCP packet with SYN and FIN flags set.	disable
tcp-flag-SR	TCP packet with SYN and RST flags set.	disable
arp-mac-mismatch	ARP packet with MAC source address mismatch between the MAC header and the ARP packet payload.	disable
macsa-eq-macda	Packet with source MAC address equal to the destination MAC address.	disable
sip-eq-dip	TCP packet with source IP address equal to the destination IP address.	disable
tcp-port-eq	TCP packet with the same source and destination TCP port.	disable
udp-port-eq	IP packet with the same source and destination UDP port.	disable
ip-pod	The IPv4/IPv6 packet length is larger than 64 kB.	disable
icmp-frag	Fragmented ICMP packet.	disable
tcp-frag-off-min	TCP non-initial fragments carry the TCP header.	disable
tcp-syn-sp-less-1024	TCP SYN packet with a source port less than 1024.	disable
invalid-ipv4-hdr-len	IPv4 packet with a header length greater than the total length. **NOTE:** This command is available only on the FS-124F, FS-124F-FPOE, FS-124F-POE, FS-148F, FS-148F-FPOE, and FS-148F-POE models.	disable
gratuitous-arp	Gratuitous ARP packet. **NOTE:** This command available only on the FS-108E, FS-108E-FPOE, FS-108E-POE, FS-108F, FS-108F-FPOE, FS-108F-POE, FS-124E, FS-124E-FPOE, FS-124E-POE, FS-148E, and FS-148E-POE models.	disable

Syntax (for all other FortiSwitch models)

```
config switch security-feature
   set sip-eq-dip {enable | disable}
   set tcp-flag {enable | disable}
   set tcp-port-eq {enable | disable}
   set tcp-flag-FUP {enable | disable}
   set tcp-flag-SF {enable | disable}
   set v4-first-frag {enable | disable}
   set udp-port-eq {enable | disable}
   set tcp-hdr-partial {enable | disable}
   set macsa-eq-macda {enable | disable}
   set allow-mcast-sa {enable | disable}
   set allow-sa-mac-all-zero {enable | disable}
end
```

Variable	Description	Default
sip-eq-dip	TCP packet with the same source IP address and destination IP address.	disable
tcp-flag	DoS attack checking for TCP flags.	disable
tcp-port-eq	TCP packet with the same source and destination TCP port.	disable
tcp-flag-FUP	TCP packet with FIN, URG, and PSH flags set, and sequence number is zero.	disable
tcp-flag-SF	TCP packet with SYN and FIN flags set.	disable
v4-first-frag	DoS attack checking for IPv4 first fragment.	disable
udp-port-eq	IP packet with the same source and destination UDP port.	disable
tcp-hdr-partial	TCP packet with partial header.	disable
macsa-eq-macda	Packet with the same source MAC address and destination MAC address.	disable
allow-mcast-sa	Ethernet packet whose source MAC address is multicast.	disable
allow-sa-mac-all-zero	Ethernet packet whose source MAC address is all zeros.	disable

Cut-through switching mode

By default, all FortiSwitch models use the store-and-forward technique to forward packets. This technique waits until the entire packet is received, verifies the content, and then forwards the packet.

Some switch models also have a cut-through switching mode to reduce latency. This technique forwards the packet as soon as the switch receives it. To see which models support this feature, refer to the FortiSwitch feature matrix.

NOTE: For the FS-3032D model, the cut-through switching mode is not supported on split ports.

To change the switching mode for the main buffer for these three models, use the following commands:

```
config switch global
   set packet-buffer-mode {store-forward | cut-through}
end
```

NOTE: Changing the switching mode might stop traffic on all ports during the change.

Enabling packet forwarding

NOTE: These commands apply only to the 200 Series and 400 Series.

If you want to use layer-3 interfaces and IGMP snooping on certain FortiSwitch models, you must enable the forwarding of reserved multicast packets and IPv6 neighbor-discovery packets to the CPU. These features are enabled by default.

```
config switch global
   set reserved-mcast-to-cpu {enable | disable}
```

```
   set neighbor-discovery-to-cpu {enable | disable}
end
```

Configuring auto-topology

Starting in FortiSwitchOS 7.2.0, auto-network is enabled by default.

After an `execute factoryreset` command is executed on a FortiSwitch unit in standalone mode, the auto-network configuration is enabled by default. If you are not using auto-network, you must manually disable it:

```
config switch auto-network
   set status disable
end
config switch global
   set auto-fortilink-discovery disable
end
```

Use the auto-topology feature to automatically form an inter-switch link (ISL) between two switches. The ISL trunk contains all physical members connected between the same two switches. If a new link is connected between two switches with an existing ISL, the link is automatically added to the existing ISL trunk. The ISL trunk mode is lacp-active.

The ISL switch interface has the following properties:

- The native VLAN is determined by the auto-network mgmt-vlan setting.
- The allowed VLANs are 1 to 4,094.
- DHCP snooping is set to trusted.
- The STP edge port is disabled.

The mgmt-vlan is the VLAN to use for the native VLAN on ISL ports and the native VLAN on the internal switch interface.

NOTE: Do not use the same VLAN for the mgmt-vlan and an existing switch virtual interface (SVI).

To configure auto-topology:

```
config switch auto-network
   set mgmt-vlan <1-4094>
   set status {enable | disable}
end
```

The following example shows how to enable auto-topology:

```
config switch auto-network
   set mgmt-vlan 101
   set status enable
end
```

The following example shows the ISL trunk:

```
config switch trunk
   edit "8DVTA18000567-0"
      set mode lacp-active
      set auto-isl 1
      set members "port8"
   next
```

```
end
```

The following example shows the ISL switch interface:

```
config switch interface
   edit "8DVTA18000567-0"
      set allowed-vlans 1-4094
      set dhcp-snooping trusted
      set edge-port disabled
      set snmp-index 12
   next
end
```

To disable auto-topology for individual ports, assign the `default` LLDP profile (or a user-defined LLDP profile with `auto-isl` disabled). For example:

```
config switch physical-port
   edit "port8"
      set lldp-profile "default"
      set speed auto
   next
end
```

Viewing port statistics

Using the GUI:

Go to *Switch* > *Monitor* > *Port Stats*.

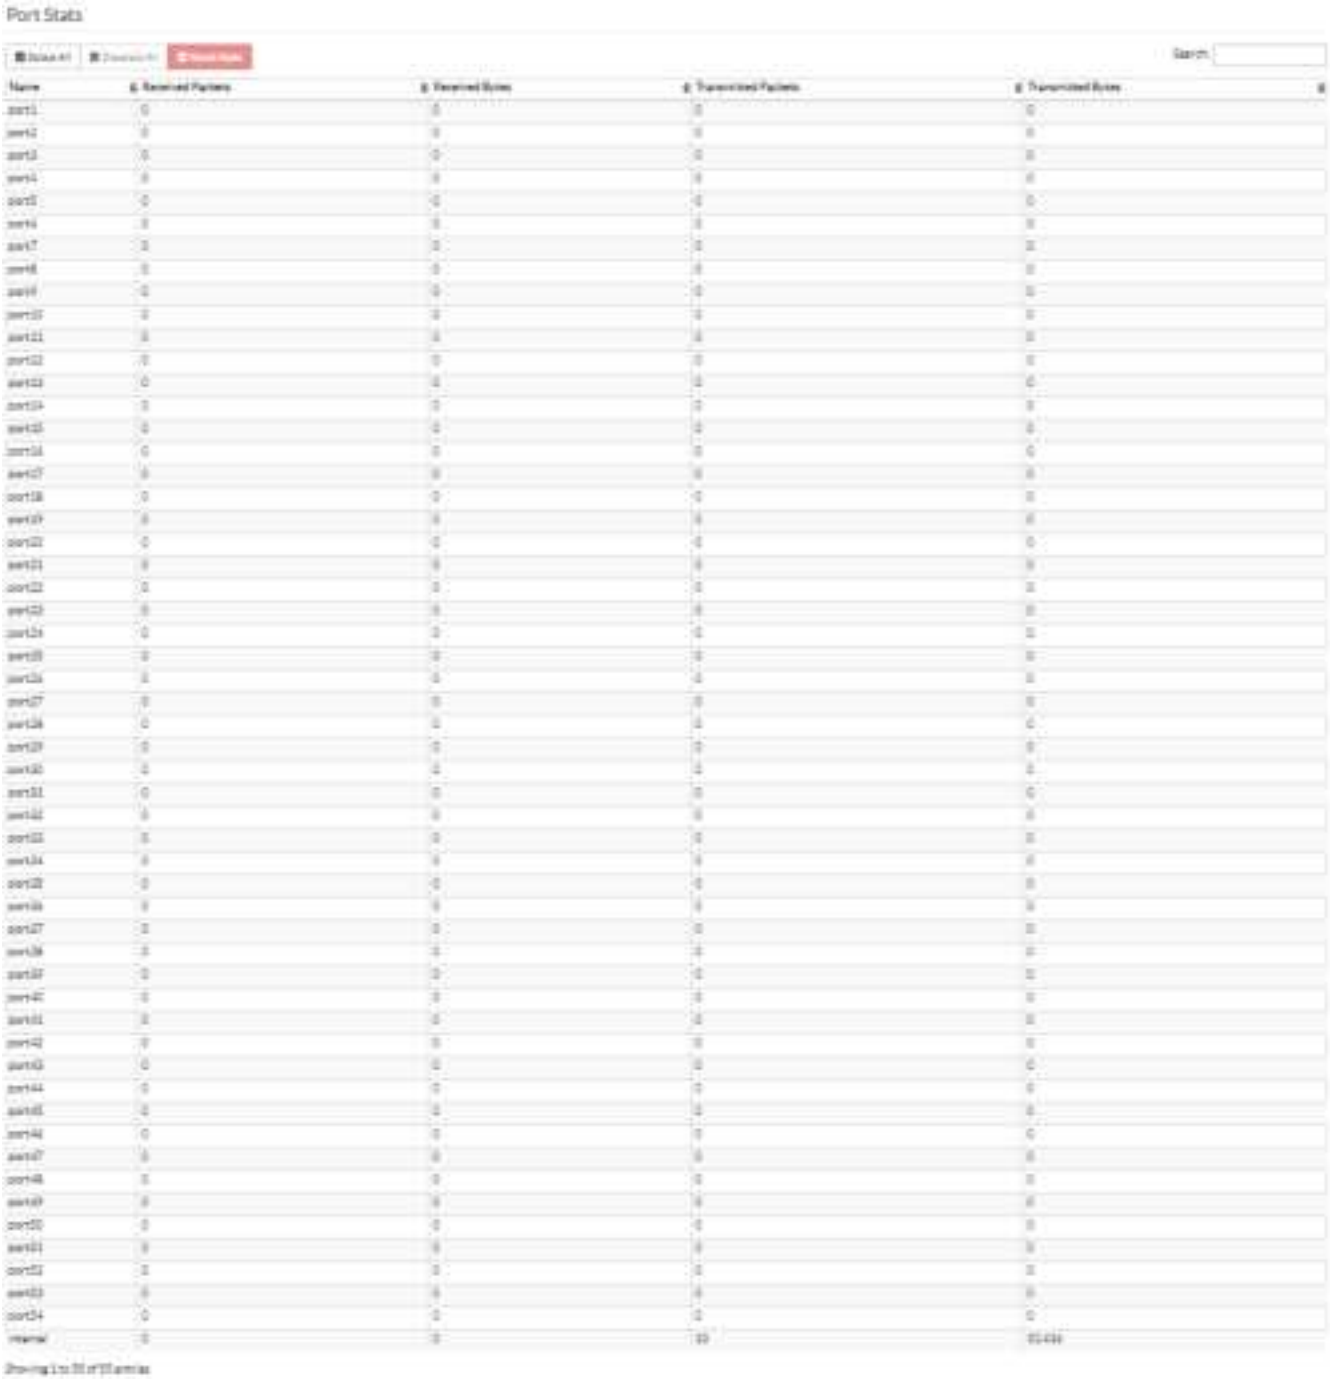

To clear the statistics on all ports, select *Select All* and then select *Reset Stats*.

To clear the statistics on some of the ports, select the ports and then select *Reset Stats*.

You can also view the incoming and outgoing traffic rates by going to the *Switch* > *Physical Ports* page:

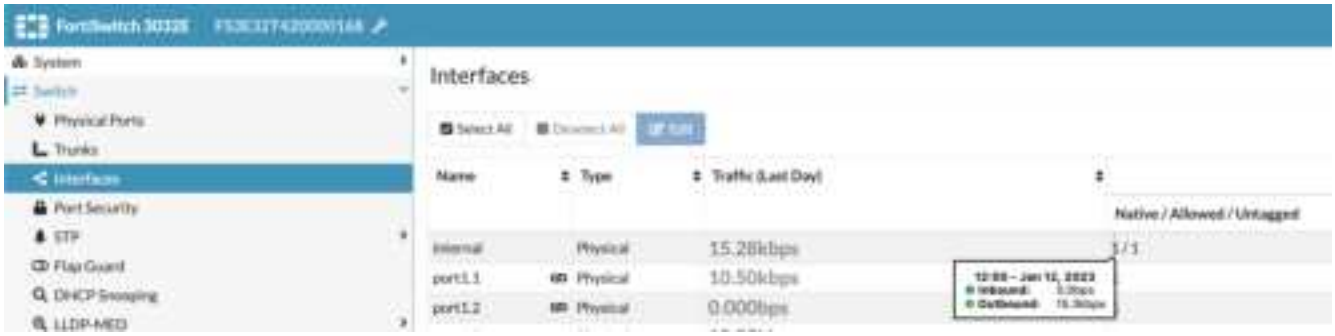

Using the CLI:

```
diagnose switch physical-ports port-stats list [<list_of_ports>]
```

For example:

```
diagnose switch physical-ports port-stats list 1,3,4-6
```

To clear all hardware counters (except for QoS, SNMP, and web GUI counters) on the specified ports:

```
diagnose switch physical-ports set-counter-zero [<list_of_ports>]
```

To restore hardware counters (except for QoS, SNMP, and web GUI counters) on the specified ports:

```
diagnose switch physical-ports set-counter-revert [<list_of_ports>]
```

Media Redundancy Protocol

A ring of Ethernet switches can use the Media Redundancy Protocol (MRP) to overcome a failure faster than with STP. An MRP network consists of a ring of switches with one manger switch; the rest of the switches are clients. The MRP ring ports must be physical ports or a single port configured as a static trunk (this requirement supports FortiLink). The MRP ring ports are disabled in STP.

If a ring has more than one switch that can be manager, MRP selects the switch with the highest priority (numerically lower number) as the automanager. If a ring has more than one switch that can be manager and the switches have the same priority, MRP selects the switch with the lowest MAC address as the automanager. Each node of the MRP network must be configured as an automanager (manager switch) or a client. The MRP network cannot contain both a manually configured manager and automanager. The MRP automanager and client switches must have matching parameters, such as MRP VLAN and domain identifier, for the MRP ring to function properly.

MRP sends three types of frames through the ring ports:

- MRP_Test frames detect a failure or recovery of a ring port link.
- MRP_LinkChange frames indicate a failure or recovery of a ring port link.
- MRP_TopologyChange frames indicate that the MRP network topology has changed.

Starting in FortiSwitchOS 7.4.4, the FortiSwitch unit supports the following:

- Two MRP rings
- One automanager per client
- Ring-check mode
- The media redundancy interconnection manager (MIM) is not supported.
- The media redundancy interconnection client (MIC) is not supported.

Refer to the FortiSwitch feature matrix to see which FortiSwitch models support this feature.

To configure two MRP rings, each ring must have unique values for `domain-id`, `ring-port1`, and `ring-port2`. The `domain-id` must match all of the FortiSwitch units in the MRP ring.

Configuring an MRP network

Configuring an MRP network requires the following steps:

1. Configure both ring ports with the MRP VLAN as the `allowed-vlans` or `native-vlan`.
2. Use the default MRP profile (`500ms`) or create a custom MRP profile.
3. Configure the settings for the MRP manager.
4. Configure the settings for the MRP client.

NOTE: The MRP VLAN identifier must be configured as `allowed-vlans` or `native-vlan` on both ring ports. If there is mismatch between the MRP `vlan-id` and the ring-ports VLAN, MRP is disabled. If MRP is disabled because of a mismatch, you need to configure both ring ports for the MRP VLAN, and then you can manually enable the MRP status.

To configure ring-port 1:

```
config switch interface
   edit "<ring_port1>"
      set allowed-vlans <1-4094>
   next
end
```

For example:

```
config switch interface
   edit "port8"
      set allowed-vlans 4094
   next
end
```

To configure ring-port 2:

```
config switch interface
   edit " <ring_port2>"
      set allowed-vlans <1-4094>
   next
```

```
end
```

For example:

```
config switch interface
   edit "port27"
      set allowed-vlans 4094
   next
end
```

To create a custom MRP profile:

```
config switch mrp profile
   edit <MRP_profile_name>
      set default-test-interval <30-50 ms>
      set short-test-interval <10-30 ms>
      set test-monitoring-count <1-5>
      set topology-change-interval <10-20 ms>
      set topology-change-repeat-count <1-5>
   next
end
```

NOTE: With a custom profile, some parameters on the lower range, such as `test-monitoring-count` and `default-test-interval`, might make the MRP ring unstable or flapping. Fortinet recommends fine-tuning these parameters in a custom profile to ensure a stable MRP ring.

To configure the settings for the MRP manager:

```
config switch mrp settings
   edit <MRP_ring_ID>
      set status enable
      set role automanager
      set domain-id <32_hexadecimal_digits>
      set domain-name <domain_name>
      set vlan-id <1-4094>
      set priority <0-65535>
      set ring-port1 <port_name>
      set ring-port2 <port_name>
      set profile-name {500ms | <custom_profile_name>}
   next
end
```

For example:

```
config switch mrp settings
   edit 1
      set status enable
      set role automanager
      set domain-id FFFFFFFF-FFFF-FFFF-FFFF-FFFFFFFFFFFF
      set domain-name domain1
      set vlan-id 4094
      set priority 40960
      set ring-port1 port7
      set ring-port2 port8
      set profile-name profile1
   next
end
```

To configure the settings for the MRP client:

```
config switch mrp settings
   edit <MRP_ring_ID>
      set status enable
      set role client
      set domain-id <32_hexadecimal_digits>
      set domain-name <domain_name>
      set vlan-id <1-4094>
      set priority <0-65535>
      set ring-port1 <port_name>
      set ring-port2 <port_name>
      set profile-name {500ms | <custom_profile_name>}
   next
end
```

For example, if you are configuring one MRP ring:

```
config switch mrp settings
   edit 1
      set status enable
      set role client
      set domain-id FFFFFFFF-FFFF-FFFF-FFFF-FFFFFFFFFFFF
      set domain-name domain1
      set vlan-id 4094
      set priority <0-65535>
      set ring-port1 port8
      set ring-port2 port27
      set profile-name profile1
   next
end
```

For example, if you are configuring two MRP rings:

```
config switch mrp settings
   edit "1"
      set status disable
      set role client
      set domain-id FFFFFFFF-FFFF-FFFF-FFFF-FFFFFFFFFFFF
      set domain-name "domain1"
      set vlan-id 1
      set priority 40960
      set ring-port1 port7
      set ring-port2 port8
      set profile-name "500ms"
   next
   edit "2"
      set status disable
      set role client
      set domain-id 9F35CD67-25D0-3B4C-1915-F8B549C47899
      set domain-name "domain2"
      set vlan-id 1
      set priority 40960
      set ring-port1 port9
      set ring-port2 port10
      set profile-name "500ms"
   next
end
```

Viewing the MRP configuration

To display the current MRP settings:

```
get switch mrp settings
```

To display the current MRP status:

```
diagnose switch mrp status {1| 2}
```

To display the statistics for the MRP manager:

```
diagnose switch mrp stats {1| 2}
```

To delete the statistics for the MRP manager:

```
diagnose switch mrp clear {1| 2}
```

To see the configuration commands for the specified MRP profile:

```
show switch mrp profile <MRP_profile_name>
```

To see the configuration commands for the MRP settings:

```
show switch mrp settings
```

Precision Time Protocol

This section covers the following topics:

- PTP node types on page 363
- PTP message types on page 364
- Packet flow in end-to-end mode on page 365
- Packet flow in peer-to-peer mode on page 366
- PTP profiles on page 366
- PTP settings on page 368
- PTP policies on page 368
- Topology examples on page 369
- Configuring PTP on page 369
- Troubleshooting PTP on page 371
- Configuration example on page 372
- PTP operation details and limitations on page 373

The Precision Time Protocol (PTP) defines packet-based time synchronization, which is described in IEEE 1588v2. You can use PTP to synchronize clocks across networks for high clock accuracy within the submicrosecond to nanosecond range. PTP is required in time-sensitive applications such as telecommunications, audiovisual, and electric power generation, transmission, and distribution.

FortiSwitchOS supports only PTP version 2.

There are two PTP modes:

- End-to-end mode

  The link delay is measured from end to end between the primary node and secondary node. The sum of the switch residence time along the path is reported to the secondary node. The Correction Field (CF) of the PTP Sync messages is then updated with the residence time.
- Peer-to-peer mode

  All link delays are measured on a peer-to-peer basis. The sum of the switch residence time and link delay along the path is reported to the secondary node. The CF of the PTP Sync messages is then updated with the link delay added to the residence time

The following table describes FortiSwitchOS support for PTP.

Supported platforms	Supported software releases	Supported PTP modes
FS-424E-Fiber, FS-448E, FS-448E-POE, and FS-448E-FPOE	• FortiSwitchOS 7.4.3 or later • FortiSwitchOS 7.2.7 or later	End-to-end and peer-to-peer modes
FSR-424F-POE	FortiSwitchOS 7.4.0 or later	End-to-end and peer-to-peer modes
FS-224D, FS-224E, FS-4xxE, 500 series, FS-1024E, FS-1048E, FS-T1024E, FS-T1024F-FPOE, FS-3032E	FortiSwitchOS 6.4.0 or later	End-to-end mode

PTP node types

FortiSwitchOS supports both the transparent clock and boundary clock.

PTP node	Description
Transparent clock	This PTP node has more than one port. It does not participate in the Best Master Clock Algorithm and does not pass the VLAN boundary. PTP messages (except peer-to-peer PDelay messages) are forwarded through the FortiSwitch unit. The transparent clock updates the CF in the Sync messages based on the PTP mode.
Boundary clock	Starting in FortiSwitchOS 7.4.3, the FSR-424F-POE, FS-424E-Fiber, FS-448E, FS-448E-POE, and FS-448E-FPOE models support the boundary clock. The boundary clock is a PTP node with more than one port. It connects to multiple networks across VLANs. The boundary clock participates in the Best Master Clock Algorithm. It synchronizes itself to a primary clock and then distributes the time.

PTP message types

The following table describes the types of PTP messages.

PTP message	Description	PTP mode	PTP node
Announce	The Announce message establishes the synchronization hierarchy and announces who is the grandmaster clock (GM).	Both	Boundary and transparent clocks
Sync	The Sync message provides the value of the GM's time to the secondary clock.	Both	Boundary and transparent clocks
Follow_Up	The Follow_Up message provides the value of the synchronized time, which is used in the two-step mode.	Both	Boundary and transparent clocks
Delay_Req	The Delay_Req message provides the secondary clock time to the GM.	End-to-end mode	Boundary and transparent clocks
Delay_Resp	The Delay_Resp message provides the GM time to the secondary clock.	End-to-end mode	Boundary and transparent clocks
PDelay_Req	The PDelay_Req message measures the PTP link delay between two PTP ports.	Peer-to-peer mode	Boundary and transparent clocks in peer-to-peer mode
PDelay_Resp	The PDelay_Resp message measures the PTP link delay between two PTP ports.	Peer-to-peer mode	Boundary and transparent clocks in peer-to-peer mode
PDelay_Resp_ Follow_Up (peer to peer)	The PDelay_Resp_Follow_Up message measures the PTP link delay between two PTP ports, which is used in the two-step mode.	Peer-to-peer mode	Boundary and transparent clocks in peer-to-peer mode

Packet flow in end-to-end mode

The following figure shows the end-to-end Sync message (two-step mode):

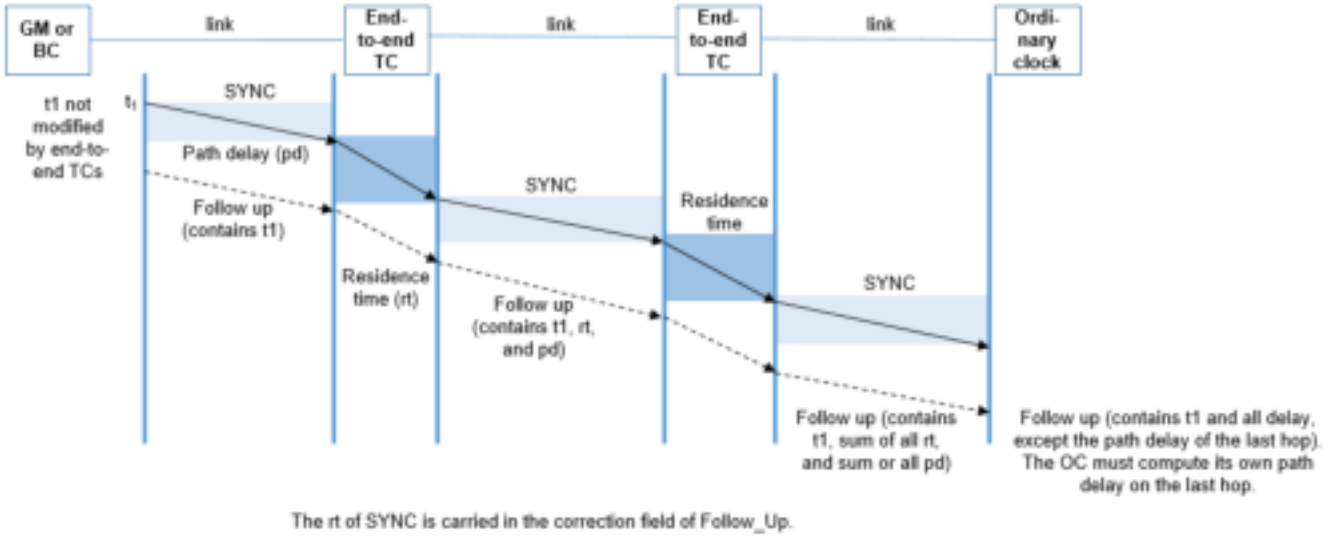

The rt of SYNC is carried in the correction field of Follow_Up.

The following figure shows the end-to-end Delay_Req and Delay_Resp messages (two-step mode):

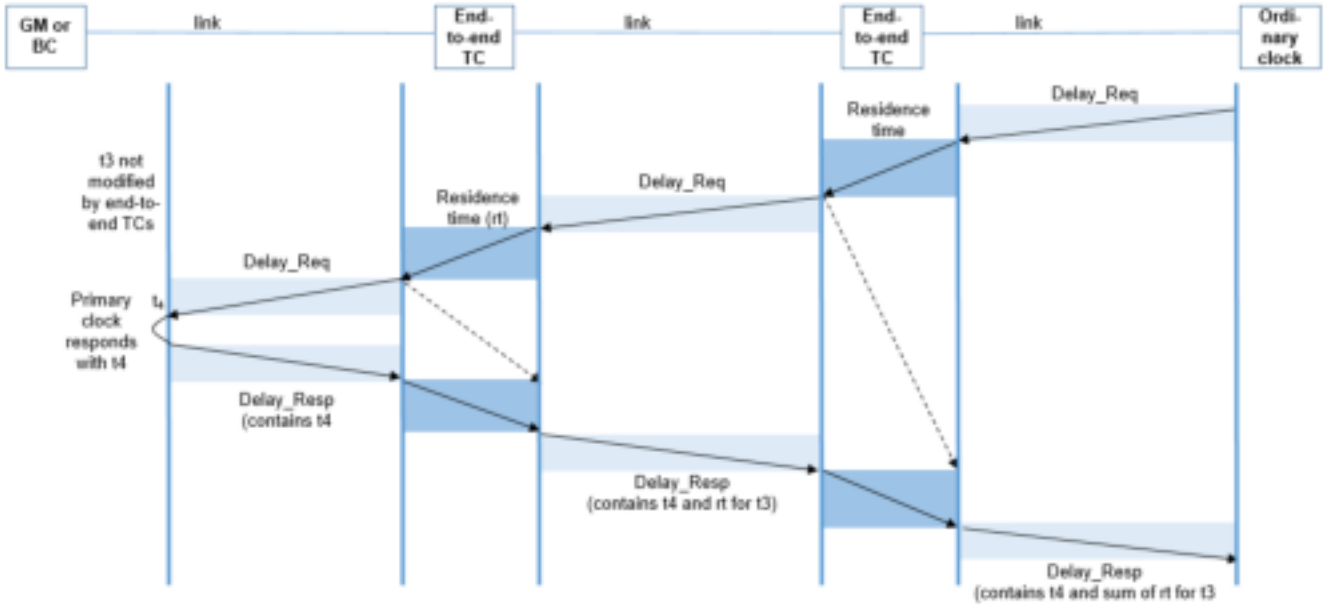

The rt of Delay_Req is carried in the correction field of Delay_Resp.

Packet flow in peer-to-peer mode

The following figure shows the Sync and Delay messages in the peer-to-peer transparent clock:

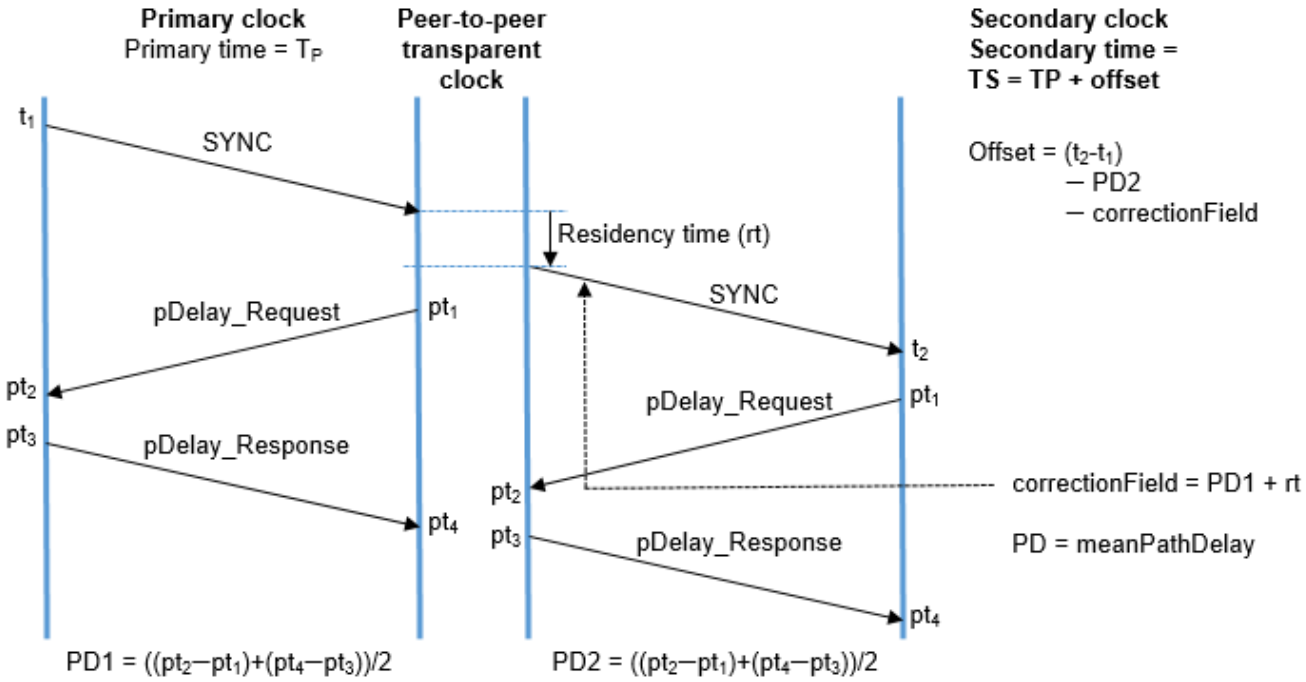

PTP profiles

If you are not using one of the defined profiles, configure a custom PTP profile.

Parameter	Description	Default value in end-to-end mode	Default value in peer-to-peer mode
{default | name_of_PTP_profile}	Name of the PTP profile.	No default	No default
announce-interval {0.25sec | 0.5sec | 1sec | 2sec | 4sec}	Select the number of seconds between Announce messages. This option is available only when `mode` is set to `boundary-e2e` or `boundary-p2p`.	1sec	1sec
announce-timeout <2-10>	Select how many seconds before the PTP Announce message expires. This option is available only when `mode` is set to `boundary-e2e` or `boundary-p2p`.	3	3

Parameter	Description	Default value in end-to-end mode	Default value in peer-to-peer mode
description <description_of_PTP_profile>	Description of the PTP profile.	No default	No default
domain <0-255>	PTP domain number. The range of values is 0-255. This option is available only when `mode` is set to `transparent-p2p`, `boundary-e2e`, or `boundary-p2p`.	1	For the transparent clock, the default value is 1 if using the `default` PTP profile or 254 if using the `power` PTP profile.
min-delay-req-interval {0.25sec | 0.5sec | 1sec | 2sec | 4sec}	Select the number of seconds between Delay_Req messages. This option is available only when `mode` is set to `boundary-e2e`.	1sec	Not applicable
mode {boundary-e2e | boundary-p2p | transparent-e2e | transparent-p2p}	PTP mode. You can select from the following modes: • `boundary-e2e`—Boundary clock using the end-to-end mode. • `boundary-p2p`—Boundary clock using the peer-to-peer mode. • `transparent-e2e`—Transparent clock using the end-to-end mode. • `transparent-p2p`—Transparent clock using the peer-to-peer mode.	transparent-e2e	Not applicable. You need to create a profile and set the mode to `boundary-p2p` or `tranparent-p2p`.
pdelay-req-interval {0.25sec | 0.5sec | 1sec | 2sec | 4sec}	The time between PDelay_Req messages. You can select 0.25, 0.5, 1, 2, or 4 seconds. The default value is 1 second. This option is available only when `mode` is set to `transparent-p2p` or `boundary-p2p`.	Not applicable	1sec
priority1 <0-255>	Set the PTP priority 1. Use a smaller number for a higher priority. This option is available only when `mode` is set to `boundary-e2e` or `boundary-p2p`.	128	128

Parameter	Description	Default value in end-to-end mode	Default value in peer-to-peer mode
priority2 <0-255>	Set the PTP priority 2. Use a smaller number for a higher priority. This option is available only when `mode` is set to `boundary-e2e` or `boundary-p2p`.	128	128
ptp-profile {default | C37.238-2017}	PTP profile. Select `default` for the IEEE 1588 default profile or `C37.238-2017` for the power profile. `C37.238-2017` is available only when `mode` is set to `transparent-p2p`.	default	default
sync-interval {0.25sec | 0.5sec | 1sec | 2sec | 4sec}	Select how many seconds between clock synchronization.	1sec	1sec
transport l2-mcast	PTP message transmission. This option is available only when `mode` is set to `transparent-p2p`, `boundary-e2e`, or `boundary-p2p`.	Layer-2 and layer-3 multicast are supported for end-to end transparent clock. All other modes support layer-2 multicast only.	Layer-2 multicast

PTP settings

Enable or disable PTP and select which PTP profile will use these PTP settings.

Parameter	Description	Default value
status	Enable or disable PTP.	disable
profile	The `default` profile is automatically selected. **NOTE:** On some legacy platforms, the `default` profile must be manually selected.	`default`

PTP policies

You can configure a custom PTP policy or use the default PTP policy.

Parameter	Description	Default value
name	Name of the PTP policy.	`default`
description	Description of the PTP policy.	No default

Parameter	Description	Default value
vlan	The VLAN that will use the PTP policy. The range of values is 0-4094. Setting `vlan` to 0 means that the native VLAN is used for PDelayXXX messages. **NOTE:** The VLAN must be a valid VLAN that the interface belongs to. Selecting an invalid VLAN can affect the performance.	0
vlan-pri	The priority of the PTP VLAN; it corresponds to the 802.1p priority. The VLAN priority is used only when there is traffic congestion. The range of values is 0-7. Set `vlan-pri` to 7 for the highest priority.	4

Topology examples

The following are three examples of supported topologies.

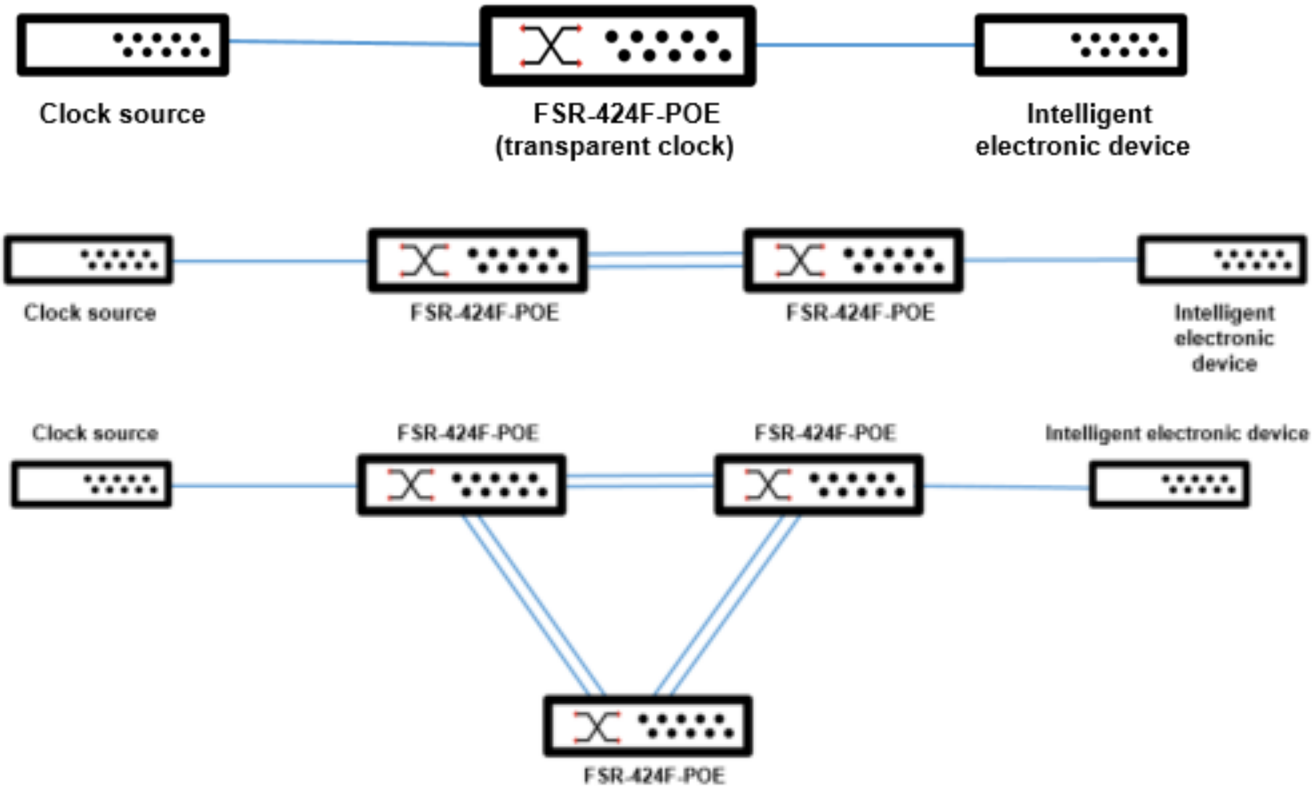

Configuring PTP

Follow these steps to configure PTP:

1. Configure a PTP profile or use the `default` profile.
2. Configure the PTP settings.

Enable or disable PTP and select which PTP profile will use these PTP settings. The `default` profile is automatically selected.

3. Configure the default PTP policy or create a custom PTP policy.

   Select which VLAN will use the PTP policy and the priority of the VLAN. The default PTP policy is applied to all ports. If you want to select which ports to apply the PTP policy to, you need to create a custom PTP policy.
4. If you are not using the default PTP policy, select which port to apply your custom PTP policy to.

   By default, the PTP status is enabled.

To configure PTP:

1. If you are not using the `default` profile, configure a PTP profile:

```
config system ptp profile
   edit {default | name_of_PTP_profile}
      set announce-interval {0.25sec | 0.5sec | 1sec | 2sec | 4sec}
      set announce-timeout <2-10>
      set description <description_of_PTP_profile>
      set domain <0-255>
      set min-delay-req-interval {0.25sec | 0.5sec | 1sec | 2sec | 4sec}
      set mode {boundary-e2e | boundary-p2p | transparent-e2e | transparent-p2p}
      set pdelay-req-interval {0.25sec | 0.5sec | 1sec | 2sec | 4sec}
      set priority1 <0-255>
      set priority2 <0-255>
      set ptp-profile {default | C37.238-2017}
      set sync-interval {0.25sec | 0.5sec | 1sec | 2sec | 4sec}
      set transport l2-mcast
   next
end
```

   For example:

```
config system ptp profile
   edit newprofile
      set description "New PTP profile"
      set domain 1
   next
end
```

2. Enable or disable PTP and select which PTP profile will use these PTP settings:

```
config switch ptp settings
   set status {enable | disable}
   set profile {default | name_of_PTP_profile}
end
```

   For example:

```
config switch ptp settings
   set status enable
   set profile newprofile
end
```

3. Configure the default PTP policy or create a custom PTP policy:

```
config system ptp interface-policy
   edit {default | PTP_policy_name}
      set description <description_of_PTP_policy>
      set vlan <0-4094>
      set vlan-pri <0-7>
   next
end
```

For example:

```
config system ptp interface-policy
   edit newPTPpolicy
      set description "PTP policy for VLAN 100"
      set vlan 100
      set vlan-pri 3
   next
end
```

4. If you are not using the default PTP policy, apply your custom PTP policy to a port:

```
config switch interface
   edit <interface_name>
      set ptp-status {enable | disable}
      set ptp-policy <PTP_policy_name>
   next
end
```

For example:

```
config switch interface
   edit port5
      set ptp-status enable
      set ptp-policy newPTPpolicy
   next
end
```

Troubleshooting PTP

Use the following command to troubleshoot your PTP configuration:

```
diagnose switch ptp port get-link-delay
```

For example:

```
SR24FPTF21000005 # diagnose switch  ptp port get-link-delay
Portname     Speed  Link-Delay

__________   _____  ___________
port1        1G      -
port2        1G      -
port3        2.5G    2286ns
port4        2.5G    2300ns
port5        1G      628ns
port6        1G      628ns
port7        2.5G    2294ns
port8        1G      -
port9        2.5G    4718ns
port10       2.5G    -
port11       2.5G    4600ns
port12       2.5G    4614ns
port13       1G      679ns
port14       1G      669ns
port15       1G      675ns
port16       1G      681ns
port17       1G      688ns
port18       2.5G    -
port19       2.5G    -
port20       2.5G    -
port21       2.5G    -
```

```
port22       1G      670ns
port23       2.5G    -
port24       2.5G    -
port25       10G     -
port26       10G     -
port27       10G     -
port28       10G     -
port29       40G     -
port30       40G     32ns
```

Configuration example

```
config system ptp profile
   edit "default"
      set mode transparent-e2e
   next
   edit "power"
      set domain 254
      set mode transparent-p2p
      set ptp-profile C37.238-2017
   next
   edit "1588-e2e"
      set mode transparent-e2e
   next
end

config switch ptp settings
   set status enable
   set profile "power"
end

config system ptp interface-policy
   edit "default"
   next
   edit "policy1"
      set vlan 100
      set vlan-pri 4
   next
   edit "policy2"
      set vlan 200
      set vlan-pri 7
   next
end

config switch interface
   edit "port3"
      set allowed-vlans 100
      set snmp-index 51
      set ptp-policy "policy1"
   next
end
```

PTP operation details and limitations

Review the following limitations before configuring PTP:

- Layer-3 peer-to-peer mode is not supported.
- When using the peer-to-peer mode, interoperability with IGMP snooping is not supported.
- When using the peer-to-peer mode, using the `diagnose switch ptp port add-link-delay` command does not work.
- When using the peer-to-peer mode, Fortinet recommends setting `pdelay-req-interval` to `1sec` because the power profile limits it to 1 second.
- When using the peer-to-peer mode, PTP events are not logged in syslog.
- Setting `ptp-policy` on a switch interface is valid only in peer-to-peer mode.
- When using the end-to-end mode, if PTP is over IPv4 multicast and IGMP snooping is enabled on the VLAN that forwards PTP, you must add an IGMP static group entry to facilitate the forwarding of the PTP packets.
- The STP blocking port does not block PTP Announce, Sync, and FollowUp messages.
- Ports send out PDelayXXX messages intermittently when other ports do not allow the PTP VLAN. To work around this issue, make sure that the configured PTP VLAN in the interface policy is allowed (or native) in all ports in the switch.
- In the HSR and PRP modes, the PDelayXXX messages are always untagged, and the policy on the interface is ignored; this limitation applies only to HSR and PRP ports.
- Under the HSR/PRP mode, the interlink ports can only send untagged PDelayXXX messages. The PTP VLAN used by the interlink ports must be the same as the native VLAN, which matches `hsr-internal-vlan` or `prp-internal-vlan`.
- The link delay for the HSR/PRP trunk ports is not computed after the FortiSwitch unit is restarted. To work around this issue, use an automation stitch to restart PTP. For example:

```
config system automation-action
   edit "restart-ptp"
      set action-type cli-script
      set script "config switch ptp setting
      set status disable
      end
      config switch ptp setting
      set status enable
      end"
      set accprofile "super_admin"
   next
end

config system automation-trigger
   edit "restart-ptp"
      set event-type reboot
   next
end

config system automation-stitch
   edit "1"
      set trigger "restart-ptp"
      set action "restart-ptp"
   next
end
```

High-Availability Seamless Redundancy

This section covers the following topics:

- Configuring the HSR settings on page 375
- Configuring the HSR ring on page 375
- Clearing the HSR statistics on page 376
- Troubleshooting HSR on page 376
- Configuration examples on page 376
- HSR operation details and limitations on page 380

High-Availability Seamless Redundancy (HSR) is defined in the international standard IEC 62439-3-2016 clause 5. HSR provides seamless communication with fault tolerance by duplicating every unicast frame sent in HSR networks. Although HSR can be used in different topologies such as ring, bus, and mesh, the most commonly used topology is a single ring topology. This document focuses on the HSR ring topology. A simple HSR network consists of doubly attached bridging nodes, each having two ring ports, interconnected by full-duplex links, as shown in the following figure:

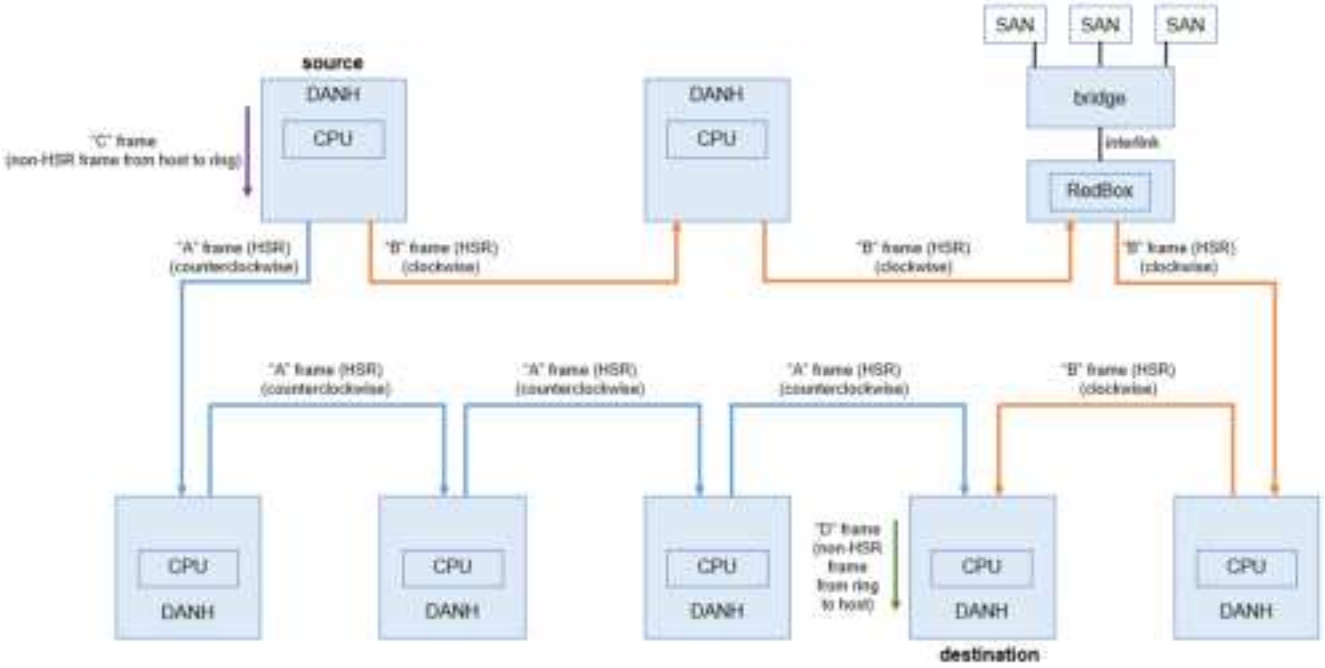

The two ring ports are referred to as port A and port B or the ring port pair. HSR ring ports are always a pair of an odd-numbered switch port and an even-numbered switch port. The pair of switch ports are hard coded, for example, port1-port2, port3-port4,…port27-port28. Port A has a lower port number than port B.

Nodes within the ring must be HSR-capable bridging nodes to avoid the use of dedicated bridges. Each node has two ports operated in parallel, which are doubly attached nodes with HSR protocol (DANHs). Singly attached nodes (SANs), such as laptops or printers, cannot be attached directly to the ring. Instead, the SANs are attached through a redundancy box (RedBox).

A source DANH sends a frame (a “C” frame) passed from its upper layers, prefixes it with an HSR tag to identify frame duplicates, and sends the frame over each port (“A” frame and “B” frame). A destination DANH receives, in the fault-free state, two identical frames from each port within a certain interval and removes the HSR tag of the first frame before passing it to its upper layers.

The simplest HSR topology contains two switches with two links between them; the ports connected to these two links serve as the HSR ring ports.

Configuring the HSR settings

Use the following commands to configure HSR settings:

```
config switch hsr settings
   set mac-da <0-255>
   set life-check-interval <2-60 seconds>
end
```

Variable	Description	Default
mac-da <0-255>	Specify the last 8 bits of the HSR supervision frame MAC destination address (DA).	0
life-check-interval <2-60 seconds>	Specify how often (in seconds) the HSR supervision frame is generated for each MAC address in the VDAN table.	2

Configuring the HSR ring

Use the following commands to configure an HSR ring. You can configure a maximum of two HSR rings (or a combination of one HSR ring and one PRP channel) on a FortiSwitch unit.

```
config switch hsr ring
   edit {1 | 2}
      set status {enable | disable}
      set ring-port-pair <physical_port_pair>
      set redbox-mode hsr-san
      set vlan-id <1-4094>
      set vlan-id-cos <0-7>
      set vlan-id-tagged {enable | disable}
      set hsr-internal-vlan <VLAN_ID>
   next
end
```

Variable	Description	Default
status {enable | disable}	Enable or disable this HSR ring.	disable
ring-port-pair <physical_port_pair>	Select which port A and port B pair to use for this HSR ring. Enter `set ring-port-pair ?` to see the available physical port pairs.	No default
redbox-mode hsr-san	HSR-SAN is currenly the only RedBox operation mode supported.	hsr-san
vlan-id <1-4094>	Enter the VLAN identifier of the supervision frame. This option is available only when `vlan-id-tagged` is set to `enable`.	1

Variable	Description	Default
vlan-id-cos <0-7>	Enter the class of service (CoS) value to be set in the VLAN tag of the supervision frame. This option is available only when `vlan-id-tagged` is set to `enable`.	0
vlan-id-tagged {enable | disable}	Enable or disable supervision frame VLAN ID tagging.	disable
hsr-internal-vlan <2-4094>	Assign all MAC addresses of this HSR ring to this internal VLAN ID. **NOTE:** If you are using an HSR ring and a PRP channel in your network, you need to change the default value so that each HSR ring and PRP channel is in a different internal VLAN.	No default

Clearing the HSR statistics

To delete the HSR statistics from the FortiSwitch unit:

```
diagnose switch hsr clear
```

Troubleshooting HSR

Use the following commands to troubleshoot your HSR configuration:

```
diagnose switch hsr {config | node-table | settings | stats | status | vdan-table}
```

Variable	Description
config	Display the current HSR configuration.
node-table	Display the HSR node table.
settings	Display the current HSR settings.
stats	Display the HSR statistics.
status	Display the current HSR status.
vdan-table	Display the HSR virtual doubly attached node (VDAN) table.

Configuration examples

Example 1: Simple HSR topology with three switches

The following figure shows a simple HSR topology that contains three switches:

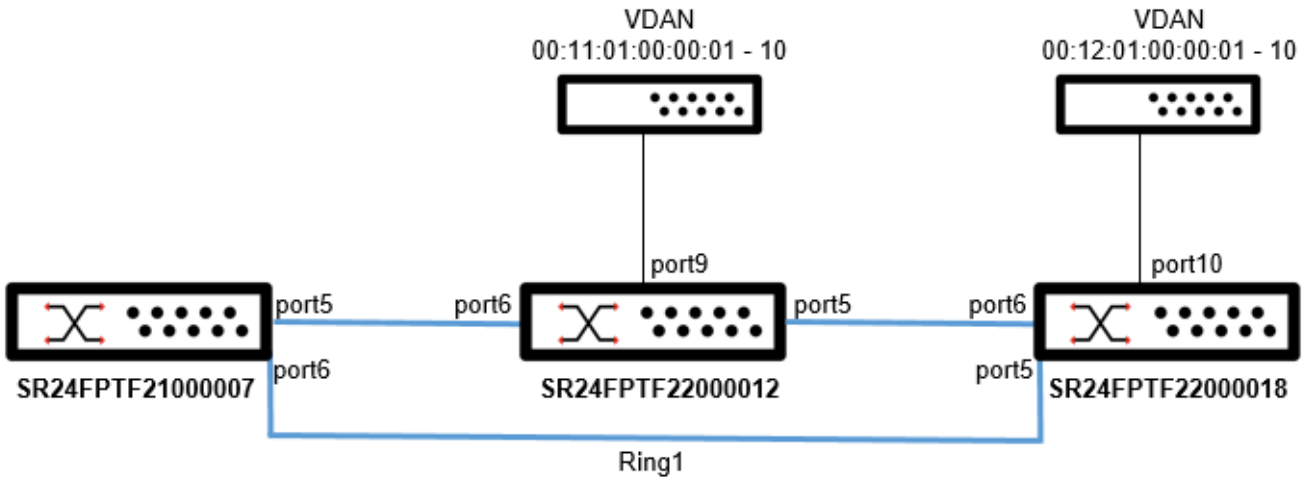

To configure SR24FPTF21000007:

```
config switch hsr ring
   edit 1
      set status enable
      set ring-port-pair port5-port6
      set hsr-internal-vlan 4093
   next
end
```

To configure SR24FPTF22000012:

```
config switch hsr ring
   edit 1
      set status enable
      set ring-port-pair port5-port6
      set hsr-internal-vlan 4093
   next
end

config switch interface
   edit "port9"
      set description "Interlink port"
      set native-vlan 4093
   next
end
```

To configure SR24FPTF22000018:

```
config switch hsr ring
   edit 1
      set status enable
      set ring-port-pair port5-port6
      set hsr-internal-vlan 4093
   next
end

config switch interface
```

```
   edit "port10"
      set description "Interlink port"
      set native-vlan 4093
   next
end
```

Example 2: Two HSR rings

The following topology contains two HSR rings. The QuadBox is a RedBox that interconnects two HSR rings. H1, H2, H3, H4, and H5 are intelligent electronic devices (IEDs).

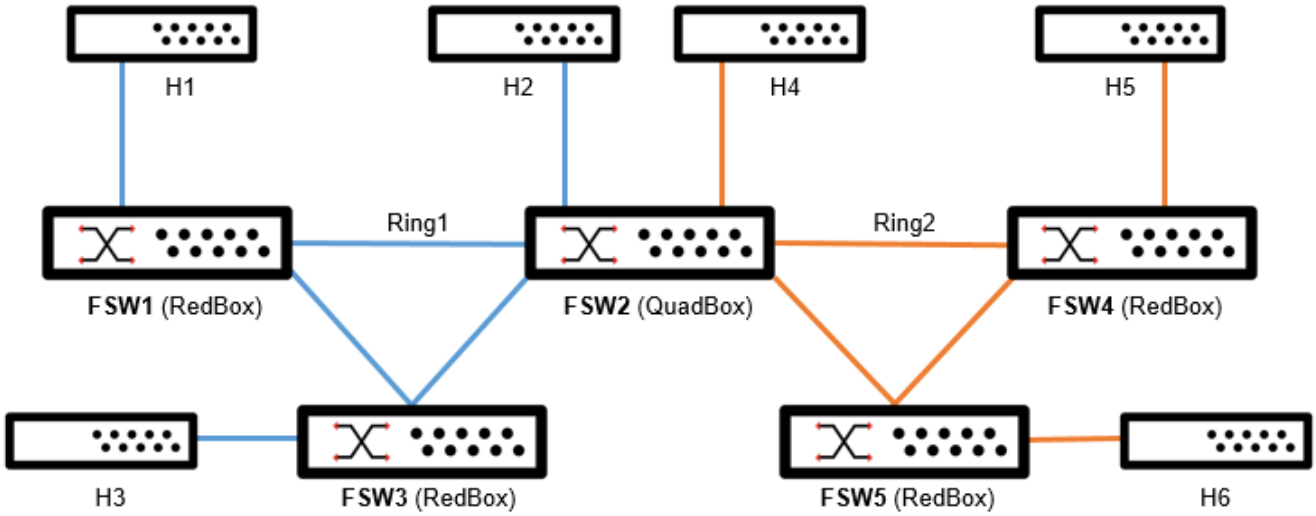

This topology has the following connections:

- For FSW1:
  - Interlink port (port1 connects to H1)
  - Ring ports (Ring1 uses port7 and port8)
- For FSW2:
  - Interlink port (port3 connects to H2; port4 connects to H4)
  - Ring ports (Ring1 uses port7 and port8; Ring2 uses port9 and port10)
- For FSW3:
  - Interlink port (port2 connects to H3)
  - Ring ports (Ring1 uses port7 and port8)
- For FSW4:
  - Interlink port (port5 connects to H5)
  - Ring ports (Ring2 uses port9 and port10)
- For FSW5:
  - Interlink port (port6 connects to H6)
  - Ring ports (Ring2 uses port9 and port10)

To configure FSW1:

```
config switch hsr ring
   edit 1
```

```
      set status enable
      set ring-port-pair port7-port8
      set hsr-internal-vlan 4093
   next
end

config switch interface
   edit "port1"
      set description "Interlink port FSW1 Ring1"
      set native-vlan 4093
   next
end
```

To configure FSW2:

```
config switch hsr ring
   edit 2
      set status enable
      set ring-port-pair port9-port10
      set hsr-internal-vlan 4093
   next
   edit 1
      set status enable
      set ring-port-pair port7-port8
      set hsr-internal-vlan 4092
   next
   end

config switch interface
   edit "port3"
      set description "Interlink port FSW2 Ring1"
      set native-vlan 4093
   next
   edit "port4"
      set description "Interlink port FSW2 Ring2"
      set native-vlan 4092
   next
end
```

To configure FSW3:

```
config switch hsr ring
   edit 1
      set status enable
      set ring-port-pair port7-port8
      set hsr-internal-vlan 4093
   next
end

config switch interface
   edit "port2"
      set description "Interlink port FSW3 Ring1"
      set native-vlan 4093
   next
end
```

To configure FSW4:

```
config switch hsr ring
   edit 2
      set status enable
      set ring-port-pair port7-port8
      set hsr-internal-vlan 4092
   next
end

config switch interface
   edit "port5"
      set description "Interlink port FSW4 Ring2"
      set native-vlan 4092
   next
end
```

To configure FSW5:

```
config switch hsr ring
   edit 2
      set status enable
      set ring-port-pair port7-port8
      set hsr-internal-vlan 4092
   next
end

config switch interface
   edit "port6"
      set description "Interlink port FSW5 Ring2"
      set native-vlan 4092
   next
end
```

HSR operation details and limitations

Review the following limitations before configuring HSR:

- Any physical port that is configured with the `hsr-internal-vlan` as the native VLAN becomes the interlink port; any nodes or hosts connected to this interlink port becomes VDANs.
- Interlink ports cannot be configured with allowed VLANs other than the `hsr-internal-vlan`.
- When the `hsr-internal-vlan` configuration is removed from the interlink port, the interlink port becomes a normal switch port.
- HSR supervision and data packets can be captured at ring ports with the help of a mirror and viewed for debugging purpose.
- When an HSR ring is enabled, a static trunk is created with the ring ports as members.
- The recommended maximum number of nodes in one HSR ring is 512. Nodes are all the DANH and VDAN devices that can be connected to the ring at the same time. When a table is full, no new MAC address is learned in the HSR node table or VDAN table until there is a free entry in the table.
- HSR ring ports can only be configured in layer-2 mode. LACP and STP are disabled.
- When HSR is enabled, the `mac-aging-interval` (under `config switch global`) is changed to 60 seconds. When HSR is disabled, the `mac-aging-interval` is changed to the default value, 300.

- Layer-3 routing is not supported in HSR.
- Interlinks ports cannot be multiple members of a LAG.
- RSTP+ is not supported in HSR.
- HSR only allows static-isl and no auto-vlan to the FortiGate device.
- You cannot share interlink ports between RedBoxes.
- The HSR traffic load cannot exceed 90% of the bandwidth of the gigabit Ethernet interface channels.

Parallel Redundancy Protocol

This section covers the following topics:

- Configuring the PRP settings on page 382
- Configuring the PRP channel on page 382
- Clearing the PRP statistics on page 383
- Troubleshooting PRP on page 383
- Configuration examples on page 384
- PRP operation details and limitations on page 386

The Parallel Redundancy Protocol (PRP) is defined in the international standard IEC 62439-3-2016 clause 4. PRP provides seamless communication with fault tolerance by duplicating every unicast frame sent in PRP networks. You can use PRP in different topologies such as ring, bus, or meshed.

A doubly attached node with PRP (DANP) is attached to two independent local area networks (LANs) with similar topologies, named LAN_A and LAN_B, which operate in parallel. A source DANP sends the same frame over both LANs, and a destination DANP receives it from both LANs within a certain time, consumes the first frame, and discards the duplicate. If a LAN fails, a DANP destination continues to operate with the frames from the other LAN.

Uncritical nodes, such as laptops or printers, are usually attached to just one LAN as single attached nodes (SANs). SANs that need to communicate with each other must be on the same LAN. If a critical node without PRP capability needs to communicate with all other nodes, it can be attached to a redundancy box (RedBox). The RedBox allows the single interface node to be attached to both networks and communicate with all other nodes. Because a node behind a RedBox appears to be a doubly attached node (DAN) to the other nodes, it is called a virtual DAN (VDAN). The RedBox itself is a DANP and acts as a proxy on behalf of its VDANs. Because both LAN A and LAB B must be independent, any connections among DANs and RedBoxes are not allowed.

The simplest PRP topology configuration is two switches with two links between them; the ports connected to these two links serve as PRP channel ports. PRP channel ports are always a pair of an odd-numbered switch port and an even-numbered switch port. The pair of switch ports are hard coded, for example, port1-port2, port3-port4,…port27-port28.

The following figure shows a PRP example of a general redundant network:

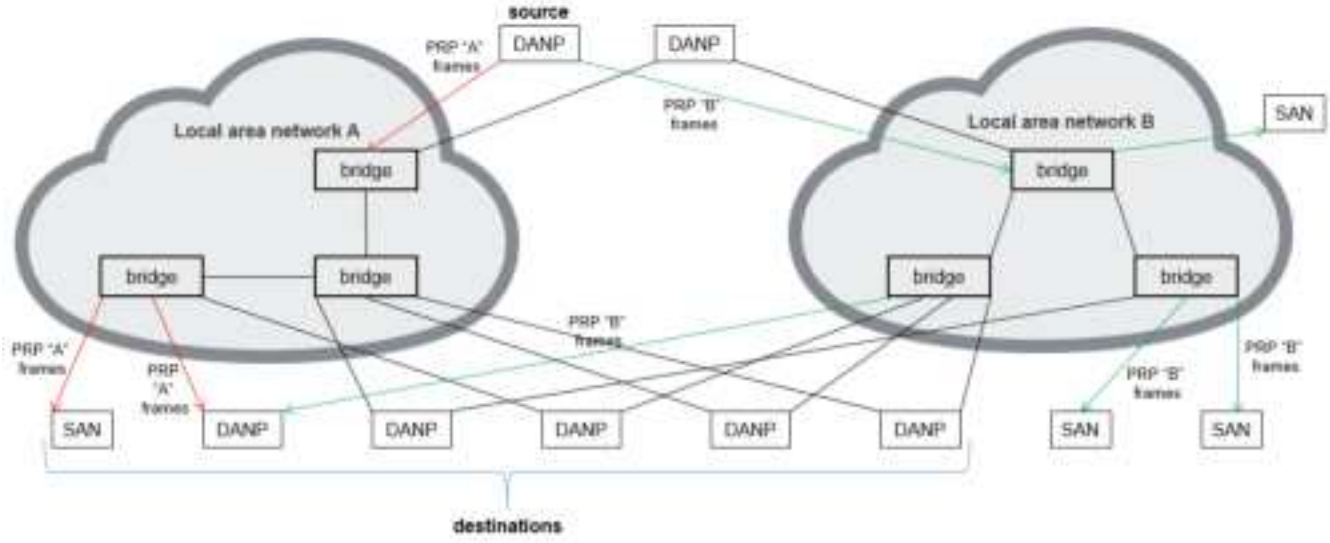

Configuring the PRP settings

Use the following commands to configure PRP settings:

```
config switch prp settings
   set mac-da <0-255>
   set life-check-interval <2-60 seconds>
end
```

Variable	Description	Default
mac-da <0-255>	Specify the last 8 bits of the PRP supervision frame MAC DA.	0
life-check-interval <2-60 seconds>	Specify how often (in seconds) the PRP supervision frame is generated for each MAC address in the VDAN table.	2

Configuring the PRP channel

Use the following commands to configure a PRP channel. You can configure a maximum of two PRP channels (or a combination of one HSR ring and one PRP channel) on an FSR-424F-POE.

```
config switch prp channel
   edit {1 | 2}
      set status {enable | disable}
      set channel-port-pair <physical_port_pair>
      set vlan-id <1-4094>
      set vlan-id-cos <0-7>
      set vlan-id-tagged {enable | disable}
      set prp-internal-vlan <2-4094>
   next
end
```

Variable	Description	Default
status {enable | disable}	Enable or disable this PRP channel.	disable
channel-port-pair <physical_port_pair>	Select which port A and port B pair to use for this PRP channel. Enter `set channel-port-pair ?` to see the available physical port pairs.	No default
vlan-id <1-4094>	Enter the VLAN identifier of the supervision frame. This option is available only when `vlan-id-tagged` is set to `enable`.	1
vlan-id-cos <0-7>	Enter the class of service (CoS) value to be set in the VLAN tag of the supervision frame. This option is available only when `vlan-id-tagged` is set to `enable`.	0
vlan-id-tagged {enable | disable}	Enable or disable supervision frame VLAN ID tagging.	disable
prp-internal-vlan <2-4094>	Assign all MAC addresses of this PRP channel to this internal VLAN ID. **NOTE:** If you are using an HSR ring and a PRP channel in your network, you need to change the default value so that each HSR ring and PRP channel is in a different internal VLAN.	No default

Clearing the PRP statistics

To delete the PRP statistics from the FortiSwitch unit:

```
diagnose switch prp clear
```

Troubleshooting PRP

Use the following commands to troubleshoot your PRP configuration:

```
diagnose switch prp {config | node-table | settings | stats | status | vdan-table}
```

Variable	Description
config	Display the current PRP configuration.
node-table	Display the PRP node table.
settings	Display the current PRP settings.
stats	Display the PRP statistics.
status	Display the current PRP status.
vdan-table	Display the PRP VDAN table.

Configuration examples

Example 1: Simple PRP topology

The following figure shows a simple PRP topology. H1 and H2 are IEDs.

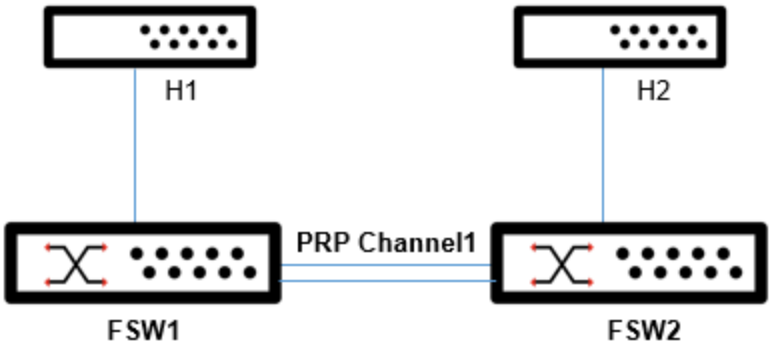

This topology has the following connections:

- For FSW1:
  - Interlink port (port1 connects to H1).
  - Channel ports(PRP Channel1 connects to port5 and port6)
- For FSW2:
  - Interlink port (port2 connects to H2).
  - Channel ports (PRP Channel1 connects to port5 and port6)

To configure FSW1:

```
config switch prp channel
   edit 1
      set status enable
      set channel-port-pair port3-port4
      set prp-internal-vlan 4093
   next
end

config switch interface
   edit "port1"
      set description "Interlink port FSW1 PRP1"
      set native-vlan 4093
   next
end
```

To configure FSW2:

```
config switch prp channel
   edit 1
      set status enable
      set channel-port-pair port3-port4
      set prp-internal-vlan 4093
```

```
   next
end

config switch interface
   edit "port2"
      set description "Interlink port FSW2 PRP1"
      set native-vlan 4093
   next
end
```

Example 2: Two PRP channels

The following topology contains two PRP channels. H1, H2, H3, and H4 are IEDs.

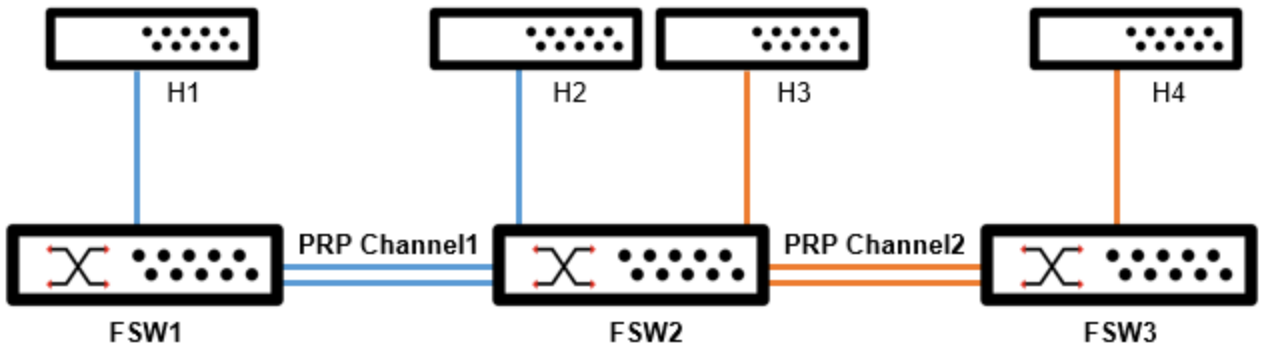

This topology has the following connections:

- For FSW1:
  - Interlink port (port1 connects to H1).
  - Channel ports (PRP Channel1 uses port5 and port6)
- For FSW2:
  - Interlink port (port2 connects to H2; port3 connects to H3).
  - Channel ports (PRP Channel1 uses port5 and port6; PRP Channel2 uses port7 and port8)
- For FSW3:
  - Interlink port (port4 connects to H4).
  - Channel ports(PRP Channel2 uses port7 and port8)

To configure FSW1:

```
config switch prp channel
   edit 1
      set status enable
      set channel-port-pair port5-port6
      set prp-internal-vlan 4093
   next
end

config switch interface
   edit "port1"
      set description "Interlink port FSW1 PRP1"
      set native-vlan 4093
```

```
   next
end
```

To configure FSW2:

```
config switch prp channel
   edit 1
      set status enable
      set channel-port-pair port5-port6
      set prp-internal-vlan 4093
   next
   edit 2
      set status enable
      set channel-port-pair port7-port9
      set prp-internal-vlan 4092
   next
end

config switch interface
   edit "port2"
      set description "Interlink port FSW2 PRP1"
      set native-vlan 4093
   next
   edit "port3"
      set description "Interlink port FSW2 PRP3"
      set native-vlan 4092
   next
end
```

To configure FSW3:

```
config switch prp channel
   edit 2
      set status enable
      set channel-port-pair port7-port8
      set prp-internal-vlan 4092
   next
end

config switch interface
   edit "port4"
      set description "Interlink port FSW3 PRP2"
      set native-vlan 4092
   next
   edit "port3"
      set description "Interlink port FSW2 PRP3"
      set native-vlan 4092
   next
end
```

PRP operation details and limitations

Review the following limitations before configuring PRP:

- Any physical port that is configured with the `prp-internal-vlan` as the native VLAN becomes the interlink port; any nodes or hosts connected to this interlink port becomes VDANs.
- Interlink ports cannot be configured with allowed VLANs other than the `prp-internal-vlan`.
- When the `prp-internal-vlan` configuration is removed from the interlink port, the interlink port becomes a normal switch port.
- PRP supervision and data packets can be captured at ring ports with the help of a mirror and viewed for debugging purpose.
- When a PRP channel is enabled, a static trunk is created with the channel ports as members.
- The recommended maximum number of nodes in one PRP channel is 512. Nodes are all the DANH and VDAN devices that can be connected to the ring at the same time. When a table is full, no new MAC address is learned in the PRP node table or VDAN table until there is a free entry in the table.
- PRP channel ports can only be configured in layer-2 mode. LACP and STP are disabled.
- When PRP is enabled, the `mac-aging-interval` (under `config switch global`) is changed to 60 seconds. When PRP is disabled, the `mac-aging-interval` is changed to the default value, 300.
- Layer-3 routing is not supported in PRP.
- Interlinks ports cannot be multiple members of a LAG.
- RSTP+ is not supported in PRP.
- PRP only allows static-isl and no auto-vlan to the FortiGate device.
- You cannot share interlink ports between RedBoxes.
- The PRP traffic load cannot exceed 90% of the bandwidth of the gigabit Ethernet interface channels.

Router

This section provides information on how to configure options related to routing protocols and packet forwarding:

- Config on page 388
- Diagnostic on page 471
- ARP table on page 475
- Monitor on page 476

Config

The following topics provide information about router configuration:

- Layer-3 routing in hardware on page 388
- Using layer-3 routing within an MCLAG on page 389
- Unicast reverse-path forwarding (uRPF) on page 404
- BGP routing on page 405
- IS-IS routing on page 424
- OSPF on page 429
- RIP on page 437
- Multicast on page 448
- Access lists on page 451
- Static and IPv6 static routing on page 453
- Link probes on page 457
- Virtual routing and forwarding on page 461
- Policy-based routing on page 464
- Key chains on page 468

Layer-3 routing in hardware

In FortiSwitchOS 3.3.0 and later, some FortiSwitch models support hardware-based layer-3 forwarding.

For FortiSwitch models that support Equal Cost Multi-Path (ECMP) (see the FortiSwitch Feature Matrix in the Fortinet Document Library), forwarding for all ECMP routes is performed in hardware.

For switch models that support hardware-based layer-3 forwarding but do not support ECMP, only one route to each destination will be hardware-forwarded. If you configure multiple routes to the same destination, you can configure a priority value for each route. Only the route with highest priority will be forwarded by the hardware. If no priority values are assigned to the routes, the most recently configured route is forwarded by the hardware.

Using layer-3 routing within an MCLAG

Starting in FortiSwitchOS 7.0.1, you can now use the Virtual Router Redundancy Protocol to make layer-3 routing in an MCLAG function as a single router.

Note:

- Only IPv4 addresses are supported.
- 250 switch virtual interfaces (SVIs) are supported.
- Both peer switches must be configured.
- Multicast (PIM) routing, policy-based routing (PBR), IS-IS routing, and RIP are not supported.

There are four use cases:

- One-tier MCLAG on page 389
- Two-tier MCLAG on page 390
- One-tier MCLAG with a southbound switch on page 392
- One-tier MCLAG without a northbound MCLAG trunk on page 393

One-tier MCLAG

To use layer-3 routing for a one-tier MCLAG, you can use a combination of VRRP with static or dynamic routing (BGP or OSPF).

The following figure shows the scenario with VRRP and BGP.

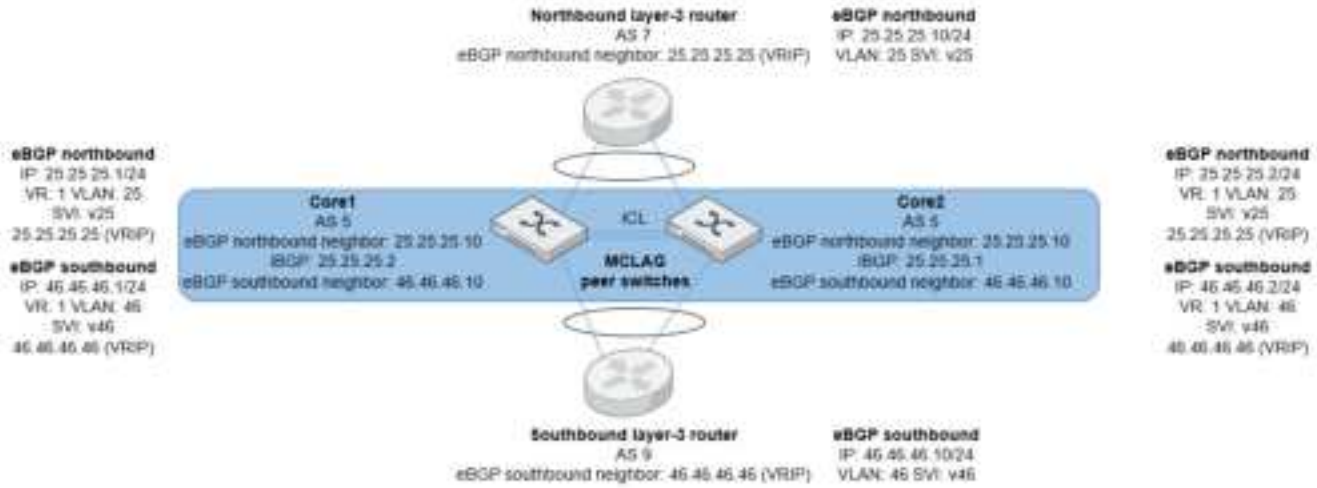

For a one-tier MCLAG topology:

- Core1 and Core2 are FortiSwitch units that form the MCLAG.
- Traffic flows between northbound and southbound routers through the MCLAG peer group. The two routers can be FortiSwitch units, but this is not mandatory

Using VRRP and BGP

Enable VRRP on the switch virtual interfaces (SVIs) towards the northbound and southbound neighboring routers on both MCLAG peers. The VRRP IP address is used as the next hop or BGP neighbor in the northbound and southbound neighboring routers.

Always enable `vrrp-virtual-mac` for VRRP. Layer-3 lookup for the VRRP virtual MAC address on the VRRP backup is enabled automatically. By virtue of MCLAG and trunk hashing, ingress packets on the VRRP backup MCLAG core are routed without crossing the ICL if the appropriate route is available.

Enable external BGP (eBGP) between the northbound router and the MCLAG VRRP IP address of the northbound SVI and between the southbound router and the MCLAG VRRP IP address of the southbound SVI. Because the eBGP neighbor is the VRRP IP address, the router establishes a connection with only the VRRP master. Enable `ebgp-enforce-multihop` and set `ebgp-multihop-ttl` to 3.

Use internal BGP (iBGP) between the MCLAG cores across the ICL. The routes from the eBGP sessions are advertised to iBGP, and the VRRP backup obtains the appropriate routes and stores them in its routing table and hardware. This achieves northbound-southbound layer-3 routing in an MCLAG topology, avoiding traffic across the ICL and using active-active forwarding across the MCLAG cores.

Using VRRP and OSPF

OSPF can also be used as the routing protocol between MCLAG peers and northbound/southbound routers. In this case, OSPF is also the IGP. It requires an active VRRP IP address in each MCLAG peer.

Use OSPF between the virtual router IP address and the router connecting the MCLAG core switches over an MCLAG link.

Always enable `vrrp-virtual-mac` for VRRP.

Configure two VRRP sessions on each SVI and configure the VRRP priorities so that there is a VRRP master on each MCLAG core.

The layer-3 lookup for the VRRP virtual MAC address is automatically enabled on the VRRP backup. Because of MCLAG and trunk hashing, ingress packets on the VRRP backup core are routed without crossing the ICL if an appropriate route is available.

The result of this topology is northbound-southbound layer-3 routing in the MCLAG topology without traffic crossing the ICL, and active-active forwarding is used across MCLAG cores.

Using VRRP, BGP (northbound), and OSPF (southbound)

- Start with the BGP configuration to configure MCLAG for nouthbound routing.
- Start with the OSPF configuration to configure MCLAG for sorthbound routing.
  - In the OSPF configuration, include the BGP subnet used for northbound routing in both the OSPF network and OSPF interface configuration.

Two-tier MCLAG

For layer-3 routing between MCLAG tiers, the configuration is similar for the tier-2 and tier-3 MCLAG peers. You can use a combination of VRRP with static or dynamic routing (BGP or OSPF). The following figure shows the scenario with VRRP and BGP.

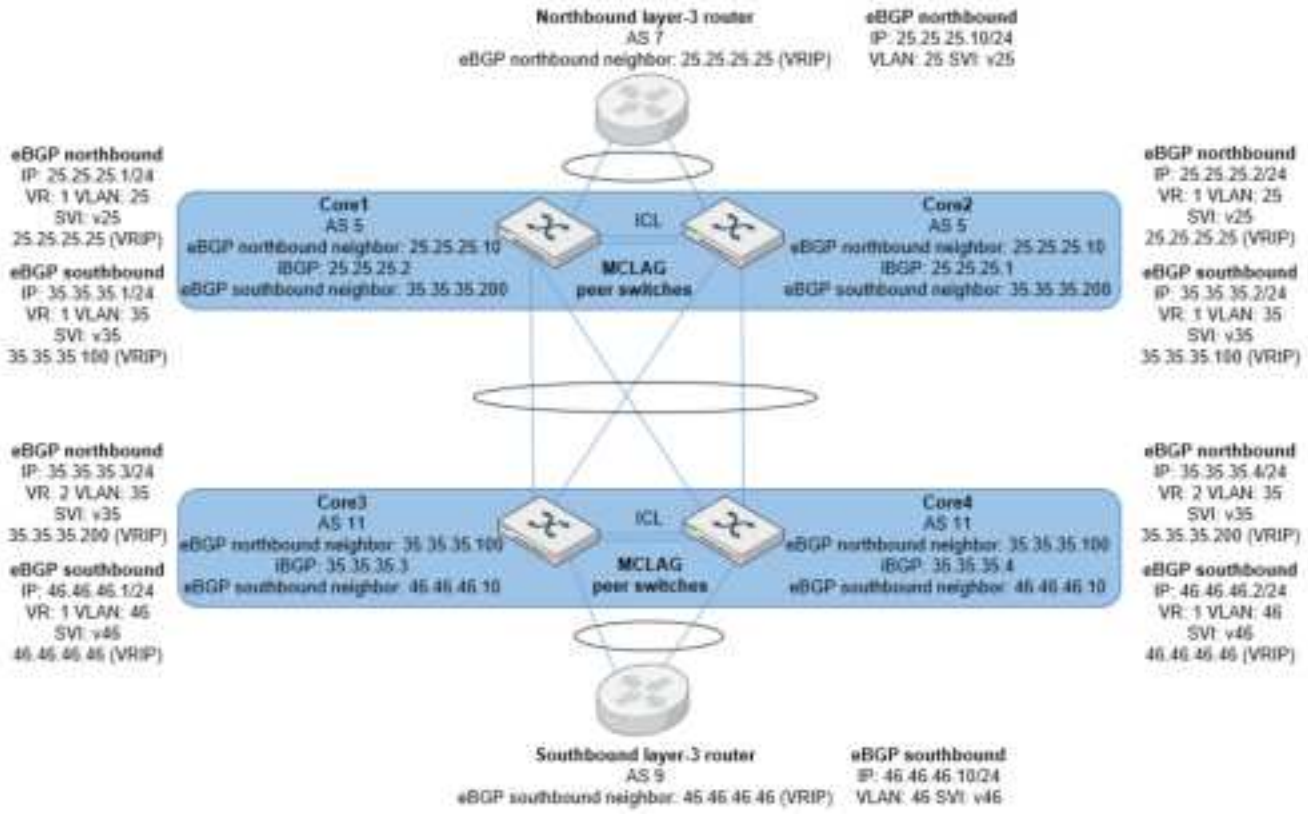

For a two-tier MCLAG topology:

- Core1 and Core2 are FortiSwitch units that form the tier-1 MCLAG. Core3 and Core4 are FortiSwitch units that form the tier-2 MCLAG.
- Traffic flows between northbound and southbound routers through the MCLAG peer groups. The two routers can be FortiSwitch units, but this is not mandatory.

Using VRRP and BGP

Each MCLAG tier has two VRRP sessions:

- One VRRP session is on the SVI that connects the router and the two core switches.
- One VRRP session is on the SVI subnet that is common between the pairs of MCLAG switches. For this subnet, the virtual router IP address belongs to the same subnet on both MCLAG pairs.

Each session has a different `vrip` value. Each session has a different virtual route identifier (VRID).

Configure eBGP for Core1, Core2, Core3, Core4, the northbound AS, and the southbound AS. You need to enable `ebgp-enforce-multihop` and set `ebgp-multihop-ttl` to 3.

Configure iBGP for Core1, Core2, Core3, and Core4.

When you configure VRRP, enable `vrrp-virtual-mac`.

Using VRRP and OSPF

Use OSPF between the virtual router IP address and the router connecting the MCLAG core switches over an MCLAG link and between the MCLAG tiers.

Always enable `vrrp-virtual-mac` for VRRP.

Configure two VRRP sessions on each SVI and configure the VRRP priorities so that there is a VRRP master on each MCLAG core.

The layer-3 lookup for the VRRP virtual MAC address is automatically enabled on the VRRP backup. Because of MCLAG and trunk hashing, ingress packets on the VRRP backup core are routed without crossing the ICL if an appropriate route is available.

The result of this topology is northbound-southbound layer-3 routing in the MCLAG topology without traffic crossing the ICL, and active-active forwarding is used across MCLAG cores.

One-tier MCLAG with a southbound switch

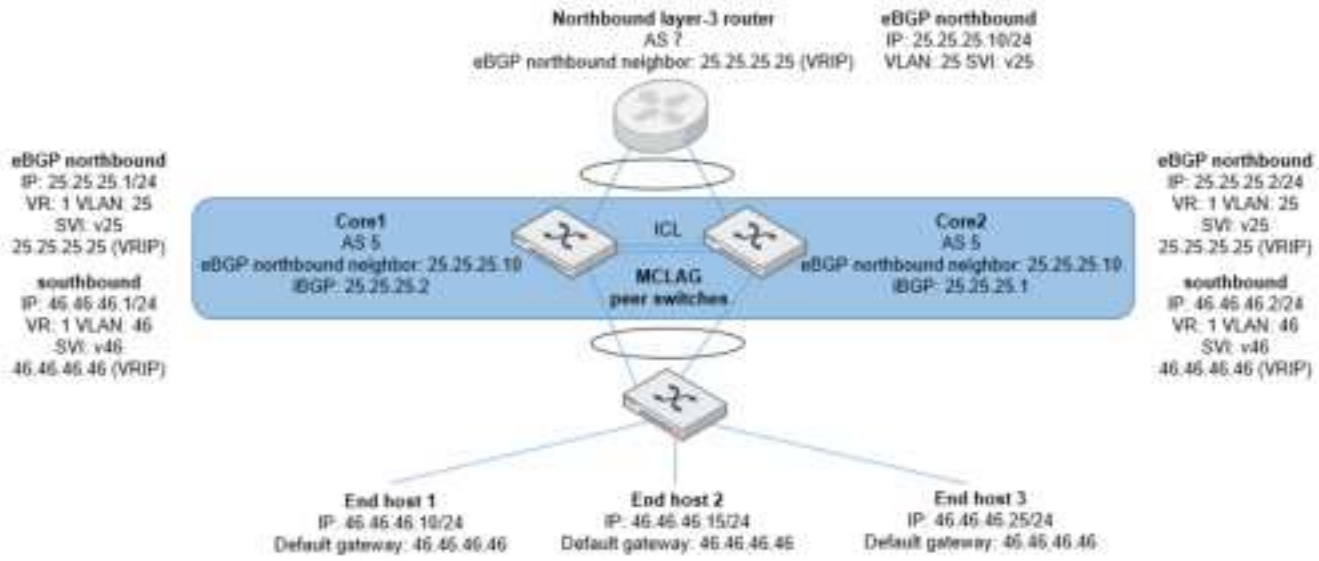

For this topology:

- Core1 and Core2 are FortiSwitch units that form the MCLAG.
- Traffic flows between the northbound router and the southbound hosts through the MCLAG peer group. The router can be a FortiSwitch unit, but this is not mandatory.
- The southbound switch or endpoint does not use eBGP with the MCLAG peer switches. The MCLAG SVI VRRP IP address is the default gateway for the endpoints.

One-tier MCLAG without a northbound MCLAG trunk

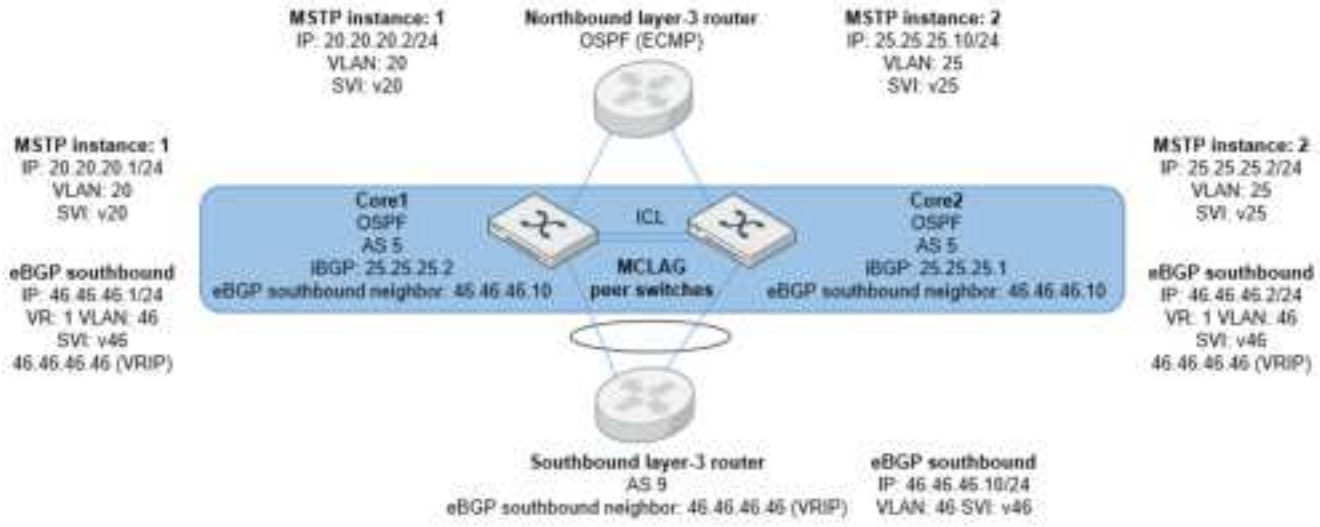

For this topology:

- Core1 and Core2 are FortiSwitch units that form the MCLAG.
- Traffic flows between northbound and southbound routers through the MCLAG peer group. The two routers can be FortiSwitch units, but this is not mandatory.
- The northbound router does not form an MCLAG trunk with the peer switches; instead, each link has its own layer-3 interface and MSTP instance. The northbound SVIs on the MCLAG peers do not need VRRP.
- Make certain that the two VLANs are on two different MSTP instances to avoid STP loops.

Using VRRP with static routing

Enable VRRP on the switch virtual interfaces (SVIs) towards the northbound and southbound neighboring routers on both MCLAG peers. The VRRP IP address is used as the next hop in the static routes in the northbound and southbound neighboring routers.

Configure static routes on both MCLAG peers pointing to the neighboring routers. In the case of tier-2 or tier-3 MCLAG, configure static routes on both MCLAG peers pointing to the VRRP IP address of the SVI on the adjacent MCLAG peers.

Always enable `vrrp-virtual-mac` for VRRP.

East-west traffic

For east-west traffic, where the eastbound router is connected to the east MCLAG and the westbound router is connected to the west MCLAG, traffic crosses the MCLAG ICL. Any routing protocol can be used between the routers and the FortiSwitch units; these routes can be redistributed to the FortiSwitch MCLAG peers using IGP (iBGP or OSPF).

Configuration example (BGP and VRRP)

Use the following steps to configure layer-3 routing in a one-tier MCLAG using BGP and VRRP:

1. Configure the trunks on page 394
2. Configure the layer-3 SVIs on page 395
3. Configure the layer-2 switch interfaces on page 396
4. Configure the layer-3 routing on page 397

Configure the trunks

To configure the northbound trunk on the northbound router:

```
config switch trunk
   edit "nb1"
      set mode lacp-active
      set members "port49" "port50"
   next
end
```

To configure the trunk for the FortiSwitch peer 1 (Core1):

```
config switch trunk
   edit "fsw2"
      set mode lacp-active
      set mclag-icl enable
      set members "port52" "port53"
   next
   edit "sb"
      set mode lacp-active
      set mclag enable
      set members "port26"
   next
   edit "nb"
      set mode lacp-active
      set mclag enable
      set members "port25"
   next
end
```

To configure the trunk for the FortiSwitch peer 2 (Core2):

```
config switch trunk
   edit "fsw1"
      set mode lacp-active
      set mclag-icl enable
   set members "port52" "port53"
   next
   edit "sb"
      set mode lacp-active
      set mclag enable
      set members "port15″
   next
   edit "nb"
      set mode lacp-active
```

```
      set mclag enable
      set members "port10"
   next
end
```

To configure the trunk on the southbound router:

```
config switch trunk
   edit "sb1"
      set mode lacp-active
      set members "port4" "port5"
   next
end
```

Configure the layer-3 SVIs

To configure the layer-3 SVI on the northbound router:

```
config system interface
   edit "nb1" <<<<< System interface used to connect to the MCLAG core VRRP IP address
      set ip 20.1.1.48 255.255.255.0
      set vlanid 20
   next
end
```

To configure the layer-3 SVI interfaces on FortiSwitch peer 1 (Core1):

```
config system interface
   edit "sb" <<<<<< connected to the southbound router
      set ip 100.1.1.21 255.255.255.0
      set vrrp-virtual-mac enable
      config vrrp
         edit 1
            set vrip 100.1.1.20
         next
      end
      set vlanid 100
   next
   edit "nb" <<<<<< connected to the northbound router
      set ip 20.1.1.11 255.255.255.0
      set vrrp-virtual-mac enable
      config vrrp
         edit 5
            set vrip 20.1.1.1
         next
      end
      set vlanid 20
   next
end
```

To configure the layer-3 SVI on the FortiSwitch peer 2 (Core2):

```
config system interface
   edit "sb" <<<<<< connected to the southbound router
      set ip 100.1.1.22 255.255.255.0
      set vrrp-virtual-mac enable
```

```
      config vrrp
         edit 1
            set vrip 100.1.1.20
         next
      end
      set vlanid 100
   next
   edit "nb" <<<<<< connected to the northbound router
      set ip 20.1.1.12 255.255.255.0
      set vrrp-virtual-mac enable
      config vrrp
         edit 5
            set vrip 20.1.1.1
         next
      end
      set vlanid 20
   next
end
```

To configure the layer-3 SVI on the southbound router:

```
config system interface
   edit "sb" <<<<<< connected to MCLAG core switches VRRP IP address
      set ip 100.1.1.10 255.255.255.0
      set vlanid 100
   next
end
```

Configure the layer-2 switch interfaces

To configure the layer-2 switch interfaces on the northbound router:

```
config switch interface
   edit "internal"
      set native-vlan 4094
      set allowed-vlans 20,4094
   next
   edit "nb1"
      set allowed-vlans 20
   next
end
```

To configure the layer-2 switch interfaces for FortiSwitch peer 1 (Core1):

```
config switch interface
   edit "internal"
      set native-vlan 4094
      set allowed-vlans 20,100,4094
   next
   edit "fsw2"
      set native-vlan 4094
      set allowed-vlans 1-4094
      set dhcp-snooping trusted
      set edge-port disabled
      set igmp-snooping-flood-reports enable
      set mcast-snooping-flood-traffic enable
```

```
   next
   edit "sb"
      set allowed-vlans 100
   next
   edit "nb"
      set allowed-vlans 20
   next
end
```

To configure the layer-2 switch interfaces for FortiSwitch peer 2 (Core2):

```
config switch interface
   edit "internal"
      set native-vlan 4094
      set allowed-vlans 20,100,4094
   next
   edit "fsw1"
      set native-vlan 4094
      set allowed-vlans 1-4094
      set dhcp-snooping trusted
      set edge-port disabled
      set igmp-snooping-flood-reports enable
      set mcast-snooping-flood-traffic enable
   next
   edit "sb"
      set allowed-vlans 100
   next
   edit "nb"
      set allowed-vlans 20
   next
end
```

To configure the layer-2 switch interfaces on the southbound router:

```
config switch interface
   edit "internal"
      set native-vlan 4094
      set allowed-vlans 100,4094
   next
   edit "sb1"
      set allowed-vlans 100
   next
end
```

Configure the layer-3 routing

To configure the routing for the northbound router:

```
config router bgp
   set as 7
   set router-id 20.1.1.48
   config neighbor
      edit "20.1.1.1" >>>>> eBGP to the MCLAG peer VRRP IP address
         set ebgp-enforce-multihop enable
         set ebgp-multihop-ttl 3
         set remote-as 5
```

```
      next
   end
```

To configure the routing for the FortiSwitch peer 1 (Core1):

```
config router bgp
   set as 5
   set router-id 100.1.1.21
   config neighbor
      edit "20.1.1.48" >>>>> eBGP to the northbound router
         set ebgp-enforce-multihop enable
         set ebgp-multihop-ttl 3
         set remote-as 7
      next
      edit "100.1.1.22" >>>>> iBGP to MCLAG peer
         set remote-as 5
      next
      edit "100.1.1.10" >>>>> eBGP to the southbound router
         set ebgp-enforce-multihop enable
         set ebgp-multihop-ttl 3
         set remote-as 9
      next
   end
```

To configure the routing the FortiSwitch peer 2 (Core2):

```
config router bgp
   set as 5
   set router-id 100.1.1.22
   config neighbor
      edit "20.1.1.48" >>>>> eBGP to the northbound router
         set ebgp-enforce-multihop enable
         set ebgp-multihop-ttl 3
         set remote-as 7
      next
      edit "100.1.1.21" >>>>> iBGP to the MCLAG peer
         set remote-as 5
      next
      edit "100.1.1.10" >>>>> eBGP to the southbound router
         set ebgp-enforce-multihop enable
         set ebgp-multihop-ttl 3
         set remote-as 9
      next
   end
```

To configure the routing for the southbound router:

```
config router bgp
   set as 9
   set router-id 100.1.1.10
   config neighbor
      edit "100.1.1.20" >>>>> eBGP to the MCLAG peer VRRP IP address
         set ebgp-enforce-multihop enable
         set ebgp-multihop-ttl 3
         set remote-as 5
      next
```

```
   end
```

Configuration example (OSPF and VRRP)

Use the following steps to configure layer-3 routing in a one-tier MCLAG using OSPF and VRRP:

1. Configure the trunks on page 399
2. Configure the layer-3 SVIs on page 400
3. Configure the layer-2 switch interfaces on page 401
4. Configure the layer-3 routing on page 403

Configure the trunks

To configure the northbound trunk on the northbound router:

```
config switch trunk
   edit "nb1"
      set mode lacp-active
      set members "port49" "port50"
   next
end
```

To configure the trunk for the FortiSwitch peer 1 (Core1):

```
config switch trunk
   edit "fsw2"
      set mode lacp-active
      set mclag-icl enable
      set members "port52" "port53"
   next
   edit "sb"
      set mode lacp-active
      set mclag enable
      set members "port26"
   next
   edit "nb"
      set mode lacp-active
      set mclag enable
      set members "port25"
   next
end
```

To configure the trunk for the FortiSwitch peer 2 (Core2):

```
config switch trunk
   edit "fsw1"
      set mode lacp-active
      set mclag-icl enable
      set members "port52" "port53"
   next
   edit "sb"
      set mode lacp-active
      set mclag enable
      set members "port15″
```

```
   next
   edit "nb"
      set mode lacp-active
      set mclag enable
      set members "port10"
   next
end
```

To configure the trunk on the southbound router:

```
config switch trunk
   edit "sb1"
      set mode lacp-active
      set members "port4" "port5"
   next
end
```

Configure the layer-3 SVIs

To configure the layer-3 SVI on the northbound router:

```
config system interface
   edit "nb1" <<<<< System interface used to connect to the MCLAG core VRRP IP address
      set ip 20.1.1.48 255.255.255.0
      set vlanid 20
   next
end
```

To configure the layer-3 SVI interfaces on FortiSwitch peer 1 (Core1):

```
config system interface
   edit "sb" <<<<<< connected to the southbound external router using VRRP
      set ip 100.1.1.21 255.255.255.0
      set vrrp-virtual-mac enable
      config vrrp
         edit 1
            set vrip 100.1.1.20
         next
         edit 3
            set priority 200
            set vrip 100.1.1.200
         next
      end
      set vlanid 100
   next
   edit "nb" <<<<<< connected to the northbound external router using VRRP
      set ip 20.1.1.11 255.255.255.0
      set vrrp-virtual-mac enable
      config vrrp
         edit 5
            set priority 200
            set vrip 20.1.1.1
         next
         edit 8
            set vrip 20.1.1.100
         next
```

```
      end
      set vlanid 20
   next
end
```

To configure the layer-3 SVI on the FortiSwitch peer 2 (Core2):

```
config system interface
   edit "sb" <<<<<< connected to the southbound external router using VRRP
      set ip 100.1.1.22 255.255.255.0
      set vrrp-virtual-mac enable
      config vrrp
         edit 1
            set priority 200
            set vrip 100.1.1.20
         next
         edit 3
            set vrip 100.1.1.200
         next
      end
      set vlanid 100
   next
   edit "nb" <<<<<< connected to the northbound external router using VRRP
      set ip 20.1.1.12 255.255.255.0
      set vrrp-virtual-mac enable
      config vrrp
         edit 5
            set vrip 20.1.1.1
         next
         edit 8
            set priority 200
            set vrip 20.1.1.100
         next
      end
      set vlanid 20
   next
end
```

To configure the layer-3 SVI on the southbound router:

```
config system interface
   edit "sb" <<<<< System interface used to connect to the MCLAG core VRRP IP address
      set ip 100.1.1.48 255.255.255.0
      set vlanid 100
   next
end
```

Configure the layer-2 switch interfaces

To configure the layer-2 switch interfaces on the northbound router:

```
config switch interface
   edit "internal"
      set native-vlan 4094
      set allowed-vlans 20,4094
   next
```

```
   edit "nb1"
      set allowed-vlans 20
   next
end
```

To configure the layer-2 switch interfaces for FortiSwitch peer 1 (Core1):

```
config switch interface
   edit "internal"
      set native-vlan 4094
      set allowed-vlans 20,100,4094
   next
   edit "fsw2"
      set native-vlan 4094
      set allowed-vlans 1-4094
      set dhcp-snooping trusted
      set edge-port disabled
      set igmp-snooping-flood-reports enable
      set mcast-snooping-flood-traffic enable
   next
   edit "sb"
      set allowed-vlans 100
   next
   edit "nb"
      set allowed-vlans 20
   next
end
```

To configure the layer-2 switch interfaces for FortiSwitch peer 2 (Core2):

```
config switch interface
   edit "internal"
      set native-vlan 4094
      set allowed-vlans 20,100,4094
   next
   edit "fsw1"
      set native-vlan 4094
      set allowed-vlans 1-4094
      set dhcp-snooping trusted
      set edge-port disabled
      set igmp-snooping-flood-reports enable
      set mcast-snooping-flood-traffic enable
   next
   edit "sb"
      set allowed-vlans 100
   next
   edit "nb"
      set allowed-vlans 20
   next
end
```

To configure the layer-2 switch interfaces on the southbound router:

```
config switch interface
   edit "internal"
      set native-vlan 4094
```

```
      set allowed-vlans 100,4094
   next
   edit "sb1"
      set allowed-vlans 100
   next
end
```

Configure the layer-3 routing

To configure the routing for the northbound router:

```
config router ospf
   set router-id 20.1.1.48
   config area
      edit 0.0.0.100
      next
   end
   config interface
      edit "nb1"
      next
   end
   config network
      edit 1 <<< connected to the MCLAG core
         set area 0.0.0.100
         set prefix 20.1.1.0 255.255.255.0
      next
   end
end
```

To configure the routing for the FortiSwitch peer 1 (Core1):

```
config router ospf
   set router-id 100.1.1.21
   config area
      edit 0.0.0.100
      next
   end
   config interface
      edit "sb" <<<<<< to the southbound router
      next
      edit "nb" <<<<<< to the northbound router
      next
   end
   config network
      edit 100 <<<<<< to the southbound router
         set area 0.0.0.100
         set prefix 100.1.1.0 255.255.255.0
      next
      edit 20 <<<<<< to the northbound router
         set area 0.0.0.100
         set prefix 20.1.1.0 255.255.255.0
      next
   end
end
```

To configure the routing the FortiSwitch peer 2 (Core2):

```
config router ospf
   set router-id 100.1.1.22
   config area
      edit 0.0.0.100
      next
   end
   config interface
      edit "sb"
      next
      edit "nb"
      next
   end
   config network
      edit 100 <<< to the southbound router
         set area 0.0.0.100
         set prefix 100.1.1.0 255.255.255.0
      next
      edit 20 <<< to the northbound router
         set area 0.0.0.100
         set prefix 20.1.1.0 255.255.255.0
      next
   end
end
```

To configure the routing for the southbound router:

```
config router ospf
   set router-id 100.1.1.48
   config area
      edit 0.0.0.100
      next
   end
   config interface
      edit "sb1"
      next
   end
   config network
      edit 1 <<< connected to the MCLAG core
         set area 0.0.0.100
         set prefix 100.1.1.0 255.255.255.0
      next
   end
end
```

Unicast reverse-path forwarding (uRPF)

RPF, also called anti-spoofing, prevents an IP packet from being forwarded if its source IP address does not belong to a locally attached subnet (local interface) or is not part of the routing between the FortiSwitch unit and another source (such as a static route, RIP, OSPF, or BGP).

In unicast RPF, the router not only looks up the destination information but it also looks up the source information to ensure that it exists. If no source is found, that packet is dropped because the router assumes it is an error or an attack on the network.

There are two uRPF modes:

- **Strict**—The packet must be received on the same interface that the router uses to forward the return packet. In this mode, asymmetric routing paths in the network might cause legitimate traffic to be dropped.
- **Loose**—The routing table must include the source IP address of the packet. If you disable the `src-check-allow-default` option, the packet is dropped if the source IP address is not found in the routing table. If you enable the `src-check-allow-default` option, the packet is allowed even if the source IP address is not found in the routing table, but the default route is found in the routing table.

By default, uRPF is disabled. You must enable it on each interface that you want protected.

```
config system interface
   edit <interface_name>
      set src-check {disable | loose | strict}
      set src-check-allow-default {enable | disable} // This option is available only when
           src-check is set to loose.
   end
```

BGP routing

You must have an advanced features license to use BGP routing.

Border Gateway Protocol (BGP) contains two distinct subsets: internal BGP (iBGP) and external BGP (eBGP). iBGP is intended for use within your own networks. eBGP is used to connect many different networks together and is the main routing protocol for the Internet backbone. FortiSwitch units support iBGP, and eBGP only for communities.

BGP was first used in 1989. The current version, BGP-4, was released in 1995 and is defined in RFC 1771. That RFC has since been replaced by RFC 4271. The main benefits of BGP-4 are classless inter-domain routing and aggregate routes. BGP is the only routing protocol to use TCP for a transport protocol. Other routing protocols use UDP.

BGP makes routing decisions based on path, network policies, and rulesets instead of the hop-count metric as RIP does, or cost-factor metrics as OSPF does.

BGP-4+ supports IPv6. It was introduced in RFC 2858 and RFC 2545.

BGP is the routing protocol used on the Internet. It was designed to replace the old Exterior Gateway Protocol (EGP) which had been around since 1982, and was very limited. BGP enabled more networks to take part in the Internet backbone to effectively decentralize it and make the Internet more robust, and less dependent on a single ISP or backbone network.

This section covers the following topics:

- Parts and terminology of BGP on page 406
- How BGP works on page 414
- Troubleshooting BGP on page 416
- Configuring BGP on page 419
- Sample configuration on page 421

Parts and terminology of BGP

In a BGP network, there are some terms that need to be explained before going ahead. Some parts of BGP are not explained here because they are common to other dynamic routing protocols. When determining your network topology, note that the number of available or supported routes is not set by the configuration but depends on the available memory on the FortiSwitch units.

This section covers the following topics:

- BGP and IPv6 on page 406
- Role of routers in BGP networks on page 406
- Speaker routers on page 406
- Peer routers or neighbors on page 406
- Route reflectors on page 409
- Confederations on page 410
- Network Layer Reachability Information on page 410
- BGP attributes on page 410
- AS_PATH on page 411
- MULTI_EXIT_DESC on page 412
- COMMUNITY on page 412
- NEXT_HOP on page 412
- ATOMIC_AGGREGATE on page 413
- ORIGIN on page 413

BGP and IPv6

FortiSwitch units support IPv6 over BGP using the same `config router bgp` CLI command as IPv4 but different subcommands.

The main CLI keywords have IPv6 equivalents that are identified by the “6” on the end of the keyword, such as `config network6` or `set allowas-in6`. For more information about IPv6 BGP keywords, see the *FortiSwitchOS CLI Reference*.

Role of routers in BGP networks

Dynamic routing has a number of different roles that routers can fill. BGP has a number of custom roles that routers can fill. These include speaker routers, peer routers or neighbors, and route reflectors.

Speaker routers

Any router that is configured for BGP is considered a BGP speaker. This means that a speaker router advertises BGP routes to its peers.

Any routers on the network that are not speaker routers are not treated as BGP routers.

Peer routers or neighbors

In a BGP network, all neighboring BGP routers or peer routers are routers that are connected to a FortiSwitch unit. A FortiSwitch unit learns about all other routers through these peers.

You need to manually configure BGP peers on a FortiSwitch unit as neighbors. Otherwise, these routers are not seen as peers but simply as other routers on the network that do not support BGP. Optionally, you can use MD5 authentication to password-protect BGP sessions with those neighbors (see RFC 2385).

You can configure up to 1000 BGP neighbors on a FortiSwitch unit. You can clear all or some BGP neighbor connections (sessions), using the `execute router clear bgp` CLI command.

For example, if you have 10 routes in the BGP routing table and you want to clear the specific route to IP address 10.10.10.1, enter the following CLI command:

```
execute router clear bgp ip 10.10.10.1
```

To remove all routes for autonomous system (AS) number 650001, enter the following CLI command:

```
execute router clear bgp as 650001
```

To remove route flap dampening information for the 10.10.0.0/16 subnet, enter the following CLI command:

```
execute router clear bgp dampening 10.10.0.0/16
```

In the following diagram, Router A is directly connected to five other routers in a network that contains 12 routers. These routers (the ones in the blue circle) are Router A's peers or neighbors.

Router A and its five peer routers

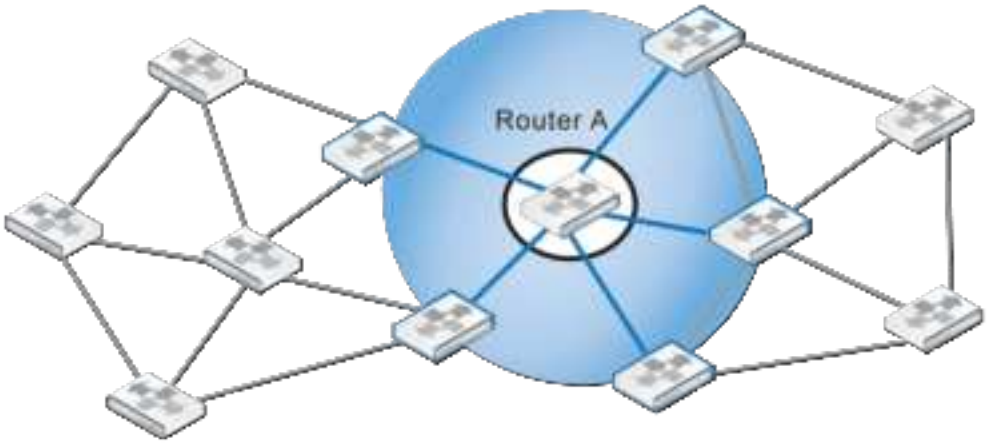

As a minimum, when configuring BGP neighbors, you must enter their IP address and the AS number (remote-as). This is all of the information the GUI allows you to enter for a neighbor.

The following BGP commands are related to neighbors:

```
config router bgp
   config neighbor
      edit "<IPv4_IPv6_address>"
         set advertisement-interval <0-600>
         set allowas-in-enable {disable | enable}
            set allowas-in <1-10>
         set allowas-in-enable6 {disable | enable}
            set allowas-in6 <1-10>
         set attribute-unchanged {as-path | MED | next-hop}
         set attribute-unchanged6 {as-path | MED | next-hop}
         set activate {disable | enable}
         set activate6 {disable | enable}
         set bfd {disable | enable}
```

```
        set capability-dynamic {disable | enable}
        set capability-orf {both | none | receive | send}
        set capability-orf6 {both | none | receive | send}
        set capability-default-originate {disable | enable}
        set capability-default-originate6 {disable | enable}
        set dont-capability-negotiate {disable | enable}
        set ebgp-enforce-multihop {disable | enable}
           set ebgp-multihop-ttl <1-255>
           set ebgp-ttl-security-hops <1-254>
        set next-hop-self {disable | enable}
        set next-hop-self6 {disable | enable}
        set override-capability {disable | enable}
        set passive {disable | enable}
        set remove-private-as {disable | enable}
        set remove-private-as6 {disable | enable}
        set route-reflector-client {disable | enable}
        set route-reflector-client6 {disable | enable}
        set route-server-client {disable | enable}
        set route-server-client6 {disable | enable}
        set shutdown {disable | enable}
        set soft-reconfiguration {disable | enable}
        set soft-reconfiguration6 {disable | enable}
        set as-override {disable | enable}
        set as-override6 {disable | enable}
        set strict-capability-match {disable | enable}
        set description <string>
        set distribute-list-in <string>
        set distribute-list-in6 <string>
        set distribute-list-out <string>
        set distribute-list-out6 <string>
        set filter-list-in <string>
        set filter-list-in6 <string>
        set filter-list-out <string>
        set filter-list-out6 <string>
        set interface <interface_name>
        set maximum-prefix <1-4294967295>
        set maximum-prefix6 <1-4294967295>
        set prefix-list-in <string>
        set prefix-list-in6 <string>
        set prefix-list-out <string>
        set prefix-list-out6 <string>
        set remote-as <MANDATORY_1-4294967295>
        set route-map-in <string>
        set route-map-in6 <string>
        set route-map-out <string>
        set route-map-out6 <string>
        set send-community {both | disable | extended | standard}
        set send-community6 {both | disable | extended | standard}
        set keep-alive-timer <0-65535>
        set holdtime-timer <0, 3-65535>
        set connect-timer <0-65535>
        set unsuppress-map <string>
        set unsuppress-map6 <string>
        set update-source {interface_name}
        set weight <0-65535>
    end
end
```

```
end
```

Route reflectors

Route reflectors (RRs) in iBGP concentrate route updates so other routers only need to talk to the RRs to get all of the updates. This results in smaller routing tables, fewer connections between routers, faster responses to network topology changes, and less administration bandwidth. iBGP RRs are defined in RFC 1966.

In an iBGP RR configuration, the AS is divided into different clusters that each include client and reflector routers. The client routers supply the reflector routers with the client's route updates. The reflectors pass this information along to other RRs and border routers. Only the reflectors need to be configured, not the clients, because the clients find the closest reflector and communicate with it automatically. The reflectors communicate with each other as peers. A FortiSwitch unit can be configured as either reflectors or clients.

Because RRs are processing more than the client routers, the reflectors should have more resources to handle the extra workload.

Smaller networks running iBGP typically do not require RRs. However, RRs are a useful feature for large companies, where their AS may include 100 routers or more. For example, a full mesh 20 router configuration within an AS, there would have to be 190 unique iBGP sessions just for routing updates within the AS. The number of sessions jumps to 435 sessions for just 30 routers, or 4950 sessions for 100 routers. Based on these numbers, updating this many sessions will quickly consume the limited bandwidth and processing resources of the routers involved.

The following diagram illustrates how RRs can improve the situation when only six routers are involved. The AS without RRs requires 15 sessions between the routers. In the AS with RRs, the two RRs receive route updates from the reflector clients (unlabeled routers in the diagram) in their cluster, as well as other RRs, and pass them on to the border router. The RR configuration requires only six sessions. This example shows a reduction of 60% for the number of required sessions.

Required sessions within an AS with and without RRs

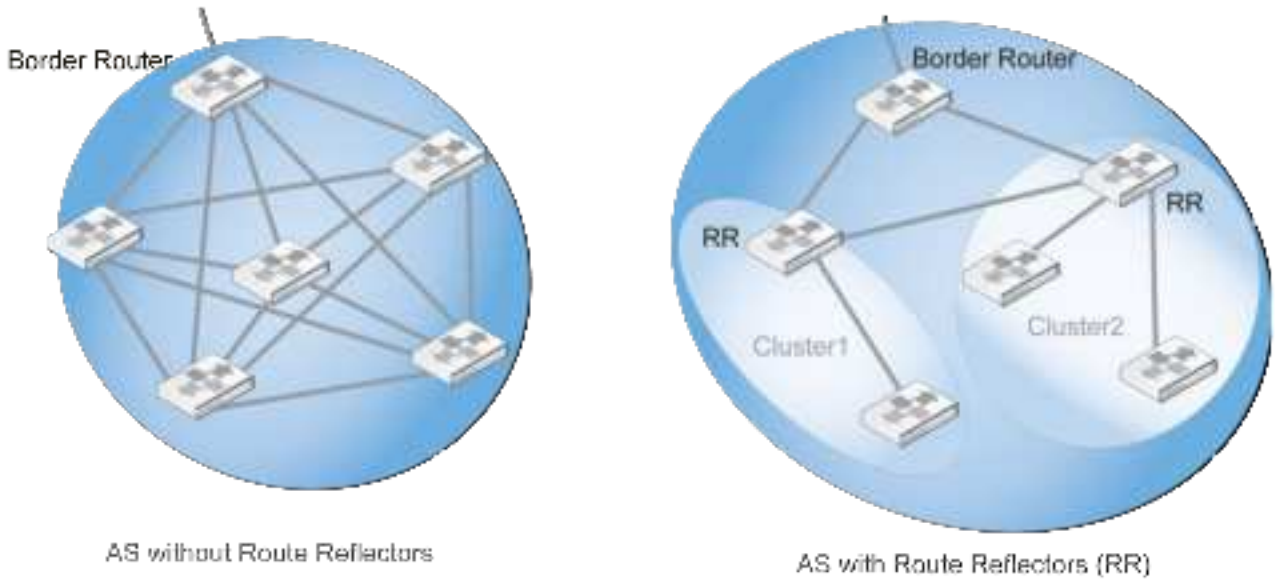

The iBGP commands related to RRs include:

```
config router bgp
   config neighbor
```

```
      edit "<IPv4_IPv6_address>"
         set route-reflector-client {disable | enable}
         set route-reflector-client6 {disable | enable}
   end
end
```

Confederations

Confederations were introduced to reduce the number of iBGP advertisements on a segment of the network and reduce the size of the routing tables. Confederations essentially break up an AS into smaller units. Confederations are defined in RFC 3065 and RFC 1965.

Within a confederation, all routers communicate with each other in a full mesh arrangement. Communications between confederations is more like inter-AS communications because many of the attributes are changed as they would be for BGP communications leaving the AS, or eBGP.

Confederations are useful when merging autonomous systems. Each AS being merged can easily become a confederation, which requires few changes. Any additional permanent changes can then be implemented over time, as required. After merging, if the border router becomes a route reflector, then each confederation only needs to communicate with one other router instead of five others.

Confederations and RRs perform similar functions: they both sub-divide large autonomous systems for more efficient operation. They differ in that route reflector clusters can include routers that are not members of a cluster, whereas routers in a confederation must belong to that confederation. Also, confederations place their confederation numbers in the AS_PATH attribute, making it easier to trace.

NOTE: While confederations essentially create sub-autonomous systems, all the confederations within an AS appear as a single AS to external autonomous systems.

Confederation related BGP commands include the following:

```
config router bgp
   set confederation-identifier <peerid_integer>
end
```

Network Layer Reachability Information

Network Layer Reachability Information (NLRI) is unique to BGP-4. It is sent as part of the update messages sent between BGP routers and contains information necessary to supernet, or aggregate route, information. The NLRI includes the length and prefix that, when combined, are the address of the aggregated routes referred to.

There is only one NLRI entry per BGP update message.

BGP attributes

Each route in a BGP network has a set of attributes associated with it. These attributes define the route and are modified, as required, along the route.

BGP can work well with mostly default settings, but if you're going to change settings you need to understand the roles of each attribute and how they affect those settings.

The BGP attributes include the ones listed in the following table.

Attribute	Description
AS_PATH	A list of autonomous systems a route has passed through. For more information, see AS_PATH on page 411.
MULTI_EXIT_DESC (MED)	Which router to use to exit an AS with more than one external connection. For more information, see MULTI_EXIT_DESC on page 412.
COMMUNITY	Used to apply attributes to a group of routes. For more information, see COMMUNITY on page 412.
NEXT_HOP	Where the IP packets should be forwarded to, like a gateway in static routing. For more information, see NEXT_HOP on page 412.
ATOMIC_AGGREGATE	Used when routes have been summarized to tell downstream routers not to de-aggregate the route. For more information, see ATOMIC_AGGREGATE on page 413.
ORIGIN	Used to determine if the route is from the local AS or not. For more information, see ORIGIN on page 413.
LOCAL_PREF	Used only within an AS to select the best route to a location (like MED).

Inbound policies on FortiSwitch units can change the NEXT-HOP, LOCAL-PREF, MED, and AS-PATH attributes of an internal BGP (iBGP) route for its local route selection purposes. However, outbound policies on the device cannot affect these attributes.

AS_PATH

AS_PATH is the BGP attribute that keeps track of each AS that a route advertisement has passed through. AS_PATH is used by confederations and by external BGP (eBGP) to help prevent routing loops. A router knows there is a loop if it receives an AS_PATH with that router's AS in it. The diagram shows the route between Router A and Router B. The AS_PATH from A to B would read 701,702,703 for each AS that the route passes through.

As of the beginning of 2010, the industry upgraded from 2-byte to 4-byte AS_PATHs. This upgrade was due to the imminent exhaustion of 2-byte AS_PATH numbers. FortiOS supports 4-byte AS_PATHs in its BGP implementation.

AS_PATH of 701,702, 703 between routers A and B

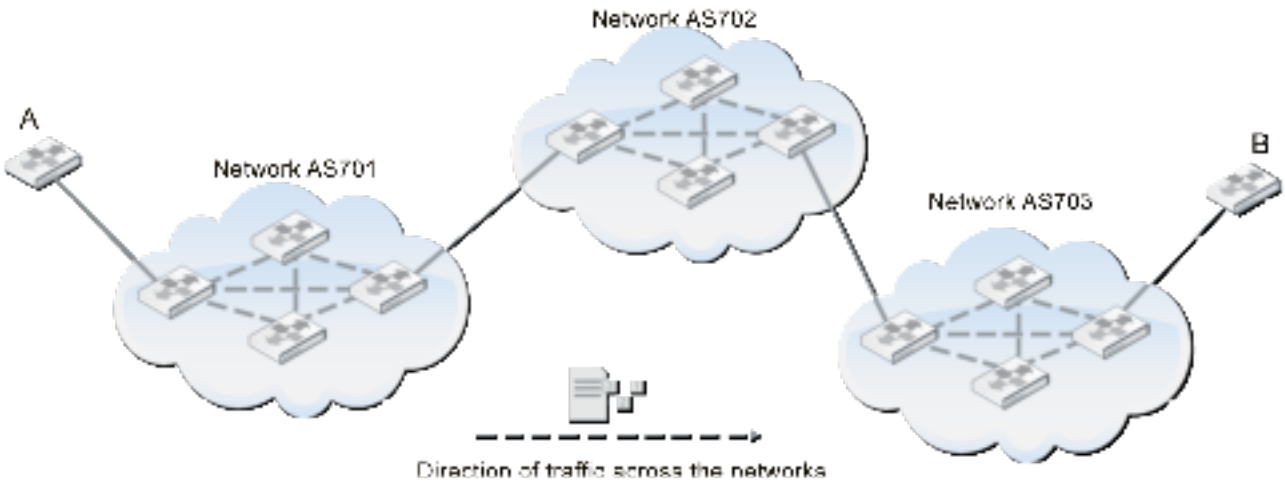

The BGP commands related to AS_PATH include the following:

```
config router bgp
   set bestpath-as-path-ignore {enable | disable}
end
```

MULTI_EXIT_DESC

BGP AS systems can have one or more routers that connect them to other autonomous systems. For autonomous systems with more than one connecting router, the Multi-Exit Discriminator (MED) lists which router is best to use when leaving the AS. The MED is based on attributes, such as delay. It is a recommendation only, as some networks may have different priorities.

BGP updates advertise the best path to a destination network. When a FortiSwitch unit receives a BGP update, the FortiSwitch unit examines the MED attribute of potential routes to determine the best path to a destination network before recording the path in the local FortiSwitch routing table.

FortiSwitch units have the option to treat any routes without an MED attribute as the worst possible routing choice. This can be useful because a lack of MED information is a lack of routing information, which can be suspicious as a possible hacking attempt or an attack on the network. At best, it signifies an unreliable route to select.

The BGP commands related to MED include the following:

```
config router bgp
   set always-compare-med {enable | disable}
   set bestpath-med-confed {enable | disable}
   set bestpath-med-missing-as-worst {enable | disable}
   set deterministic-med {enable | disable}
   config neighbor
      edit "<IPv4_IPv6_address>"
         set attribute-unchanged [as-path] [med] [next-hop]
         set attribute-unchanged6 {as-path | MED | next-hop}
      end
   end
end
```

COMMUNITY

A community is a group of routes that have the same routing policies applied to them. This saves time and resources. A community is defined by the COMMUNITY attribute of a BGP route.

A FortiSwitch unit can set the COMMUNITY attribute of a route to assign the route to predefined paths (see RFC 1997). The FortiSwitch unit can examine the COMMUNITY attribute of learned routes to perform local filtering and/or redistribution.

The BGP commands related to COMMUNITY include the following:

```
config router bgp
   set send-community {both | disable | extended | standard}
   set send-community6 {both | disable | extended | standard}
end
```

NEXT_HOP

The NEXT_HOP attribute says what IP address the packets should be forwarded to next. Each time the route is advertised, this value is updated. The NEXT_HOP attribute is much like a gateway in static routing.

FortiSwitch units allow you to change the advertising of the FortiSwitch unit's IP address (instead of the neighbor's IP address) in the NEXT_HOP information that is sent to iBGP peers. This is changed with the `config neighbor`, `set next-hop-self` command.

The BGP commands related to NEXT_HOP include the following:

```
config router bgp
   config neighbor
      edit "<IPv4_IPv6_address>"
         set attribute-unchanged [as-path] [med] [next-hop]
         set attribute-unchanged6 {as-path | MED | next-hop}
         set next-hop-self {enable | disable}
         set next-hop-self6 {disable | enable}
      next
   end
end
```

ATOMIC_AGGREGATE

The ATOMIC_AGGREGATE attribute is used when routes have been summarized. It indicates which AS and which router summarize the routes. It also tells downstream routers not to de-aggregate the route. Summarized routes are routes with similar information that have been combined, or aggregated, into one route that is easier to send in updates for. When it reaches its destination, the summarized routes are split back up into the individual routes.

The FortiSwitch unit does not specifically set this attribute in the BGP router command, but it is used in the route map command.

The CLI commands related to ATOMIC_AGGREGATE include the following:

```
config router route-map
   edit <route_map_name>
      set protocol bgp
      config rule
         edit <route_map_rule_id>
            set set-aggregator-as <id_integer>
            set set-aggregator-ip <address_ipv4>
            set set-atomic-aggregate {enable | disable}
         end
      end
   end
```

ORIGIN

The ORIGIN attribute records where the route came from. The options can be iBGP, eBGP, or incomplete. This information is important because internal routes (iBGP) are, by default, higher priority than external routes (eBGP). However, incomplete ORIGINs are the lowest priority of the three.

The CLI commands related to ORIGIN include the following:

```
config router route-map
   edit <route_map_name>
      set protocol bgp
      config rule
         edit <route_map_rule_id>
            set match-origin {egp | igp | incomplete | none}
         end
      end
   end
```

How BGP works

BGP is a link-state routing protocol and keeps link-state information about the status of each network link it has connected. A BGP router receives information from its peer routers that have been defined as neighbors. BGP routers listen for updates from these configured neighboring routers on TCP port 179.

A BGP router is a finite state machine with six various states for each connection. As two BGP routers discover each other and establish a connection, they go from the idle state and through the various states until they reach the established state. An error can cause the connection to drop and the state of the router to reset to either active or idle. These errors can be caused by TCP port 179 not being open, a random TCP port above port 1023 not being open, the peer address being incorrect, or the AS number being incorrect.

When BGP routers start a connection, they negotiate which (if any) optional features will be used, such as multiprotocol extensions, that can include IPv6 and VPNs.

This section covers the following topics:

- iBGP versus eBGP on page 414
- BGP path determination: Which route to use on page 414
- Decision phase 1 on page 415
- Decision phase 2 on page 416
- Decision phase 3 on page 416
- Aggregate routes and addresses on page 416

iBGP versus eBGP

When you read about BGP, you often see eBGP or iBGP mentioned. These are both BGP routing, but BGP used in different roles. eBGP involves packets crossing multiple autonomous systems, and iBGP involves packets that stay within a single AS. For example, the AS_PATH attribute is only useful for eBGP where routes pass through multiple autonomous systems.

These two modes are important because some features of BGP are used only for one of eBGP or iBGP. For example, both confederations and RRs are used only in iBGP.

FortiSwitch units have some commands that are specific to eBGP, including the following:

- automatically resetting the session information to external peers if the connection goes down:`set fast-external-failover {enable | disable}`
- setting an administrative distance for all routes learned from external peers (you must also configure local and internal distances if this is set):`set distance-external <distance_integer>`
- enforcing eBGP multihops and their TTL (number of hops): `set ebgp-enforce-multihop {enable | disable}` and `set ebgp-multihop-ttl <seconds_integer>`

BGP path determination: Which route to use

Firstly, recall that the number of available or supported routes is not set by the configuration but depends on the available memory on the FortiSwitch unit. All learned routes and their attributes come into the BGP router in raw form. Before routes are installed in the routing table or are advertised to other routers, three levels of decisions must be made.

The three phases of BGP best path determination do not change. However, some manufacturers have added more information to the process, such as Cisco's WEIGHT attribute, to allow an administrator to force one route's selection over another.

There is one Adj-RIB-IN and Adj-RIB-OUT for each configured neighbor. They are updated when the FortiSwitch unit receives BGP updates or when the FortiSwitch unit sends out BGP updates.

The three phases of a BGP routing decision

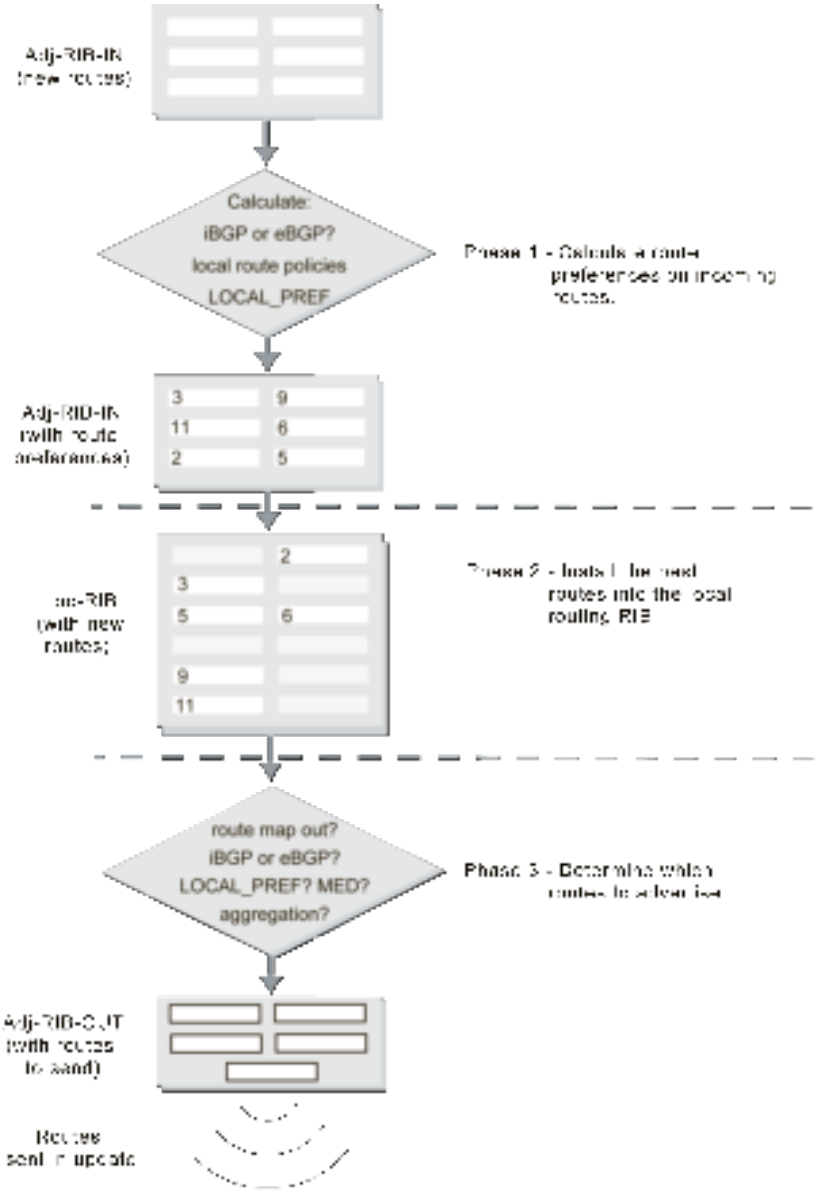

Decision phase 1

At this phase, the decision is to calculate how preferred each route and its NRLI are the Adjacent Routing Information Base Incoming (Adj-RIBs-In) compared to the other routes. For internal routes (iBGP), policy information or LOCAL_PREF is used. For external peer learned routes, it is based strictly on policy. These rules set up a list of which routes are most preferred going into Phase 2.

Decision phase 2

Phase 2 involves installing the best route to each destination into the local Routing Information Base (Loc-RIB). Effectively, the Loc-RIB is the primary routing table. Each route from Phase 1 has their NEXT_HOP checked to ensure the destination is reachable. If it is reachable, the AS_PATH is checked for loops. After that, routes are installed based on the following decision process:

- If there is only one route to a location, it is installed.
- If there are multiple routes to the same location, use the most preferred route from Level 1.
- If there is a tie, break the tie based on the following, in descending order of importance: shortest AS_PATH, smallest ORIGIN number, smallest MED, eBGP over iBGP, smallest metric or cost for reaching the NEXT_HOP, BGP identifier, and lowest IP address.

Note that the new routes that are installed into the Loc-RIB are in addition to any existing routes in the table. Once Phase 2 is completed, the Loc-RIB will consist of the best of both the new and older routes.

Decision phase 3

Phase 3 is route distribution or dissemination. This is the process of deciding which routes the router will advertise. If there is any route aggregation or summarizing, it happens here. Also, any route filtering from route maps happens here.

Once Phase 3 is complete, an update can be sent out to update the neighbor of new routes.

Aggregate routes and addresses

BGP-4 allows classless routing, which uses netmasks as well as IP addresses. This classless routing allows the configuration of aggregate routes by stating the address bits the aggregated addresses have in common.

The ATOMIC_AGGREGATE attribute informs routers that the route has been aggregated and should not be de-aggregated. An associated AGGREGATOR attribute include the information about the router that did the aggregating including its AS.

The BGP commands associated with aggregate routes and addresses are the following:

```
config router bgp
   config aggregate-address
      edit <aggr_addr_id>
         set as-set {enable | disable}
         set prefix <address_ipv4mask>
         set summary-only {enable | disable}
      end
   end
   config aggregate-address6
      edit <aggr_addr_id>
         set prefix6 <address_ipv6mask>
         set summary-only {enable | disable}
      end
   end
```

Troubleshooting BGP

There are some features in BGP that are used to deal with problems that may arise. Typically, the problems with a BGP network that has been configured involve routes going offline frequently. This is called route flap and causes problems for the routers using that route.

This section covers the following topics:

- Clearing routing table entries on page 417
- Route flap on page 417

Clearing routing table entries

To see if a new route is being properly added to the routing table, you can clear all or some BGP neighbor connections (sessions) using the `execute router clear bgp` command.

For example, if you have 10 routes in the BGP routing table and you want to clear the specific route to IP address 10.10.10.1, enter the following CLI command:

```
execute router clear bgp ip 10.10.10.1
```

To remove all routes for AS number 650001, enter the following CLI command:

```
execute router clear bgp as 650001
```

Route flap

When routers or hardware along a route go offline and back online that is called a route flap. Flapping is the term that is used if these outages continue, especially if they occur frequently.

Route flap is a problem in BGP because each time a peer or a route goes down, all the peer routers that are connected to that out-of-service router advertise the change in their routing tables. This creates a lot of administration traffic on the network and the same traffic re-occurs when that router comes back online. If the problem is something like a faulty network cable that wobbles online and offline every 10 seconds, there could easily be an overwhelming amount of routing updates sent out unnecessarily.

Another possible reason for route flap occurs with multiple FortiSwitch units in HA mode. When an HA cluster fails over to the secondary unit, other routers on the network may see the HA cluster as being offline, resulting in route flap. While this does not occur often, or more than once at a time, it can still result in an interruption in traffic that is unpleasant for network users. The easy solution for this problem is to increase the timers on the HA cluster, such as TTL timers, so they do not expire during the failover process. Also, configuring graceful restart on the HA cluster helps with a smooth failover.

The first method of dealing with route flap is to check your hardware. If a cable is loose or bad, it can easily be replaced and eliminate the problem. If an interface on the router is bad, either avoid using that interface or swap in a functioning router. If the power source is bad on a router, either replace the power supply or use a power conditioning backup power supply. These quick and easy fixes can save you from configuring more complex BGP options. However, if the route flap is from another source, configuring BGP to deal with the outages will ensure your network users uninterrupted service.

Some methods of dealing with route flap in BGP include the following:

- Holdtime timer
- Dampening
- BFD

Holdtime timer

The first line of defense to a flapping route is the holdtime timer. This timer reduces how frequently a route going down will cause a routing update to be broadcast.

After it is activated, the holdtime timer does not allow the FortiSwitch unit to accept any changes to that route for the duration of the timer. If the route flaps five times during the timer period, only the first outage is recognized by the

FortiSwitch unit. For the duration of the other outages, there will not be changes because the FortiSwitch unit is essentially treating this router as down. If the route is still flapping after the timer expires, it'll happen all over again.

Even if the route is not flapping (for example, if it goes down, comes up, and stays back up) the timer still counts down and the route is ignored for the duration of the timer. In this situation, the route is seen as down longer than it really is but there will be only the one set of route updates. This is not a problem in normal operation because updates are not frequent.

Also, the potential for a route to be treated as down when it is really up can be viewed as a robustness feature. Typically, you do not want most of your traffic being routed over an unreliable route. So if there is route flap going on, it is best to avoid that route if you can. This is enforced by the holdtime timer.

How to configure the holdtime timer

There are three different route flapping situations that can occur: the route goes up and down frequently, the route goes down and back up once over a long period of time, or the route goes down and stays down for a long period of time. These can all be handled using the holdtime timer.

For example, your network has two routes that you want to set the timer for. One is your main route (to 10.12.101.4) that all of your Internet traffic goes through, and it cannot be down for long if it is down. The second is a low speed connection to a custom network that is used infrequently (to 10.13.101.4). The timer for the main route should be fairly short (for example, 60 seconds). The second route timer can be left at the default because it is rarely used. In your BGP configuration, this looks like the following:

```
config router bgp
   config neighbor
      edit 10.12.101.4
         set holdtime-timer 60
      next
      edit 10.13.101.4
         set holdtime-timer 180
      next
   end
end
```

Dampening

Dampening is a method that is used to limit the amount of network problems due to flapping routes. With dampening, the flapping still occurs but the peer routers pay less and less attention to that route as it flaps more often. One flap does not start dampening, but the second flap starts a timer where the router will not use that route because it is considered unstable. If the route flaps again before the timer expires, the timer continues to increase. There is a period of time called the reachability half-life, after which a route flap will be suppressed for only half the time. This half-life comes into effect when a route has been stable for a while but not long enough to clear all the dampening completely. For the flapping route to be included in the routing table again, the suppression time must expire.

If the route flapping was temporary, you can clear the flapping or dampening from the FortiSwitch unit's cache by using one of the `execute router clear bgp` CLI commands:

```
execute router clear bgp dampening {<ip_address> | <ip/netmask>}
```

For example, to remove route flap dampening information for the 10.10.0.0/16 subnet, enter the following CLI command:

```
execute router clear bgp dampening 10.10.0.0/16
```

The BGP commands related to route dampening are the following:

```
config router bgp
   set dampening {enable | disable}
```

```
   set dampening-max-suppress-time <minutes_integer>
   set dampening-reachability-half-life <minutes_integer>
   set dampening-reuse <reuse_integer>
   set dampening-suppress <limit_integer>
end
```

BFD

Bidirectional Forwarding Detection (BFD) is a protocol that you can use to quickly locate hardware failures in the network. Routers running BFD communicate with each other and if a timer runs out on a connection then that router is declared down. BFD then communicates this information to the routing protocol and the routing information is updated. For more information about BFD, see Bidirectional forwarding detection on page 460.

Configuring BGP

Starting in FortiSwitchOS 7.2.2, the `set ebgp-requires-policy` command (under `config router bgp`) is set to `enable` by default, which prevents the BGP router from learning or advertising prefixes from or to its eBGP peers.

Configuring BGP on the FortiSwitch unit includes the following major steps:

1. Enter the BGP configuration mode on page 419.
2. Set the autonomous system and router identifier on page 419.
3. Configure the BGP neighbors on page 419.
4. Redistribute non-BGP routes on page 420.
   Advertise these non-BGP routes within BGP.

Enter the BGP configuration mode

Enter the BGP configuration mode to access all of the BGP configuration commands:

```
# config router bgp
```

Set the autonomous system and router identifier

Set the autonomous system. For iBGP, the AS value needs to match the `remote-as` value in the neighbor router. For eBGP, the AS value differs from the `remote-as` value in the neighbor router. You also need to specify a fixed router identifier for the FortiSwitch unit. These two commands are mandatory.

```
# set as <AS number>
# set router-id <IP_address>
```

Configure the BGP neighbors

Configure the BGP neighbors.

NOTE: For iBGP, if the IP address of the BGP neighbor is a loopback address, you must use the `set update-source cmd` command to specify which interface address will be used as the source IP address in the outgoing BGP packet.

```
config neighbor
   edit "<IPv4_or_IPv6 address>"
      set remote-as <1-4294967295>
```

```
   end
```

Redistribute non-BGP routes

Redistribute non-BGP IPv4 or IPv6 routes within BGP:

```
   config redistribute {connected | isis | ospf | rip | static}
      set status enable
      set route-map <string>
   end

   config redistribute6 {connected | isis | ospf | rip | static}
      set status {disable | enable}
      set route-map <string>
   end
```

Other BGP commands

Clearing the BGP routes

Use the following commands to clear the BGP routes:

```
execute router clear bgp all
execute router clear bgp ip <IPv4_or_IPv6_address>
execute router clear bgp ipv6 <IPv4_or_IPv6_address>
execute router clear bgp as <AS_number>
execute router clear bgp dampening <IP_address>
```

Checking the BGP configuration

The `get router info bgp` and `get router info6 bgp` commands have options to display different aspects of the BGP configuration and status.

For example:

```
get router info bgp neighbors
get router info bgp network
get router info6 bgp filter-list
get router info6 bgp route-map
```

Changing the maximum number of paths for ECMP

If you are using equal-cost multi-path (ECMP) routing with the eBGP or iBGP, the maximum number of paths is 1 by default. Use the following commands to change the default:

```
config router bgp
   set maximum-paths-ebgp <1-64>
   set maximum-paths-ibgp <1-64>
end
```

Sample configuration

Here is an example of a BGP routing configuration:

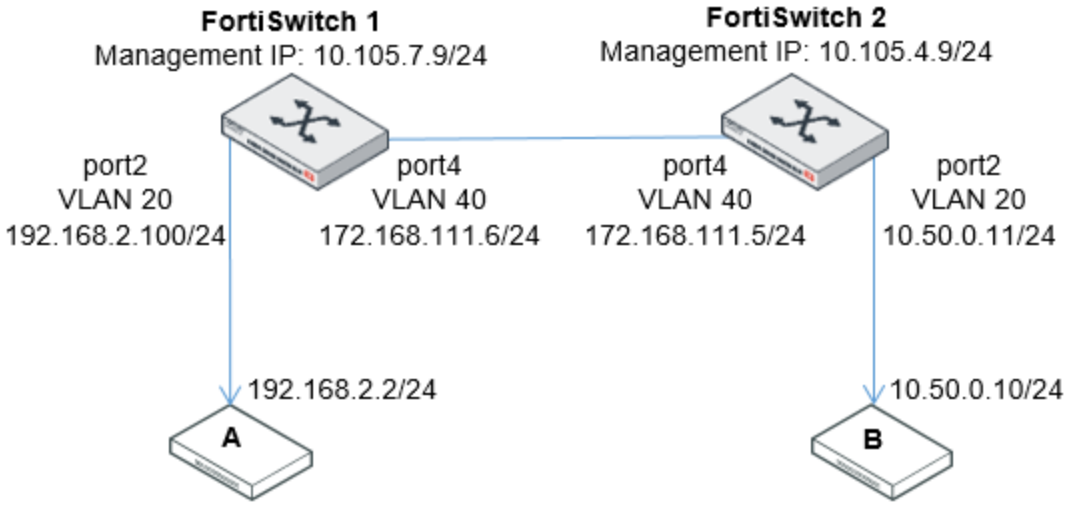

Configure system interfaces

Interface configuration for FortiSwitch 1:

```
config system interface
   edit mgmt
      set ip 10.105.7.9 255.255.255.0
      set allowaccess ping https http ssh telnet
      set type physical
   next
   edit internal
      set type physical
   next
   edit vlan20-p2
      set ip 192.168.2.100 255.255.255.0
      set allowaccess ping https http ssh telnet
      set vlanid 20
      set interface internal
   next
   edit vlan40-p4
      set ip 172.168.111.6 255.255.255.0
      set allowaccess ping https http ssh telnet
      set vlanid 40
      set interface internal
end
config switch interface
   edit "port2"
      set native-vlan 20
      set stp-state disabled
   next
   edit "port4"
      set native-vlan 40
      set stp-state disabled
   next
   edit "internal"
```

```
      set allowed-vlans 1,20, 40, 4094
      set stp-state disabled
   next
end
```

Internal BGP

In this example, the two neighboring switches are in the same autonomous system.

Configuration for FortiSwitch 1:

```
config router bgp
   set as 6500
   set router-id 1.2.3.4
      config neighbor
         edit "172.168.111.5"
            set remote-as 6500
         next
      end
      config network
         edit 1
            set prefix 192.168.2.0 255.255.255.0
         next
      end
      config redistribute "connected"
      end
   end
end
```

Configuration for FortiSwitch 2:

```
config router bgp
   set as 6500
   set router-id 5.6.7.8
      config neighbor
         edit "172.168.111.6"
            set remote-as 6500
         next
      end
      config network
         edit 1
            set prefix 10.50.2.0 255.255.255.0
         next
      end
      config redistribute "connected"
      end
   end
end
```

External BGP

In this example, the two neighboring switches are in separate autonomous systems.

Configuration for FortiSwitch 1:

```
config router bgp
   set as 6500
   set router-id 1.2.3.4
```

```
   set ebgp-requires-policy disable
      config neighbor
         edit "172.168.111.5"
            set remote-as 7500
         next
      end
      config network
         edit 1
            set prefix 192.168.2.0 255.255.255.0
         next
      end
      config redistribute "connected"
      end
   end
end
```

Configuration for FortiSwitch 2:

```
config router bgp
   set as 7500
   set router-id 5.6.7.8
   set ebgp-requires-policy disable
      config neighbor
         edit "172.168.111.6"
            set remote-as 6500
         next
      end
      config network
         edit 1
            set prefix 10.50.2.0 255.255.255.0
         next
      end
      config redistribute "connected"
      end
   end
end
```

Checking the BGP configuration

Using the following command, you can check the BGP status on the local switch:

```
# get router info bgp summary
```

To check the details about the BGP neighbors:

```
# get router info bgp neighbors
```

To check the routes learned by BGP, use the following command:

```
# get router info routing-table details
```

IS-IS routing

- You must have an advanced features license to use IS-IS routing.
- This feature is supported only on the SVI.

The Intermediate System to Intermediate System Protocol (IS-IS) allows routing of ISO's OSI protocol stack Connectionless Network Service (CLNS). IS-IS is an Interior Gateway Protocol (IGP) that is not intended to be used between autonomous systems.

IS-IS is a link state protocol that is well-suited to smaller networks. It is in widespread use and has near universal support on routing hardware. It is quick to configure and works well if there are no redundant paths. However, IS-IS updates are sent out node-by-node, so it can be slow to find a path around network outages. IS-IS also lacks good authentication, cannot choose routes based on different quality-of-service methods, and can create network loops if you are not careful. IS-IS uses Djikstra's algorithm to find the best path, like OSPF.

While OSPF is more widely known, IS-IS is a viable alternative to OSPF in enterprise networks and ISP infrastructures, largely due to its native support for IPv6 and its nondisruptive methods for splitting, merging, migrating, and renumbering network areas.

IS-IS uses type-length-value (TLV) parameters to carry information in Link-State PDUs (LSPs). The TLV field consists of one octet of type (T), one octet of length (L), and "L" octets of value (V). The LSP contains information about each router in an area and its connected interfaces. The complete sequence number PDU (CSNP) contains a list of all LSPs in the current LSDB.

Using the GUI:

1. Go to *Router* > *Config* > *ISIS* > *Settings*.

2. In the *Router ID* field, enter the IPv4 address for the router identifier.
3. In the *Type* dropdown list, select how to distribute the default route into the level's link-state packet (LSP): *Level-1*, *Level-1-2*, or *Level-2 Only*.
4. In the *Default Information Metric* field, enter the default information metric.
5. Under *Nets*, click + to configure the IS-IS network.
   a. In the *ID* field, enter an integer identifier (1-63).
   b. In the *Net* field, enter the IS-IS network.
   c. To add another row, click + again.

   NOTE: The maximum number of rows is 3.
6. Under *Redistribute*, select the *Enable* checkbox to redistribute connected, static, RIP, BGP, or OSPF routes within IS-IS and then enter the redistribution metric for each protocol that you enabled.
7. Click *Update*.
8. Go to *Router* > *Config* > *ISIS* > *Interfaces*.

9. Click *Configure* for the IS-IS interface that you want to configure.
   a. In the *Circuit Type* dropdown list, select the IS-IS circuit type to use for this interface:
      - Select *Level-1* for intra-area.
      - Select *Level-1-2* for both intra-area and inter-area.
      - Select *Level-2* for inter-area.
   b. Select the *Passive Interface* checkbox to set this interface as passive.
   c. Select the *Bidirectional Forwarding Detection* checkbox to enable bidirectional forwarding detection (BFD).
   d. Under *Hello Authentication*, select *Password* or *MD5* from the *Mode* dropdown list.

      NOTE: To use MD5 authentication, you must first configure a key chain. For more details, see Creating a key chain on page 468.
      - If you selected *Password* for the authentication mode, enter the password in the *Password* field.
      - If you selected *MD5* for the authentication mode, select a key chain from the *MD5 Key* dropdown list.
   e. Click *Add*.
10. Go to *Router* > *Monitor* > *Routing* to check your IS-IS configuration.

Using the CLI:

1. Enter the IS-IS configuration mode to access all of the IS-IS configuration commands:
   ```
   # config router isis
   ```
2. Under the `config router isis` command, enable the `status` option for IPv4 traffic or the `status6` option for IPv6 traffic on the specified interface:
   ```
   config interface
      edit <IS-IS interface name>
         set auth-keychain-hello <string>
         set auth-mode-hello {md5 | password}
         set auth-password-hello <password>
         set bfd {enable | disable}
         set bfd6 {enable | disable}
         set circuit-type {level-1 | level-1-2 | level-2}
         set csnp-interval-l1 <1-65535 seconds>
         set csnp-interval-l2 <1-65535 seconds>
         set hello-interval-l1 <1-65535 seconds; 0 to use 1-second hold time>
         set hello-interval-l2 <1-65535 seconds; 0 to use 1-second hold time>
         set hello-multiplier-l1 <2-100>
         set hello-multiplier-l2 <2-100>
         set hello-padding {disable | enable}
         set metric-l1 <1-63>
         set metric-l2 <1-63>
         set passive {disable | enable}
         set priority-l1 <0-127>
   ```

```
      set priority-l2 <0-127>
      set status {disable | enable}
      set status6 {disable | enable}
      set wide-metric-l1 <1-16777214>
      set wide-metric-l2 <1-16777214>
   end
```

3. You can use BFD for the IS-IS routing protocol with IPv4 or IPv6 addresses:

```
config router isis
   config interface
      edit <IS-IS interface name>
         set bfd {enable| disable}
         set bfd6 {enable| disable}
      next
   end
end
```

For example, if you want to enable IPv4 BFD on vlan100:

```
config router isis
   config interface
      edit "vlan100"
         set bfd enable
      next
   end
end
```

4. Under the `config router isis` command, configure the IS-IS network:

```
config net
   edit <1-63>
      set <IS-IS net xx.xxxx. ... .xxxx.xx>
   end
```

5. Under the `config router isis` command, redistribute non-IS-IS routes within IS-IS for IPv4 traffic or for IPv6 traffic:

```
config redistribute {bgp | connected | ospf | rip | static}
   set status {disable | enable}
   set metric <0-4261412864>
   set metric-type {external | internal}
   set level {level-1 | level-1-2 | level-2}
   set routemap <string>
end

config redistribute6 {bgp6 | connected | ospf6 | ripng | static}
   set status {disable | enable}
   set metric <0-4261412864>
   set level {level-1 | level-1-2 | level-2}
   set routemap <string>
end
```

6. Check your IS-IS configuration:

```
get router info isis interface
get router info isis route
get router info isis summary
get router info isis topology
get router info6 isis interface
get router info6 isis route
get router info6 isis summary
get router info6 isis topology
```

Configuration examples

The following is an example of an IS-IS configuration for IPv4 traffic:

```
config router isis
   set default-information-metric 60
      config interface
         edit "vlan100"
            set circuit-type level-1
            set priority-l1 80
            set wide-metric-l1 200
         next
         edit "vlan102"
            set circuit-type level-2
         next
      end
      config net
         edit 1
            set net 49.0002.0000.0000.1048.00
      next
   end
   set metric-style wide
      config redistribute "connected"
            set status enable
         end
         config redistribute "rip"
         end
         config redistribute "ospf"
         end
         config redistribute "bgp"
         end
         config redistribute "static"
      end
   end
```

The following is an example of an IS-IS configuration for IPv6 traffic:

```
config router isis
   config interface
      edit "vlan10"
      next
   end
   config net
      edit 1
         set net 49.0000.0010.0100.1001.00
      next
   end
   config redistribute "connected"
   end
   config redistribute "rip"
   end
   config redistribute "ospf"
   end
   config redistribute "bgp"
   end
   config redistribute "static"
   end
   config redistribute6 "connected"
```

```
    end
    config redistribute6 "static"
    end
    config redistribute6 "ospf6"
    end
    config redistribute6 "ripng"
    end
end
```

OSPF

NOTE: You must have an advanced features license to use OSPF routing.

Open shortest path first (OSPF) is a link-state interior routing protocol that is widely used in large enterprise organizations. OSPF provides routing within a single autonomous system (AS). This differs from BGP, which provides routing between autonomous systems.

An OSPF AS can contain only one area, or it can consist of a group of areas connected to a backbone area. A router connected to more than one area is an area border router (ABR). An autonomous system boundary router (ASBR) is located between an OSPF autonomous system and a non-OSPF network. Routing information is contained in a link-state database. Routing information is communicated between routers using link-state advertisements (LSAs).

The main benefit of OSPF is that it detects link failures in the network quickly and converges network traffic successfully within seconds without any network loops. Also, OSPF has features to control which routes are propagated to contain the size of the routing tables.

You can enable bidirectional forwarding detection (BFD) with OSPF. BFD is used to quickly locate hardware failures in the network. Routers running BFD communicate with each other, and, if a timer runs out on a connection, that router is declared to be down. BFD then communicates this information to OSPF, and the routing information is updated.

NOTE: OSPF MIBs are not supported in this release.

For additional information about OSPF routing, see the OSPF section of the FortiOS Administration Guide.

How OSPF works

Areas

An OSPF implementation consists of one or more areas. An area consists of a group of contiguous networks. If you configure more than one area, Area Zero is always the backbone area. An ABR links one or more areas to the OSPF backbone area.

The FortiSwitch unit supports different types of areas—stub areas, Not So Stubby areas (NSSA), and regular areas. A stub area is an interface without a default route configured. NSSA is a type of stub area that can import AS external routes and send them to the backbone but cannot receive AS external routes from the backbone or other areas. All other areas are considered regular areas.

Adjacencies

When an OSPF router boots up, it sends OSPF Hello packets to find neighbors on the same network. Neighbors exchange information, and the link-state databases of both neighbors are synchronized. At this point, these neighbors are said to be adjacent.

For two OSPF routers to become neighbors, the following conditions must be met:

- The subnet number and subnet mask for the interface must match in both routers.
- The Hello interval and Dead interval values must match.
- The routers must have the same OSPF area ID.
- If authentication is used, they must pass authentication checks.

In OSPF, routing protocol packets are only passed between adjacent routers.

Route summarization

Using route summarization reduces the number of LSAs being sent between routers. OSPF offers two types of route summarization:

- Between areas through an ABR. This method summarizes routes in the area configuration.

```
config area
   edit <area_IPv4_address>
      config range
         edit <id>
            set prefix <xxx.xxx.xxx.xxx> <xxx.xxx.xxx.xxx>
         next
      end
   next
end
```

- Between an OSPF AS and a non-OSPF network through an ASBR. This method summarizes external routes when you redistribute them.

```
config summary-address
   edit <id>
      set prefix <xxx.xxx.xxx.xxx> <xxx.xxx.xxx.xxx>
   next
end
```

Graceful restart helper mode

Starting in FortiSwitchOS 6.4.3, the FortiSwitch unit enters the helper (neighbor) mode when a neighboring router sends a grace LSA before it restarts. The FortiSwitch unit keeps the restarting router in the forwarding path for OSPF routing, as long as there are no network topology changes. After the restarting router completes its graceful restart, the FortiSwitch unit exits the helper mode.

This feature is always enabled.

Database overflow protection

When the OSPF link-state database is large, some routers do not have enough resources to store the complete link-state database. To prevent database overflow, you can limit the number of AS-external-LSAs in the link-state database. When the maximum number of AS-external-LSAs is reached, the router deletes all AS-external-LSAs that it originated and stops originating AS-external-LSAs for the specified number of seconds.

By default, this feature is disabled.

Use the following commands to configure database overflow protection:

```
config router ospf
```

```
    set database-overflow enable
    set database-overflow-max-external-lsa <0-2147483647>
    set database-overflow-time-to-recover <0-65535>
end
```

Configuring OSPF

If you want to use virtual routing and forwarding (VRF) with OSPF, you need to create a VRF instance before configuring OSPF. See Virtual routing and forwarding on page 461.

Using the GUI:

1. Create a switch virtual interface. See Switch virtual interfaces on page 24.
2. Go to *Router* > *Config* > *OSPF* > *Settings*.
   a. If you want to use a VRF instance, select it from the *VRF* dropdown list.
   b. Select the *Enable* checkbox.
   c. Enter a unique 32-bit number in dotted decimal format for the router identifier. **NOTE:** Without a router identifier, OSPF routing will not work.
   d. If you are going to advertise default routes within OSPF, configure the default route option and enter the routing metric (cost) for other routing protocols.
   e. If you want to redistribute non-OSPF routes, select *Enabled* under Connected, Static, RIP, BGP, or ISIS and then enter the routing metric in the Metric field.
   f. Select *Update*.
3. Go to *Router* > *Config* > *OSPF* > *Areas*, select the VRF instance or *None*, and then select *Add OSPF Area*.
   a. Enter the area IP address.
   b. Select if the area is a stub area, NSSA, or a regular area.
   c. Select *Add*.
4. Go to *Router* > *Config* > *OSPF* > *Networks*, select the VRF instance or *None*, and then select *Add Network*.
   a. Enter the network identifier.
   b. Enter the IP address and netmask, separated with a space. Use an IP address that includes the switch virtual interface.
   c. Select the area that you created.
   d. Select *Add*.
5. Go to *Router* > *Config* > *OSPF* > *Interfaces*, select the VRF instance or *None*, and then select *Configure OSPF Interface*.
   a. Select the same type of authentication that you selected for the area.
   b. If you want static bidirectional forwarding detection, select *Enable* or *Global*.
   c. Enter the maximum transmission unit.
   d. Enter the cost.
   e. Enter the number of seconds between Hello packets being sent.
   f. Enter the number of seconds that a Hello packet is not received before the OSPF router decides that a neighbor has failed.
   g. Select *Add*.

Using the CLI:

Configuring OSPF using IPv4 on the FortiSwitch unit includes the following major steps:

1. Entering the OSPF configuration mode on page 432.
2. Setting the router identifier on page 432.
   Each router must have a unique 32-bit number. **NOTE:** Without a router identifier, OSPF routing will not work.
3. Creating an area on page 432.
   You must create at least one area.
4. Configuring the network on page 433.
   Attach one or more networks to each area.
5. Configuring the OSPF interface on page 433.
6. Redistributing non-OSPF routes on page 433.
   Advertise these non-OSPF routes within OSPF.
7. Checking the OSPF configuration on page 434.

NOTE:

- You can also configure OSPF using IPv6 with the `config router ospf6` command.
- Starting in FortiSwitchOS 7.0.0, OSPF supports VRF. To create multiple routing tables within the same router, use the `config vrf` command under `config router ospf`.

Entering the OSPF configuration mode

Enter the OSPF configuration mode to access all of the OSPF configuration commands:

```
# config router ospf
```

Setting the router identifier

Each router within an area must have a unique 32-bit number. The router identifier is written in dotted decimal format, but it is not an IPv4 address. **NOTE:** Without a router identifier, OSPF routing will not work.

```
set router-id <router-id>
```

For example:

```
# config router ospf
(ospf) # set router-id 1.1.1.2
```

Creating an area

You must create at least one area. The area number is written in dotted decimal format (for example, configure area 100 as 0.0.0.100).

```
config area
   edit <area number>
      set shortcut (default | disable | enable)
      set type {nssa | regular | stub}
end
```

For example:

```
   (ospf) # config area
   (area) # edit 0.0.0.4
   (0.0.0.4) # set type nssa
```

Configuring the network

Use this subcommand to identify the OSPF-enabled interfaces. The prefix length in the interface must be equal or larger than the prefix length in the network statement.

```
config network
   edit <network number>
      set area <area>
      set prefix <network prefix> <mask>
```

For example:

```
(ospf) # config network
(network) # edit 1
(1) # set area 0.0.0.4
(1) # set prefix 10.1.1.0 255.255.255.0
```

Configuring the OSPF interface

Configure interface-related OSPF settings. Enter a descriptive name for the OSPF interface name.

```
config interface
   edit <OSPF_interface_name>
      set priority <1-255>
```

For example:

```
(ospf) # config interface
(ospf-interface) # edit oi1
(oi1) # set priority 255
```

NOTE: The following values must match for an adjacency to form:

- area type and number
- interface subnet and mask
- hello interval
- dead interval

Redistributing non-OSPF routes

Redistribute non-OSPF routes (directly connected or static routes) within OSPF:

```
config redistribute {bgp | connected | isis | rip | static}
   set status enable
   set metric <integer>
   set metric-type {1 | 2}
end
```

Add route summarization:

```
config summary-address
   edit <id>
      set prefix <xxx.xxx.xxx.xxx> <xxx.xxx.xxx.xxx>
   next
end
```

For example:

```
(ospf) # config redistribute connected
(connected) # set status enable
```

```
(connected) # end

(ospf) # config summary-address
(summary-address) # edit 1
new entry '1' added
(1) # set prefix 10.1.0.0 255.255.0.0
(1) # next
(summary-address) # end
```

Checking the OSPF configuration

The `get router info ospf` command has options to display different aspects of the OSPF configuration and status. For example:

```
get router info ospf neighbor {<neighbor_ID> | all | detail | detail all | <interface_IP_
     address>}
get router info ospf database {brief | self-originate | router | network | summary | asbr-
     summary| external | nssa-external | opaque-link | opaque-area | opaque-as | max-age}
```

Example configuration

The following example shows a very simple OSPF network with one area. FortiSwitch 1 has one OSPF interface to FortiSwitch 2:

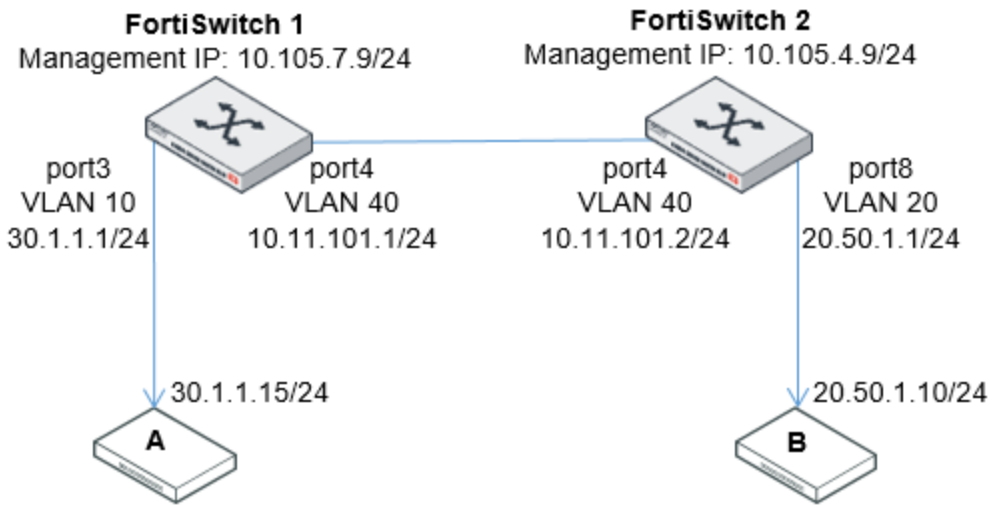

Configuring system interfaces

These are the same configuration steps as for static routing.

Switch 1

```
config system interface
   edit vlan10-p3
      set ip 30.1.1.1 255.255.255.0
      set allowaccess ping https http ssh telnet
      set vlanid 10
   next
```

```
   edit vlan40-p4
      set ip 10.11.101.1 255.255.255.0
      set allowaccess ping https http ssh telnet
      set vlanid 40
end
config switch interface
   edit "port3"
      set native-vlan 10
   next
   edit "port4"
      set native-vlan 40
   next
end
```

Switch 2

```
config system interface
   edit vlan20-p8
      set ip 20.50.1.1 255.255.255.0
      set allowaccess ping https http ssh telnet
      set vlanid 20
   next
   edit vlan40-p4
      set ip 10.11.101.2 255.255.255.0
      set allowaccess ping https http ssh telnet
      set vlanid 40
end
config switch interface
   edit "port8"
      set native-vlan 20
   next
   edit "port4"
      set native-vlan 40
   next
end
```

Configuring the OSPF router

Configure OSPF with the following:

1. Set the router ID.
2. Create the area.
3. Create the network (set network prefix and associate with an area).
4. Configure the OSPF interface.

Switch 1

```
config router ospf

   set router-id 10.11.101.1

   config area
      edit 0.0.0.0
      next
   end
```

```
   config network
      edit 1
         set area 0.0.0.0
         set prefix 10.11.101.0 255.255.255.0
      next
   end

   config interface
      edit vlan40
         set cost 100
         set priority 100
      next
   end

   config redistribute connected
      set status enable
   end

end
```

Switch 2

```
config router ospf
   set router-id 10.11.101.2

   config area
      edit 0.0.0.0
      next
   end

   config network
      edit 1
         set area 0.0.0.0
         set prefix 10.11.101.0 255.255.255.0
      next
   end

   config interface
      edit vlan40
         set cost 100
         set priority 100
      next
   end

   config redistribute connected
      set status enable
   end

end
```

Verifying OSPF neighbors

```
get router info ospf neighbor all
```

Verifying OSPF routes

```
get router info ospf route
```

RIP

NOTE: You must have an advanced features license to use RIP routing.

The Routing Information Protocol (RIP) is a distance-vector routing protocol that works best in small networks that have no more than 15 hops. Each router maintains a routing table by sending out its routing updates and by asking neighbors for their routes. RIP is relatively simple to configure on FortiSwitch units but slow to respond to network outages. RIP routing is better than static routing but less scalable than open shortest path first (OSPF) routing.

The FortiSwitch unit supports RIP version 1 and RIP version 2:

- RIP version 1 uses classful addressing and broadcasting to send out updates to router neighbors. It does not support different sized subnets or classless inter-domain routing (CIDR) addressing.
- RIP version 2 supports classless routing and subnets of various sizes. Router authentication supports MD5 and authentication keys. Version 2 uses multicasting to reduce network traffic.

RIP uses three timers:

- The update timer determines the interval between routing updates. The default setting is 30 seconds.
- The timeout timer is the maximum time that a route is considered reachable while no updates are received for the route. The default setting is 180 seconds. The timeout timer setting should be at least three times longer than the update timer setting.
- The garbage timer is the is the how long that the FortiSwitch unit advertises a route as being unreachable before deleting the route from the routing table. The default setting is 120 seconds.

You can enable bidirectional forwarding detection (BFD) with RIP. BFD is used to quickly locate hardware failures in the network. Routers running BFD communicate with each other, and, if a timer runs out on a connection, that router is declared to be down. BFD then communicates this information to RIP, and the routing information is updated.

When you configure RIP routing, you can choose the strategy the access list uses to permit or deny IP addresses:

- *Prefix*—Specify the IP address and bit mask to allow or block.
- *Wildcard*—Specify the Cisco-style filter to allow or block.

For additional information about RIP routing, see the RIP section of the FortiOS Administration Guide.

Terminology

Active RIP interface: Each RIP router sends and receives updates by actively communicating with its neighbors.

Metric: RIP uses hop count as the metric for choosing the best route. A hop count of 1 represents a network that is connected directly to the FortiSwitch unit. A hop count of 16 represents a network that cannot be reached.

Passive RIP interface: The RIP router listens to updates from other routers but does not send out updates. A passive RIP interface reduces network traffic.

Prefix list: A more powerful prefix-based filtering mechanism. A prefix is an IP address and netmask.

Split horizon: A way to avoid routing loops.

Configuring RIP routing

Using the GUI and the prefix strategy:

1. Create a switch virtual interface (SVI). See Switch virtual interfaces on page 24.
2. Go to *Router* > *Config* > *RIP* > *Settings*.

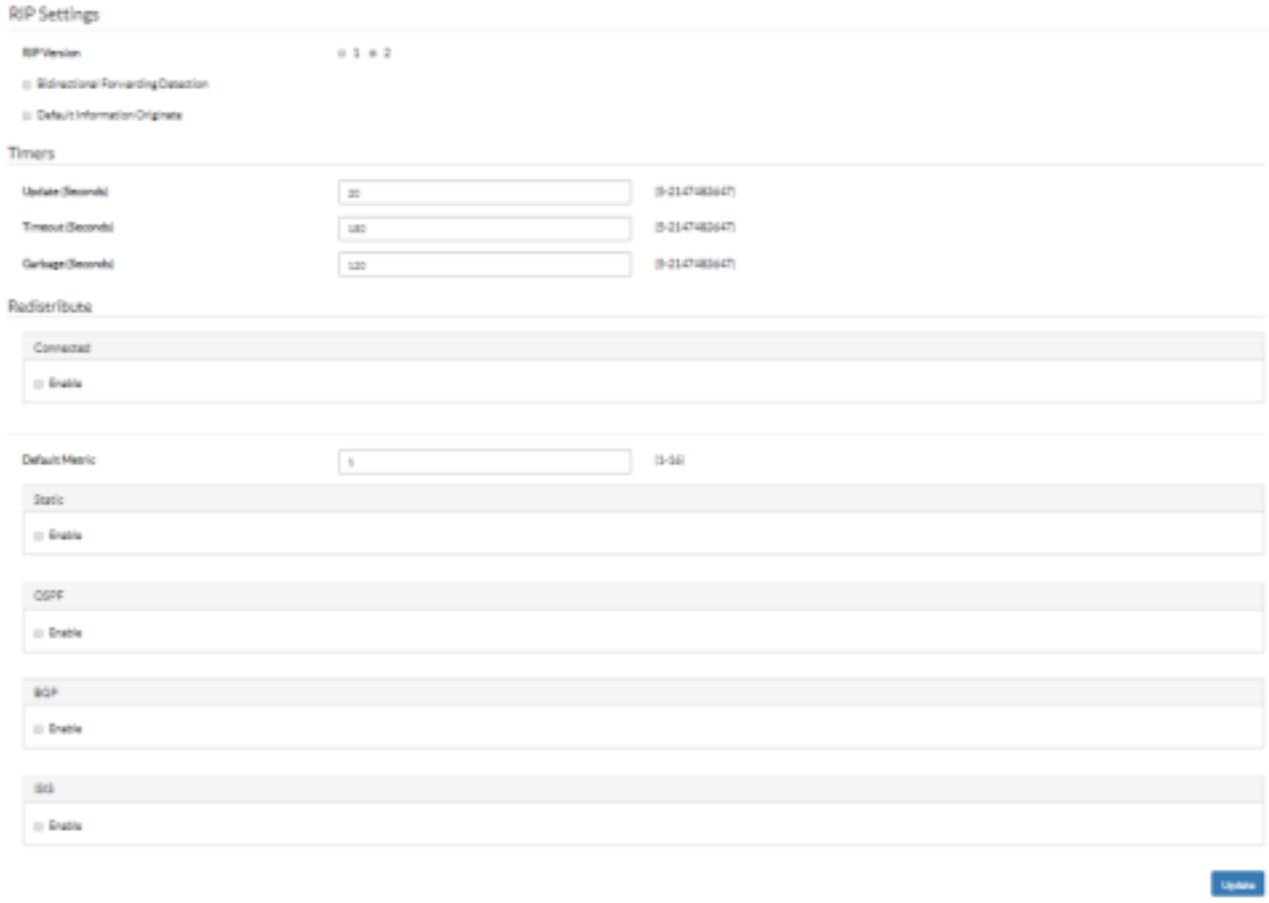

   a. Select whether you want to use RIP version 1 or RIP version 2. RIP version 2 is the default.
   b. If you want to use BFD, select *Bidirectional Forwarding Detection*.
   c. If you want to use a default route, select *Default Information Originate*.
   d. If you want to change the default timer values, enter the number of seconds in the *Update*, *Timeout*, and *Garbage* fields.
   e. If you want to redistribute non-RIP routes, select *Enable* under Connected, Static, OSPF, BGP, or ISIS.
      - If you select *Enable* under Connected, enter the routing metric to use.
      - If you select *Enable* under Static, OSPF, BGP, or ISIS, select *Override Metric* if you do not want to use the default routing metric and then enter the routing metric to use.
   f. Enter the default routing metric to use for static routing, OSPF, BGP, and ISIS.

3. Go to *Router* > *Config* > *Access Lists* and select *Add Access List*.

   a. Enter an identifier with one or more alphabetic characters.
   b. Enter an optional description of the access list.
   c. Select *Add*.
   d. Select *Config Rules* in the row for the access list that you just created.

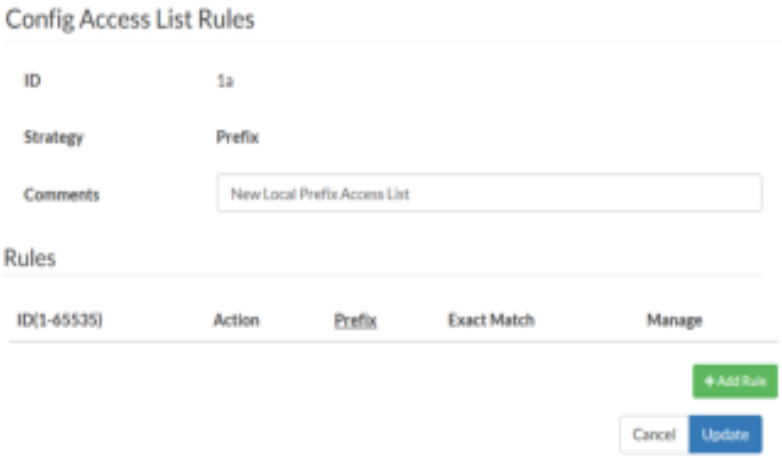

   e. Select *Add Rule*.
   f. Enter an identifier (1-65535), select *Deny* or *Permit* to specify if the rule will block or allow the specified IP addresses, and enter the prefix.
   g. If you entered the complete IP address, select the *Exact Match* checkbox.
   h. Select *Add Rule* if you want to add more rules.
   i. After you have added all of the rules that you want in the access list, select *Update* to save the rules you added.

4. Go to *Router* > *Config* > *RIP* > *Distances* and select *Add RIP Distance*.

   a. Enter the distance identifier in the Distance ID field.
   b. Enter the distance.
   c. Select the access list that you added in the previous step.

d. Enter the IP address and netmask, separated with a space or with a slash. For example, enter `1.2.3.4/5` or `1.2.3.4 248.0.0.0`.
e. Select *Add*.

5. Go to *Router* > *Config* > *RIP* > *Networks* and select *Add Network*.

a. Enter a unique value to identify this network configuration.
b. Enter an IP address and netmask for your RIP network, separated with a slash, and select *Add*. For example, enter 172.168.200.0/255.255.255.0. **NOTE:** Select an IP address for a network that includes all SVIs that you want to use. You can configure multiple network ranges to cover all SVIs that will be using RIP routing.

6. Go to *Router* > *Config* > *RIP* > *Interfaces* and select *Configure RIP* for the appropriate interface.

a. If you want to change the RIP version used to send and receive routing updates, select from the *Send Version* and *Receive Version* drop-down menus.
b. If you do not want to send RIP updates from this interface, select *Passive Interface*.
c. If you want to use authentication, select *Text* or *MD5*.
NOTE: To use MD5 authentication, you must first configure a key chain. For more details, see Creating a key chain on page 468.
   - If you select *Text*, enter the password to use for authentication in the *Password* field.
   - If you select *MD5*, you can select which key chain to use if you have more than one configured.

d. Select *Add*.

Using the GUI and the wildcard strategy:

1. Create a switch virtual interface (SVI). See Switch virtual interfaces on page 24.
2. Go to *Router* > *Config* > *RIP* > *Settings*.

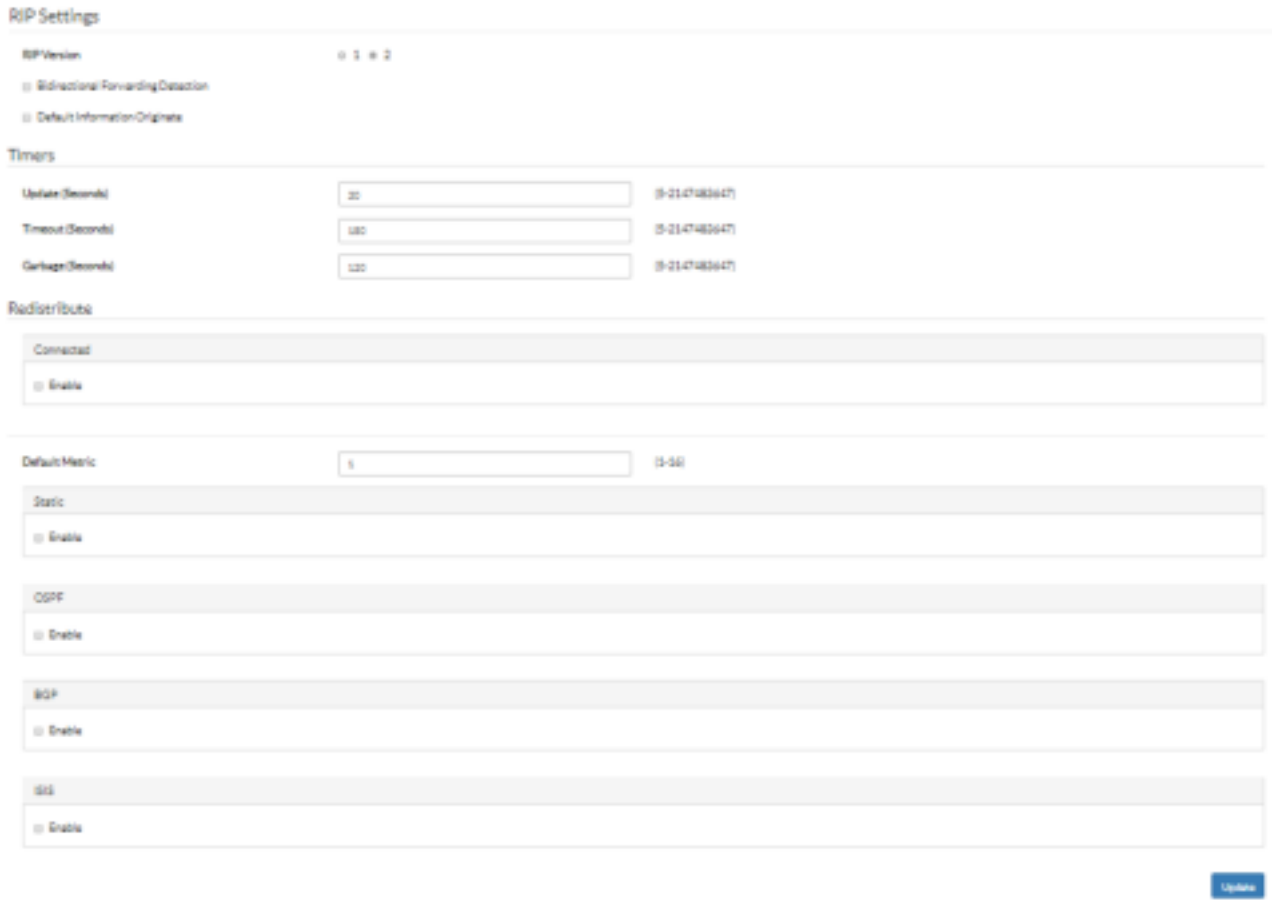

   a. Select whether you want to use RIP version 1 or RIP version 2. RIP version 2 is the default.
   b. If you want to use BFD, select *Bidirectional Forwarding Detection*.
   c. If you want to use a default route, select *Default Information Originate*.
   d. If you want to change the default timer values, enter the number of seconds in the *Update*, *Timeout*, and *Garbage* fields.
   e. If you want to redistribute non-RIP routes, select *Enable* under Connected, Static, OSPF, BGP, or ISIS.
      - If you select *Enable* under Connected, enter the routing metric to use.
      - If you select *Enable* under Static, OSPF, BGP, or ISIS, select *Override Metric* if you do not want to use the default routing metric and then enter the routing metric to use.
   f. Enter the default routing metric to use for static routing, OSPF, BGP, and ISIS.

3. Go to *Router* > *Config* > *Access Lists* and select *Add Access List*.

   a. Enter an identifier with all digits (in the range of 1-99).
   b. Enter an optional description of the access list.
   c. Select *Add*.
   d. Select *Config Rules* in the row for the access list that you just created.

   e. Select *Add Rule*.
   f. Enter an identifier (1-65535), select *Deny* or *Permit* to specify if the rule will block or allow the specified IP addresses, and enter the Cisco-style wildcard filter.
   g. Select *Add Rule* if you want to add more rules.
   h. After you have added all of the rules that you want in the access list, select *Update* to save the rules you added.

4. Go to *Router* > *Config* > *RIP* > *Distances* and select *Add RIP Distance*.

   a. Enter the distance identifier in the Distance ID field.
   b. Enter the distance.
   c. Select the access list that you added in the previous step.
   d. Enter the IP address and netmask, separated with a space or with a slash. For example, enter `1.2.3.4/5` or `1.2.3.4 248.0.0.0`.
   e. Select *Add*.

5. Go to *Router* > *Config* > *RIP* > *Networks* and select *Add Network*.

   a. Enter a unique value to identify this network configuration.
   b. Enter an IP address and netmask for your RIP network, separated with a slash, and select *Add*. For example, enter 172.168.200.0/255.255.255.0. **NOTE:** Select an IP address for a network that includes all SVIs that you want to use. You can configure multiple network ranges to cover all SVIs that will be using RIP routing.

6. Go to *Router* > *Config* > *RIP* > *Interfaces* and select *Configure RIP* for the appropriate interface.

   a. If you want to change the RIP version used to send and receive routing updates, select from the *Send Version* and *Receive Version* drop-down menus.
   b. If you do not want to send RIP updates from this interface, select *Passive Interface*.
   c. If you want to use authentication, select *Text* or *MD5*.
   NOTE: To use MD5 authentication, you must first configure a key chain. For more details, see Creating a key chain on page 468.
      - If you select *Text*, enter the password to use for authentication in the *Password* field.
      - If you select *MD5*, you can select which key chain to use if you have more than one configured.
   d. Select *Add*.

Using the CLI for IPv4 traffic:

```
config router access-list
   edit <access_list_name>
      set comments <comments>
      config rule
         edit <rule_int>
            set action {deny | permit}
            set prefix {<xxx.xxx.xxx.xxx> <xxx.xxx.xxx.xxx> | any}
            set wildcard <IP_address>
            set exact-match {enable | disable}
         end
```

```
      end

config router rip
   set bfd {disable | enable}
   set default-information-originate {disable | enable}
   set garbage-timer <5-2147483647 seconds>
   set timeout-timer <5-2147483647 seconds>
   set update-timer <5-2147483647 seconds>
   set default-metric <1-16>
   config redistribute {bgp | connected | isis | ospf | static}
      set status {disable | enable}
      set metric <0-16>
   end
   config distance
      edit <distance_ID>
         set access-list <access_list_name>
         set distance <1-255>
         set prefix <IPv4_address> <netmask>
      end
   config network
      edit <network identifier>
         set prefix <IPv4_address> <netmask>
      end
   config interface
      edit <interface_name>
         set auth-keychain <key_chain_str>
         set auth-mode {md5 | none |text}
         set auth-string <password_str>
         set receive-version {1 | 2 | both | global}
         set send-version {1 | 2 | both | global}
      end
   end
end
```

Using the CLI for IPv6 traffic:

```
config router access-list6
   edit <access_list_name>
      set comments <comments>
      config rule
         edit <rule_int>
            set action {deny | permit}
            set prefix6 {<xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx> | any}
            set exact-match {enable | disable}
         end
      end

config router ripng
   set bfd {disable | enable}
   set default-information-originate {disable | enable}
   set garbage-timer <5-2147483647 seconds>
   set timeout-timer <5-2147483647 seconds>
   set update-timer <5-2147483647 seconds>
   set default-metric <1-16>
   config redistribute {bgp | connected | isis | ospf6 | static}
      set status {disable | enable}
      set metric <0-16>
```

```
    end
    config offset-list
      edit <offset-list_name>
        set access-list6 <access-list_name>
        set direction {in | out}
        set interface {in | out}
        set offset <1-16>
        set status {disable | enable}
      end
    config aggregate-address
      edit <aggregate-address_entry_ID>
        set prefix6 <xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx/xxx>
      end
    config interface
      edit <interface_name>
        set passive {disable | enable}
        set split-horizon-statsus {disable | enable}
        set split-horizon {poisoned |regular}
      end
    end
end
```

Checking the RIP configuration

The `get router info rip` and `get router info6 rip` commands have options to display different aspects of the RIP configuration and status. For example, there are options to display the RIP general information and the RIP database:

```
get router info rip status
get router info6 rip status
get router info rip database
get router info6 rip database
```

Example configuration

The following example shows a very simple RIP network:

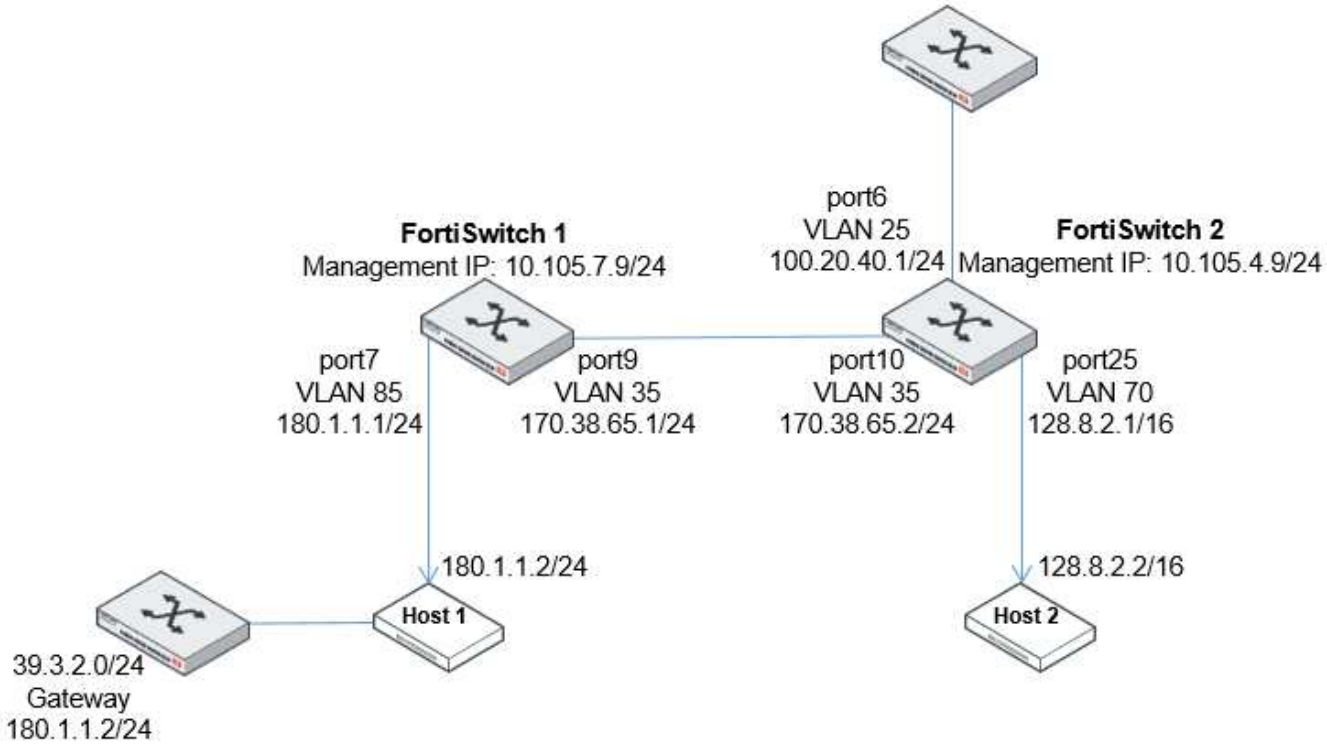

Switch 1: Configure the switch interface

```
config switch interface
   edit "port9"
      set allowed-vlans 35
   next
   edit "port7"
      set allowed-vlans 85
   next
end
```

Switch 1: Configure the system interface

```
config system interface
   edit "vlan35"
      set ip 170.38.65.1/24
      set allowaccess ping https http ssh snmp telnet
      set vlanid 35
   next
   edit "vlan85"
      set ip 180.1.1.1/24
      set allowaccess ping https http ssh snmp telnet
      set vlanid 85
   next
end
```

Switch 1: Configure the RIP router; add authentication between FortiSwitch 1 and FortiSwitch 2

```
config router rip
   config network
      edit 1
         set prefix 170.38.65.0/24
      next
      edit 2
         set prefix 180.1.1.0/24
      next
   end
   config interface
      edit "vlan35"
         set auth-mode text
         set auth-string simplepw1
      next
   end
end
```

Switch 1: Add a static route and redistribute it

```
config router static
   edit 1
      set dst 39.3.2.0 255.255.255.0
      set gateway 180.1.1.2
      set status enable
   next
end

config router rip
   config redistribute "static"
      set status enable
   next
end
```

Switch 2: Configure the switch interface

```
config switch interface
   edit "port10"
      set allowed-vlans 35
   next
   edit "port25"
      set allowed-vlans 70
   next
end
```

Switch 2: Configure the system interface

```
config system interface
   edit "vlan35"
      set ip 170.38.65.2/24
      set allowaccess ping https http ssh snmp telnet
      set vlanid 35
   next
   edit "vlan70"
```

```
      set ip 128.8.2.1/16
      set allowaccess ping https http ssh snmp telnet
      set vlanid 70
   next
end
```

Switch 2: Configure the RIP router; add authentication between FortiSwitch 1 and FortiSwitch 2

```
config router rip
   config network
      edit 1
         set prefix 170.38.65.0/24
      next
      edit 2
         set prefix 128.8.0.0/16
      next
   end
   config interface
      edit "vlan35"
         set auth-mode text
         set auth-string simplepw1
      next
   end
end
```

Switch 2: Add a connected route and redistribute it

```
config switch interface
   edit "port6"
      set allowed-vlans 25
   next
end
config system interface
   edit "vlan25"
      set ip 100.20.40.1/24
      set allowaccess ping https http ssh snmp telnet
      set vlanid 25
   next
end

config router rip
   config redistribute "connected"
      set status enable
   next
end
```

Multicast

A FortiSwitch unit can operate as a Protocol Independent Multicast (PIM) version-2 router. FortiSwitchOS supports PIM source-specific multicast (SSM) and version 3 of Internet Group Management Protocol (IGMP).

You can configure a FortiSwitch unit to support PIM using the `config router multicast` CLI command. When PIM is enabled, the FortiSwitch unit allocates memory to manage mapping information. The FortiSwitch unit communicates

with neighboring PIM routers to acquire mapping information and, if required, processes the multicast traffic associated with specific multicast groups.

NOTE:

- You must have an advanced features license to use PIM routing.
- This feature is supported only on the SVI.
- Access lists, prefix lists, and route maps are not supported.
- Bidirectional forwarding detection (BFD) is not supported.
- You cannot use PIM and the IGMP querier at the same time on the same switch virtual interface.
- PIM and IGMP snooping work independently.
- IPv6 is not supported.
- IGMP version-3 explicit membership tracking is not supported.
- SSM mapping is not supported.
- The multicast routing information base (MRIB) is not supported.
- The PIM management information base (MIB) is not supported.
- Starting in FortiSwitchOS 7.2.0, PIM is supported on RVIs.

Configuring PIM

A PIM domain is a logical area comprising a number of contiguous networks. The domain contains at least one Boot Strap Router (BSR) and a number of Rendezvous Points (RPs) and Designated Routers (DRs). An RP represents the root of a non-source-specific distribution tree to a multicast group.

To configure a PIM domain:

1. Determine the appropriate paths for multicast packets.
2. Make a note of the interfaces that will be PIM enabled. These interfaces can run a unicast routing protocol.
3. Enable PIM version 2 on all participating routers between the source and receivers. Use the `config router multicast` command to set global operating parameters. See Enabling PIM on page 449.
4. If required, adjust the default settings of PIM-enabled interface(s). See Configuring a PIM-enabled interface on page 450.

Enabling PIM

Using the GUI:

1. Go to *Router* > *Config* > *Multicast* > *Settings*.
2. Select the *Enable Multicast* checkbox.
3. Click *Update*.

Using the CLI:

```
config router multicast
   set multicast-routing enable
end
```

Creating a multicast flow

Starting in FortiSwitchOS 7.0.2, you can specify a range of multicast group addresses (IPv4) when configuring a PIM multicast flow in the CLI. Setting `group-addr-end` is optional, and the range must not overlap other defined ranges.

Using the GUI:

1. Go to *Router* > *Config* > *Multicast* > *Flows*.
2. Select *Add Multicast Flow*.
3. In the *Name* field, enter the name of the multicast flow.
4. In the *Comments* field, enter an optional description of the multicast flow.
5. Click +.
6. In the *ID* field, enter a number between 1 and 4294967295 to identify the multicast-flow entry.
7. In the *Group Address* field, enter the multicast group IPv4 address.
8. In the *Source Address* field, enter an IPv4 address for the multicast source.
9. Click *Add*.

Using the CLI:

```
config router multicast-flow
   edit <name>
      set comments <string>
      config flows
         edit <muliticast-flow_entry_identifier>
            set group-addr <224-239.xxx.xxx.xxx>
            set group-addr-end <224-239.xxx.xxx.xxx>
            set source-addr <IP_address>
         end
      end
```

Configuring a PIM-enabled interface

Using the GUI:

1. Go to *Router* > *Config* > *Multicast* > *Interfaces*.
2. Select *Configure* for one of the PIM-enabled interfaces.
3. In the *Multicast Flow* dropdown list, select a multicast flow.
   To create a multicast flow, see Creating a multicast flow on page 450.
4. In the *Hello Interval (Seconds)* field, enter the number of seconds that the FortiSwitch unit waits between sending hello messages to neighboring PIM routers.
5. In the *Designated Router Priority* field, enter a priority to the FortiSwitch unit Designated Router (DR) candidacy. The value is compared to that of other DR interfaces connected to the same network segment, and the router having the highest DR priority is selected to be the DR. If two DR priority values are the same, the interface having the highest IP address is selected.
6. In the *Response Time (Seconds)* field, enter the number of seconds between queries to IGMP hosts.
7. In the *Interval (Seconds)* field, enter the maximum number of seconds to wait for an IGMP query response.
8. Click *Add*.

Using the CLI:

```
config router multicast
```

```
config interface
   edit {interface_name | internal | mgmt}
      set pim-mode ssm-mode
      set multicast-flow <string>
      set hello-interval <1-180 seconds>
      set dr-priority <1-4294967295>
      config igmp
         set query-max-response-time <1-25 seconds>
         set query-interval <1-1800 seconds>
   next
end
```

Checking the PIM configuration

Use the following commands to check your PIM configuration:

```
get router info multicast config
get router info multicast igmp {groups | sources | interface <interface_name> | join}
get router info multicast pim {neighbour | interface <interface_name>}
```

Access lists

An access list is a list of IP addresses and the action to take for each one. Access lists provide basic route and network filtering. You can use the GUI or CLI to create an access list.

Using the GUI:

1. Go to *Router* > *Config* > *Access Lists* and select *Add Access List*.

2. Enter an identifier with one or more alphabetic characters.
3. Enter an optional description of the access list.
4. Select *Add*.
5. Select *Config Rules* in the row for the access list that you just created.

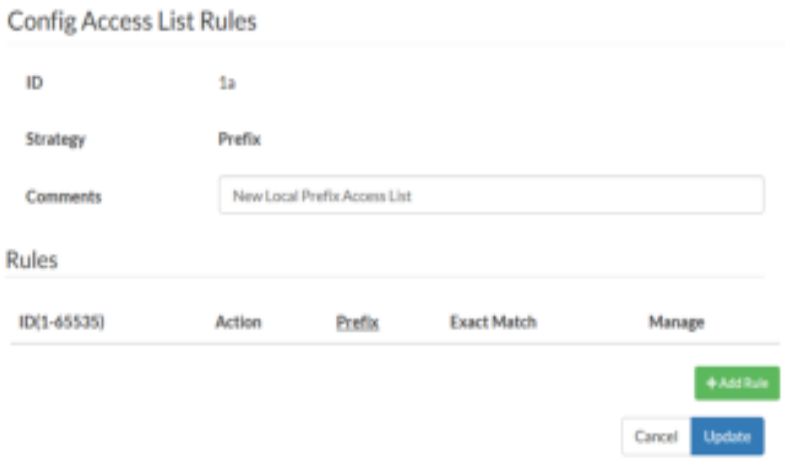

6. Select *Add Rule*.
7. Enter an identifier (1-65535), select *Deny* or *Permit* to specify if the rule will block or allow the specified IP addresses, and enter the prefix.
8. If you entered the complete IP address, select the *Exact Match* checkbox.
9. Select *Add Rule* if you want to add more rules.
10. After you have added all of the rules that you want in the access list, select *Update* to save the rules you added.

Using the CLI for IPv4 traffic:

```
config router access-list
   edit <access_list_name>
      set comments <comments>
      config rule
         edit <rule_int>
            set action {deny | permit}
            set prefix {<xxx.xxx.xxx.xxx> <xxx.xxx.xxx.xxx> | any}
            set wildcard <IP_address>
            set exact-match {enable | disable}
         end
      end
```

Using the CLI for IPv6 traffic:

```
config router access-list6
   edit <access_list_name>
      set comments <comments>
      config rule
         edit <rule_int>
            set action {deny | permit}
            set prefix6 {<xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx> | any}
            set exact-match {enable | disable}
         end
      end
```

Static and IPv6 static routing

For static routes in standalone, MCLAG, or layer-3 MCLAG network topologies, Fortinet recommends using a link monitor or BFD to detect whether the gateway is available.

This section covers the following topics:

- Remote access to the management port on page 453
- Equal cost multi-path (ECMP) routing on page 455

Remote access to the management port

To provide remote access to the management port, configure an IPv4 or IPv6 static route. Set the gateway address to the IPv4 or IPv6 address of the router.

Using the GUI for an IPv4 static route:

1. Go to *Router* > *Config* > *Static* and click *Add Route*.

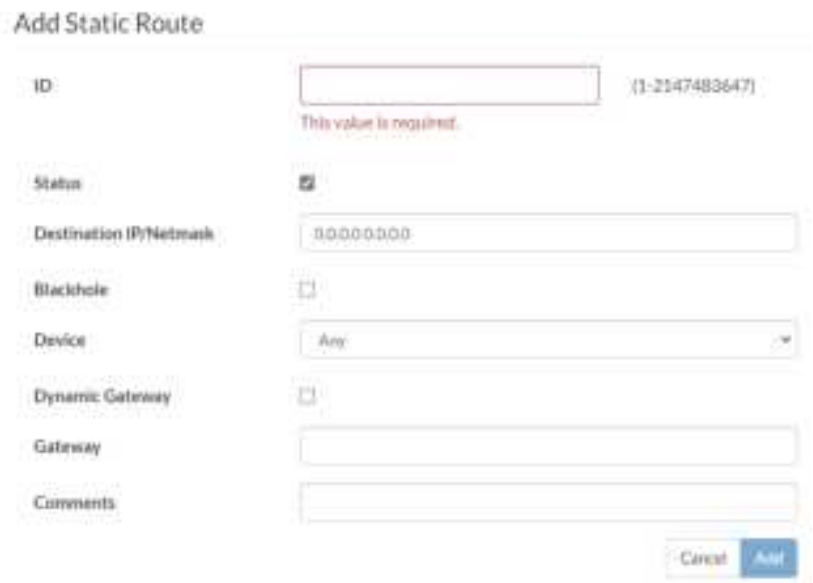

2. In the *ID* field, enter an identifier. This is a unique number to identify the static route.
3. Select the *Status* checkbox if it is not selected.
4. In the *Device* dropdown list, select *mgmt*.
5. In the *Gateway* field, enter the gateway router IPv4 address.
6. Click *Add*.

Using the CLI for an IPv4 static route:

```
config router static
   edit 1
      set device mgmt
```

```
    set gateway <router IPv4 address>
    set status enable
  next
end
```

Using the GUI for an IPv6 static route:

1. Go to *Router* > *Config* > *IPv6 Static* and click *Add Route*.

2. In the *Seq Num* field, enter an identifier. This is a unique number to identify the static route.
3. Select the *Enabled* checkbox if it is not selected.
4. In the *Device* dropdown list, select *mgmt*.
5. In the *Gateway* field, enter the gateway router IPv6 address.
6. Click *Add*.

Using the CLI for an IPv6 static route:

```
config router static6
  edit 1
    set device mgmt
    set gateway <router IPv6 address>
    set status enable
  next
```

```
end
```

Equal cost multi-path (ECMP) routing

ECMP is a forwarding mechanism that enables load-sharing of traffic to multiple paths of equal cost. An ECMP set is formed when the routing table contains multiple next-hop address for the same destination with equal cost. Routes of equal cost have the same preference and metric value. If there is an ECMP set for an active route, the switch uses a hash algorithm to choose one of the next-hop addresses. As input to the hash, the switch uses one or more of the following fields in the packet to be routed:

- Source IP
- Destination IP
- Input port

Configuring ECMP

The switch automatically uses ECMP to choose between equal-cost routes.

This configuration value is system-wide. The source IP address is the default value.

NOTE:

- There is a maximum of eight alternative paths (that is, ECMP paths).
- When you configure a static route with a gateway, the gateway must be in the same IP subnet as the device. Also, the destination subnet cannot match any of device IP subnets in the switch.
- When you configure a static route without a gateway, the destination subnet must be in the same IP subnet as the device.

Using the CLI:

```
config system settings
   set ip-ecmp-mode [ source-ip-based ] [ dst-ip-based ] [ port-based ]
end
```

Example ECMP configuration

The following is an example CLI configuration for ECMP forwarding.

In this configuration, ports 2 and 6 are routed ports. Interfaces I-RED and I-GREEN are routed VLAN interfaces. The remaining ports in the switch are normal layer-2 ports.

1. Configure native VLANs for ports 2, 6, and 9. Also configure the “internal” interface to allow native VLANs for ports 2, 6, and 9:

```
config switch interface
   edit port2
      set native-vlan 10
   edit port6
      set native-vlan 20
   edit port9
      set native-vlan 30
   edit internal
      set allowed-vlans 10,20,30
   end
```

2. Configure the system interfaces:

```
config system interface
   edit "internal"
      set type physical
   next
      edit "i-blue"
         set ip 1.1.1.1 255.255.255.0
         set allowaccess ping https http ssh snmp telnet
         set vlanid 10
         set interface internal
   next
      edit "i-red"
         set ip 172.16.11.1 255.255.255.0
         set allowaccess ping ssh telnet
         set vlanid 20
         set interface internal
   next
      edit "i-green"
         set ip 172.168.13.1 255.255.255.0
         set allowaccess ping https http ssh snmp telnet
         set vlanid 30
         set interface internal
   next
end
```

3. Configure static routes. This code configures multiple next-hop gateways for the same network:

```
config router static
   edit 1
      set device "mgmt"
      set gateway 10.105.0.1
      set status enable
   next
   edit 2
      set device "i-red"
      set dst 8.8.8.0/24
      set gateway 172.16.11.2
      set status enable
   next
   edit 3
      set device "i-green"
```

```
    set dst 8.8.8.0/24
    set gateway 172.168.13.2
    set status enable
  next
```

Viewing ECMP configuration

Display the status of the ECMP configuration using following command:

```
show system interface [ <system interface name> ]
```

Link probes

This section covers the following topics:

- Link monitor on page 457
- Bidirectional forwarding detection on page 460

Link monitor

You can create a probe to monitor the link to a server. The FortiSwitch unit sends periodic ping messages to test that the server is available. You can use both IPv4 and IPv6 addresses.

Starting in FortiSwitchOS 7.4.3, you can specify multiple servers for the link probe.

The link monitor configuration in the GUI and CLI was changed in FortiSwitchOS 7.4.2:

- You must now set the server IP address when you create a link monitor. This option supports remote peer monitoring.
- Ping is now the default protocol when using IPv4 addresses, instead of Address Resolution Protocol (ARP).

  When you upgrade to FortiSwitchOS 7.4.2, a link monitor's gateway IP address is moved to the server IP address. The link monitor protocol does not change.
- The gateway IP address (IPv4 and IPv6) is no longer mandatory. The gateway and interface are still required for the link monitor to modify the static routing table when a failure is detected. If the server and interface are configured but not the gateway IP address, the link monitor tries to automatically derive the gateway IP address from the static routing table.

The link monitor will only update static routes if the `set device` command under `config router static` is set.

Configuring the link probe

The link monitor is supported only in the default virtual routing and forwarding (VRF) instance.

Using the GUI:

1. Go to *Router* > *Config* > *Link Probes*.

2. Click *Add Probe* to create a new probe.

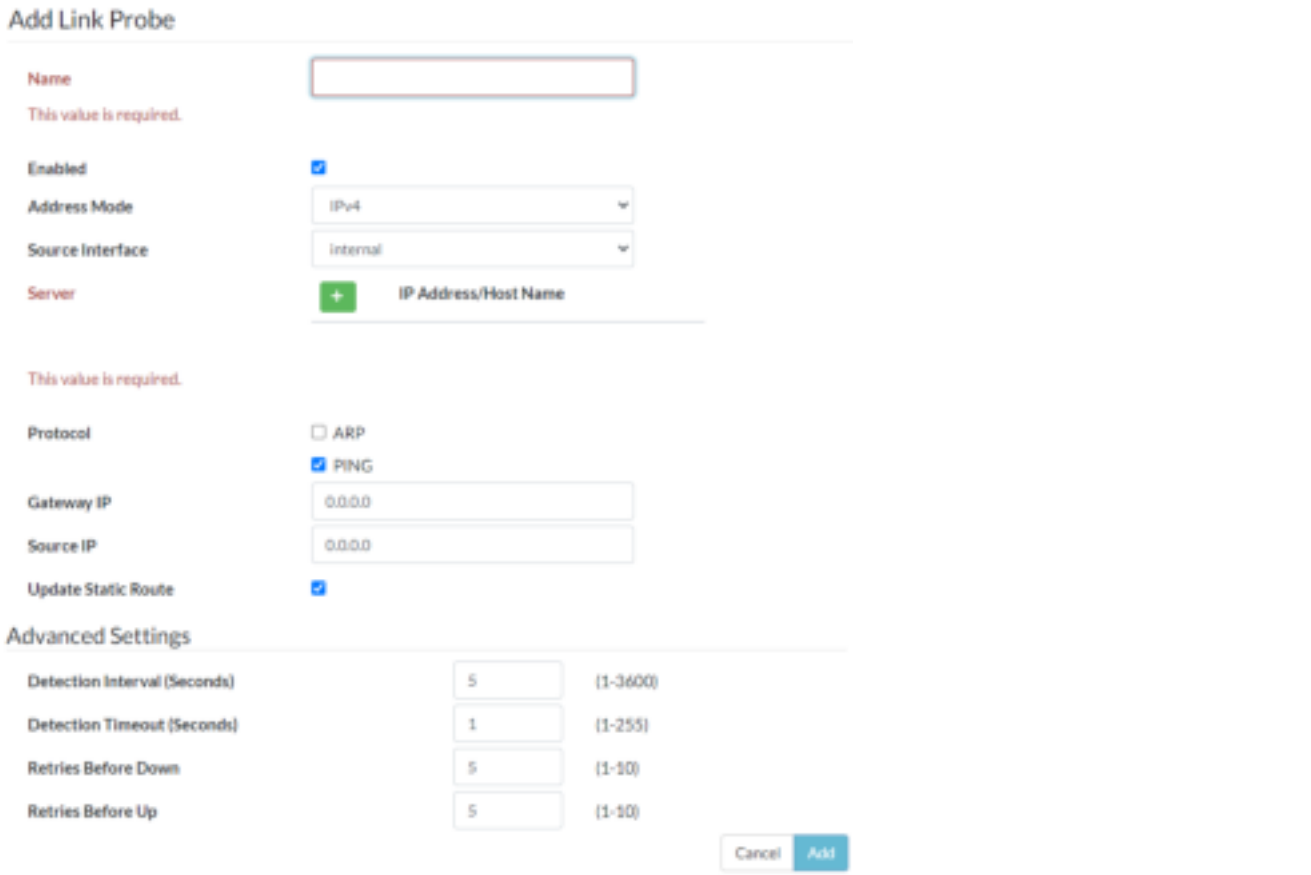

3. In the *Name* field, enter a name for the new probe.
4. By default, the *Enabled* checkbox is selected. If you do not want the probe to be active, clear the *Enabled* checkbox.
5. In the *Address Mode* dropdown list, select whether the address mode is IPv4 or IPv6.
6. In the *Source Interface* dropdown list, select the source interface.
7. For the *Server*, click + and then enter the IP address or host name of the server.
   To add another server, click + again.
8. For the *Protocol*, select the *PING* checkbox, *ARP* checkbox, or both.
   By default, the *PING* checkbox is selected.
9. In the *Gateway IP* field, enter an IP address for the *Gateway IP*.
   The gateway IP address is not mandatory, but it must be different from the source interface.

10. In the *Source IP* field, enter the source IP address.
11. By default, the *Update Static Route* checkbox is selected, but you can clear the checkbox if you do not want the static route updated.
12. In the *Detection Interval (Seconds)* field, enter the detection interval.
13. In the *Detection Timeout (Seconds)* field, enter the detection timeout.
14. In the *Retries Before Down* field, enter the number of retry attempts before the server is brought down.
15. In the *Retries Before Up* field, enter the number of retry attempts before bringing the server up.
16. Click *Add* to create the probe.

Using the CLI:

```
config system link-monitor
   edit <link monitor name>
      set addr-mode {ipv4 | ipv6}
      set srcintf <string>
      set server <IP_address1>, <IP_address2>, ...
      set protocol {arp | ping}
      set gateway-ip <IPv4 address>
      set gateway-ip6 <IPv6 address>
      set source-ip <IPv4 address>
      set source-ip6 <IPv6 address>
      set interval <integer>
      set timeout <integer>
      set failtime <integer>
      set recoverytime <integer>
      set update-static-route {enable | disable}
      set status {enable | disable}
   next
end
```

Variable	Description	Default
<link monitor name>	Enter the link monitor name.	No default
addr-mode {ipv4 | ipv6}	Select whether to use IPv4 or IPv6 addresses. The default is IPv4 addresses.	ipv4
srcintf <string>	Interface where the monitor traffic is sent.	No default
server <IP_ address1>, <IP_ address2>, ...	The IP address(es) of the server(s). Use a comma to separate multiple IP addresses.	No default
protocol {arp | ping}	Protocols used to detect the server. Select ARP or ping.	ping
gateway-ip <IPv4 address>	Gateway IPv4 address used to PING the server. This option is available only when `addr-mode` is set to `ipv4`.	0.0.0.0
gateway-ip6 <IPv6 address>	Gateway IPv6 address used to PING the server. This option is available only when `addr-mode` is set to `ipv6`.	::

Variable	Description	Default
source-ip <IPv4 address>	Source IPv4 address used in packet to the server. This option is available only when `addr-mode` is set to `ipv4`.	0.0.0.0
source-ip6 <IPv6 address>	Source IPv6 address used in packet to the server. This option is available only when `addr-mode` is set to `ipv6`.	::
interval <integer>	Detection interval in seconds. The range is 1-3600.	5
timeout <integer>	Detect request timeout in seconds. The range is 1-255.	1
failtime <integer>	Number of retry attempts before bringing the server down. The range is 1-10.	5
recoverytime <integer>	Number of retry attempts before bringing the server up. The range is 1-10.	5
update-static-route {enable | disable}	Enable or disable whether the static route is updated. The default is enabled.	enable
status {enable | disable}	Enable or disable link monitor administrative status. The default is enabled.	enable

Viewing the link monitor

To display the list of link probes:

Go to *Router* > *Config* > *Link Probes* or enter `get system link-monitor`.

To view the link monitor:

Go to *System* > *Link Monitor* or enter `diagnose sys link-monitor status all`.

Bidirectional forwarding detection

FortiSwitchOS v3.4.2 and later supports static bidirectional forwarding detection (BFD), a point-to-point protocol to detect faults in the datapath between the endpoints of an IETF-defined tunnel (such as IP, IP-in-IP, GRE, and MPLS LSP/PW).

BFD defines demand mode and asynchronous mode operation. The FortiSwitch unit supports asynchronous mode. In this mode, the systems periodically send BFD control packets to one another, and if a number of those packets in a row are not received by the other system, the session is declared to be down.

BFD packets are transported using UDP/IP encapsulation and BFD control packets are identified using well-known UDP destination port 3784 (**NOTE:** BFD echo packets are identified using 3785).

BFD packets are not visible to the intermediate nodes and are generated and processed by the tunnel end systems only.

Configuring BFD

Use the following steps to configure BFD:

1. Configure the following values in the system interface:
   - *Enable BFD*: Set to *enable* or set to *global* to inherit the global configuration value.
   - *Desired min TX interval*: This is the minimum interval that the local system would like to use between transmission of BFD control packets. Value range is 200 ms – 30,000 ms. Default value is 250.
   - *Required min RX interval*: This is the minimum interval that the local system can support between receipt of BFD control packets. If you set this value to zero, the remote system will not transmit BFD control packets. The value range is 200 ms – 30000 ms. The default value is 250.
   - *Detect multi*: This is the detection time multiplier. The negotiated transmit interval multiplied by this value is the Detection Time for the receiving system. The value range is 1 – 20. The default is 3.
2. Enable BFD in the static router configuration.

Using the CLI:

```
config system interface
   edit <system interface name>
      set bfd {enable| disable | global}
      set bfd-desired-min-tx <number of ms>
      set bfd-required-min-rx <number of ms>
      set bfd-detect-multi [1…20]
next
config router static
   edit 1
      set bfd enable
      set status enable
```

Viewing the BFD configuration

Using the GUI:

Go to *Router* > *Monitor* > *BFD Neighbor*.

Using the CLI:

To display the status of BFD sessions:

```
get router info bfd neighbor [ <IP address of neighbor>]
```

```
OurAddr         NeighAddr       LD/RD   State   Int
192.168.15.2    192.168.15.1    1/4     UP      vlan2000
192.168.16.2    192.168.16.1    2/2     UP      vlan2001
```

To filter the command output:

```
get router info bfd neighbor [<BFD_local_IPv4_address>] [<BFD_peer_interface>]
```

Virtual routing and forwarding

NOTE: You must have an advanced features license to use virtual routing and forwarding (VRF).

You can use the VRF feature to create multiple routing tables within the same router.

Use the following steps to configure VRF:

1. Creating a VRF instance on page 462
2. Assigning the VRF instance to a SVI on page 463
3. Assigning the VRF instance to a static route on page 463
4. Checking the VRF configuration on page 464

Starting in FortiSwitchOS 7.0.0, OSPF supports VRF. To use VRF with OSPF, create a VRF instance and then use the same VRF identifier in the `config vrf` commands under `config router ospf`.

Starting in FortiSwitchOS 7.2.1, you can configure port-based VRF for an RVI.

NOTE: This feature is supported only on the switch virtual interface (SVI).

Creating a VRF instance

You create a VRF instance by assigning a name and an identifier.

- The VRF name cannot match any SVI name.
- The VRF identifier is a number in the range of 1-1023, except for 252, 253, 254, and 255. You cannot assign the same VRF identifier to more than one VRF instance. After the VRF instance is created, the VRF identifier cannot be changed.

Using the GUI:

1. Go to *Router* > *Config* > *VRF*.
2. Click *Add VRF*.
3. In the *Name* field, enter a name for your VRF instance.
4. In the *VRFID* field, enter a VRF identifier.
5. Click *Add*.

Using the CLI:

```
config router vrf
     edit <string>
        set vrfid <VRF_ID>
     end
```

For example:

```
config router vrf
   edit vrfv4
      set vrfid 1
   next
   edit vrfv6
      set vrfid 2
   next
end
```

Assigning the VRF instance to a SVI

You assign the VRF instance to an SVI when you create the SVI. After the SVI is created, the VRF instance cannot be changed or unset.

You can assign the same VRF instance to more than one SVI. The VRF instance cannot be assigned to an internal SVI.

Using the GUI:

1. Go to *System* > *Network* > *Interface* > *VLAN*.
2. Click *Add VLAN*.
3. From the *VRF* dropdown list, select the VRF instance.
4. Configure the other fields as required.
5. Click *Add*.

Using the CLI:

```
config system interface
   edit <interface_name>
      set vrf <string>
   end
```

For example:

```
config system interface
   edit v40
      set vlanid 40
      set vrf vrfv4
   next
   edit v50
      set vlanid 50
      set vrf vrfv4
   next
end
```

Assigning the VRF instance to a static route

You assign the VRF instance to an IPv4 or IPv6 static route when you create the static route. After the static route is created, the VRF instance cannot be changed or unset.

You can assign the same VRF instance to more than one static route.

Using the GUI:

1. Go to *Router* > *Config* > *Static* or *Router* > *Config* > *IPv6 Static*.
2. Click *Add Route*.
3. From the *VRF* dropdown list, select the VRF instance.
4. Configure the other fields as required.
5. Click *Add*.

Using the CLI:

```
config router static
   edit <seq-num>
      set vrf <string>
   end

config router static6
   edit <seq-num>
      set vrf <string>
   end
```

For example:

```
config router static
   edit 1
      set device mgmt
      set gateway 192.168.0.10
      set status enable
      set vrf vrfv4
   end

config router static6
   edit 2
      set dst 5555::/64
      set gateway 4000::2
      set status enable
      set vrf vrfv6
   end
```

Checking the VRF configuration

Using the GUI:

Go to *Router* > *Config* > *Static*, *Router* > *Config* > *IPv6 Static*, or *System* > *Network* > *Interface* > *VLAN*.

Using the CLI:

- `get router info routing-table vrf <VRF_name>`
- `get router info6 routing-table vrf <VRF_name>`

Policy-based routing

NOTE: You must have an advanced features license to use policy-based routing.

Policy-based routing (PBR) allows users to define the next hop for packets based on the packet's source or destination IP addresses. You can specify the virtual routing and forwarding (VRF) instance that the next hop belongs to or the default VRF instance is used. You can assign the next hop to a next-hop group to use equal-cost multi-path (ECMP) routing.

Starting in FortiSwitchOS 7.2.3, you can use the GUI to configure policy-based routing.

To see which models support this feature, refer to the FortiSwitch feature matrix.

Configuring policy-based routing

Using the GUI:

1. Go to *Router* > *Config* > *Policy* > *Next Hop Groups* to configure the next-hop group using ECMP routing.
   a. Click *Add Next Hop Group*.
   b. In the *Name* field, enter the name of the VRF instance.
   c. To configure the next hop, click +.
   d. In the *IP* field, enter the IPv4 address of the next hop.
   e. From the *VRF* dropdown list, select the VRF instance.
   f. Click *Add* to save the next-hop group.
2. Go to *Router* > *Config* > *Policy* > *PBR Maps* to configure the PBR map.
   a. Click *Add PBR Map*.
   b. In the *Name* field, enter the name of the PBR map.
   c. In the *Comments* field, enter a description of the PBR map.
   d. To configure the PBR rule, click +.
   e. In the *Source* field, enter the source IPv4 address and mask.
   f. In the *Destination* field, enter the destination IPv4 address and mask.
   g. In the *IP* field, enter the IPv4 address of the next hop.
   h. From the *VRF* dropdown list, select the VRF instance that the next-hop address belongs to. If you do not select a VRF instance, the default VRF is used.
   i. From the *Group* dropdown list, select a next-hop group. This setting is used for ECMP.
   j. Click *Add* to save the PBR map.
3. Go to *Router* > *Config* > *Policy* > *Interfaces.*
   a. To configure the interface, click +.
   b. From the *Name* dropdown list, select the interface to configure.
   c. From the *PBR Map Name* dropdown list, select the PBR map.
   d. Click *Update* to save your changes.

Using the CLI:

```
config router policy
   config nexthop-group
      edit <name_of_next-hop_group>
         config nexthop
            edit <configuration_identifier>
               set nexthop-ip <IPv4_address>
               set nexthop-vrf-name <VRF_name>
            next
         end
      next
   end
   config pbr-map
      edit <PBR_map_name>
         set comments <string>
         config rule
            edit <rule_sequence_number>
               set src <IPv4_address_mask>
               set dst <IPv4_address_mask>
```

```
            set nexthop-ip <IPv4_address>
            set nexthop-vrf-name <VRF_name>
            set nexthop-group name <next-hop_group_name>
          next
        end
      next
    end
    config interface
      edit <interface_name>
        set pbr-map-name <PBR_policy_map_name>
      next
    end
end
```

Variable	Description	
config nexthop-group	Configure the next-hop group using ECMP routing.	
<name_of_next-hop_group>	Enter the name of the next-hop group.	No default
config nexthop	Configure the next hop.	
<configuration_identifier>	Enter the configuration identifier.	No default
nexthop-ip <IPv4_address>	Enter the IPv4 address of the next hop.	0.0.0.0
nexthop-vrf-name <VRF_name>	Enter the VRF instance name.	No default
config pbr-map	Configure the PBR map.	
<PBR_map_name>	Enter the name of the PBR map.	No default
comments <string>	Enter a descriptive comment.	No default
config rule	Configure the PBR rule.	
<rule_sequence_number>	Enter a rule identifier. The range of values is 1-10000.	No default
src <IPv4_address_mask>	Enter the source IPv4 address and mask.	0.0.0.0 0.0.0.0
dst <IPv4_address_mask>	Enter the destination IPv4 address and mask.	0.0.0.0 0.0.0.0
nexthop-ip <IPv4_address>	Enter the IPv4 address of the next hop.	0.0.0.0
nexthop-vrf-name <VRF_name>	Enter the name of the VRF instance that the next-hop address belongs to. If the name is not specified, the default VRF is used.	No default

Variable	Description	
nexthop-group name <next-hop_group_name>	Enter the next-hop group name. This setting is used for ECMP.	No default
config interface	Configure the interface.	
<interface_name>	Enter the name of the interface to configure.	No default
pbr-map-name <PBR_map_name>	Enter the name of the PBR map. The PBR map is created with the `config pbr-map` command.	No default

Example

This example creates the “pbrmap1” policy for vlan10, which is an ingress switch virtual interface (SVI). The policy has three rules:

- Rule 1 finds packets with a source address of 22.1.1.0/24 and forwards them to the next hop, 12.1.1.2, which belongs to the default VRF instance.
- Rule 2 finds packets with a destination address of 33.1.1.0/24 and forwards them to the ECMP route with the two next-hop IP addresses in the next-hop group . Both next hops belong to the default VRF instance.
- Rule 3 finds packets with a destination address of 11.1.1.0/24 and forwards them to the next hop, 13.1.1.2, which belongs to the “vrfv4” VRF instance.

```
config router policy
   config nexthop-group
      edit "nhgroup1"
         config nexthop
            edit 1
               set nexthop-ip 12.1.1.4
            next
            edit 2
               set nexthop-ip 12.1.1.5
            next
         end
      next
   end
   config pbr-map
      edit "pbrmap1"
         config rule
            edit 1
               set src 22.1.1.0 255.255.255.0
               set nexthop-ip 12.1.1.2
            next
            edit 2
               set dst 33.1.1.0 255.255.255.0
               set nexthop-group-name "nhgroup1"
            next
            edit 3
               set src 11.1.1.0 255.255.255.0
               set nexthop-ip 13.1.1.2
               set nexthop-vrf-name "vrfv4"
```

```
            next
         end
      next
   end
   config interface
      edit "vlan10"
         set pbr-map-name "pbrmap1"
      next
   end
end
```

Checking the PBR configuration

Use the following command get information about the specified PBR rule. If the PBR rule is not specified , all rules are returned.

```
get router info pbr map ["<map-name> <sequence-number> <interface-name>"]
```

For example:

```
get router info pbr map "pbrmap1 1 vlan10"
```

Use the following command to get information about the PBR next-hop group:

```
get router info pbr nexthop-group
```

Key chains

You can create, edit, or delete key chains in the GUI and CLI.

A key chain is a list of one or more authentication keys including its lifetime, which is how long each key is valid. Key chains are used for MD5 authentication for Routing Information Protocol (RIP) or Intermediate System to Intermediate System Protocol (IS-IS) routing.

Creating a key chain

You must create a key chain first before you can use MD5 authentication with RIP version 2 or IS-IS routing.

Using the GUI:

1. Go to *Router* > *Config* > *Key Chains*.

2. Click *Add Key Chain*.

3. In the *Name* field, enter a name for the new key chain.
4. Under *Keys*, click +.
5. In the *ID* column, enter an identifier for the new key chain.
6. In the *Accept Lifetime* column, click to select the range of dates that the received authentication key is valid.
7. In the *Send Lifetime* column, click to select the range of dates that the sent authentication key is valid.
8. In the *Key String* column, enter a password string for the key.
9. Click *Add*.
   The new key chain is listed on the *Key Chains* page.

Using the CLI:

```
config router key-chain
   edit <key_chain_name>
      config key
         edit <key_chain_int>
            set key-string <key_str>
            set accept-lifetime <START> <END>
            set send-lifetime <START> <END>
         next
      end
   next
end
```

For example:

```
config router key-chain
   edit "newkeychain"
      config key
         edit 1
```

```
            set accept-lifetime 12:00:00 09 01 2023 12:00:00 09 01 2024
            set key-string "keychain1"
            set send-lifetime 12:00:00 09 01 2025 12:00:00 09 01 2026
        next
    end
  next
end
```

Editing a key chain

1. Go to *Router* > *Config* > *Key Chains*.
2. In the row of the key chain that you want to change, click *Edit*.

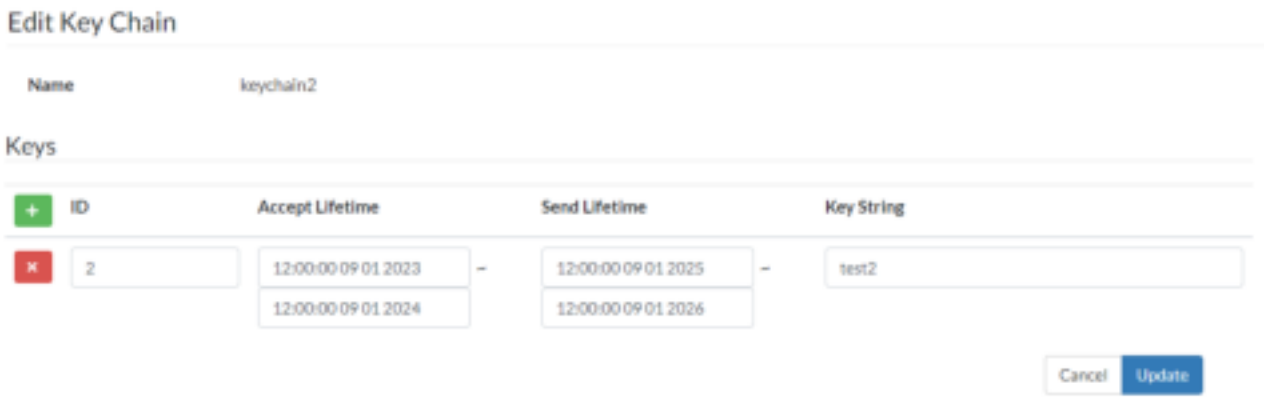

3. When you are finished making your changes, click *Update*.

Deleting a key chain

1. Go to *Router* > *Config* > *Key Chains*.
2. Select a key chain.
3. Click *Delete*.
4. In the *Confirm Key Chain Deletion* dialog, click *Delete*.

Diagnostic

Use the *Route Diagnostic* page to display a summary of existing routes for a specific IP address or host name and to view the network hops to the specified IP address or host name.

To display route diagnostics in the GUI:

1. Go to *Router* > *Diagnostic*.

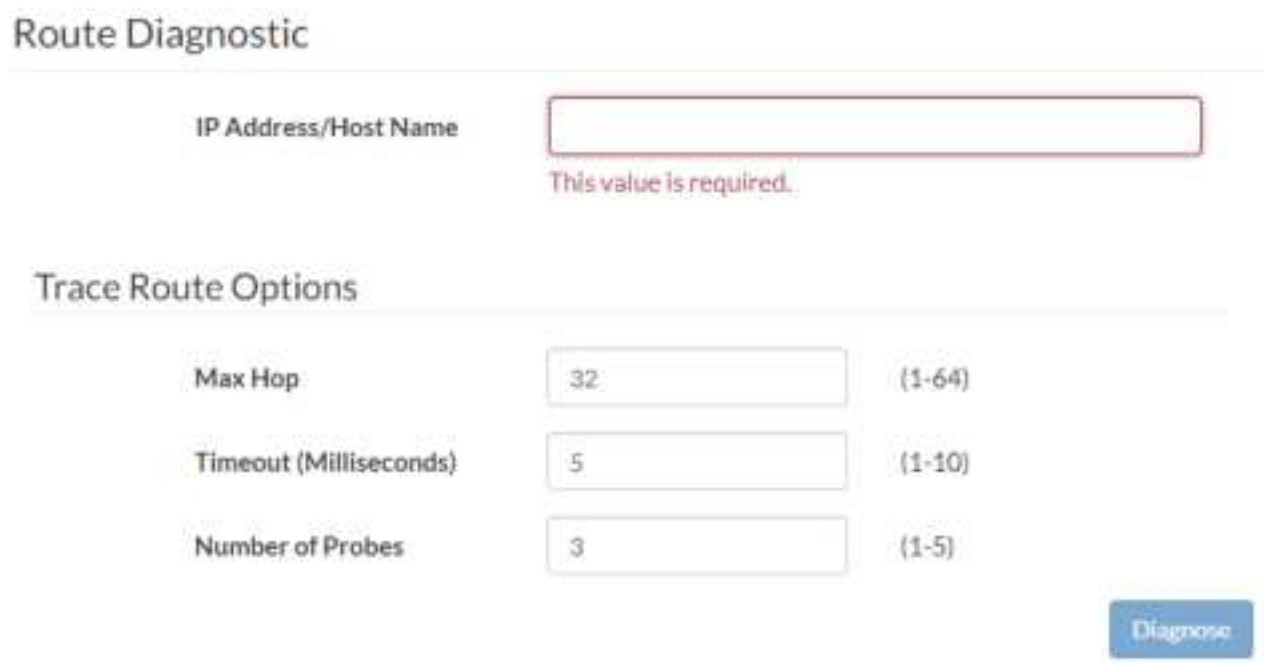

2. In the *IP Address/Host Name* field, enter an IPv4 address or host name.
3. You can use the default values for the *Trace Route Options* or change them:
   - In the *Max Hop* field, enter the maximum number of hops that the trace route can take.
   - In the *Timeout (Milliseconds)* field, enter how many milliseconds a route can take before the trace route is stopped.
   - In the *Number of Probes* field, enter the maximum number of probes to use to trace the route.

4. Click *Diagnose*.

Do not close your browser while the route diagnostic is running.

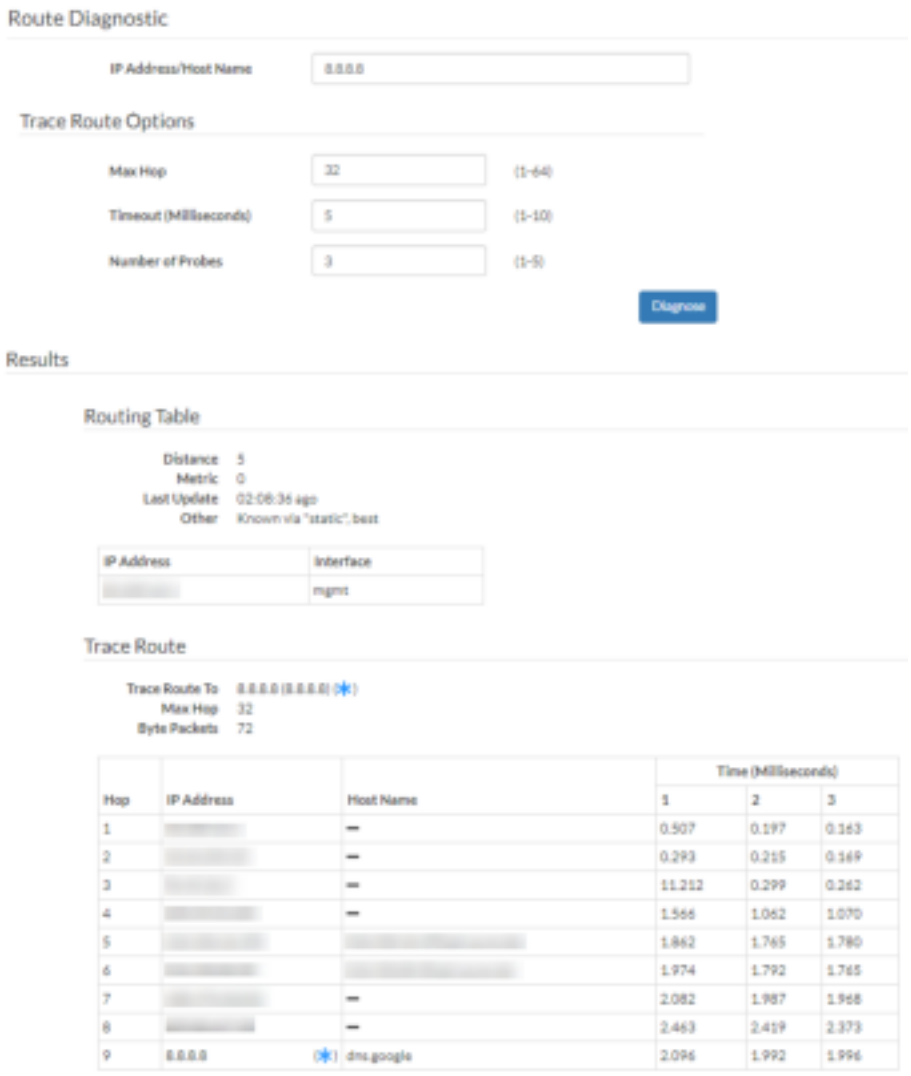

To display route diagnostics in the CLI:

```
get router info routing-table <IPv4_address>
execute traceroute <IPv4_address_or_host_name> <maximum_number_of_hops> <number_of_probes>
     <maximum_number_of_milliseconds>
```

For example:

```
get router info routing-table 8.8.8.8
execute traceroute 8.8.8.8 16 5 15
```

Multi-traceroute

Starting in FortiSwitchOS 7.2.0, you can use the CLI for multiple path traceroute, which allows you to find all the routers that perform load balancing between the FortiSwitch unit and destination.

Starting in FortiSwitchOS 7.2.1, you can use the *Router* > *Multi-Traceroute* page for multiple path traceroute.

To run multiple path traceroute in the GUI:

1. Go to *Router* > *Multi-Traceroute*.

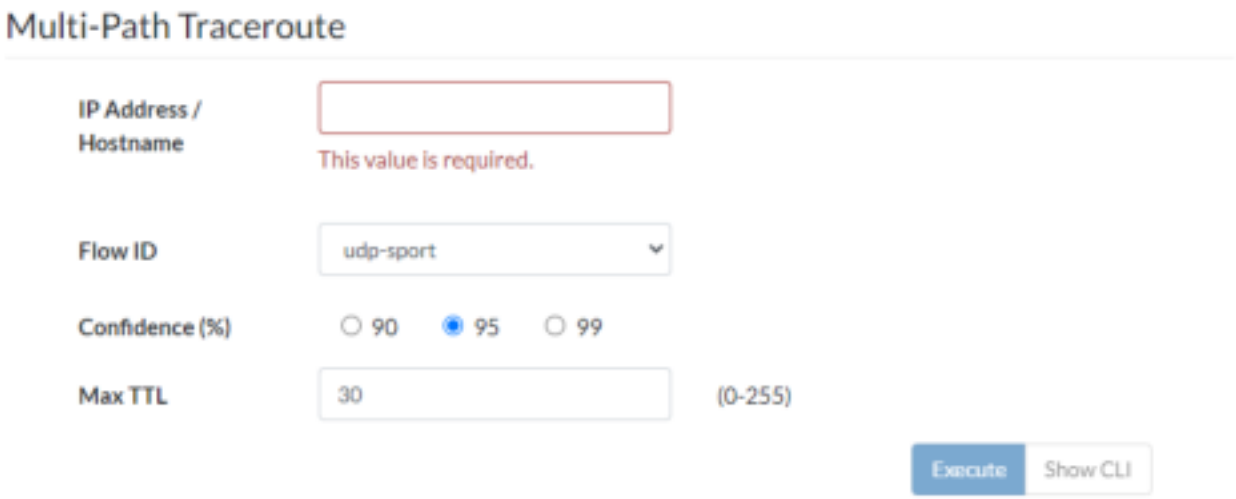

2. In the *IP/Address/Hostname* field, enter the IPv4 or IPv6 address or the hostname to test the connection to.
3. In the *Flow ID* field, select the flow identifier to use.
   If you entered an IPv4 address to test, you can select *icmp-chk*, *icmp-dst*, *tcp-dst*, *tcp-sport*, *udp-dst*, or *udp-sport*. The default value is *udp-sport*.
   If you entered an IPv6 address to test, you can select *icmp-chk*, *icmp-dst*, *icmp-fl*, *icmp-tc*, *tcp-dst*, *tcp-fl*, *tcp-sport*, *tcp-tc*, *udp-dst*, *udp-fl*, *udp-sport*, or *udp-tc*. The default value is *udp-sport*.
4. Select *90*, *95*, or *99* for the confidence percentage.
5. In the *Max TTL* field, enter the maximum number of hops to test.
   The range of values is 0-255. The default is 30.

6. Click *Execute*.

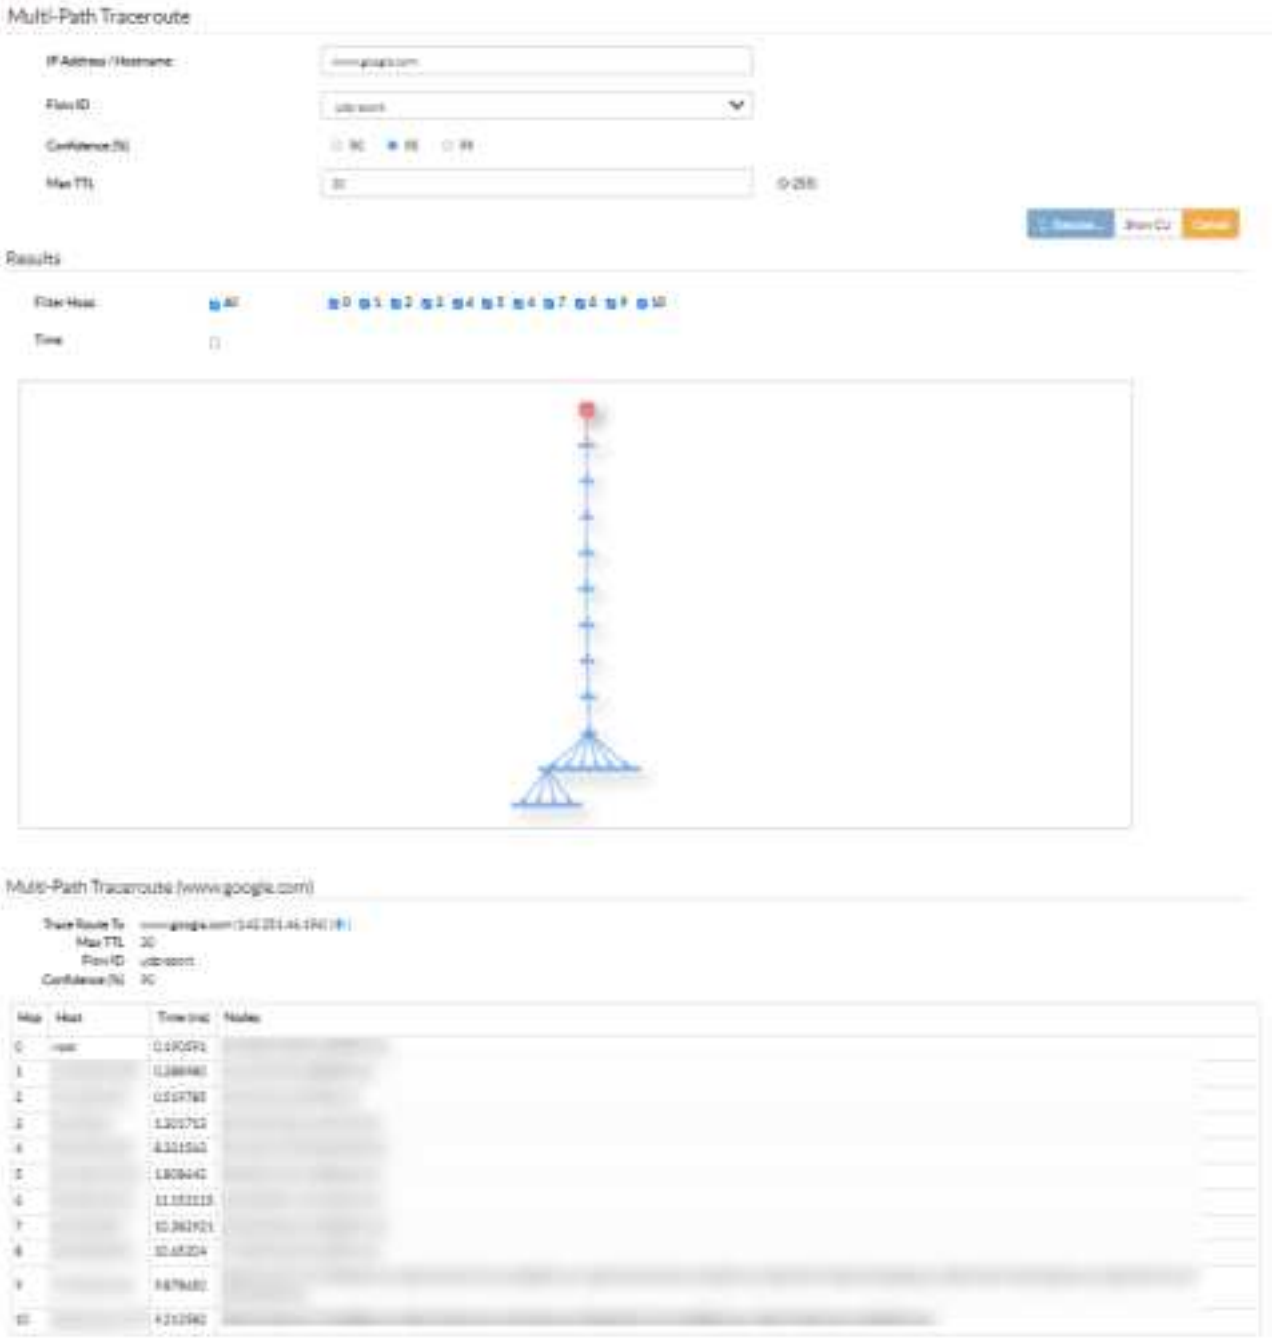

7. Click *Show CLI* if you want to see the results in the CLI.
8. Click *Cancel* if you want to stop the multiple path traceroute.

To run multiple path traceroute in the CLI:

```
execute mtraceroute <IP_address> <confidence_level> <flow_ID> <maximum_hops>
```

Variable	Description
<IP_address>	Enter the IP address to test the connection to.
<confidence_level>	Select the confidence level in percent. You can select `90`, `95`, or `99`. The default value is `95`.
<flow_ID>	Select the flow identifier to use. If you selected an IPv4 address to test, you can select `icmp-chk`, `icmp-dst`, `udp-sport`, `udp-dst`, `tcp-sport`, or `tcp-dst` as the flow identifier with `udp-sport` as the default value. If you selected an IPv6 address to test, you can select `icmp-chk`, `icmp-dst`, `icmp-fl`, `icmp-tc`, `udp-sport`, `udp-dst`, `udp-fl`, `udp-tc`, `tcp-sport`, `tcp-dst`, `tcp-fl`, or `tcp-tc` as the flow identifier with `udp-sport` as the default value.
<maximum_hops>	Enter the maximum number of hops to test. The range of values is 0-255. The default is 30.

For example:

```
S108FFTV21000010 # execute mtraceroute 1.2.3.4 90 icmp-chk 50
Run mtraceroute to 1.2.3.4 - max-ttl: 50, flow-id: icmp-chk, confidence: 90
0  root:  10.105.201.133 (0.767220 ms)
1  10.105.201.133:  192.168.201.1 (0.296219 ms)
2  192.168.201.1:  10.64.254.33 (0.306219 ms)
3  10.64.254.33:  96.45.36.3 (0.501219 ms)
4  96.45.36.3:  *
...
```

ARP table

The *ARP Table* page lists the IP address, number of minutes that the ARP entry has been in the ARP table, MAC address, and interface for each ARP table entry. The ARP table entries are manually added with the `config system artp-table` command or provided by dynamic ARP inspection (DAI).

To view the ARP table entries in the GUI:

Go to *Router* > *ARP Table*.

To view the ARP table entries in the CLI:

```
get system arp
```

Monitor

Logging allows you to review all router activity.

NOTE: Router logs are available only on supported platforms if you have the advanced features license.

To enable router logging:

1. Go to *Log* > *Config*.
2. Under *Event Type*, select *Enable* and *Router*.
3. Select *Apply*.

To view router logs:

Go to *Router* > *Monitor* > *Routing* or *Router* > *Monitor* > *IPv6 Routing*.

Starting in FortiSwitchOS 7.2.2, you can display IPv4 and IPv6 routes by VRF instance on the *Router* > *Monitor* > *Routing* and *Router* > *Monitor* > *IPv6 Routing* pages.

Starting in FortiSwitchOS 7.4.0, the Route Monitor and IPv6 Route Monitor display the routes graphically, as well as in a table. Hover your cursor over a route to highlight it and to see the destination and interface of the route.

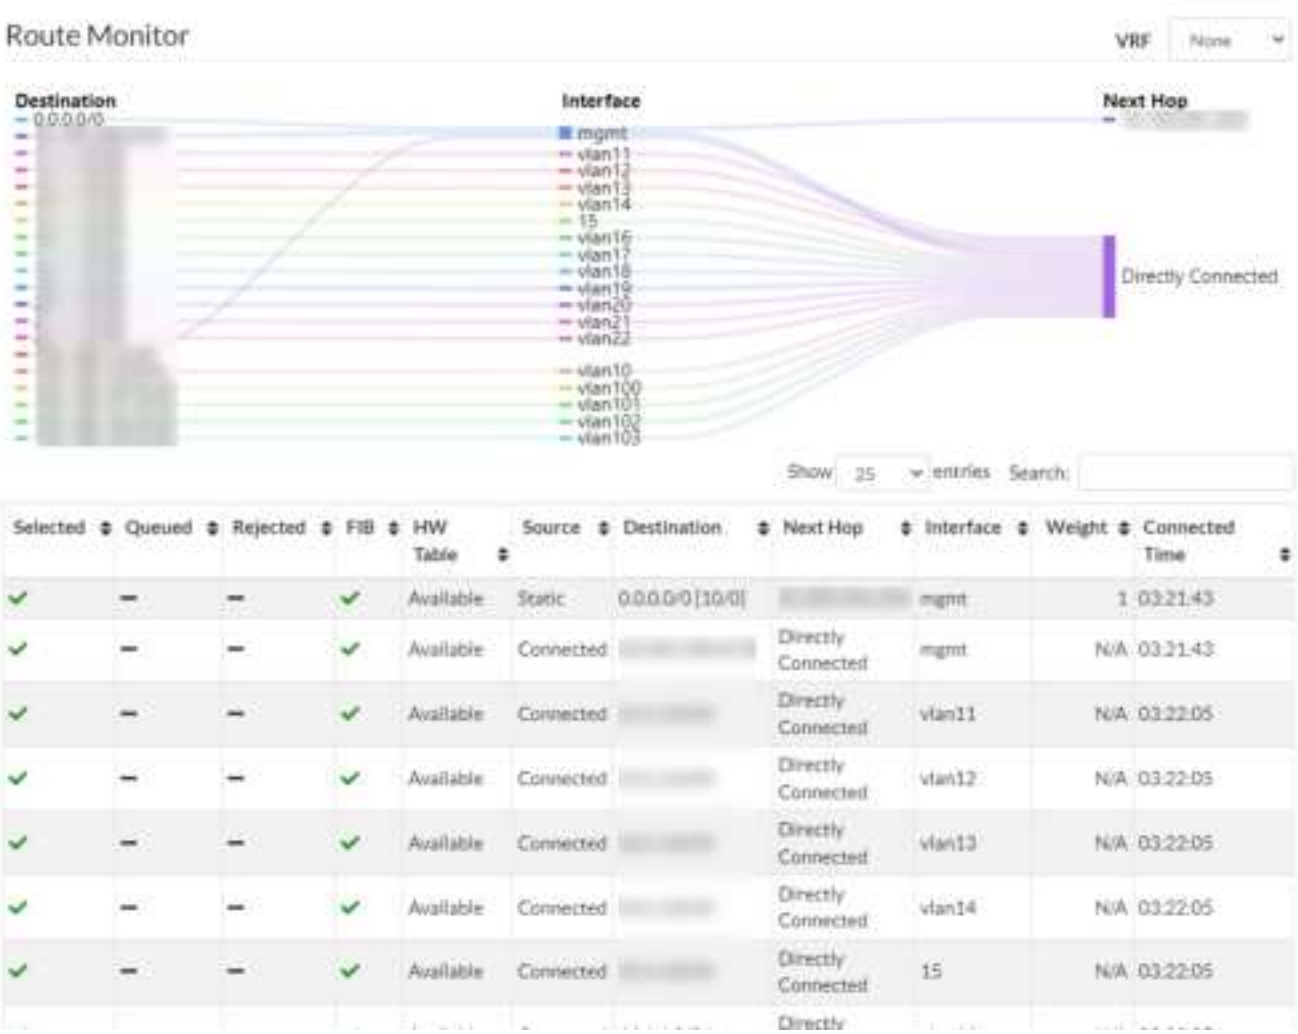
Route Monitor
VRF
None
Destination
0.0.0.0/0
Interface
mgmt
vlan11
vlan12
vlan13
vlan14
15
vlan16
vlan17
vlan18
vlan19
vlan20
vlan21
vlan22
vlan10
vlan100
vlan101
vlan102
vlan103
Next Hop
Directly Connected
Show 25 entries Search:
Selected
Queued
Rejected
FIB
HW Table
Source
Destination
Next Hop
Interface
Weight
Connected Time
Available Static 0.0.0.0/0 [10/0] mgmt 1 03:21:43
Available Connected Directly Connected mgmt N/A 03:21:43
Available Connected Directly Connected vlan11 N/A 03:22:05
Available Connected Directly Connected vlan12 N/A 03:22:05
Available Connected Directly Connected vlan13 N/A 03:22:05
Available Connected Directly Connected vlan14 N/A 03:22:05
Available Connected Directly Connected 15 N/A 03:22:05

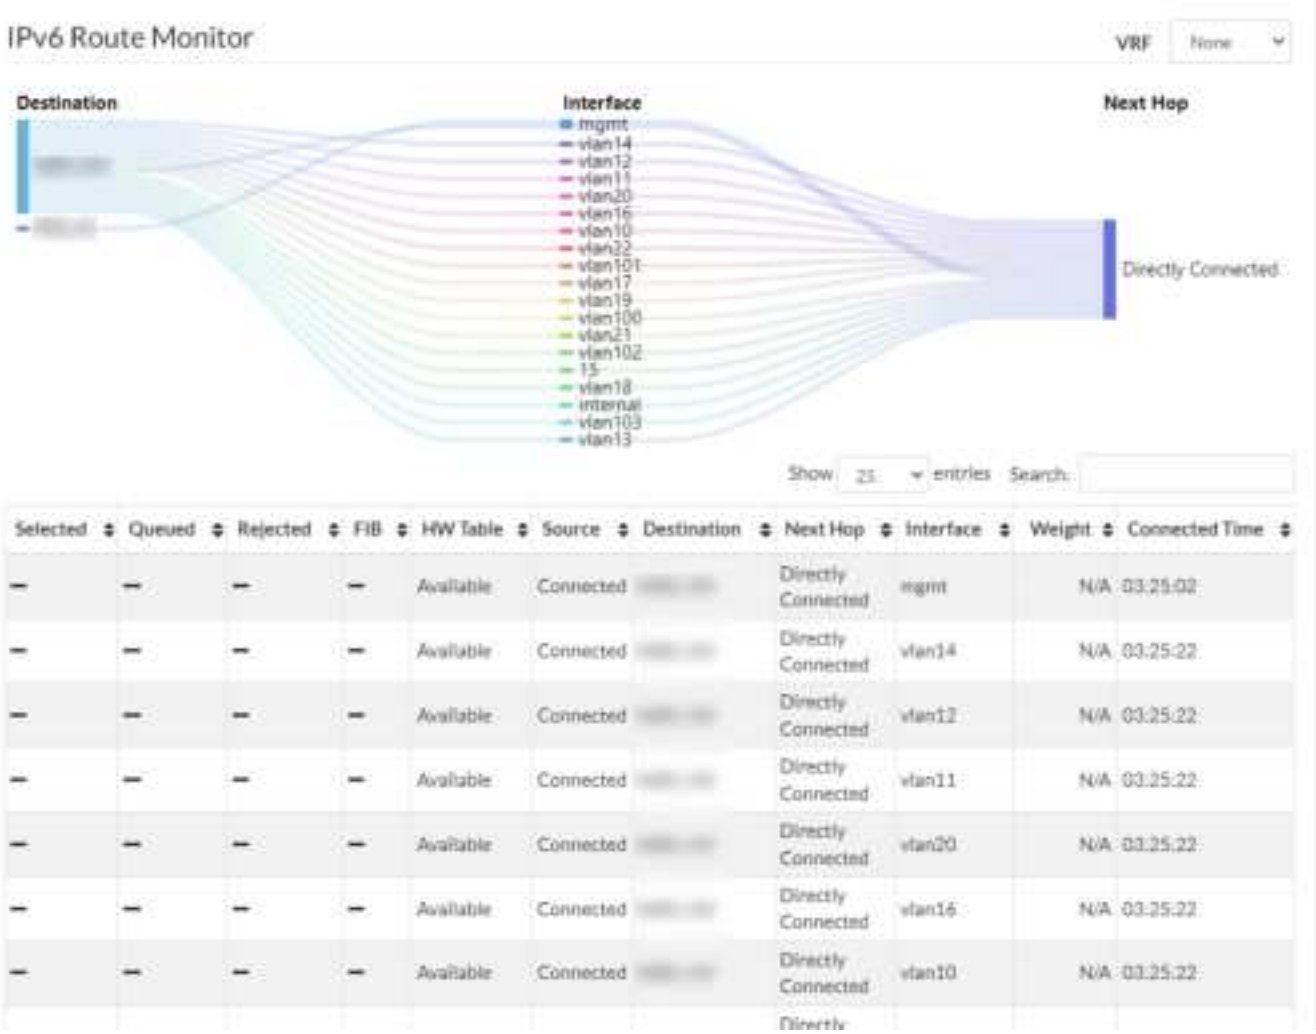

Selected	Queued	Rejected	FIB	HW Table	Source	Destination	Next Hop	Interface	Weight	Connected Time
–	–	–	–	Available	Connected	[redacted]	Directly Connected	mgmt	N/A	03:25:02
–	–	–	–	Available	Connected	[redacted]	Directly Connected	vlan14	N/A	03:25:22
–	–	–	–	Available	Connected	[redacted]	Directly Connected	vlan12	N/A	03:25:22
–	–	–	–	Available	Connected	[redacted]	Directly Connected	vlan11	N/A	03:25:22
–	–	–	–	Available	Connected	[redacted]	Directly Connected	vlan20	N/A	03:25:22
–	–	–	–	Available	Connected	[redacted]	Directly Connected	vlan16	N/A	03:25:22
–	–	–	–	Available	Connected	[redacted]	Directly Connected	vlan10	N/A	03:25:22
							Directly			

Log

FortiSwitchOS provides a robust logging environment that enables you to monitor, store, and report traffic information and FortiSwitch events, including attempted log ins and hardware status. Depending on your requirements, you can log to a number of different hosts.

To configure event logging using the GUI:

1. Go to *Log* > *Config*.

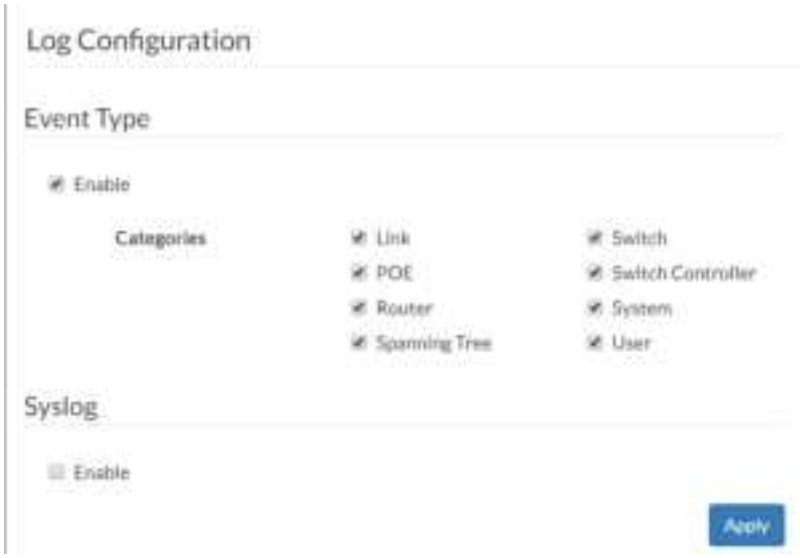

2. Under *Event Type*, select *Enable*.
3. Under *Event Type*, select the categories of events that you want logged.
4. Select *Apply*.

To configure event logging using the CLI:

```
config log eventfilter
   set event {enable | disable}
   set link {enable | disable}
   set poe {enable | disable}
   set router {enable | disable}
   set spanning_tree {enable | disable}
   set switch {enable | disable}
   set switch_controller {enable | disable}
   set system {enable | disable}
   set user {enable | disable}
end
```

To view the event logs in the GUI:

1. Go to *Log* > *Entries*.
2. From the *Subtype* dropdown list, select the type of log entries to view.

3. From the *Level* dropdown list, select the severity of events to view.
4. From the *User* dropdown list, select which user or process generated the log entry.
5. From the *User Interface* dropdown list, select the IP network service that applies to the log entry.
6. From the *Action* dropdown list, select the event to view.
7. From the *Status* dropdown list, select the event result to view.

The MAC event log does not report any MAC changes on the port when 802.1X authentication is enabled.

To view the event logs in the CLI:

```
show log eventfilter
```

Syslog server

Sysog is an industry standard for collecting log messages for off-site storage. You can send logs to a single syslog server. The syslog server can be configured in the GUI or CLI. Reliable syslog (RFC 6587) can be configured only in the CLI.

To configure a syslog server in the GUI:

1. Go to *Log* > *Config*.

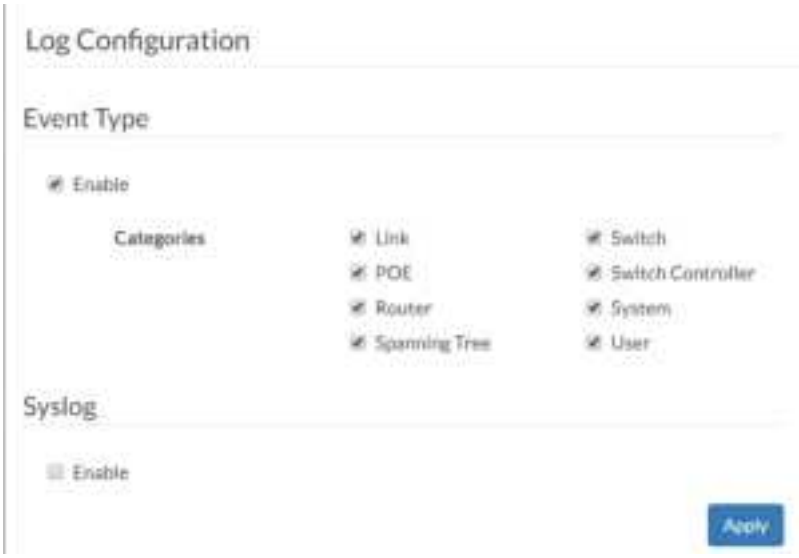

2. Under *Syslog*, select *Enable*.
3. Select the severity of events to log.
4. Enter the IP address or fully qualified domain name in the *Server* field.
5. Enter the port number that the syslog server will use. By default, port 514 is used.
6. Select *Apply*.

To configure a syslog server in the CLI:

```
config log syslogd setting
   set status enable
   set server <IP address or FQDN of the syslog server>
   set port <port number that the syslog server will use for logging traffic>
   set facility <facility used for remote syslog>
   set source-ip <source IP address of the syslog server>
end
```

For example, to set the source IP address of a syslog server to have an IP address of 192.168.4.5:

```
config log syslogd setting
   set status enable
   set source-ip 192.168.4.5
end
```

To configure a reliable syslog server in the CLI:

```
config log syslogd setting
   set status enable
   set server <IP address or FQDN of the syslog server>
   set mode reliable
   set port <port number that the syslog server will use for logging traffic>
   set enc-algorithm {high | high-medium | low}
   set certificate <certificate_used_to_communicate_with_syslog_server>
end
```

For example:

```
config log syslogd setting
   set status enable
   set source-ip 192.168.4.5
   set mode reliable
   set port 6514 // This is the default port used for reliable syslog.
   set enc-algorithm high-medium
   set certificate "155-sub-client"
end
```

Deployment scenario

Working configuration for PC and phone for 802.1X authentication using MAC

Summary

1. Configure all devices.
   - PC
   - Phone
   - FortiSwitch
   - FortiAuthenticator
   - DHCP server
2. Authenticate phone using MAB and using LLDP-MED.
3. Authenticate PC using EAP 802.1X.

A. Configure all devices

I. Configure the PC, phone, FortiSwitch, FortiAuthenticator [RADIUS server], and DHCP server)

Phone configuration (file: macmode_phone_pc_ping_work)

i. On the phone, enable the WAN port and leave the VLAN ID at the default to allow LLDP-Med (Policy) designate for voice VLAN assignment.
ii. On the phone, enable the LAN port and assign the VLAN ID for data matching the RADIUS VLAN assignment.

PC configuration

i. Install the supplicant software.
ii. Launch the supplicant software, type the user name and password, and enable DHCP on the interface.

FortiSwitch configuration

1. Configure the LLDP profile for voice.

```
# show switch lldp
config switch lldp profile

edit "pexa" <<<<<<<<<<<<<<<<
```

```
set 802.1-tlvs port-vlan-id
   config med-network-policy
      edit "voice"
         set status enable
         set vlan 21
      next
      edit "voice-signaling"
         set status enable
         set vlan 31
      next
      edit "guest-voice"
      next
      edit "quest-voice-signaling"
      next
      edit "softphone-voice"
         set status enable
         set vlan 41
      next
      edit "video-conferencing"
      next
      edit "streaming-video"
      next
      edit "video-signaling"
      next
   end
set med-tlvs inventory-management network-policy
```

2. Apply the LLDL profile on a dot1x port.

```
# show switch physical-port port4
config switch physical-port

   edit "pexa" <<<<<<<<<<<<<<<<
   set lldp-profile "pexa"
   set speed auto
   next
end
```

3. Configure a user group.

```
# show user group
config user group

   edit "Corp_Grp_10"
   set member "FAC_LAB"
   next
end
```

4. Configure the RADIUS server.

```
# show user radius
config user radius

   edit "FAC_LAB" <<<<<<<<
   set secret
```

```
      ENCW82jBg06XhKD/4Dugqm8QF2f7D1B4bfFdDSZaLUQPwZXv4F8zMc5sWHRl9suwmbmzNnAnyqPaarAYcSL
           uT8kVjFSRO0znx+TXVWTqdSeLCpbMv
           +HYFNOHMbYlfES8wTYYD40InCgrYr2johvr2vfa5KG4g8XMwKSIM0LurR//1WqT0fH
      set server
      next
   end
```

5. Configure port security on the dot1x port.

   a. Configure mac-mode port-security.
   b. Add voice VLAN on allowed list (for example, 21).
   c. Apply the security group.

   Interface port4 configuration:

```
   # show switch interface port4
   config switch interface

      edit "port4"
      set allowed-vlans 20-21,31,41
      set security-groups "Corp_Grp_10"
      set snmp-index 4
   configure port-security
      set auth-fail-vlan disable
      set guest-auth-delay 120
      set guest-vlan disable
      set mac-auth-bypass enable
      set port-security-mode 802.1X-mac-based
      set radius-timeout-overwrite disable
      set auth-fail-vlanid 40
      set guest-vlanid 30
   end
```

RADIUS configuration

MAB Authentication:

- Add phone MAC address to MAB list.

802.1X Authentication

1. Create a local user.
2. Create a user group with "Attributes" and enable PEAP and MSChapv2.

DHCP configuration

1. On the DHCP server, configure a pool for phone and a pool for the PC.

```
   !
   ip dhcp pool PC
   network 10.1.1.0 255.255.255.0
   default-router 10.1.1.1
   dns-server 10.1.1.1
   !
   ip dhcp pool PC
   network 20.1.1.0 255.255.255.0
```

```
default-router 20.1.1.1
dns-server 20.1.1.5
```

2. Configure exclude lists for pools for both gateway and DNS.

```
ip dhcp excluded-address 20.1.1.1 20.1.1.1.5
<<<<gateway and dns server
ip dhcp excluded-address 10.1.1.1 10.1.1.1.5
<<<<gateway and dns server
!
ip dhcp pool PC
network 20.1.1.0 255.255.255.0
default-router 20.1.1.1
dns-server 20.1.1.5
```

3. Configure the switch port VLAN interface as a gateway for the phone.

```
# show run
Building configuration

Current configuration
!
interface vlan21 <<<<<<
ip address 20.1.1.1
end
```

4. Configure the switch port VLAN interface as a gateway for the PC.

```
# show run
Building configuration

Current configuration
!
interface vlan10 <<<<<<
ip address 10.1.1.1
end

#
```

5. Configure the l2 port and associate the voice VLAN.

```
# show run
Building configuration

Current configuration
!
interface GigabitEthernet g1/0/1 <<<<<<
switchport access vlan 21
switchport trunk encapsulation dot1q
switchport trunk all
switchport mode trunk
end
```

6. Configure the l2 port and associate the data VLAN.

```
# show run
Building configuration
```

```
Current configuration
!
interface GigabitEthernet g1/0/2 <<<<<<
switchport access vlan 10
switchport trunk encapsulation dot1q
switchport trunk all
switchport mode trunk
end
```

II. Connect a link between the FortiSwitch unit and the DHCP server and assign matching VLAN for the phone for both ports

III. Connect a link between the FortiSwitch unit and the DHCP server and assign a matching VLAN for the PC for both ports

B. Authenticate phone using MAB

1. Connect the phone to the switch to authenticate with RADIUS through the MAB (mac-bypass).
2. Once authenticated:
   a. On the FortiSwitch unit, verify that the port is authorized and that the voice VLAN is on the allowed list.

```
# diagnose switch 8 status
Signal 10 received - config reload scheduled

wrdapd_hostapd_dump_state_console Hostapd own address 90:6c:ac:18:6f:2f
dump_diag:1:
receive dump diagnostic 802_1x/MAB sessions. ifname :port4: dump_diag:1:

port4 : Mode: mac-based (mac-by-pass enable)
        Link: Link up
        Port State: authorized ( ) <<<<<<
        Native Vlan : 1
        Allowed Vlan list: 1,10,20-21,31,41 <<<<<<
        Untagged Vlan list:
        Guest VLAN:

        Client MAC Type Vlan Dynamic-Vlan
        68:f7:28:fb:c0:0f 802.1x 1 10
<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<phone

        Sessions info:
        68:f7:28:fb:c0:0f Type=802.1x,PEAP,state=AUTHENTICATED
        params:reAuth=3600
        00:a8:59:d8:f1:f6 Type=MAB,,state=AUTHENTICATED
        params: reAuth=3600

        edited on: 2016-11-29 17:25

        edited on: 2016-11-29 17:59
```

b. On the FortiSwitch unit, verify that the lldp neighbor detail accurately reflects the phone and voice VLAN designation.

```
        Neighbor learned on port4 by LLDP protocol
Last change 140 seconds ago
Last packet received 13 seconds ago

Chassis ID: 20.1.1.10 (ip) <<<<<<<<<<
System Name: FON-670i
System Description
V12.740.335.12.B

Time To Live: 60 seconds
System Capabilities: BT
Enabled Capabilities: BT
MED type: Communication Device Endpoint (Class III)
MED Capabilities: CP
Management IP Address: 20.1.1.10

Port ID: 00:a8:59:d8:f1:f6 (mac) <<<<<<<<<<<<<<
Port description: WAN Port 10M/100M/1000M
IEEE802.3, Power via MDI:
Power devicetype: PD
PSE MDI Power: Not Supported
PSE MDI Power Enabled: No
PSE Pair Selection: Can not be controlled
PSE power pairs: Signal
Power class: 1
Power type: 802.3at off
Power source: Unknown
Power priority: Unknown
Power requested: 0
Power allocated: 0
LLDP-MED, Network Policies:
voice: VLAN: 21 (tagged), Priority: 0 DSCP: 0 <<<<<<<<<<<<
voice-signaling: VLAN: 21 (tagged), Priority: 0 DSCP: 0
streaming-video: VLAN: 21 (tagged), Priority: 0 DSCP: 0

# Checking STA 00:a8:59:d8:f1:f6 inactivity:
Station has been active
```

c. On the phone, verify that the DHCP address is assigned.

d. On the DHCP server, check binding and ping from gateway to verify that the phone is reachable.

```
# show ip dhcp binding
IP address Client-ID/ Lease expiration Type
Hardware address
20.1.1.10 00a8.59d8.f1f6 Mar 20 1993 01:52 AM Automatic
#
#
#
# show ip dhcp binding
IP address Client-ID/ Lease expiration Type
Hardware address
10.1.1.7 0168.f728.fbc0.0f Mar 11 1993 01:54 AM Automatic <<<<<< pc
20.1.1.10 00a8.59d8.f1f6 Mar 20 1993 01:52 AM Automatic <<<<< phone
```

```
# ping 10.1.1.7

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.7, timeout is 2
!!!!!
seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/8 ms
# ping 10.1.1.7

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.7, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/8 ms
# ping 10.1.1.7

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.7, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/4/9 ms
# ping 20.1.1.10

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.1.1.10, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/8 ms
#
```

C. Authenticate the PC using EAP dot1x

1. Connect the PC to the phone for EAP authentication and VLAN assignment (for data)
2. After authentication:
   a. On the FortiSwitch unit, verify that the port is authorized and that the data VLAN assigned to dynamic has been placed on the allowed list.

```
# diagnose switch 8 status
Signal 10 received - config reload scheduled

wrdapd_hostapd_dump_state_console Hostapd own address 90:6c:ac:18:6f:2f
dump_diag:1:
receive dump diagnostic 802_1x/MAB sessions. ifname :port4: dump_diag:1:

port4 : Mode: mac-based (mac-by-pass enable)
        Link: Link up
        Port State: authorized ( ) <<<<<<
        Native Vlan : 1
        Allowed Vlan list: 1,10,20-21,31,41
   <<<<<<
        Untagged Vlan list:
        Guest VLAN:

        Client MAC Type Vlan Dynamic-Vlan
        68:f7:28:fb:c0:0f 802.1x 1 10
<<<<<<<<<<<<<<<<<<<<<< PC
```

```
00:a8:59:d8:f1:f6 MAB 1 0

Sessions info:
68:f7:28:fb:c0:0f Type=802.1x,PEAP,state=AUTHENTICATED
params:reAuth=3600
00:a8:59:d8:f1:f6 Type=MAB,,state=AUTHENTICATED

params:reAuth=3600

edited on: 2016-11-29 17:25

edited on: 2016-11-29 17:59
```

b. On the PC, verify that the DHCP address is assigned.

c. From the DHCP server, check the binding and a ping from gateway to verify that the PC is reachable.

Appendix A: FortiSwitch-supported RFCs

FortiSwitchOS supports the following RFCs:

- BFD on page 490
- BGP on page 490
- DHCP on page 491
- IP/IPv4 on page 491
- IP multicast on page 491
- IPv6 on page 491
- IS-IS on page 492
- MIB on page 492
- OSPF on page 493
- Other protocols on page 493
- RADIUS on page 493
- RIP on page 494
- SNMP on page 494
- Syslog on page 494
- VXLAN on page 494

BFD

- RFC 5880: Bidirectional Forwarding Detection (BFD)
- RFC 5881: Bidirectional Forwarding Detection (BFD) for IPv4 and IPv6 (Single Hop)
- RFC 5882: Generic Application of Bidirectional Forwarding Detection (BFD)

BGP

- RFC 1771: A Border Gateway Protocol 4 (BGP-4)
- RFC 1965: Autonomous System Confederations for BGP
- RFC 1997: BGP Communities Attribute
- RFC 2545: Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing
- RFC 2796: BGP Route Reflection - An Alternative to Full Mesh IBGP
- RFC 2842: Capabilities Advertisement with BGP-4
- RFC 2858: Multiprotocol Extensions for BGP-4
- RFC 4271: A Border Gateway Protocol 4 (BGP-4)
- RFC 6286: Autonomous-System-Wide Unique BGP Identifier for BGP-4
- RFC 6608: Subcodes for BGP Finite State Machine Error
- RFC 6793: BGP Support for Four-Octet Autonomous System (AS) Number Space

- RFC 7432: BGP MPLS-Based Ethernet VPN
- RFC 7606: Revised Error Handling for BGP UPDATE Messages
- RFC 7607: Codification of AS 0 Processing
- RFC 7705: Autonomous System Migration Mechanisms and Their Effects on the BGP AS_PATH Attribute
- RFC 8212: Default External BGP (EBGP) Route Propagation Behavior without Policies
- RFC 8654: Extended Message Support for BGP

DHCP

- RFC 2131: Dynamic Host Configuration Protocol
- RFC 3046: DHCP Relay Agent Information Option
- RFC 7513: Source Address Validation Improvement (SAVI) Solution for DHCP

IP/IPv4

- RFC 2697: A Single Rate Three Color Marker
- RFC 3168: The Addition of Explicit Congestion Notification (ECN) to IP
- RFC 5227: IPv4 Address Conflict Detection
- RFC 5517: Cisco Systems' Private VLANs: Scalable Security in a Multi-Client Environment
- RFC 7039: Source Address Validation Improvement (SAVI) Framework

IP multicast

- RFC 2710: Multicast Listener Discovery (MLD) for IPv6 (MLDv1)
- RFC 4541: Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches
- RFC 4605: Internet Group Management Protocol (IGMP)/Multicast Listener Discovery (MLD)-Based Multicast Forwarding (“IGMP/MLD Proxying”)
- RFC 4607: Source-Specific Multicast for IP

IPv6

- RFC 2464: Transmission of IPv6 Packets over Ethernet Networks: Transmission of IPv6 Packets over Ethernet Networks
- RFC 2474: Definition of the Differentiated Services Field (DS Field) in the and IPv6 Headers (DSCP)
- RFC 2893: Transition Mechanisms for IPv6 Hosts and Routers
- RFC 4213: Basic Transition Mechanisms for IPv6 Hosts and Router
- RFC 4291: IP Version 6 Addressing Architecture

- RFC 4443: Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification
- RFC 4861: Neighbor Discovery for IP version 6 (IPv6)
- RFC 4862: IPv6 Stateless Address Auto configuration
- RFC 5095: Deprecation of Type 0 Routing Headers in IPv6
- RFC 6724: Default Address Selection for Internet Protocol Version 6 (IPv6)
- RFC 7113: Implementation Advice for IPv6 Router Advertisement Guard (RA-Guard)
- RFC 8200: Internet Protocol, Version 6 (IPv6) Specification
- RFC 8201: Path MTU Discovery for IP version 6

IS-IS

- RFC 1195: Use of OSI IS-IS for Routing in TCP/IP and Dual Environments
- RFC 5308: Routing IPv6 with IS-IS

MIB

- RFC 1213: Management Information Base for Network Management of TCP/IP-based internets: MIB-II
- RFC 1354: IP Forwarding Table MIB
- RFC 1493: Definitions of Managed Objects for Bridges
- RFC 1573: Evolution of the Interfaces Group of MIB-II
- RFC 1643: Definitions of Managed Objects for the Ethernet-like Interface Types
- RFC 1724: RIP Version 2 MIB Extension
- RFC 1850: OSPF Version 2 Management Information Base
- RFC 2233: The Interfaces Group MIB using SMIv2
- RFC 2618: RADIUS Authentication Client MIB
- RFC 2620: RADIUS Accounting Client MIB
- RFC 2665: Definitions of Managed Objects for the Ethernet-like Interface Types
- RFC 2674: Definitions of Managed Objects for Bridges with Traffic Classes, Multicast Filtering and Virtual LAN Extensions
- RFC 2787: Definitions of Managed Objects for the Virtual Router Redundancy Protocol
- RFC 2819: Remote Network Monitoring Management Information Base
- RFC 2863: The Interfaces Group MIB
- RFC 2932: IPv4 Multicast Routing MIB
- RFC 2934: Protocol Independent Multicast MIB for IPv4
- RFC 3289: Management Information Base for the Differentiated Services Architecture
- RFC 3433: Entity Sensor Management Information Base
- RFC 3621: Power Ethernet MIB
- RFC 6933: Entity MIB (Version 4)

OSPF

- RFC 1583: OSPF Version 2
- RFC 1765: OSPF Database Overflow
- RFC 2328: OSPF Version 2
- RFC 2370: The OSPF Opaque LSA Option
- RFC 2740: OSPF for IPv6
- RFC 3101: The OSPF Not-So-Stubby Area (NSSA) Option
- RFC 3137: OSPF Stub Router Advertisement
- RFC 3623: Graceful OSPF Restart
- RFC 5340: OSPF for IPv6
- RFC 5709: OSPFv2 HMAC-SHA Cryptographic Authentication
- RFC 6549: OSPFv2 Multi-Instance Extensions
- RFC 6845: OSPF Hybrid Broadcast and Point-to-Multipoint Interface Type
- RFC 6860: Hiding Transit-Only Networks in OSPF
- RFC 7474: Security Extension for OSPFv2 When Using Manual Key Management
- RFC 7503: OSPFv3 Autoconfiguration
- RFC 8042: OSPF Two-Part Metric
- RFC 8362: OSPFv3 Link State Advertisement (LSA) Extensibility

Other protocols

- RFC 854: Telnet Protocol Specification
- RFC 2030: Simple Network Time Protocol (SNTP) Version 4 for IPv4, IPv6 and OSI
- RFC 2362: Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification
- RFC 3176: InMon Corporation's sFlow: A Method for Monitoring Traffic in Switched and Routed Networks
- RFC 3768: Virtual Router Redundancy Protocol (VRRP)
- RFC 3954: Cisco Systems NetFlow Services Export Version 9
- RFC 5101: Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information
- RFC 5798: Virtual Router Redundancy Protocol (VRRP) Version 3 for IPv4 and IPv6

RADIUS

- RFC 2865: Remote Authentication Dial In User Service (RADIUS)
- RFC 2866: RADIUS Accounting
- RFC 3580: IEEE 802.1X Remote Authentication Dial In User Service (RADIUS) Usage Guidelines
- RFC 4675: RADIUS Attributes for Virtual LAN and Priority Support
- RFC 5176: Dynamic Authorization Extensions to Remote Authentication Dial In User Service (RADIUS)

RIP

- RFC 1058: Routing Information Protocol
- RFC 2080: RIPng for IPv6
- RFC 2082: RIP-2 MD5 Authentication
- RFC 2453: RIP Version 2
- RFC 4822: RIPv2 Cryptographic Authentication

SNMP

- RFC 1157: A Simple Network Management Protocol (SNMP)
- RFC 2571: An Architecture for Describing SNMP Management Frameworks
- RFC 2572: Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)
- RFC 2573: SNMP Applications
- RFC 2576: Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework

Syslog

- RFC 5424: The Syslog Protocol
- RFC 5426: Transmission of Syslog Messages over UDP

VXLAN

- RFC 7348: Virtual eXtensible Local Area Network (VXLAN)

NOTE: For FortiSwitch hardware support, broadcast traffic uses unicast replication. Not supported in software.

Appendix B: Supported attributes for RADIUS CoA and RSSO

Attributes sent from the FortiSwitch unit to the RADIUS server during MAB (Access-Request)

Attribute	AVP Type	Type	Description
NAS-Identifier	32	text	Host name of switch
User-Name	1	alphanumeric	User name of supplicant or MAC address
User -Password	2	string	User password of supplicant

Attribute	AVP Type	Type	Description
Service-Type	6	enum	Optional. The following settings are available: - administrative—The user granted access to the administrative interface. - authenticate-only—Authentication is requested, and no authentication information needs to be returned. - call-check—This setting is used by the NAS in an Access-Request packet or Access-Accept packet to answer the call. - callback-administrative—The user disconnected, called back, and granted access to the administrative interface. - callback-framed—The user disconnected and called back and then used a Framed-Protocol attribute. - callback-login—The user disconnected and called back. - callback-nas-prompt—The user disconnected and called back and then provided a command prompt. - framed—The user used a Framed-Protocol attribute. - login—The user should be connected to a host. - nas-prompt—The user provided a command prompt on the NAS. - none—Disable the Service-Type AVP. - outbound—The user granted access to outgoing devices. The default is `none` for 802.1X authentication. MAC Authentication Bypass (MAB) always uses the `call-check` setting, no matter what is configured.
Framed-MTU	12	integer	Configurable (size of bytes). The range of values is 600-1500. The default value is 1500.
NAS-Port-Id	87	text	Port connected to supplicant
NAS-Port	5	integer	Value of port ID; for example, 12 means port12
NAS-Port-Type	61	enum	Ethernet (15)
Calling-Station-ID	31	text	MAC address of supplicant
Message-Authenticator	80	string	The Message-Authenticator attribute is a checksum of the entire Access-Request packet, containing the Type, ID, Length, and Authenticator field; the shared secret is used as the key.

Attributes sent from the FortiSwitch unit to the RADIUS server during 802.1X authentication (Access-Request)

Attribute	AVP Type	Type	Description
NAS-Identifier	32	text	Host name of switch
User-Name	1	alphanumeric	User name of supplicant or MAC address
EAP-Message	79	concat	Include EAP content
Framed-MTU	12	integer	Configurable (size of bytes). The range of values is 600-1500. The default value is 1500.
NAS-Port-Id	87	text	Port connected to supplicant
NAS-Port	5	integer	Value of port ID; for example, 12 means port12
NAS-Port-Type	61	enum	Ethernet (15)
Calling-Station-ID	31	text	MAC address of supplicant
Message-Authenticator	80	string	The Message-Authenticator attribute is a checksum of the entire Access-Request packet, containing the Type, ID, Length, and Authenticator field; the shared secret is used as the key.
Service-Type	6	enum	Optional. The following settings are available: - administrative—The user granted access to the administrative interface. - authenticate-only—Authentication is requested, and no authentication information needs to be returned. - call-check—This setting is used by the NAS in an Access-Request packet or Access-Accept packet to answer the call. - callback-administrative—The user disconnected, called back, and granted access to the administrative interface. - callback-framed—The user disconnected and called back and then used a Framed-Protocol attribute. - callback-login—The user disconnected and called back. - callback-nas-prompt—The user disconnected and called back and then provided a command prompt. - framed—The user used a Framed-Protocol attribute. - login—The user should be connected to a host. - nas-prompt—The user provided a command prompt on the NAS.

Attribute	AVP Type	Type	Description
			- none—Disable the Service-Type AVP. - outbound—The user granted access to outgoing devices. The default is `none` for 802.1X authentication. MAC Authentication Bypass (MAB) always uses the `call-check` setting, no matter what is configured.

Attributes sent from the RADIUS server to the FortiSwitch unit during 802.1X authentication (Access-Accept)

Attribute	AVP Type	Type	Description
User-Name	1	alphanumeric	User name of supplicant (MAC address of host in MAB)
Class	25	string	Whatever the server returns
Tunnel-Type	64	enum	Optional. Set to 13 for VLAN.
Tunnel-Medium-Type	65	vsa	Optional. Set to 6 for IEEE-802.
Tunnel-Private-Group-ID	81	text	VLAN number or VLAN name
Egress-VLANID	56	integer	Provides the VLAN identifier and controls whether egress packets are tagged.
Egress-VLAN-Name	58	text	Provides the VLAN name and controls whether egress packets are tagged.
Ingress-Filters	57	enum	Enables (1) the use of ingress filters. The use of ingress filters cannot be disabled.
Vendor-Specific	26	vsa	Fortinet-Group-Name
Filter-Id	11	text	Relayed from the server
Session-Timeout	27	integer	How many seconds before the session times out

RADIUS attributes in the Accounting Start message

Attribute	AVP Type	Description
Acct-Status-Type	40	1 for Start
Acct-Session-Id	44	802.1X or MAB session ID generated by the switch. For example: 0000004b
User-Name	1	Host login name or MAC address. For example: host01
Acct-Multi-Session-Id	50	For example, e81cba8e8146 in MAC mode. This attribute cannot be used in port mode. The minimum value is 1; the maximum value is 1.
NAS-Identifier	32	For example, S148EP591900009 for the host name of the switch.

Attribute	AVP Type	Description
Framed-IP-Address	8	This value is the host IP address if is found in the switch; otherwise, the switch does not send this attribute. For example: 100.1.0.3
NAS-Port-Id	87	This value is a text string that identifies the port of the NAS connected to the host. For example: port48
NAS-Port	5	This value indicates the physical port number of the NAS. For example: 48
NAS-Port-Type	61	0 for asynchronous
Called-Station-Id	30	MAC address of the 802.1X port. For example: E8-1C-BA-8E-81-46
Calling-Station-Id	31	MAC address of host. For example: 00-12-01-00-00-01
Event-Timestamp	55	Time when the event occurred. For example: May 31, 2019 12:25:03.00000000 Pacific Daylight Time
Filter-Id	11	Relayed from the server
Vendor-Specific	26	Fortinet-Group-Name. Authentication fails if this value does not match.
Class	25	Whatever the server returns

RADIUS attributes in the Accounting Interim Update message

Attribute	AVP Type	Description
Acct-Status-Type	40	3 for Interim-Update
Acct-Session-Id	44	802.1X or MAB session ID generated by the switch. For example: 0000004b
User-Name	1	Host login name or MAC address. For example: host01
Acct-Multi-Session-Id	50	For example, e81cba8e8146 in MAC mode. This attribute cannot be used in port mode.
Acct-Link-Count	51	2 for two sessions on the port. This attribute is only valid for MAC mode.
NAS-Identifier	32	For example, S148EP591900009 for the host name of the switch.
Framed-IP-Address	8	This value is the host IP address if is found in the switch; otherwise, the switch does not send this attribute. For example: 100.1.0.3
NAS-Port-Id	87	This value is a text string that identifies the port of the NAS connected to the host. For example: port48
NAS-Port	5	This value indicates the physical port number of the NAS. For example: 48
NAS-Port-Type	61	15 for Ethernet
Called-Station-Id	30	MAC address of the 802.1X port. For example: E8-1C-BA-8E-81-46
Calling-Station-Id	31	MAC address of host. For example: 00-12-01-00-00-01

Attribute	AVP Type	Description
Event-Timestamp	55	Time when the event occurred. For example: May 31, 2019 12:25:03.00000000 Pacific Daylight Time
Filter-Id	11	Eng-Group. If Filter-Id is received during authentication, it is included in accounting.
Class	25	Whatever the server returns
Vendor-Specific	26	Fortinet-Group-Name. Authentication fails if this value does not match.

RADIUS attributes in the Accounting Stop message

Attribute	AVP Type	Description
Acct-Status-Type	40	2 for Stop
Acct-Session-Id	44	802.1X or MAB session ID generated by the switch. For example: 0000004b
User-Name	1	Host login name or MAC address. For example: host01
Acct-Multi-Session-Id	50	For example, e81cba8e8146 in MAC mode. This attribute cannot be used in port mode.
Acct-Link-Count	51	2 for two sessions on the port
NAS-Identifier	32	For example, S148EP591900009 for the host name of the switch.
Framed-IP-Address	8	This value is the host IP address if is found in the switch; otherwise, the switch does not send this attribute. For example: 100.1.0.3
NAS-Port-Id	87	This value is a text string that identifies the port of the NAS connected to the host. For example: port48
NAS-Port	5	This value indicates the physical port number of the NAS. For example: 48
NAS-Port-Type	61	15 for Ethernet
Called-Station-Id	30	MAC address of the 802.1X port. For example: E8-1C-BA-8E-81-46
Calling-Station-Id	31	MAC address of host. For example: 00-12-01-00-00-01
Acct-Input-Octets	42	3200
Acct-Output-Octets	43	16050448
Acct-Input-Packets	47	20
Acct-Output-Packets	48	93606
Acct-Terminate-Cause	49	6 for Admin-Reset
Event-Timestamp	55	Time when the event occurred. For example: May 31, 2019 12:25:03.00000000 Pacific Daylight Time
Filter-Id	11	Eng-Group. If Filter-Id is received during authentication, it is included in accounting.

Attribute	AVP Type	Description
Class	25	Whatever the server returns
Vendor-Specific	26	Fortinet-Group-Name. Authentication fails if this value does not match.

RADIUS attributes in the Disconnect-Request message

Attribute	AVP Type	Description
Calling-Station-ID	31	MAC address of host
Framed-IP-Address	8	IP address of host
User-Name	1	Host login name
NAS-IP-Address	4	NAS IP address
Message-Authenticator	80	The Message-Authenticator attribute is a checksum of the entire Access-Request packet, containing the Type, ID, Length, and Authenticator field; the shared secret is used as the key.
Event-Timestamp	55	Time when the event occurred. For example: May 31, 2019 12:25:03.00000000 Pacific Daylight Time

RADIUS attributes in the Disconnect-ACK message

Attribute	AVP Type	Description
Event-Timestamp	55	Time when the event occurred. For example: May 31, 2019 12:25:03.00000000 Pacific Daylight Time
Message-Authenticator	80	The Message-Authenticator attribute is a checksum of the entire Access-Request packet, containing the Type, ID, Length, and Authenticator field; the shared secret is used as the key.

RADIUS attributes in the Disconnect-NAK message

Attribute	AVP Type	Description
Calling-Station-ID	31	MAC address of host
NAS-Port	5	Port that the host is connected to
Acct-Session-Id	44	802.1X or MAB session identifier generated by the switch
Framed-IP-Address	8	IP address of host
User-Name	1	Host login name
Error-Cause	101	Refer to the “Error-Cause codes in RADIUS CoA-NAK and Disconnect-NAK messages” table in this appendix for a listing of error causes, error codes, and descriptions.

RADIUS attributes in the CoA-Request message (reauth-port)

Attribute	AVP Type	Description
Calling-Station-ID	31	MAC address of host
Message-Authenticator	80	The Message-Authenticator attribute is a checksum of the entire Access-Request packet, containing the Type, ID, Length, and Authenticator field; the shared secret is used as the key.
Vendor-Specific	26	Fortinet-Group-Name
Event-Timestamp	55	Time when the event occurred. For example: May 31, 2019 12:25:03.00000000 Pacific Daylight Time
User-Name	1	Host login name

RADIUS attributes in the CoA-Request message (disable-port)

Attribute	AVP Type	Description
Calling-Station-ID	31	MAC address of host
User-Name	1	Host login name
NAS-IP-Address	4	NAS IP address
Message-Authenticator	80	The Message-Authenticator attribute is a checksum of the entire Access-Request packet, containing the Type, ID, Length, and Authenticator field; the shared secret is used as the key.
Vendor-Specific	26	Fortinet-Group-Name
Event-Timestamp	55	Time when the event occurred. For example: May 31, 2019 12:25:03.00000000 Pacific Daylight Time
Class	25	Whatever the server returns
Filter-Id	11	Relayed from the server

RADIUS attributes in the CoA-Request message (bounce-port)

Attribute	AVP Type	Description
Calling-Station-ID	31	MAC address of host
User-Name	1	Host login name
Message-Authenticator	80	The Message-Authenticator attribute is a checksum of the entire Access-Request packet, containing the Type, ID, Length, and Authenticator field; the shared secret is used as the key.
Vendor-Specific	26	Fortinet-Group-Name

Attribute	AVP Type	Description
Event-Timestamp	55	Time when the event occurred. For example: May 31, 2019 12:25:03.00000000 Pacific Daylight Time
Class	25	Whatever the server returns
Filter-Id	11	Relayed from the server

RADIUS attributes in the CoA-Request message (session-timeout)

Attribute	AVP Type	Description
Calling-Station-ID	31	MAC address of host
NAS-Port	5	Port that the host is connected to
Acct-Session-Id	44	802.1X or MAB session identifier generated by the switch
Framed-IP-Address	8	IP address of host
User-Name	1	Host login name

RADIUS attributes in the CoA-ACK message

Attribute	AVP Type	Description
Event-Timestamp	55	Time when the event occurred. For example: May 31, 2019 12:25:03.00000000 Pacific Daylight Time
Message-Authenticator	80	The Message-Authenticator attribute is a checksum of the entire Access-Request packet, containing the Type, ID, Length, and Authenticator field; the shared secret is used as the key.

RADIUS attributes in the CoA-NAK message

Attribute	AVP Type	Description
Error-Cause	101	Refer to the “Error-Cause codes in RADIUS CoA-NAK and Disconnect-NAK messages” table in this appendix for a listing of error causes, error codes, and descriptions.
Event-Timestamp	55	Time when the event occurred. For example: May 31, 2019 12:25:03.00000000 Pacific Daylight Time
Message-Authenticator	80	The Message-Authenticator attribute is a checksum of the entire Access-Request packet, containing the Type, ID, Length, and Authenticator field; the shared secret is used as the key.

Error-Cause codes in RADIUS CoA-NAK and Disconnect-NAK messages

Error Cause	Error Code	Description
Unsupported Attribute	401	This error is a fatal error, which is sent if a request contains an attribute that is not supported.
NAS Identification Mismatch	403	This error is a fatal error, which is sent if one or more NAS-Identifier Attributes do not match the identity of the NAS receiving the request.
Invalid Attribute Value	407	This error is a fatal error, which is sent if a CoA-Request or Disconnect-Request message contains an attribute with an unsupported value.
Session Context Not Found	503	This error is a fatal error if the session context identified in the CoA-Request or Disconnect-Request message does not exist on the NAS.

Stop error codes for RADIUS accounting

Error Message	Error Code	Description
ACCT_TERM_CAUSE_IDLE_TIMEOUT	4	The system has been idle for too long.
ACCT_TERM_CAUSE_USER_REQUEST	1	The user requested the service to be stopped.
ACCT_TERM_CAUSE_SESSION_TIMEOUT	5	The session has timed out.
ACCT_TERM_CAUSE_ADMIN_RESET	6	The administrator has reset the session or port.

Appendix C: SNMP OIDs for FortiSwitch models

The following table lists the SNMP object identifiers (OIDs) for FortiSwitch models. The SNMP OIDs correspond to the `sysObjectID` OID that is defined in RFC 1213, “Management Information Base for Network Management of TCP/IP-based internets: MIB-II.” The `sysObjectID` is different for each FortiSwitch model.

You can check the latest OIDs by downloading the MIB from any FortiSwitch unit. Go to *System* > *Config* > *SNMP* > *Settings* and clicking *FortiSwitch MIB File*.

FortiSwitch model	SNMP OID
FS-108D-POE	1081
FS-108E	1082
FS-108E-POE	1083
FS-108E-FPOE	1084
FS-108F	1086
FS-108F-POE	1087
FS-108F-FPOE	1088
FS-110G-FPOE	1102
FS-124D	1241
FS-124D-POE	1242
FS-124E	1244
FS-124E-POE	1245
FS-124E-FPOE	1246
FS-124F	12410
FS-124F-POE	12411
FS-124F-FPOE	1249
FS-148E	1247
FS-148E-POE	1248
FS-148F	1484
FS-148F-POE	1485
FS-148F-FPOE	1486
FS-224D-POE	2241
FS-224D-FPOE	2242
FS-224E	2243

FortiSwitch model	SNMP OID
FS-224E-POE	2244
FS-248D	2483
FS-248D-POE	2481
FS-248D-FPOE	2482
FS-248E-POE	2485
FS-248E-FPOE	2484
FS-424D	4241
FS-424D-POE	4242
FS-424D-FPOE	4243
FS-424E	42401
FS-424E-POE	42402
FS-424E-FPOE	42403
FS-424E-Fiber	42404
FS-M426E-FPOE	42405
FS-448D	4482
FS-448D-FPOE	4483
FS-448E	4485
FS-448E-POE	4486
FS-448E-FPOE	4487
FS-524D	5242
FS-524D-FPOE	5241
FS-548D	5482
FS-548D-FPOE	5481
FS-624F	6241
FS-624F-FPOE	6242
FS-648F	6481
FS-648F-FPOE	6482
FS-1024D	10241
FS-1024E	10242
FS-T1024E	10243

FortiSwitch model	SNMP OID
FS-T1024F-FPOE	10244
FS-1048D	10481
FS-1048E	10482
FS-2048F	20481
FS-3032D	30321
FS-3032E	30322
FSR-112D-POE	1121
FSR-124D	1243
FSR-216F-POE	2161
FSR-424F-POE	42406

www.fortinet.com

Copyright© 2025 Fortinet, Inc. All rights reserved. Fortinet®, FortiGate®, FortiCare® and FortiGuard®, and certain other marks are registered trademarks of Fortinet, Inc., and other Fortinet names herein may also be registered and/or common law trademarks of Fortinet. All other product or company names may be trademarks of their respective owners. Performance and other metrics contained herein were attained in internal lab tests under ideal conditions, and actual performance and other results may vary. Network variables, different network environments and other conditions may affect performance results. Nothing herein represents any binding commitment by Fortinet, and Fortinet disclaims all warranties, whether express or implied, except to the extent Fortinet enters a binding written contract, signed by Fortinet's Chief Legal Officer, with a purchaser that expressly warrants that the identified product will perform according to certain expressly-identified performance metrics and, in such event, only the specific performance metrics expressly identified in such binding written contract shall be binding on Fortinet. For absolute clarity, any such warranty will be limited to performance in the same ideal conditions as in Fortinet's internal lab tests. Fortinet disclaims in full any covenants, representations, and guarantees pursuant hereto, whether express or implied. Fortinet reserves the right to change, modify, transfer, or otherwise revise this publication without notice, and the most current version of the publication shall be applicable.